



# Improved financial forecasting via quantum machine learning

Sohum Thakkar<sup>1</sup> · Skander Kazdagli<sup>2</sup> · Natansh Mathur<sup>2,3</sup> · Iordanis Kerenidis<sup>2,3</sup> · André J. Ferreira–Martins<sup>4</sup> · Samurai Brito<sup>4</sup>

Received: 13 July 2023 / Accepted: 25 March 2024  
© The Author(s) 2024

## Abstract

Quantum algorithms have the potential to enhance machine learning across a variety of domains and applications. In this work, we show how quantum machine learning can be used to improve financial forecasting. First, we use classical and quantum Determinantal Point Processes to enhance Random Forest models for churn prediction, improving precision by almost 6%. Second, we design quantum neural network architectures with orthogonal and compound layers for credit risk assessment, which match classical performance with significantly fewer parameters. Our results demonstrate that leveraging quantum ideas can effectively enhance the performance of machine learning, both today as quantum-inspired classical ML solutions, and even more in the future, with the advent of better quantum hardware.

**Keywords** Computational finance · Machine learning · Quantum computing · Credit risk · Churn prediction

## 1 Introduction

Quantum computing is a rapidly evolving field that promises to revolutionize various domains, and finance is no exception. There is a variety of computationally hard financial problems for which quantum algorithms can potentially offer advantages (Herman et al. 2022; Egger et al. 2020; McKinsey & Company 2021; Bouland et al. 2020), for example in combinatorial optimization (Leclerc et al. 2022; Rebentrost and Lloyd 2018), convex optimization (Kerenidis et al. 2019; Rebentrost et al. 2022), Monte Carlo simulations (Doriguello et al. 2022; Suzuki et al. 2020; Giurgica-Tiron et al. 2022), and machine learning (Pistoia et al. 2021; Emmanoulopoulos and Dimoska 2022; Alcazar et al. 2020; Nguyen and Chen 2022).

In this work, we explore the potential of quantum machine learning methods in improving the performance of forecasting in finance, specifically focusing on two use cases within

the business of Itaú Unibanco, the largest bank in Latin America.

In the first use case, we aim to improve the performance of Random Forest methods for churn prediction. We introduce quantum algorithms for Determinantal Point Processes (DPP) sampling (Kerenidis and Prakash 2022), and develop a method of DPP sampling to enhance Random Forest models. We evaluate our model on the churn dataset using classical DPP sampling algorithms and perform experiments on a scaled-down version of the dataset using quantum algorithms. Our results demonstrate that, in the classical setting, the proposed algorithms outperform the baseline Random Forest in precision, efficiency, and bottom line, and also offer a precise understanding of how quantum computing can impact this kind of problem in the future. The quantum algorithm run on an IBM quantum processor gives similar results as the classical DPP on small batch dimensions but falters as the dimensions grow bigger due to hardware noise.

In the second use case, we aim to explore the performance of neural network models for credit risk assessment by incorporating ideas from quantum compound neural networks (Landman et al. 2022). We start by using quantum orthogonal neural networks (Landman et al. 2022), which add the property of orthogonality for the trained model weights to avoid redundancy in the learned features (Arjovsky et al. 2016). These orthogonal layers, which can be trained efficiently on a classical computer, are the simplest case of what we call compound neural networks, which explore an exponential

✉ Sohum Thakkar  
sohum.thakkar@qcware.com

✉ André J. Ferreira–Martins  
andre.ferreira-martins@itau-unibanco.com.br

<sup>1</sup> QC Ware Corp, Palo Alto, USA

<sup>2</sup> QC Ware Corp, Paris, France

<sup>3</sup> IRIF, Université Paris Cité and CNRS, Paris, France

<sup>4</sup> Itaú Unibanco, São Paulo, Brazil

space in a structured way. For our use case, we design compound neural network architectures that are appropriate for financial data. We evaluate their performance on a real-world dataset and show that the quantum compound neural network models both have far fewer parameters and achieve better accuracy and generalization than classical fully connected neural networks.

This paper is organized as follows: In Sections 2–4, we focus on the churn prediction use case and present the DPP-based quantum machine learning methods. In Sections 5–7, we present quantum neural network models for risk assessment. Finally, in Section 8, we conclude the paper and discuss potential future research directions.

## 2 DPP-enhanced Random Forest models for churn prediction

### 2.1 DPP-Random Forest model

The Random Forest algorithm was introduced in 2001 by Breiman (2001) and has easily become one of the most popular supervised machine learning algorithms in use. It consists of an ensemble of decision trees, each trained on a uniform subsample of rows and columns from the dataset.

In this section, we propose an extension of the Random Forest, called the DPP-Random Forest (DPP-RF), which utilizes Determinantal Point Processes (DPPs) instead of uniform sampling to subsample rows and columns for individual decision trees. In the original RF algorithm, subsampling makes the model more robust to variance in the training data; however, the use of uniform sampling runs the risk of improperly representing the dataset and missing under-sampled areas. DPP sampling better preserves the diversity of the dataset, and corrects for sampling bias (Kulesza and Taskar 2012). We first introduce the theory and techniques of DPP sampling, then present the algorithm.

### 2.2 Determinantal point processes

We will now introduce the Determinantal Point Process (DPP), which lies at the core of the methodology behind our solution to the churn problem. DPPs are a class of probabilistic models that can be used to sample diverse subsets of items from a larger set. They were first formalized by Macchi in 1975 as a way to model fermions in quantum mechanics (Macchi 1975). More recently, these models are showing increasing promise in the context of machine learning (Kulesza and Taskar 2012), where they can be used for a variety of tasks, such as building unbiased estimators for linear regression (Derezinski and Mahoney 2021), performing Monte Carlo estimation (Bardenet and Hardy 2020), and promoting diversity in model outputs (Elfeki et al. 2019).

### 2.2.1 Definitions

A point process  $P$  on a set  $Y$  is a probability measure over the subsets of  $Y$ . Sampling from a point process on  $Y$  will produce some subset  $S \subseteq Y$  with probability  $P(S)$ . A repulsive point process is a point process in which points that are more similar to each other are less likely to be selected together.

A determinantal point process (DPP) is a particular case of a repulsive point process, in which the selection probability of a subset of items  $T \subseteq Y$  is given by a determinant. Given a real, symmetric  $n \times n$  matrix  $K$  indexed by the elements of  $Y$ :

$$P\{T \subseteq S\} = \det(K_{T,T}),$$

where  $K_{T,T}$  denotes the  $|T| \times |T|$  submatrix indexed by the set  $T$  and  $n$  is the cardinality of  $Y$ . In other words, the marginal distribution  $P\{T \subseteq S\}$  is defined by the subdeterminants of  $K$ .

The above is the most general definition, but in machine learning, we typically focus on a slightly more restrictive class of DPPs called  $L$ -ensembles. In  $L$ -ensembles, the whole distribution, not just the marginals, is given by the subdeterminant of a real, symmetric  $n \times n$  matrix  $L$ .

$$P\{S\} \propto \det(L_{S,S}).$$

Just like  $K$ ,  $L$  is indexed by the elements of  $Y$ . Because of some convenient properties of the determinant (Kulesza and Taskar 2012), we can explicitly write down the distribution of an  $L$ -ensemble:

$$P\{S\} = \frac{\det(L_{S,S})}{\det(L + I)}.$$

In machine learning literature, DPPs are typically defined over a set of points  $\mathbf{X}$ , with each item  $\mathbf{x}_i$  a row in the data matrix  $\mathbf{X}$ . If we preprocess  $\mathbf{X}$  such that its columns are orthonormal and choose  $\mathbf{L}$  to be the inner-product similarity matrix, i.e.,  $L = \mathbf{X}\mathbf{X}^T$ , then the distribution becomes even simpler to write down. Instead of explicitly computing the  $\mathbf{L}$  matrix, we can write the distribution in terms of the data matrix  $\mathbf{X}$  itself, courtesy of the Cauchy-Binet formula,

$$P\{S\} = \frac{\det(\mathbf{X}_S)^2}{\det(\mathbf{X}\mathbf{X}^T + \mathbf{I})}. \quad (1)$$

Moreover, the distribution will almost surely produce samples with size  $d$ , the rank of the orthogonalized data matrix  $\mathbf{X}$ . This kind of DPP is denoted  $d$ -DPP. We will focus here on an application of sampling from a  $d$ -DPP from a data matrix  $\mathbf{X}$ .

## 2.2.2 Unbiased least squares regression

One unique feature of the DPP compared to i.i.d sampling techniques is that it can lead to provably unbiased estimators for least squares linear regression (Derezinski 2018; Derezinski et al. 2018). Given an  $n \times d$  data matrix  $\mathbf{X}$  and a target vector  $\mathbf{y} \in \mathbb{R}^n$ , where  $n \gg d$ , we wish to approximate the least squares solution  $\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|$ .  $\mathbf{w}^*$  represents the best-fit parameters to a linear model to predict  $\mathbf{y}$ .

Surprisingly, if we sample  $d$  points  $S$  from  $\operatorname{DPP}(\mathbf{X}\mathbf{X}^T)$  and solve the reduced system of equations  $\mathbf{y}_S = \mathbf{X}_S\mathbf{w}$ , we get an unbiased estimate of  $\mathbf{w}^*$ . Formally, if  $S \sim d\text{-DPP}_{\mathbf{L}}(\mathbf{X}\mathbf{X}^T)$ ,

$$\mathbb{E}[\mathbf{X}_S^{-1}\mathbf{y}_S] = \operatorname{argmin}_{\mathbf{w}} \|\mathbf{X}\mathbf{w} - \mathbf{y}\| = \mathbf{w}^*. \quad (2)$$

This allows us to create an “ensemble” of unbiased linear regressors, each trained on a DPP sample. In some regard, this was the inspiration for trying an ensemble of *decision trees* trained on DPP samples, as detailed in Section 2.1.

## 2.2.3 Algorithms for sampling

There are several efficient algorithms for sampling from DPPs and computing their properties. The naive sampling method — calculating all subdeterminants and performing  $l2$  sampling — takes exponential time. The first major leap in making DPP sampling feasible on today’s computers was the “spectral method” (Kulesza and Taskar 2011; Hough et al. 2006). This algorithm performs an eigendecomposition of the kernel matrix before applying a projection-based iterative sampling approach. Thus, the first sample takes  $O(nd^2)$  time, and subsequent samples take  $O(d^3)$ .

Monte Carlo methods have been proposed to approximate the DPP distribution (Anari et al. 2016; Li et al. 2016), though they are not exact, and are often still prohibitively slow with a runtime of  $O(n\operatorname{poly}(d))$  per sample.

In a counter-intuitive result, Derezinski et al. (2019) and Calandriello et al. (2020) proposed methods that avoid performing the full DPP sampling procedure on large parts of the basis set. This approach resulted in a significant reduction in runtime, making DPPs more practical for mid-to-large-scale datasets. These techniques allow exact sampling of subsequent  $d$ -DPP samples in  $O(\operatorname{poly}(d))$ , independent of the size of the full basis set  $n$ . Many of these algorithms are implemented in the open-source DPPy library (Gautier et al. 2019), which we used in the experiments in this paper.

Recent work has shown that quantum computers are in principle able to sample from DPPs in even lower complexity in some cases. We describe this quantum algorithm in Section 4. This and several other algorithms which arise from the techniques introduced in Kerenidis and Prakash (2022) are a budding area of research in the quantum computing

space, and will hopefully inspire more applications like the one we describe in this paper. For example, in Kazdaghli et al. (2023), DPPs and deterministic DPPs were used to improve the methods for the imputation of clinical data.

## 2.3 DPP-RF algorithm outline

In principle, the DPP-RF uses DPP sampling on the whole dataset to select diverse subsets of data on which to train decision trees. However, sampling from a DPP on large datasets (like the entire churn dataset of 174,000 points) can take copious time, especially when using current open-source implementations of DPP sampling. To be able to test these techniques quickly, a novel sampling procedure was developed which preserves many of the benefits of DPPs, but does not require sampling from the full dataset. The procedure can be summarized as follows:

1. Divide the training set uniformly into smaller batches;
2. Sample  $S_1 \sim d\text{-DPP}(\mathbf{X}_{batch}\mathbf{X}_{batch}^T)$  data points from every batch;
3. Sample  $S_2 \sim d\text{-DPP}(\mathbf{X}_{S_1}^T\mathbf{X}_{S_1})$  features;
4. Train a first group  $G_1$  of  $N_1$  decision trees on these small patches of data;
5. Aggregate the patches of data resulting from step 2 to create a larger dataset  $\mathbf{X}_{agg}$ ;
6. Repeat for  $N_2$  times: sample  $S_3 \sim d\text{-DPP}(\mathbf{X}_{agg}^T\mathbf{X}_{agg})$  features to create a long matrix;
7. Train a second group  $G_2$  of  $N_2$  decision trees on these new datasets;
8. Combine  $G_1$  and  $G_2$  by aggregating them to make predictions (similar to the classical Random Forest algorithm).

## 3 Classical DPP-RF results

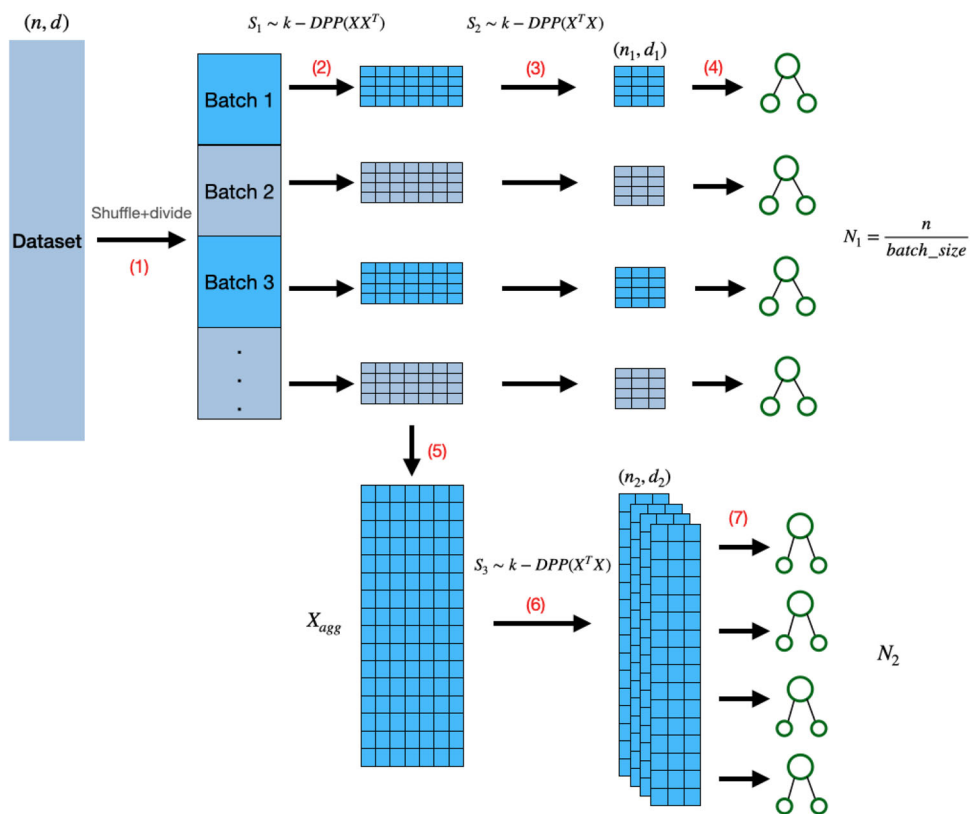
The DPP-RF algorithm was designed for the purpose of predicting customer churn in the bank. In this section, we define this use case and present the results. In addition, we benchmark our proposed DPP-RF method by constructing models on public tabular classification tasks (Fig. 1).

### 3.1 Use case introduction: churn prediction

Churn, defined as a customer withdrawing more than a certain amount of money in a single month, is a significant concern for retail banks. Our objective is to predict which customers are most likely to churn in the next three months using customer data from the previous six months.

The primary dataset used in this study consists of 304,000 datapoints, with 153 features for each datapoint. Each datapoint represents a banking customer at a particular month in time, with the features representing various aspects of their

**Fig. 1** Steps 1 to 7 of the DPP-Random Forest algorithm



activity over the previous six months. The target variable is a binary flag indicating whether or not the customer churned in at least one of the following three months. The data was anonymized and standardized before being split into training and test sets based on time period, with 130,000 datapoints being set aside as the test set and 174,000 datapoints used for training. The data was split in a way that did not produce any significant covariate shift between the train and test sets.

With the end goal of preventing churn, the model works by flagging customers with the highest risk of potential churn. For these flagged customers, the bank can deploy a representative to intervene and better understand their needs. However, resource limitations make it necessary to flag a relatively small number of customers with high confidence. The focus of this exploration was to reduce false positives in the flagged customers to increase the efficiency of bank interventions. In terms of the precision-recall trade-off, our model should be tuned to provide the highest possible precision for low recall values. Despite this simplification to a classification problem, the primary business KPI is the amount of withdrawal money correctly captured by the model, as discussed more in Section 3.4.

This use case already had a solution in production: a Random Forest classifier (Breiman 2001), whose performance was used as a benchmark. The model in production already captured a significant amount of churn, but there was clear room for improvement in the amount of withdrawals captured

(see Fig. 3). Moreover, given the large number of customers in the dataset and the relative homogeneity of the population of interest, there existed an opportunity to employ techniques that explicitly try to explore diversity in the data.

We focused on three key performance indicators (KPIs): the precision-recall curve, the training time and the bottom line.


### 3.2 Precision-recall

To evaluate the performance of our proposed method, we optimized hyperparameters and measured the precision for a low fixed recall (6% in this case). As seen in Fig. 2, our method showed an improvement in precision from 71.6% for the benchmark model to 77.5% with the new model. Our method also provided similar improvements in precision for the relevant range of small recall.

### 3.3 Training time

The DPP-RF model has a longer training time compared to the traditional random forest on a classical computer: it took 54 min to train the model with the best hyperparameters using, compared to 311 s for the benchmark model. The models were trained on a computer with an Intel® Core™ i5-8350U CPU running at 1.70 GHz, 24 GB of RAM and Windows 10 version 21H2, compilation 19044.2604.

**Fig. 2** Precision-recall curve for the test set. Using DPP with the Random Forest algorithm shows an improvement of 5.9%



The computational bottleneck in this algorithm is the DPP sampling. Instead of simulating quantum DPP circuits (which is infeasible for large datasets), we used a classical SVD-based sampling algorithm (Hough et al. 2006) implemented in the dppy library (Gautier et al. 2019). We believe that improved classical sampling techniques (Calandriello et al. 2020) and future quantum techniques (Section 4) can reduce the runtime dramatically.

Hyperparameters were selected using a grid search with 5-fold cross-validation over the typical RF parameters `n_estimators`, `max_depth`, `min_samples_leaf`, `min_samples_split`, `max_features`, `max_samples`, and the DPP-RF-specific parameter `batch_size`. The training time of a DPP-RF model depends heavily on this `batch_size` parameter, which is the size of the batches from which we take DPP samples. Choosing a batch size higher than 1000 can increase runtime dramatically. Thus, in our hyperparameter search, we limited the batch size to less than 1000.

Within the bank, the churn model is retrained just once every few months, so the training time was not prohibitive. However, faster sampling algorithms still serve to increase the range of feasible hyperparameters (especially the batch size).

### 3.4 Bottom line — withdrawals captured

From a business perspective, the most direct indicator of the success of the model is the amount of assets under management (AUM) that can be salvaged via interventions. Thus, we evaluated the amount of money withdrawn every month by the 500 customers flagged by the model, i.e., the 500 customers which had the highest predicted probability of churning in one of the following 3 months. As seen in Figs. 3, 4 and 5, our model showed substantial overall

improvements. The true financial impact of these predictions is dependent on the success of the interventions as well as the bank's profit-per-dollar-AUM.


### 3.5 Summary of results

The proposed DPP-Random Forest model provides significant improvements in precision and bottom line, while taking significantly longer to train. The results are summarized in (Table 1).

### 3.6 Further benchmarks


We further benchmarked our model on various classification datasets from OpenML. All except one (madelon) of these datasets were used in Grinsztajn et al. (2022) and pre-processed accordingly. They were chosen to be representative of a wide variety of classification tasks. Each dataset was split into train, validation, and test sets. For each model, 400 sets of hyperparameters were randomly chosen and evaluated on the validation set. Both models used the same hyperparameter space, except for the addition of the `batch_size` parameter for the DPP-RF. The hyperparameters which gave the best results on the validation set were evaluated on the test set, and the results are reported in (Table 2). Models were evaluated with the ROC-AUC metric<sup>1</sup>.

<sup>1</sup> The area under the receiver operating characteristic curve (ROC-AUC) is a common metric for two-class classification tasks, and evaluates the ability of the model to produce a proper ranking of data-points by likelihood of being class 1.




**Fig. 3** Classical benchmark (BM) vs DPP-RF solution: money withdrawn per month by the flagged 500 customers, comparing the benchmark model (blue line) to the DPP-RF one (orange line). On the y-axis, we have monetary values (not shown). The green line represents the total amount of money withdrawn by all customers in each month.

The purple line is the sum of the 500 largest withdrawals, which is the maximum value that the model could capture. The red line represents the withdrawals captured by randomly flagging 500 observations. The y-axis units are omitted for confidentiality



**Fig. 4** Classical benchmark vs DPP-RF solution — percentage of total withdrawals captured per month, that is, relative to the green line in Fig. 3. On average over the 11 test months, the BM model captures 61.42% of the total, while the DPP-RF model captures 62.77% — an improvement of 1.35%



**Fig. 5** Classical benchmark vs DPP-RF solution — the percentage of maximum money possible to be captured (given  $n\_flags = 500$  customers flagged every month), that is, relative to the purple line in Fig. 3.

On average over the 11 test months, the BM model captures 69.18% of the total, while the DPP-RF model captures 70.72% — an improvement of 1.54%

**Table 1** Summarized comparison between models

Metric	Benchmark model	Proposed model
Precision	71.6%	77.5%
% total withdrawals captured	61.42%	62.77%
% maximum possible withdrawals captured	69.18%	70.72%
Train time	311 s	54 min

## 4 Quantum DPP-RF

### 4.1 Quantum circuits for determinantal point processes

Classical DPP sampling algorithms have improved significantly since their inception, but there may still be room for improvement. Recent work by Kerenidis and Prakash (2022) has shown that a quantum computer can more natively perform DPP sampling, achieving a gate complexity of  $O(nd)$  and a circuit depth of  $O(d \log(n))$  for an orthogonal matrix of size  $n \times d$ . The classical time complexity for sampling is  $O(d^3)$  (Hough et al. 2006). Note that when  $n$  is very large, then one can reduce the number of rows to  $O(d^2)$  before performing the sampling (Calandriello et al. 2020).

For a thorough review of the quantum methods and circuits, we refer the reader to Kerenidis and Prakash (2022). The circuit is described in brief below.

Given an orthogonal matrix  $\mathbf{X} = (\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n) \in \mathbb{R}^{n \times d}$ , the quantum DPP circuit applied on  $\mathbf{X}$  performs the following operation:

$$\mathcal{D}(\mathbf{X})|0^n\rangle = \sum_{\substack{|S|=d \\ S \in \{0,1\}^n}} \det(\mathbf{X}_S) |1_S\rangle, \quad (3)$$

**Table 2** Comparison of DPP-Random Forest and Random Forest models for different datasets. Superior results in bold

Dataset	Random forest	DPP-Random Forest
madelon	0.916	<b>0.941</b>
credit-default	0.856	0.856
house-pricing	<b>0.948</b>	0.939
jannis	0.866	<b>0.870</b>
eye movements	0.704	<b>0.710</b>
bank-marketing	0.881	0.881
wine	0.901	<b>0.906</b>
california	0.962	<b>0.963</b>

The results are reported via the ROC-AUC metric

where  $\mathbf{X}_S$  is the  $\mathbb{R}^{d \times d}$  submatrix obtained after sampling the rows of  $\mathbf{X}$  indexed by  $S$ ;  $1_S$  is the characteristic vector of  $S$  (with 1's in the positions indexed by the elements of  $S$ ) and  $\mathcal{D}(\mathbf{X})$  represents the quantum  $d$ -DPP circuit, as detailed below.

Thus, the probability of sampling  $S$ , i.e., of measuring  $|1_S\rangle$ , is:  $Pr(S) = \det(\mathbf{X}_S)^2 = \det(L_{S,S})$ , where  $L = \mathbf{X}\mathbf{X}^T$ . This draws the link between the quantum determinantal sampling circuit and the classical  $d$ -DPP model as seen in Eq. 1.

To construct the quantum  $d$ -DPP circuit, we need to first introduce a circuit known as a *Clifford loader*, which performs the following operation:

$$\mathcal{C}(\mathbf{x}) = \sum_{i=1}^n x_i Z^{i-1} X I^{n-i}, \quad \text{for } \mathbf{x} \in \mathbb{R}^n. \quad (4)$$

The Clifford loader was shown to have a log-depth circuit in Kerenidis and Prakash (2022), and is shown for  $n = 8$  in Fig. 6, in which the gates represented by vertical lines are RBS gates — parameterized, hamming weight-preserving two-qubit gates.

The full quantum  $d$ -DPP circuit is a series of  $d$  Clifford loaders, one for each orthogonal column of  $\mathbf{X}$ :

$$\mathcal{D}(\mathbf{X}) = \mathcal{C}(\mathbf{x}^1)\mathcal{C}(\mathbf{x}^2) \dots \mathcal{C}(\mathbf{x}^d). \quad (5)$$

An example of a  $d$ -DPP circuit as a series of Cliffords for  $n = 4$  is shown in Fig. 7.


### 4.2 Hardware experiment results

As a hardware experiment, we aimed to implement a simplified version of our algorithm on a quantum processor. We chose to use the “ibmq\_guadalupe” 16-qubit chip, which is only capable of running small quantum DPP circuits for matrices of certain dimensions, such as (4, 2), (5, 2), (5, 3), (6, 2), (8, 2). As a result, we had to reduce the size of our problem.

To accomplish this, we defined reduced train/test sets: a train set of  $\sim 1000$  points from 03/2019 and a test set of  $\sim 10,000$  points from 04/2019. The quantum hardware-ready simplified algorithm is outlined in Fig. 8. It includes the following steps:

1. Applying PCA to reduce the number of columns from 153 to  $d = 2, 3$ ;
2. Dividing the dataset into batches of  $n = 4, 5, 6, 8$  points;
3. Sampling  $S \sim d$ -DPP( $\mathbf{X}_{batch}\mathbf{X}_{batch}^T$ ) rows from each batch, resulting in small  $d \times d$  patches of data;
4. Aggregating these patches to form a larger dataset, then training one decision tree on this dataset.

**Fig. 6** Clifford loader circuit  $\mathcal{C}(x)$  for  $x \in \mathbb{R}^8$



We repeated this process for a number of trees and estimated the  $F1^2$  score for every tree. We then compared the results for different sampling methods: uniform sampling, quantum DPP sampling using a simulator, and quantum DPP sampling using a quantum processor.

The IBM quantum computer only allows using RBS gates on adjacent qubits, so we cannot use the circuit described in Section 4.1. Instead, we use two different Clifford loader architectures which only use adjacent-qubit connectivity. The diagonal Clifford loader Fig. 9 is explained in Kerenidis and Prakash (2022), and the semi-diagonal loader (Fig. 10) is a modification that halves the circuit depth. As an error mitigation measure, we disregarded all results that did not have the expected hamming weight ( $d$ ). The results are shown in the violin plots in Figs. 11 and 12.

The results indicate that for small matrix dimensions — up to (6,2) — the IBM quantum processor gives results similar to the ones achieved with the simulator. However, as the dimensions grow bigger, the samples from the quantum DPP circuits lead to worse classifier performance. This highlights the limitations of the available quantum hardware, which is prone to errors.

## 5 Quantum neural networks for credit risk assessment

### 5.1 Quantum neural networks with orthogonal and compound layers


In recent years, variational/parameterized quantum circuits (Benedetti et al. 2019) have become very prominent as NISQ-

<sup>2</sup> We chose the  $F1$  score as the evaluation metric for two reasons. Firstly, a single decision tree, unlike the random forest, does not provide an estimate of the likelihood, which is required for the computation of ROC-AUC as we had used before. Secondly, we had an imbalanced dataset and thus needed a metric that balanced precision and recall. The  $F1$  score is particularly effective when used in these scenarios.

friendly QML techniques. When applied to classification problems, they are commonly known as Variational Quantum Classifiers (VQC) (Havlíček et al. 2018). The quantum circuits associated with VQCs may be schematically thought of as composed of three layers: the *feature map*  $\mathcal{U}_{\Phi(\vec{x})}$ , which encodes classical data  $\vec{x}$  into quantum states; the *variational layer*  $W(\theta)$ , which is the part of the circuit parameterized by a set of parameters  $\theta$  which are learned in the training process; and finally, the *measurement layer*, which measures the quantum registers and produces classical information used in training and inference.

The feature map and variational layers can take different forms, called *ansätze*, consisting of many possible different quantum gates in different configurations. Such immense freedom raises an important question: how should one choose an architecture for a given problem, and can it be expected to yield a quantum advantage? This question is of major practical importance, and although benchmark results have been shown for very particular datasets (Havlíček et al. 2018; Liu et al. 2021), there is little consensus on which ansätze are good choices for machine learning.

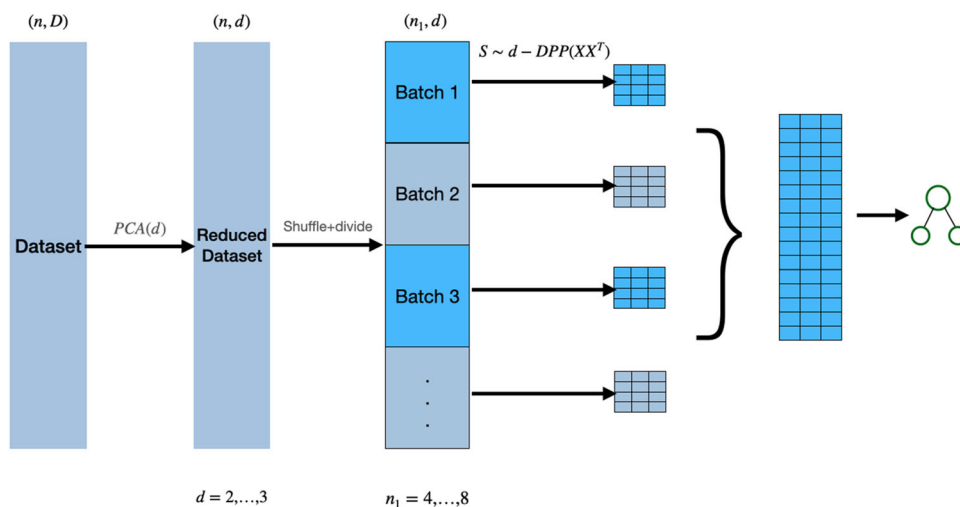
In our work, we use quantum neural networks with orthogonal and compound layers. Although these neural networks roughly match the general VQC construction, they produce well-defined linear algebraic operations, which not only



**Fig. 7** DPP circuit as a series of Clifford loaders



**Fig. 8** Quantum hardware-ready procedure for DPP sampling



makes them much more interpretable but gives us the ability to analyze their complexity and scalability. Because we understand the actions of these layers precisely, we are able to identify instances for which we can design efficient classical simulators, allowing us to classically train and test the models on real-scale datasets.

A standard feed-forward neural network layer modifies an input vector by first multiplying it by a weight matrix and then applying a non-linearity to the result. Feed-forward neural networks usually use many such layers and learn to predict a target variable by optimizing the weight matrices to minimize a loss function. Enforcing the orthogonality of these weight matrices, as proposed in Jia et al. (2019), brings theoretical and practical benefits: it reduces the redundancy in the trained weights and can avoid the age-old problem of vanishing gradients. However, the overhead of typical projection-based methods to enforce orthogonality prevents mainstream adoption.

In Landman et al. (2022), an improved method of constructing orthogonal neural networks using quantum ideas was developed. We describe it below in brief.


### 5.2 Data loaders

In order to perform a machine learning task with a quantum computer, we need to first load classical data into the quantum circuit.

#### Unary data loading circuits

The first way we will load classical data is an example of *amplitude encoding*, which means that we load the (normalized) vector elements as the amplitudes of a quantum state. In Johri et al. (2021), three different circuits to load a vector  $\mathbf{x} \in \mathbb{R}^d$  using  $d - 1$  gates are proposed. The circuits range in depth from  $O(\log(d))$  to  $O(d)$ , with varying qubit connectivity (see Fig. 13). They use the *unary* amplitude encoding, where a vector  $\mathbf{x} = (x_1, \dots, x_d)$  is loaded in the quantum state  $|\mathbf{x}\rangle = \frac{1}{\|\mathbf{x}\|} \sum_{i=1}^d x_i |e_i\rangle$ , where  $|e_i\rangle$  is the quantum state with all qubits in  $|0\rangle$  except the  $i^{th}$  qubit in state  $|1\rangle$  (e.g.,  $|e_3\rangle = |00100000\rangle$ ). The circuit uses *RBS* gates: a parameterized two-qubit hamming weight-preserving gate

**Fig. 9** Diagonal Clifford loader






Fig. 10 Semi-diagonal Clifford loader

implementing the unitary given by Eq. 6:

$$RBS(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta & \sin \theta & 0 \\ 0 & -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{6}$$

The parameters  $\theta_i : i \in \{1, \dots, d - 1\}$  of the  $d - 1$  RBS gates are classically pre-computed to ensure they encode the correct vector  $|\mathbf{x}\rangle$ .

### RY-loading circuits

We will also use data loading procedures beyond the unary basis. In particular, for a normalized input vector  $\mathbf{x} \in \mathbb{R}^d$ , we use  $d$  qubits, where on each of the qubits, we apply an  $RY(\theta)$  rotation gate where the angle parameter on the  $i^{th}$  qubit is  $\theta_i = 2\pi x_i$ , according to Eq. 7. Multiplication with  $2\pi$  allows us to cover the entire range of the sin and cos functions. This technique loads the data in the entire  $2^d$ -dimensional Hilbert space encompassing all the hamming weights from 0 to  $d$ . This loading technique has constant depth independent of  $d$ , and we refer to it as the *RY loading*, whose circuit for  $d = 8$

is illustrated in Fig. 14.

$$RY(\theta)|0\rangle = \cos \frac{\theta}{2}|0\rangle + \sin \frac{\theta}{2}|1\rangle \tag{7}$$

### H-loading circuits

Lastly, we define a different technique for loading the data in the entire  $2^d$ -dimensional Hilbert space, which loads the vector in the unary basis and then applies a Hadamard gate on each qubit. This operation applies a Fourier transform on  $\mathbb{Z}_2$  and gives us a state encompassing all the hamming weights from 0 to  $d$  at no additional cost to the circuit depth. We call this the *H-loading*, whose circuit for  $d = 8$  is illustrated in Fig. 14.

$$H|0\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \quad H|1\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}} \tag{8}$$

The  $RY$  and  $H$  loading circuits spread the data over all  $2^n$  bases to allow future RBS-based neural network layers to utilize an exponential space. In contrast, unary loaders spread the data across the  $n$  bases with hamming weight 1, and since RBS gates are hamming weight-preserving (known as *match gates* (Jozsa and Miyake 2008)), they cannot change this. Their action when using unary loaders is thus restricted to a much smaller space.

### 5.3 Quantum orthogonal and compound layers

Quantum orthogonal layers consist of a unary data loader plus a parametrized quantum circuit made of  $RBS$  gates, while quantum compound layers consist of a general data loader plus a parametrized quantum circuit made of  $RBS$  gates.

$RBS$  gates and circuits preserve the hamming weight of the input state, and thus if we use a unary data loader, then the output of the layer will be another vector in unary amplitude encoding. Similarly, if the loaded quantum state is a superposition of only basis states of hamming weight  $k$ , so is the output state. More generally, we can think of such hamming weight-preserving circuits with  $n$  qubits as block-diagonal




Fig. 11 Decision trees performance using quantum DPP sampling with diagonal Clifford loaders




Fig. 12 Decision trees performance using quantum DPP sampling with semi-diagonal Clifford loaders


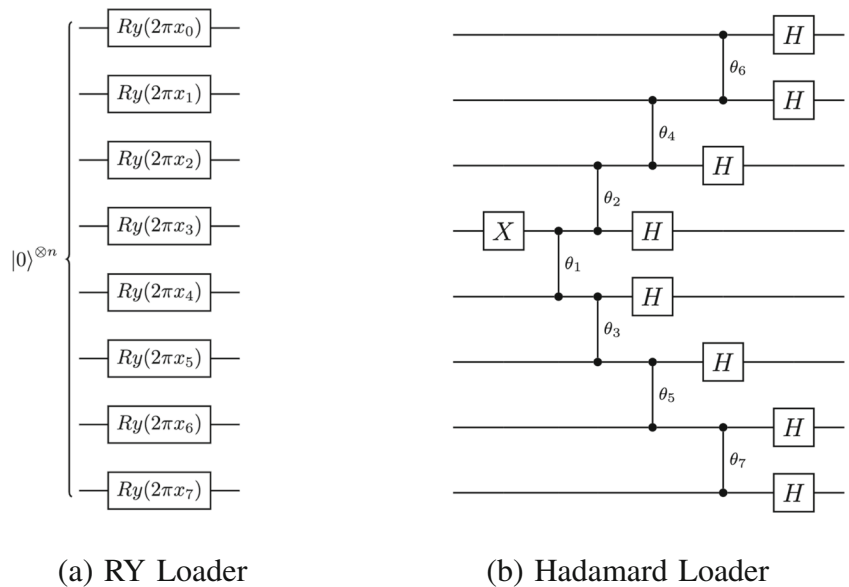


Fig. 13 Three possible unary data loaders for  $d$ -dimensional vectors ( $d = 8$ ). From left to right: the parallel, diagonal, and semi-diagonal circuits have respectively a circuit depth of  $\log(d)$ ,  $d$ , and  $d/2$ . The X gate represents the Pauli X gate, and the vertical lines represent RBS gates with tunable parameters

Fig. 14 Non-unary loaders



(a) RY Loader

(b) Hadamard Loader

unitaries that act separately on  $n + 1$  subspaces, where the  $k^{th}$  subspace is defined by all computational basis states with hamming weight equal to  $k$ . The dimension of these subspaces is equal to  $\binom{n}{k}$ . The first block of this unitary is an  $n \times n$  orthogonal matrix, such that when a vector is loaded in the unary basis, this circuit simply performs orthogonal matrix multiplication. In general, the  $k$ -th block of this unitary applies a *compound matrix of order  $k$*  of the  $n \times n$  unary matrix. The dimension of this  $k$ -th order compound matrix is  $\binom{n}{k} \times \binom{n}{k}$ . We refer to the layers that use bases beyond the unary as *compound layers*.

There exist many possibilities for building a parametrized quantum circuit made of *RBS* gates which can be used in a quantum orthogonal or compound layer, each with different properties.

The **Pyramid circuit** (Fig. 15), proposed in Landman et al. (2022), is a parameterized quantum circuit composed of exactly  $n(n - 1)/2$  *RBS* gates. This circuit requires only adjacent-qubit connectivity, which makes it suitable for most superconducting qubit hardware. In addition, when restricted to the unary basis, the pyramid circuit expresses exactly the Special Orthogonal Group, i.e., orthogonal matrices with the determinant equal to  $+1$ . To allow this circuit to express the entire orthogonal group, we can add a final *Z* gate on the last qubit. This allows us to express orthogonal matrices with a  $-1$  determinant as well. The pyramid circuit is, therefore, very general and covers all the possible orthogonal matrices of size  $n \times n$ .

The **X circuit** (Fig. 15), introduced in Cherrat et al. (2022), uses just  $O(n)$  gates and has nearest-neighbor connectivity. Due to reduced depth and gate complexity, it accumulates less hardware noise.

The **Butterfly circuit** (Fig. 15) is inspired by the classical fast Fourier transform algorithm, and uses  $O(n \log(n))$  gates. It was also introduced in Cherrat et al. (2022), and despite having reduced expressivity compared to the Pyramid circuit, it often performs just as well.


In Landman et al. (2022), a method is proposed to train orthogonal layers for the unary basis by computing the

gradient of each parameter  $\theta_i$  using backpropagation. This backpropagation method for the pyramid circuit (which is the same for any circuit with *RBS* gates) takes time  $O(n^2)$ , corresponding to the number of gates, and provides a polynomial improvement in runtime compared to the previously known orthogonal neural network training algorithms which relied on an  $O(n^3)$  SVD operation (Jia et al. 2019). Since the runtime corresponds to the number of gates, it is lower for the butterfly and X circuits. See Table 3 for full details on the comparison between the three types of circuits. For the compound layers, we need to consider the entire  $2^n \times 2^n$  space and thus train an exponential size weight matrix, which takes exponential time on a classical computer. In principle, a compound layer can also be trained using the parameter shift rule for quantum circuits, which can be more efficient since the number of parameters is polynomial in the input size, though noise in current quantum hardware makes this impractical for the time being.

### 5.4 Expectation-per-subspace compound layer

We describe here a compound layer that we call the Expectation-per-subspace compound layer. This layer involves loading the input vector using a non-unary basis which could be done either via the *RY*-loading or the *H*-loading circuit as previously defined. Then, we apply a parameterized quantum circuit with *RBS* gates, e.g., a pyramid circuit, which performs the compound matrix operation on all the fixed hamming weight subspaces. More precisely, we can think of the operation as performing the matrix-vector multiplication of an  $\binom{n}{k} \times \binom{n}{k}$  matrix with an  $\binom{n}{k}$ -dimensional vector for each hamming weight  $k$  from 0 to  $n$ . Note that for 0 and  $n$ , the dimension is 1, and hence the unitary acts as identity.

If we look at the output quantum state, it defines a distribution over a domain of size  $2^n$ . Given the exponential size of the distribution, it is not advisable to try and train the entire distribution, since that would take exponential time.




**Fig. 15** Parameterized quantum circuits for orthogonal and compound layers. Vertical lines represent two-qubit *RBS* gates, parameterized with independent angles  $\theta$ , which are shown as the same in the figures for simplicity

**Table 3** Comparison of different parameterized quantum circuits for orthogonal and compound layers with  $n$  qubits

Circuit	Hardware Connectivity	Depth	# Gates
Pyramid	Nearest-Neighbor	$2n - 3$	$\frac{n(n-1)}{2}$
X	Nearest-Neighbor	$n - 1$	$2n - 3$
Butterfly	All-to-all	$\log(n)$	$\frac{n}{2} \log(n)$

However, one can still try to use a loss function that contains some information about the distribution. For example, one can use the expectation of the distribution, which is what normally happens in variational quantum algorithms where one approximates this expectation by using a number of measurement outcomes. Given the fact that our unitary is block-diagonal, one can try to define a more complex loss function that contains more information about the distribution. In particular, one can split the domain of the distribution into  $n + 1$  subdomains, one for each subspace, and then train on all these expectations.

This is what we do in the Expectation-per-subspace compound layer, where for each  $k$  from 0 to  $n$ , we take the outputs corresponding to the hamming weight  $k$  strings and sort them. Now, for each  $k$ , we assign values which are equally spaced between two bounds  $a$  and  $b$  (which are 0 and 10, in our models) to the  $\binom{n}{k}$  strings. We normalize the outputs using the  $L1$ -norm to correspond to a probability distribution over the  $\binom{n}{k}$  values between  $a$  and  $b$ , and then we calculate the expectation value for that hamming weight. This gives us a set of  $n + 1$  values corresponding to each hamming weight. Since for hamming weight 0 and  $n$  the dimension of the subspace is 1 (the all-zero and all-one strings), we combine them and calculate the expectation for these two together and make the layer have  $n$  outputs. The entire operation is illustrated in Fig. 16.



**Fig. 16** Expectation-per-subspace compound layer. In the final step, we combine the  $\binom{n}{0}$  and the  $\binom{n}{n}$  subspaces and calculate their overall expectation

While these compound layers do increase classical simulation complexity, they do not increase quantum complexity. The advent of better quantum hardware will allow us to test larger compound layers that explore much larger portions of the Hilbert space.

## 6 QNN results with classical simulation

In this second study, we focus on the problem of credit-default prediction associated with credit applications from Small and Medium Enterprises (SMEs). We report the results of our neural network architectures on this use case and on public datasets.

### 6.1 Use case introduction: credit risk prediction

The credit operation is one of the largest and most important operations in a retail banking institution. Throughout the credit journey (life-cycle) of a customer within the bank, several different models are used at different points of the journey and for different purposes, such as the determination of interest rates, offering of different products, etc.

The credit granting model is a particularly important one since it determines whether or not a credit relationship will be established. It is also particularly challenging in the case of SMEs, where the relationship with the bank often starts only when the SME submits an application for credit, so very little data is available.

Given these challenges, we propose the use of quantum techniques aiming at improving the predictive performance of the credit granting model.

The credit granting decision may be seen as a binary classification problem, in which the objective is to predict if the SME will default on credit. More specifically, we are

interested in calculating the so-called probability of default (PD), which is given by  $P(\hat{y} = 1|\mathbf{x})$ . The PD information is used internally for other pipelines primarily concerned with the determination of credit ratings for the SMEs (though in this study, we focus solely on the PD model), so the PD distribution is the main output of interest from the model. For this reason, we do not threshold the probability outputs from the model — thus, we use threshold-independent classification metrics to evaluate its predictive performance. Namely, the main Key Performance Indicator (KPI) that we use is the Gini score, constructed from the Area Under the Curve (AUC) of the ROC (Receiver Operating Characteristic) curve as  $\text{Gini} = 2 \times \text{AUC} - 1$ . The Gini score is easily interpreted by the business team and allows for a holistic estimation of the model's impact.

In this study, we chose to focus on the development of an “internal model” of credit default, which only uses features collected by Itaú, without considering any external information (from credit bureaus, for instance). The dataset used consists of  $\approx 141,500$  observations, each one represented by 32 features: 31 numerical and 1 categorical. Each observation represents a given SME customer in a specified reference month, whose observed target indicates its default behavior, and whose features consist of internal information about the company. The data was anonymized, standardized, and split into training and test sets based on the time period: the training set consists of  $\approx 74,700$  observations covering 12 months of data, while in the test set, we have  $\approx 66,800$  observations covering the subsequent 8 months.

## 6.2 Neural network architectures for credit risk

To compare the performance of the orthogonal and compound layers to the classical baseline, we designed three neural network architectures. Each architecture had three layers: an *encoding layer*, an *experimental layer*, and a *classification head*. The encoding layer was a standard linear layer of size  $32 \times 8$  followed by a tanh activation. Its purpose is to bring the dimension of the features down to 8, which is a reasonable simulation size for both proposed quantum layers. The second layer was the experimental layer of size  $8 \times 8$  (described below). Finally, the third layer, the classification head, was a linear layer of size  $8 \times 2$  followed by softmax to predict the probabilities.

The first quantum neural network architecture, named **OrthoResNN**, uses an  $8 \times 8$  orthogonal experimental layer implemented with a semi-diagonal loader and  $X$  circuit. Note that the final output of the layer is provided by measurements. We add a skip connection by adding the input of the

orthogonal layer to the output.<sup>3</sup> The layer is followed by a tanh activation function. This architecture is illustrated in Fig. 17.

Our second architecture, **ExpResNN**, replaces the experimental layer with an  $8 \times 8$  Expectation-per-subspace compound layer. We use the  $H$ -loader to encode our data. The layer is again followed by a tanh activation function. Figure 18 illustrates the ExpResNN architecture.

And finally, the classical architecture, **ResNN**, used an  $8 \times 8$  linear residual layer followed by tanh.

## 6.3 Methods and training

The training of the networks was performed using the JAX package by Google. We train our models for 500 epochs. To identify suitable hyperparameters, we performed a search over learning rate, learning-rate-halving points, and batch size. The hyperparameter search was performed with the `ray-tune` library.

The dataset contains a large number of missing values, which motivated the experimentation of different imputation techniques such as zero-filling, round-robin imputation (implemented in the Python package `scikit-learn`), and MICE (van Buuren and Groothuis-Oudshoorn 2011). The best results were achieved with round-robin imputation using `scikit-learn`'s `IterativeImputer` with Bayesian ridge regression. This was the pre-processing employed in all the results on the SME dataset.


## 6.4 Results

In our experimental setup, we consider the fully connected residual layer (ResNN) as the classical benchmark. We performed the same experiment with an orthogonal layer using the semi-diagonal loader and the  $X$  circuit (OrthoResNN). Finally, we tried the expectation-per-subspace compound layer with the Hadamard loader and  $X$  circuit (ExpResNN). While the performance of the OrthoResNN and ExpResNN remained nearly the same as the FNN layer, these new layers learn the angles of  $2n$  RBS gates instead of  $n^2$  elements of the weight matrix, dramatically reducing the number of parameters needed. The results are shown in Table 4.

The results show that quantum orthogonal and compound layers can preserve the performance of fully connected layers on this dataset while using a fraction of the trainable parameters. We note that the ExpResNN did not show advantages over the OrthoResFNN.

<sup>3</sup> The model cannot learn to “ignore” the orthogonal layer via the skip connection, because the orthogonal layer cannot change the magnitude of the input.

**Fig. 17** Architecture of the OrthoResNN model




**6.5 Further benchmarks**

We compared the orthogonal and linear layers on public classification datasets from OpenML as we did for the Random Forests (Section 3.6). Again, we used a three-layer architecture: a linear encoding layer to 16 dimensions with GeLU

activation, a  $16 \times 16$  experiment layer with tanh activation, and a binary classification head. These architectures did not have residual connections.

We compared two models: **OrthoFNN** and **FNN**. OrthoFNN used a  $16 \times 16$  pyramid circuit as the experiment layer,

**Fig. 18** Architecture of the ExpResNN model using the *H*-loader



**Table 4** Comparison between different architectures

Model	Trainable parameters	Test Gini
ResNN	64	54.20%
OrthoResNN	13	54.29%
ExpResNN	13	53.95%

for a total of 120 trainable parameters. FNN used a feed-forward linear layer with 256 trainable parameters.

For all the datasets and models, we use a batch size of 128 and a learning rate of  $10^{-4}$ . Each is trained for 500 epochs and evaluated with the ROC-AUC metric. The results are summarized in (Table 5).

## 7 QNN results on quantum hardware

### 7.1 Implementation of quantum circuits


Using a classical computer, inference using an orthogonal layer takes time  $O(n^2)$ , while for a general compound layer, this time is exponential in  $n$ . Using a quantum computer, inference with an orthogonal or compound layer uses a quantum circuit that has depth  $O(n)$  (Pyramid or X) or  $O(\log(n))$  (Butterfly), and  $O(n^2)$  gates. Therefore, one may find a further advantage if the inference is performed on a quantum computer. This motivated us to test the inference step for classically trained OrthoResNN and ExpFNN models (ExpResNN from the classical experiments without a residual connection) on currently available quantum hardware.

The data loader and orthogonal/compound layer circuits employed in our model architectures are NISQ-friendly and particularly suitable for superconducting qubits, with low depth and nearest-neighbors qubit connectivity. Thus, we chose to use IBM's 27-qubit machine *ibm\_hanoi* (see Fig. 19).

**Table 5** Comparison of OrthoFNN and FNN for different datasets. Superior results in bold

Dataset	FNN	OrthoFNN
madelon	0.624	<b>0.636</b>
credit-default	<b>0.647</b>	0.634
house-pricing	0.780	<b>0.781</b>
jannis	0.791	<b>0.797</b>
eye movements	<b>0.580</b>	0.578
bank-marketing	0.847	<b>0.850</b>
wine	0.553	<b>0.640</b>
california	0.634	<b>0.668</b>

The results are reported via the ROC-AUC metric

**Fig. 19** Topology graph of the 27-qubit *ibm\_hanoi* machine used to perform our hardware experiments. The colors in the qubits indicate readout assignment error; and in the connections the CNOT error — dark blue is low, purple is high

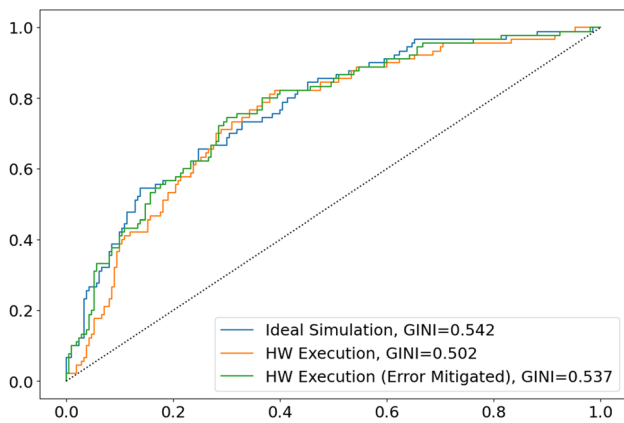
To perform inference on *ibm\_hanoi*, we used the semi-diagonal data loader and X circuit to implement the OrthoResNN model; and the Hadamard loader and X circuit for the ExpFNN model — the same architectures described in Section 6. Both neural networks were trained classically, and the trained parameters were used to construct our quantum circuits for inference.

Given the large size of the test dataset (66,750 data points), we decided to perform inference using the trained models on a small test subsample of 300 test points, corresponding to the maximum number of circuits we could send in one job to the IBM machine. After testing different subsamples with the OrthoResNN model,<sup>4</sup> we selected one for which we achieved a subsample test Gini score of 54.19% using a noiseless simulator (blue ROC curve in Fig. 20). The same was done for the ExpFNN experiment, yielding a subsample test Gini of 53.90% with the noiseless simulator. These values were taken as the best possible Gini scores if the inference was performed on noiseless quantum hardware, which could then be compared with the values actually achieved with *ibm\_hanoi*.

The circuits were then run on the quantum processor. Due to its limited Hilbert space of size  $n$ , the OrthoResNN has a natural error-rejection procedure: any measurements outside of the unary basis can be disregarded as errors. As a result, the inference yielded a Gini score of 50.19%, as shown in the orange ROC curve in Fig. 20. The achieved Gini was not too far from the noise-free simulation result (54.19%), but there was clearly a room for improvement in order to close the 4 pp difference. We also attempted inference with

<sup>4</sup> We made sure to pick a representative subsample of observations, for which we would have neither a heavily under nor overestimated Gini score compared to the one for the full test set. By selecting 50 different subsamples of randomly chosen 300 observations and performing the classical inference, we found Gini scores between 45.9% and 64.3%, with an average of 54.44%. This further supports the fact that the subsample that we chose correctly yields the (approximate) expected value for this Gini score distribution, which then yields an unbiased subsample Gini score. The same procedure was performed for the ExpFNN experiment.





**Fig. 20** ROC curves with Gini score for the ideal simulation, hardware execution, and the error-mitigated hardware execution

the more complex ExpFNN, which yielded a Gini score of 40.20%, much farther from the noiseless simulation Gini of 53.90%. Since the ExpFNN uses the entire  $2^n$ -dimensional Hilbert space, it is more prone to errors due to noise, as the error-rejection procedure used for the OrthoResNN cannot be employed.

## 7.2 Improving the hardware results with error mitigation techniques

Error mitigation and error suppression techniques undoubtedly play a very important role in NISQ-era quantum computing. While these techniques alone may not be sufficient to fully overcome the imperfections of current quantum systems, they can push the practical limits of what can be achieved. As a next step, for the OrthoResNN model, we experimented with various error mitigation and suppression approaches, going beyond the simple hamming weight post-selection procedure, in an attempt to close the gap of 4 pp between the Gini score from the noisy simulation and the one from hardware execution.

The first approach that we tried was a correlated readout mitigator. This is a purely classical post-processing technique which demands the construction of a calibration circuit for each one of the possible  $2^N$  states of the full  $N$  qubits Hilbert space. The calibration circuits' execution (simulated using `ibm_hanoi`'s backend information, in our case) yields a  $2^N \times 2^N$  assignment matrix, which is used to understand how errors might occur during readout. One can see that this method rapidly becomes intractable as the number of qubits  $N$  increases. In our case, for  $N = 8$ , the Gini score improved to 50.24%, a small improvement of only 0.05 pp.

Thus, in order to investigate the effect of more robust error suppression and mitigation techniques in our results, we moved on to a new round of hardware experiments,

performing the inference by executing the exact same OrthoResNN circuits via the Qiskit Runtime (contributors 2021) service using the Sampler primitive, which allows one to use circuit optimization as well as error suppression and mitigation techniques, as detailed below.

Firstly, we used circuit optimization at the point of circuit transpilation and compilation by setting the `optimization_level` parameter to the highest possible value, 3. This performs the following circuit optimization routines: Layout selection and routing (VF2 layout pass and SABRE layout search heuristics (Li et al. 2018); 1 qubit gate optimization (chains of single-qubit `u1`, `u2`, `u3` gates are combined into a single gate); commutative cancellation (cancelling of redundant self-adjoint gates); 2 qubit KAK optimization (decomposition of 2-qubit unitaries into the minimal number of uses of 2-qubit basis gates).

Secondly, we used the Dynamical Decoupling error suppression technique (Viola and Lloyd 1998; Ezzell et al. 2022). This technique works as a pulse schedule by inserting a DD pulse sequence into periods of time in which qubits are idle. The DD pulses effectively behave as an identity gate, thus not altering the logical action of the circuit, but having the effect of mitigating decoherence in the idle periods, reducing the impact of errors.

Thirdly, we used the M3 (Matrix-free Measurement Mitigation) error mitigation technique (Nation et al. 2021) by setting the Sampler `resilience_level` parameter to 1 (the only option available for the Sampler primitive). This provides mitigated quasi-probability distributions after the measurement. M3 works in a reduced subspace defined by the noisy input bitstrings supposed to be corrected, which is often much smaller than the full  $N$  qubits Hilbert space. For this reason, this method is much more efficient than the matrix-based readout mitigator technique mentioned above. M3 provides a matrix-free preconditioned iterative solution method, which removes the need to construct the full reduced assignment matrix but rather computes individual matrix elements, which uses orders of magnitude less memory than direct factorization.

By employing these three techniques, we were able to achieve a Gini score of 53.68% for the OrthoResNN (Fig. 20). This is a 3.49 pp improvement from the initial 50.19% Gini of the unmitigated run, falling only 0.53 pp behind the ideal noiseless execution (54.21% Gini)! This remarkable result underscores the NISQ-friendliness of the orthogonal layer and highlights the importance of error suppression and mitigation techniques in the NISQ era.

It is important to note that circuit optimization, error suppression, and mitigation techniques typically result in some classical/quantum pre/post-processing overhead to the overall circuit runtime. Some of these techniques are based on heuristics and/or do not have efficient scaling at larger circuit

sizes. It is important to balance the desired levels of optimization and resilience with the required time for the full execution, especially as the circuit sizes increase.

## 8 Conclusion

In this work, we have explored the potential of quantum machine learning methods in improving forecasting in finance, with a focus on two specific use cases within the Itaú business: churn prediction and credit risk assessment. Our results demonstrate that the proposed algorithms, which leverage quantum ideas, can effectively enhance the performance of Random Forest and neural network models, achieving better accuracy and training with fewer parameters.

In the present day, quantum hardware is not powerful enough to provide real improvements or conclusive large-scale benchmarks. Performance enhancements can be achieved today by turning these quantum ideas into classical ML solutions run on GPUs. However, with the advent of better quantum hardware, we expect these methods to run faster and produce even better results when run on quantum computers.

The general nature of the proposed methods makes them applicable to other use cases in finance and beyond, although they must be tuned to specific datasets and tasks. We hope this work inspires confidence that QML research holds promise both for today as well as for the coming era of scaled, fault-tolerant quantum hardware.

**Author contribution** S.T., N.M., S.K., and A.F. wrote the main manuscript text. N.M., S.K., and A.F. performed the quantum hardware experiments. All authors reviewed the manuscript.

**Availability of data and materials** The primary datasets and model used in this study are proprietary to Itaú and QC Ware and are not publicly available. Additional benchmark datasets are available via openML.

## Declarations

**Ethics approval** Not applicable as this study did not involve any human or animal subjects.

**Conflict of interest** S.T., S.K., N.M., and I.K. are employed by QC Ware, while A.F. and S.B. are employed by Itaú, both of which may stand to benefit from the results of this research. The authors declare no other competing interests.

**Disclaimer** This paper is a research collaboration between Itaú Unibanco and QC Ware. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Itaú Unibanco. This paper is not and does not constitute or intend to constitute investment advice or any investment service. It is not and should not be deemed to be an offer to purchase or sell, or a solicitation of an offer to purchase or sell, or a recommendation to purchase or sell any securities or other financial instruments. Moreover, all data used in this study is compliant with the Brazilian General Data Protection Law.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Herman D, Googin C, Liu X, Galda A, Safro I, Sun Y, Pistoia M, Alexeev Y (2022) A survey of quantum computing for finance. *Papers* 2201.02773, [arXiv.org](https://arxiv.org/abs/2201.02773). <https://ideas.repec.org/p/arx/papers/2201.02773.html>
- Egger DJ, Gambella C, Marecek J, McFaddin S, Mevissen M, Raymond R, Simonetto A, Woerner S, Yndurain E (2020) Quantum computing for finance: State-of-the-art and future prospects. *IEEE Transactions on Quantum Engineering* 1:1–24. <https://doi.org/10.1109/TQE.2020.3030314>
- McKinsey & Company (2021) Quantum computing: an emerging ecosystem and industry use cases. Accessed 16 Feb 2023
- Bouland A, Dam W, Joorati H, Kerenidis I, Prakash A (2020) Prospects and challenges of quantum finance. *arXiv*. <https://doi.org/10.48550/ARXIV.2011.06492>
- Leclerc L, Ortiz-Guitierrez L, Grijalva S, Albrecht B, Cline JRK, Elfving VE, Signoles A, Henriot L, Del Bimbo G, Sheikh UA, Shah M, Andrea L, Ishtiaq F, Duarte A, Mugel S, Caceres I, Kurek M, Orus R, Seddik A, Hammammi O, Isselane H, M'tamon D (2022) Financial Risk Management on a Neutral Atom Quantum Processor. *arXiv*. <https://doi.org/10.48550/ARXIV.2212.03223>
- Rebentrost P, Lloyd S (2018) Quantum computational finance: quantum algorithm for portfolio optimization. [arXiv:1811.03975](https://arxiv.org/abs/1811.03975)
- Kerenidis I, Prakash A, Szilágyi D (2019) Quantum algorithms for portfolio optimization. In: *Proceedings of the 1st ACM conference on advances in financial technologies*. AFT '19, pp. 147–155. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3318041.3355465>
- Rebentrost P, Luongo A, Bosch S, Lloyd S (2022) Quantum computational finance: martingale asset pricing for incomplete markets. *arXiv*. <https://doi.org/10.48550/ARXIV.2209.08867>
- Doriguello JaF, Luongo A, Bao J, Rebentrost P, Santha M (2022) Quantum algorithm for stochastic optimal stopping problems with applications in Finance. In: Le Gall F, Morimae T (eds) 17th conference on the theory of quantum computation, Communication and Cryptography (TQC 2022). *Leibniz International Proceedings in Informatics (LIPIcs)*, vol 232, pp 2–1224. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany. <https://doi.org/10.4230/LIPIcs.TQC.2022.2>. <https://drops.dagstuhl.de/opus/volltexte/2022/16509>
- Suzuki Y, Uno S, Raymond R, Tanaka T, Onodera T, Yamamoto N (2020) Amplitude estimation without phase estimation. *Quantum Inf Process* 19(2):75. <https://doi.org/10.1007/s11128-019-2565-2>
- Giurgica-Tiron T, Kerenidis I, Labib F, Prakash A, Zeng W (2022) Low depth algorithms for quantum amplitude estimation. *Quantum* 6:745. <https://doi.org/10.22331/q-2022-06-27-745>
- Pistoia M, Ahmad SF, Ajagekar A, Buts A, Chakrabarti S, Herman D, Hu S, Jena A, Minssen P, Niroula P, Rattew A, Sun Y, Yalovetzky R (2021) Quantum Machine Learning for Finance. *arXiv*. <https://doi.org/10.48550/ARXIV.2109.04298>

- Emmanoulopoulos D, Dimoska S (2022) Quantum machine learning in finance: time series forecasting. arXiv e-prints, 2202
- Alcazar J, Leyton-Ortega V, Perdomo-Ortiz A (2020) Classical versus quantum models in machine learning: insights from a finance application. *Mach Learn Sci Technol* 1(3):035003. <https://doi.org/10.1088/2632-2153/ab9009>
- Nguyen N, Chen K-C (2022) Bayesian quantum neural networks. *IEEE*. Access 10:54110–54122. <https://doi.org/10.1109/ACCESS.2022.3168675>
- Kerenidis I, Prakash A (2022) Quantum machine learning with subspace states. [arXiv:2202.00054](https://arxiv.org/abs/2202.00054)
- Landman J, Mathur N, Li YY, Strahm M, Kazdagli S, Prakash A, Kerenidis I (2022) Quantum methods for neural networks and application to medical image classification. *Quantum* 6:881. <https://doi.org/10.22331/q-2022-12-22-881>
- Arjovsky M, Shah A, Bengio Y (2016) Unitary evolution recurrent neural networks. In: International conference on machine learning, PMLR, pp 1120–1128
- Breiman L (2001) Random forests. *Mach Learn* 45(1):5–32. <https://doi.org/10.1023/A:1010933404324>
- Kulesza A, Taskar B (2012) Determinantal point processes for machine learning. *Found Trends Mach Learn* 5(2–3):123–286. <https://doi.org/10.1561/22000000044>
- Macchi O (1975) The coincidence approach to stochastic point processes. *Adv Appl Probab* 7(1):83–122. <https://doi.org/10.2307/1425855>
- Derezinski M, Mahoney MW (2021) Determinantal point processes in randomized numerical linear algebra. *Not Am Math Soc* 68(1):34–45
- Bardenet R, Hardy A (2020) Monte carlo with determinantal point processes. *Ann Appl Probab* 30(1):368–417
- Elfeki M, Couprie C, Riviere M, Elhoseiny M (2019) Gdpp: learning diverse generations using determinantal point processes. In: International conference on machine learning, PMLR, pp 1774–1783
- Derezinski M (2018) Volume sampling for linear regression. UC Santa Cruz Electronic Theses and Dissertations 36
- Derezinski M, Warmuth MK, Hsu D (2018) Leveraged volume sampling for linear regression. In: Proceedings of the 32nd international conference on neural information processing systems. NIPS'18, pp 2510–2519. Curran Associates Inc., Red Hook, NY, USA
- Kulesza A, Taskar B (2011) K-dpps: fixed-size determinantal point processes. In: Proceedings of the 28th international conference on international conference on machine learning. ICML'11, pp 1193–1200. Omnipress, Madison, WI, USA
- Hough JB, Krishnapur M, Peres Y, Virág B (2006) Determinantal Processes and Independence. *Probab Surv* 3(none), 206–229 <https://doi.org/10.1214/154957806000000078>
- Anari N, Oveis Gharan S, Rezaei A (2016) Monte carlo markov chain algorithms for sampling strongly rayleigh distributions and determinantal point processes. In: Feldman V, Rakhlin A, Shamir O (eds) 29th Annual conference on learning theory. proceedings of machine learning research, vol 49, pp 103–115. PMLR, Columbia University, New York, New York, USA. <https://proceedings.mlr.press/v49/anari16.html>
- Li C, Sra S, Jegelka S (2016) Fast mixing markov chains for strongly rayleigh measures, dpps, and constrained sampling. In: Lee DD, Sugiyama M, Luxburg U, Guyon I, Garnett R (eds) Advances in neural information processing systems 29: annual conference on neural information processing systems 2016, December 5–10, 2016, Barcelona, Spain, pp 4188–4196. <https://proceedings.neurips.cc/paper/2016/hash/850af92f8d9903e7a4e0559a98ecc857-Abstract.html>
- Derezinski M, Calandriello D, Valko M (2019) Exact sampling of determinantal point processes with sublinear time preprocessing. In: Wallach H, Larochelle H, Beygelzimer A, Alché-Buc F, Fox E, Garnett R (eds) Advances in neural information processing systems, vol 32. Curran Associates, Inc., Vancouver, Canada. <https://proceedings.neurips.cc/paper/2019/file/fa3060edb66e6ff4507886f9912e1ab9-Paper.pdf>
- Calandriello D, Derezinski M, Valko M (2020) Sampling from a k-dpp without looking at all items. In: Larochelle H, Ranzato M, Hadsell R, Balcan MF, Lin H (eds) Advances in Neural Information Processing Systems, vo 33, pp 6889–6899. Curran Associates, Inc., Online. <https://proceedings.neurips.cc/paper/2020/file/4d410063822cd9be28f86701c0bc3a31-Paper.pdf>
- Gautier G, Polito G, Bardenet R, Valko M (2019) Dppy: Dpp sampling with python. *J Mach Learn Res* 20(180):1–7
- Kazdagli S, Kerenidis I, Kieckbusch J, Teare P (2023) Improved clinical data imputation via classical and quantum determinantal point processes. [arXiv:2303.17893](https://arxiv.org/abs/2303.17893)
- Grinsztajn L, Oyallon E, Varoquaux G (2022) Why do tree-based models still outperform deep learning on typical tabular data? In: Thirty-sixth conference on neural information processing systems datasets and benchmarks track. [https://openreview.net/forum?id=Fp7\\_\\_phQszn](https://openreview.net/forum?id=Fp7__phQszn)
- Benedetti M, Lloyd E, Sack S, Fiorentini M (2019) Parameterized quantum circuits as machine learning models. *Quantum Science and Technology* 4(4):043001. <https://doi.org/10.1088/2058-9565/ab4eb5>, [arXiv:1906.07682](https://arxiv.org/abs/1906.07682) [quant-ph]
- Havlíček V, Córcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, Gambetta JM (2018) Supervised learning with quantum-enhanced feature spaces. *Nature* 567:209–212
- Liu Y, Arunachalam S, Temme K (2021) A rigorous and robust quantum speed-up in supervised machine learning. *Nat Phys* 17(9):1013–1017. <https://doi.org/10.1038/s41567-021-01287-z>, [arXiv:2010.02174](https://arxiv.org/abs/2010.02174) [quant-ph]
- Jia K, Li S, Wen Y, Liu T, Tao D (2019) Orthogonal deep neural networks. *IEEE Trans Pattern Anal Mach Intell*
- Johri S, Debnath S, Mocherla A, Singh A, Prakash A, Kim J, Kerenidis I (2021) Nearest centroid classification on a trapped ion quantum computer. *npj Quantum Information* (to appear), [arXiv:2012.04145](https://arxiv.org/abs/2012.04145)
- Jozsa R, Miyake A (2008) Matchgates and classical simulation of quantum circuits. *Proc Royal Soc A: Math Phys Eng Sci* 464(2100):3089–3106. <https://doi.org/10.1098/rspa.2008.0189>
- Cherrat EA, Kerenidis I, Mathur N, Landman J, Strahm M, Li YY (2022) Quantum vision transformers. [arXiv:2209.08167](https://arxiv.org/abs/2209.08167)
- Buuren S, Groothuis-Oudshoorn K (2011) mice: multivariate imputation by chained equations in r. *J Stat Softw* 45(3):1–67. <https://doi.org/10.18637/jss.v045.i03>
- contributors Q (2021) Qiskit: an open-source framework for quantum computing. <https://doi.org/10.5281/zenodo.2573505>
- Li G, Ding Y, Xie Y (2018) Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices. arXiv e-prints, 1809–02573. <https://doi.org/10.48550/arXiv.1809.02573> [cs.ET]
- Viola L, Lloyd S (1998) Dynamical suppression of decoherence in two state quantum systems. *Phys Rev A* 58, 2733. <https://doi.org/10.1103/PhysRevA.58.2733>, [arXiv:quant-ph/9803057](https://arxiv.org/abs/quant-ph/9803057)
- Ezzell N, Pokharel B, Tewala L, Quiroz G, Lidar DA (2022) Dynamical decoupling for superconducting qubits: a performance survey. [arXiv:2207.03670](https://arxiv.org/abs/2207.03670) [quant-ph]
- Nation PD, Kang H, Sundaresan N, Gambetta JM (2021) Scalable mitigation of measurement errors on quantum computers. *PRX Quantum* 2, 040326 <https://doi.org/10.1103/PRXQuantum.2.040326>, [arXiv:2108.12518](https://arxiv.org/abs/2108.12518) [quant-ph]