

Mitigating the measurement overhead of ADAPT-VQE with optimized informationally complete generalized measurements

Anton Nykänen ^{*}, Matteo A. C. Rossi , Elsi-Mari Borrelli, Sabrina Maniscalco, Guillermo García-Pérez, and Adam Glos
Algorithmiq Ltd, Kanavakatu 3 C, FI-00160 Helsinki, Finland



(Received 8 August 2023; accepted 16 September 2025; published 29 October 2025)

Adaptive derivative-assembled problem-tailored ansatz variational quantum eigensolver (ADAPT-VQE) stands out as a robust algorithm for constructing compact ansätze for molecular simulation. It enables to significantly reduce the circuit depth with respect to other methods, such as unitary coupled cluster singles and doubles, while achieving higher accuracy and not suffering from so-called barren plateaus that hinder the variational optimization of many hardware-efficient ansätze. In its standard implementation, however, it introduces a considerable measurement overhead in the form of gradient evaluations through estimations of many commutator operators. In this work, we mitigate this measurement overhead by exploiting a recently introduced method for energy evaluation relying on adaptive informationally complete generalized measurements (AIMs). Besides offering an efficient way to measure the energy itself, informationally complete (IC) measurement data can be reused to estimate all the commutators of the operators in the operator pool of ADAPT-VQE, using only classically efficient postprocessing. The framework is generic for any IC positive operator-valued measures (POVMs) implementation, and in this work, we demonstrate its effectiveness with dilation POVMs. We present the AIM-ADAPT-VQE scheme in detail, and investigate its performance with H_4 , N_2 , and 1,3,5,7-octatetraene Hamiltonians and several operator pools. Our numerical simulations indicate that the measurement data obtained to evaluate the energy can be reused to implement ADAPT-VQE with no additional measurement overhead for the systems considered here. In addition, we show that, if the energy is measured within chemical precision, the CNOT count in the resulting circuits is close to the ideal one. With scarce measurement data, AIM-ADAPT-VQE still converges to the ground state with high probability, albeit with an increased circuit depth in some cases.

DOI: [10.1103/t1hr-y7c8](https://doi.org/10.1103/t1hr-y7c8)

I. INTRODUCTION

Quantum computing holds the potential to solve problems that are intractable on classical computers. To unlock the full potential of the technology, error-corrected fault-tolerant quantum computing is needed, but even in the near term, the continuously improving quantum computing hardware may be able to demonstrate practical computational advantage for certain important tasks [1,2]. One of the fields in which near-term quantum computing hardware has the potential to provide advantage over classical state of the art is the simulation of quantum systems themselves [3]. Especially, obtaining quantum advantage in the context of electronic structure problems in quantum chemistry [4–6] and materials research [7] would carry significant scientific and industrial impact.

The general goal in quantum chemistry is to accurately estimate relevant properties of a molecular system described by a Hamiltonian H_f . To determine the relevant chemical properties, the energy of the ground state of H_f needs to be

accurately estimated. However, strong electronic correlations result in complex ground-state wave functions that are difficult to describe with classical means. For this reason, it was proposed that near-term quantum computers with the ability to prepare highly entangled states could be used to overcome these limitations [8]. Within this context, the most commonly employed tool for the quantum simulation of molecular systems is the hybrid quantum-classical algorithm variational quantum eigensolver (VQE), designed to find eigenvalues of quantum system Hamiltonians [9].

In VQE, one considers a family of states that depend on some parameters and are encoded in a parametrized quantum circuit. The circuit should be shallow enough to survive the short-coherence times of near-term quantum computers, and yet expressive enough to approximate the ground state for some values of its parameters [10]. For every value of the parameters, the resulting state represents the state of the electrons in a molecule and, as such, has an associated energy. To find the ground state of the system, the parameters of the circuit are changed iteratively as to minimize the energy of the state.

VQE has been successfully applied to small electronic structure problems [4,8,11], but there are several roadblocks that have so far prevented the scaling to practically relevant problem sizes [12]. Employing VQE for larger systems requires a robust methodology to design parametrized quantum circuits with the aforementioned properties.

^{*}Contact author: anton.nykanen@algorithmiq.fi

Early attempts relied on so-called chemistry-inspired ansätze, such as the unitary coupled cluster singles and doubles (UCCSD) [8], which can approximate the ground state relatively well. However, these circuits are generally too deep to be executed within the coherence times of current devices, and UCCSD and its many variations have been shown to perform well only for systems with relatively low correlations [13], which can be solved classically efficiently.

In order to limit the circuit depth, it was later on proposed to use hardware-efficient ansätze [4,14]; parametrized circuits that, by construction, are shallow enough to be executed on hardware. Yet, it has become apparent that only very limited accuracy can be achieved with these circuits. They typically exhibit landscapes of the energy cost function that are essentially flat almost everywhere in parameter space, rendering the optimization unfeasible; a phenomenon better known as barren plateaus [15,16].

The most promising approaches offering a solution to these problems are based on the adaptive derivative-assembled problem-tailored ansatz variational quantum eigensolver (ADAPT-VQE) [17]. In ADAPT-VQE, the circuit is built iteratively such that, in each iteration, a new unitary operator is chosen from a pool and added to the ansatz based on its energy gradient. Multiple variations of this algorithm have been studied, with different operator pools [18–21], operator selection methods [20,22,23], symmetry breaking and projections [24–26], and mutual-information assisted operator screening [27].

The adaptive ansatz construction both substantially limits the circuit depth and avoids the emergence of barren plateaus [28]. Moreover, ADAPT-VQE algorithms have been used for finding highly excited states [29] and quantum Gibbs thermal states [30]. However, identifying the best operator to be added at each step introduces a noticeable measurement cost, since the gradients of all the possible new operators must be estimated. Several works address this measurement overhead [22,31,32], in some cases with significant cost reductions, although with the caveat that they result in deeper circuits. Even in those cases, the overhead remains significant.

To use VQE for larger system sizes, besides the efficient construction of representative ansätze, it is necessary to perform the energy estimation within a reasonable runtime. In the standard formulation, VQE algorithms require a prohibitively high number of measurements necessary for accurate energy estimation [33]. In order to address this problem, many alternative measurement strategies have been proposed [4,34–49]. As a possible solution, some of the authors of this manuscript introduced an algorithm based on adaptive informationally complete measurements (AIMs) [50]. The method exploits the informational completeness of generalized measurements to optimize them as to minimize the statistical error of the estimation of the energy. In addition to reducing the total number of measurement rounds, also called shots, required to reach good accuracy in the estimation of the energy, the resulting IC data can be reused to estimate other expectation values using only classical postprocessing, that is, without the need to perform additional measurements.

In this article, we propose to exploit the informational completeness of the measurement data to mitigate the measurement overhead of ADAPT-VQE. More precisely, we

consider the situation in which, at some iteration, the energy is measured using AIM in order to optimize the parameters in the current ansatz and the data are then reused to compute all the gradients in ADAPT-VQE. While this is in principle possible by the nature of the IC data, we address the question of whether the precision of the gradient estimation is sufficient for this to perform in practice. To that end, we introduce a framework to combine adaptive IC measurements with ADAPT-VQE, which we name AIM-ADAPT-VQE. Our numerical simulations show that, with AIM-ADAPT-VQE, the measurement data obtained to determine the energy within chemical precision can be reused to implement ADAPT-VQE without increasing the CNOT counts for the systems analyzed. Interestingly, the ADAPT-VQE routine converges even if much less data are used, although sometimes with deeper circuits.

This article is organized as follows. We first briefly review ADAPT-VQE and the AIM methodology in Sec. II, and then introduce in detail the AIM-ADAPT-VQE algorithm in Sec. III. In Sec. IV, we test the algorithm with H_4 , N_2 , and 1,3,5,7-octatetraene Hamiltonians and discuss the performance of the method for different measurement accuracies. Section V discusses some implications of the results and possible future work.

II. BACKGROUND

In this section, we give a brief overview of the different algorithms that compose AIM-ADAPT-VQE. We first summarize the general idea behind VQE, introduce ADAPT-VQE and some of its variants and, then, discuss the energy-measurement cost problem along with the AIM method.

A. VQE in a nutshell

The general goal in quantum chemistry is to accurately estimate relevant properties of a molecular system described by a Hamiltonian:

$$H_f = \sum_{p,q} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{p,q,r,s} h_{pqrs} a_p^\dagger a_q^\dagger a_s a_r, \quad (1)$$

where $a_p^\dagger$ and a_p are the fermionic creation and annihilation operators, h_{pq} and h_{pqrs} are one- and two-electron integrals, and p, q, r , and s are spin-orbital indices. As stated in the previous section, a Hamiltonian of this form can present highly complex ground states, which makes them hard to simulate for large system sizes, that is, when they involve a large number of spin-orbitals N .

In order to use VQE to tackle the problem, one first maps the fermionic problem into qubit space using so-called fermion-to-qubit mappings to construct a multiqubit Hamiltonian H with the same spectrum as H_f , so that the problem can be solved by a multiqubit processor. Many fermion-to-qubit mappings exist [51–61], and they generally lead to different qubit operators H for the same fermionic Hamiltonian H_f , but the resulting operator can always be written as $H = \sum_k c_k S_k$, where $c_k \in \mathbb{R}$ and each $S_k \in \{\mathbb{1}, X, Y, Z\}^{\otimes N}$ is a Pauli string. The choice of mapping impacts several properties of the qubit operator H , which will be briefly discussed throughout the paper.

Once H has been obtained, one then chooses a parametrized quantum circuit, a family of states $|\psi(\vec{\theta})\rangle$ that depends on some parameters $\vec{\theta}$. The circuit can be executed on a quantum computer, and its corresponding energy $E(\vec{\theta}) = \langle \psi(\vec{\theta}) | H | \psi(\vec{\theta}) \rangle$ can thus be evaluated through a sequence of physical measurements. The energy $E(\vec{\theta})$ is then minimized with respect to parameters $\vec{\theta}$ until convergence. According to the variational principle, $E(\vec{\theta}) \geq E_0$, where E_0 is the ground state energy, so the minimization process brings the state $|\psi(\vec{\theta})\rangle$ closer to the ground state. In particular, if $E(\vec{\theta}_{\text{opt}}) = E_0$, then $|\psi(\vec{\theta}_{\text{opt}})\rangle$ is the ground state of H . If successfully implemented, the algorithm returns the ground-state energy E_0 , and further physical properties of the molecular system can be extracted from $|\psi(\vec{\theta}_{\text{opt}})\rangle$.

Chemical accuracy is a term used to describe measurement precision required to make realistic chemical predictions. Given that the basis sets used in our simulations are not accurate enough to completely describe the molecules to reach chemical accuracy, we prefer to use the term *chemical precision* to describe 1.6 mHa measurement error.

B. ADAPT-VQE

The key feature of all ADAPT-based strategies is that the circuit is built adaptively for the problem at hand, instead of having a fixed ansatz structure. At iteration n , the ansatz construction step in ADAPT-VQE adds a unitary operator $e^{\theta_i P_i}$ to the current circuit $|\psi^{(n-1)}(\vec{\theta}^{(n-1)})\rangle$, so that the resulting circuit prepares the states $|\psi^{(n)}(\vec{\theta}^{(n)})\rangle = e^{\theta_n P_n} |\psi^{(n-1)}(\vec{\theta}^{(n-1)})\rangle$, where $\vec{\theta}^{(n)}$ indicates the set of n parameters in the ansatz at iteration n . The operator P_i is chosen from a pool of operators that is defined beforehand. To determine which element should be chosen from the pool, one needs to perform a series of measurements for each of them, which results in the aforementioned overhead. Once the operator is added, the new variational circuit,

$$|\psi^{(n)}(\vec{\theta}^{(n)})\rangle = \prod_{i=1}^n e^{\theta_i P_i} |\psi^{(0)}\rangle, \quad (2)$$

where $|\psi^{(0)}\rangle$ is an initial reference state (usually the Hartree-Fock state), is variationally optimized with respect to the set of parameters $\vec{\theta}^{(n)}$. Many aspects of the algorithm can be varied. In the following, we will shortly describe some of the features and variants of ADAPT-VQE.

1. Operator pool

Defining a representative pool is very important, as it determines how well the circuit is able to explore the Hilbert space. The pool in the original ADAPT algorithm follows chemical intuition and consists of fermionic single and double excitations [17]. Moreover, the operators are spin complemented, which means that each operator excites both α and β spin-orbitals (up and down spins) and thus conserves spin symmetry. The fermionic pool used in this paper is the spin-dependent version with operators $\hat{\tau}_p^q = a_q^\dagger a_p - a_p^\dagger a_q$ and $\hat{\tau}_{rs}^{pq} = a_p^\dagger a_q^\dagger a_r a_s - a_s^\dagger a_r^\dagger a_q a_p$ as follows:

$$\mathcal{P} = \{ \tau_{p\alpha}^{q\alpha}, \tau_{p\beta}^{q\beta}, \tau_{r\alpha s\alpha}^{p\alpha q\alpha}, \tau_{r\beta s\beta}^{p\beta q\beta}, \tau_{r\alpha s\beta}^{p\alpha q\beta}, \tau_{r\beta s\alpha}^{p\beta q\alpha} \}, \quad (3)$$

where α and β indices indicate the spin orientation of the spin-orbital. This pool conserves S_z and N symmetries but can break the S^2 symmetry.

To apply fermionic operators in a quantum circuit, they need to be mapped into operators in qubit space using fermion-to-qubit transformations. The mappings considered in this paper are the Jordan-Wigner (JW) [51], Bravyi-Kitaev (BK) [52], and the JKMN [54] mappings. Due to the structure of the JW mapping, the fermionic operators in the pool result in qubit operators involving Pauli strings with many Z operators, the purpose of which is to keep track of the fermionic antisymmetries. It has been found, however, that removing these has a minimal effect on the performance. The resulting operator pool is called qubit-excitation-based (QEB) pool [19]. Note that this pool works only for the JW mapping and similar simplifications for other mappings are not straightforward.

A further simplification to the operators can be made by splitting the sums of the mapped fermionic operators (QEB operators with JW mapping), $\tau \mapsto \sum_k c_k S_k$, with $\tau \in \mathcal{P}$ and $S_k \in \{\mathbb{1}, X, Y, Z\}^{\otimes N}$, into separate terms. This creates the qubit-ADAPT pool, defined in qubits space in terms of the operators $P_k = iS_k$ [18]. This pool simplifies the gate complexity but increases the number of variational parameters in the ansatz.

2. Operator selection

The operator selection can be done in many ways. Ideally, the aim is to add gates that have maximum effect in taking the state toward the ground state of the system. A heuristic to this end is based on the gradient of the energy with respect to the parameter in the unitary generated by operator P_i when the parameter value is set to zero, which can be calculated by evaluating the commutator of P_i with the Hamiltonian,

$$\begin{aligned} \left. \frac{\partial E}{\partial \theta_i} \right|_{\theta_i=0} &= \left[\frac{\partial}{\partial \theta_i} \langle \psi^{(n)} | e^{-\theta_i P_i} H e^{\theta_i P_i} | \psi^{(n)} \rangle \right] \Big|_{\theta_i=0} \\ &= \langle \psi^{(n)} | [H, P_i] | \psi^{(n)} \rangle. \end{aligned} \quad (4)$$

Once the gradients for all the operators P_i in the pool have been obtained, the operator with the highest absolute value is added to the ansatz. Optionally, different strategies for gate ordering exist, when multiple gates are added in each step [22,23].

Instead of maximum gradient, another possible selection criterion is the maximal energy reduction. One can consider how much each operator decreases the energy by adding the gate to the current ansatz and optimizing the whole circuit. A more resource-efficient option is to only optimize the parameters of the trial operators. In both cases, the operator or operators that lower the energy the most are added to the circuit.

3. Optimization

There is further freedom in choosing the timing of the optimizations. In the original ADAPT algorithm, the circuit is optimized after adding each gate. To save resources, this can be relaxed and optimization of the whole circuit can be done only between certain periods. In order to advance the

optimization in the energy landscape and identify new operators to add, each newly added gate must be optimized individually. This can be done efficiently with *Rotosolve* or *ExcitationSolve* [62,63], which finds the optimal parameter for the gate with few measurements based on the sine wave form of the energy curve as a function of the gate parameter.

4. Measurement overhead

As discussed previously, the measurement cost associated with the adaptive strategies is large, as each iteration of ADAPT requires information about the expectation values of many operators. To analyze this in a more quantitative basis, notice that the number of operators in a pool, like the one in Eq. (3), contains $O(N^4)$ operators and, for each of them, an expectation value of the form in Eq. (4) must be measured. Moreover, the naive evaluation of an expectation value of an N -qubit operator requires measuring each of the Pauli strings conforming it. Thus, if we write the commutator in Eq. (4) as a linear combination of Pauli strings, $[H, P_i] = \sum_k c_k S_k$, we need to measure the expectation value of each S_k in the expansion

$$\langle \psi^{(n)} | [H, P_i] | \psi^{(n)} \rangle = \sum_k c_k \langle \psi^{(n)} | S_k | \psi^{(n)} \rangle. \quad (5)$$

Since the Hamiltonian H typically contains $O(N^4)$ Pauli strings itself, evaluating all the expectation values required in a single iteration of ADAPT-VQE can require measuring up to $O(N^8)$ Pauli strings. The energy selection criterion discussed previously is of course much more demanding than the gradient based in terms of measurements, because of the need to optimize the circuit for each operator candidate.

Restricting the operator pool size and adding the operators according to different batching strategies have been suggested as possible ways to reduce the number of measurements in ADAPT-VQE [18,22,31,32]. Such approaches allow up to a quadratic improvement in the number of measurements, but may result in deeper ansatz circuits if applied too greedily. It has also been shown that the $O(N^4)$ observables required in operator selection can be ordered into $O(N)$ mutually commuting sets to drastically reduce the required measurement resources [64]. While these approaches seem promising, they still require a dedicated measurement for the measured operators. Our approach allows for unbiased estimations of many observables by reusing informationally complete measurement data obtained primarily for estimating the energy. Next subsection explains this in detail.

C. Adaptive informationally complete measurements

The simple measurement method introduced in the previous subsection is expensive in terms of shots. Even if one is interested in estimating a single operator average, such as the energy of the system $\langle H \rangle$, the number of measurements required scales rapidly with the system size, especially for operators involving many Pauli strings. Since any VQE-based algorithm requires measuring the energy repeatedly, the excessive number of shots has been pointed out as a potential roadblock for real-world applications [33]. To overcome this issue, several ways to reduce the amount of measurements necessary to estimate an observable have been proposed

[4,36,38,65,66]. Using adaptive IC generalized measurements to improve the measurement scheme on the fly has been introduced as one solution [50]. The latter measurement algorithm, which we now review, was not only shown to reduce the energy estimation cost by orders of magnitude, but also to enable reusing the measurement data to evaluate other expectation values. Importantly, multiple implementations of informationally complete positive operator-valued measures (IC POVM) have been proposed for gate-based quantum computers, either using ancillary qubits [50], exploiting higher energy levels of the physical qubits [67], or by using randomized projective measurements [68]. In this paper, we consider only the class of POVMs used in Ref. [50].

For a single qubit i , these POVMs are generalized quantum measurements described by four so-called POVM effects, $\{\Pi_{m_i}\}$. That is, if the measurement is performed on a qubit in state $\rho^{(i)}$, the probability of obtaining outcome m_i is $p_{m_i} = \text{Tr}[\rho^{(i)} \Pi_{m_i}]$. These POVM effects are positive operators $\Pi_{m_i} > 0$, $m_i = 0, \dots, 3$ that add up to identity $\sum_{m_i} \Pi_{m_i} = \mathbb{1}$. If they are linearly independent, and therefore span the space of linear operators in the Hilbert space of the qubit, the POVM is called informationally complete. As such, any one qubit operator can be decomposed into the basis of the POVM effects.

Now, suppose every qubit in an N -qubit system is measured with a local POVM. The probability to obtain an outcome $\mathbf{m} = (m_1, \dots, m_N)$, where $m_i \in \{0, 1, 2, 3\}$, given a state ρ , is $\text{Tr}[\rho \Pi_{\mathbf{m}}]$, where $\Pi_{\mathbf{m}} = \bigotimes_{i=1}^N \Pi_{m_i}^{(i)}$. If each of the single-qubit POVMs is IC, the set of 4^N effects $\{\Pi_{\mathbf{m}}\}$ is IC in the space of linear operators of the N -qubit Hilbert space $\mathcal{H}^{\otimes N}$. In such case, any observable can be estimated with measurement data obtained from these POVMs, as we now show.

Since the effects $\{\Pi_{\mathbf{m}}\}$ form a basis, an arbitrary observable O can be decomposed as $O = \sum_{\mathbf{m}} \omega_{\mathbf{m}} \Pi_{\mathbf{m}}$, and its expectation value reads

$$\langle O \rangle = \text{Tr}[\rho O] = \sum_{\mathbf{m}} \omega_{\mathbf{m}} \text{Tr}[\rho \Pi_{\mathbf{m}}] = \sum_{\mathbf{m}} \omega_{\mathbf{m}} p_{\mathbf{m}}, \quad (6)$$

where $p_{\mathbf{m}}$ is the probability of obtaining outcome $\mathbf{m}$. This allows the estimation to be obtained with a Monte Carlo approach: One can sample from the probability distribution $\{p_{\mathbf{m}}\}$ with the quantum computer and average the $\omega_{\mathbf{m}}$ corresponding to each result.

The measurement is thus repeated S times and measurement outcomes $\mathbf{m}_1, \dots, \mathbf{m}_S$ are obtained. The expectation value of the operator can then be estimated as

$$\bar{O} = \frac{1}{S} \sum_{s=1}^S \omega_{\mathbf{m}_s}. \quad (7)$$

In essence, any observable can be measured in this manner with the same measurement data. Of course, the accuracy for the measurement will not be the same for all observables. The standard error in the estimation is given by $\sigma_O = [(\langle \omega_{\mathbf{m}}^2 \rangle_{\{p_{\mathbf{m}}\}} - \langle \omega_{\mathbf{m}} \rangle_{\{p_{\mathbf{m}}\}}^2) / S]^{1/2}$, which depends on the probability distribution $\{p_{\mathbf{m}}\}$ and the operator decomposition. Regarding the latter point, the standard error depends on the so-called Pauli weight of the operator (the number of nonidentity Pauli operators in the Pauli strings in its decomposition), and different

fermion-to-qubit mappings lead to different Pauli weights for the same fermionic operator [50,54]. Importantly, this approach enables producing accurate estimates of multiple observables with the same measurement data, even when the number of measurement shots S is much smaller than 4^N , the number of degrees of freedom in the states. In fact, as shown in Ref. [50], only a polynomial amount of measurements is required for typical few-electron operators, like the ones considered in this work, if an appropriate fermion-to-qubit mapping is used.

To decrease the standard error for a given observable, the probability distribution should be modified, which can be achieved by modifying the POVM. By parametrizing the POVM, one can optimize over the parameter space to find a measurement that lowers the variance of the estimation, and hence becomes increasingly precise. Moreover, this can be done by reusing the same measurement data that is used for the estimation of the observable. In practice, a measurement procedure is split into iterations, where after each small batch of measurements, the data are used to adjust the POVM and estimate desired observables on-the-fly. The process is repeated, and the number of shots in each round is updated according to a schedule, so that more measurements are used with more precise POVM [50].

Another similar measurement method, called classical shadows, can also be used to estimate many observables from very few measurements on the quantum state [45,46]. For these measurement schemes, many useful and rigorous performance guarantees have been presented. In addition, concrete bounds for shots required for specific estimations are given. Interestingly, local classical shadows are a special case of nondilation POVM, which can be optimized further to decrease the variance as shown in Ref. [68]. On the other hand, while global classical shadows are a different promising variant of IC-POVM, implementing this global variant requires applying a complicated Clifford operation, which is beyond the capabilities of NISQ devices, contrary to local IC-POVMs [69]. Classical shadows have been used in other works to estimate large number of observables required, for example, in quantum subspace expansion [70]. The authors of this paper compared randomized measurement techniques against other measurement techniques that utilize Hamiltonian data and wave function information to minimize the number of measurements and found the randomized classical shadow technique based on Majorana operator algebra to outperform these.

III. ESTIMATING ADAPT-VQE COMMUTATORS WITH IC DATA

In this section, we introduce AIM-ADAPT-VQE. As discussed in the previous sections, the fact that IC measurement data can be reused to evaluate multiple expectation values of observables using classically efficient postprocessing suggests a way to tackle the overhead from the measurements of the gradients of many operators in each iteration of ADAPT-VQE. In particular, if AIMS are used in the optimization of the parameters in the ansatz at each step (which would lead to a more efficient energy evaluation than with many other approaches), the IC data are automatically generated at no additional cost.

The idea is to then use Eq. (7) to estimate all the commutators required to implement the next step of ansatz construction in ADAPT-VQE. In cases in which the ADAPT-VQE schedule does not require the optimization of the current ansatz, some measurements will be needed, but we proceed in the same manner, that is, by optimizing the AIM to measure the energy.

An important aspect of the routine is how to determine whether enough measurements have been conducted. Typically, it is assumed that chemically precise (1.6 mHa) estimates of the energy are needed, since this is the level of precision required in real-world applications. However, in the operator selection and at intermediate stages of the variational circuit optimization, such precision might not be required. In order to investigate this question, we propose two different strategies: energy-based and gradient-based target precision. In the former, the AIM routine stops when the estimation of the energy is accurate enough. In the latter, the criterion to stop is based on the relative error in the estimation of the ADAPT commutators.

The steps of the AIM-ADAPT-VQE algorithm with energy-based measurement target precision, as shown also in Fig. 1, are the following:

- (1) Prepare the operator pool and choose the fermion-to-qubit mapping.
- (2) Initialize the ansatz. The most common choice is the Hartree-Fock state but other choices, like the Separable Pair Ansatz [71], may be used.
- (3) Measure the state with an initial amount of shots.
- (4) Estimate the energy and the standard error of the estimation.
- (5) Estimate every commutator $[H, \tau_i]$ and their corresponding standard errors.
- (6) Check if the standard error of the energy measurement is below the target threshold T_E . If so, skip to step 8.
- (7) Optimize the POVM, update the measurement strategy according to the schedule, and go back to step 3.
- (8) Check for convergence of the ADAPT-VQE routine based on the gradient norm of the operator pool or the variance of the energy, and stop the algorithm if it has converged.
- (9) Choose the operator with the largest gradient.
- (10) Find an optimal parameter for the new operator using *Rotosolve* or *ExcitationSolve*.
- (11) Add the gate to the ansatz.
- (12) If assigned in the schedule, optimize the whole circuit. Go to step 3.

In the gradient-based stopping version of the algorithm, to select the next operator, the goal is to be able to choose the one with the highest gradient while minimising the risk of choosing a wrong gate. Often, many operators have gradients that are very close to each other and, in these cases, choosing any of these usually gives similar results. The problem arises when the standard error of the estimation is close to or larger than the typical difference between gradients, as it is then likely for a wrong choice to be made. In order to prevent these situations, we continue measuring until the relative error in the largest gradient is smaller than the target threshold. In this alternative measurement strategy, step 6 is changed to the following:




FIG. 1. Flowchart of the AIM-ADAPT-VQE algorithm. At each step of the ADAPT algorithm, the state is prepared on the quantum computer and measured with an IC POVM. A first batch of shots is collected and it is used to estimate the energy of the state, as well as the gradients of the operators in the pool. The POVM is optimized, a new batch of shots is collected and the process is repeated until the target precision is reached. The operator with the largest gradient is added to the ansatz, with an initial parameter chosen with *Rotosolve* or *ExcitationSolve*. All the parameters in the ansatz are then optimized using a classical optimizer. Notice that this step requires additional measurements. However, in this work, the latter task is performed using exact simulations (that is, without shot noise), as the focus is in the estimation of the gradients of the operators in the ADAPT pool.

(6) Check if the relative error in the estimation of the commutator with the current highest mean is below the target threshold T_r . If so, skip to step 8.

The measurements in the AIM-ADAPT-VQE algorithm are performed while the POVM are continuously optimized, using the same shot schedule as in Ref. [50]. The performance of this optimization is evaluated after each small batch of measurements by considering the standard deviation of the energy estimation, defined as $\epsilon_E = [\langle \omega_{\mathbf{m}}^2 \rangle_{\{p_{\mathbf{m}}\}} - \langle \omega_{\mathbf{m}} \rangle_{\{p_{\mathbf{m}}\}}^2]^{1/2}$, where $\omega_{\mathbf{m}}$ are defined through the relation $H = \sum_{\mathbf{m}} \omega_{\mathbf{m}} \Pi_{\mathbf{m}}$, and $\{\Pi_{\mathbf{m}}\}$ are the POVM effects used with the last batch. Between optimization rounds, if ϵ_E increases twice in a row, the POVM is reset back two iterations. In addition, if there is an increase in ϵ_E for a total of 10 times during a measurement procedure, the POVM optimization is deemed converged and it is thus turned off, and the best POVM until that point is chosen for following measurements. In AIM-ADAPT-VQE, we also modify the gradient-based POVM optimization strategy. Instead of having a defined gradient step size, after the gradient direction is estimated, a line search is done in order to find the step size that decreases the variance the most. These modifications to the POVM optimization procedure result in an improved performance over the method employed in Ref. [50]. Finally, further resource savings may be obtained by recycling the POVM between ADAPT iterations [72]. The assumption is that, if the state has not changed significantly between two consecutive iterations, the optimal POVM for the previous round may be close to the optimal POVM for the next round. Therefore, we initialize the POVM based on this strategy.

While in the rest of our paper, we will use the dilation POVM [50], let us stress that AIM-ADAPT-VQE can be com-

bined with any measurement protocol that allows estimating multiple observables from a few measurements. This includes among other local and global classical shadows, which are a particular instances of IC-POVM [45,46], as well as a generalization of local classical shadows, i.e., nondilation POVM [68]. Eventually, we found choice of dilation POVM a well balanced choice because they simultaneously are applicable in the near-term quantum devices due to low-CNOT count, and they are optimizable for decreasing the variance of the estimation of multiple Paulis.

IV. RESULTS

In order to test AIM-ADAPT-VQE, we consider a 4-hydrogen chain H_4 with bond distance 1.5 Å, dinitrogen N_2 with equilibrium bond distance 1.098 Å and 1,3,5,7-octatetraene. These are good case studies, as they are fairly complex molecules, with nontrivial correlations, so that ADAPT-VQE does not converge in too few iterations.

The fermionic Hamiltonians are obtained with the EOS-PySCF implementation in Algorithmiq's software framework *Aurora* using STO-3G orbital basis set for the first two molecules and cc-pVDZ orbital basis set for 1,3,5,7-octatetraene, resulting in four, six and eight active spatial orbitals, respectively. The 1,3,5,7-octatetraene structure has been optimized with a DFT/PBE0/TZVP computational setup.

The H_4 fermionic Hamiltonian is then mapped into qubit space with JW, BK, and JKMN fermion-to-qubit transformations, which results in three different 8-qubit Hamiltonians. For N_2 , only JW mapping is used, which results in a 12-qubit Hamiltonian. For H_4 with JW mapping, the simulations are

run with QEB and qubit-ADAPT operator pools, and for BK and JKMN mappings, with spin-dependent fermionic and qubit-ADAPT pools. For N_2 and 1,3,5,7-octatetraene, only the spin-dependent fermionic pool and QEB-pool are used, respectively, to focus on the scaling of the method.

Each spin-dependent fermionic and qeb gate with JW mapping is implemented with an efficient circuit from Ref. [73] leading to a decrease in CNOT count. Spin-dependent fermionic operators for BK and JKMN mappings are implemented with the operator evolution in *Qiskit* [74] with one Trotter step. Each of the blocks in the Trotter approximation and the gates from qubit-ADAPT pool are implemented by *Qiskit*, resulting in the usual exponential map circuit containing staircase-like sequences of CNOT gates.

As previously explained, in this work, we focus on the ADAPT-VQE operator selection cost. Therefore, in order to avoid the computational cost of simulating the intermediate ansatz optimization steps with shot noise, the ansatz is optimized with exact simulations where the state vector is used to compute expectation values (in the following, we will refer to this kind of simulations as “statevector” simulations for conciseness) using the limited-memory Broyden-Fletcher-Goldfarb-Shanno bound (L-BFGS-B) optimizer [75]. To improve the speed and precision of the gradient calculations, the analytic gradient method introduced in the Supplementary Information of Ref. [17] is used. All ansatz parameters are optimized after each AIM-ADAPT-VQE iteration. Once the ansatz parameters are optimized, we perform the IC measurements to determine the next operator to be added with AIM-ADAPT-VQE, and record the number of measurements needed. The simulations presented in this paper have been obtained with the implementation of AIM-ADAPT-VQE in *Aurora*.

A. H_4 molecule

In Figs. 2(a)–2(c), simulations of AIM-ADAPT-VQE for all three fermion-to-qubit mappings, QEB and spin-dependent fermionic operator pools and different convergence criteria, as well as target measurement precision, are shown. Each line is an average over ten realizations, where each realization is an AIM-ADAPT-VQE simulation run until the error to the exact ground-state energy decreased below 1 mHa or the number of iterations reached 50. The relation between the measurement precision and the rate of convergence can be clearly seen. The simulations with higher-precision measurements converge faster, while the most imprecise ones tend to require more iterations.

Importantly, the ground state is reached up to chemical precision by nearly all simulations. In addition, by comparing them to the exact, infinite-statistics statevector simulation represented by the dashed black line, one can see that most of the realizations follow it quite closely. In all cases, a threshold of $T_E = 5$ mHa or lower for energy-based target measurement precision, or $T_r = 20\%$ or lower for relative error target precision in the gradient-based case, are required to converge in the same number of iterations as the exact simulation.

Measuring less precisely introduces randomness into the operator selection, making the runs behave differently, even amongst different simulations with the same target measurement precision. This can be seen from the increasing opacity

of some of the lines, which indicates that only a few runs fail to converge. In Fig. 2(a), the 10 and 8 mHa lines, as well as the 30% one in panel (c), seem to suffer the most from this.

The same simulations with qubit-ADAPT operator pools can be seen in Figs. 3(a)–3(c). The effect of the target measurement precision on the rate of convergence is much larger here. In all cases, an energy-based target measurement precision of $T_E = 1.6$ mHa or $T_r = 20\%$ relative error of operator gradient target measurement precision suffice to converge as with a statevector simulation. From all simulations with QEB, spin-dependent fermionic, and qubit-ADAPT operators pools, it can be concluded that the measurement scheme with relative error of operator gradient as target measurement precision seems to be a safer choice in ensuring that the convergence rate remains optimal.

The QEB and spin-dependent fermionic pools, which are chemically inspired, behave better with less precise measurements than the qubit-ADAPT pool. As a qubit-ADAPT pool operator is constructed by splitting a single excitation term into two or eight separate operators acting on the same qubits, one would expect significant redundancy in choosing the optimal gate and, therefore, that nonoptimal choices would not affect the performance much. Instead, as Fig. 3 shows, this is clearly not the case. Choosing different qubit-ADAPT operators that originate from the same fermionic excitation leads to different phases in statevector terms. The results indicate that nonoptimal choices lead to worse convergence rates.

An important metric to consider when constructing VQE ansätze is the number of gates required, as this affects both the execution time and, in near-term devices, the amount of noise introduced in the computation. In particular, the number of CNOT gates is a relevant figure of merit as these are typically noisier than single-qubit gates. In Figs. 2(g)–2(i), the resulting CNOT counts of each run are shown. Each CNOT count is calculated after compiling the circuit with *Qiskit*’s transpile function with the highest optimization level with all-to-all connectivity [74]. The performance can be evaluated by comparing the counts to the green line representing the CNOT count of an exact, statevector run. In all cases, measuring until chemical precision is sufficient to reach the same CNOT count as the statevector simulation up to small deviations.

In many cases, measuring until chemical precision is not required, and one can match the statevector simulation result with less precise measurements. However, measuring less precisely does introduce randomness and increases the probability of having a worse run. The level of this effect differs for different fermion-to-qubit mappings.

Interestingly, many of the runs reach lower CNOT counts than the exact statevector simulation. While the stochastic nature of the transpilation function employed may cause small deviations, this alone does not explain the result. Instead, this shows that the gate ordering obtained from an ideal ADAPT-VQE algorithm with gradient-based selection might not be optimal, which is somewhat expected given that ADAPT-VQE is a greedy algorithm. In any case, the improvement in the CNOT counts obtained here arises only from randomness in the operator selection due to imprecision in the gradient evaluations.

The analysis for different measurement precision levels allows us to examine how robust the AIM-ADAPT-VQE algorithm is to measurement data scarcity. In Figs. 2(d)–2(f)




FIG. 2. AIM-ADAPT-VQE simulations for different target measurement precisions for a hydrogen chain H_4 with interatomic distance $1.5 \text{ \AA}$, using different fermion-to-qubit mappings and pools: (a), (d), (g) JW mapping and QEB pool; (b), (e), (h) BK mapping and spin-dependent fermionic pool; (c), (f), (i) JKMN mapping and spin-dependent fermionic pool. All the results are obtained from ten realizations of each experiment. In this and following figures, the horizontal dotted line shows the chemical precision (1.6 mHa) with energies below it shaded gray. The top panels, (a), (b), (c), show the average error as a function of the ADAPT iteration. The dashed black line shows an exact statevector simulation. The error bars show one standard error distance from the average and the diamonds display outliers. The opacity of the lines is proportional to the number of runs that have not met the stopping criteria up until that point. The second row of panels, (d), (e), (f), show the average error as a function of the cumulative number of shots used by AIM-ADAPT-VQE. The bottom panels, (g), (h), (i), are box-and-whisker plots of the CNOT counts of the constructed circuits at the end of the AIM-ADAPT-VQE simulations. The boxes represent the distribution of CNOT counts over the realizations of each experiment. The dashed lines show the number of CNOTs obtained with a statevector simulation of ADAPT-VQE.

[as well as in Figs. 3(d)–3(f) for the qubit-ADAPT pool], the error in the energy as a function of the total number of measurements used is shown. Importantly, it can be seen that, if the measurement threshold, either T_E or T_τ , is set to a high value, the ground state can be reached with much fewer shots than needed simply for energy evaluation within chemical precision (which is indicated by the $T_E = 1.6 \text{ mHa}$ curve). However, this typically results in deeper circuits. In any case, the gradient-based target measurement threshold seems to consistently perform better, since resources are spared in the first iterations, where the gradients are large and hence precise measurements are not needed to identify the largest ones.

It is interesting to discuss the effect of the initialization of the POVM based on the optimality of the last round. In

Fig. 4, we depict the error in the estimation of the energy as a function of the number of shots for different ADAPT iterations in which the POVM at each iteration is initialized to match the optimal previous one. The results show that, after the first few iterations, the error does not decrease faster than the standard statistical inference curve, indicating that the optimization of the POVM no longer improves its performance. The erratic behavior in the first iteration in all cases, as well as in the second and third iterations in Fig. 4(c), are caused by the fact that some of the qubits remain in the computational basis in the first iterations, and after optimising the POVM corresponding to those qubits, they perform very poorly in the next iteration when they are recycled to be used on entangled qubits. In these cases, the POVMs are reset.




FIG. 3. AIM-ADAPT-VQE simulations for H_4 with interatomic distance $1.5 \text{ \AA}$, using the qubit-ADAPT pool and different fermion-to-qubit mappings: (a), (d), (g) JW mapping; (b), (e), (h) BK mapping; (c), (f), (i) JKMN mapping. The same quantities as in Fig. 2 are shown. Top panels: average error vs number of iterations; mid panels: average error vs cumulative number of measurement shots; and bottom panels: distribution of the CNOT counts of the transpiled final circuit.




FIG. 4. Optimization of the IC POVM for different AIM-ADAPT-VQE iterations. The plots show the estimated standard error in Ha as a function of the total number of shots used during steps 3–7 of the algorithm, for one run of the experiment. The panels correspond to those of Fig. 2, for the strategy $T_E = 1.6 \text{ mHa}$. The horizontal dotted line represents the chemical precision and the dashed line the shot-noise $1/\sqrt{\text{shots}}$ scaling. At each iteration, the optimal POVM of the previous iteration is used as the initial condition. In the first iterations, the lines decrease faster than the shot noise, meaning that the initial POVM is not optimal for the current ansatz and can be further refined by the algorithm. In later iterations, the lines follow the shot-noise line, i.e., the POVM of the previous iteration is already quite close to optimal.




FIG. 5. AIM-ADAPT-VQE simulations for N_2 with interatomic distance $1.098 \text{ \AA}$, using the fermionic pool and Jordan-Wigner fermion-to-qubit mapping. The same quantities as in Fig. 2 are shown: (a) average error vs number of iterations; (b) average error vs cumulative number of measurement shots; and (c) distribution of the CNOT counts of the transpiled final circuit.

The optimization of the POVM only requires classical resources. However, in the case when there is noise affecting the POVM implementation, quantum detector tomography can be used in order to reduce the bias induced by noise on the expectation values [68,76]. When optimising the POVM, detector tomography needs to be performed each time the POVM is changed. These results show that, since after a certain number of ADAPT iterations the POVM stops benefiting from the optimization, and hence the POVM optimization can be stopped, the overall cost of detector tomography could be heavily reduced.

B. N_2 molecule

A similar analysis is performed for the larger example, N_2 , with the results shown in Fig. 5. Again, most of the simulations reach chemical precision within a number of iteration close to the statevector simulation [as seen in Figs. 5(a) and 5(c)]. A threshold of $T_E = 3 \text{ mHa}$ or lower for energy-based target measurement precision is required for minimal iteration overhead. The relative-error-based target precision is again more efficient and $T_r = 20\%$ or below is required to match the statevector simulation. The $T_r = 20\%$ simulation convergences slightly faster than the $T_E = 3 \text{ mHa}$ one, but requires less measurements as seen in Fig. 5(b).

A critical point here is that the correct operator can be found without performing full-state tomography on the state each iteration. Full-state tomography requires inferring all 4^N components in the system's density matrix, and therefore the number of measurements needed in state tomography is typically much larger than 4^N . In the case of a 12-qubit system like the one considered here, there are approximately 10^7 elements in the system's density matrix. Yet, from Fig. 5(b), it can be seen that each ADAPT iteration requires much fewer shots than that. While for H_4 , the measurement cost for the energy measurement is in the order of 4^N ; for N_2 , it is already considerably below it. This indicates that as the system size increases, the required measurements cost for operator selection remains well below the tomographic regime, in agreement with theoretical expectations [50].

C. 1,3,5,7-octatetraene molecule

Third analysis is performed on the 1,3,5,7-octatetraene molecule with the results shown in Fig. 6. Instead of showing

a full simulation of the AIM-ADAPT-VQE process, the results show a single-operator selection procedure. The procedure was performed on an ansatz generated with 20 iterations of regular ADAPT-VQE with statevector simulation using Jordan-Wigner mapping and QEB pool. The statevector error of the ansatz with 20 pool elements is 2.28 mHa , which is close to chemical precision. This approach was chosen to save computational resources. In the first column, results for $T_E = 1.6 \text{ mHa}$ energy-based target measurement precision are shown. The top figure shows a close match between operator energy gradients estimated with state vector and AIM procedure, and the bottom figure confirms that the operator rankings ordered using state vector and AIM estimations are almost the same. We conclude that AIM estimation can be used to reliably choose an operator that statevector estimations would also rank highly. Same conclusion can be obtained from the second column, which shows the results for $T_r = 10\%$ gradient-based target measurement precision. Slightly larger deviation can be seen in the third column, which shows results for $T_r = 20\%$ gradient-based target measurement precision. Still, the top four- ranked operators are the same set in both rankings.

V. CONCLUSIONS AND OUTLOOK

ADAPT-VQE algorithms provide an efficient toolbox to construct accurate and shallow circuits specific to the electronic structure problem. A noticeable drawback of the original algorithm is the measurement overhead arising from the gradient calculations that need to be performed in each iteration. Here, we have shown that this overhead can be mitigated, or even bypassed, by measuring with an adaptive IC POVM.

By reusing the IC data obtained in energy measurement steps needed to optimize the ansatz, it is possible to estimate all the operator gradients with no additional measurement cost. Furthermore, for increasing system size, this selection can be done with a number of shots that is well below the tomographic regime. For the examples studied in this work, the data collected for energy estimations within chemical precision sufficed to implement different versions of ADAPT-VQE without changing the CNOT counts of the original algorithm. In some cases, the counts were even lower than with exact simulations.




FIG. 6. AIM-ADAPT-VQE simulations for 1,3,5,7-octatetraene, using the QEB pool and Jordan-Wigner fermion-to-qubit mapping. Results are shown for $T_E = 1.6$ mHa energy-based target measurement precision, $T_r = 10\%$ and $T_r = 20\%$ gradient-based target measurement precisions in their respective columns. Top row compares the energy gradients estimated with the different methods for top ten operator candidates ranked. Bottom row compares the rankings of candidate operators when ordered by the gradients estimated by the different methods. We can see particularly a very good agreement between the ranks computed with AIM and statevector.

In addition, we have shown that AIM-ADAPT-VQE is rather robust in terms of measurement precision and performs reasonably well even with less precise measurements. Some of the fermion-to-qubit mappings and operator pools considered were found to be more robust than others. Interestingly, even when very few measurements are performed, the AIM-ADAPT-VQE routine eventually converges, albeit with higher CNOT counts. This trade-off between measurement overhead and circuit complexity can be exploited to optimize the implementation of algorithms of the ADAPT family on slow but noise resilient hardware platforms, such as ion-trap-based systems.

Using dilation POVMs doubles the number of required qubits. Since AIM-ADAPT-VQE is well defined for any informationally complete POVM, one could consider other measurement protocols instead like nondilation POVM [68]. We leave the numerical investigation of whether other POVM used would show similar reduction in the shots required for future investigation.

The IC-based implementation of ADAPT algorithms offers other interesting and potentially useful ways to further optimize the ansatz construction. For example, if IC data are available, it is possible to evaluate the energy of the system once a given gate is appended to the circuit using only classical postprocessing. This could open the path to a different operator selection strategy without additional mea-

surement cost: instead of choosing the operators based on energy gradients, they could be chosen based on their impact on the energy once optimised.

Further, for AIM-ADAPT-VQE, the only computational cost that increases with the size of the pool is, by construction, classical (assuming similar algebraic properties in the pool operators). This property motivates the search of larger operator pools that may be able to produce accurate ansätze with even smaller CNOT counts. This approach is in stark contrast to most works related to operator pool selection, where the aim has been to reduce measurement costs by reducing the pool sizes without compromising performance.

The combination of ADAPT-VQE algorithms and adaptive informationally complete measurements provides a recipe for performing accurate quantum simulations with shallow circuits and reasonable measurement cost scaling. We see this as an important step toward performing quantum simulations of systems with increasing system size, and therefore toward near-term quantum advantage for quantum chemistry.

DATA AVAILABILITY

The data and code that support the findings of this article will be provided upon reasonable request.

- [1] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [2] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, The future of quantum computing with superconducting qubits, *J. Appl. Phys.* **132**, 160902 (2022).
- [3] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, *Nature (London)* **607**, 667 (2022).
- [4] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
- [5] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, and I. Tavernelli, Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions, *Phys. Rev. A* **98**, 022322 (2018).
- [6] P. J. Ollitrault, A. Kandala, C.-F. Chen, P. K. Barkoutsos, A. Mezzacapo, M. Pistoia, S. Sheldon, S. Woerner, J. M. Gambetta, and I. Tavernelli, Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor, *Phys. Rev. Res.* **2**, 043140 (2020).
- [7] V. Lordi and J. M. Nichol, Advances and opportunities in materials science for scalable quantum computing, *MRS Bull.* **46**, 589 (2021).
- [8] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [9] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
- [10] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [11] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [12] L. Bittel and M. Kliesch, Training variational quantum algorithms is NP-hard, *Phys. Rev. Lett.* **127**, 120502 (2021).
- [13] J. Lee, W. J. Huggins, M. Head-Gordon, and K. B. Whaley, Generalized unitary coupled cluster wave functions for quantum computation, *J. Chem. Theory Comput.* **15**, 311 (2019).
- [14] K. Mitarai, T. Yan, and K. Fujii, Generalization of the output of a variational quantum eigensolver by parameter interpolation with a low-depth ansatz, *Phys. Rev. Appl.* **11**, 044087 (2019).
- [15] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [16] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [17] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer, *Nat. Commun.* **10**, 3007 (2019).
- [18] H. L. Tang, V. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou, Qubit-ADAPT-VQE: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor, *PRX Quantum* **2**, 020310 (2021).
- [19] Y. S. Yordanov, V. Armaos, C. H. W. Barnes, and D. R. M. Arvidsson-Shukur, Qubit-excitation-based adaptive variational quantum eigensolver, *Commun. Phys.* **4**, 228 (2021).
- [20] H. G. A. Burton, D. Marti-Dafcik, D. P. Tew, and D. J. Wales, Exact electronic states with shallow quantum circuits from global optimisation, *npj Quantum Inf.* **9**, 75 (2023).
- [21] J. S. Van Dyke, G. S. Barron, N. J. Mayhall, E. Barnes, and S. E. Economou, Scaling adaptive quantum simulation algorithms via operator pool tiling, *Phys. Rev. Res.* **6**, L012030 (2024).
- [22] Z. Lan and W. Liang, Amplitude reordering accelerates the adaptive variational quantum eigensolver algorithms, *J. Chem. Theory Comput.* **18**, 5267 (2022).
- [23] P. G. Anastasiou, Y. Chen, N. J. Mayhall, E. Barnes, and S. E. Economou, TETRIS-ADAPT-VQE: An adaptive algorithm that yields shallower, denser circuit Ansätze, *Phys. Rev. Res.* **6**, 013254 (2024).
- [24] L. W. Bertels, H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, Symmetry breaking slows convergence of the adapt variational quantum eigensolver, *J. Chem. Theory Comput.* **18**, 6656 (2022).
- [25] V. O. Shkolnikov, N. J. Mayhall, S. E. Economou, and E. Barnes, Avoiding symmetry roadblocks and minimizing the measurement overhead of adaptive variational quantum eigensolvers, *Quantum* **7**, 1040 (2023).
- [26] T. Tsuchimochi, M. Taii, T. Nishimaki, and S. L. Ten-No, Adaptive construction of shallower quantum circuits with quantum spin projection for fermionic systems, *Phys. Rev. Res.* **4**, 033100 (2022).
- [27] Z.-J. Zhang, T. H. Kyaw, J. S. Kottmann, M. Degroote, and A. Aspuru-Guzik, Mutual information-assisted adaptive variational quantum eigensolver, *Quantum Sci. Technol.* **6**, 035001 (2021).
- [28] H. R. Grimsley, G. S. Barron, E. Barnes, S. E. Economou, and N. J. Mayhall, Adaptive, problem-tailored variational quantum eigensolver mitigates rough parameter landscapes and barren plateaus, *npj Quantum Inf.* **9**, 19 (2023).
- [29] F. Zhang, N. Gomes, Y. Yao, P. P. Orth, and T. Iadecola, Adaptive variational quantum eigensolvers for highly excited states, *Phys. Rev. B* **104**, 075159 (2021).
- [30] A. Warren, L. Zhu, N. J. Mayhall, E. Barnes, and S. E. Economou, Adaptive variational algorithms for quantum Gibbs state preparation, *arXiv:2203.12757*.
- [31] J. Liu, Z. Li, and J. Yang, An efficient adaptive variational quantum solver of the Schrödinger equation based on reduced density matrices, *J. Chem. Phys.* **154**, 244112 (2021).
- [32] M. D. Sapova and A. K. Fedorov, Variational quantum eigensolver techniques for simulating carbon monoxide oxidation, *Commun. Phys.* **5**, 199 (2022).
- [33] J. F. Gonthier, M. D. Radin, C. Buda, E. J. Dosekocil, C. M. Abuan, and J. Romero, Measurements as a roadblock to near-term practical quantum advantage in chemistry: Resource analysis, *Phys. Rev. Res.* **4**, 033154 (2022).
- [34] A. Jena, S. Genin, and M. Mosca, Pauli partitioning with respect to gate sets *arXiv:1907.07859*.
- [35] T.-C. Yen, V. Verteletskyi, and A. F. Izmaylov, Measuring all compatible operators in one series of single-qubit measurements using unitary transformations, *J. Chem. Theory Comput.* **16**, 2400 (2020).

- [36] W. J. Huggins, J. R. McClean, N. C. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, *npj Quantum Inf.* **7**, 23 (2021).
- [37] P. Gokhale, O. Angiuli, Y. Ding, K. Gui, T. Tomesh, M. Suchara, M. Martonosi, and F. T. Chong, Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families, [arXiv:1907.13623](https://arxiv.org/abs/1907.13623).
- [38] O. Crawford, B. v. Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, Efficient quantum measurement of Pauli operators in the presence of finite sampling error, *Quantum* **5**, 385 (2021).
- [39] M. Pains and A. Kalev, An approximate description of quantum states, [arXiv:1910.10543](https://arxiv.org/abs/1910.10543).
- [40] X. Bonet-Monroig, R. Babbush, and T. E. O'Brien, Nearly optimal measurement scheduling for partial tomography of quantum states, *Phys. Rev. X* **10**, 031064 (2020).
- [41] J. Cotler and F. Wilczek, Quantum overlapping tomography, *Phys. Rev. Lett.* **124**, 100401 (2020).
- [42] I. Hamamura and T. Imamichi, Efficient evaluation of quantum observables using entangled measurements, *npj Quantum Inf.* **6**, 56 (2020).
- [43] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [44] G. Torlai, G. Mazzola, G. Carleo, and A. Mezzacapo, Precise measurement of quantum observables with neural-network estimators, *Phys. Rev. Res.* **2**, 022060 (2020).
- [45] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [46] H.-Y. Huang, R. Kueng, and J. Preskill, Efficient estimation of Pauli observables by derandomization, *Phys. Rev. Lett.* **127**, 030503 (2021).
- [47] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, Measurements of quantum Hamiltonians with locally-biased classical shadows, *Commun. Math. Phys.* **391**, 951 (2022).
- [48] Y. Quek, S. Fort, and H. K. Ng, Adaptive quantum state tomography with neural networks, *npj Quantum Inf.* **7**, 105 (2021).
- [49] D. Miller, L. E. Fischer, K. Levi, E. J. Kuehnke, I. O. Sokolov, P. Kl. Barkoutsos, J. Eisert, and I. Tavernelli, Hardware-tailored diagonalization circuits, *npj Quantum Inf.* **10**, 122 (2024).
- [50] G. García-Pérez, M. A. Rossi, B. Sokolov, F. Tacchino, P. K. Barkoutsos, G. Mazzola, I. Tavernelli, and S. Maniscalco, Learning to measure: Adaptive informationally complete generalized measurements for quantum algorithms, *PRX Quantum* **2**, 040342 (2021).
- [51] P. Jordan and E. Wigner, Über das Paulische Äquivalenzverbot, *Z. Phys.* **47**, 631 (1928).
- [52] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, *Ann. Phys.* **298**, 210 (2002).
- [53] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, Tapering off qubits to simulate fermionic Hamiltonians, [arXiv:1701.08213](https://arxiv.org/abs/1701.08213).
- [54] Z. Jiang, A. Kalev, W. Mruczkiewicz, and H. Neven, Optimal fermion-to-qubit mapping via ternary trees with applications to reduced quantum states learning, *Quantum* **4**, 276 (2020).
- [55] M. Chiew and S. Strelchuk, Discovering optimal fermion-qubit mappings through algorithmic enumeration, *Quantum* **7**, 1145 (2023).
- [56] K. Setia and J. D. Whitfield, Bravyi-Kitaev superfast simulation of electronic structure on a quantum computer, *J. Chem. Phys.* **148**, 164104 (2018).
- [57] M. Steudtner and S. Wehner, Fermion-to-qubit mappings with varying resource requirements for quantum simulation, *New J. Phys.* **20**, 063010 (2018).
- [58] V. Havlíček, M. Troyer, and J. D. Whitfield, Operator locality in the quantum simulation of fermionic models, *Phys. Rev. A* **95**, 032332 (2017).
- [59] R. W. Chien and J. D. Whitfield, Custom fermionic codes for quantum simulation, [arXiv:2009.11860](https://arxiv.org/abs/2009.11860).
- [60] M. Steudtner and S. Wehner, Quantum codes for quantum simulation of fermions on a square lattice of qubits, *Phys. Rev. A* **99**, 022308 (2019).
- [61] A. Miller, Z. Zimborás, S. Knecht, S. Maniscalco, and G. García-Pérez, Bonsai algorithm: Grow your own fermion-to-qubit mappings, *PRX Quantum* **4**, 030314 (2023).
- [62] M. Ostaszewski, E. Grant, and M. Benedetti, Structure optimization for parameterized quantum circuits, *Quantum* **5**, 391 (2021).
- [63] J. Jäger, T. N. Kaldenbach, M. Haas, and E. Schultheis, Fast gradient-free optimization of excitations in variational quantum eigensolvers, [arXiv:2409.05939](https://arxiv.org/abs/2409.05939) [quant-ph].
- [64] P. G. Anastasiou, N. J. Mayhall, E. Barnes, and S. E. Economou, How to really measure operator gradients in ADAPT-VQE, [arXiv:2306.03227](https://arxiv.org/abs/2306.03227) [quant-ph].
- [65] A. F. Izmaylov, T.-C. Yen, R. A. Lang, and V. Verteletskiy, Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method, *J. Chem. Theory Comput.* **16**, 190 (2020).
- [66] A. Zhao, A. Tranter, W. M. Kirby, S. F. Ung, A. Miyake, and P. J. Love, Measurement reduction in variational quantum algorithms, *Phys. Rev. A* **101**, 062322 (2020).
- [67] L. E. Fischer, D. Miller, F. Tacchino, P. K. Barkoutsos, D. J. Egger, and I. Tavernelli, Ancilla-free implementation of generalized measurements for qubits embedded in a qudit space, *Phys. Rev. Res.* **4**, 033027 (2022).
- [68] A. Glos, A. Nykänen, E.-M. Borrelli, S. Maniscalco, M. A. C. Rossi, Z. Zimborás, and G. García-Pérez, Adaptive POVM implementations and measurement error mitigation strategies for near-term quantum devices, [arXiv:2208.07817](https://arxiv.org/abs/2208.07817).
- [69] J. Malmi, K. Korhonen, D. Cavalcanti, and G. García-Pérez, Enhanced observable estimation through classical optimization of informationally overcomplete measurement data: Beyond classical shadows, *Phys. Rev. A* **109**, 062412 (2024).
- [70] S. Choi and A. F. Izmaylov, Measurement optimization techniques for excited electronic states in near-term quantum computing algorithms, [arXiv:2302.11421](https://arxiv.org/abs/2302.11421) [quant-ph].
- [71] J. S. Kottmann and A. Aspuru-Guzik, Optimized low-depth quantum circuits for molecular electronic structure using a separable-pair approximation, *Phys. Rev. A* **105**, 032449 (2022).
- [72] K. Korhonen, Exploring the use of adaptive POVM measurements in the VQE (2022), <http://hdl.handle.net/10138/345159>.
- [73] Y. S. Yordanov, D. R. M. Arvidsson-Shukur, and C. H. W. Barnes, Efficient quantum circuits for quantum computational chemistry, *Phys. Rev. A* **102**, 062612 (2020).
- [74] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W.

- Cross, B. R. Johnson, and J. M. Gambetta, Quantum computing with Qiskit, [arXiv:2405.08810](#) [quant-ph].
- [75] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization, *ACM Trans. Math. Softw.* **23**, 550 (1997).
- [76] M. Cattaneo, M. A. C. Rossi, K. Korhonen, E.-M. Borrelli, G. García-Pérez, Z. Zimborás, and D. Cavalcanti, Self-consistent quantum measurement tomography based on semidefinite programming, *Phys. Rev. Res.* **5**, 033154 (2023).