



Finding eigenvectors with a quantum variational algorithm

Juan Carlos Garcia-Escartin¹

Received: 26 February 2024 / Accepted: 10 June 2024 / Published online: 25 June 2024
© The Author(s) 2024

Abstract

This paper presents a hybrid variational quantum algorithm that finds a random eigenvector of a unitary matrix with a known quantum circuit. The algorithm is based on the SWAP test on trial states generated by a parametrized quantum circuit. The eigenvector is described by a compact set of classical parameters that can be used to reproduce the found approximation to the eigenstate on demand. This variational eigenvector finder can be adapted to solve the generalized eigenvalue problem, to find the eigenvectors of normal matrices and to perform quantum principal component analysis on unknown input mixed states. These algorithms can all be run with low-depth quantum circuits, suitable for an efficient implementation on noisy intermediate-scale quantum computers and, with some restrictions, on linear optical systems. In full-scale quantum computers, where there might be optimization problems due to barren plateaus in larger systems, the proposed algorithms can be used as a primitive to boost known quantum algorithms. Limitations and potential applications are discussed.

Keywords Quantum algorithms · Variational algorithms · Hybrid quantum algorithms · Eigenvectors · Quantum principal component analysis · SWAP test

1 Introduction: Noisy intermediate-scale quantum computers and hybrid algorithms

Scalable general-purpose quantum computers could run algorithms that are more efficient than any classical alternative [1, 2]. However, at the present moment, the available technology is restricted to computers with a moderate number of qubits with a varying degree of noise. These computers are usually dubbed noisy intermediate-scale quantum computers, NISQC [3, 4].

In this scenario, there has been a growing interest on hybrid quantum-classical algorithms [5–7] where part of the work is shifted to a classical computer. In most cases there is a continuous feedback between the classical and the quantum computer,

✉ Juan Carlos Garcia-Escartin
juagar@tel.uva.es

¹ Dpto. Teoría de la Señal e Ing. Telemática, Universidad de Valladolid, Paseo Belén no 15, 47011 Valladolid, Spain

which has quantum circuits that are a function of a few parameters which are updated in the classical part of the algorithm according to the results of the quantum measurements from previous stages.

These algorithms are particularly useful in problems dealing with large state spaces which can be sampled quickly on a quantum computer without the need to explicitly write out the whole state, which has a size that grows exponentially with the size of the problem.

A good example is quantum simulation, one of the most promising applications of quantum computing. For instance, in molecular simulation, the variational quantum eigensolver can efficiently search for a particular state in the Hilbert space of n qubits, which grows exponentially with the size of the problem, and still achieve reasonable results even in the presence of noise [8–10].

This approach based on using quantum circuits with a classical control also plays a key role in quantum machine learning [11–15].

In this paper, I present an application of these concepts to the search for an eigenvalue of a known unitary which can be generalized to different kinds of matrices with some restrictions. The algorithm is based on a modified version the SWAP test for state comparison and can be realized with circuits of a low depth (with few consecutive elementary gates). We only require a state preparation phase, a circuit for U and one CNOT gate and one Hadamard gate for each qubit on which U acts. At the end of the process, we obtain a list of classical parameters that can produce a good approximation to a random eigenvector of the desired matrix using a known classically controlled quantum circuit.

The same method can be adapted to perform quantum principal component analysis of an unknown mixed state, which can be cast as a problem of finding the eigenvectors of a Hermitian matrix.

The paper starts with a brief reminder of the SWAP test for quantum state comparison and some hardware-efficient optimizations in Sect. 2, followed by an explanation of the notation and methods of hybrid algorithms in Sect. 3.


The core quantum variational eigenvector finder algorithm is presented in Sect. 4, including modifications to solve the generalized eigenvalue problem (Sect. 4.1) and to find the eigenstates of general normal matrices instead of just unitaries (Sect. 4.2). Section 5 shows how to adapt the variational eigenvector finder for quantum principal component analysis.

Section 6 proposes an alternative implementation in optical systems instead of the usual quantum circuit model. Section 7 compares the proposed quantum variational algorithms to similar variational algorithms for the same problems commenting the advantages of each method. Finally, Sect. 8 discusses the strong and weak points of the variational approach for eigenvector determination and comments some potential applications in the short and long terms.

2 The SWAP test

While it is impossible to tell apart with certainty two quantum states which are not orthogonal, there are methods that can tell us if they are equal or not with a certain

Fig. 1 Quantum circuit realizing the SWAP test. The crossed lines represent qubit registers with n qubits



probability. The SWAP test [16, 17] permits a simple comparison. We start with two arbitrary n qubit states $|\psi\rangle$ and $|\phi\rangle$ and an ancillary qubit initially in the state $|0\rangle$. The quantum circuit corresponding to the test is shown in Fig. 1.

The evolution in that circuit is

$$\begin{aligned}
 &|0\rangle |\phi\rangle |\psi\rangle \xrightarrow{H} \frac{|0\rangle + |1\rangle}{\sqrt{2}} |\phi\rangle |\psi\rangle \\
 &\xrightarrow{CSWAP} \frac{1}{\sqrt{2}} (|0\rangle |\phi\rangle |\psi\rangle + |1\rangle |\psi\rangle |\phi\rangle) \xrightarrow{H} \quad (1) \\
 &|\Phi_{out}\rangle = \frac{|0\rangle (|\phi\rangle |\psi\rangle + |\psi\rangle |\phi\rangle) + |1\rangle (|\phi\rangle |\psi\rangle - |\psi\rangle |\phi\rangle)}{2}.
 \end{aligned}$$

If we measure the ancillary qubit after the last Hadamard gate, the probability of finding the state $|0\rangle$ is associated with the eigenvalue 1 of the observable

$$Z = |0\rangle \langle 0| - |1\rangle \langle 1| = P_0 - P_1 \quad (2)$$

for the projectors to the $|0\rangle$ and $|1\rangle$ states P_0 and P_1 . For the output state $|\Phi_{out}\rangle$ the probability of measuring the $|0\rangle$ state in the ancillary qubit is

$$P(0) = \langle \Phi_{out} | P_0 | \Phi_{out} \rangle = \frac{1 + |\langle \psi | \phi \rangle|^2}{2} \quad (3)$$

and the average value of the Z observable for the first qubit becomes

$$\langle Z \rangle = |\langle \psi | \phi \rangle|^2. \quad (4)$$

For two identical states, the SWAP operation

$$SWAP |\psi\rangle |\phi\rangle = |\phi\rangle |\psi\rangle \quad (5)$$

does not change the state and the value of the control qubit is irrelevant. The two Hadamard gates on the ancillary qubit cancel and the whole circuit is equivalent to the identity. We can see from Eq. (1) that when both states are the same $|\psi\rangle = |\phi\rangle$, there is a destructive interference between the terms entangled to the ancillary qubit state $|1\rangle$. The result from the measurement is always $|0\rangle$ and whenever we get that result we say the SWAP test has been passed.

However, if the states are different, each part of the uniform superposition of the ancillary qubit becomes entangled to a different state. In that case, there is always a probability greater than zero of measuring the state $|1\rangle$. In a noiseless system, failing

to pass the test shows with certainty the input states are different. The probability of finding a $|1\rangle$ ancillary qubit on measurement grows as the input states become more different. If we have k copies of the same states and repeat the test k times, even for very close states with an overlap $|\langle\psi|\phi\rangle| = 1 - \epsilon$ with $\epsilon \ll 1$, $P(0) = \frac{1+(1-\epsilon)^2}{2} \approx (1 - \epsilon)$ and we have a probability of passing the test $P(0)^k \approx 1 - k\epsilon$. With enough repetitions, if we never fail the test, we can be confident that the two input states are the same or, at least, very close.

As expected, global phases do not change the result of the test. Two inputs $|\psi\rangle$ and $e^{i\Phi}|\psi\rangle$, for an arbitrary phase Φ , will act exactly in the same way under the test. This is the desired operation. While a larger circuit with a reference can distinguish between these two states, if these inputs were isolated states, there is no physical measurement that could tell them apart.

The comparison in the SWAP test is also valid for general inputs with mixed states. We consider an ensemble given by the density matrix

$$\rho = \sum_i^m p_i |\psi_i\rangle \langle\psi_i| \tag{6}$$

where we have a statistical mixture of finding a pure state $|\psi_i\rangle$ with a probability p_i for each. For two mixed states with density matrices ρ and σ the probability of passing the test becomes:

$$P(0) = \frac{1 + \text{tr}(\rho\sigma)}{2} \tag{7}$$

where $\text{tr}(\rho\sigma)$ replaces $|\langle\psi|\phi\rangle|^2$ [18]. When both states are pure states with $\rho = |\psi\rangle \langle\psi|$ and $\sigma = |\phi\rangle \langle\phi|$ we recover the original overlap.

2.1 The destructive SWAP test

The depth of the SWAP test circuit can be reduced if we notice that, after the test, the original inputs become entangled and they are discarded. The total number of elementary gates needed to perform a SWAP test can be reduced to $2n$ for n qubits. Most importantly, we only need two stages.

This is the destructive SWAP test explained in [19], which was later rediscovered by an AI driven search [20]. We can check the test is valid starting from the SWAP circuit for two qubits shown in Fig. 2a. The SWAP operation is performed with three CNOT gates. A sequence of 3 CNOT gates can perform classical XOR swapping [21] for two bits in the computational basis $x, y \in \mathbb{F}_2$ so that

$$\begin{aligned} |x\rangle |y\rangle &\xrightarrow{\text{CNOT}(2,1)} |x \oplus y\rangle |y\rangle \xrightarrow{\text{CNOT}(1,2)} |x \oplus y\rangle |y \oplus x \oplus y\rangle = |x \oplus y\rangle |x\rangle \\ &\xrightarrow{\text{CNOT}(2,1)} |x \oplus y \oplus x\rangle |x\rangle = |y\rangle |x\rangle \end{aligned} \tag{8}$$




Fig. 2 Equivalent circuits for the SWAP test on single qubit states

where $CNOT(i, j)$ is controlled by the i th qubit and has the j th qubit as target. The result follows from two properties the XOR operation \oplus (equivalent to modulo 2 addition): $x \oplus x = 0$ and $x \oplus 0 = x$.

If the middle CNOT is replaced by a Toffoli gate controlled by the ancillary qubit in the SWAP test, we have a controlled SWAP operation. When the ancilla has a state $|1\rangle$ the middle Toffoli gate is equivalent to a CNOT gate on the inputs and we recover the 3 CNOTs of the SWAP circuit. When it is $|0\rangle$ the two consecutive CNOT gates cancel each other and we have the identity.

The NOT operation can be decomposed as $X = HZH$ and we can convert CNOT gates into CZ gates surrounded by two Hadamard gates. The Toffoli gate becomes a doubly controlled Z gate with operation $CCZ|x\rangle|y\rangle|z\rangle = (-1)^{x \cdot y \cdot z}|x\rangle|y\rangle|z\rangle$ where the sign change only happens if all the qubits have a 1. In that sense, any of the qubits can be considered to be the target. Putting all together, we can find a series of equivalent circuits shown in Fig. 2b–d.

Now, starting from Fig. 2d, we can notice that changes on the input qubits after the Toffoli gate do not affect the state of the ancillary qubit. We might just as well measure both qubits right after the Toffoli. Finally, there is a further simplification if we notice that the ancillary qubit will result in a $|0\rangle$ unless both control qubits from the inputs are in state $|1\rangle$. The result of measuring the two input qubits at this point perfectly determines the state of the ancillary qubit, which can be removed. If we measure in the computational basis (observable Z), we can associate a measurement finding a $|0\rangle$ state to a bit 0 (Z 's eigenvalue 1) and finding state $|1\rangle$ to a bit 1 (Z 's eigenvalue -1). The logical AND of the measurement results reproduces the result of measuring the ancillary qubit. A measurement on the input qubits and classical computation (the AND operation) return the same statistics as measuring the Z observable for the ancilla.

The same reasoning can be applied to multiple qubits. In that case, the SWAP operation has the same decomposition for each pair of qubits of the multiqubit state. We group the first qubit of $|\phi\rangle$, q_ϕ^1 , with the first qubit of $|\psi\rangle$, q_ψ^1 , and continue forming all the n (q_ϕ^i, q_ψ^i) pairs for $1 \leq i \leq n$. For n qubit states, the simplified circuit in Fig. 2d would have the same CNOT and Hadamard gates in the corresponding qubit pairs and a Toffoli gate controlled by them and acting on the ancillary qubit. There is a total of n Toffolis that change the state of the input from $|0\rangle$ to $|1\rangle$ or from $|1\rangle$ to $|0\rangle$ when both controls are $|1\rangle$. The final value depends on the parity of the bitwise AND, \wedge , of

the strings coming from measuring the n qubits of each input state. An even number of ones means the output state in the ancillary qubit would have been $|0\rangle$ (a pass). An odd number of ones corresponds to failing the test.

In terms of the Z observable of the ancillary qubit, the measurement result is exactly the same as the product of the measurement results for the CZ observable on each pair. For two qubits, the CZ operation has eigenvectors $|00\rangle, |01\rangle, |10\rangle$ associated with eigenvalue 1 and an eigenvector $|11\rangle$ associated with eigenvalue -1 . The product of these CZ observables has the same value as a measurement of the observable Z in the SWAP test. Considering all the qubits, the observable with the result of the destructive SWAP test is $CZ^{\otimes n}$ if we write our inputs in the order $q_\phi^1 q_\psi^1 q_\phi^2 q_\psi^2 \cdots q_\phi^n q_\psi^n$. Finding a 1 eigenvalue corresponds to passing the test. Measuring eigenvalue -1 indicates a failure.

3 Hybrid algorithms: variational methods

Physical quantum computers are difficult to build. They have a high cost of operation when compared to classical computers and, in their current implementations, are still too prone to decoherence to be able to carry out long calculations. For all these reasons, whenever possible, we would like to replace the quantum parts of the algorithm by classical computation. For instance, in Shor's algorithm for integer factorization [1], the only part that needs a quantum computer is the order finding subroutine. As shown by Ekerå [22], a classical computer can perform a longer preprocessing stage and the quantum computer can be reserved for a single use of the non-classical routine for order finding. Even this simplified form of factoring still requires a number of qubits and consecutive stages beyond the capabilities of current noisy quantum computers.

The first promising results of real-world quantum computers come from hybrid algorithms where the quantum and the classical stages feed each other. One particularly successful family of algorithms are variational methods taking advantage of the native capacity of a quantum computer to describe states in a large space without needing to store an exponential amount of data [8, 10].

For instance, if we want to compute the ground state energy E_g of a particular Hamiltonian, H , classical variational methods propose a series of trial states. For any state $|\psi\rangle$, the expected energy value $\langle \psi | H | \psi \rangle \geq E_g$ gives an upper bound to the minimal energy of the ground state. If we choose with care the states for which we measure the energy, we can approximate E_g with good precision. Quantum methods offer a compact way to perform variational algorithms. We no longer need to write down the whole state vector of the ground state, which can have a number of complex elements growing exponentially with the size of the problem. For instance, a simple system with n electrons in separate sites can have 2^n spin configurations.

Most variational hybrid algorithms can be described using four main blocks. We will describe them in parallel with an example application for ground state estimation for a Hamiltonian made of a combination of Pauli operators (Fig. 3). The four blocks are:




Fig. 3 Example for three qubits $\langle \psi(\vec{\theta}) | Z \otimes X \otimes Z | \psi(\vec{\theta}) \rangle$

- A parametrized quantum circuit $\mathbf{P}(\vec{\theta})$:

The gates in the circuit depend on a list of classical parameters $\vec{\theta} = (\theta_1, \dots, \theta_k)$ to prepare a trial state, or *ansatz*, $|\psi(\vec{\theta})\rangle$. We assume a fixed initial state, usually with the n qubits initialized to $|0\rangle$ so that $\mathbf{P}(\vec{\theta}) |0\rangle^{\otimes n} = |\psi(\vec{\theta})\rangle$.

In the example of Fig. 3, the classical parameters appear in a series of rotation gates RX and RY which rotate a single qubit with respect to the X and Y axis of the Bloch sphere, respectively, by an angle determined by a classical parameter. There are also two CNOT gates that introduce entanglement in the final state. This circuit follows the approach of what is called a hardware-efficient Ansatz [23], where the gates are selected to give the shallower possible circuit that still can approximate the desired state.

- A processing circuit:

Depending on the desired target, ψ we need to act on the ansatz in different ways. The measurement stage must return values related to an objective function scoring how close we are to our goal.

The example in Fig. 3 shows a typical variational simulation scenario where the gates perform a change of basis so that the final measurements correspond to one term of the desired Hamiltonian in a particular encoding. In this case, the H gate in the middle qubit and the identity in the rest make the final measurement in the computational basis give the exact same results as measuring the $Z \otimes X \otimes Z$ observable for the generated *ansatz* $|\psi(\vec{\theta})\rangle$.

- Measurement:

We usually consider measurement in the computational basis $\{|0\rangle, |1\rangle\}$ for each qubit. As a single measurement is not enough to estimate the desired averages, we need to prepare multiple copies of the same trial state using the exact same classical parameters and processing circuit so that we can get meaningful statistics. This stage has been considered together with the processing circuit in the example.

- Classical processing and feedback:

After a few repetitions, the statistics from the measurement are used to compute an *objective function* $f(\vec{\theta})$ which takes a value which depends on the chosen classical parameters and is related to the problem we want to solve. This dependence is, in general, quite complex. Using a quantum circuit we reduce an explicit description of the problem to measuring a small experimental setup.

The global process is a loop that takes the estimated value of the objective function $f(\vec{\theta})$ at each iteration as a guide to tweak the classical parameters $\vec{\theta}$. In most

cases, the problem is posed as a minimization problem where we use classical optimization methods (such as nonlinear optimization or gradient descent methods) changing the input parameters at stage k to new values $\vec{\theta}_k$ until we find a stable value $f(\vec{\theta}_s)$ at iteration s which is assumed to be the minimum.

In the example, the expected value $\langle \psi(\vec{\theta}) | Z \otimes X \otimes Z | \psi(\vec{\theta}) \rangle$, which approximates the energy of $|\psi(\vec{\theta})\rangle$ for the simulated Hamiltonian, is the objective function to be minimized. The classical part starts by choosing a random initial list of parameters $\vec{\theta}_1$ that gives a first value $f(\vec{\theta}_1)$. The chosen optimization method then generates a new set of classical parameters $\vec{\theta}_2$ that are fed into the parametrized quantum circuit. After we complete the second round of measurements, the new value of the objective function $f(\vec{\theta}_2)$ is compared to the previous one. The changes in the parameters that lead to a smaller final value are kept and the ones that increase the objective function are modified until we can no longer find a better solution. This part of the hybrid algorithm draws heavily on classical optimization and there are multiple available software libraries, like Python’s scikit-learn [24].

Notice that, in many cases, we have no guarantee that the algorithm will succeed. Still, we can provide educated guesses for the solution, similar to what happens in many classical machine learning methods.

The main advantage of this family of hybrid algorithms is the possibility to probe a large state space with a quantum circuit that grows linearly instead of exponentially with the size of the system. This is particularly interesting in the simulation of quantum systems, where the states of the simulated system can be easily mapped into the quantum computer. For instance, chains of n fermions with spin $1/2$ or $-1/2$ can be directly represented by n qubits in the state $|0\rangle$ or $|1\rangle$ and we can directly explore the Hilbert space of dimension 2^n with a compact system and study, among others, phase transitions in the Ising model [25]. A good review of existing applications can be found in [10].

4 A variational method for the eigenvectors of unitary matrices

The SWAP test gives a suitable objective function for a hybrid variational algorithm that finds the eigenstates $|e_i\rangle$ of a unitary evolution U such that $U |e_i\rangle = \lambda_i |e_i\rangle$. The method can be generalized to multiple scenarios.

For an $N \times N$ unitary U , we can find N orthogonal eigenvectors $|e_i\rangle$ which form an orthonormal basis and their corresponding eigenvalues $\lambda_i = e^{i\phi_i}$ are complex numbers of modulus 1.

The circuit for the quantum part of the algorithm is shown in Fig. 4. We use the same parametrized circuit $\mathbf{P}(\vec{\theta})$ with the same classical parameters twice to generate a trial state $|\psi(\vec{\theta})\rangle|\psi(\vec{\theta})\rangle$, which is transformed by the unitary into $|\psi(\vec{\theta})\rangle U |\psi(\vec{\theta})\rangle$ and then proceeds to a destructive SWAP test circuit.

The objective function in the algorithm comes from the statistics of the SWAP test and can be written as

$$f(\vec{\theta}) = P(0) - P(1) = |\langle \psi(\vec{\theta}) | U | \psi(\vec{\theta}) \rangle|^2. \tag{9}$$




Fig. 4 Variational quantum circuit for the eigenvectors of U

We approximate the probabilities of passing and failing the SWAP test, $P(0)$ and $P(1)$, respectively, from counting how many times we get each outcome and dividing by the total number of times we run the test. The 0 and 1 results correspond to the parity of the bitwise AND of the measurement results in the computational basis (assuming 0 for $|0\rangle$ and 1 for $|1\rangle$) and they have the same statistics that would appear if we had measured the ancillary qubit in the complete SWAP test. The evaluation of the objective function is reduced to computing the expected value of the Z observable for the ancillary qubit, albeit in a roundabout way to reduce the final number of gates.

The objective function is an inner product of normalized states, with $f(\vec{\theta}) \leq 1$, and we can see that the maximum is obtained only if the trial state is an eigenvector of U , $|\psi(\vec{\theta})\rangle = |e_i\rangle$. Then

$$f(\vec{\theta}) = |\langle e_i | U | e_i \rangle|^2 = |\langle e_i | \lambda_i | e_i \rangle|^2 = 1. \tag{10}$$

Most optimization methods and software libraries work by minimizing a function. For these, we can consider:

$$f'(\vec{\theta}) = P(1) - P(0) = -|\langle \psi(\vec{\theta}) | U | \psi(\vec{\theta}) \rangle|^2 \tag{11}$$

or

$$f''(\vec{\theta}) = \frac{1}{f(\vec{\theta})} = \frac{1}{P(0) - P(1)}. \tag{12}$$

With these functions that can be computed from measurements on the circuit, we can use a classical optimization method to update the parameters $\vec{\theta}$ until the parametrized circuit giving $\mathbf{P}(\vec{\theta})$ produces a good approximation to an eigenstate of the target unitary U .




Fig. 5 Variational quantum circuit for the generalized eigenvector problem $U |e_i\rangle = \lambda V |e_i\rangle$

4.1 Solving the generalized eigenvalue problem

The circuit can be modified to find generalized eigenvectors in what is usually called the generalized eigenvalue problem of finding $|\psi\rangle$ and λ such that

$$U |e_i\rangle = \lambda V |e_i\rangle \tag{13}$$

for two given matrices U and V (see Fig. 5). The hybrid algorithm is exactly the same as before, now maximizing the objective function $f(\vec{\theta}) = |\langle \psi(\vec{\theta}) | U^\dagger V | \psi(\vec{\theta}) \rangle|^2$ until $f(\vec{\theta}) = |\langle e_i | U^\dagger V | e_i \rangle|^2 = |\langle e_i | \lambda^* | e_i \rangle|^2 = 1$.

4.2 Generalization for normal matrices

The variational eigenvector finder can be modified to work on normal matrices for both the usual and the generalized eigenvalue problems. Consider a normal matrix N such that $NN^\dagger = N^\dagger N$. For an $M \times M$ normal matrix the spectral theorem guarantees we have a decomposition:

$$N = \sum_{i=1}^M \lambda_i |e_i\rangle \langle e_i| \tag{14}$$

for M orthogonal eigenvectors. This spectral decomposition might not be unique if there are degenerate eigenvalues.

For any square $M \times M$ complex matrix A , the matrix $H_A = A + A^\dagger$ is Hermitian ($H_A = H_A^\dagger$). If A admits a spectral decomposition like in Eq. (14), this Hermitian

matrix shares the eigenvectors of A with

$$\begin{aligned}
 H_A &= A + A^\dagger = \sum_{i=1}^M \lambda_i |e_i\rangle \langle e_i| + \sum_{i=1}^M \lambda_i^* |e_i\rangle \langle e_i| = \sum_{i=1}^M (\lambda_i + \lambda_i^*) |e_i\rangle \langle e_i| \\
 &= \sum_{i=1}^M 2 \operatorname{Re}(\lambda_i) |e_i\rangle \langle e_i|.
 \end{aligned}
 \tag{15}$$

If we consider the unitary $U_A = e^{iH_A}$, it will have a decomposition

$$U_A = \sum_{i=1}^M e^{i2\operatorname{Re}(\lambda_i)} |e_i\rangle \langle e_i|.
 \tag{16}$$

The matrices A , H_A and U_A have the same eigenvectors with eigenvector $|e_i\rangle$ associated with eigenvalues λ_i , $2 \operatorname{Re}(\lambda_i)$ and $e^{i2\operatorname{Re}(\lambda_i)}$, respectively.

If we can construct the unitary evolution U_A , we can obtain the eigenvectors of any A admitting a spectral decomposition, including all normal matrices. Finding an efficient quantum gate sequence that realizes the exponential e^{iH_A} might be difficult. This is a well-studied problem that appears in Hamiltonian simulation [26, 27] and in the HHL quantum algorithm for linear systems of equations [28]. All the tricks used there can be recycled in this variational method, including efficient approximation for sparse matrices [29] or designing evolutions directly for the particular quantum computer architecture in which the variational circuit is run [30].

5 Principal component analysis of unknown mixed states

The proposed variational method can be also extended to statistical mixtures, which can appear, among other situations, after a pure state is subject to decoherence. Any mixed state can be described by a density matrix


$$\rho = \sum_j^m p_j |\psi_j\rangle \langle \psi_j|,
 \tag{17}$$

with a variable number of terms m where the p_j are the probabilities associated with finding a pure state $|\psi_j\rangle$ and sum to one. This decomposition is not unique and the states are not necessarily orthogonal. The density matrix ρ is still Hermitian and it will have a spectral decomposition

$$\rho = \sum_{j=1}^N \lambda_j |\psi_j\rangle \langle \psi_j|,
 \tag{18}$$

where N is the dimension of the state space and the $|\psi_i\rangle$ are now orthogonal states. For non-degenerate eigenvalues, this decomposition is unique.

Fig. 6 Quantum circuit for variational quantum principal component analysis



One important problem for mixed states is quantum principal component analysis (QPCA): determining the leading terms in the spectral decomposition. The eigenvalues of the Hermitian ρ are real and can be ordered from the largest to the smallest as $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$. The principal component is the eigenstate associated with the largest eigenvalue which, in certain scenarios, gives a good approximation to the full state. This mirrors classical principal component analysis methods [31], which have applications in multiple fields, like in machine learning, where it is generally used as a way to reduce the dimensionality of the data [32, 33].

The general QPCA problem is thought to be hard for classical computers. Lloyd, Mohseni and Rebentrost showed that there is an efficient quantum algorithm for QPCA as long as we can prepare multiple copies of a state with a density matrix AA^\dagger for any given matrix A [34]. The proposed quantum method was later dequantized by Tang who showed that if we assume there is an equivalent black box giving this state preparation, there are classical methods which are only quadratically worse than the quantum proposal (as opposed to the exponential quantum advantage in the case of low rank matrices) [35].

The variational quantum algorithm proposed in the previous sections can also be used to perform principal component analysis with the same caveats regarding state preparation. We assume we are given multiple copies of a mixed state with density matrix ρ . While for general tasks it might be hard mapping a classical matrix to a state, the method can still be useful, for instance, to approximate the input of a noisy channel within certain precision when the principal component does indeed give a good estimation of the original input [36]. When compared to the QPCA algorithm of [34], we just need the raw input states instead of using them to build the evolution $e^{i\rho}$, which makes the variational algorithm more suitable for current noisy quantum computers. However, we are subject to the usual problems of variational methods: they are heuristic and we have no guarantee of finding a solution, among other challenges.

We will use the circuit in Fig. 6, which also admits a gate efficient realization with a destructive SWAP test like the circuits in Figs. 4 and 5. For an unknown input mixed state ρ , we perform a series of SWAP tests comparing the input to the pure state $\sigma(\vec{\theta}) = |\psi(\vec{\theta})\rangle\langle\psi(\vec{\theta})|$ generated from a fixed input and a parametrized quantum circuit controlled by the classical parameters $\vec{\theta}$. In the optimization phase, we will tweak these classical parameters until the *ansatz* $\sigma(\vec{\theta})$ maximizes the value of the measured observable.

From Eqs. (7) and (18), we see the expected value we approximate is

$$\text{tr}(\rho\sigma(\vec{\theta})) = \text{tr}\left(\sum_{i=1}^N \lambda_i |\psi(\vec{\theta})\rangle\langle\psi(\vec{\theta})|\right) = \sum_i \lambda_i |\langle\psi(\vec{\theta})|\psi_i\rangle|^2. \quad (19)$$

The global maximum is achieved when the ansatz is the principal component (the eigenstate $|\psi_1\rangle$ associated with the largest eigenvalue λ_1). In that case the expected value becomes λ_1 . For any other trial state, with at least some part of the state orthogonal to $|\psi_1\rangle$ and $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \dots \geq \lambda_N$:

$$\text{tr}(\rho\sigma(\vec{\theta})) = \sum_i \lambda_i |\langle\psi(\vec{\theta})|\psi_i\rangle|^2 \leq \lambda_1 \sum_i |\langle\psi(\vec{\theta})|\psi_i\rangle|^2 = \lambda_1. \quad (20)$$

In general, the final state found at the end of the algorithm will depend on the optimization method and the input state. If we start close to a principal component the objective function will find a local maximum with a trial state close to the desired state. However, the search algorithm can also converge to local maxima that do not correspond to a principal component. The global maximum will be more prominent for mixed states with a strong principal component.

There are two important differences with respect to eigenvector estimation for unitary and normal matrices. In those cases, we could know for sure the algorithm had converged to a true eigenvector by checking the expected value, which should converge to 1. Here, we cannot be sure we have found the principal component and not converged to some local maximum (likely a second or third component, or a superposition of the first components). However, we get additional information if we did. In QPCA, it is useful finding both the state and the eigenvalue, which can be learnt from the expected value. The value will tell us how important the principal component is with respect to other terms and how well it can be used as a compact description of the whole state. Larger values correspond to principal components that represent the whole mixed state better.

Using multiobjective optimization, we can also search for successive components. When the algorithm converges after k iterations, we obtain a set of parameters $\vec{\theta}_k^1$ that give the estimated principal component $|\tilde{\psi}_1\rangle = |\psi(\vec{\theta}_k^1)\rangle$. We can now alternate two series of rounds for each iteration j in order to evaluate the parameters $\vec{\theta}_j^2$ defining our trial state for the second component. First we can perform the usual comparison of the *ansatz* with the input state ρ to estimate $\text{tr}(\rho\sigma(\vec{\theta}_j^2))$. Then, we generate as input states $\sigma(\vec{\theta}_k^1)$ and $\sigma(\vec{\theta}_j^2)$ and use a SWAP test to estimate the overlap $\text{tr}(\sigma(\vec{\theta}_k^1)\sigma(\vec{\theta}_j^2)) = |\langle\psi(\vec{\theta}_k^1)|\psi(\vec{\theta}_j^2)\rangle|^2$. We want two orthogonal states such that $\text{tr}(\sigma(\vec{\theta}_k^1)\sigma(\vec{\theta}_j^2)) = 0$. Using these two estimated values, the classical parameters must be optimized according to a new objective function that the classical algorithm will use to propose the vector $\vec{\theta}_{j+1}^2$ for the next iteration. Assuming we minimize $f(\vec{\theta})$ during optimization, we can, for instance, define an objective function

$$f(\vec{\theta}) = 1/\text{tr}(\rho\sigma(\vec{\theta})) + C|\langle\psi(\vec{\theta})|\tilde{\psi}_1\rangle|^2 \quad (21)$$

with a weight constant C that penalizes the previously found principal component and directs the search toward the next maximum eigenvalue in the decomposition of ρ . This procedure can be repeated as many times as needed with additional terms to search for the third, fourth and successive components, much like molecular search for excited energy levels above the ground state in molecules [37].

6 Optical implementation

The variational algorithms for eigenvectors and quantum principal component analysis we have seen are not restricted to near-term universal quantum computers. They can also be carried out in linear optical setups with certain restrictions.

We consider passive linear optical systems with input states that are a superposition of terms $|n_1 n_2 \cdots n_m\rangle$ where we have $n = \sum_{i=1}^m n_i$ photons distributed into m orthogonal modes (like different paths or orthogonal polarizations). A linear optics system acting on m modes can be described classically by its scattering matrix S which is an $m \times m$ unitary matrix [38, 39]. The corresponding quantum evolution acting on n photons can be computed from S as $U = \varphi_{m,n}(S)$ using different known methods [40–43]. The unitary U has a size $M \times M$, with $M = \binom{n+m-1}{n}$. For one photon the quantum evolution is exactly the scattering matrix $S = U$ [44].

There are multiple constructive methods to map any desired unitary S into a physical setup using only linear optical devices such as beamsplitters and phase shifters [45–50]. These implementations have a circuit size that grows quadratically with the number of modes in the worst case and can generate any desired unitary S . There exist many successful experimental realizations of configurable optical circuits with integrated optics [51–57] that offer a linear optical system that can be controlled electronically. These hybrid systems have configurable phase shifters and beamsplitters and give quantum evolutions that depend on a few classical parameters. They can be designed to provide any available scattering matrix for the size input, but there are also simpler configurable circuits that can produce good enough *ansätze* in variational quantum algorithms [58].

These systems give an efficient way to realize the first half of the proposed quantum circuits which has the parametrized unitaries that generate the trial states for eigenvector estimation and principal component analysis. For the second half which performs the state comparison, we need to restrict our setups.

The SWAP test admits a simple experimental realization for two photons. We start from two separate single photons in the same large Hilbert space. We just need to direct the two photons encoding the states to the two inputs of a balanced beamsplitter. At the output of the beamsplitter, we place two binary photodetectors (that click for one or more photons and do nothing for the vacuum). In such a setup, we call a coincidence to the simultaneous detection of one photon at each output port of the beamsplitter. For two input single photon states described by density matrices ρ and σ , the probability of finding a coincidence at the output is [59]:

$$P_C = \frac{1 - \text{tr}(\rho\sigma)}{2}, \quad (22)$$

which is exactly the probability of failing a SWAP test. That way, optical measurements at a beamsplitter can be used for state comparison [19].

Notice we need to encode the whole state from a large Hilbert space into a single photon. There are many alternatives, like taking advantage of the high dimensionality provided by orbital angular momentum [60–62]. A good option compatible with current optical network technology is using a time-bin encoding for the orthogonal




Fig. 7 Time bin encoding for optical qubits. The time location of an optical pulse can serve as a basis for a qubit. With different amplitude combinations and relative phases between the optical pulses, we can produce any desired qubit

modes. Figure 7 shows different states of a qubit in time-bin encoding. The photon wavefunction can be confined to two different time bins representing the $|0\rangle$ and $|1\rangle$ qubits. If we use a common phase reference that is assigned to the pulses in the $|0\rangle$ time bin, we can use amplitude modulators to distribute the probability amplitude in the two time bins and phase modulators to introduce a relative phase with respect to the reference. The result is an optical qubit in any state of choice. This approach has been used, for instance, in quantum key distribution protocols over optical fiber [63, 64].

In the proposed variational algorithms, we need d -dimensional systems, or qudits, which can be generated by dividing the photon into d time bins. With the current technology, we can expect a reasonable coherence time of microseconds and amplitude and phase modulators in the GHz range, which could produce systems with 1024 bins, equivalent to 10 qubits. Existing modulators, beamsplitters and detectors could give results comparable to what can be done in a quantum computer with 20 qubits. Each of the two photons could also be distributed into multiple separate spatial modes, possibly with added time-bin encoding. The measurement would need a balanced beamsplitter for each pair of modes corresponding to the same position for each photon and one detector at each of their outputs.

In all the variational algorithms proposed, both for finding eigenvectors and for quantum principal component analysis, one of the subsystems would be one photon directed to a parametrized linear optics multipoint. The second system can either come from a general optical quantum channel, which could produce a mixed state we want to characterize, or a known optical transformation if we want to map a known matrix

S into the optical system in order to find a compact description of its eigenvectors, as, in this case with a single photon, $U = S$.

While the scalability of this approach is limited, it can be an interesting alternative to full quantum computers in the near term, particularly in the cases where we need a low noise implementation. Photons have long coherence times, relatively low noise and do not need advanced cooling systems unlike many implementations of quantum computers, especially those based on superconducting qubits. Optical systems have their own problems like losses, synchronization or producing good approximations to single photon states, but, for most of these problems, there are acceptable technical solutions that have been used in quantum key distribution systems [65, 66].

7 Comparison to other variational quantum algorithms for eigenvector determination

There exist various proposals for variational quantum algorithms that learn an eigenstate of a Hamiltonian, a density matrix or a general matrix, each focusing on different objectives. The main motivation of the algorithms put forward in this paper is to reduce as much as possible the depth of final quantum circuit.

Noisy intermediate-scale quantum computers are the natural hardware to run the quantum part of hybrid variational quantum algorithms. However, variational methods are severely limited by the noise in each quantum gate.

A first strong limitation is that there is a bound to the maximum depth for which variational algorithms give genuine quantum advantage [67]. Above a certain depth threshold, which depends on the noise level, classical Monte Carlo techniques offer an efficient simulation of the quantum variational method. Even if the results are correct despite the noise, the problems that can be solved in this case can also be solved, without all the specialized hardware, using a classical computer.

A second important limitation comes from the classical optimization phase. Noise introduces barren plateaus in the optimization landscape [68]: the gradient vanishes exponentially with the depth of the parametrized circuit generating the ansatz. This makes classical optimization unfeasible. The search for a global minimum of the cost function will get stuck at some sub-optimal parameters.

We can compare the proposed variational methods to previous algorithms in terms of the depth of the circuits.

We have approximated the overlap $\langle |\psi(\vec{\theta})\rangle | U | \psi(\vec{\theta})\rangle \rangle$ comparing two states $U | \psi(\vec{\theta})\rangle = U \mathbf{P}(\vec{\theta}) | 0\rangle^{\otimes n}$ and $\langle |\psi(\vec{\theta})\rangle = \mathbf{P}(\vec{\theta}) | 0\rangle^{\otimes n}$ using the SWAP test. However, we can just take an initial all-zeroes state and make it evolve under $\mathbf{P}^\dagger(\vec{\theta}) U \mathbf{P}(\vec{\theta})$ [69]. If we project the resulting state to the $|0\rangle^{\otimes n}$ output, the statistics for finding all the output qubits in $|0\rangle$ are ${}^{n\otimes}\langle 0 | \mathbf{P}^\dagger(\vec{\theta}) U \mathbf{P}(\vec{\theta}) | 0\rangle^{\otimes n}$ which is exactly $\langle |\psi(\vec{\theta})\rangle | U | \psi(\vec{\theta})\rangle \rangle$.

This approach has advantages and disadvantages with respect to using the SWAP test. The biggest advantage is that it only uses n qubits for $2^n \times 2^n$ unitaries instead of $2n$ qubits and does not require any additional gates. In the proposed destructive SWAP test the overhead of the SWAP test reduces to two gates per qubit pair, which is not

so important. The qubit advantage is not trivial, though, particularly for intermediate-scale applications where qubits are scarce. However, the additional circuit depth due to including $\mathbf{P}(\vec{\theta})$ can be important, as the performance of variational algorithms can degrade exponentially with the total circuit depth [68] and deep noisy circuits can lose quantum advantage [67].

There are also other variational methods that assume a known Hamiltonian that can be efficiently written in the Pauli basis [70, 71]. For eigenstates, the cost function

$$C(\vec{\theta}) = \langle \psi(\vec{\theta}) | H^2 | \psi(\vec{\theta}) \rangle - \langle \psi(\vec{\theta}) | H | \psi(\vec{\theta}) \rangle^2 \quad (23)$$

is equal to zero and can be used to tune the parameters in $\mathbf{P}(\vec{\theta})$. This only requires n qubits for a Hamiltonian corresponding to a $2^n \times 2^n$ matrix. The additional cost of measuring the corresponding Pauli operators is small (at most a change of basis gate on each qubit). This can be more efficient than the SWAP test eigenvector finder, but it only works for known Hamiltonians which can be efficiently decomposed into Pauli operators (in some encoding).

Notice that the SWAP test algorithms proposed in this paper would also work on arbitrary unitaries U , even if we just have a description in terms of the quantum circuit giving the evolution. We do not need to compute a closed form for the matrix U or find the corresponding Hamiltonians.

The application of the SWAP test for variational quantum principal analysis works on unknown input density matrices ρ and has the simplest circuit from all we have seen in this paper. Two notable variational algorithms are given in [72, 73], both of which follow the same philosophy of the method in Sect. 5: after the procedure we have a set of classical parameters that allow the reconstruction of the desired principal component with a small error.

The algorithm for variational quantum state diagonalization in [72] takes two copies of the unknown state per each quantum iteration. Both copies go through the same parametrized diagonalization circuit $\mathbf{P}(\vec{\theta})$ and they are compared using variations of the destructive SWAP test or inner product measurements. The cost function is chosen so that, after successful optimization, $\mathbf{P}^\dagger(\vec{\theta})\rho\mathbf{P}(\vec{\theta})$ is diagonal. Later, this diagonalizing operation can be used to recover the eigenvalues and eigenvectors of ρ . The complexity is essentially the same as the circuit in Fig. 6. The circuits have the same depth, but the method needs two copies of the parametrized quantum circuit and the input state instead of the one copy in Fig. 6.

The alternative variational quantum state eigensolver in [73] only requires one copy of ρ and the parametrized circuit and, instead of the SWAP test approach, it only uses n qubits (for states in a Hilbert space of dimension 2^n). This algorithm produces the same result as the method in [72] using a cost function $C(\vec{\theta}) = \text{tr}(\mathbf{P}^\dagger(\vec{\theta})\rho\mathbf{P}(\vec{\theta})H)$ for a given Hamiltonian H which can be adapted during training. This method is more versatile and compact at the cost of adding a few gates to include the Hamiltonian in the measurement.

In general, these two algorithms are more flexible than the method in Sect. 5 and they give a full solution to principal component analysis with a slightly higher gate

cost. If the full spectrum is needed they can be more interesting, but they both require more sophisticated processing and interpretation.

8 Strong and weak points, applications and outlook

Classically, finding the eigenvalues of a matrix is an efficient task. For unitary matrices, the QR decomposition [74–76] gives a robust algorithm with efficient and stable numerical software implementations in reference suites like the LAPACK library [77]. For an $N \times N$ matrix, the number of operations grows as $O(N^3)$, with alternative algorithms with complexities between quadratic and cubic depending on special cases [78–80]. This includes bisection methods that can search only for one or a limited subset of eigenvectors instead of the whole spectrum at a smaller cost in terms of the number of operations.


However, in a quantum setting, the size of the state space grows exponentially in the number of qubits. For n qubits, this means that any algorithm that needs to explicitly write down the $N = 2^n$ complex entries in a general eigenvector will suffer from this growth. The variational quantum algorithms introduced in this paper return a compressed version of the state so that we can produce the desired eigenstate on demand. The classical parameters in $\vec{\theta}$ are a short list that contains all the information we need to recreate the desired eigenvector using the parametrized circuit we chose for the search. This native encoding into a quantum state allows for an efficient preparation for future uses.

The algorithm for variational quantum principal component analysis can also be useful for channel characterization, for instance, in the analysis of mixed states coming from a quantum optical channel with decoherence or to study the decay of a quantum state with time in a noisy quantum computer.

The main potential application for these methods is finding eigenvectors as a stepping stone for more complex routines. Take for instance the quantum phase estimation algorithm of Kitaev [81] where a quantum computer can give efficient approximations to an eigenvalue of a given unitary provided one of the inputs is the corresponding eigenvector. The variational quantum eigenvector finder can be used in conjunction with quantum phase estimation to obtain the full spectrum of any unitary U , including all the eigenvalues.

The variational eigenvector finder also combines well with Abram and Lloyd's algorithm that gives an efficient approximation to the eigenvectors of a Hamiltonian with arbitrary precision as long as it has a "good enough" initial guess state [82]. If we can produce a state $|\tilde{e}_i\rangle$ that approximates the desired exact eigenstate $|e_i\rangle$ and has a non-negligible overlap $|\langle \tilde{e}_i | e_i \rangle|^2$ with it, the eigenstate can be refined and taken as close to the exact value as desired. This fits particularly well with the output of the variational eigenvector finder. Even if the optimization does not converge to an exact eigenvector, we can determine from $P(0) - P(1)$ (the Z observable from the SWAP test) whether the trial state is close or not to an eigenstate. Larger expected values correspond to a greater overlap. With the correction from Abrams–Lloyd algorithm, even poor *ansätze* can give a clean eigenstate.

Fig. 8 Quantum circuit realizing the Hadamard test



While these are interesting applications, both require long circuits and use the quantum Fourier transform which, for large systems, is still impractical due to accumulated noise. A more realistic application in the short term would be optimizing algorithms for noisy intermediate-scale quantum computers.

Take, for instance, the Hadamard test of Fig. 8, which appears in many proposals for quantum machine learning [12, 83] and in the quantum algorithm for Jones polynomials [84, 85]. This circuit is a simpler version of the quantum phase estimation algorithm that still can give results that no classical system can.

A key element in that circuit is the controlled unitary cU which applies a unitary U to the target when the control qubit is $|1\rangle$ and is the identity for a control qubit $|0\rangle$. If we are given U as a black box, this task is impossible in the usual quantum circuit model without describing the complete cU operator and decomposing it [86]. The total number of gates in those decompositions can be large, introducing too much noise, and there have been various attempts to give concise descriptions for particular cases in order to reduce the depth of these circuits [87–90]. In particular, if we can produce an eigenstate of the evolution there are restricted efficient circuits for cU [81]. The variational quantum eigenvector finder can be part of general methods to compute a Hadamard test with a smaller total number of gates.

A preliminary simulation of the proposed quantum variational eigenvector finder [91] shows that, for moderate sizes of U , the algorithm converges to valid eigenvectors both for random input unitaries and for defective matrices with high degeneracy like the quantum Fourier transform. As the size of the problem grows, the found state has a smaller overlap with a true eigenstate. This is a common problem in variational quantum algorithms. For an ideal algorithm, we can just use more complex parametrized circuits to cover a larger part of the state space up to a point. However, in noisy execution, this approach can be counterproductive. In the simulations, even for the smallest problems, noise introduces an appreciable deviation from true eigenstates.

On this account, the proposed algorithms share the drawbacks of most hybrid variational methods. First of all, we need a good *ansatz*. There are different proposals for parametrized circuits that could be used. In any variational method, there is a tension between expressivity and depth. Expressivity describes how much of the whole state space can be reached from outputs of the parametrized circuit. Depth is given by the number of consecutive elementary gates we need to build the circuit. We would like to have expressive enough circuits that can generate an *ansatz* close to our target state (the exact eigenvector we are searching for). However, highly expressive parametrized circuits tend to have larger depths and the noise can build up to levels that make the output unusable. Apart from that, in order to sample more of the space they need more classical parameters. Optimizing a large set of parameters is less efficient and the classical part of the algorithm can converge to suboptimal solutions. Among other challenges, if we can represent very small changes, there can appear barren plateaus

(the optimization algorithm gets stuck in a region of the state space with small changes in the objective function but far from the real minimum) [92]. While this can limit the usefulness of the method as the state space grows, there are methods to optimize the parametric circuits so that the number of parameters is small but we are still able to produce states close to the target [93, 94]. The search for better parametrized circuits and optimization methods is a vibrant area of research and most of the results for other variational quantum algorithms are likely to be useful in the eigenvector algorithms.

One advantage with respect to other problems, like searching for the ground state in molecular simulation, is that we do have a confirmation whether the algorithm has succeeded or not. We might not converge to a true solution, but, when we do, the expected value will be one.

In general, hybrid algorithms are an interesting solution to middle sized problems, but have challenges scaling to really large state spaces. The classical optimization phase becomes more involved with the number of parameters and we might not converge to the optimal solution. Even in this case, if scalable full quantum computers become available, the initial approximation the variational eigenvector finder gives can be used to kickstart the Abrams–Lloyd algorithm [82].

Despite these limitations, the presented variational eigenvector finder and quantum principal component analysis algorithms can become a useful addition to the toolbox of quantum computation in the near future and they can be incorporated into practical, more complex hybrid quantum-classical algorithms.

Acknowledgements This work has been funded by the Spanish Ministerio de Ciencia e Innovación (MCIN), project PID2020-119418GB-I00 and by the European Union NextGeneration UE/MICIU/Plan de Recuperación, Transformación y Resiliencia/Junta de Castilla y León.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature.

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

1. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **26**, 1484 (1997). <https://doi.org/10.1137/S0097539795293172>
2. Montanaro, A.: Quantum algorithms: an overview. *npj Quantum Inf.* **2**, 1–8 (2016). <https://doi.org/10.1038/npjqi.2015.23>
3. Preskill, J.: Quantum Computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018). <https://doi.org/10.22331/q-2018-08-06-79>
4. Bharti, K., Cervera-Lierta, A., Kyaw, T.H., Haug, T., Alperin-Lea, S., Anand, A., Degroote, M., Heimonen, H., Kottmann, J.S., Menke, T., Mok, W.-K., Sim, S., Kwok, L.-C., Aspuru-Guzik, A.: Noisy intermediate-scale quantum algorithms. *Rev. Modern Phys.* **94**, 015004 (2022). <https://doi.org/10.1103/RevModPhys.94.015004>

5. McClean, J.R., Romero, J., Babbush, R., Aspuru-Guzik, A.: The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **18**, 023023 (2016). <https://doi.org/10.1088/1367-2630/18/2/023023>
6. Endo, S., Cai, Z., Benjamin, S.C., Yuan, X.: Hybrid quantum-classical algorithms and quantum error mitigation. *J. Phys. Soc. Jpn.* **90**, 032001 (2021). <https://doi.org/10.7566/JPSJ.90.032001>
7. Callison, A., Chancellor, N.: Hybrid quantum-classical algorithms in the noisy intermediate-scale quantum era and beyond. *Phys. Rev. A* **106**, 010101 (2022). <https://doi.org/10.1103/PhysRevA.106.010101>
8. McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S.C., Yuan, X.: Quantum computational chemistry. *Rev. Mod. Phys.* **92**, 015003 (2020). <https://doi.org/10.1103/RevModPhys.92.015003>
9. Bauer, B., Bravyi, S., Motta, M., Chan, G.K.-L.: Quantum algorithms for quantum chemistry and quantum materials science. *Chem. Rev.* **120**, 12685–12717 (2020). <https://doi.org/10.1021/acs.chemrev.9b00829>
10. Tilly, J., Chen, H., Cao, S., Picozzi, D., Setia, K., Li, Y., Grant, E., Wossnig, L., Rungger, I., Booth, G.H., Tennyson, J.: The variational quantum eigensolver: A review of methods and best practices. *Phys. Rep.* **986**, 1–128 (2022). <https://doi.org/10.1016/j.physrep.2022.08.003>
11. Mitarai, K., Negoro, M., Kitagawa, M., Fujii, K.: Quantum circuit learning. *Phys. Rev. A* **98**, 032309 (2018). <https://doi.org/10.1103/PhysRevA.98.032309>
12. Benedetti, M., Lloyd, E., Sack, S., Fiorentini, M.: Parameterized quantum circuits as machine learning models. *Quantum Sci. Technol.* **4**, 043001 (2019). <https://doi.org/10.1088/2058-9565/ab4eb5>
13. Schuld, M., Killoran, N.: Quantum machine learning in feature hilbert spaces. *Phys. Rev. Lett.* **122**, 040504 (2019). <https://doi.org/10.1103/PhysRevLett.122.040504>
14. Schuld, M., Bocharov, A., Svore, K.M., Wiebe, N.: Circuit-centric quantum classifiers. *Phys. Rev. A* **101**, 032308 (2020). <https://doi.org/10.1103/PhysRevA.101.032308>
15. Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S.C., Endo, S., Fujii, K., McClean, J.R., Mitarai, K., Yuan, X., Cincio, L., et al.: Variational quantum algorithms. *Nat. Rev. Phys.* **3**, 625–644 (2021). <https://doi.org/10.1038/s42254-021-00348-9>
16. Buhrman, H., Cleve, R., Watrous, J., de Wolf, R.: Quantum fingerprinting. *Phys. Rev. Lett.* **87**, 167902 (2001). <https://doi.org/10.1103/PhysRevLett.87.167902>
17. Barenco, A., Berthiaume, A., Deutsch, D., Ekert, A., Jozsa, R., Macchiavello, C.: Stabilization of quantum computations by symmetrization. *SIAM J. Comput.* **26**, 1541–1557 (1997). <https://doi.org/10.1137/S0097539796302452>
18. Kobayashi, H., Matsumoto, K., Yamakami, T.: “Quantum Merlin-Arthur proof systems: are multiple Merlins more helpful to Arthur?”. In Ibaraki, T., Katoh, N. and Ono, H. editors, *Algorithms and Computation*. Pages 189–198. Berlin, Heidelberg (2003). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-24587-2_21
19. Juan Carlos Garcia-Escartin and Pedro Chamorro-Posada: SWAP test and Hong-Ou-Mandel effect are equivalent. *Phys. Rev. A* **87**, 052330 (2013). <https://doi.org/10.1103/PhysRevA.87.052330>
20. Cincio, L., Subaşı, Y., Sornborger, A.T., Coles, P.J.: Learning the quantum algorithm for state overlap. *New J. Phys.* **20**, 113022 (2018). <https://doi.org/10.1088/1367-2630/aae94a>
21. Warren, H.S.: “Hacker’s delight”. Addison-Wesley Professional. (2012). 2nd edition
22. Ekerå, M.: On completely factoring any integer efficiently in a single run of an order-finding algorithm. *Quantum Inf. Process.* **20**, 1–14 (2021). <https://doi.org/10.1007/s11128-021-03069-1>
23. Kandala, A., Mezzacapo, A., Temme, K., Takita, M., Brink, M., Chow, J.M., Gambetta, J.M.: Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017). <https://doi.org/10.1038/nature23879>
24. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)
25. Smith, A., Jobst, B., Green, A.G., Pollmann, F.: Crossing a topological phase transition with a quantum computer. *Phys. Rev. Res.* **4**, L022020 (2022). <https://doi.org/10.1103/PhysRevResearch.4.L022020>
26. Lloyd, S.: Universal quantum simulators. *Science* **273**, 1073–1078 (1996). <https://doi.org/10.1126/science.273.5278.1073>
27. Berry, D.W., Childs, A.M., Cleve, R., Kothari, R., Somma, R.D.: Simulating hamiltonian dynamics with a truncated taylor series. *Phys. Rev. Lett.* **114**, 090502 (2015). <https://doi.org/10.1103/PhysRevLett.114.090502>
28. Harrow, A.W., Hassidim, A., Lloyd, S.: Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **103**, 150502 (2009). <https://doi.org/10.1103/PhysRevLett.103.150502>

29. Berry, D.W., Childs, A.M., Cleve, R., Kothari, R., Somma, R.D.: Exponential improvement in precision for simulating sparse hamiltonians. In Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing. Page 283–292. STOC '14New York, NY, USA (2014). Association for Computing Machinery. <https://doi.org/10.1145/2591796.2591854>
30. Clinton, L., Bausch, J., Cubitt, T.: Hamiltonian simulation algorithms for near-term quantum hardware. *Nat. Commun.* **12**, 4989 (2021). <https://doi.org/10.1038/s41467-021-25196-0>
31. Karl Pearson, F.R.S.: LIII. On lines and planes of closest fit to systems of points in space. *London Edinburgh Dublin Philos. Mag. J. Sci.* **2**, 559–572 (1901). <https://doi.org/10.1080/14786440109462720>
32. Jolliffe, I.T., Cadima, J.: Principal component analysis: a review and recent developments. *Philos. Transact. Royal Soc. A: Math. Phys. Eng. Sci.* **374**, 20150202 (2016). <https://doi.org/10.1098/rsta.2015.0202>
33. Jolliffe, I.T.: *Principal component analysis* (2nd ed). Springer. (2002). <https://doi.org/10.1007/b98835>
34. Lloyd, S., Mohseni, M., Rebentrost, P.: Quantum principal component analysis. *Nat. Phys.* **10**, 631–633 (2014). <https://doi.org/10.1038/nphys3029>
35. Tang, E.: Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions. *Phys. Rev. Lett.* **127**, 060503 (2021). <https://doi.org/10.1103/PhysRevLett.127.060503>
36. Koczor, B.: The dominant eigenvector of a noisy quantum state. *New J. Phys.* **23**, 123047 (2021). <https://doi.org/10.1088/1367-2630/ac37ae>
37. Higgott, O., Wang, D., Brierley, S.: Variational quantum computation of excited states. *Quantum* **3**, 156 (2019). <https://doi.org/10.22331/q-2019-07-01-156>
38. Haus, H.A.: From classical to quantum noise. *J. Opt. Soc. Am. B* **12**, 2019–2036 (1995). <https://doi.org/10.1364/JOSAB.12.002019>
39. Pozar, D.M.: *Microwave engineering*. Wiley. (2004). Fourth edition
40. Caianiello, E.R.: On quantum field theory – I: explicit solution of Dyson’s equation in electrodynamics without use of Feynman graphs. *Il Nuovo Cimento* **1943–1954**(10), 1634–1652 (1953). <https://doi.org/10.1007/BF02781659>
41. Skaar, J., Escartin, J.C.G., Landro, H.: Quantum mechanical description of linear optics. *Am. J. Phys.* **72**, 1385–1391 (2004). <https://doi.org/10.1119/1.1775241>
42. Scheel, S.: Permanents in linear optical networks (2004). [arXiv:quant-ph/0406127](https://arxiv.org/abs/quant-ph/0406127)
43. Aaronson, S., Arkhipov, A.: The computational complexity of linear optics. In Proceedings of the 43rd Annual ACM Symposium on Theory of Computing. Pages 333–342. STOC '11New York, NY, USA (2011). ACM. <https://doi.org/10.1145/1993636.1993682>
44. Cerf, N.J., Adami, C., Kwiat, P.G.: Optical simulation of quantum logic. *Phys. Rev. A* **57**, 1477–1480 (1998). <https://doi.org/10.1103/PhysRevA.57.R1477>
45. Reck, M., Zeilinger, A., Bernstein, H.J., Bertani, P.: Experimental realization of any discrete unitary operator. *Phys. Rev. Lett.* **73**, 58–61 (1994). <https://doi.org/10.1103/PhysRevLett.73.58>
46. Bouland, A., Aaronson, S.: Generation of universal linear optics by any beam splitter. *Phys. Rev. A* **89**, 062316 (2014). <https://doi.org/10.1103/PhysRevA.89.062316>
47. Sawicki, A.: Universality of beam splitters. *Quantum Inf. Comput.* **16**, 0291–0312 (2016). <https://doi.org/10.26421/QIC16.3-4-6>
48. Clements, W.R., Humphreys, P.C., Metcalf, B.J., Steven Kolthammer, W., Walmsley, I.A.: Optimal design for universal multiport interferometers. *Optica* **3**, 1460–1465 (2016). <https://doi.org/10.1364/OPTICA.3.001460>
49. de Guise, H., Di Matteo, O., Sánchez-Soto, L.L.: Simple factorization of unitary transformations. *Phys. Rev. A* **97**, 022328 (2018). <https://doi.org/10.1103/PhysRevA.97.022328>
50. Bell, B.A., Walmsley, I.A.: Further compactifying linear optical unitaries. *APL Photon.* **6**, 070804 (2021). <https://doi.org/10.1063/5.0053421>
51. Carolan, J., Harrold, C., Sparrow, C., Martín-López, E., Russell, N.J., Silverstone, J.W., Shadbolt, P.J., Matsuda, N., Oguma, M., Itoh, M., Marshall, G.D., Thompson, M.G., Matthews, J.C.F., Hashimoto, T., O’Brien, J.L., Laing, A.: Universal linear optics. *Science* **349**, 711–716 (2015). <https://doi.org/10.1126/science.aab3642>
52. Mennea, P.L., Clements, W.R., Smith, D.H., Gates, J.C., Metcalf, B.J., Bannerman, R.H.S., Burgwal, R., Renema, J.J., Steven Kolthammer, W., Walmsley, I.A., Smith, P.G.R.: Modular linear optical circuits. *Optica* **5**, 1087–1090 (2018). <https://doi.org/10.1364/OPTICA.5.001087>

53. Bogaerts, W., Pérez, D., Capmany, J., Miller, D.A.B., Poon, J., Englund, D., Morichetti, F., Melloni, A.: Programmable photonic circuits. *Nature* **586**, 207–216 (2020). <https://doi.org/10.1038/s41586-020-2764-0>
54. Elshaari, A.W., Pernice, W., Srinivasan, K., Benson, O., Zwiller, V.: Hybrid integrated quantum photonic circuits. *Nat. Photon.* **14**, 285–298 (2020). <https://doi.org/10.1038/s41566-020-0609-x>
55. Arrazola, J.M., Bergholm, V., Brádler, K., Bromley, T.R., Collins, M.J., Dhand, I., Fumagalli, A., Gerrits, T., Goussev, A., Helt, L.G., et al.: Quantum circuits with many photons on a programmable nanophotonic chip. *Nature* **591**, 54–60 (2021). <https://doi.org/10.1038/s41586-021-03202-1>
56. Taballione, C., van der Meer, R., Snijders, H.J., Hooijschuur, P., Epping, J.P., de Goede, M., Kassenberg, B., Venderbosch, P., Toebes, C., van den Vlekkert, H., Pinkse, P.W.H., Renema, J.J.: A universal fully reconfigurable 12-mode quantum photonic processor. *Mater. Quantum Technol.* **1**, 035002 (2021). <https://doi.org/10.1088/2633-4356/ac168c>
57. Hoch, F., Piacentini, S., Giordani, T., Tian, Z.-N., Iuliano, M., Esposito, C., Camillini, A., Carvacho, G., Ceccarelli, F., Spagnolo, N., Crespi, A., Sciarrino, F., Osellame, R.: Reconfigurable continuously-coupled 3D photonic circuit for boson sampling experiments. *npj Quantum Inf.* **8**, 1–7 (2022). <https://doi.org/10.1038/s41534-022-00568-6>
58. Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.-H., Zhou, X.-Q., Love, P.J., Aspuru-Guzik, A., Jeremy, O.L.: A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 1–7 (2014). <https://doi.org/10.1038/ncomms5213>
59. Schwarz, L., van Enk, S.J.: Detecting the drift of quantum sources: not the de Finetti theorem. *Phys. Rev. Lett.* **106**, 180501 (2011). <https://doi.org/10.1103/PhysRevLett.106.180501>
60. Allen, L., Beijersbergen, M.W., Spreeuw, R.J.C., Woerdman, J.P.: Orbital angular momentum of light and the transformation of laguerre-gaussian laser modes. *Phys. Rev. A* **45**, 8185–8189 (1992). <https://doi.org/10.1103/PhysRevA.45.8185>
61. Allen, L., Barnett, S.M., Padgett, M.J.: *Optical angular momentum*. Institute of Physics Publishing, Bristol, UK. (2003). <https://doi.org/10.1201/9781482269017>
62. Willner, A.E., Pang, K., Song, H., Zou, K., Zhou, H.: Orbital angular momentum of light for communications. *Appl. Phys. Rev.* **8**, 041312 (2021). <https://doi.org/10.1063/5.0054885>
63. Muller, A., Herzog, T., Huttner, B., Tittel, W., Zbinden, H., Gisin, N.: Plug and play’ systems for quantum cryptography. *Appl. Phys. Lett.* **70**, 793–795 (1997). <https://doi.org/10.1063/1.118224>
64. Woodward, R.I., Lo, Y.S., Pittaluga, M., Minder, M., Paraíso, T.K., Lucamarini, M., Yuan, Z.L., Shields, A.J.: Gigahertz measurement-device-independent quantum key distribution using directly modulated lasers. *npj Quantum Inf.* **7**, 58 (2021). <https://doi.org/10.1038/s41534-021-00394-2>
65. Scarani, V., Bechmann-Pasquinucci, H., Cerf, N.J., Dušek, M., Lütkenhaus, N., Peev, M.: The security of practical quantum key distribution. *Rev. Mod. Phys.* **81**, 1301–1350 (2009). <https://doi.org/10.1103/RevModPhys.81.1301>
66. Feihu, X., Ma, X., Zhang, Q., Lo, H.-K., Pan, J.-W.: Secure quantum key distribution with realistic devices. *Rev. Mod. Phys.* **92**, 025002 (2020). <https://doi.org/10.1103/RevModPhys.92.025002>
67. Stilck França, D., García-Patron, R.: Limitations of optimization algorithms on noisy quantum devices. *Nat. Phys.* **17**(11), 1221–1227 (2021). <https://doi.org/10.1038/s41567-021-01356-3>
68. Wang, S., Fontana, E., Cerezo, M., Sharma, K., Sone, A., Cincio, L., Coles, P.J.: Noise-induced barren plateaus in variational quantum algorithms. *Nat. Commun.* **12**(1), 6961 (2021). <https://doi.org/10.1038/s41467-021-27045-6>
69. Khatri, S., LaRose, R., Poremba, A., Cincio, L., Sornborger, A.T., Coles, P.J.: Quantum-assisted quantum compiling. *Quantum* **3**, 140 (2019). <https://doi.org/10.22331/q-2019-05-13-140>
70. Kokail, C., Maier, C., van Bijnen, R., Brydges, T., Joshi, M.K., Jurcevic, P., Muschik, C.A., Silvi, P., Blatt, R., Roos, C.F., Zoller, P.: Self-verifying variational quantum simulation of lattice models. *Nature* **569**(7756), 355–360 (2019). <https://doi.org/10.1038/s41586-019-1177-4>
71. Kardašin, A., Uvarov, A., Yudin, D., Biamonte, J.: Certified variational quantum algorithms for eigenstate preparation. *Phys. Rev. A* **102**, 052610 (2020). <https://doi.org/10.1103/PhysRevA.102.052610>
72. LaRose, R., Tikku, A., O’Neel-Judy, É., Cincio, L., Coles, P.J.: Variational quantum state diagonalization. *Quantum*. *npj Inform.* **5**(1), 57 (2019). <https://doi.org/10.1038/s41534-019-0167-6>
73. Cerezo, M., Sharma, K., Arrasmith, A., Coles, P.J.: Variational quantum state eigensolver. *npj Quantum Inf.* **8**(1), 113 (2022). <https://doi.org/10.1038/s41534-022-00611-6>
74. Francis, J.G.F.: The QR Transformation A Unitary Analogue to the LR Transformation—Part I. *Comput. J.* **4**, 265–271 (1961). <https://doi.org/10.1093/comjnl/4.3.265>

75. Francis, J.G.F.: The QR Transformation–Part 2. *Comput. J.* **4**, 332–345 (1962). <https://doi.org/10.1093/comjnl/4.4.332>
76. Kublanovskaya, V.N.: On some algorithms for the solution of the complete eigenvalue problem. *USSR Comput. Math. Math. Phys.* **1**, 637–657 (1962). [https://doi.org/10.1016/0041-5553\(63\)90168-X](https://doi.org/10.1016/0041-5553(63)90168-X)
77. Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen, D.: “LAPACK users’ guide”. Society for Industrial and Applied Mathematics. Philadelphia, PA (1999). Third edition. <https://doi.org/10.1137/1.97808987196604>
78. Demmel, J.W.: “Applied numerical linear algebra”. SIAM. (1997)
79. Stewart, Gilbert W.: “Matrix algorithms: Volume II: Eigensystems”. SIAM. (2001)
80. Björck, Åke: “Numerical methods in matrix computations”. Springer. (2015). <https://doi.org/10.1007/978-3-319-05089-8>
81. Kitaev, A.: Quantum measurements and the abelian stabilizer problem (1995). [arXiv:quant-ph/951102](https://arxiv.org/abs/quant-ph/951102)
82. Abrams, D.S., Lloyd, S.: Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Lett.* **83**, 5162–5165 (1999). <https://doi.org/10.1103/PhysRevLett.83.5162>
83. Cao, S., Wossnig, L., Vlastakis, B., Leek, P., Grant, E.: Cost-function embedding and dataset encoding for machine learning with parametrized quantum circuits. *Phys. Rev. A* **101**, 052309 (2020). <https://doi.org/10.1103/PhysRevA.101.052309>
84. Aharonov, D., Jones, V., Landau, Z.: A polynomial quantum algorithm for approximating the Jones polynomial”. In Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing. Page 427–436. STOC ’06New York, NY, USA (2006). Association for Computing Machinery. <https://doi.org/10.1145/1132516.1132579>
85. Passante, G., Moussa, O., Ryan, C.A., Laflamme, R.: Experimental approximation of the Jones polynomial with one quantum bit. *Phys. Rev. Lett.* **103**, 250501 (2009). <https://doi.org/10.1103/PhysRevLett.103.250501>
86. Araújo, M., Feix, A., Costa, F., Brukner, Č: Quantum circuits cannot control unknown operations. *New J. Phys.* **16**, 093026 (2014). <https://doi.org/10.1088/1367-2630/16/9/093026>
87. Liu, Y., Long, G.L., Sun, Y.: Analytic one-bit and CNOT gate constructions of general n-qubit controlled gates. *Int. J. Quantum Inf.* **6**, 447–462 (2008). <https://doi.org/10.1142/S0219749908003621C>
88. Mitarai, K., Fujii, K.: Methodology for replacing indirect measurements with direct measurements. *Phys. Rev. Res.* **1**, 013006 (2019). <https://doi.org/10.1103/PhysRevResearch.1.013006>
89. Alonso-Linaje, G., Atchade-Adelomou, P.: EVA: a quantum exponential value approximation algorithm (2021). <https://doi.org/10.48550/arXiv.2106.08731>
90. Bujiao, W., Ray, M., Zhao, L., Sun, X., Reberstrost, P.: Quantum-classical algorithms for skewed linear systems with an optimized Hadamard test. *Phys. Rev. A* **103**, 042422 (2021). <https://doi.org/10.1103/PhysRevA.103.042422>
91. Garcia-Escartin, J.C.: Example code for the algorithm. <https://github.tel.uva.es/juagar/QVEF> (2023)
92. Holmes, Z., Sharma, K., Cerezo, M., Coles, P.J.: Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum* **3**, 010313 (2022). <https://doi.org/10.1103/PRXQuantum.3.010313>
93. Yuxuan, D., Zhuozhuo, T., Yuan, X., Tao, D.: Efficient measure for the expressivity of variational quantum algorithms. *Phys. Rev. Lett.* **128**, 080506 (2022). <https://doi.org/10.1103/PhysRevLett.128.080506>
94. Funcke, L., Hartung, T., Jansen, K., Kühn, S., Stornati, P.: Dimensional expressivity analysis of parametric quantum circuits. *Quantum* **5**, 422 (2021). <https://doi.org/10.22331/q-2021-03-29-422>

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.