

A Stabilizer Framework for the Contextual Subspace Variational Quantum Eigensolver and the Noncontextual Projection Ansatz

Tim Weaving,* Alexis Ralli, William M. Kirby, Andrew Tranter, Peter J. Love, and Peter V. Coveney*



Cite This: *J. Chem. Theory Comput.* 2023, 19, 808–821



Read Online

ACCESS |



Metrics & More

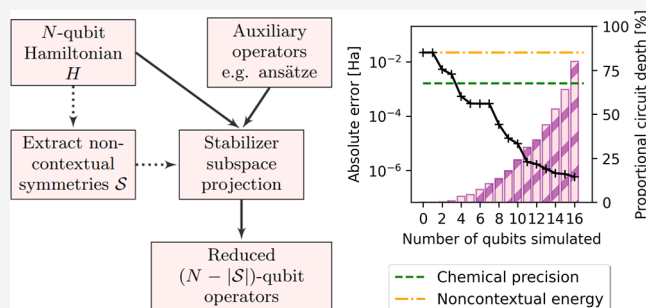


Article Recommendations



Supporting Information

ABSTRACT: Quantum chemistry is a promising application for noisy intermediate-scale quantum (NISQ) devices. However, quantum computers have thus far not succeeded in providing solutions to problems of real scientific significance, with algorithmic advances being necessary to fully utilize even the modest NISQ machines available today. We discuss a method of ground state energy estimation predicated on a partitioning of the molecular Hamiltonian into two parts: one that is *noncontextual* and can be solved classically, supplemented by a *contextual* component that yields quantum corrections obtained via a Variational Quantum Eigensolver (VQE) routine. This approach has been termed *Contextual Subspace VQE* (CS-VQE); however, there are obstacles to overcome before it can be deployed on NISQ devices. The problem we address here is that of the ansatz, a parametrized quantum state over which we optimize during VQE; it is not initially clear how a splitting of the Hamiltonian should be reflected in the CS-VQE ansatz. We propose a “noncontextual projection” approach that is illuminated by a reformulation of CS-VQE in the stabilizer formalism. This defines an ansatz restriction from the full electronic structure problem to the contextual subspace and facilitates an implementation of CS-VQE that may be deployed on NISQ devices. We validate the noncontextual projection ansatz using a quantum simulator and demonstrate chemically precise ground state energy calculations for a suite of small molecules at a significant reduction in the required qubit count and circuit depth.



1. INTRODUCTION

Quantum computers promise to yield solutions to complex problems that have previously been unattainable by classical means, yet experimental demonstration remains challenging. To date, some of the largest molecules simulated on noisy intermediate-scale quantum (NISQ) hardware are H_{12} —albeit only a Hartree–Fock calculation—conducted by Google using just 12 of the 53 qubits available on their superconducting quantum processor *Sycamore*,¹ and H_2O performed independently by IonQ using 3 qubits of an unspecified proprietary trapped ion device² and by IBM using 5 of the 27 qubits on the now-decommissioned *ibmq_dublin* superconducting device.³

Due to the limitations of short coherence times, restrictive qubit connectivity and high noise floors that characterize the NISQ era, we are not able to harness the full state-space afforded to these machines. To circumvent the above issues, we turn to the class of variational quantum algorithms, of which the Variational Quantum Eigensolver (VQE)⁴ is most widely studied. In contrast with eigenvalue-finding algorithms requiring fault-tolerant machines such as Quantum Phase Estimation (QPE),⁵ which necessitates state evolution over an extended period of coherence, VQE executes a large ensemble of comparatively shallow parametrized circuits to estimate energy expectation values, informing a classical optimizer that updates the parameter settings before reinitialization of the


quantum circuit. Its success is predicated on the variational principle, meaning the ground state energy of the system bounds expectation values from below.⁶

However, VQE is not without its challenges. First of all, the parametrized quantum state mentioned above, known as an *ansatz*, needs to be constructed carefully: It must be sufficiently expressible so the subspace of quantum states it spans contains the true ground state. On the other hand, if the ansatz is too expressible, we run into the problem of barren plateaus⁷ where we observe vanishing gradients. This is more often a symptom of “hardware efficient” ansätze,^{1,8–12} which aim to access the largest possible region of Hilbert space for the fewest number of native quantum gates.

To avoid barren plateaus, one must take into account some of the underlying problem structure to define ansatz circuits whose images are confined to a smaller but more targeted region of Hilbert space. Within this category are “chemically

Received: September 5, 2022

Published: January 23, 2023



inspired" ansätze that represent sequences of electronic excitation operators in circuit; unitary coupled cluster (UCC)^{13,14} is widely acknowledged as the gold standard for electronic structure simulations, albeit computationally very expensive in practice.

More recently, we have seen the development of hybrid ansätze that bridge the gap between hardware efficiency and chemical motivation. For example, Gard et al.¹⁵ designed a compact circuit designed to conserve molecule symmetries such as particle number and spin, while Adaptive Derivative-Assembled Pseudo-Trotter (ADAPT) VQE^{16–19} describes a more complete approach to scalable quantum chemistry simulations by defining selection criteria of ansatz terms from a pool of excitation operators.

Second, the energy estimation procedure in VQE invokes the measurement problem; in order to mitigate statistical error, many prepare-and-measure cycles are necessary to achieve sufficient precision in the estimate. The advances made in recent years toward measurement reduction techniques are expansive^{20–29} and range from classical pre/postprocessing of the measurement information, such as in classical shadow tomography,^{30,31} to Hamiltonian term-grouping schemes and reductions in the number of Hamiltonian terms at a cost of coherent resource, such as in unitary partitioning.^{32–36} Combined with techniques of error mitigation,^{37–44} one can optimize VQE with the objective of maximal NISQ resource utilization.

In this work, we are concerned with Contextual Subspace VQE (CS-VQE),⁴⁵ which describes a method of partitioning the molecular Hamiltonian into disjoint parts so that an electronic structure problem may be simulated to some degree on the available quantum device, even when the dimension of the full problem is too great to be encoded on the number of qubits available. This is supplemented by some classical overhead, but this often permits one to achieve chemical precision (to within 1.6 mHa \approx 4 kJ/mol of the full configuration interaction (FCI) energy) at a saving of qubits, as indicated by Kirby et al.⁴⁵ We highlight the relevance of chemical precision over accuracy here; since a minimal basis set (STO-3G) will be used throughout our benchmark, one should not expect agreement with experimentally obtained molecular energies and thus chemical accuracy is not an appropriate phrase. A finer basis set such as cc-pVxZ where $x = D, T, Q$, etc. should be used if one wishes to assess accuracy; however, this comes at the cost of increased qubits.

There has since been further research into the use of classical estimates of the electronic structure problem to reduce the resource requirements on quantum hardware. In particular, Classically Boosted VQE (CB-VQE)⁴⁶ identifies classically tractable states and excludes them from the quantum simulation, alleviating some measurement and fidelity requirements of the VQE routine. CS-VQE also bears a resemblance to the qubit reduction technique of qubit tapering,^{47,48} which exploits \mathbb{Z}_2 symmetries of the Hamiltonian; the differences and similarities are highlighted herein and by Kirby et al.⁴⁵

There are still a number of problems to address before CS-VQE may be successfully deployed on real quantum hardware, most notably with regard to the ansatz, which is the principal focus of this work. To aid this objective, we place the method on a strong theoretical footing of stabilizer subspaces and projections therein; this reformulation is better suited to efficient implementation, which is being addressed through the Symmer project.⁴⁹ This rephrasing of CS-VQE illuminates the

matter of constructing ansätze for the contextual subspace and renders this method compatible with contemporary approaches to ansatz construction such as ADAPT-VQE.

2. PRELIMINARIES

The notation used throughout shall be to write operators in standard capital font (A, B, C , etc.), with the exception of single-qubit Pauli operators being written in the form

$$\sigma_p := \begin{pmatrix} \delta_{p,0} + \delta_{p,3} & \delta_{p,1} - i\delta_{p,2} \\ \delta_{p,1} + i\delta_{p,2} & \delta_{p,0} - \delta_{p,3} \end{pmatrix} \quad (1)$$

for $p \in \{0, 1, 2, 3\}$. Sets are denoted by script letters ($\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots$) and vector spaces by bold script typeface ($\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots$). The state space of N qubits may be identified with the 2^N dimensional Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes N}$, with the space of (bounded) linear operators acting upon \mathcal{H} denoted $\mathcal{B}(\mathcal{H})$.

We introduce the Pauli group, $\mathcal{P}_N \subset \mathcal{B}(\mathcal{H})$, consisting of operators $\sigma = \otimes_{i=0}^{N-1} \sigma_{p_i}$ for $p_i \in \{0, 1, 2, 3\}$, up to multiplication by $\pm 1, \pm i$. Note the distinction between the bold font σ denoting tensor products and σ_p a single-qubit Pauli operator; we will sometimes write $\sigma_p^{(i)}$ to index explicitly the qubit position $i \in \mathbb{Z}_N$ on which it acts. We shall also make use of the commutator $[A, B] := AB - BA$ and anticommutator $\{A, B\} := AB + BA$, defined for operators $A, B \in \mathcal{B}(\mathcal{H})$, which are zero when A and B commute/anticommute, respectively.

An N -qubit Hamiltonian can be written in the form

$$H_{\mathcal{T}} = \sum_{\sigma \in \mathcal{T}} h_{\sigma} \sigma, \quad h_{\sigma} \in \mathbb{R} \quad (2)$$

for a set of Pauli operators $\mathcal{T} \subset \mathcal{P}_N$; specifying real coefficients ensures that $H_{\mathcal{T}}$ is Hermitian. The objective of quantum chemistry simulations is to estimate the ground state energy

$$\varepsilon_0 := \min_{|\psi\rangle \in \mathcal{H}} \langle H_{\mathcal{T}} \rangle_{\psi} \quad (3)$$

where $\langle H_{\mathcal{T}} \rangle_{\psi} := \langle \psi | H_{\mathcal{T}} | \psi \rangle$ is the expectation value of $H_{\mathcal{T}}$ with respect to some quantum state $|\psi\rangle \in \mathcal{H}$. Many physical properties of the target system are determined by the ground state, motivating this goal.

The Variational Quantum Eigensolver (VQE) quantum-classical hybrid algorithm⁴ is the most widely studied means of achieving this on NISQ hardware. VQE requires a parametrized ansatz state, $|\psi_{\text{anz}}(\theta)\rangle \in \mathcal{H}$, whose parameters θ are manipulated within a classical optimization scheme that aims to minimize the energy expectation value

$$E(\theta) := \langle H_{\mathcal{T}} \rangle_{\psi_{\text{anz}}(\theta)} \quad (4)$$

evaluated via many prepare-and-measure cycles. The choice of ansatz restricts us to a subspace of quantum states and therefore must be carefully designed to be sufficiently expressible so as to capture the true ground state of the system.

A common form of ansatz state, particularly in relation to the electronic structure problem, is

$$|\psi_{\text{anz}}(\theta)\rangle = e^{iA(\theta)} |\psi_{\text{ref}}\rangle \quad (5)$$

where $|\psi_{\text{ref}}\rangle \in \mathcal{H}$ is some fixed reference state in which the quantum circuit is initialized and $A(\theta) = \sum_{\sigma \in \mathcal{A}} \theta_{\sigma} \sigma$ for parameters $\theta_{\sigma} \in \mathbb{R}$ and Pauli operators $\sigma \in \mathcal{A} \subset \mathcal{P}_N$; the unitary $e^{iA(\theta)}$ effects excitations above the reference state. Such ansätze as unitary coupled cluster (UCC)^{13,14} may be expressed by our choice of A (taking as reference the Hartree–Fock state), in addition to any others based on the theory of excitation operators such as ADAPT-VQE.^{16–19} The quantum advantage in VQE stems from the ability to prepare classically intractable states from our parametrized ansatz circuits.

3. PROJECTIONS ONTO STABILIZER SUBSPACES

Given an operator $\sigma \in \mathcal{P}_N$, the space of quantum states $|\psi\rangle \in \mathcal{H}$ that it stabilizes are those satisfying $\sigma|\psi\rangle = |\psi\rangle$, the +1-eigenspace of σ . Extending this notion to an Abelian subgroup of Pauli operators $\mathcal{Q} \subset \mathcal{P}_N$, there is an induced vector space $\mathcal{V}_{\mathcal{Q}}$ of states stabilized by the elements of \mathcal{Q} .

A particularly useful definition is that of a Hamiltonian symmetry, taken here to mean a set $\mathcal{S} \subset \mathcal{P}_N$ of Pauli operators such that

$$[\sigma, \sigma'] = 0 \quad \forall \sigma \in \mathcal{S}, \quad \sigma' \in \mathcal{T} \quad (6)$$

In other words, a symmetry of $H_{\mathcal{T}}$ is any set of Pauli operators that commute universally among \mathcal{T} , which we may extend to an Abelian group $\bar{\mathcal{S}}$, the closure of \mathcal{S} under operator multiplication, which we shall call a *symmetry group*.

Note the setting in which we present symmetries here is stricter than the conventional definition, which considers any operator S that commutes with the Hamiltonian, i.e., $[S, H_{\mathcal{T}}] = 0$, to be a symmetry. Such an operator need not commute with the individual terms as we require here. For example, in the Fermionic picture, the number operator $\sum_i a_i^{\dagger} a_i$ (where a is the Fermionic annihilation operator and its Hermitian conjugate a^{\dagger} represents the creation operator) commutes with the full second-quantized molecular Hamiltonian, but not with an arbitrary excitation term.

The operators of \mathcal{S} will in general be algebraically dependent, but the theory of stabilizers⁵⁰ ensures the existence of a set of independent generators \mathcal{G} such that $\bar{\mathcal{S}} = \bar{\mathcal{G}}$. Now, recall that the Clifford group consists of unitary operators $U \in \mathcal{B}(\mathcal{H})$ (meaning $UU^{\dagger} = U^{\dagger}U = \mathbf{1}$) with the property $U\sigma U^{\dagger} \in \mathcal{P}_N \quad \forall \sigma \in \mathcal{P}_N$, i.e., U normalizes the Pauli group. We may construct a Clifford operation U mapping each symmetry generator to distinct single-qubit Pauli operators σ_p , where we are free to choose $p \in \{1, 2, 3\}$. More precisely, there exists a subset of qubit positions $\mathcal{I}_{\text{stab}} \subset \mathbb{Z}_N$ satisfying $|\mathcal{I}_{\text{stab}}| = |\mathcal{G}|$ and a bijective map $f: \mathcal{G} \rightarrow \mathcal{I}_{\text{stab}}$ such that

$$UGU^{\dagger} = \sigma_p^{(f(G))} \quad \forall G \in \mathcal{G} \quad (7)$$

This is a powerful concept that provides a mechanism for reducing the number of qubits in the Hamiltonian while preserving its energy spectrum. This is at the core of qubit tapering,^{47,48} in which it is observed that

$$[G, H_{\mathcal{T}}] = 0 \Rightarrow [\sigma_p^{(f(G))}, H'_{\mathcal{T}}] = 0 \quad \forall G \in \mathcal{G} \quad (8)$$

implying the rotated Hamiltonian $H'_{\mathcal{T}} := UH_{\mathcal{T}}U^{\dagger}$ consists solely of identity or Pauli σ_p operators in the qubit positions

indexed by $\mathcal{I}_{\text{stab}}$. Taking expectation values, one may replace the qubits $\mathcal{I}_{\text{stab}}$ by their eigenvalues $\nu_i = \pm 1$; each assignment

$$\nu = (\nu_i)_{i \in \mathcal{I}_{\text{stab}}} \in \{\pm 1\}^{\times \mathcal{I}_{\text{stab}}} \quad (9)$$

defines a symmetry sector and at least one such sector will contain the true solution to the eigenvalue problem. Note the other sectors still have physical significance and may, for example, relate to solutions with different particle numbers or to excited states. Ancillary data files are provided in which we report the symmetry generators and corresponding sector for the Hamiltonians representing the molecular systems listed in Table 1.

Table 1. Systems Investigated to Benchmark the Noncontextual Projection Ansatz^a

molecular systems			number of qubits		
name	charge	mult.	full	taper	CS-VQE ^b
Be	0	1	10	5	3
B	0	2	10	5	3
LiH	0	1	12	8	4
BeH	+1	1	12	8	6
HF	0	1	12	8	4
BeH ₂	0	1	14	9	7
H ₂ O	0	1	14	10	7
F ₂	0	1	20	16	10
HCl	0	1	20	17	4

^aAll in the STO-3G basis. ^bIndicates the fewest number of qubits required to achieve chemical precision.

A quantum state consistent with any such sector must be stabilized by the operators $\nu_i \sigma_p^{(i)}$, and we may define a projection onto the corresponding stabilizer subspace. In general, a projection is defined to be an *idempotent* operator $P \in \mathcal{B}(\mathcal{H})$, i.e., $P^2 = P$; the projection onto the ± 1 -eigenspace of a single-qubit Pauli operator σ_p for $p \in \{1, 2, 3\}$ may be written

$$P_p^{\pm} := \frac{1}{2}(I \pm \sigma_p) \quad (10)$$

States with no component inside the chosen eigenspace are mapped to zero and observe that

$$P_p^{\pm} \sigma_q P_p^{\pm} = \pm \delta_{p,q} P_p^{\pm} \quad (11)$$

for $q \in \{1, 2, 3\}$.

Let $\mathcal{H}_{\text{stab}}$ be the reduced Hilbert space supported by the stabilized qubits $\mathcal{I}_{\text{stab}}$ and \mathcal{H}_{red} its complement such that $\mathcal{H} = \mathcal{H}_{\text{stab}} \otimes \mathcal{H}_{\text{red}}$. Given an assignment of eigenvalues $\nu \in \{\pm 1\}^{\times \mathcal{I}_{\text{stab}}}$, we may project onto the corresponding sector via

$$P_{\nu} := \bigotimes_{i \in \mathcal{I}_{\text{stab}}} P_p^{\nu_i} \quad (12)$$

and subsequently perform a *partial trace* over the stabilized qubits $\mathcal{I}_{\text{stab}}$. This is effected by the unique linear map $\text{Tr}_{\text{stab}}: \mathcal{H} \rightarrow \mathcal{H}_{\text{red}}$ satisfying the property $\text{Tr}_{\text{stab}}(A \otimes B) = \text{Tr}(A)B$ for all $A \in \mathcal{B}(\mathcal{H}_{\text{stab}})$ and $B \in \mathcal{B}(\mathcal{H}_{\text{red}})$.

Finally, we may define the full stabilizer subspace projection map

$$\pi_\nu^U: \mathcal{B}(\mathcal{H}) \rightarrow \mathcal{B}(\mathcal{H}_{\text{red}}); \quad A \rightarrow \text{Tr}_{\text{stab}}(P_\nu U A U^\dagger P_\nu) \quad (13)$$

which, using the linearity of Tr_{stab} , yields a reduced Hamiltonian

$$\begin{aligned} H_{\mathcal{T}}^{\text{red}} &:= \pi_\nu^U(H_{\mathcal{T}}) \\ &= \sum_{\sigma \in \mathcal{T}} h'_\sigma \sigma'_{\text{red}} \end{aligned} \quad (14)$$

where $\sigma' = U\sigma U^\dagger = \bigotimes_{i=0}^{N-1} \sigma_{q_i}$ and we have written $\sigma' = \sigma'_{\text{stab}} \otimes \sigma'_{\text{red}}$. The new coefficients $h'_\sigma := h_\sigma \prod_{q_i \neq 0}^{i \in \mathcal{I}_{\text{stab}}} \nu_i$ differ from h_σ by a sign dependent on the chosen symmetry sector.

In qubit tapering, U is taken as eq 7 with the corresponding basis \mathcal{G} a generating set for a full Hamiltonian symmetry.^{47,48} Assuming identification of the correct sector, the ground state energy of the $(N - |\mathcal{G}|)$ -qubit reduced Hamiltonian $H_{\mathcal{T}}^{\text{red}}$ will coincide with the true value of the full system $H_{\mathcal{T}}$.

This stabilizer projection procedure is straightforward with respect to the Hamiltonian, since the stabilized qubits contain only operators with nonzero image under conjugation with P_ν . However, suppose we were to take another observable $A \in \mathcal{B}(\mathcal{H})$ and wish to determine a reduced form on $\mathcal{B}(\mathcal{H}_{\text{red}})$ that is consistent with the reduced Hamiltonian $H_{\mathcal{T}}^{\text{red}}$. This may be achieved by following precisely the same process that was applied to $H_{\mathcal{T}}$, but the symmetry \mathcal{S} will not in general be a symmetry of A and therefore the “symmetry-breaking” terms (those which anticommute with the generators \mathcal{G}) will vanish under projection onto the stabilizer subspace, as per eq 11. Letting $\mathcal{A} \subset \mathcal{P}_N$ be the set of terms in the Pauli-basis expansion of A , observe that

$$\begin{aligned} A^{\text{red}} &:= \pi_\nu^U(A) \\ &= \sum_{\sigma \in \mathcal{A}} h_\sigma \text{Tr}(P_\nu \sigma'_{\text{stab}} P_\nu) \sigma'_{\text{red}} \\ &= \sum_{\sigma \in \mathcal{A}} h_\sigma \text{Tr} \left(\bigotimes_{\substack{i \in \mathcal{I}_{\text{stab}} \\ q_i \neq 0}} P_p^{\nu_i} \otimes \bigotimes_{\substack{i \in \mathcal{I}_{\text{stab}} \\ q_i \neq 0}} \underbrace{P_p^{\nu_i} \sigma_{q_i} P_p^{\nu_i}}_{=\nu_i \delta_{p,q_i}} \right) \sigma'_{\text{red}} \\ &= \sum_{\sigma \in \mathcal{A}} h_\sigma \sigma'_{\text{red}} \underbrace{\text{Tr}(P_\nu)}_{=1} \prod_{\substack{i \in \mathcal{I}_{\text{stab}} \\ q_i \neq 0}} \nu_i \delta_{p,q_i} \\ &= \sum_{\substack{\sigma \in \mathcal{A} \\ q_i \in \{0,p\} \\ \forall i \in \mathcal{I}_{\text{stab}}}} h'_\sigma \sigma'_{\text{red}}, \end{aligned} \quad (15)$$

recalling that q_i indicates the type of single-qubit Pauli acting on qubit position $i \in \mathbb{Z}_N$ in some tensor product σ , defined in Section 2.

The resulting form is identical to eq 14, except we are explicit that the terms surviving projection are only those whose qubit positions indexed by $\mathcal{I}_{\text{stab}}$ consist exclusively of identity and Pauli σ_p operators; this is trivially true for the Hamiltonian by construction. Most importantly, this extends the stabilizer subspace projection to ansätze defined on the full system for use in variational algorithms. It should be noted that the above operations are classically tractable and can be implemented efficiently in the symplectic representation of Pauli operators.^{51,52}

It would be remiss of us not to draw attention to the likeness of eq 13 to Positive Operator-Valued Measures (POVMs);⁵³ indeed, the projectors (eq 12) define a complete set of Kraus

operators.⁵⁴ The stabilizer subspace projection procedure is reduced to a matter of enforcing a partial measurement over some subsystem of the full problem, for which the relevant outcomes have been determined via an auxiliary method. For example, this could involve identifying a quantum state with a known nonzero overlap with the true ground state; measuring the symmetry generators \mathcal{G} in this state will yield the correct sector.

Hartree–Fock often provides such a state for electronic structure problems, although it is not immune to failure; this is particularly true in the strongly correlated regime. In these cases, we should defer to more effective reference states such as those obtained from Møller–Plesset perturbation theory (MP), coupled-cluster (CC) methods, and so on. One can imagine a hierarchy of increasingly precise ground state approximations, for which we should hope to obtain at some point a nonzero overlap with the true ground state.

4. CS-VQE IN THE STABILIZER FORMALISM

We now describe the *Contextual Subspace VQE* (CS-VQE) method in the stabilizer setting introduced in Section 3. CS-VQE partitions the Hamiltonian (2) into two disjoint components, one that is noncontextual and another that is contextual, which provides quantum corrections to the former via VQE.⁴⁵ Explicitly, this allows us to write

$$H_{\mathcal{T}} = H_{\mathcal{T}_{\text{nc}}} + H_{\mathcal{T}_{\text{c}}} \quad (16)$$

where \mathcal{T}_{nc} is a noncontextual set of Pauli operators and $\mathcal{T}_{\text{c}} := \mathcal{T} \setminus \mathcal{T}_{\text{nc}}$ is what remains, which will in general be contextual.

CS-VQE differs from qubit tapering (described in Section 3) in the following way: the latter exploits existing (i.e., physical) symmetries of the Hamiltonian, whereas in CS-VQE, we impose additional “pseudosymmetries” derived from the noncontextual Hamiltonian. This results in a loss of information, since any terms of \mathcal{T} not commuting with the symmetry generators will vanish under projection.

4.1. The Noncontextual Problem. The notion of contextuality goes back to the Bell–Kochen–Specker theorem.^{55–57} Here we use an explicit condition for the noncontextuality of a set of Pauli operators, developed by Kirby and Love⁵⁸ and independently by Raussendorf et al.⁵⁹ Strictly speaking, this condition tests for strong measurement contextuality. In this setting, a set \mathcal{T}_{nc} is understood to be noncontextual if and only if commutation forms an equivalence relation on $\mathcal{T}_{\text{nc}} \setminus \mathcal{S}$, where we have defined the sub-Hamiltonian symmetry $\mathcal{S} := \{\sigma \in \mathcal{T}_{\text{nc}} \mid [\sigma, \sigma'] = 0 \forall \sigma' \in \mathcal{T}_{\text{nc}}\}$. There is an implied structure

$$\mathcal{T}_{\text{nc}} = \mathcal{S} \cup C_1 \cup \dots \cup C_M \quad (17)$$

where the C_i are equivalence classes with respect to commutation—in other words, elements of the same class commute and across classes they anticommute. Conversely, such a set of Pauli operators is contextual if and only if commutation fails to be transitive on $\mathcal{T}_{\text{nc}} \setminus \mathcal{S}$.

The symmetry \mathcal{S} can be expanded by taking pairwise products within equivalence classes, since $\{C_i, C_j\} = 0$ for $C_i \in C_i, C_j \in C_j$ with $i \neq j$, it is the case that $\{C_i C_i', C_j C_j'\} = 0$ and we may define $\mathcal{S}' = \mathcal{S} \cup \bigcup_{i=1}^M \{C_i C_i' \mid C_i, C_i' \in C_i\}$. As before, in Section 3, \mathcal{S}' induces a symmetry group for which one may define independent generators \mathcal{G} and a Clifford

operation $U_{\mathcal{G}}$ mapping the generators to single-qubit Pauli operators; the expectation value over these qubits will again be determined by an assignment $\nu \in \{\pm 1\}^{|\mathcal{G}|}$ of eigenvalues, analogous to the selection of a symmetry sector in qubit tapering.

From each equivalence class C_i , we select a representative C_i and construct an observable $C(\mathbf{r}) := \sum_{i=1}^M r_i C_i$ where $\mathbf{r} \in \mathbb{R}^M$ and $|r_i| = 1$. Kirby and Love⁶⁰ found that quantum states $|\psi_{(\nu, \mathbf{r})}\rangle \in \mathcal{H}$ stabilized by the operators $\{\nu_{f(G)} G | G \in \mathcal{G}\} \cup \{C(\mathbf{r})\}$ are consistent with a classical objective function $\eta(\nu, \mathbf{r})$ (derived in the [Supporting Information](#)), in the sense that $\eta(\nu, \mathbf{r})$ coincides with the noncontextual energy expectation value $\langle H_{\mathcal{T}_{nc}} \rangle_{\psi_{(\nu, \mathbf{r})}}$ for all parametrizations (ν, \mathbf{r}) . This is a consequence of the joint probability distribution chosen over the phase-space points of their (epistricted) model.^{60,61}

The noncontextual energy spectrum is therefore parametrized by two vectors: the ± 1 eigenvalue assignments ν , determining the contribution of the universally commuting terms, and \mathbf{r} , encapsulating the remaining pairwise anti-commuting classes. In this sense, we may refer to (ν, \mathbf{r}) as a state of the noncontextual Hamiltonian itself, abstracted from quantum states of the corresponding stabilizer subspace. Optimizing over these parameters, we obtain the non-contextual ground state energy

$$\epsilon_0^{\text{nc}} := \min_{\substack{\nu \in \{\pm 1\}^{|\mathcal{G}|} \\ \mathbf{r} \in \mathbb{R}^M, |\mathbf{r}|=1}} \eta(\nu, \mathbf{r}) \quad (18)$$

and call an element (ν, \mathbf{r}) of the preimage $\eta^{-1}(\epsilon_0^{\text{nc}})$ a noncontextual ground state of $H_{\mathcal{T}_{nc}}$. Let us denote by $\Delta_{\text{nc}} := |\epsilon_0^{\text{nc}} - \epsilon_0|$ the absolute error with respect to the true ground state energy.

As a classical estimate to the ground state energy of the full Hamiltonian $H_{\mathcal{T}}$, in [Section 5](#), we found the difference between the noncontextual ground state and Hartree–Fock energy to be negligible for each of the molecules simulated, since the heuristic used to choose $H_{\mathcal{T}_{nc}}$ prioritizes diagonal Hamiltonian terms. In principle, it may be an improvement upon Hartree–Fock as the noncontextual set can also take into account an off-diagonal contribution within the anticommuting classes. This is highly dependent on the chosen form of noncontextual set; a reformulation in terms of graphs, e.g., representing Pauli operators as nodes with (non)adjacency indicating (anti)commutation, will allow one to identify what the equivalent problem(s) are in computer science and therefore draw upon the vast body of existing research and select the best algorithms designed to solve such computational problems of graph theory. It should be noted that the “optimal” noncontextual subset will not necessarily be that which minimizes the noncontextual ground state energy and some consideration of the resulting quantum corrections must inform this choice, which remains an open question.

4.2. Quantum Corrections. Our simulation approach has thus far been strictly classical—now we arrive at the quantum element of CS-VQE. We have derived a classical estimate of the ground state energy from the noncontextual part of the Hamiltonian $H_{\mathcal{T}_{nc}}$; however, the contextual component $H_{\mathcal{T}_c}$ has so far been neglected.

While $C(\mathbf{r})$ is not a stabilizer in the strict sense (it is not an element of the Pauli group), it is unitarily equivalent to one as a linear combination of anticommuting Pauli elements. Similar to the symmetry generators \mathcal{G} , it is possible to define a unitary operation U_C mapping $C(\mathbf{r})$ onto a single-qubit Pauli operator, following the approach of unitary partitioning.^{32–36} However, unlike the $U_{\mathcal{G}}$ rotation, U_C is not Clifford as it collapses M terms onto a single Pauli operator and can therefore introduce additional terms to the Hamiltonian. Kirby et al.⁴⁵ cautioned that, in principle, this increase in Hamiltonian complexity could be exponential in the number of equivalence classes M , namely, a scaling of $O(2^M)$. However, Ralli et al.³⁶ demonstrated that the general scaling for this sequence of rotations (SeqRot) method is $O(x^{M-1})$ where $x \in [1, 2]$; that is, still exponential, yet the necessary conditions to obtain the worst-case $x = 2$ are contrived and have not been observed for any molecular Hamiltonians investigated to date. Regardless, one may circumvent this potentially adverse scaling entirely by implementing the linear combination of unitaries (LCU) approach to unitary partitioning,^{33,35} which is only quadratic in the number of equivalence classes $O(M^2)$.³⁶

Appending $C(\mathbf{r})$ to our set of generators $\tilde{\mathcal{G}} := \mathcal{G} \cup \{C(\mathbf{r})\}$ and defining $U := U_C U_{\mathcal{G}}$, there exists a subset of qubit indices $\mathcal{I}_{\text{stab}}$ satisfying $|\mathcal{I}_{\text{stab}}| = |\tilde{\mathcal{G}}|$ and a bijective map $f: \tilde{\mathcal{G}} \rightarrow \mathcal{I}_{\text{stab}}$ such that $U G U^\dagger = \sigma_p^{f(G)}$ for each $G \in \tilde{\mathcal{G}}$. We reiterate that $p \in \{1, 2, 3\}$ may be chosen at will; the approach taken by Kirby et al.⁴⁵ is to select $p = 3$ to enforce diagonal generators.

Suppose we have a quantum state $|\psi_{(\nu, \mathbf{r})}\rangle$ that is consistent with ϵ_0^{nc} ; since the rotated state $|\psi'_{(\nu, \mathbf{r})}\rangle = U |\psi_{(\nu, \mathbf{r})}\rangle$ must be stabilized by $\sigma_p^{(i)} \forall i \in \mathcal{I}_{\text{stab}}$, the qubit positions $\mathcal{I}_{\text{stab}}$ must be fixed. This implies a decomposition

$$|\psi'_{(\nu, \mathbf{r})}\rangle = |b_{(\nu, \mathbf{r})}\rangle_{\text{stab}} \otimes |\varphi\rangle_{\text{red}} \quad (19)$$

where $|b_{(\nu, \mathbf{r})}\rangle$ represents a single basis state of $\mathcal{H}_{\text{stab}}$ and $|\varphi\rangle \in \mathcal{H}_{\text{red}}$ is independent of the parameters (ν, \mathbf{r}) . The expectation value of the full Hamiltonian may be expressed as

$$\langle H_{\mathcal{T}} \rangle_{\psi_{(\nu, \mathbf{r})}} = \epsilon_0^{\text{nc}} + \langle \pi_{\nu}^U(H_{\mathcal{T}_c}) \rangle_{\varphi} \quad (20)$$

where $\pi_{\nu}^U(H_{\mathcal{T}_c})$ contains only the terms of the contextual Hamiltonian that commute with all the noncontextual generators, just as in eq 15. It was observed by Kirby et al.⁴⁵ that any term which anticommutes with at least one noncontextual generator must have zero expectation value, and our stabilizer subspace projection captures this fact.

Inspecting eq 20, we may optimize freely over quantum states φ , i.e., we are not constrained by the noncontextual ground state within \mathcal{H}_{red} . In fact, we may absorb the noncontextual ground state energy into the reduced contextual Hamiltonian

$$\tilde{H}_{\mathcal{T}_c} := \epsilon_0^{\text{nc}} \cdot \mathbf{1} + \pi_{\nu}^U(H_{\mathcal{T}_c}) \quad (21)$$

defining the *contextual subspace Hamiltonian*; this form is obtained naturally when applying the stabilizer subspace projection to the full Hamiltonian, which automatically includes the noncontextual energy by fixing the corresponding eigenvalue assignments.

Now, we may perform unconstrained VQE to obtain a quantum-corrected estimate

$$\varepsilon_0^c := \min_{|\phi\rangle \in \mathcal{H}_{\text{red}}} \langle \tilde{H}_{\mathcal{T}_c} \rangle_{\phi} \quad (22)$$

of the true ground state energy with absolute error $\Delta_c := |\varepsilon_0^c - \varepsilon_0| \leq \Delta_{\text{nc}}$. We have equality when the stabilizers span every qubit position, which is the case when $|\tilde{\mathcal{G}}| = N$ since the generators must be algebraically independent: this means the initial quantum correction is trivial as the noncontextual part determines the entire system.

For instances of the electronic structure problem, there is no guarantee that ε_0^c will achieve chemical precision ($\Delta_c < 1.6$ mHa ≈ 4 kJ/mol) and, indeed, it might not improve upon the noncontextual estimate (although it will never be worse, due to the variational principle applying in this case). However, one can easily define a subset of \mathcal{T}_{nc} that is again noncontextual; this is achieved by discarding one of the noncontextual generators $G \in \tilde{\mathcal{G}}$, along with the operators that it generates. We now append the discarded operators to the contextual Hamiltonian, relaxing the stabilizer constraint on the qubit position $f(G)$ and permitting a search over its Hilbert space. This process may be iterated until the noncontextual set is exhausted and we recover full VQE. This means that, unless the ground state energy of $H_{\mathcal{T}_{\text{nc}}}$ and H coincides, CS-VQE will improve upon the noncontextual energy using less quantum resources than full VQE; this is more rigorously defined in the next section.

In summary, what we have described here is a technique of scaling the relative sizes of the noncontextual (read classical) and contextual (read quantum) simulations in a reciprocal manner. We can therefore trade-off quantum and classical workloads in CS-VQE.

4.3. Expanding the Contextual Subspace. Now we describe the process of growing the contextual subspace more rigorously. We select a subset of noncontextual generators $\mathcal{F} \subset \tilde{\mathcal{G}}$ whose stabilizer constraints we mean to enforce and construct a new noncontextual set $\mathcal{T}'_{\text{nc}} := \mathcal{T}_{\text{nc}} \cap \bar{\mathcal{F}}$; the contextual set is expanded accordingly by appending the terms not generated by \mathcal{F} , i.e., $\mathcal{T}'_c := \mathcal{T}_c \cup (\mathcal{T}_{\text{nc}} \setminus \bar{\mathcal{F}})$. As before, there exists a unitary operation $U_{\mathcal{F}}$, a subset of qubit indices $\mathcal{I}_{\text{fix}} \subset \mathcal{I}_{\text{stab}}$, and a bijective map $f: \mathcal{F} \rightarrow \mathcal{I}_{\text{fix}}$ satisfying $U_{\mathcal{F}} G U_{\mathcal{F}}^\dagger = \sigma_p^{f(G)} \forall G \in \mathcal{F}$ (the rotation $U_{\mathcal{F}}$ may or may not be Clifford depending on whether $C(\mathbf{r})$ is among the stabilizers we wish to fix).

Denote by $\varepsilon_0^{\text{nc}}(\mathcal{F})$ the ground state energy of the new noncontextual Hamiltonian \mathcal{T}'_{nc} with absolute error $\Delta_{\text{nc}}(\mathcal{F}) \geq \Delta_{\text{nc}}$. While this is weaker as an estimate of the true ground state energy of the full system, at the very least we are guaranteed to recover the initial noncontextual ground state energy from performing a simulation of the expanded contextual subspace,⁴⁵ which we describe below.

The stabilizer constraints of \mathcal{F} are enforced over the Hilbert space $\mathcal{H}_{\text{fix}} = (\mathbb{C}^2)^{\otimes \mathcal{I}_{\text{fix}}}$ of qubits indexed by \mathcal{I}_{fix} , whereas we may perform a VQE simulation over $\mathcal{H}_{\text{sim}} = (\mathbb{C}^2)^{\otimes \mathcal{I}_{\text{sim}}}$, the Hilbert space of the remaining $N - |\mathcal{F}|$ qubits indexed by $\mathcal{I}_{\text{sim}} = \mathbb{Z}_N \setminus \mathcal{I}_{\text{fix}}$. Invoking the stabilizer subspace projection map $\pi_{\nu'}^{U_{\mathcal{F}}}$ with the eigenvalue assignments $\nu' = (\nu_i)_{i \in \mathcal{I}_{\text{fix}}}$ yields an expanded contextual subspace Hamiltonian

$$\tilde{H}_{\mathcal{T}'_c} := \varepsilon_0^{\text{nc}}(\mathcal{F}) \cdot \mathbf{1} + \pi_{\nu'}^{U_{\mathcal{F}}}(H_{\mathcal{T}'_c}) \quad (23)$$

Performing an $|\mathcal{I}_{\text{sim}}|$ -qubit VQE simulation over the contextual subspace, we obtain a new quantum-corrected estimate

$$\varepsilon_0^c(\mathcal{F}) := \min_{|\phi\rangle \in \mathcal{H}_{\text{sim}}} \langle \tilde{H}_{\mathcal{T}'_c} \rangle_{\phi} \quad (24)$$

with an error satisfying $\Delta_c(\mathcal{F}) \leq \Delta_c$. Recall that $\Delta_c = \Delta_c(\tilde{\mathcal{G}})$ corresponds with the contextual error when we enforce the full set of noncontextual stabilizers.

Observe that, when $|\mathcal{I}_{\text{sim}}| = N$, we are simply performing full VQE over the entire system; this occurs when we do not enforce the stabilizer constraint for any of the noncontextual generators, i.e., $\mathcal{F} = \emptyset$. Therefore, it must be the case that

$$\Delta_c(\mathcal{F}) \rightarrow 0 \text{ as } |\mathcal{F}| \rightarrow 0 \quad (25)$$

Furthermore, given a nested sequence of generator subsets $(\mathcal{F}_i)_i$ with $\mathcal{F}_{i+1} \subset \mathcal{F}_i$, then $\Delta_c(\mathcal{F}_{i+1}) \leq \Delta_c(\mathcal{F}_i)$ and the convergence is monotonic. In this way, CS-VQE describes an interpolation between a purely classical estimate of the ground state energy and a full VQE simulation of the Hamiltonian. In the context of electronic structure calculations, this often permits one to achieve chemical precision at a saving of qubit resources, as indicated by Kirby et al.⁴⁵ for a suite of tapered test molecules of up to 18 qubits. We note in eq 25 that the quality of the chosen ansatz and optimization procedure will limit the actual error one may achieve in practice. This statement instead indicates that, for an appropriate level of contextual subspace approximation, it is possible to construct a reduced Hamiltonian whose exact ground state lies within some error threshold of the true value.

Suppose we wish to find the optimal contextual subspace Hamiltonian of size $N' < N$. The problem reduces to minimizing the error $\Delta_c(\mathcal{F})$ over the $\binom{|\tilde{\mathcal{G}}|}{N - N'}$ generator subsets $\mathcal{F} \subset \tilde{\mathcal{G}}$ satisfying $|\mathcal{F}| = N - N'$. CS-VQE is highly sensitive to this choice and remains a vital open question for the continued success of the technique. For chemistry applications, we grow the contextual subspace until the CS-VQE error attains chemical precision, which means finding the minimal \mathcal{F} such that $\Delta_c(\mathcal{F}) < 1.6$ mHa. In general, we will not have access to a target energy and so will not necessarily know when the desired precision is achieved; instead, we might choose the largest contextual subspace accommodated by the available quantum resource or iterate until the VQE convergence is within some fixed bound.

Greedy selecting combinations of $d \leq N$ generators that yield the greatest reduction in error is an effective stabilizer relaxation ordering heuristic, where iterating $k < N/d$ involves a search of depth d over $N - dk$ elements, thus necessitating $\sum_{k=0}^{\lfloor N/d \rfloor} \binom{N - dk}{d} = O(N^{d+1})$ CS-VQE simulations. Taking $d = 2$ produces a good balance between efficiency and efficacy,⁴⁵ but there is room for more targeted approaches that exploit some structure of the underlying problem. For example, in quantum chemistry problems, it could be that one should relax the stabilizers that have nontrivial action near the Fermi level, between the highest occupied molecular orbital (HOMO) and lowest unoccupied molecular orbital (LUMO). Excitations clustered around this gap are more likely to appear in the true ground state and should therefore not be assigned definite values under the noncontextual projection. This idea comes from the theory of pseudopotential approximations,⁶² in which it is observed that chemically relevant electrons are

predominantly those of the valence space, whereas the core may be “frozen”, thus reducing the electronic complexity.

Alternatively, one might define a Hamiltonian term-importance metric that considers coefficient magnitudes⁶³ or second-order response with respect to a perturbation of the Hartree–Fock state.⁶⁴ In relation to this, it is also not clear which features of a molecular system mean that it might be more or less amenable to CS-VQE; additional insight here would allow one to predict how many qubits will be required to simulate a given problem to chemical precision.

It is not fully understood how CS-VQE relates to active space techniques more generally, but this would be an interesting pursuit for future work. For example, the down-folding technique of subsystem embedding subalgebra coupled cluster (SES-CC)⁶⁵ presents a compelling approach that iteratively decouples excitations $\sigma = \sigma_{\text{int}} + \sigma_{\text{ext}}$ into an “internal” part that belongs to a chosen excitation subalgebra and its “external” complement that may additionally be combined with the double unitary coupled cluster (DUCC) ansatz.⁶⁶ This yields an effective Hamiltonian $H_{\text{ext}}^{\text{eff(DUCC)}} = (P + Q_{\text{int}}) e^{-\sigma_{\text{ext}} H} e^{\sigma_{\text{ext}} (P + Q_{\text{int}})}$ where P projects onto the reference state and Q_{int} onto the subspace of excitations generated by σ_{int} . This has a similar form to our stabilizer subspace projection (eq 13); indeed, it might be possible to reproduce SES-CC under a qubit mapping within the contextual subspace framework by identifying an appropriate noncontextual sub-Hamiltonian and stabilizer subspace.

A benchmark of this and other dimensionality reduction methodologies such as projection-based embedding (PBE)⁶⁷ would be valuable. Furthermore, CS-VQE can be layered on top of these techniques to yield hybrid methods that might outperform any of them on their own; this is a consideration that we plan to take forward into further work, with the goal of deployment on larger molecular systems and basis sets.

4.4. The Noncontextual Projection Ansatz. CS-VQE has thus far not been applied to systems exceeding 18 qubits, and the resulting reduced Hamiltonians (eq b23) have been solved by direct diagonalization;⁴⁵ clearly, this will not scale to larger systems, with the required classical memory increasing exponentially. Instead, they must be simulated by performing VQE routines, but defining an ansatz for the contextual subspace provided an obstacle to achieving this in practice.

However, having now placed the problem within the stabilizer formalism described in Section 3, we have already introduced (in Sections 4.1–4.3) the tools necessary to restrict an ansatz of the form in eq 5—defined over the full system—to the contextual subspace (eq 23). The approach adopted here is equivalent to that which we defined for qubit tapering in eq 15. To restrict a parametrized ansatz operator

$$A(\theta) = \sum_{\sigma \in \mathcal{A}} \theta_{\sigma} \sigma \rightarrow \tilde{A}(\theta) \in \mathcal{B}(\mathcal{H}_{\text{sim}}) \quad (26)$$

in line with the stabilizer constraints $\mathcal{F} \subset \tilde{\mathcal{G}}$, we may simply call upon the stabilizer subspace projection map $\pi_{\nu'}^{U_{\mathcal{F}}}$ once more, which yields a restricted ansatz state

$$|\tilde{\psi}_{\text{anz}}(\theta)\rangle := e^{i\tilde{A}(\theta)} |\tilde{\psi}_{\text{ref}}\rangle \in \mathcal{H}_{\text{sim}} \quad (27)$$

where

$$\tilde{A}(\theta) := \pi_{\nu'}^{U_{\mathcal{F}}}(A(\theta)) \quad (28)$$

Any rotated ansatz term $U_{\mathcal{F}} \sigma U_{\mathcal{F}}^{\dagger}$ that is not identity or a Pauli σ_p on some subset of the qubit positions indexed by \mathcal{I}_{fix} will vanish.

The restricted reference state $|\tilde{\psi}_{\text{ref}}\rangle$ is obtained from an effective partial projective measurement of $U_{\mathcal{F}} |\psi_{\text{ref}}\rangle$ (see the discussion on POVMs in Section 3) with outcomes defined by ν' , which yields a product state

$$\frac{P_{\nu'} U_{\mathcal{F}} |\psi_{\text{ref}}\rangle}{\langle \psi_{\text{ref}} | U_{\mathcal{F}}^{\dagger} P_{\nu'} U_{\mathcal{F}} | \psi_{\text{ref}} \rangle} = |b_{(\nu, \mathbf{r})}\rangle_{\text{fix}} \otimes |\tilde{\psi}_{\text{ref}}\rangle_{\text{sim}} \quad (29)$$

where we have explicitly demarcated the separability across \mathcal{H}_{fix} and \mathcal{H}_{sim} . The postmeasurement state $|b_{(\nu, \mathbf{r})}\rangle \in \mathcal{H}_{\text{fix}}$ on the noncontextual subspace represents a single basis vector and can therefore be disregarded, leaving just the state of the contextual subspace; this we take as reference for our restricted ansatz. If the unitary partitioning rotations are *not* to be applied, then the $U_{\mathcal{F}}$ rotation is trivial over \mathcal{H}_{sim} and we incur no expense in coherent resource. However, if one does enforce the operator $C(\mathbf{r})$ over the contextual subspace, there might be some nontrivial component of the rotation that must be applied in-circuit to ensure that the ansatz lies within the correct subspace; referring to Section 4.2, for the SeqRot approach this will consist of at most $O(N^2)$ CNOT operations in-circuit, whereas LCU is probabilistic due to the nature of block-encoding.³⁵ Given a hardware-efficient ansatz, one may neglect this since the optimizer should compensate the parameters accordingly.

We may now define the contextual subspace energy expectation function

$$\tilde{E}(\theta) := \langle \tilde{H}_{\mathcal{F}_c} \rangle_{\tilde{\psi}_{\text{anz}}(\theta)} \quad (30)$$

with $\tilde{H}_{\mathcal{F}_c}$ as in eq 23, at which point we have reduced the problem to standard VQE, performed over a subspace of the full problem.

5. SIMULATION RESULTS

The molecular systems that were simulated to benchmark the noncontextual projection ansatz for CS-VQE are given in Table 1. The molecule geometries were obtained from the Computational Chemistry Comparison and Benchmark Database (CCCBDB)⁶⁸ and their Hamiltonians were constructed using IBM's Qiskit Nature⁶⁹ with PySCF as the underlying quantum chemistry package.⁷⁰

Before we evaluate the efficacy of our noncontextual projection ansatz, there are a few features of eq 27 that should be highlighted. First of all, from the discussion following eq 29, we potentially apply some component of the operation $U_{\mathcal{F}}$ in-circuit, introducing further gates that will contribute additional noise. However, when the reference state is taken to be that of Hartree–Fock, we observed $U_{\mathcal{F}} \psi_{\text{ref}}$ to coincide with the noncontextual ground state. This is an artifact of the noncontextual set construction heuristic prioritizing diagonal entries, used within both this work and that of Kirby et al.⁴⁵ This need not always be the case, but for the molecular systems investigated, this allows us to avoid performing $U_{\mathcal{F}}$ in-circuit and instead take the noncontextual ground state as our reference. Since we choose to rotate the noncontextual symmetry generators onto Pauli σ_3 operators here, this may be prepared by applying a Pauli σ_1 in each of the qubit

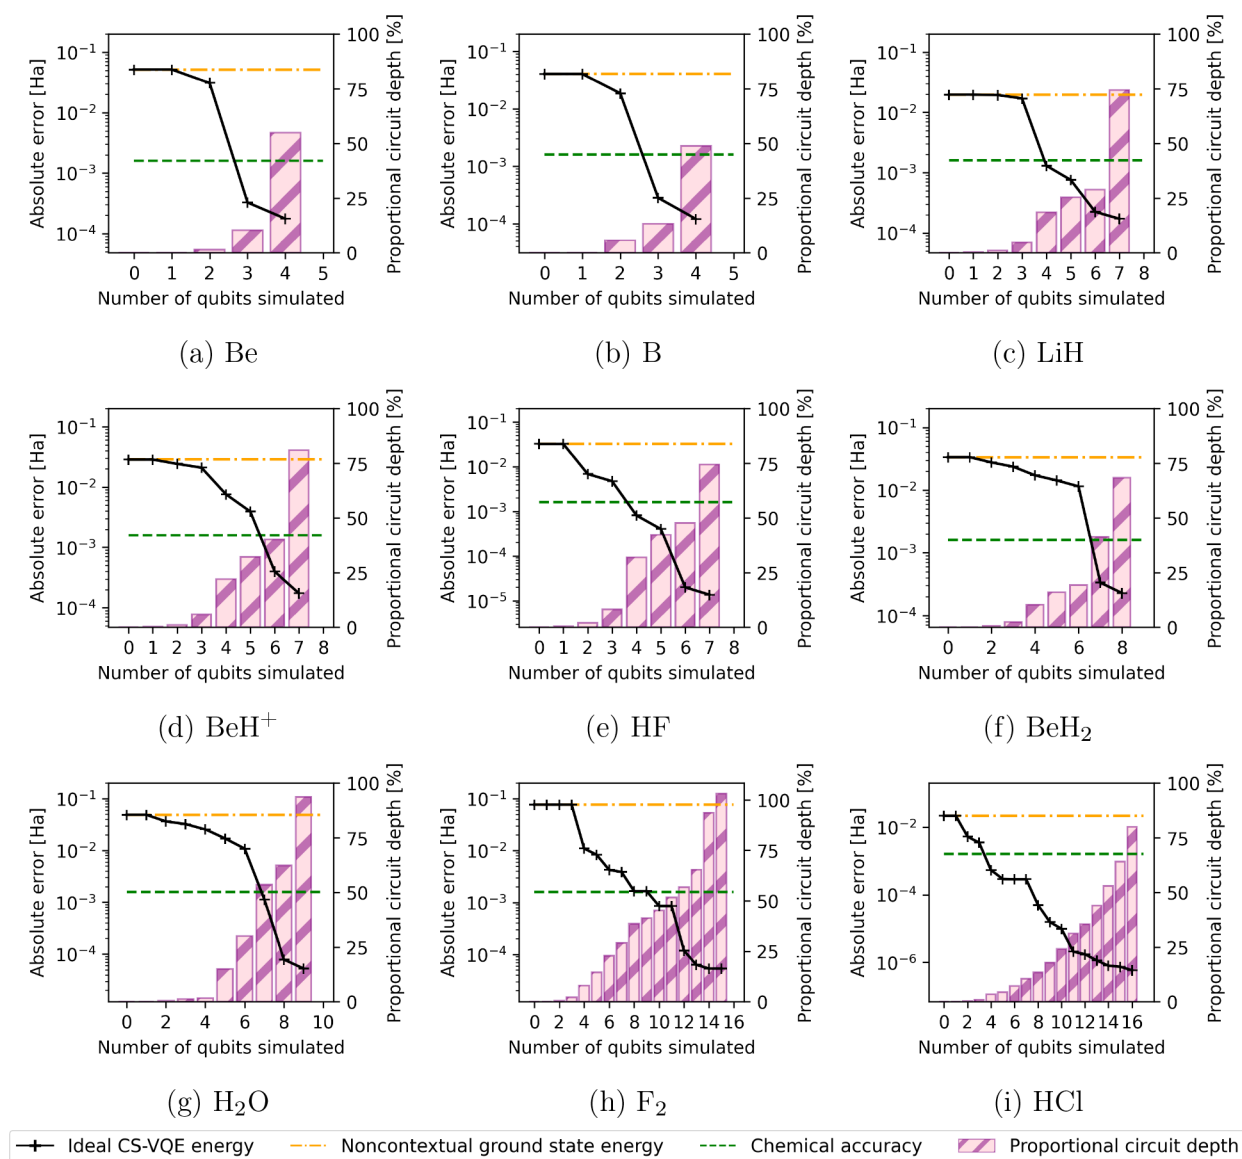


Figure 1. Ideal CS-VQE errors (left-hand axis) and corresponding noncontextual projection ansatz circuit depths as a proportion of the full UCCSD operator from which it is derived (right-hand axis) against the number of qubits simulated.

positions $i \in \mathcal{I}_{\text{sim}}$ such that $\nu_i = -1$ so that the corresponding reference state is stabilized by the relevant operators $\nu_i \sigma_3^{(i)}$. This is visible in Figure 2, in which the VQE routine is initiated with the optimization parameters zeroed, i.e., $\theta = \mathbf{0}$, and since $e^{i\hat{A}(0)} = \mathbf{1}$, optimization begins at the noncontextual ground state energy.

Second, application of the unitary partitioning rotations U_C to the ansatz operator $A(\theta)$ may introduce additional terms by a worst-case scaling factor of $O(2^{M-1})$ where M is the number of equivalence classes in eq 17, although the true scaling is unlikely to be this severe as discussed in Section 4.2. We obtained $M = 2$ for all of the molecules tested, in which case SeqRot is identical to the asymptotically favorable LCU method. In fact, for small $M \ll N$, SeqRot may generate fewer terms than LCU (Ralli et al. presented a toy problem with $M = 3$ in which this was the case³⁶) and therefore our choice of SeqRot here is valid given that the noncontextual set \mathcal{T}_{nc} construction heuristic prioritizes the universally commuting terms \mathcal{S} in eq 17. Different heuristics may lead to larger values

for M , in which case, we recommend an adoption of LCU for implementations of CS-VQE.

Despite this, upon the subsequent projection of $A(\theta)$, it is possible that a significant number of terms will vanish. This is highly dependent on the quality of the initial ansatz and how heavily it is supported on the stabilized qubit positions \mathcal{I}_{fix} . Figure 1 presents circuit depths of the noncontextual projection ansatz as a proportion of the base ansatz from which it is derived, in this case the unitary coupled-cluster singles and doubles (UCCSD) operator. A net reduction in circuit depth is observed, which is quite dramatic up to the point of reaching chemical precision in the CS-VQE routine; in Table 2, we give the specific number of ansatz terms before and after application of the noncontextual projection to UCCSD and UCCSDT for the fewest number of qubits permitting chemical precision.

In order to identify a compact ansatz that closely captures the underlying chemistry with minimal redundancy, we employ the ADAPT-VQE methodology.^{16–19} The algorithm centers around an operator pool from which terms are selected in line

Table 2. Number of Pauli Terms $|\mathcal{A}|$ for a Selection of (Tapered) Ansatz^a

molecule	$ \mathcal{I}_{\text{sim}} $	number of terms in ansatz operator		
		UCCSDT (full/proj) ^b	UCCSD (full/proj) ^b	ADAPT-VQE ^c
Be	3	(48/6)	(48/6)	5
B	3	(48/12)	(32/4)	3
LiH	4	(704/53)	(192/53)	5
BeH ⁺	6	(646/191)	(166/79)	11
HF	4	(92/57)	(92/57)	4
BeH ₂	7	(1312/352)	(224/96)	10
H ₂ O	7	(1892/942)	(324/238)	21
F ₂	10	(176/114)	(176/114)	12
HCl	4	(348/40)	(348/40)	4

^aEach simulation is plotted in Figure 2. ^bThe number of qubits in the contextual subspace over which the ansatz is projected; each tuple (full/proj) gives the number of terms pre- and postprojection. ^cThe number of ADAPT-VQE cycles required to achieve chemical precision, with the operator pool consisting of the projected UCCSD terms.

with a gradient-based argument and appended to a dynamically expanding ansatz whose parameters are optimized at each cycle via VQE. The particular approach we implement here is that of qubit-ADAPT-VQE,¹⁷ following on from iterative qubit coupled cluster,⁷¹ which searches at the level of Jordan–Wigner encoded Pauli operators; the seminal ADAPT-VQE paper¹⁶ instead defines its operator pool over Fermionic excitations.

The Jordan–Wigner transformation⁷² maps a single Fermionic annihilation operator onto two Pauli operators

$$a_i \mapsto \frac{1}{2}(\sigma_1^{(i)} + i\sigma_2^{(i)}) \otimes \bigotimes_{j<i} \sigma_3^{(j)}, \quad (31)$$

with the creation operator given by its Hermitian conjugate a_i^\dagger . Therefore, an excitation on $s \in \mathbb{N}$ spin orbitals of the form

$$\mathbf{a} = a_{i_1}^\dagger \dots a_{i_s}^\dagger a_{j_1} \dots a_{j_s} \quad (32)$$

is represented by 2^{2s} Pauli operators under this encoding. In the unitary coupled cluster theory, we are interested rather in the operator $\mathbf{a} - \mathbf{a}^\dagger$ to ensure unitarity upon exponentiation; this may be expressed by 2^{2s-1} Pauli terms.

As such, after a mapping onto qubits via the Jordan–Wigner transformation, single, double, and triple excitations account for 2, 8, and 32 Pauli operator terms, respectively; while these are required to enforce various electronic symmetries in the ansatz state, not all are necessary to reach chemical precision. This idea lies behind qubit-ADAPT-VQE, which will select only the necessary Pauli terms and therefore yields considerably reduced circuit depths.¹⁷

To leverage ADAPT-VQE in the context of CS-VQE, we define an operator pool $\mathcal{O} \subset \mathcal{P}_N$ and apply to it the stabilizer subspace projection (eq 13) to define a reduced pool $\pi_{\mathcal{V}}^{U_{\mathcal{F}}}(O)$ for the corresponding contextual subspace. Projecting the full pool in this way will ensure that any symmetries S present will be preserved, since $[U_{\mathcal{F}}HU_{\mathcal{F}}^\dagger, U_{\mathcal{F}}SU_{\mathcal{F}}^\dagger] = U_{\mathcal{F}}[H, S]U_{\mathcal{F}}^\dagger = 0$, allowing us to incorporate some chemical intuition into the contextual subspace despite an abstraction from the original problem; one could define a reduced pool directly, but care should be taken to avoid the inclusion of symmetry-breaking terms that may needlessly increase the complexity of the

ADAPT-VQE procedure. The algorithm is then executed as normal, only terminating once the ADAPT-VQE energy is chemically precise with respect to the FCI energy; for scalability, one should terminate computation when the largest gradient in magnitude falls below some predefined threshold, since the true ground state energy will not in general be known. In the Supporting Information, we provide a detailed description of the specific ADAPT-VQE implementation used within this work.

For the following, we take our pool \mathcal{O} to be the terms of the UCCSD operator for each of the molecules in Table 1 before tapering and projecting into the relevant contextual subspace. In Figure 2, we present the ADAPT-VQE convergence data with expectation values obtained via exact wave function (statevector) calculations (i.e., no statistical/hardware noise); chemical precision is achieved in each instance. We used the adaptive moment estimation (Adam)⁷³ classical optimizer and computed parameter gradients as per the parameter shift rule.⁷⁴ Adam has been adopted for previous research in VQE for its resilience to noise, although it exhibits relatively slow convergence compared with other optimizers^{75,76} such as Broyden–Fletcher–Goldfarb–Shanno (BFGS)⁷⁷ and quantum natural gradient (NatGrad);⁷⁸ the latter might be preferable for future work.

The number of ADAPT-VQE cycles (and therefore the number of terms in the resulting ansatz operator) are presented in Table 2, alongside the size of the projected UCCSD operator pool used; one observes a significant reduction in the number of terms. The optimized ADAPT-VQE ansatz operators are reported in ancillary data files, along with a description of the smallest CS-VQE problem permitting chemical precision. This includes the optimal noncontextual generator subset \mathcal{F} , the resulting noncontextual projection ansatz (eq 27), the restricted reference state $|\tilde{\psi}_{\text{ref}}\rangle$ (eq 29), the target error $\Delta_c(\mathcal{F})$ (eq 25), and that which was actually achieved in our VQE simulations (Figure 2). We also include the corresponding contextual subspace Hamiltonians for reproducibility.

Extracting the optimal parameter configuration θ_{min} from the wave function simulations in Figure 2, we subsequently assess the effect of sampling noise on the simulation error with our ansatz circuit preparing the optimal quantum state $|\tilde{\psi}_{\text{anz}}(\theta_{\text{min}})\rangle$. Note that, for each of the molecular systems in Table 1, θ_{min} is given explicitly in the ancillary data files.

To achieve an absolute error of $\Delta > 0$, one should expect to perform $\mathcal{O}\left(\frac{1}{\Delta^2}\right)$ shots (for each term of the Hamiltonian).⁴ Conversely, suppose we are allocated a quantity $S \in \mathbb{N}$ of shots; the obtained error should be of the order $\mathcal{O}\left(\frac{1}{\sqrt{S}}\right)$. In order to increase estimate accuracy, we collected the Pauli terms into qubit-wise commuting (QWC) groups²⁵ using the graph-coloring functionality of NetworkX;⁷⁹ such groups may be measured simultaneously.

In Figure 3, the number of shots $S = 2^n$ for $n = 0, \dots, 20$ carried out per QWC group is varied, and we observe the root mean-square error (RMSE) over 20 realizations of the ground state energy estimate, plotted on a log–log scale. For clarity, note the **only** source of noise here is that which arises from statistical variation of the quantum circuit sampling; we have **not** introduced hardware noise in the form of imperfect quantum gates or decoherence.

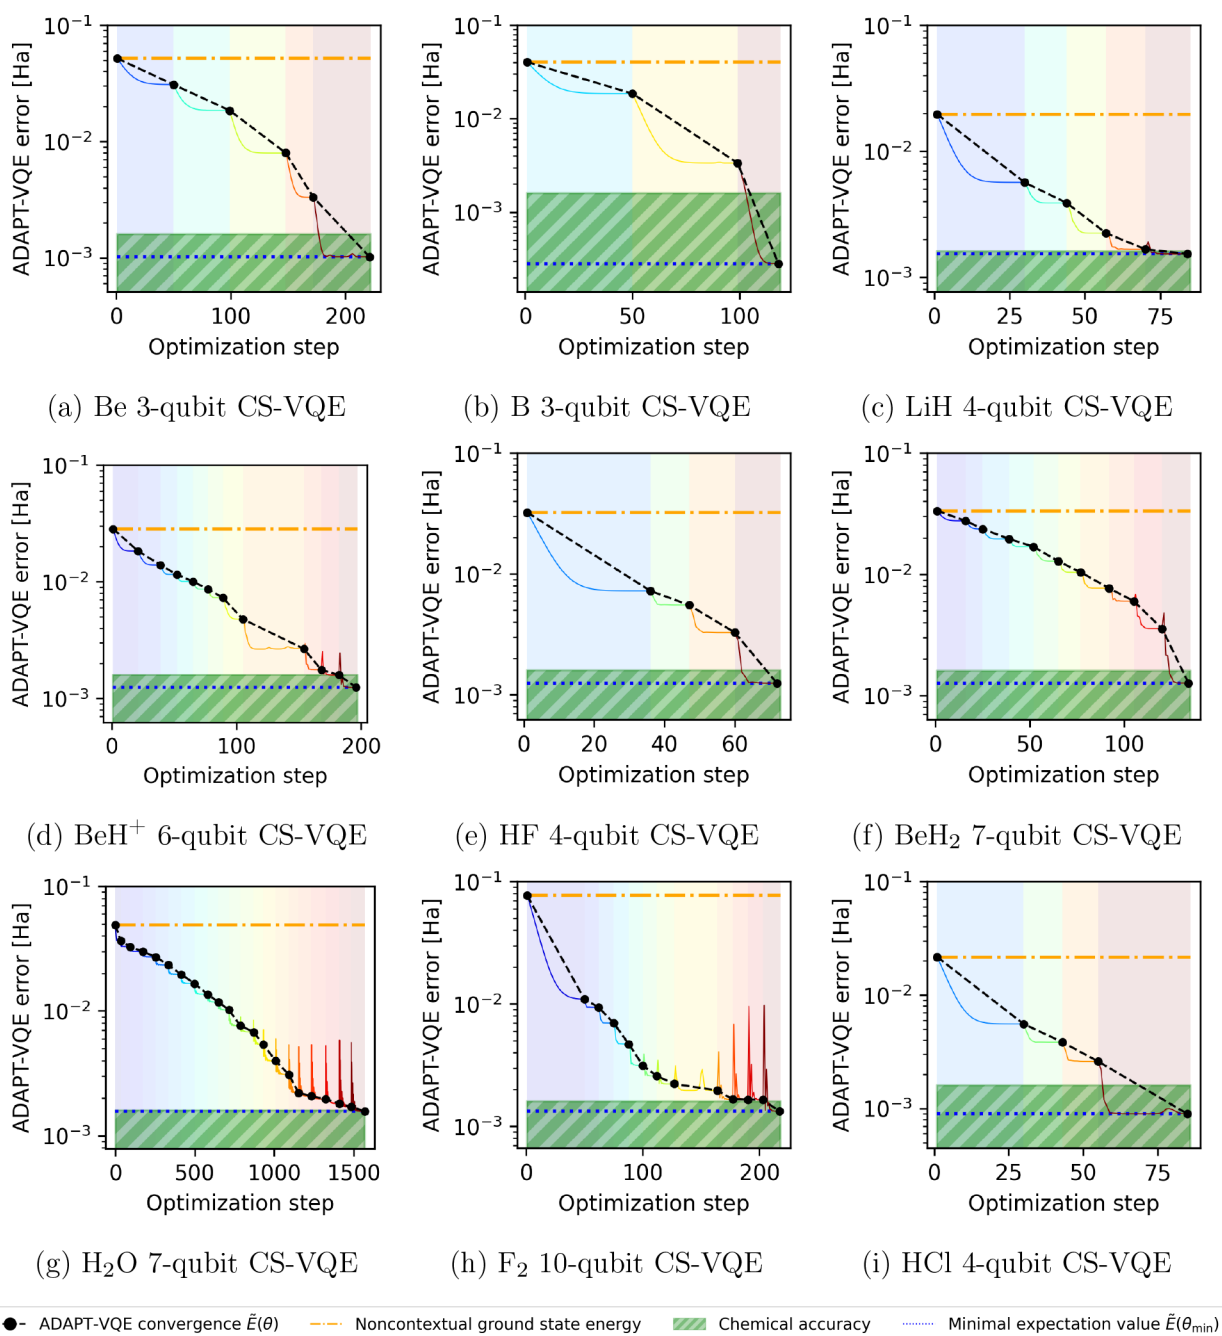


Figure 2. Validation of the noncontextual projection approach to ansatz construction for CS-VQE (eq 27), used here in conjunction with ADAPT-VQE.^{16–19} We plot (on a \log_{10} scale) the absolute error of wave function simulations conducted for the suite of trial molecules outlined in Table 1, each shown to achieve chemical precision; the horizontal axis indicates the algorithm step counter with each shaded region a separate ADAPT-VQE cycle. Adaptive moment estimation (Adam)⁷³ is the classical optimizer taken in the VQE routine performed over the contextual subspace for each ADAPT-VQE cycle, and the settings used are as follows: tolerance = 10^{-4} , learning rate = 10^{-2} , $\beta_1 = 0.4$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$. The parameter gradients $\partial \tilde{E}(\theta)/\partial \theta_i$, required for both operator pool term selection and VQE, were computed using the parameter shift rule.⁷⁴

Two error regimes are observed, one of which is quite trivial: at high shot counts, we see a plateau resulting from the optimal error $|\tilde{E}(\theta_{\min}) - \epsilon_0|$ being recovered. To assess the convergence properties outside of this limiting region, we plot a line of best fit $m \cdot \log_{10}(S) + c$ among the data not exhibiting such behavior; since the data is represented on a log–log scale, this corresponds with a decay in error of $O(S^m)$. In each plot of Figure 3, we obtain $m \approx -0.5$, meaning the RMSE follows the predicted decay of $O\left(\frac{1}{\sqrt{S}}\right)$.

In every simulation bar F_2 , chemical precision was achieved within $S = 2^{20} \approx 10^6$ shots per QWC group. However, our shot budget could be reduced by implementing more advanced allocation strategies, for example, according to the magnitude of Hamiltonian term coefficients⁸⁰ or a classical shadow tomography approach.^{30,31}

6. CONCLUSIONS

We have placed CS-VQE on the theoretical footing of stabilizer subspace projections, which allows one to compare

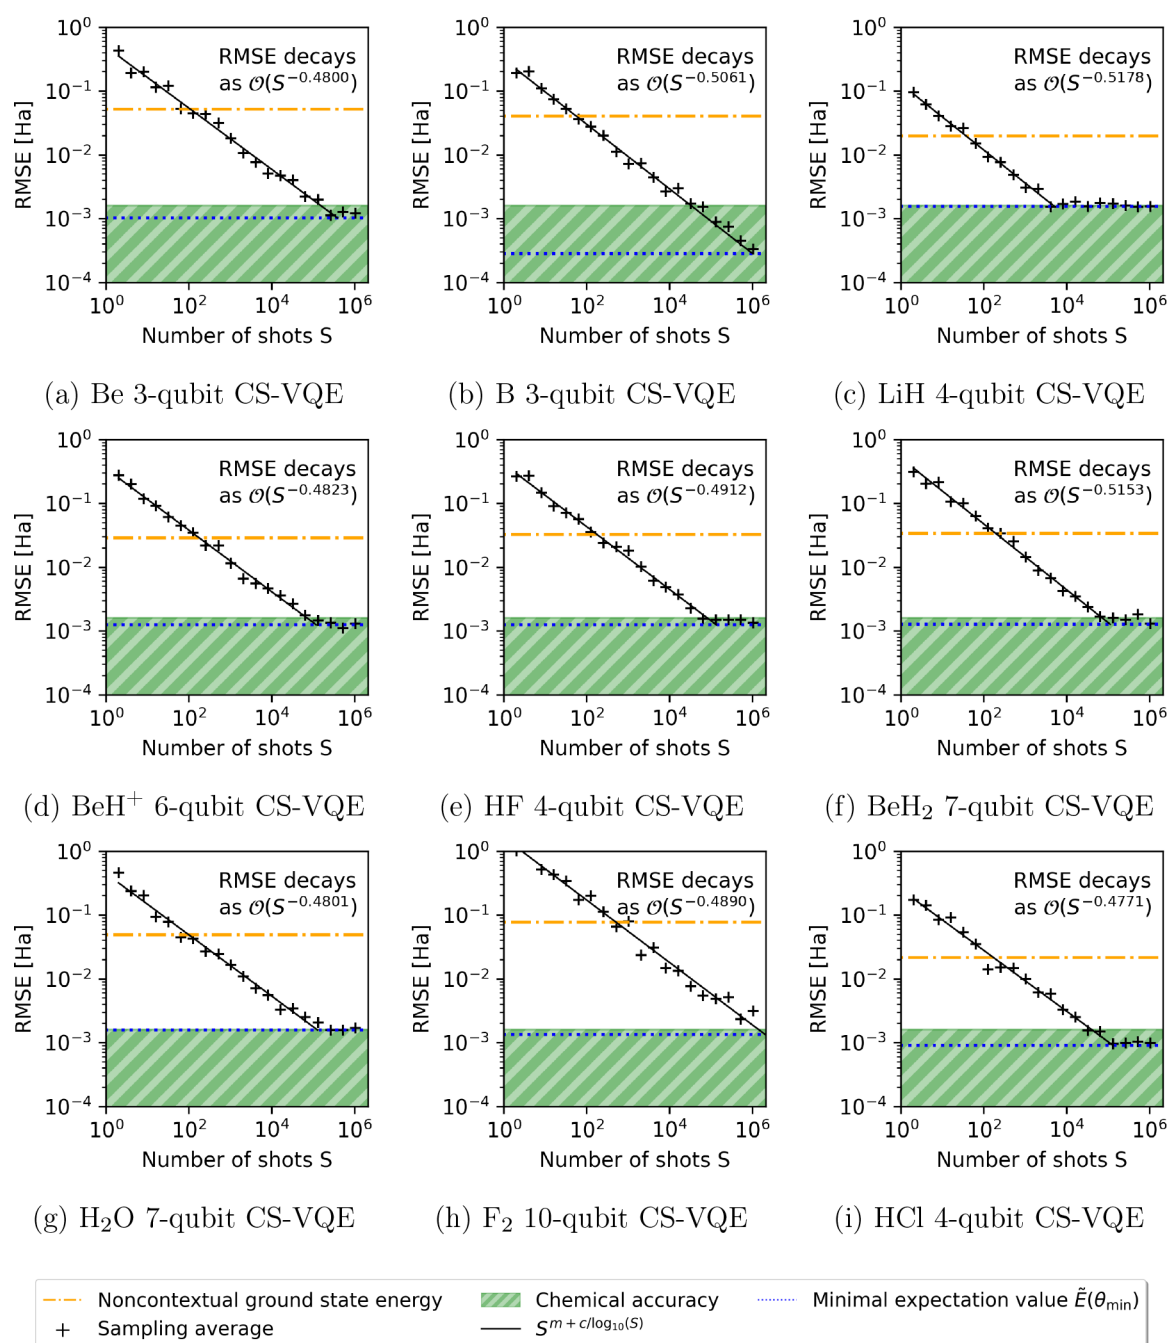


Figure 3. Each of the plots in panels a–i correspond with Figure 2a–i and illustrate the statistical effect of sampling noise at the optimal parametrization θ_{\min} determined from the ADAPT-VQE statevector simulations in Figure 2. We plot the root-mean-square error (RMSE) for 20 “realizations” of the ground state energy estimate with $S \leq 2^{20}$ shots executed via IBM’s QASM simulator; determining the line of best fit $m \cdot \log_{10}(S) + c$ with respect to the log–log data indicates a decay in error of $\mathcal{O}(S^m)$.

it against other qubit reduction techniques such as qubit tapering.^{47,48} Tapering defines a projection dependent on a symmetry of the full Hamiltonian and preserves the ground state energy exactly, whereas CS-VQE is approximate and projects onto a contextual subspace consistent with the symmetry of the noncontextual sub-Hamiltonian, augmented by an anticommuting contribution. In combination, the two techniques can effect a significant reduction in quantum resource requirements, as illustrated by Kirby et al.⁴⁵ and in Figure 1.

Previously, the only obstacle to building a CS-VQE framework that would be faithful to deployment on quantum

devices was that of the ansatz, which has been addressed within this work. Furthermore, we demonstrated how CS-VQE may be combined with the ADAPT-VQE^{16–19} ansatz construction framework by applying our noncontextual projection to the operator pool; validation was presented in Figure 2 in which we achieved chemical precision for the suite of small molecules outlined in Table 1. This combination provides considerable flexibility in both qubit count and circuit depth, allowing one to identify a reduced problem that may be simulated on the available quantum resource.

A number of research questions concerning the scalability of CS-VQE remain; we recapitulate these here. First, the success

of CS-VQE is sensitive to the generator subset \mathcal{F} one chooses to constrain in the stabilizer subspace projection. To date, the most effective method for choosing this subset has been a greedy-search heuristic necessitating $O(N^{d+1})$ VQE simulations where $d \leq N$ is the search depth; this is expensive for NISQ hardware, and there is room for more targeted heuristics. For example, we may draw on chemical intuition to inform the selection of a contextual subspace that captures information about the underlying electronic structure problem. The second obstacle lies in the approach taken to construct the noncontextual sub-Hamiltonian. There is currently no intuition as to what constitutes an effective choice here, although it should be noted that the “optimal” noncontextual subset will not necessarily be that which minimizes the noncontextual ground state energy; some consideration of the resulting contextual subspaces must come into the construction of the noncontextual problem. We leave these issues for future work.

The natural next step is to execute this method on a NISQ computer, challenging the current best-in-class electronic structure simulations from Google, IonQ, and IBM.^{1–3} To achieve this goal, CS-VQE could be combined with techniques of measurement reduction^{20–35} and error mitigation.^{37–44}

Finally, we have written an open-source Python package that facilitates the stabilizer subspace projection techniques of this work, with in-built tapering and CS-VQE functionality. We welcome the reader to make use of our code,⁴⁹ which is freely available on GitHub.

■ ASSOCIATED CONTENT

Data Availability Statement

In the interest of reproducibility, we supply in the [Supporting Information](#) the tapering parameters, CS-VQE model data, and noncontextual projection ansätze which permit chemical precision for the fewest number of qubits with respect to the molecular systems listed in [Table 1](#); the raw data for these results are supplied in ancillary files hosted at <https://arxiv.org/abs/2204.02150> along with an explanatory notebook. This will provide sufficient information for the reader to reproduce [Figures 1, 2 and 3](#).

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.2c00910>.

Further information concerning circuit construction, deriving the noncontextual objective function one must solve to perform CS-VQE and a closer look at the ADAPT-VQE algorithm ([PDF](#))

Tapering parameters, CS-VQE model data and noncontextual projection ansätze which permit chemical precision for the fewest number of qubits with respect to the molecular systems listed in [Table 1](#)—also includes the raw data required to reproduce [Figures 1, 2 and 3](#), with an accompanying notebook ([ZIP](#))

■ AUTHOR INFORMATION

Corresponding Authors

Tim Weaving – Centre for Computational Science, Department of Chemistry, University College London, London WC1H 0AJ, United Kingdom; Email: timothy.weaving.20@ucl.ac.uk

Peter V. Coveney – Centre for Computational Science, Department of Chemistry, University College London,

London WC1H 0AJ, United Kingdom; UCL Centre for Advanced Research Computing, London WC1E 6BT, United Kingdom; Informatics Institute, University of Amsterdam, Amsterdam 1098 XH, the Netherlands; orcid.org/0000-0002-8787-7256; Email: p.v.coveney@ucl.ac.uk

Authors

Alexis Ralli – Centre for Computational Science, Department of Chemistry, University College London, London WC1H 0AJ, United Kingdom

William M. Kirby – Department of Physics and Astronomy, Tufts University, Medford, Massachusetts 02155, United States

Andrew Tranter – Cambridge Quantum Computing, Cambridge CB2 1UB, United Kingdom

Peter J. Love – Department of Physics and Astronomy, Tufts University, Medford, Massachusetts 02155, United States; Computational Science Initiative, Brookhaven National Laboratory, Upton, New York 11973, United States

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.2c00910>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

T.W. and A.R. acknowledge support from the Unitary Fund and the Engineering and Physical Sciences Research Council (EP/S021582/1 and EP/L015242/1, respectively). T.W. also acknowledges support from CBKSciCon Ltd., Atos, Intel and Zapata. W.K. and P.J.L. acknowledge support by the NSF STAQ project (PHY-1818914). W. K. acknowledges support from the National Science Foundation, Grant No. DGE-1842474. P.V.C. is grateful for funding from the European Commission for VECMA (800925) and EPSRC for SEAVEA (EP/W007711/1). We would like to thank both Atos and the Leibniz Supercomputing Centre (LRZ), who each provided access to separate Atos Quantum Learning Machine (QLM) simulators that aided with the computational workload.

■ REFERENCES

- (1) Arute, F.; Arya, K.; Babbush, R.; Bacon, D.; Bardin, J. C.; Barends, R.; Boixo, S.; Broughton, M.; Buckley, B. B.; Buell, D. A.; et al. Hartree-Fock on a superconducting qubit quantum computer. *Science* **2020**, *369*, 1084–1089.
- (2) Nam, Y.; Chen, J.-S.; Plesch, N. C.; Wright, K.; Delaney, C.; Maslov, D.; Brown, K. R.; Allen, S.; Amini, J. M.; Apisdorf, J.; et al. Ground-state energy estimation of the water molecule on a trapped-ion quantum computer. *npj Quantum Information* **2020**, *6*, 33.
- (3) Eddins, A.; Motta, M.; Gujarati, T. P.; Bravyi, S.; Mezzacapo, A.; Hadfield, C.; Sheldon, S. Doubling the size of quantum simulators by entanglement forging. *arXiv preprint* 2021, <https://arxiv.org/abs/2104.10220>.
- (4) Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, X.-Q.; Love, P. J.; Aspuru-Guzik, A.; O'Brien, J. L. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **2014**, *5*, 4213.
- (5) Kitaev, A. Y. Quantum measurements and the Abelian stabilizer problem. *arXiv preprint* 1995, <https://arxiv.org/abs/quant-ph/9511026>.
- (6) Griffiths, D. J.; Schroeter, D. F. *Introduction to quantum mechanics*; Cambridge University Press, 2018.
- (7) McClean, J. R.; Boixo, S.; Smelyanskiy, V. N.; Babbush, R.; Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **2018**, *9*, 4812.

- (8) Wecker, D.; Hastings, M. B.; Troyer, M. Progress towards practical quantum variational algorithms. *Phys. Rev. A* **2015**, *92*, 042303.
- (9) 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* **2017**, *549*, 242–246.
- (10) Rattew, A. G.; Hu, S.; Pistoia, M.; Chen, R.; Wood, S.A domain-agnostic, noise-resistant, hardware-efficient evolutionary variational quantum eigensolver. *arXiv preprint2019*, <https://arxiv.org/abs/1910.09694>.
- (11) Wiersema, R.; Zhou, C.; de Sereville, Y.; Carrasquilla, J. F.; Kim, Y. B.; Yuen, H. Exploring entanglement and optimization within the hamiltonian variational ansatz. *PRX Quantum* **2020**, *1*, 020319.
- (12) Harrigan, M. P.; Sung, K. J.; Neeley, M.; Satzinger, K. J.; Arute, F.; Arya, K.; Atalaya, J.; Bardin, J. C.; Barends, R.; Boixo, S.; et al. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat. Phys.* **2021**, *17*, 332–336.
- (13) Taube, A. G.; Bartlett, R. J. New perspectives on unitary coupled-cluster theory. *Int. J. Quantum Chem.* **2006**, *106*, 3393–3401.
- (14) Romero, J.; Babbush, R.; McClean, J. R.; Hempel, C.; Love, P. J.; Aspuru-Guzik, A. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Science and Technology* **2019**, *4*, 014008.
- (15) Gard, B. T.; Zhu, L.; Barron, G. S.; Mayhall, N. J.; Economou, S. E.; Barnes, E. Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm. *npj Quantum Information* **2020**, *6*, 10.
- (16) Grimsley, H. R.; Economou, S. E.; Barnes, E.; Mayhall, N. J. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.* **2019**, *10*, 3007.
- (17) Tang, H. L.; Shkolnikov, V.; Barron, G. S.; Grimsley, H. R.; Mayhall, N. J.; Barnes, E.; Economou, S. E. qubit-adapt-vqe: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor. *PRX Quantum* **2021**, *2*, 020310.
- (18) Shkolnikov, V.; Mayhall, N. J.; Economou, S. E.; Barnes, E. Avoiding symmetry roadblocks and minimizing the measurement overhead of adaptive variational quantum eigensolvers. *arXiv preprint2021*, <https://arxiv.org/abs/2109.05340>.
- (19) Fedorov, D. A.; Alexeev, Y.; Gray, S. K.; Otten, M. Unitary Selective Coupled-Cluster Method. *arXiv preprint2021*, <https://arxiv.org/abs/2109.12652>.
- (20) Babbush, R.; Wiebe, N.; McClean, J.; McClain, J.; Neven, H.; Chan, G. K.-L. Low-depth quantum simulation of materials. *Physical Review X* **2018**, *8*, 011044.
- (21) Wang, D.; Higgott, O.; Brierley, S. Accelerated variational quantum eigensolver. *Phys. Rev. Lett.* **2019**, *122*, 140504.
- (22) Verteletskyi, V.; Yen, T.-C.; Izmaylov, A. F. Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *J. Chem. Phys.* **2020**, *152*, 124114.
- (23) Jena, A.; Genin, S.; Mosca, M. Pauli partitioning with respect to gate sets. *arXiv preprint2019*, <https://arxiv.org/abs/1907.07859>.
- (24) Gokhale, P.; Angiuli, O.; Ding, Y.; Gui, K.; Tomesh, T.; Suchara, M.; Martonosi, M.; Chong, F. T. Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families. *arXiv preprint2019*, <https://arxiv.org/abs/1907.13623>.
- (25) Yen, T.-C.; Verteletskyi, V.; Izmaylov, A. F. Measuring all compatible operators in one series of single-qubit measurements using unitary transformations. *J. Chem. Theory Comput.* **2020**, *16*, 2400–2409.
- (26) Huggins, W. J.; McClean, J. R.; Rubin, N. C.; Jiang, Z.; Wiebe, N.; Whaley, K. B.; Babbush, R. Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers. *npj Quantum Information* **2021**, *7*, 23.
- (27) Gokhale, P.; Angiuli, O.; Ding, Y.; Gui, K.; Tomesh, T.; Suchara, M.; Martonosi, M.; Chong, F. T. $O(N^3)$ measurement Cost for Variational Quantum Eigensolver on Molecular Hamiltonians. *IEEE Transactions on Quantum Engineering* **2020**, *1*, 1–24.
- (28) Torlai, G.; Mazzola, G.; Carleo, G.; Mezzacapo, A. Precise measurement of quantum observables with neural-network estimators. *Physical Review Research* **2020**, *2*, 022060.
- (29) Crawford, O.; van Straaten, B.; Wang, D.; Parks, T.; Campbell, E.; Brierley, S. Efficient quantum measurement of Pauli operators in the presence of finite sampling error. *Quantum* **2021**, *5*, 385.
- (30) Huang, H.-Y.; Kueng, R.; Preskill, J. Predicting many properties of a quantum system from very few measurements. *Nat. Phys.* **2020**, *16*, 1050–1057.
- (31) Hadfield, C.; Bravyi, S.; Raymond, R.; Mezzacapo, A. Measurements of quantum Hamiltonians with locally-biased classical shadows. *arXiv preprint2020*, <https://arxiv.org/abs/2006.15788>.
- (32) Izmaylov, A. F.; Yen, T.-C.; Lang, R. A.; Verteletskyi, V. Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method. *J. Chem. Theory Comput.* **2020**, *16*, 190–195.
- (33) Zhao, A.; Tranter, A.; Kirby, W. M.; Ung, S. F.; Miyake, A.; Love, P. J. Measurement reduction in variational quantum algorithms. *Phys. Rev. A* **2020**, *101*, 062322.
- (34) Bonet-Monroig, X.; Babbush, R.; O'Brien, T. E. Nearly optimal measurement scheduling for partial tomography of quantum states. *Physical Review X* **2020**, *10*, 031064.
- (35) Ralli, A.; Love, P. J.; Tranter, A.; Coveney, P. V. Implementation of measurement reduction for the variational quantum eigensolver. *Physical Review Research* **2021**, *3*, 033195.
- (36) Ralli, A.; Weaving, T.; Tranter, A.; Kirby, W. M.; Love, P. J.; Coveney, P. V. Unitary Partitioning and the Contextual Subspace Variational Quantum Eigensolver. *arXiv preprint2022*, <https://arxiv.org/abs/2207.03451>.
- (37) Temme, K.; Bravyi, S.; Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **2017**, *119*, 180509.
- (38) Li, Y.; Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Physical Review X* **2017**, *7*, 021050.
- (39) Endo, S.; Benjamin, S. C.; Li, Y. Practical quantum error mitigation for near-future applications. *Physical Review X* **2018**, *8*, 031027.
- (40) Kandala, A.; Temme, K.; Córcoles, A. D.; Mezzacapo, A.; Chow, J. M.; Gambetta, J. M. Error mitigation extends the computational reach of a noisy quantum processor. *Nature* **2019**, *567*, 491–495.
- (41) Giurgica-Tiron, T.; Hindy, Y.; LaRose, R.; Mari, A.; Zeng, W. J. Digital zero noise extrapolation for quantum error mitigation. *2020 IEEE International Conference on Quantum Computing and Engineering (QCE); IEEE: 2020*; pp 306–316.
- (42) He, A.; Nachman, B.; de Jong, W. A.; Bauer, C. W. Zero-noise extrapolation for quantum-gate error mitigation with identity insertions. *Phys. Rev. A* **2020**, *102*, 012426.
- (43) Endo, S.; Cai, Z.; Benjamin, S. C.; Yuan, X. Hybrid quantum-classical algorithms and quantum error mitigation. *J. Phys. Soc. Jpn.* **2021**, *90*, 032001.
- (44) Huggins, W. J.; McArdle, S.; O'Brien, T. E.; Lee, J.; Rubin, N. C.; Boixo, S.; Whaley, K. B.; Babbush, R.; McClean, J. R. Virtual distillation for quantum error mitigation. *Physical Review X* **2021**, *11*, 041036.
- (45) Kirby, W. M.; Tranter, A.; Love, P. J. Contextual Subspace Variational Quantum Eigensolver. *Quantum* **2021**, *5*, 456.
- (46) Radin, M. D.; Johnson, P. Classically-Boosted Variational Quantum Eigensolver. *arXiv preprint2021*, <https://arxiv.org/abs/2106.04755>.
- (47) Bravyi, S.; Gambetta, J. M.; Mezzacapo, A.; Temme, K. Tapering off qubits to simulate fermionic Hamiltonians. *arXiv preprint2017*, <https://arxiv.org/abs/1701.08213>.
- (48) Setia, K.; Chen, R.; Rice, J. E.; Mezzacapo, A.; Pistoia, M.; Whitfield, J. D. Reducing qubit requirements for quantum simulations using molecular point group symmetries. *J. Chem. Theory Comput.* **2020**, *16*, 6091–6097.

- (49) Ralli, A.; Weaving, T. symmer. <https://github.com/UCL-CCS/symmer>, 2022.
- (50) Gottesman, D. Stabilizer codes and quantum error correction. Ph.D. Thesis, California Institute of Technology, 1997.
- (51) Gottesman, D. An introduction to quantum error correction and fault-tolerant quantum computation. *Quantum information science and its contributions to mathematics*; Proceedings of Symposia in Applied Mathematics; American Mathematical Society: 2010; pp 13–58.
- (52) Dehaene, J.; De Moor, B. Clifford group, stabilizer states, and linear and quadratic operations over GF (2). *Phys. Rev. A* **2003**, *68*, 042318.
- (53) Nielsen, M. A.; Chuang, I. L. *Quantum Computation and Quantum Information*, 10th Anniversary edition; Cambridge University Press, 2010.
- (54) Kraus, K.; Böhm, A.; Dollard, J. D.; Wootters, W. *States, effects, and operations: fundamental notions of quantum theory*. Lectures in Mathematical Physics at the University of Texas at Austin; Springer-Verlag: 1983, Vol. 190.
- (55) Bell, J. S. On the Einstein Podolsky Rosen paradox. *Physique Fizika* **1964**, *1*, 195.
- (56) Bell, J. S. On the problem of hidden variables in quantum mechanics. *Rev. Mod. Phys.* **1966**, *38*, 447.
- (57) Kochen, S.; Specker, E. P. *The logico-algebraic approach to quantum mechanics*; Springer: 1975; pp 293–328.
- (58) Kirby, W. M.; Love, P. J. Contextuality test of the nonclassicality of variational quantum eigensolvers. *Phys. Rev. Lett.* **2019**, *123*, 200501.
- (59) Raussendorf, R.; Bermejo-Vega, J.; Tyhurst, E.; Okay, C.; Zurek, M. Phase-space-simulation method for quantum computation with magic states on qubits. *Phys. Rev. A* **2020**, *101*, 012350.
- (60) Kirby, W. M.; Love, P. J. Classical simulation of noncontextual Pauli Hamiltonians. *Phys. Rev. A* **2020**, *102*, 032418.
- (61) Spekkens, R. W. *Quantum Theory: Informational Foundations and Foils*; Springer Netherlands: Dordrecht, 2016; pp 83–135.
- (62) Schwerdtfeger, P. The pseudopotential approximation in electronic structure theory. *ChemPhysChem* **2011**, *12*, 3143–3155.
- (63) Wecker, D.; Bauer, B.; Clark, B. K.; Hastings, M. B.; Troyer, M. Gate-count estimates for performing quantum chemistry on small quantum computers. *Phys. Rev. A* **2014**, *90*, 022305.
- (64) Poulin, D.; Hastings, M. B.; Wecker, D.; Wiebe, N.; Doherty, A. C.; Troyer, M. The Trotter step size required for accurate quantum simulation of quantum chemistry. *arXiv preprint*2014, <https://arxiv.org/abs/1406.4920>.
- (65) Kowalski, K. Properties of coupled-cluster equations originating in excitation sub-algebras. *J. Chem. Phys.* **2018**, *148*, 094104.
- (66) Bauman, N. P.; Bylaska, E. J.; Krishnamoorthy, S.; Low, G. H.; Wiebe, N.; Granade, C. E.; Roetteler, M.; Troyer, M.; Kowalski, K. Downfolding of many-body Hamiltonians using active-space models: Extension of the sub-system embedding sub-algebras approach to unitary coupled cluster formalisms. *J. Chem. Phys.* **2019**, *151*, 014107.
- (67) Ralli, A.; Williams, M. I.; Coveney, P. V. A Scalable Approach to Quantum Simulation via Projection-based Embedding. *arXiv preprint*, arXiv:2203.01135, 2022, <https://arxiv.org/abs/2203.01135>.
- (68) Johnson, R. D., IINIST Computational Chemistry Comparison and Benchmark Database. *NIST Standard Reference Database Number 101*, 2020, <https://cccbdb.nist.gov/>.
- (69) Anis, M. S.; Mitchell, A.; Abraham, H.; Agarwal, R.; Agliardi, G.; Aharoni, M.; Ajith, V.; Akhalwaya, I. Y.; Aleksandrowicz, G.; Alexander, T.; et al. *Qiskit: An Open-source Framework for Quantum Computing*, 2021, DOI: 10.5281/zenodo.2573505.
- (70) Sun, Q.; Berkelbach, T. C.; Blunt, N. S.; Booth, G. H.; Guo, S.; Li, Z.; Liu, J.; McClain, J. D.; Sayfutyarova, E. R.; Sharma, S.; et al. PySCF: the Python-based simulations of chemistry framework. *Wiley Interdisciplinary Reviews: Computational Molecular Science* **2018**, *8*, No. e1340.
- (71) Ryabinkin, I. G.; Lang, R. A.; Genin, S. N.; Izmaylov, A. F. Iterative qubit coupled cluster approach with efficient screening of generators. *J. Chem. Theory Comput.* **2020**, *16*, 1055–1063.
- (72) Jordan, P.; Wigner, E. P. *The Collected Works of Eugene Paul Wigner*; Springer, 1993; pp 109–129.
- (73) Kingma, D. P.; Ba, J. Adam: A method for stochastic optimization. *arXiv preprint*2014, <https://arxiv.org/abs/1412.6980>.
- (74) Parrish, R. M.; Hohenstein, E. G.; McMahon, P. L.; Martinez, T. J. Hybrid quantum/classical derivative theory: Analytical gradients and excited-state dynamics for the multistate contracted variational quantum eigensolver. *arXiv preprint*2019, <https://arxiv.org/abs/1906.08728>.
- (75) Wierichs, D.; Gogolin, C.; Kastoryano, M. Avoiding local minima in variational quantum eigensolvers with the natural gradient optimizer. *Physical Review Research* **2020**, *2*, 043246.
- (76) Pellow-Jarman, A.; Sinayskiy, I.; Pillay, A.; Petruccione, F. A comparison of various classical optimizers for a variational quantum linear solver. *Quantum Information Processing* **2021**, *20*, 202.
- (77) Fletcher, R. *Practical methods of optimization*; John Wiley & Sons: 2013.
- (78) Stokes, J.; Izaac, J.; Killoran, N.; Carleo, G. Quantum natural gradient. *Quantum* **2020**, *4*, 269.
- (79) Hagberg, A.; Swart, P.; S Chult, D. Exploring network structure, dynamics, and function using NetworkX. *Proceedings of the 7th Python in Science Conference*; 2008.
- (80) Arrasmith, A.; Cincio, L.; Somma, R. D.; Coles, P. J. Operator sampling for shot-frugal optimization in variational algorithms. *arXiv preprint*2020, <https://arxiv.org/abs/2004.06252>.

Recommended by ACS

Improved Algorithms of Quantum Imaginary Time Evolution for Ground and Excited States of Molecular Systems

Takashi Tsuchimochi, Kazuki Sasasako, et al.

JANUARY 13, 2023

JOURNAL OF CHEMICAL THEORY AND COMPUTATION

READ 

Molecular Energy Landscapes of Hardware-Efficient Ansätze in Quantum Computing

Boy Choy and David J. Wales

FEBRUARY 07, 2023

JOURNAL OF CHEMICAL THEORY AND COMPUTATION

READ 

Resolution of Electronic States in Heisenberg Cluster Models within the Unitary Group Approach

Giovanni Li Manni, Niklas Liebermann, et al.

FEBRUARY 03, 2023

JOURNAL OF CHEMICAL THEORY AND COMPUTATION

READ 

Eliminating Artificial Boundary Conditions in Time-Dependent Density Functional Theory Using Fourier Contour Deformation

Jason Kaye, Angel Rubio, et al.

FEBRUARY 14, 2023

JOURNAL OF CHEMICAL THEORY AND COMPUTATION

READ 

Get More Suggestions >