Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Benchmarking a Hardware-Efficient Variational Quantum Eigensolver for Ground-State Energies and Potential Energy Curves of Nine Small Molecules against Exact Active-Space and Full Configuration Interaction References

Nambury Surendra Babu

nsbabusk@gmail.com

The University of Dodoma

Research Article

Keywords: variational quantum eigensolver, hardware-efficient ansatz, quantum computational chemistry, potential energy curves, full configuration interaction, active space, Qiskit, PySCF

Posted Date: November 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11453036/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: No competing interests reported.

Abstract

The variational quantum eigensolver (VQE) is the leading near-term quantum algorithm for molecular ground states, but its accuracy is often reported with inconsistent reference values and metrics. We present a reproducible benchmark of a hardware-efficient VQE ansatz (R_Y–R_Z rotations with linear CNOT entanglement) for H_2, HeH^+, LiH, BeH_2, H_2O, NH_3, CH_4, CO_2 and N_2 in the STO-3G basis. Hamiltonians were generated with PySCF and Qiskit Nature, encoded with the parity mapping and two-qubit reduction on 2, 4 or 6 qubits, and minimised with SLSQP and analytic gradients in noiseless statevector simulations, with a penalty term enforcing the correct electron number. Energies are compared with exact active-space (CASCI) and all-electron full configuration interaction (FCI) references in the same basis, at experimental geometries and along potential energy curves (144 geometries in total). For the two- and four-qubit systems, VQE reproduces the exact energies, potential energy curves, bond lengths and harmonic frequencies to numerical precision. For the six-qubit systems the equilibrium errors are 0.3–7.5 mE_h and reach 51.1 mE_h along the curves, so chemical accuracy (1.6 mE_h) is reached only for H_2O and CH_4; increasing the depth to reps = 4–6 (20–30 CNOTs) restores chemical accuracy at equilibrium, and UCCSD reaches chemical accuracy for all nine molecules but requires 476 instead of 15 CNOT gates. Active-space truncation errors (1–139 mE_h) exceed the VQE errors, and deviations of STO-3G bond lengths from experiment (up to 0.14 Å) are basis-set effects shared by FCI. The results show why VQE errors should be reported in absolute terms against exact references for the same Hamiltonian, and they provide a baseline for evaluating noise-aware and symmetry-preserving ansätze.

1. Introduction

Simulating molecular electronic structure is widely regarded as one of the most natural applications of quantum computers, because the exponential growth of the many-electron Hilbert space that limits exact classical methods can, in principle, be represented with a number of qubits that grows only linearly with the number of spin orbitals [1–5]. Whether, and for which problems, this translates into a practical advantage remains an open question [6, 7], and fault-tolerant hardware capable of running phase-estimation-type algorithms on chemically relevant systems is not yet available. In the current noisy intermediate-scale quantum (NISQ) era [8], the variational quantum eigensolver (VQE) [9, 10] has become the most widely studied algorithm for molecular ground states, and it has been demonstrated experimentally for H_2, LiH, BeH_2, H_2O and related systems on superconducting and trapped-ion processors [11–15]. Comprehensive reviews of the method and its many variants are available [16–18].

The accuracy of a VQE calculation is governed by three separate factors: (i) the one-electron basis set and active space, which define the Hamiltonian being solved; (ii) the expressibility of the parameterised circuit (ansatz) within that Hamiltonian; and (iii) the ability of the classical optimiser to locate the global minimum of a non-convex cost landscape. Chemically inspired ansätze such as unitary coupled cluster (UCCSD) [19] and its adaptive variants [20] conserve particle number and spin by construction but lead to deep circuits, whereas hardware-efficient ansätze (HEA) [12] built from native single-qubit rotations and nearest-neighbour CNOT gates are shallow but carry no physical symmetry, may require many parameters, and are prone to local minima and barren plateaus [21–23]. Because an HEA does not conserve the electron number, an unconstrained minimisation can also converge to states of the wrong particle-number sector, which must be prevented with constraints or penalty terms [24, 25].

Many VQE studies report results for one or two molecules at a single geometry, and accuracy is frequently expressed as a percentage ratio of total energies. Such a ratio is not a meaningful measure in chemistry: total energies of first-row molecules are tens to hundreds of hartree, so a ratio of 99% can conceal errors of several hartree, i.e. thousands of kJ mol^{-1}. The chemically relevant measure is the absolute error relative to the exact solution of the *same* Hamiltonian, compared with the conventional chemical-accuracy threshold of 1.6 mE_h (≈ 1 kcal mol^{-1}), together with the fraction of

correlation energy recovered and the error in molecular properties such as equilibrium bond lengths and vibrational frequencies [3, 18, 26].

In this work we present a systematic and fully reproducible benchmark of a hardware-efficient VQE for nine small molecules of increasing complexity—H_2, HeH^+, LiH, BeH_2, H_2O, NH_3, CH_4, CO_2 and N_2—using 2, 4 or 6 qubits. The specific contributions are: (1) ground-state energies at experimental geometries, compared with exact active-space (CASCI) and all-electron full configuration interaction (FCI) energies computed in the same STO-3G basis, which cleanly separates ansatz/optimiser error from active-space truncation error; (2) complete potential energy curves (16 geometries per molecule, 144 independent VQE optimisations) together with maximum errors and non-parallelity errors; (3) equilibrium bond lengths and, for the diatomics, harmonic vibrational frequencies extracted from the VQE curves; (4) a circuit-depth convergence study and a comparison with the UCCSD ansatz; and (5) a full account of quantum resources (qubits, Pauli terms, parameters and CNOT counts). All energies are reported in hartree (E_h) and errors in millihartree (mE_h).

2. Theoretical Background

2.1 Electronic Hamiltonian and active-space reduction

Within the Born–Oppenheimer approximation the non-relativistic electronic Hamiltonian, expressed in an orthonormal basis of molecular spin orbitals $\{\varphi_p\}$, takes the second-quantised form [27, 28]

$$\hat{H} = E_{nuc} + \Sigma_{pq}\, h_{pq}\, \hat{a}^\dagger_p\, \hat{a}_q + \tfrac{1}{2}\, \Sigma_{pqrs}\, g_{pqrs}\, \hat{a}^\dagger_p\, \hat{a}^\dagger_q\, \hat{a}_r\, \hat{a}_s \quad (1)$$

where E_{nuc} is the nuclear repulsion energy, h_{pq} are one-electron (kinetic and nuclear-attraction) integrals, g_{pqrs} are two-electron repulsion integrals, and the fermionic operators obey $\{\hat{a}_p, \hat{a}^\dagger_q\} = \delta_{pq}$. The molecular orbitals are obtained from a restricted Hartree–Fock (RHF) calculation in the STO-3G minimal basis [29]. Diagonalising Eq. (1) in the full determinant space gives the FCI energy, the exact solution for the chosen basis [30]. To limit the number of qubits, we partition the orbitals into doubly occupied inactive (core) orbitals, an active space of n_o orbitals containing n_e electrons, and unoccupied external orbitals. The inactive orbitals are folded into an effective one-electron operator and a constant core energy E_{core}, giving an active-space Hamiltonian whose exact ground-state eigenvalue is the complete-active-space configuration-interaction (CASCI) energy. CASCI is therefore the exact target of any VQE performed on the active-space Hamiltonian, whereas the difference $E_{CASCI} - E_{FCI}$ measures the error introduced by the active-space truncation itself.

2.2 Fermion-to-qubit mapping

Fermionic operators are mapped to Pauli operators so that the Hamiltonian becomes a weighted sum of Pauli strings,

$$\hat{H}_q = \Sigma_k\, c_k\, P_k,\ P_k \in \{I, X, Y, Z\}^{\otimes n} \quad (2)$$

In the Jordan–Wigner (JW) encoding [31] each qubit stores the occupation of one spin orbital, so $2n_o$ qubits are needed. In the parity encoding [32, 33] qubit j stores the parity of the occupations of spin orbitals 0…j. When spin orbitals are ordered with all α orbitals first, the parity encoding makes the conserved α-electron and total-electron parities explicit on two qubits, which can be replaced by their known eigenvalues (two-qubit Z_2 reduction) [34]. The resulting Hamiltonian acts on $n = 2n_o - 2$ qubits and is exactly isospectral with the original Hamiltonian within the correct parity sectors. This is the encoding used here; it is the reason why, for example, H_2 in STO-3G (two spatial orbitals, four spin orbitals) requires two qubits rather than four.

2.3 Variational quantum eigensolver

VQE prepares a trial state $|\Psi(\theta)\rangle = U(\theta)|\Phi_0\rangle$ on the quantum processor, measures the energy expectation value as a sum of Pauli expectation values, and uses a classical optimiser to minimise it [9, 10]:

$E(\theta) = \langle\Psi(\theta)|\hat{H}_q|\Psi(\theta)\rangle = \Sigma_k\, c_k\, \langle\Psi(\theta)|P_k|\Psi(\theta)\rangle \geq E_0$ (3)

By the variational principle, $E(\theta)$ is an upper bound to the lowest eigenvalue E_0 of $\hat{H}_q$ in the space accessible to the ansatz. Consequently, a correctly implemented VQE can never give an energy below the exact (CASCI) eigenvalue of the Hamiltonian it is minimising; any such result indicates an inconsistency between the Hamiltonian, the reference value or the electron-number sector. For Pauli-rotation gates $R_P(\theta) = \exp(-i\theta P/2)$, exact gradients are given by the parameter-shift rule [35, 36]:

$\partial E/\partial\theta_i = ½\,[E(\theta + (\pi/2)e_i) - E(\theta - (\pi/2)e_i)]$ (4)

2.4 Hardware-efficient ansatz and particle-number constraint

The hardware-efficient ansatz used here (Qiskit *n_local*) consists of reps + 1 layers of single-qubit R_Y and R_Z rotations on every qubit, separated by reps layers of CNOT gates in a linear (nearest-neighbour) chain [12]:

$U(\theta) = L_{rot}(\theta^{(reps)})\, \Pi_{l=reps-1}^{0}\, [\, U_{ent}\, L_{rot}(\theta^{(l)})\,],\ L_{rot}(\theta^{(l)}) = \otimes_i\, R_Z(\theta^{(l)}_{i,2})\, R_Y(\theta^{(l)}_{i,1})$ (5)

The number of variational parameters is $N_\theta = 2n(reps + 1)$ and the number of CNOT gates is $(n - 1)\cdot reps$. The ansatz acts on the Hartree–Fock reference $|\Phi_0\rangle = |HF\rangle$, which in the parity encoding is prepared with X gates only. Because the HEA does not conserve particle number, the cost function minimised in practice is

$C(\theta) = \langle\Psi(\theta)|\hat{H}_q + \mu\,(\mathcal{N} - N)^2|\Psi(\theta)\rangle$ (6)

where $\mathcal{N}$ is the qubit representation of the electron-number operator in the active space, N is the target number of active electrons and $\mu = 1\ E_h$ [24,25]. The reported VQE energy is the expectation value of $\hat{H}_q$ alone (plus E_{core}) evaluated at the optimal parameters, and $\langle\mathcal{N}\rangle$ is monitored to confirm that the correct sector is reached.

2.5 Error metrics

We report the absolute VQE error $\Delta E = E_{VQE} - E_{CASCI}$ in mE_h; the active-space truncation error $E_{CASCI} - E_{FCI}$; the percentage of active-space correlation energy recovered, $100 \times (E_{VQE} - E_{RHF})/(E_{CASCI} - E_{RHF})$; and, along each potential energy curve, the maximum error and the non-parallelity error $NPE = \max(\Delta E) - \min(\Delta E)$. Errors below 1.6 mE_h are classified as chemically accurate.

3. Computational Details

3.1 Software

Molecular integrals, RHF, CASCI and FCI calculations were performed with PySCF 2.14.0 [37, 38]. Qubit Hamiltonians, the parity mapping with two-qubit reduction, the Hartree–Fock initial state and the UCCSD ansatz were generated with Qiskit 2.6.0 [39] and Qiskit Nature 0.8.0 [40]; the VQE driver and parameter-shift gradients of Qiskit Algorithms 0.4.0 [41] were used for validation. Energies were evaluated by exact (noiseless, infinite-shot) statevector simulation. For the large number of optimisations required for the potential energy curves, the same statevector simulation and gradient were implemented in NumPy [42] using an adjoint (reverse-mode) evaluation that is mathematically identical to the parameter-shift rule; this implementation reproduces the Qiskit statevector to within 10^{-15} in fidelity and the parameter-

shift gradient to within 2×10^{-15} E_h. Optimisation used SLSQP [43] as implemented in SciPy [44]. Figures were prepared with Matplotlib [45]. All calculations were run on a Linux workstation with Python 3.13.16.

3.2 Molecular geometries

Equilibrium calculations used experimental geometries [46, 47]: r(H–H) = 0.7414 Å; r(He–H^+) = 0.7743 Å; r(Li–H) = 1.5957 Å; r(Be–H) = 1.3264 Å ($D_{\infty h}$); r(O–H) = 0.9572 Å, ∠HOH = 104.52°; r(N–H) = 1.012 Å, ∠HNH = 106.67°; r(C–H) = 1.087 Å (T_d); r(C = O) = 1.160 Å ($D_{\infty h}$); r(N ≡ N) = 1.0977 Å. For the polyatomic molecules the potential energy curve is a cut along the totally symmetric stretching coordinate, in which all equivalent bonds are stretched simultaneously with bond angles fixed at their experimental values. Each curve contains 16 points at R/R_{exp} = 0.75, 0.80, …, 1.20, 1.30, 1.40, 1.55, 1.70, 1.85 and 2.00. Equilibrium distances r_e were obtained by minimising a cubic-spline interpolant of the computed points, and harmonic frequencies of the diatomics from its second derivative, $\omega_e = (k/\mu)^{1/2}/(2\pi c)$.

3.3 Active spaces

Active orbitals were chosen from symmetry-adapted RHF orbitals by irreducible representation, so that the same chemically equivalent orbitals are selected at every geometry and degenerate partners (π_x/π_y, e and t_2 components) are never split arbitrarily. The resulting active spaces are summarised in Table 1. For H_2 and HeH^+ the active space comprises all orbitals, so CASCI and FCI coincide. For LiH and BeH_2, the 1s core is frozen and the σ-bonding valence orbitals are correlated with σ-type virtual orbitals. For H_2O and NH_3, the bonding orbitals that change most upon symmetric stretching and their antibonding partners are active ($3a_1$, $1b_2$ → $4a_1$, $2b_2$ for H_2O; the degenerate 1e → 2e pair for NH_3); for CH_4, two components of the bonding t_2 set and the corresponding t_2^* components; for N_2, the bonding π_u and antibonding π_g pairs; and for CO_2, the highest occupied π_g and lowest unoccupied π_u pairs.

3.4 VQE protocol

The ansatz repetitions were reps = 2 for the two-qubit systems and reps = 3 for the four- and six-qubit systems. At each equilibrium geometry the optimisation was started from the Hartree–Fock point (θ = 0) and from 20 random parameter vectors drawn uniformly from [−π, π]; the lowest converged energy is reported. Along each potential energy curve, four starts were used at each point in a forward sweep from short to long bond lengths (the optimum of the previous geometry, θ = 0, and two random vectors). For the six-qubit systems a second, backward sweep from long to short bond lengths was performed with the same number of starts, together with 30 additional random starts at the equilibrium geometry; at every geometry the lowest energy from all optimisations is reported. SLSQP was run with a function tolerance of 10^{-12} and at most 2000 iterations. The penalty weight μ = 1 E_h was sufficient to keep $|\langle\hat{N}\rangle - N|$ below 4×10^{-5} at every point. For comparison, UCCSD-VQE energies were obtained at the equilibrium geometries with the same mapping, initial state and optimiser, and an ansatz-depth study (reps = 1–6, 11 starts each) was performed for H_2O and N_2.

4. Results and Discussion

4.1 Hamiltonians and quantum resources

Table 1 summarises the active spaces and the quantum resources required. The qubit Hamiltonians contain between 5 (H_2) and 165 Pauli strings. The two linear six-qubit molecules (CO_2 and N_2) have fewer terms (85) than the non-linear ones (165) because the π-type active orbitals of a linear molecule have higher symmetry, so more two-electron integrals vanish. These term counts determine the number of distinct measurement settings needed per energy evaluation on hardware, before any grouping of commuting terms. The hardware-efficient circuits are very shallow: 2, 9 and 15 CNOT

gates for 2, 4 and 6 qubits, respectively, with 12, 32 and 48 parameters (Fig. 2). For comparison, the UCCSD ansatz for the same active spaces has far fewer parameters (3, 8 and 26) but, after compilation to CNOT and single-qubit gates, needs 72 CNOTs for LiH and 476 CNOTs for the six-qubit systems.

Table 1
Active spaces and quantum resources. N_e: total electrons; N_{MO}: STO-3G spatial orbitals; (n_e, n_o): active electrons and orbitals; N_q: qubits after parity mapping with two-qubit reduction; N_P: Pauli strings in $\hat{H}_q$; N_θ and N_{CX}: parameters and CNOT gates of the hardware-efficient ansatz (HEA) and of UCCSD (CNOTs after transpilation to {CX, U}).

Molecule	N_e	N_{MO}	(n_e, n_o)	N_q	N_P	HEA reps	HEA N_θ	HEA N_{CX}	UCCSD N_θ	UCCSD N_{CX}
H_2	2	2	(2, 2)	2	5	2	12	2	3	2
HeH^+	2	2	(2, 2)	2	9	2	12	2	3	2
LiH	4	6	(2, 3)	4	100	3	32	9	8	72
BeH_2	6	7	(4, 3)	4	52	3	32	9	8	70
H_2O	10	7	(4, 4)	6	165	3	48	15	26	476
NH_3	10	8	(4, 4)	6	165	3	48	15	26	476
CH_4	10	9	(4, 4)	6	165	3	48	15	26	476
CO_2	22	15	(4, 4)	6	85	3	48	15	26	476
N_2	14	10	(4, 4)	6	85	3	48	15	26	476

4.2 Ground-state energies at the experimental geometries

Table 2 lists the RHF, FCI, CASCI and VQE energies at the experimental geometries. Every VQE energy lies above (or equal to) the corresponding CASCI eigenvalue, as required by the variational principle, and the expectation value of the electron-number operator equals the target value to within 4×10^{-5}. For the two- and four-qubit systems (H_2, HeH^+, LiH and BeH_2) the HEA reproduces the exact CASCI energy to better than 10^{-3} mE_h, i.e. it is exact for practical purposes. For H_2 and HeH^+, whose active space contains all orbitals, the VQE energy is therefore identical to the FCI/STO-3G energy (−1.137270 and − 2.851468 E_h).

For the six-qubit systems the picture is different. The VQE errors relative to CASCI range from 0.34 mE_h (CH_4) to 7.45 mE_h (N_2), and H_2O and CH_4 reach chemical accuracy at this geometry, even though the 48 parameters exceed the 35 independent real amplitudes needed to specify a real ground state in the relevant sector (36 determinants with two α and two β electrons in four orbitals). The ansatz recovers between 90.1% and 98.2% of the active-space correlation energy. The difficulty is therefore not simply the number of free parameters but how they are arranged in the layered circuit and how easily they can be optimised: different initialisations converge to different local minima (Section 4.5), a behaviour known for symmetry-agnostic hardware-efficient circuits [21–23, 48]. In contrast, the particle- and spin-conserving UCCSD ansatz reaches the CASCI energy to within 0.38 mE_h (well inside chemical accuracy) for all nine molecules, but at the price of many more CNOT gates (476 versus 15 for the six-qubit systems; Table 1).

Table 2 also shows that, for every molecule beyond H_2 and HeH^+, the error introduced by the active space ($E_{CASCI} - E_{FCI}$ = 1.3–139.0 mE_h) is larger than the VQE error itself. The chosen active spaces capture between 24% (CH_4) and 94% (LiH)

of the STO-3G correlation energy. This distinction is essential when VQE results are compared with literature FCI values: a VQE calculation can only be expected to approach the exact eigenvalue of the Hamiltonian that it actually minimises, and in a reduced active space that eigenvalue is CASCI, not all-electron FCI.

Table 2
Total energies (E_h) at the experimental geometries in the STO-3G basis and errors (mE_h). $\Delta E_{VQE} = E_{VQE} - E_{CASCI}$; $\Delta E_{AS} = E_{CASCI} - E_{FCI}$ (active-space truncation error); Rec.: percentage of active-space correlation energy recovered by VQE; $\Delta E_{UCCSD} = E_{UCCSD} - E_{CASCI}$.

Molecule	E_{RHF}	E_{FCI}	E_{CASCI}	E_{VQE}	ΔE_{VQE}	ΔE_{AS}	Rec. (%)	ΔE_{UCCSD}
H_2	−1.116684	−1.137270	−1.137270	−1.137270	0.000	0.00	100.0	0.000
HeH^+	−2.841838	−2.851468	−2.851468	−2.851468	0.000	0.00	100.0	0.000
LiH	−7.862002	−7.882392	−7.881135	−7.881135	0.000	1.26	100.0	0.000
BeH_2	−15.560312	−15.595177	−15.569979	−15.569979	0.000	25.20	100.0	0.000
H_2O	−74.962928	−75.012404	−74.995431	−74.994056	1.374	16.97	95.8	0.039
NH_3	−55.454046	−55.519120	−55.481592	−55.478866	2.726	37.53	90.1	0.004
CH_4	−39.726810	−39.805681	−39.745389	−39.745050	0.339	60.29	98.2	0.002
CO_2	−185.064696	−185.276125	−185.137124	−185.132931	4.193	139.00	94.2	0.063
N_2	−107.495893	−107.652829	−107.598356	−107.590903	7.453	54.47	92.7	0.383

All VQE values are noiseless statevector results; chemical accuracy corresponds to $\Delta E < 1.6$ mE_h.

4.3 Potential energy curves

Figure 3 shows the RHF, FCI, CASCI and VQE potential energy curves, and Table 3 summarises the VQE errors along each curve. For the two- and four-qubit molecules the VQE curves coincide with the CASCI curves at all 16 geometries (maximum error 0.0000 mE_h), including the dissociation region of H_2 and LiH where RHF fails qualitatively. For the six-qubit molecules, the mean error along the curve ranges from 0.95 mE_h (CH_4) to 12.33 mE_h (H_2O) and the non-parallelity error from 3.94 to 49.78 mE_h. The error is not constant along the curve (Fig. 4a), and its behaviour is molecule-specific. It grows steadily with bond stretching for CH_4 and NH_3, as static correlation increases; it peaks near 1.3–1.5 R_{exp} for N_2, where the π bonds are partly broken and several determinants have comparable weights, and then decreases again; for H_2O it jumps abruptly beyond 1.3 R_{exp} to 40–50 mE_h, because all optimisations, including the forward and backward continuation sweeps, remained trapped in a higher-energy solution; and for CO_2 it decreases at long distances because the π active space contributes little correlation there, so that CASCI approaches RHF and the target state becomes nearly a single determinant. These patterns show that a fixed, shallow hardware-efficient circuit does not deliver uniform accuracy along a reaction coordinate, which is a more demanding test than a single equilibrium calculation.

The comparison with all-electron FCI in Fig. 3 also shows the limits of the minimal active spaces at stretched geometries. For CO_2, the (4e, 4o) π active space captures little of the correlation that develops when both C = O σ bonds are stretched, so the CASCI and VQE curves remain close to RHF at large distances while FCI lies well below. Larger active spaces that include the σ and σ^* orbitals would be required for a qualitatively correct description of bond breaking in CO_2, H_2O, NH_3 and CH_4; with the parity mapping such spaces need 10–14 qubits.

Table 3
VQE errors relative to CASCI along the potential energy curves (16 geometries per molecule, R/R_{exp} = 0.75–2.00), in mE_h. NPE: non-parallelity error. N_{CA}: number of geometries within chemical accuracy (1.6 mE_h).

Molecule	Mean ΔE_{VQE}	Max ΔE_{VQE}	NPE	N_{CA} / 16
H_2	0.000	0.000	0.000	16
HeH^+	0.000	0.000	0.000	16
LiH	0.000	0.000	0.000	16
BeH_2	0.000	0.000	0.000	16
H_2O	12.333	51.067	49.776	7
NH_3	4.880	13.685	11.552	0
CH_4	0.949	4.127	3.939	13
CO_2	3.788	9.267	9.265	4
N_2	4.607	13.197	11.967	1

4.4 Equilibrium bond lengths and vibrational frequencies

Table 4 compares the equilibrium bond lengths obtained from the spline minima with experiment. Because all methods share the STO-3G basis, the dominant deviation from experiment is the basis-set error, which is visible already at the FCI level: $|r_e(\mathrm{FCI}) - r_e(\mathrm{exp})|$ ranges from 0.007 Å to 0.139 Å (HeH^+). The minimal basis lengthens the bonds of the first-row hydrides and of N_2 and CO_2, and gives a much too long bond for HeH^+ because the He 1s function is poorly described. These deviations are properties of the Hamiltonian, not of the quantum algorithm, and must not be attributed to VQE. The relevant measure of VQE performance is the deviation from the CASCI bond length obtained with the same Hamiltonian: this is below 10^{-4} Å for the two- and four-qubit systems and at most 0.002 Å for the six-qubit systems, where the residual comes from the geometry-dependent optimisation error discussed above.

For the diatomic molecules, harmonic frequencies were obtained from the curvature of the spline at its minimum (Table 4). VQE reproduces the CASCI frequencies of H_2, HeH^+ and LiH exactly, and the N_2 frequency to within 40 cm^{-1}. As expected for a minimal basis, the FCI/STO-3G frequencies differ substantially from experiment (e.g. 5001 versus 4401 cm^{-1} for H_2 and 2096 versus 2359 cm^{-1} for N_2) [46].

Table 4
Equilibrium bond lengths r_e (Å) from the minima of the computed STO-3G curves, compared with experiment, and harmonic frequencies ω_e (cm^{-1}) for the diatomic molecules. For polyatomic molecules r_e refers to the symmetric-stretch cut with experimental bond angles.

Molecule	r_e RHF	r_e FCI	r_e CASCI	r_e VQE	r_e exp.	ω_e FCI	ω_e VQE	ω_e exp.
H_2	0.712	0.735	0.735	0.735	0.7414	5001	5001	4401.2
HeH^+	0.930	0.914	0.914	0.914	0.7743	2740	2740	–
LiH	1.511	1.548	1.550	1.550	1.5957	1682	1677	1405.6
BeH_2	1.291	1.316	1.297	1.297	1.3264	–	–	–
H_2O	0.987	1.024	1.024	1.024	0.9572	–	–	–
NH_3	1.029	1.058	1.039	1.038	1.0120	–	–	–
CH_4	1.083	1.108	1.088	1.088	1.0870	–	–	–
CO_2	1.188	1.234	1.188	1.189	1.1600	–	–	–
N_2	1.134	1.194	1.178	1.175	1.0977	2096	2270	2358.6

Experimental values from refs. [46, 47].

4.5 Effect of circuit depth and optimisation

To test whether the six-qubit errors are caused by insufficient circuit depth, the equilibrium calculations for H_2O and N_2 were repeated with reps = 1–6 (Fig. 4b; 11 optimisations per depth). The errors with respect to CASCI are 19.92, 15.71, 1.41, 0.90, 0.03, 0.00 mE_h for H_2O and 66.00, 47.15, 9.92, 1.19, 0.00, 0.00 mE_h for N_2 for reps = 1, 2, …, 6. The error decreases steeply between reps = 2 and reps = 4, chemical accuracy is reached at reps = 4 (60 parameters, 20 CNOTs) for both molecules, and the CASCI energy is recovered essentially exactly with reps = 5–6 (72–84 parameters, 25–30 CNOTs). Even the deepest of these circuits uses about 16 times fewer CNOT gates than UCCSD. The reps = 3 circuit adopted for the six-qubit systems is therefore at the edge of sufficient expressibility, which explains both its residual errors and their sensitivity to the starting point: along the potential energy curves, a second (backward) continuation sweep lowered several six-qubit energies by up to tens of mE_h at stretched geometries, showing that individual optimisations can become trapped in local minima. These observations are consistent with earlier analyses of the expressibility–trainability trade-off of hardware-efficient circuits [21–23] and argue for deeper or symmetry-preserving and adaptive ansätze [20, 49] when chemical accuracy is required beyond four qubits.

4.6 Limitations and outlook

The present results are noiseless, infinite-shot statevector simulations and therefore represent the best-case performance of the algorithm for the chosen ansatz and optimiser. On present hardware, gate errors, decoherence, readout errors and finite sampling would add further errors that typically exceed chemical accuracy without error mitigation [12, 50, 51]. A shot budget sufficient to resolve 1.6 mE_h for Hamiltonians with 85–165 Pauli terms is also substantial [26, 48]. Second, the STO-3G basis and small active spaces limit the physical accuracy of all methods, including exact diagonalisation, as shown by the bond lengths and frequencies in Table 4. Third, only symmetric-stretch cuts of the polyatomic potential energy surfaces were computed. Natural extensions are noisy simulations with error

mitigation, measurement grouping, symmetry-preserving ansätze, larger basis sets and active spaces, and execution on quantum hardware.

5. Conclusions

We have benchmarked a hardware-efficient VQE (R_Y–R_Z rotations with linear CNOT entanglement) for nine small molecules using 2–6 qubits, with all reference energies computed consistently in the same STO-3G basis. Four conclusions follow. (i) For two- and four-qubit problems (H_2, HeH^+, LiH, BeH_2) the ansatz is exact: energies, potential energy curves, bond lengths and vibrational frequencies agree with the exact active-space results to within numerical precision. (ii) For six-qubit problems (H_2O, NH_3, CH_4, CO_2, N_2) the same ansatz with reps = 3 gives equilibrium errors of 0.3–7.5 mE_h and maximum errors along the curves of up to 51.1 mE_h; H_2O and CH_4 reach chemical accuracy at equilibrium; increasing the depth to reps = 4–6 (20–30 CNOTs) recovers chemical accuracy at equilibrium, and UCCSD reaches chemical accuracy for all molecules but requires 476 CNOT gates for the six-qubit systems. (iii) For all molecules except H_2 and HeH^+, the active-space truncation error is larger than the VQE error, so VQE energies must be compared with CASCI in the same active space rather than with all-electron FCI. (iv) Deviations of STO-3G bond lengths and frequencies from experiment are basis-set effects shared by FCI, and should not be interpreted as errors of the quantum algorithm. Reporting absolute errors in mE_h against exact references for the same Hamiltonian, together with full resource counts, provides a transparent basis for judging progress in near-term quantum chemistry.

Declarations

Conflict of interest:

The author declares no competing interests.

Ethics approval

, **consent to participate and consent to publish**: Not applicable.

Funding:

This research received no specific grant from any funding agency.

Author Contribution

N.S.B. conceived and designed the study, performed all quantum chemical and variational quantum eigensolver calculations, wrote the analysis code, analysed and interpreted the data, prepared all figures and tables, and wrote and revised the manuscript. N.S.B. read and approved the final manuscript.

Data and code availability:

All input geometries, active-space definitions, Python scripts used to generate the Hamiltonians and run the VQE, CASCI and FCI calculations, and the raw energies for every geometry are provided as Supplementary Material and will be deposited in a public repository (e.g. Zenodo) upon acceptance.

References

1. Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* 1982, 21, 467–488.
2. Aspuru-Guzik, A.; Dutoi, A. D.; Love, P. J.; Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* 2005, 309, 1704–1707.
3. McArdle, S.; Endo, S.; Aspuru-Guzik, A.; Benjamin, S. C.; Yuan, X. Quantum computational chemistry. *Rev. Mod. Phys.* 2020, 92, 015003.
4. Cao, Y.; Romero, J.; Olson, J. P.; Degroote, M.; Johnson, P. D.; Kieferová, M.; Kivlichan, I. D.; Menke, T.; Peropadre, B.; Sawaya, N. P. D.; Sim, S.; Veis, L.; Aspuru-Guzik, A. Quantum chemistry in the age of quantum computing. *Chem. Rev.* 2019, 119, 10856–10915.
5. Bauer, B.; Bravyi, S.; Motta, M.; Chan, G. K.-L. Quantum algorithms for quantum chemistry and quantum materials science. *Chem. Rev.* 2020, 120, 12685–12717.
6. Reiher, M.; Wiebe, N.; Svore, K. M.; Wecker, D.; Troyer, M. Elucidating reaction mechanisms on quantum computers. *Proc. Natl. Acad. Sci. USA* 2017, 114, 7555–7560.
7. Lee, S.; Lee, J.; Zhai, H.; Tong, Y.; Dalzell, A. M.; Kumar, A.; Helms, P.; Gray, J.; Cui, Z.-H.; Liu, W.; Kastoryano, M.; Babbush, R.; Preskill, J.; Reichman, D. R.; Campbell, E. T.; Valeev, E. F.; Lin, L.; Chan, G. K.-L. Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nat. Commun.* 2023, 14, 1952.
8. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* 2018, 2, 79.
9. 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.
10. McClean, J. R.; Romero, J.; Babbush, R.; Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* 2016, 18, 023023.
11. O'Malley, P. J. J.; Babbush, R.; Kivlichan, I. D.; Romero, J.; McClean, J. R.; Barends, R.; Kelly, J.; Roushan, P.; Tranter, A.; Ding, N.; et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* 2016, 6, 031007.
12. 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.
13. Hempel, C.; Maier, C.; Romero, J.; McClean, J.; Monz, T.; Shen, H.; Jurcevic, P.; Lanyon, B. P.; Love, P.; Babbush, R.; Aspuru-Guzik, A.; Blatt, R.; Roos, C. F. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X* 2018, 8, 031022.
14. Nam, Y.; Chen, J.-S.; Pisenti, 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 Inf.* 2020, 6, 33.
15. Google AI Quantum and Collaborators; Arute, F.; Arya, K.; Babbush, R.; et al. Hartree–Fock on a superconducting qubit quantum computer. *Science* 2020, 369, 1084–1089.
16. Cerezo, M.; Arrasmith, A.; Babbush, R.; Benjamin, S. C.; Endo, S.; Fujii, K.; McClean, J. R.; Mitarai, K.; Yuan, X.; Cincio, L.; Coles, P. J. Variational quantum algorithms. *Nat. Rev. Phys.* 2021, 3, 625–644.
17. Bharti, K.; Cervera-Lierta, A.; Kyaw, T. H.; Haug, T.; Alperin-Lea, S.; Anand, A.; Degroote, M.; Heimonen, H.; Kottmann, J. S.; Menke, T.; Mok, W.-K.; Sim, S.; Kwek, L.-C.; Aspuru-Guzik, A. Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.* 2022, 94, 015004.
18. Tilly, J.; Chen, H.; Cao, S.; Picozzi, D.; Setia, K.; Li, Y.; Grant, E.; Wossnig, L.; Rungger, I.; Booth, G. H.; Tennyson, J. The variational quantum eigensolver: a review of methods and best practices. *Phys. Rep.* 2022, 986, 1–128.
19. 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 Sci. Technol.* 2019, 4, 014008.

20. 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.
21. Sim, S.; Johnson, P. D.; Aspuru-Guzik, A. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Adv. Quantum Technol.* 2019, 2, 1900070.
22. 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.
23. Larocca, M.; Thanasilp, S.; Wang, S.; Sharma, K.; Biamonte, J.; Coles, P. J.; Cincio, L.; McClean, J. R.; Holmes, Z.; Cerezo, M. Barren plateaus in variational quantum computing. *Nat. Rev. Phys.* 2025, 7, 174–189.
24. Ryabinkin, I. G.; Genin, S. N.; Izmaylov, A. F. Constrained variational quantum eigensolver: quantum computer search engine in the Fock space. *J. Chem. Theory Comput.* 2019, 15, 249–255.
25. Kuroiwa, K.; Nakagawa, Y. O. Penalty methods for a variational quantum eigensolver. *Phys. Rev. Res.* 2021, 3, 013197.
26. Kühn, M.; Zanker, S.; Deglmann, P.; Marthaler, M.; Weiß, H. Accuracy and resource estimations for quantum chemistry on a near-term quantum computer. *J. Chem. Theory Comput.* 2019, 15, 4764–4780.
27. Szabo, A.; Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*; Dover: Mineola, NY, 1996.
28. Helgaker, T.; Jørgensen, P.; Olsen, J. *Molecular Electronic-Structure Theory*; Wiley: Chichester, 2000.
29. Hehre, W. J.; Stewart, R. F.; Pople, J. A. Self-consistent molecular-orbital methods. I. Use of Gaussian expansions of Slater-type atomic orbitals. *J. Chem. Phys.* 1969, 51, 2657–2664.
30. Knowles, P. J.; Handy, N. C. A new determinant-based full configuration interaction method. *Chem. Phys. Lett.* 1984, 111, 315–321.
31. Jordan, P.; Wigner, E. Über das Paulische Äquivalenzverbot. *Z. Phys.* 1928, 47, 631–651.
32. Bravyi, S. B.; Kitaev, A. Y. Fermionic quantum computation. *Ann. Phys.* 2002, 298, 210–226.
33. Seeley, J. T.; Richard, M. J.; Love, P. J. The Bravyi–Kitaev transformation for quantum computation of electronic structure. *J. Chem. Phys.* 2012, 137, 224109.
34. Bravyi, S.; Gambetta, J. M.; Mezzacapo, A.; Temme, K. Tapering off qubits to simulate fermionic Hamiltonians. arXiv:1701.08213, 2017.
35. Mitarai, K.; Negoro, M.; Kitagawa, M.; Fujii, K. Quantum circuit learning. *Phys. Rev. A* 2018, 98, 032309.
36. Schuld, M.; Bergholm, V.; Gogolin, C.; Izaac, J.; Killoran, N. Evaluating analytic gradients on quantum hardware. *Phys. Rev. A* 2019, 99, 032331.
37. 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.; Wouters, S.; Chan, G. K.-L. PySCF: the Python-based simulations of chemistry framework. *WIREs Comput. Mol. Sci.* 2018, 8, e1340.
38. Sun, Q.; Zhang, X.; Banerjee, S.; Bao, P.; Barbry, M.; Blunt, N. S.; Bogdanov, N. A.; Booth, G. H.; Chen, J.; Cui, Z.-H.; et al. Recent developments in the PySCF program package. *J. Chem. Phys.* 2020, 153, 024109.
39. Javadi-Abhari, A.; Treinish, M.; Krsulich, K.; Wood, C. J.; Lishman, J.; Gacon, J.; Martiel, S.; Nation, P. D.; Bishop, L. S.; Cross, A. W.; Johnson, B. R.; Gambetta, J. M. Quantum computing with Qiskit. arXiv:2405.08810, 2024.
40. Qiskit Nature developers and contributors. *Qiskit Nature*, version 0.8.0; open-source software, https://github.com/qiskit-community/qiskit-nature (accessed 2026).
41. Qiskit Algorithms developers and contributors. *Qiskit Algorithms*, version 0.4.0; open-source software, https://github.com/qiskit-community/qiskit-algorithms (accessed 2026).
42. Harris, C. R.; Millman, K. J.; van der Walt, S. J.; et al. Array programming with NumPy. *Nature* 2020, 585, 357–362.

43. Kraft, D. *A Software Package for Sequential Quadratic Programming*; Technical Report DFVLR-FB 88 – 28; DLR German Aerospace Center: Köln, 1988.
44. Virtanen, P.; Gommers, R.; Oliphant, T. E.; et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* 2020, 17, 261–272.
45. Hunter, J. D. Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* 2007, 9, 90–95.
46. Huber, K. P.; Herzberg, G. *Molecular Spectra and Molecular Structure IV: Constants of Diatomic Molecules*; Van Nostrand Reinhold: New York, 1979.
47. Johnson, R. D., III (Ed.). *NIST Computational Chemistry Comparison and Benchmark Database*, NIST Standard Reference Database 101; National Institute of Standards and Technology: Gaithersburg, MD, https://cccbdb.nist.gov (accessed 2026).
48. Wecker, D.; Hastings, M. B.; Troyer, M. Progress towards practical quantum variational algorithms. *Phys. Rev. A* 2015, 92, 042303.
49. 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 Inf.* 2020, 6, 10.
50. Temme, K.; Bravyi, S.; Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* 2017, 119, 180509.
51. 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.

Figures

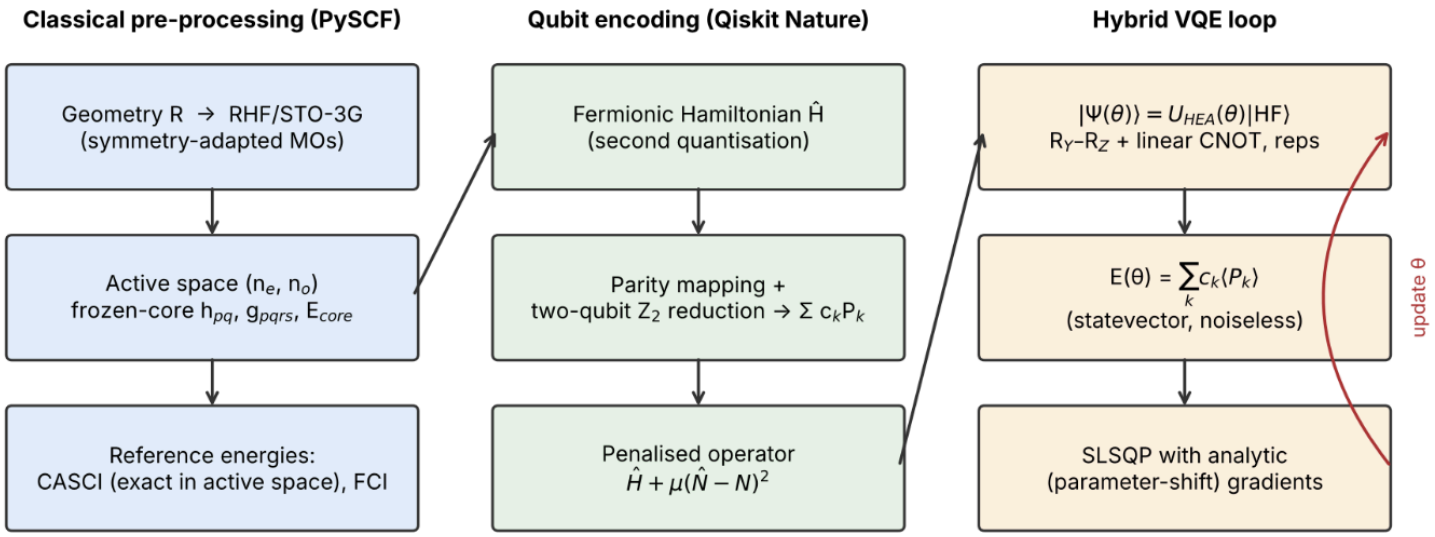


Figure 1

Computational workflow. Classical pre-processing (PySCF) provides the molecular orbitals, active-space integrals and exact reference energies; Qiskit Nature builds the qubit Hamiltonian with the parity mapping and two-qubit reduction; the hybrid VQE loop minimises the penalised energy with SLSQP using analytic gradients.

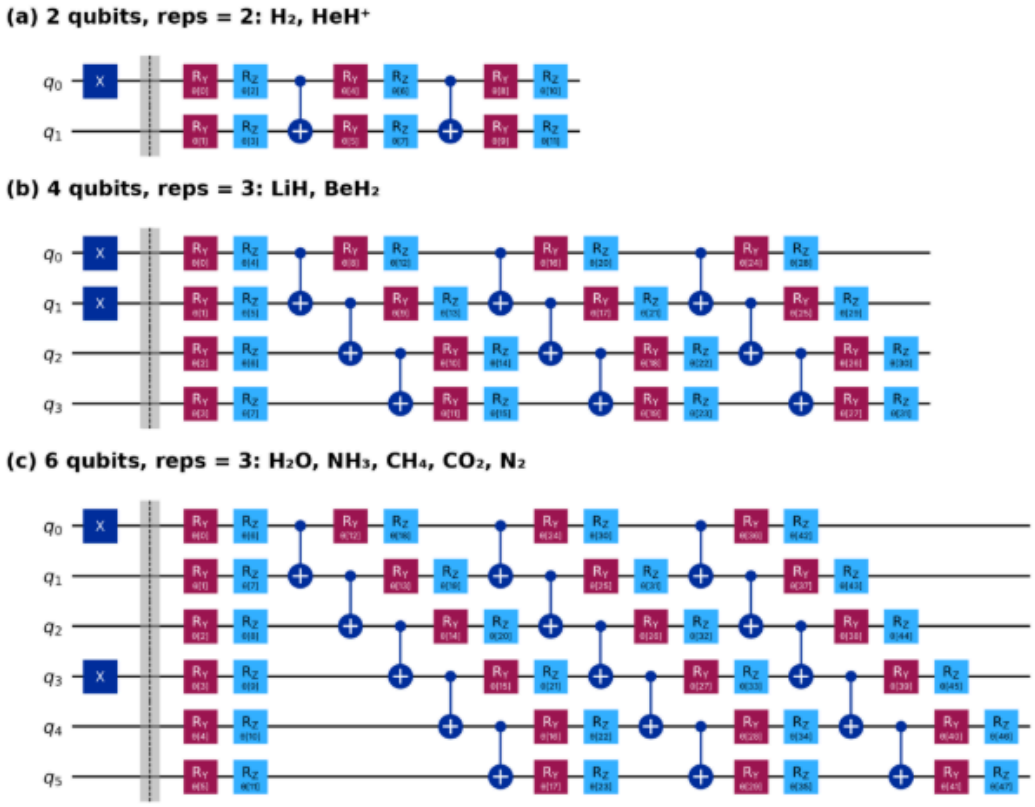


Figure 2

Hardware-efficient ansatz circuits used in this work, generated with Qiskit. The X gates prepare the Hartree–Fock reference in the parity encoding; each repetition consists of R_Y and R_Z rotations on every qubit followed by a linear chain of CNOT gates, and a final rotation layer closes the circuit. (a) 2 qubits, reps = 2 (12 parameters, 2 CNOTs); (b) 4 qubits, reps = 3 (32 parameters, 9 CNOTs); (c) 6 qubits, reps = 3 (48 parameters, 15 CNOTs).

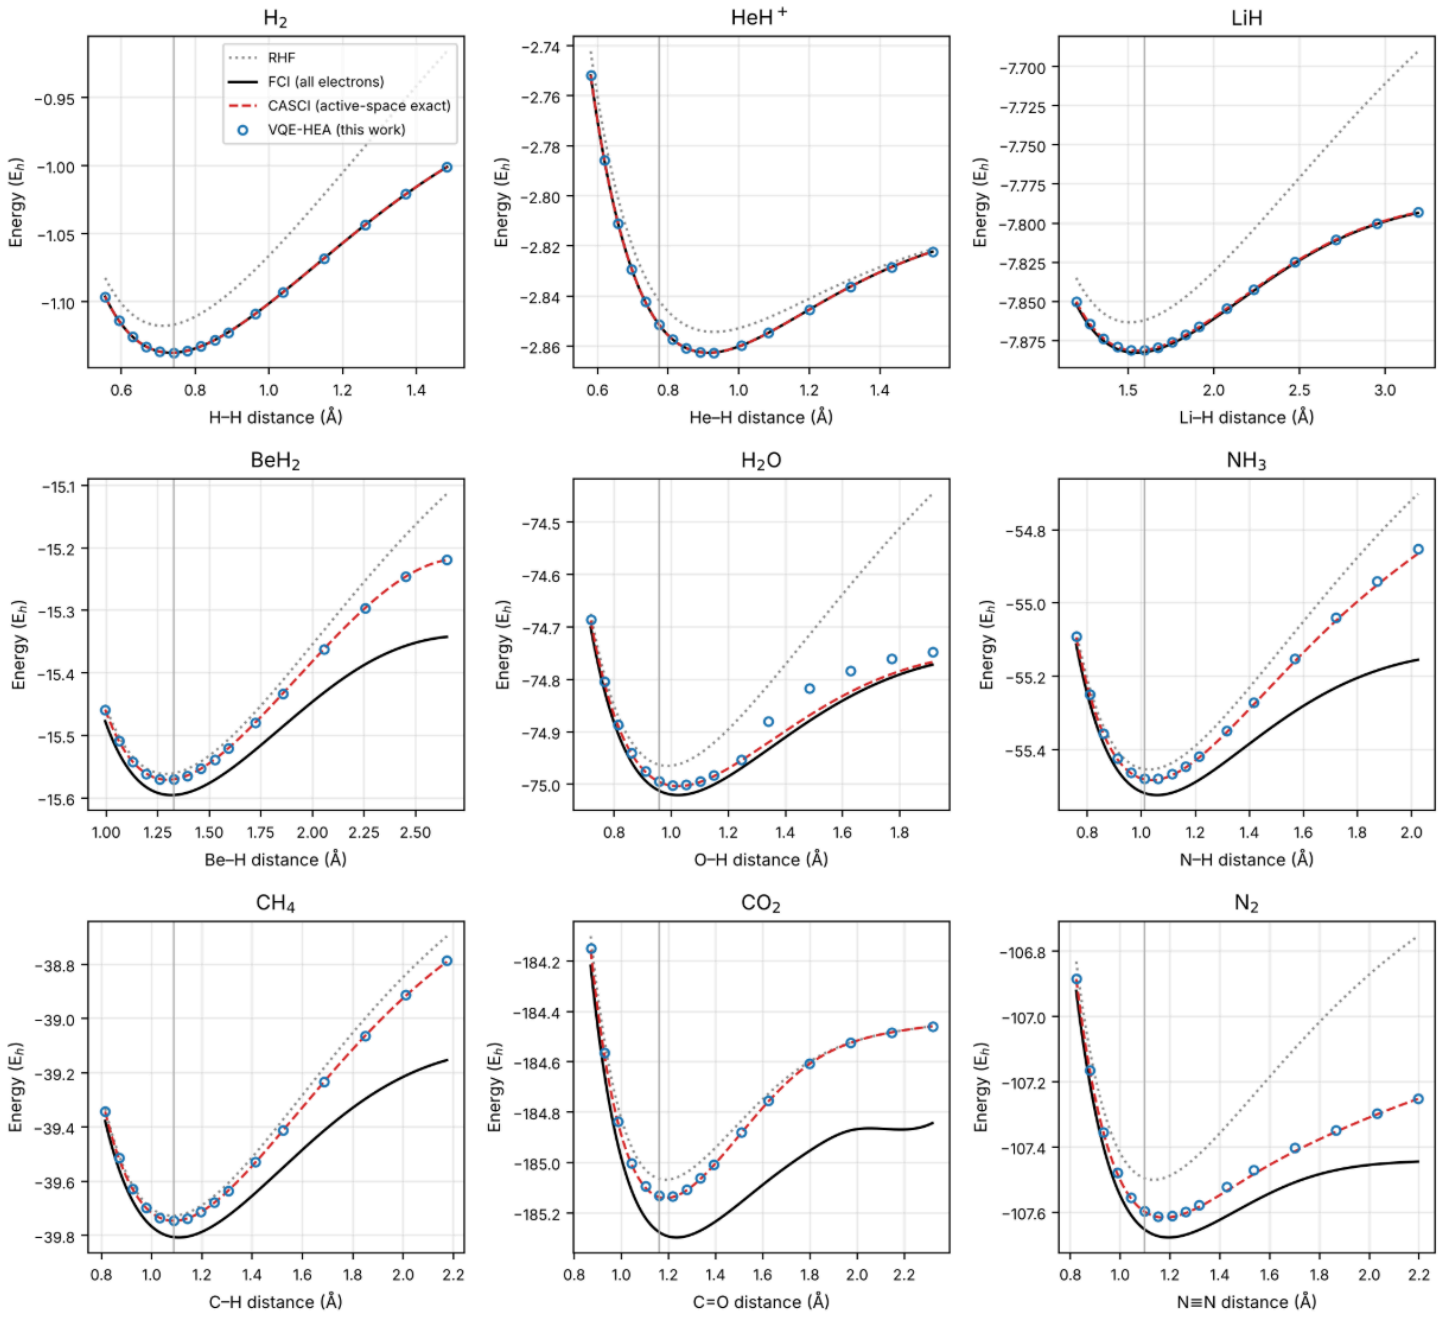


Figure 3

Potential energy curves in the STO-3G basis: RHF (dotted grey), all-electron FCI (solid black), active-space CASCI (dashed red) and VQE with the hardware-efficient ansatz (open circles, this work). For the polyatomic molecules the coordinate is the symmetric stretch with bond angles fixed at their experimental values. The vertical line marks the experimental equilibrium distance. FCI and CASCI denote the lowest $M_S = 0$ eigenstate; at the longest C=O distances several electronic states become nearly degenerate, which produces the flattening of the CO_2 FCI curve.

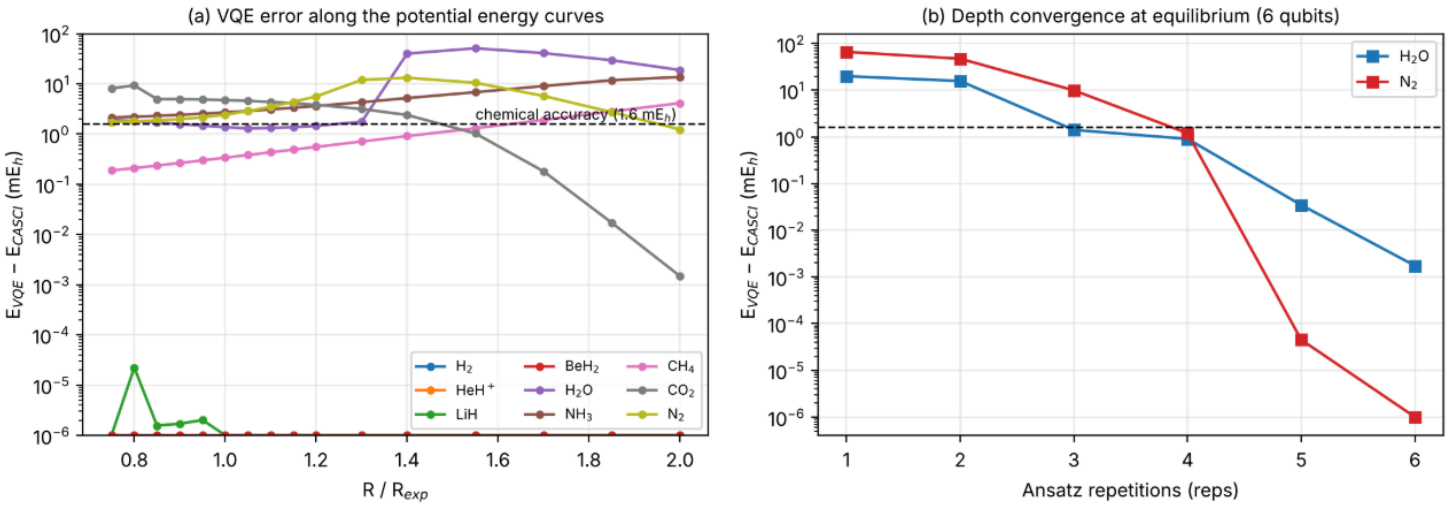


Figure 4

(a) VQE error with respect to CASCI along the potential energy curves (logarithmic scale; values below 10^{-6} mE_h are drawn at 10^{-6}). The dashed line marks chemical accuracy. (b) Dependence of the equilibrium VQE error on the number of ansatz repetitions for H_2O and N_2 (11 optimisations per depth, lowest energy shown).