

**Prediction by linear regression on a quantum computer**Maria Schuld,<sup>1,\*</sup> Ilya Sinayskiy,<sup>1,2</sup> and Francesco Petruccione<sup>1,2</sup><sup>1</sup>*Quantum Research Group, School of Chemistry and Physics, University of KwaZulu-Natal, Durban 4000, South Africa*<sup>2</sup>*National Institute for Theoretical Physics (NITheP), KwaZulu-Natal 4001, South Africa*

(Received 28 January 2016; revised manuscript received 4 August 2016; published 30 August 2016)

We give an algorithm for prediction on a quantum computer which is based on a linear regression model with least-squares optimization. In contrast to related previous contributions suffering from the problem of reading out the optimal parameters of the fit, our scheme focuses on the machine-learning task of guessing the output corresponding to a new input given examples of data points. Furthermore, we adapt the algorithm to process nonsparse data matrices that can be represented by low-rank approximations, and significantly improve the dependency on its condition number. The prediction result can be accessed through a single-qubit measurement or used for further quantum information processing routines. The algorithm's runtime is logarithmic in the dimension of the input space provided the data is given as quantum information as an input to the routine.

DOI: [10.1103/PhysRevA.94.022342](https://doi.org/10.1103/PhysRevA.94.022342)**I. INTRODUCTION**

A central problem of machine learning is pattern recognition [1], in which a machine is supposed to infer from a set of training data how to map new inputs of the same type to corresponding outputs. An important application lies in image recognition, where example images are presented to the computer together with a class label describing their content, and the task is to generalize from the training data in order to classify a new input image correctly. Nowadays, image recognition by machines can compete with human abilities for selected problems [2]. Pattern recognition plays an increasingly important role in information technology and robotics, but also in fields such as medicine and finance—in short, everywhere a decision has to be derived from data. If the class label in a pattern-recognition task is continuous, one speaks of a *regression* problem, on which we will focus here.

A widely used method adopted from statistics is linear regression, which predicts new labels based on the linear fit of the data. *Learning* in this context refers to the process of estimating optimal parameters for the linear fit by reducing the least-squares error of the data points. In computational terms, this reduces to finding the (pseudo)inverse of a matrix which represents the data, a task that can become very expensive considering the size of datasets in industrial applications today.

Quantum computing established itself as a promising extension to classical computation [3] and was theoretically shown to perform significantly better in selected computational problems. Recently, a number of proposals apply quantum information processing to methods in machine learning [4–8]. One line of approaches [9–11] focuses on the numerous machine-learning methods formulated as matrix inversion problems, which can be tackled by the quantum algorithm to solve linear systems of equations introduced by Harrow, Hassidim, and Lloyd (HHL) [12]. In particular, Wiebe, Braun and Lloyd (WBL) [11] suggested a quantum algorithm for data

fitting based on HHL. The authors approach linear regression from a statistics perspective and their goal is to find a quantum state of which the amplitudes represent the optimal parameters of the linear fit learned from data. They show that this can be done in time logarithmic in the dimension of the input data  $N$  (e.g., the number of pixels in image recognition) for a  $s$  sparse data matrix given in the form of quantum information, and with a sensible dependency on its condition number  $\kappa$  as well as the desired accuracy  $\epsilon$  [roughly  $O(\log N s^3 \kappa^6 \epsilon^{-1})$ ]. They conclude that “[a]lthough the algorithm yields [the state encoding the parameters] efficiently, it may be exponentially expensive to learn via tomography” [11] and shift their focus on a routine which estimates the quality of a fit. Another attempt in the same direction, but based on the singular-value decomposition that we will also use here, is given in [13] and requires a low-rank data matrix instead of the sparseness condition. The author proposes various algorithms to analyze the parameter vector using known quantum routines.

In this article, we want to add an important missing piece to the literature of quantum machine learning by reapproaching the rich problem of linear regression on a quantum computer from a machine-learning perspective, in which the focus lies on the prediction of new inputs based on a dataset. We demonstrate a routine for prediction which is roughly in  $O(\log N \kappa^2 \epsilon^{-3})$ , provided the information is given encoded into a quantum state. Using the techniques of quantum principal component analysis proposed in [14], our algorithm does not require the data matrix  $\mathbf{X}$  (defined below) to be sparse. Instead, we only need to be able to represent  $\mathbf{X}^T \mathbf{X}$  by a low-rank approximation, meaning that it is dominated by a few large eigenvalues.

**II. PREDICTION WITH LINEAR REGRESSION**

The problem of supervised pattern recognition can be described as follows. Given a set of data points  $\mathcal{D} = \{(\mathbf{x}^{(1)}, y^{(1)}), \dots, (\mathbf{x}^{(M)}, y^{(M)})\}$  with  $\mathbf{x}^{(m)} \in \mathbb{R}^N$  and  $y^{(m)} \in \mathbb{R}$  for  $m = 1, \dots, M$ , the goal is to predict the output  $\tilde{y} \in \mathbb{R}$  to a new input  $\tilde{\mathbf{x}} \in \mathbb{R}^N$ . A linear model proposes a map,

$$f(\mathbf{x}, \mathbf{w}) = \mathbf{x}^T \mathbf{w}, \quad (1)$$

\*schuld@ukzn.ac.za




FIG. 1. Left: Illustration of linear regression given training data points  $(x^{(m)}, y^{(m)})$ . In contrast to statistical analysis, the task in machine learning is not to determine the fit itself, but to guess the class label  $\tilde{y}$  of a new input  $\tilde{x}$  based on the fitted function. Right: The inputs can be generated by a nonlinear map  $\mathbf{x} \rightarrow \tilde{\mathbf{x}}$  into a higher dimensional space, dramatically increasing the power of linear regression to fit nonlinear functions.

where  $\mathbf{w} = (w_1, \dots, w_N)^T \in \mathbb{R}^N$  is a vector of parameters that have to be “fitted” or learned from the data. (Note that we can omit a potential bias  $w_0$  by including it in  $\mathbf{w}$  and extending  $\mathbf{x}$  by a corresponding entry  $x_0 = 1$ .) It is important to emphasize that the inputs  $\mathbf{x}$  can be generated by a nonlinear map from some original space in which the data is given into  $\mathbb{R}^N$ , which enables linear regression to fit nonlinear functions (e.g., in polynomial curve fitting; see Fig. 1).

The most common method to estimate the parameters (yielding the minimum variance with zero bias for the estimate) is to minimize the least-squares error between the predicted values by the model,  $f(\mathbf{x}^{(m)}, \mathbf{w})$ , and the actual target outputs  $y^{(m)}$  for each data point,

$$\min_{\mathbf{w}} \sum_{m=1}^M [f(\mathbf{x}^{(m)}, \mathbf{w}) - y^{(m)}]^2.$$

As is apparent from this objective function, we focus on unregularized linear regression here, and the matter of regularization is an open question in quantum machine-learning research [15].

Introducing the compact matrix notation  $\mathbf{X} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)})^T$  (which we refer to as the data matrix containing all training inputs as rows) and the target vector  $\mathbf{y} = (y^{(1)}, \dots, y^{(M)})^T$ , the least-squares error becomes  $|\mathbf{X}\mathbf{w} - \mathbf{y}|^2$ , and it is a widely used result [16] that the solution to the least-squares problem is given by

$$\mathbf{w} = \mathbf{X}^+ \mathbf{y},$$

where  $\mathbf{X}^+$  is the Moore-Penrose pseudoinverse of  $\mathbf{X}$ . Formally, one can write  $\mathbf{X}^+ = (\mathbf{X}^\dagger \mathbf{X})^{-1} \mathbf{X}^\dagger$  if  $(\mathbf{X}^\dagger \mathbf{X})^{-1}$  exists, and the approach in WBL solves the resulting equation via quantum computation. We will use a different form of the Moore-Penrose pseudoinverse based on the reduced singular-value decomposition  $\mathbf{X} = \mathbf{U} \Sigma \mathbf{V}^\dagger$  [17] in order to allow for  $\mathbf{X}$  to be non-square. Here,  $\Sigma$  is a diagonal matrix containing the real singular values  $\sigma_1, \sigma_2, \dots, \sigma_R > 0$  and the  $r$ th orthogonal column of  $\mathbf{U} \in \mathbb{R}^{M \times R}$  ( $\mathbf{V} \in \mathbb{R}^{R \times R}$ ) is the  $r$ th left (right) eigenvector  $\mathbf{v}_r$  ( $\mathbf{u}_r$ ) to the singular value  $\sigma_r$ . As opposed to the eigenvalue decomposition, the singular-value decomposition can always be found, and in particular for nonsquare matrices [18]. The pseudoinverse can then be defined as  $\mathbf{X}^+ = \mathbf{V} \Sigma^{-1} \mathbf{U}^\dagger$ . The singular values of  $\mathbf{X}$  are at the same time the square roots of the nonzero eigenvalues  $\{\sqrt{\lambda_1}, \dots, \sqrt{\lambda_R}\}$  of  $\mathbf{X}^\dagger \mathbf{X}$  and

$\mathbf{X} \mathbf{X}^\dagger$ , and the left and right singular vectors are their respective eigenvectors; we make use of this fact below.

Using this mathematical background in the case of real spaces, as well as the convenient alternative formulation  $\mathbf{X} = \sum_r \sigma_r \mathbf{u}_r \mathbf{v}_r^T$  [18], the least-squares solution can be written as

$$\mathbf{w} = \sum_{r=1}^R \sigma_r^{-1} \mathbf{v}_r \mathbf{u}_r^T \mathbf{y}. \quad (2)$$

Moreover, according to Eq. (1), the output  $\tilde{y}$  for a new input  $\tilde{\mathbf{x}}$  is given by

$$\tilde{y} = \sum_{r=1}^R \sigma_r^{-1} \tilde{\mathbf{x}}^T \mathbf{v}_r \mathbf{u}_r^T \mathbf{y}. \quad (3)$$

The desired output (“prediction”) of the pattern classification algorithm is thus a single scalar value, which will circumvent the problem of expensive state tomography notorious for quantum machine-learning algorithms [19]. The problem stated here can be solved in polynomial time  $O(N^d)$  on a classical computer, where for the best current algorithms  $2 \leq d \leq 3$ . The aim of this work is to find a quantum algorithm that can reproduce the result from Eq. (3) more efficiently assuming that the data is given as quantum information in the form specified below.

### III. THE QUANTUM LINEAR REGRESSION ALGORITHM

The quantum linear regression algorithm is based on the method of encoding classical information such as a  $2^n$  dimensional vector  $\mathbf{a} = (a_0, \dots, a_{2^n-1})^T$  into the  $2^n$  amplitudes  $a_0, \dots, a_{2^n-1}$  of a  $n$ -qubit quantum system,  $|\psi_{\mathbf{a}}\rangle = \sum_{i=0}^{2^n-1} a_i |i\rangle$ , where  $\{|i\rangle\}$  is a convenient notation for the computational basis  $\{|0 \dots 0\rangle \triangleq |0\rangle, \dots, |1 \dots 1\rangle \triangleq |2^n - 1\rangle\}$ . In other words, the probabilistic description of a set of two-level quantum systems is used to store and manipulate classical information. Here we will refer to this method as *amplitude encoding* and denote every such quantum state by a  $\psi$  with a subscript referring to the classical vector it encodes. (In other words,  $|a\rangle$  is a quantum state where some mathematical object  $a$  is encoded into the basis state in a way that is to specify in detail, while  $|\psi_{\mathbf{a}}\rangle$  is a quantum state representing the real vector  $\mathbf{a}$  via amplitude encoding.)

The general idea of the quantum pattern-recognition algorithm is to create a quantum state representing the data matrix  $\mathbf{X} = \sum_r \sigma_r \mathbf{u}_r \mathbf{v}_r^T$  via amplitude encoding (step A) and use tricks from [12,14] to invert the unknown singular values efficiently (steps B and C). In step D, quantum state representations of  $\mathbf{y}, \tilde{\mathbf{x}}$  are used to write the desired prediction from Eq. (3) into the off-diagonal elements of an ancilla qubit, where it can be read out by a simple  $\sigma_x, \sigma_y$  measurement. Details regarding the computational complexity will be considered further below.

#### A. State preparation

The quantum algorithm takes copies of the quantum states representing each of the objects  $\mathbf{X}, \mathbf{y}$ , and  $\tilde{\mathbf{x}}$  from above in

amplitude encoding,

$$|\psi_{\mathbf{X}}\rangle = \sum_{j=0}^{N-1} \sum_{m=0}^{M-1} x_j^{(m)} |j\rangle |m\rangle, \quad (4)$$

$$|\psi_{\mathbf{Y}}\rangle = \sum_{\mu=0}^{M-1} y^{(\mu)} |\mu\rangle, \quad (5)$$

$$|\psi_{\tilde{\mathbf{X}}}\rangle = \sum_{\gamma=0}^{N-1} \tilde{x}_{\gamma} |\gamma\rangle, \quad (6)$$

with  $\sum_{m,j} |x_j^{(m)}|^2 = \sum_{\mu} |y^{(\mu)}|^2 = \sum_{\gamma} |\tilde{x}_{\gamma}|^2 = 1$ . Note that the algorithm thus works with normalized data and the results have to be rescaled accordingly. Using the Gram-Schmidt decomposition, we can formally write

$$|\psi_{\mathbf{X}}\rangle = \sum_{r=1}^R \sigma^r \sum_{j=1}^J v_j^r |j\rangle \sum_{m=1}^M u_m^r |m\rangle,$$

or, in short,

$$|\psi_{\mathbf{X}}\rangle = \sum_{r=1}^R \sigma^r |\psi_{\mathbf{V}^r}\rangle |\psi_{\mathbf{U}^r}\rangle. \quad (7)$$

Here,  $|\psi_{\mathbf{V}^r}\rangle = \sum_{j=1}^J v_j^r |j\rangle$  and  $|\psi_{\mathbf{U}^r}\rangle = \sum_{m=1}^M u_m^r |m\rangle$  are quantum states representing the orthogonal sets of left and right singular vectors of  $\mathbf{X}$  (and, at the same time, the eigenvectors of  $\mathbf{X}^\dagger \mathbf{X}$  as well as  $\mathbf{X} \mathbf{X}^\dagger$ ) via amplitude encoding, and  $\sigma^r$  are the corresponding singular values. The number of qubits needed to construct states (4)–(6) are  $\lceil \log N \rceil + \lceil \log M \rceil$ ,  $\lceil \log M \rceil$ , and  $\lceil \log N \rceil$ , respectively, and the aim of the following algorithm is to remain linear in the number of qubits or logarithmic in the problem size. The states (4)–(6) can be understood as the result of previous quantum computation or simulation, as argued in [11]. If given a ‘‘classical’’ dataset, techniques for the efficient preparation of arbitrary initial quantum states are a nontrivial and still controversially discussed topic, although some ideas have been brought forward of how to prepare certain states linear in the number of qubits [20] or via quantum random access memories [21]. Clader *et al.* [22] show a way to prepare a quantum state efficiently for the HHL algorithm if an oracle is given that conditioned on a register in uniform superposition ‘‘loads’’ the vector entries into an entangled register, ‘‘writes’’ these entries into the amplitude of an ancilla, and performs a conditional measurement to prepare the desired state in amplitude encoding. This conditional measurement has a high probability to succeed for relatively uniform entries, while sparse amplitudes underlie the Grover bound and cannot be prepared in time logarithmic in the dimension of the classical vector we want to encode into a quantum state. We acknowledge that overall, the question of state preparation is an outstanding challenge for quantum machine-learning algorithmic design.

### B. Extracting the singular values

In order to transform Eq. (7) into a ‘‘quantum representation’’ of the result (3), we need to first invert the singular values of  $\mathbf{X}$ . For this, we will ‘‘extract’’ the eigenvalues  $\lambda_r$

of  $\mathbf{X}^\dagger \mathbf{X}$  to eigenvectors  $\mathbf{v}_r$  and use the inversion procedure from [12] in the following step. In order to access the eigenvalues, we use copies of the state (4) in which on the level of description we ignore the  $|m\rangle$  register in order to obtain a mixed state  $\rho_{\mathbf{X}^\dagger \mathbf{X}} = \text{tr}_m\{|\psi_{\mathbf{X}}\rangle\langle\psi_{\mathbf{X}}|\}$  which represents the positive Hermitian matrix  $\mathbf{X}^\dagger \mathbf{X}$ ,

$$\rho_{\mathbf{X}^\dagger \mathbf{X}} = \sum_{j,j'=1}^N \sum_{m=1}^M x_j^{(m)} x_{j'}^{(m)*} |j\rangle\langle j'|.$$

From here, we use the ideas of *quantum principal component analysis* introduced in [14] to ‘‘apply’’  $\rho_{\mathbf{X}^\dagger \mathbf{X}}$  to  $|\psi_{\mathbf{X}}\rangle$ , resulting in

$$\sum_{k=0}^K |k \Delta t\rangle\langle k \Delta t| \otimes e^{-ik\rho_{\mathbf{X}^\dagger \mathbf{X}} \Delta t} |\psi_{\mathbf{X}}\rangle\langle\psi_{\mathbf{X}}| e^{ik\rho_{\mathbf{X}^\dagger \mathbf{X}} \Delta t},$$

for some large  $K$ . As outlined in [14], the quantum phase-estimation algorithm results in

$$\sum_{r=1}^R \sigma^r |\psi_{\mathbf{V}^r}\rangle |\psi_{\mathbf{U}^r}\rangle |\lambda^r\rangle,$$

in which the eigenvalues  $\lambda^r = (\sigma^r)^2$  of  $\rho_{\mathbf{X}^\dagger \mathbf{X}}$  are approximately encoded in the  $\tau$  qubits of an extra third register that was initially in the ground state.

### C. Inverting the singular values

Adding an extra qubit and rotating it conditional on the eigenvalue register [12] then yields

$$\sum_{r=1}^R \sigma^r |\psi_{\mathbf{V}^r}\rangle |\psi_{\mathbf{U}^r}\rangle |\lambda^r\rangle \left[ \sqrt{1 - \left(\frac{c}{\lambda^r}\right)^2} |0\rangle + \frac{c}{\lambda^r} |1\rangle \right]. \quad (8)$$

The constant  $c$  is chosen so that the inverse eigenvalues are not larger than 1, which is given if it is smaller than the smallest nonzero eigenvalue  $\lambda^{\min}$  of  $\mathbf{X}^\dagger \mathbf{X}$ , or, equivalently, the smallest nonzero squared singular value  $(\sigma^{\min})^2$  of  $\mathbf{X}$ . We perform a conditional measurement on the ancilla qubit, only continuing the algorithm (‘‘accepting’’) if the ancilla is in state  $|1\rangle$  (otherwise the entire procedure has to be repeated). Note that amplitude amplification can boost the probability of accepting quadratically, as discussed in the runtime analysis below. Uncomputing and discarding the eigenvalue register results in

$$|\psi_1\rangle := \frac{1}{\sqrt{p(1)}} \sum_{r=1}^R \frac{c}{\sigma^r} |\psi_{\mathbf{V}^r}\rangle |\psi_{\mathbf{U}^r}\rangle, \quad (9)$$

where the probability of acceptance is given by  $p(1) = \sum_r \left|\frac{c}{\lambda^r}\right|^2$ .

### D. Executing the inner products

The goal of the last step is to write the desired result  $\sum_{r=1}^R (\sigma^r)^{-1} \langle\psi_{\tilde{\mathbf{X}}}|v^r\rangle \langle\psi_{\mathbf{Y}}|u^r\rangle$  into selected entries of an ancilla’s single-qubit density matrix, from which it can be accessed by

a simple measurement. Consider the result (9) of the previous step, as well as  $|\psi_2\rangle = |\psi_y\rangle|\psi_{\bar{x}}\rangle$  from (5) and (6). The inner product cannot be extracted by a direct SWAP test between  $|\psi_1\rangle$  and  $|\psi_2\rangle$  if we want the output domain to extend into the negative numbers, since this would lead to an acceptance probability of  $\frac{1}{2} + \frac{1}{2}|\langle\psi_1|\psi_2\rangle|^2$ , which is ambiguous in the sign of the desired result  $\langle\psi_1|\psi_2\rangle$ . The following “prediction routine” circumvents these problems: Conditionally prepare the two states so that they are entangled with an ancilla qubit,

$$1/\sqrt{2}(|\psi_1\rangle|0\rangle + |\psi_2\rangle|1\rangle), \quad (10)$$

and trace out all registers except from the ancilla. The off-diagonal elements  $\rho_{12}, \rho_{21}$  of the ancilla’s density matrix read

$$\frac{c}{2\sqrt{p(1)}} \sum_r (\sigma^r)^{-1} \sum_j v_j^r \tilde{x}_j \sum_m u_m^r y^{(m')},$$

and contain the desired result (3) up to a known normalization factor. Conditionally preparing (10) requires us to execute the entire algorithm including state preparation conditioned on the state of the ancilla qubit and might not be easy to implement. In that case, one can adapt the algorithm so that the  $|0\dots 0\rangle$  basis state in  $|\psi_1\rangle$  and  $|\psi_2\rangle$  is “excluded” from all operations and remains with a constant amplitude  $\frac{1}{\sqrt{2}}$  throughout the algorithm (while the other  $2^n - 1$  amplitudes are renormalized accordingly). This prepares states of the general form  $|a\rangle = \frac{1}{\sqrt{2}}(|0\dots 0\rangle + \sum_{i=1}^N a_i|i\rangle)$ ,  $|b\rangle = \frac{1}{\sqrt{2}}(|0\dots 0\rangle + \sum_{j=1}^N b_j|j\rangle)$ . A common SWAP test effectively shifts the inner product by 1/2 and thus reveals  $|\langle a|b\rangle|^2 = |\frac{1}{2} + \frac{1}{2}\sum_{i=1}^N a_i b_i|^2$  from which the sign of  $\sum_{i=1}^N a_i b_i$  can be extracted.

#### IV. RUNTIME ANALYSIS

According to [14], one needs temporal resources  $t = k\Delta t$  in  $O(\log N)$  and of the order of  $O(\epsilon^{-3})$  copies of  $\rho_{\mathbf{X}^\dagger\mathbf{X}}$  to “exponentiate” a density matrix in step B, where  $\epsilon$  is the error and  $N$  is the dimension of the inputs in our data set. The method requires the density matrix  $\rho_{\mathbf{X}^\dagger\mathbf{X}}$  to be close to a low-rank approximation which is dominated by a few large eigenvalues in order to maintain the exponential speedup. In general, it takes time  $t = O(1/\delta)$  to simulate  $e^{iHt}$  for a Hamiltonian  $H$  up to error  $\delta$  [12], and with the trick from [14] it takes time  $t^2$  to do the same for  $e^{i\rho t}$ . This means that if we want to resolve relatively uniform eigenvalues of the order of  $1/N$ , time grows quadratically with  $N$  and the exponential speedup is lost. Hence, the method is only efficient if the density matrix is dominated by a few large eigenvalues.

The singular-value inversion procedure in step C determines the runtime’s dependency on the condition number of  $\mathbf{X}$ ,  $\kappa = \sigma_{\max}(\sigma_{\min})^{-1}$ . The probability to measure the ancilla in the excited state is

$$p(1) = \sum_r \left| \frac{c}{\lambda^r} \right|^2 \leq R \left| \frac{\lambda_{\min}}{\lambda_{\max}} \right|^2 = R\kappa^{-4},$$

which means one needs, on average, less than  $\kappa^4$  tries to accept the conditional measurement. Amplitude amplification as in [9,12] reduces this to a factor of  $O(\kappa^2)$  in the runtime, which can become significant for matrices that are close to being singular. The condition number is a measure for how easy it is to invert a matrix, and it is no surprise that the runtime of the quantum algorithm also depends on it. The amended SWAP routine in step D is also linear in the number of qubits, and the final measurement only accounts for a constant factor. The upper bound for the runtime can thus be roughly estimated as  $O(\log N\kappa^2\epsilon^{-3})$  if we have sufficient copies of  $\rho_{\mathbf{X}^\dagger\mathbf{X}}$  available, which is required to be close to a low-rank matrix. Remember that this does not include the costs of quantum state preparation in case the algorithm processes classical information. Compared to the previous result in WBL, this is an improvement of a factor  $\kappa^{-4}$ , whereas the dependence on the accuracy is worse by a factor  $\epsilon^{-2}$ . However, in contrast to other proposals, our algorithm tackles the problem of pattern recognition or prediction. Furthermore, we can apply our algorithm efficiently to nonsparse, but low-rank approximations of the matrix  $\mathbf{X}^\dagger\mathbf{X}$ .

#### V. CONCLUSION

We described an algorithm for a universal quantum computer to implement a linear regression model for supervised pattern recognition. The quantum algorithm reproduces the prediction result of a classical linear regression method with (unregularized) least-squares optimization, thereby covering an important area of machine learning. It runs in time logarithmic in the dimension  $N$  of the feature vectors as well as independent of the size of the training set if the inputs are given as quantum information. Instead of requiring the matrix containing the training inputs,  $\mathbf{X}$ , to be sparse, it merely needs  $\mathbf{X}^\dagger\mathbf{X}$  to be representable by a low-rank approximation. One can furthermore transform the input data by a nonlinear feature map known as the “kernel trick” (discussed in [9] for polynomial kernels) to increase the potential power of the method. The application of different kernels as well as the question of how to include regularization terms is still open for further research. The sensitive dependency on the accuracy as well as the unresolved problem of state preparation (which appears in any of the numerous quantum algorithms encoding classical information into the amplitudes of quantum states) illustrate how carefully one needs to treat “magic” exponential speedups for pattern recognition. However, as demonstrated here, quantum information can make a contribution to certain problems of machine learning, promising further fruitful results in the emerging discipline of quantum machine learning.

#### ACKNOWLEDGMENTS

The authors would like to thank Patrick Rebentrost for fruitful discussions. This work is based upon research supported by the South African Research Chair Initiative of the Department of Science and Technology and National Research Foundation, as well as the European Commission.

- [1] Christopher M. Bishop, *Pattern Recognition and Machine Learning* (Springer, New York, 2006), Vol. 1.
- [2] See, for example, Yann LeCun's list of the performance of different machine-learning methods on the MNIST handwritten digits database [yann.lecun.com/exdb/mnist](http://yann.lecun.com/exdb/mnist) (unpublished).
- [3] S. Arora and B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, Cambridge, 2009).
- [4] N. Wiebe, A. Kapoor, and K. Svore, Quantum algorithms for nearest-neighbor methods for supervised and unsupervised learning, *Quantum Inf. Comput.* **15**, 0318 (2014) [[arXiv:1401.2142v2](https://arxiv.org/abs/1401.2142v2)].
- [5] N. Wiebe, A. Kapoor, and K. Svore, Quantum deep learning, [arXiv:1412.3489v1](https://arxiv.org/abs/1412.3489v1).
- [6] B. O'Gorman, R. Babbush, A. Perdomo-Ortiz, A. Aspuru-Guzik, and V. Smelyanskiy, Bayesian network structure learning using quantum annealing, *Eur. Phys. J. Spec. Top.* **224**, 163 (2015).
- [7] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum algorithms for supervised and unsupervised machine learning, [arXiv:1307.0411](https://arxiv.org/abs/1307.0411).
- [8] M. Schuld, I. Sinayskiy, and F. Petruccione, How to simulate a perceptron using quantum circuits, *Phys. Lett. A* **379**, 660 (2015).
- [9] P. Rebentrost, M. Mohseni, and S. Lloyd, Quantum Support Vector Machine for Big Data Classification, *Phys. Rev. Lett.* **113**, 130503 (2014).
- [10] Z. Zhao, J. K. Fitzsimons, and J. F. Fitzsimons, Quantum assisted Gaussian process regression, [arXiv:1512.03929](https://arxiv.org/abs/1512.03929).
- [11] N. Wiebe, D. Braun, and S. Lloyd, Quantum Algorithm for Data Fitting, *Phys. Rev. Lett.* **109**, 050505 (2012).
- [12] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [13] G. Wang, Quantum algorithms for curve fitting, [arXiv:1402.0660](https://arxiv.org/abs/1402.0660).
- [14] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum principal component analysis, *Nat. Phys.* **10**, 631 (2014).
- [15] In machine-learning applications, a regularization term favoring small parameters is usually added to avoid overfitting. It is an interesting question for further research whether the normalization condition for quantum states (i.e., containing the parameters as in WBL) can be harvested to introduce an effective regularization.
- [16] G. H. Golub and C. Reinsch, Singular value decomposition and least squares solutions, *Numer. Math.* **14**, 403 (1970).
- [17] L. Hogben, *Handbook of Linear Algebra* (CRC, Boca Raton, FL, 2006).
- [18] L. N. Trefethen and D. Bau III, *Numerical Linear Algebra* (Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997), Vol. 50.
- [19] S. Aaronson, Read the fine print, *Nat. Phys.* **11**, 291 (2015).
- [20] L. Grover and T. Rudolph, Creating superpositions that correspond to efficiently integrable probability distributions, [arXiv:quant-ph/0208112v1](https://arxiv.org/abs/quant-ph/0208112v1).
- [21] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Random Access Memory, *Phys. Rev. Lett.* **100**, 160501 (2008).
- [22] B. D. Clader, B. C. Jacobs, and C. R. Sprouse, Preconditioned Quantum Linear System Algorithm, *Phys. Rev. Lett.* **110**, 250504 (2013).