

Resource-Efficient Simulation of Molecular Ground and Excited States Based on Contextual Subspace and Quantum Subspace Expansion

Chao Liu¹ and Yuxin Deng^{*2, 1}

¹Shanghai Key Laboratory of Trustworthy Computing, East China Normal University, Shanghai 200062, China

²MoE Key Laboratory of Interdisciplinary Research of Computation and Economics, Shanghai University of Finance and Economics, Shanghai 200433, China

*Email: yxdeng@msg.sufe.edu.cn

Abstract

Predicting spectra and photochemical pathways relies on the computation of molecular excited states. However, the practicality of near-term variational quantum algorithms is threatened by the prohibitive growth of measurement overheads. Integrating Quantum Subspace Expansion (QSE) with the Contextual Subspace (CS) method, we propose a CS-QSE framework to resolve the scaling bottlenecks in excited-state energy estimation. By confining excitation operators to a compact contextual subspace, the framework removes redundant degrees of freedom and reduces Pauli string counts, thereby easing the measurement burden. The primary strength of CS-QSE lies in its substantial reduction of the operator-pool size. Benchmarking on LiH, HF, H₂O, and HCl shows that CS-QSE achieves errors within the target tolerance of 1.6×10^{-3} Ha relative to full configuration interaction benchmarks in the same basis set while mitigating the prohibitive scaling inherent in the full-space method. Numerical simulations reveal that the operator pool size is consistently pruned by over 90% across all systems, with the reduction reaching as high as 99.7% for molecules such as HCl. This CS-QSE framework establishes a resource-efficient route for molecular simulations tailored to near-term noisy intermediate-scale quantum devices.

1 Introduction

Electronic structure simulation is central to understanding the properties of molecules and materials, underpinning advancements across drug de-

sign [1], materials development [2], and catalysis [3]. While classical simulation is increasingly hampered by steep scaling overheads, quantum computing provides a viable route to addressing problems that remain computationally prohibitive for conventional methods [4, 5]. Consequently, researchers have increasingly focused on hybrid quantum-classical algorithms to bridge the gap between theoretical potential and near-term hardware capabilities [6, 7, 8]. As a representative of this approach, the Variational Quantum Eigensolver (VQE) has become one of the most widely studied approaches for quantum chemistry in the noisy intermediate-scale quantum (NISQ) era [9, 10, 11]. To address hardware limitations, various resource-efficient techniques have been developed, including adaptive ansatz construction methods for variational circuit optimization [12, 13, 14] and quantum embedding and fragmentation approaches for reducing the effective problem size [15, 16, 17].

Among these strategies, the contextual subspace (CS) approach provides an effective pathway for resource-efficient VQE calculations near the conventional chemical-precision threshold relative to full configuration interaction (FCI) benchmarks in the same basis set. Pioneered by Kirby et al. [18, 19, 20], this approach decomposes a Hamiltonian into a classically solvable non-contextual component and a quantum contextual correction, significantly reducing qubit resources and measurement costs. To address the challenge of efficient ansatz construction, Weaving et al. [21] introduced a stabilizer-based non-contextual projection ansatz, while Ralli et al. [22] further minimized measurement overhead via unitary partitioning. For hardware implementation,

Liang et al. [23] integrated parameterized pulse control into the framework, effectively reducing circuit depth and complexity. Recent experiments on superconducting processors demonstrate that this framework can obtain ground-state energies close to FCI results in strongly correlated regimes, outperforming benchmarked single-reference methods and remaining competitive with selected multiconfigurational approaches [24]. While the CS-VQE has been successfully applied to calculate ground-state energy, excited states are equally essential for photochemical dynamics and optoelectronic properties, yet their efficient computation on NISQ hardware remains substantially more challenging than ground-state simulations.

Existing techniques for excited states, such as Variational Quantum Deflation (VQD), rely on enforcing orthogonality via overlap penalty terms [25]. Although recently generalized within the CS-VQD framework [26], VQD is susceptible to sequential error accumulation, where inaccuracies in lower-lying states systematically bias higher energy levels, and its requirement for serial optimization imposes a prohibitive computational burden [27].

To address these limitations, Quantum Subspace Expansion (QSE), originally proposed by McClean et al. [28] for calculating molecular excited states and mitigating hardware decoherence, offers a robust alternative. By mapping the electronic structure problem onto a generalized eigenvalue equation, QSE effectively bypasses the iterative state-by-state optimization. Nevertheless, standard QSE formulations incur prohibitive measurement overheads [29, 30], rendering its direct application to large-scale molecular systems computationally intractable on contemporary NISQ architectures [31].

Technical Challenge. Despite the promise of hybrid architectures, there are still critical challenges to the efficient simulation of excited states within resource-constrained environments. Compared with ground-state VQE, variational excited-state algorithms typically incur substantially higher resource costs [25, 32]. In the VQD, enforcing orthogonality through sequential state-by-state deflation leads to a complicated optimization landscape and a large number of circuit evaluations [25, 27]. In the traditional QSE, one must measure a large set of Hamiltonian and overlap matrix elements [28], resulting in a quantum measurement overhead that is much greater than for ground-state VQE and grows rapidly with system size, which severely limits accurate spectral calculations on near-term noisy devices [29, 30].

Contributions. We propose a hybrid CS-QSE framework to overcome the above challenges. The

scaling bottlenecks are mitigated by integrating the Quantum Subspace Expansion within a contextual subspace, and the excited-state simulation is possible with much reduced qubits. More specifically, our contributions are twofold:

- *A resource-efficient CS-QSE framework.* We generalize the contextual subspace approach to the excited-state computation by projecting a full-space excitation-operator pool into the contextual subspace, thereby substantially reducing the dimension of the operator space. This compression alleviates both the measurement overhead and qubit requirements, enabling excitation-energy estimates with errors approaching the chemical-precision scale relative to FCI benchmarks for molecular systems such as H_2O and HCl , where standard QSE-based schemes already exhibit rapidly increasing measurement requirements.
- *An algebraically consistent projection strategy for symmetry preservation.* We address the key challenge of operator definition in reduced space by defining the excitation pool in the full Hilbert space and applying a uniform projection to the excitation operators, Hamiltonian operators, and UCCSD generators. This consistent mapping enforces algebraic equivalence across all operator spaces, effectively removing symmetry contamination from unphysical sectors, eliminating spurious eigenpairs, and stabilizing the generalized eigenvalue problem.

The rest of the article is structured as follows. Section 2 briefly reviews the theoretical foundations of the Variational Quantum Eigensolver, Contextual Subspace, and Quantum Subspace Expansion algorithms. Section 3 details the proposed hybrid CS-QSE framework. Section 4 presents and evaluates numerical simulations for representative molecular systems. Section 5 summarizes the study.

2 Theoretical Background

2.1 Variational Quantum Eigensolver

The VQE approximates the ground-state energy E_0 by iteratively optimizing a parameterized trial wavefunction $|\psi(\theta)\rangle$ [9, 10]. Based on the variational principle, the expectation value serves as an upper bound to the ground-state energy, $E_0 \leq \langle \psi(\theta) | \hat{H} | \psi(\theta) \rangle$. The trial state is prepared via a parameterized unitary ansatz $U(\theta)$ acting on an initial reference state, typically $|\psi(\theta)\rangle = U(\theta) |0\rangle^{\otimes N}$.

To evaluate the energy on quantum hardware, the

Hamiltonian is decomposed into a linear combination of Pauli strings, $\hat{P}_a \in \{I, X, Y, Z\}^{\otimes N}$:

$$\hat{H} = \sum_a w_a \hat{P}_a, \quad (1)$$

where w_a is a real-valued scalar coefficient.

Consequently, the VQE optimization problem is reduced to minimizing a cost function $E_{\text{VQE}}(\theta)$ composed of measurable expectation values,

$$\begin{aligned} E_{\text{VQE}} &= \min_{\theta} E_{\text{VQE}}(\theta) \\ &= \min_{\theta} \sum_a w_a \langle 0|^{\otimes N} U^\dagger(\theta) \hat{P}_a U(\theta) |0\rangle^{\otimes N}. \end{aligned} \quad (2)$$

A classical optimizer updates the set of parameters θ to minimize the cost function $E_{\text{VQE}}(\theta)$, thereby providing a variational upper-bound estimate of the ground-state energy.

2.2 Contextual Subspace Method

The contextual subspace (CS) method is a hybrid quantum-classical strategy for approximating the eigenenergies of a molecular Hamiltonian [18, 19]. The theoretical foundation derives from the Kochen-Specker theorem [33, 34]. The central idea is to partition the qubit Hamiltonian $\hat{H}$ into a noncontextual part $\hat{H}_{\text{nc}}$, which is solved efficiently on a classical computer, and a contextual part $\hat{H}_{\text{c}}$ that captures the remaining quantum correlations and is treated on a quantum device within the contextual subspace framework [21]:

$$\hat{H} = \hat{H}_{\text{nc}} + \hat{H}_{\text{c}}, \quad (3)$$

The construction of $\hat{H}_{\text{nc}}$ is governed by closure under inference and a rigorous algebraic consistency criterion. For a set of Pauli operators $\mathcal{P}$, let $\overline{\mathcal{P}}$ denote its closure, the minimal set containing $\mathcal{P}$ such that for any commuting pair $A, B \in \overline{\mathcal{P}}$, their product AB is also in $\overline{\mathcal{P}}$. A set is defined as noncontextual if there exists a globally consistent value assignment $v(P) \in \{\pm 1\}$ for every $P \in \overline{\mathcal{P}}$. This assignment must satisfy multiplicativity, whereby the value of a product operator equals the product of the values assigned to each individual component. It must also guarantee that each operator is assigned a unique, context-independent value, remaining invariant across all measurement contexts. For instance, the set $\mathcal{P} = \{X_1 X_2, Z_1 Z_2\}$ is noncontextual; its closure $\{X_1 X_2, Z_1 Z_2, -Y_1 Y_2, I\}$ admits a consistent assignment, such as $v(X_1 X_2) = 1$ and $v(Z_1 Z_2) = 1$, which uniquely determines $v(-Y_1 Y_2) = 1$.

Table 1: Verification of quantum contextuality via algebraic constraints in the Mermin-Peres square.

	Operator 1	Operator 2	Operator 3	Product
Row 1	XI	IX	XX	$+\mathbb{I}$
Row 2	IZ	ZI	ZZ	$+\mathbb{I}$
Row 3	XZ	ZX	YY	$+\mathbb{I}$
Product	$+\mathbb{I}$	$+\mathbb{I}$	$-\mathbb{I}$	Contradiction

In contrast, contextuality signifies that no such global assignment can be maintained across all commuting contexts. This is illustrated by the Mermin-Peres magic square in Table 1 [35]. The Mermin-Peres square arranges nine two-qubit Pauli observables selected from the products of the four Pauli generators $\mathcal{S}_{\text{MP}} = \{XI, IX, ZI, IZ\}$. As shown in Table 1, while the operators in each row and the first two columns satisfy a product identity of $+\mathbb{I}$, the third column yields a product of $-\mathbb{I}$. These incompatible parity constraints preclude any globally consistent, context-independent assignment of ± 1 values to the nine observables, thus confirming the presence of quantum contextuality.

To avoid such contradictions and ensure classical simulability, one must restrict attention to operator sets that are closed under inference and admit globally consistent value assignments, such as those used in constructing $\hat{H}_{\text{nc}}$ [20].

Within the stabilizer formalism, the noncontextual component $\hat{H}_{\text{nc}}$ is associated with a set of M independent, pairwise commuting Pauli operators $\mathcal{S}_{\text{stab}} = \{S_1, S_2, \dots, S_M\}$. These generators commute with all terms in $\hat{H}_{\text{nc}}$ and represent emergent $\mathbb{Z}_2$ pseudosymmetries of the system. By classically solving the noncontextual Hamiltonian, one determines the optimal eigenvalue assignment $\nu_j \in \{+1, -1\}$ for these generators, which defines a specific stabilizer sector. Let $|\psi\rangle$ denote the noncontextual reference state obtained from the classical solution of $\hat{H}_{\text{nc}}$, which lies in this sector and therefore satisfies

$$S_j |\psi\rangle = \nu_j |\psi\rangle, \quad \forall j \in \{1, 2, \dots, M\}. \quad (4)$$

To project the problem onto this sector, a Clifford unitary U_S is employed to simultaneously diagonalize the M generators onto single-qubit Z operators acting on the designated stabilizer qubit set $I_{\text{stab}} = \{q_1, q_2, \dots, q_M\}$, such that:

$$U_S S_j U_S^\dagger = Z_{q_j}, \quad q_j \in I_{\text{stab}}, \quad (5)$$

where q_j denotes the index of the j -th stabilizer qubit. The existence of the Clifford unitary U_S follows from the stabilizer normal form: any independent set of mutually commuting Pauli generators whose generated stabilizer group does not contain $-\mathbb{I}$ can be mapped by a Clifford transformation to single-qubit

Z operators [36]. In the selected sector, the eigenvalue signs are absorbed into the stabilizer generators, ensuring that the associated stabilizer group is consistent. The corresponding CS-VQE construction is described in Ref. [21]. For the eigenvalue vector $\nu = (\nu_1, \nu_2, \dots, \nu_M)$, the sector projection operator P_ν acting on the qubits indexed by I_{stab} is defined as

$$P_\nu = \prod_{j=1}^M \frac{\mathbb{I} + \nu_j Z_{q_j}}{2}. \quad (6)$$

For any operator A , the projection map $\pi_\nu^u(A)$ yields an effective operator acting on the reduced simulation space of $u = N - M$ qubits,

$$\pi_\nu^u(A) = \text{Tr}_{\text{stab}} \left(P_\nu U_S A U_S^\dagger P_\nu \right), \quad (7)$$

where Tr_{stab} denotes the partial trace over the M stabilizer qubits. The resulting effective Hamiltonian $\hat{H}_{\text{CS}}$ acting on the u -qubit contextual subspace is synthesized by applying the projection map π_ν^u to the entire Hamiltonian,

$$\hat{H}_{\text{CS}} = \pi_\nu^u(\hat{H}_{\text{nc}} + \hat{H}_c) = E_{\text{nc}} \cdot \mathbb{I} + \pi_\nu^u(\hat{H}'_c), \quad (8)$$

where E_{nc} is the noncontextual ground-state energy contribution within the chosen sector, $\mathbb{I}$ denotes the identity operator acting on the u -qubit contextual subspace, and $\hat{H}'_c$ denotes the collection of terms not fixed to constants by the stabilizer-sector assignment, including the projected contextual terms and any residual diagonal Pauli- Z strings from $\hat{H}_{\text{nc}}$. This formulation incorporates the noncontextual energy as a constant shift while collecting all remaining non-trivial terms into $\hat{H}'_c$. The same projection map π_ν^u is applied to variational ansatzes and excitation operator pools in the CS-VQE and CS-QSE stages to ensure that the quantum dynamics remain strictly within the u -qubit contextual subspace [21].

2.3 Quantum Subspace Expansion

QSE is a hybrid quantum-classical method that approximates molecular excited-state energies by diagonalizing the electronic Hamiltonian within a subspace spanned by configurations generated from an approximate ground-state wavefunction $|\Psi_0\rangle$ [28]. In the widely used linear-response (LR) approximation, this subspace is spanned by basis states of the form $\hat{c}_k^\dagger \hat{c}_j |\Psi_0\rangle$. For a single-determinant Hartree-Fock reference, distinct single excitations correspond to orthogonal Slater determinants. However, when $|\Psi_0\rangle$ is a correlated multi-determinant VQE reference, the corresponding operator-generated states

are generally non-orthogonal, and their overlap matrix elements can be expressed in terms of the one- and two-body reduced-density matrices of the reference [32, 37]. The framework can be systematically extended to include double and higher-order excitations, and as the maximum excitation rank approaches the total number of electrons, the resulting expansion can, in principle, approach the corresponding FCI energies within the chosen basis set.

Since the subspace basis states are generally non-orthogonal, the QSE reduces the energy problem to a generalized eigenvalue equation. Writing the trial state $|\Phi\rangle = \sum_\mu C_\mu |\Psi_\mu\rangle$ as a linear combination of w subspace basis vectors $\{|\Psi_\mu\rangle\}_{\mu=0}^{w-1}$, and applying the Rayleigh-Ritz variational principle with respect to the coefficient C_μ one has

$$\sum_\lambda \langle \Psi_\nu | \hat{H} | \Psi_\lambda \rangle C_\lambda = E \sum_\lambda \langle \Psi_\nu | \Psi_\lambda \rangle C_\lambda. \quad (9)$$

By defining the subspace Hamiltonian and the overlap matrix as

$$\begin{aligned} (\mathbf{H}_{\text{sub}})_{\nu\lambda} &= \langle \Psi_\nu | \hat{H} | \Psi_\lambda \rangle, \\ (\mathbf{S}_{\text{sub}})_{\nu\lambda} &= \langle \Psi_\nu | \Psi_\lambda \rangle, \end{aligned} \quad (10)$$

the resulting generalized eigenvalue problem can be written compactly as

$$\mathbf{H}_{\text{sub}} \mathbf{C} = E \mathbf{S}_{\text{sub}} \mathbf{C}, \quad (11)$$

where $\mathbf{C}$ is the matrix of expansion coefficients for the ground and the excited states within the subspace. The equation can be solved on a classical computer to obtain approximate ground- and excited-state energies and the corresponding wavefunction coefficients. Since QSE and related quantum equation-of-motion approaches solve the eigenvalue problem within an operator-generated response subspace, the accuracy of the resulting excited-state estimates depends on the overlap between this subspace and the target excited states [28, 32].

3 Method

High-precision quantum chemistry simulations on the NISQ devices remain constrained by qubit scarcity and noise-induced limits on circuit depth [38]. Variational quantum algorithms, such as the VQE, provide a practical route for near-term electronic-structure calculations [9, 10]. The coherent resources and Hamiltonian measurement overhead required by the VQE scale rapidly with the orbital space size, restricting applications to small molecules or simplified active spaces [12, 39].




Figure 1: Hybrid quantum-classical workflow for calculating ground- and low-lying excited-state energies.

To mitigate these resource bottlenecks, we present a hybrid quantum-classical workflow designed to determine ground-state energies and low-lying excitations with reduced quantum requirements. Implemented within the Symmer framework [40], the workflow sequentially applies the symmetry-based tapering and the contextual subspace projection to compress the problem into a reduced effective subspace. Within this reduced effective subspace, QSE extracts the excited-state information from a single variational reference state.

The overall framework is illustrated in Fig. 1. The workflow begins with lossless compression via qubit tapering, exploiting the $\mathbb{Z}_2$ Pauli symmetries that commute with the molecular Hamiltonian to obtain a reduced Hamiltonian together with consistently tapered ansatz and excitation operators [41, 42]. We then apply CS-VQE to prepare a variational ground state in the reduced space and use projected QSE to estimate the ground and low-lying excited states from the same reference state.

3.1 Tapering Process in the CS-QSE Framework

Qubit tapering is employed as the first preprocessing step in the CS-QSE to reduce qubit overhead exactly within a selected symmetry sector. The method exploits an Abelian subgroup $\mathcal{S}$ of Pauli symmetries generated by $\{s_i\}_{i=1}^k$ such that $[\hat{H}, s_i] = 0$ for all i . Using the stabilizer formalism, one can construct

a Clifford unitary U_C that brings these commuting generators to a standard form, e.g., $U_C s_i U_C^\dagger = \pm Z_{q_i}$, thereby localizing the $\mathbb{Z}_2$ symmetries onto k designated qubits.

By defining the conjugated Hamiltonian $\hat{H}_{\text{taper}} = U_C \hat{H} U_C^\dagger$, where $\hat{H}$ acts on n qubits, we have $[\hat{H}_{\text{taper}}, Z_{q_i}] = 0$ on k symmetry qubits.

Consequently, every Pauli term appearing in $\hat{H}_{\text{taper}}$ acts on those qubits only as I or Z . Fixing each sector eigenvalue $v_i \in \{+1, -1\}$ in the symmetry sector of the Hartree-Fock reference state, we replace Z_{q_i} with v_i and eliminate the k symmetry qubits, obtaining an effective Hamiltonian on $n_{\text{taper}} = n - k$ qubits that preserves the spectrum restricted to that sector.

In the QSE, the subspace is generated from a set of M operators $\mathcal{O}_{\text{full}} = \{O_j\}_{j=0}^{M-1}$ acting on a reference state, and the projected problem is solved via the generalized eigenvalue equation

$$\mathbf{H}_{\text{sub}} \mathbf{C} = E \mathbf{S}_{\text{sub}} \mathbf{C}, \quad (12)$$

where $[\mathbf{H}_{\text{sub}}]_{ij} = \langle \psi | O_i^\dagger \hat{H} O_j | \psi \rangle$ and $[\mathbf{S}_{\text{sub}}]_{ij} = \langle \psi | O_i^\dagger O_j | \psi \rangle$.

If the Hamiltonian is tapered within a selected symmetry sector while the QSE expansion operators are constructed directly as Pauli operators in the reduced space, these Pauli operators are not guaranteed to correspond to the reduced forms of the original electronic excitation operators within the same symmetry sector. Consequently, they may generate QSE basis states that are inconsistent with the electronic

excitation structure of the original Hilbert space, because the corresponding Pauli operators are not obtained through the same symmetry-reduction map used for the Hamiltonian and ansatz generators.

To address this, we introduce a Congruent Operator Projection (COP) protocol. Rather than constructing excitation operators heuristically in the reduced space, the COP starts from the full-space fermionic excitation pool $\mathcal{O}_{\text{full}}$ as defined in Sec. 3.3 and derives all QSE expansion operators via the same symmetry-reduction map $\mathcal{T}_\nu$ used for the Hamiltonian. The map $\mathcal{T}_\nu$ effectively implements the same reduction used for the Hamiltonian: each operator is conjugated by the tapering Clifford unitary U_C , restricted to the selected symmetry sector, and acts in the reduced qubit space:

$$\mathcal{O}_{\text{taper}} = \left\{ \tilde{O}_j \mid \tilde{O}_j = \mathcal{T}_\nu(O_j) \equiv P_\nu U_C O_j U_C^\dagger P_\nu, \right. \\ \left. O_j \in \mathcal{O}_{\text{full}} \right\}, \quad (13)$$

where P_ν denotes the projector onto the selected symmetry sector. Once the symmetry sector is specified, the corresponding Pauli factors are replaced by their sector eigenvalues, so that the associated qubits are tapered off and the operators act on the reduced qubit space. Moreover, we apply the same symmetry-reduction map $\mathcal{T}_\nu$ to the UCCSD generators. This map-consistent construction ensures that the Hamiltonian, reference state, ansatz generators, and QSE expansion operators preserve the same tapered symmetry sector, avoiding symmetry-breaking terms and improving the conditioning and numerical stability of the QSE generalized eigenvalue problem.

3.2 CS-VQE process in the CS-QSE Framework

Although qubit tapering reduces the Hilbert space dimension of the Hamiltonian, the resulting tapered Hamiltonian may still contain many Pauli terms, leading to a substantial measurement bottleneck. To further reduce the quantum resource requirements, the CS method applies a stabilizer-subspace projection to the tapered Hamiltonian $\hat{H}_{\text{taper}}$ acting on n_{taper} qubits, yielding a reduced Hamiltonian $\hat{H}_{\text{CS}}$ defined on a compact contextual subspace involving only u simulation qubits, where $u = n_{\text{taper}} - m$, and m is the number of qubits fixed by the selected stabilizers. Within this framework, the Hamiltonian is partitioned into a classically solvable noncontextual part $\hat{H}_{nc}$ and a contextual part $\hat{H}_c$ that must be treated on a quantum processor.

In the standard CS procedure, the stabilizers are determined from the noncontextual structure of

$\hat{H}_{nc}$ [20, 21]; however, a stabilizer set selected only from this structure may project out operator components that are important for describing electronic correlation. To address the corresponding ansatz-construction problem in CS-VQE, previous studies used a UCCSD-compatible stabilizer selection, in which the selected stabilizers are chosen to commute with the UCCSD generators before the ansatz operators are projected into the contextual subspace. In this construction, UCCSD terms that are incompatible with the selected stabilizer sector vanish under projection, so the stabilizer choice directly determines which parts of the UCCSD operator pool are retained in the reduced problem. Building on this UCCSD-compatible stabilizer selection, the same stabilizer-subspace projection used for the Hamiltonian and UCCSD ansatz is applied to the tapered QSE excitation-operator pool, yielding projected excitation operators defined in the same contextual subspace. Excitation operators whose projected forms vanish are discarded, so the retained pool consists only of operators compatible with the selected stabilizer sector.

As a result, the projected excitation operators, the Hamiltonian terms, and UCCSD generators are defined within the same contextual subspace, which helps avoid symmetry-sector inconsistencies and rank-deficient QSE overlap matrices. After the stabilizer set and the corresponding projection map are determined, the tapered Hamiltonian and the tapered QSE excitation-operator pool are projected into the same contextual subspace. The original pool of M_{exc} tapered excitation operators is reduced to a projected pool of M_u operators by applying the same CS projection map π_ν^u to each operator and retaining only the non-vanishing projected operators:

$$\{\tilde{O}_k^u\}_{k=1}^{M_u} = \left\{ \pi_\nu^u(O_l^{\text{taper}}) \mid l \in \{1, 2, \dots, M_{\text{exc}}\}, \right. \\ \left. \pi_\nu^u(O_l^{\text{taper}}) \neq 0 \right\}, \quad (14)$$

Here, $M_u \leq M_{\text{exc}}$ denotes the number of non-vanishing excitation operators after the CS projection.

3.3 Constructing the CS-QSE Excitation Operator Pool

With the tapering map $\mathcal{T}_\nu$ and the contextual projection map π_ν^u established, we construct the CS-QSE excitation-operator pool by applying the same projection maps used for the Hamiltonian and the UCCSD generators. The pool is first generated in the full fermionic space from single and double exci-

tations and is then reduced by the same tapering and contextual-subspace maps.

For a system with n_{so} spin-orbitals, let i, j denote the occupied spin-orbitals and a, b denote the virtual spin-orbitals with respect to the chosen reference determinant. The single and double excitation operators are respectively defined as

$$T_{ai}^{(1)} = a_a^\dagger a_i, \quad a \in \text{virt}, i \in \text{occ}, \quad (15)$$

$$T_{abij}^{(2)} = a_a^\dagger a_b^\dagger a_j a_i, \quad a, b \in \text{virt}, i, j \in \text{occ}. \quad (16)$$

Accordingly, the full excitation-operator pool $\mathcal{O}_{\text{full}}$ is defined as the union of the single-excitation set $\mathcal{O}^{(1)}$ and the double-excitation set $\mathcal{O}^{(2)}$. To eliminate the redundant double-excitation terms, an upper-triangular constraint such that $a > b$ and $i > j$ is adopted. All fermionic operators in $\mathcal{O}_{\text{full}}$ are then mapped to Pauli strings via the Jordan–Wigner transformation.

To ensure symmetry-sector consistency, the exact tapering map described in Section 3.1 is applied consistently to both the Hamiltonian and the excitation-operator pool, yielding the tapered pool $\mathcal{O}_{\text{taper}}$ in the selected physical sector.

For a target simulation register of size u , a stabilizer sector is chosen according to the contextual-subspace construction described in Section 2.2. Using the Clifford unitary U_S defined in Eq. (5) and the sector projector P_ν defined in Eq. (6), we apply the contextual-subspace projection map π_ν^u defined in Eq. (7).

To construct the CS-QSE operator pool, we apply the contextual-subspace projection map π_ν^u to each excitation operator in the tapered excitation-operator pool $\mathcal{O}_{\text{taper}}$. Operators that vanish after projection are discarded, and the remaining projected operators form the effective excitation-operator pool defined in Eq. (14). Consequently, each CS-QSE basis state is constructed within the contextual subspace associated with the chosen stabilizer sector, substantially reducing the expansion subspace dimension and the corresponding measurement overheads.

The dimensionality reduction in the CS-QSE framework is mathematically established by applying the projection map to each Pauli operator in the tapered pool. Many excitation terms without satisfying the stabilizer constraints vanish upon projection, which is critical since the QSE computational costs scale with the dimension of the subspace. If an operator A anti-commutes with a stabilizer generator S_k , the resulting projection $\pi_\nu^u(A)$ defined in Eq. (7) is zero. For the transformed operators $A' = U_S A U_S^\dagger$ and $S'_k = U_S S_k U_S^\dagger$, the unitary transformation U_S

preserves the commutation or anti-commutation relations of the original operators, yielding $\{A', S'_k\} = 0$. According to the diagonalization in Eq. (5), $S'_k = Z_{q_k}$, thus the local Pauli operator of A' on qubit q_k must be X or Y to anti-commute with Z .

After the Clifford transformation, the sector projector P_ν is a product of projectors on stabilizer qubits as defined in Eq. (6). This configuration results in a factor $P_{\nu_k} \sigma_{\text{anti}} P_{\nu_k}$ on qubit q_k . Due to the anti-commutation $\{\sigma_{\text{anti}}, Z_{q_k}\} = 0$, the operator σ_{anti} flips the projector to the orthogonal sector via $\sigma_{\text{anti}} P_{\nu_k} = P_{-\nu_k} \sigma_{\text{anti}}$. The inherent orthogonality $P_{\nu_k} P_{-\nu_k} = 0$ ensures that $P_{\nu_k} \sigma_{\text{anti}} P_{\nu_k} = 0$, yielding a total projection of $\pi_\nu^u(A) = \text{Tr}_{I_{\text{stab}}}(P_\nu A' P_\nu) = 0$.

Consequently, even though the initial operator pool $\mathcal{O}_{\text{taper}}$ is obtained by tapering the full pool of single and double excitation operators, a significant fraction of these operators are automatically eliminated for violating stabilizer-sector constraints.

3.4 Ground- and Excited-State Calculations in the CS-QSE Framework

The CS-VQE reference state is obtained by minimizing the expectation value of the contextual-subspace Hamiltonian $\hat{H}_{\text{CS}}$:

$$E_0^u = \min_\theta \langle \Psi^u(\theta) | \hat{H}_{\text{CS}} | \Psi^u(\theta) \rangle, \quad (17)$$

leading to an optimal parameter set θ^* and the variational reference state $|\Psi_{\text{ref}}^u\rangle = |\Psi^u(\theta^*)\rangle$. After obtaining the optimized reduced-space reference state $|\Psi_{\text{ref}}^u\rangle$, the linear-response subspace is constructed by generating a non-orthogonal basis set $\mathcal{B} = \{|\phi_1\rangle, |\phi_2\rangle, \dots, |\phi_{M_u}\rangle\}$ from the symmetry-consistent projected operator pool $\tilde{\mathcal{O}}_{\text{CS}}^u$ defined in Eq. (14), where

$$|\phi_k\rangle = \tilde{\mathcal{O}}_k^u |\Psi_{\text{ref}}^u\rangle, \quad k \in \{1, 2, \dots, M_u\}. \quad (18)$$

Within this subspace, the Schrödinger equation is projected onto the following generalized eigenvalue problem:

$$\mathbf{H}_{\text{sub}} \mathbf{C} = E \mathbf{S}_{\text{sub}} \mathbf{C}, \quad (19)$$

where the elements of the projected Hamiltonian matrix $\mathbf{H}_{\text{sub}}$ and the overlap matrix $\mathbf{S}_{\text{sub}}$ are defined as:

$$\begin{aligned} [\mathbf{H}_{\text{sub}}]_{ij} &= \langle \Psi_{\text{ref}}^u | \tilde{\mathcal{O}}_i^\dagger \hat{H}_{\text{CS}} \tilde{\mathcal{O}}_j | \Psi_{\text{ref}}^u \rangle, \\ [\mathbf{S}_{\text{sub}}]_{ij} &= \langle \Psi_{\text{ref}}^u | \tilde{\mathcal{O}}_i^\dagger \tilde{\mathcal{O}}_j | \Psi_{\text{ref}}^u \rangle. \end{aligned} \quad (20)$$

The projection of the full operator pool into the compact Contextual Subspace may introduce linear dependencies, which can render the overlap matrix $\mathbf{S}_{\text{sub}}$ ill-conditioned (i.e., condition number $\kappa(\mathbf{S}_{\text{sub}}) \gg 1$).

The Canonical Orthogonalization procedure is employed to improve numerical stability. The spectral decomposition $\mathbf{S}_{\text{sub}} = U\Sigma U^\dagger$ is performed, with a truncation threshold ϵ applied to discard eigenvectors associated with near-zero eigenvalues ($\Sigma_{ii} \leq \epsilon$). A transformation matrix $X = U_{\text{valid}}\Sigma_{\text{valid}}^{-1/2}$ is then constructed to convert the generalized eigenvalue problem into a better-conditioned standard eigenvalue problem:

$$\mathbf{H}_{\text{reg}}\mathbf{C}_{\text{reg}} = \mathbf{C}_{\text{reg}}\mathbf{E}, \quad (21)$$

where $\mathbf{H}_{\text{reg}} = X^\dagger\mathbf{H}_{\text{sub}}X$. The resulting set of eigenvalues $\{E_0, E_1, \dots, E_k\}$ constitutes the estimated energy spectrum, where the lowest eigenvalue E_0 provides a linear-response correction to the ground state energy, while a higher-lying eigenvalue E_k (with $k > 0$) resolves the low-energy excitation spectrum, enabling the simultaneous calculation of multiple excited states.

4 Numerical Results

This section presents the numerical simulations for the CS-QSE workflow described in Sec. 3. Using the FCI energies within the STO-3G basis set as references, we benchmark the accuracy of ground and low-lying excited states obtained from a single CS-VQE reference state followed by projected QSE. The results are organized as follows: (i) We analyze energy-error trends versus the contextual-subspace size u at fixed molecular geometries; (ii) We evaluate the state tracking along bond-length scans through potential-energy curves; (iii) We quantify the reduction in the number of excitation operators induced by the $\mathbb{Z}_2$ symmetry tapering and the contextual-subspace projection, and compare the resulting resource profile with CS-VQD.

4.1 CS-QSE Using UCCSD Ansatz

In this work, the Hamiltonians of LiH, HF, H₂O, and HCl are constructed using the PySCF program [43] to perform the initial electronic structure calculation and generate molecular integrals. These integrals are then processed via the Symmer framework and Qiskit library, integrated with the OpenFermion program [40, 44, 45]. We use the FCI result within the STO-3G basis set as the benchmark and report the absolute error relative to that FCI reference. Following common usage in the quantum-computing literature, we take 1.6×10^{-3} Ha as the threshold relative to this STO-3G/FCI reference; however, this criterion is more precisely understood as the chemical-precision

threshold relative to a fixed basis set [46], representing the algorithmic convergence to the FCI baseline rather than an absolute deviation from the experimental limit.

Using a symmetry-based qubit-tapering approach, the qubit counts of these molecular systems are reduced from their original values of 12, 12, 14 and 20 to 8, 8, 10, and 17 qubits, respectively. Considering the substantial computational expense of large-scale simulations, the simulation scale is limited to 10 qubits as long as it shows the performance of the CS-QSE. Among all proposed VQE algorithms, variational wave function is constructed using the UCCSD ansatz. The excitation operator $\hat{T}_{UCCSD}$ comprises single-excitation and double-excitation terms,

$$\begin{aligned} \hat{T}_{UCCSD} &= \hat{T}_1 + \hat{T}_2 \\ &= \sum_{i \in \text{occ}} \sum_{a \in \text{virt}} t_i^a a_a^\dagger a_i \\ &\quad + \sum_{i > j \in \text{occ}} \sum_{a > b \in \text{virt}} t_{ij}^{ab} a_a^\dagger a_b^\dagger a_j a_i. \end{aligned} \quad (22)$$

The corresponding unitary operator is applied to the reference state $|\psi_{\text{ref}}\rangle$ to build the parameterized quantum circuit. Within the Stabilizer Formalism, the parameterized ansatz is restricted from the full electronic structure problem to the contextual subspace. As illustrated in Fig. 2, all energy errors for both ground and excited states are plotted versus the number of qubits retained in the contextual subspace at fixed molecular geometries.

For the diatomic systems (LiH, HF, and HCl), the bond lengths are set to 1.57473 Å, 0.9945 Å, and 1.3413 Å, respectively, while for H₂O we adopt the equilibrium H–O–H bond angle and O–H bond length of 104.48° and 0.9578 Å, respectively [47, 48]. At fixed geometries, energy errors decline consistently with the increasing number of qubits retained in the contextual subspace. Notably, energy errors for both ground and excited states in the CS-QSE exhibit a sharp order-of-magnitude reduction at a specific qubit threshold, followed by rapid convergence toward the accuracy limit. When the simulated qubit count N_{sim} reaches the tapered qubit limit, the Hamiltonian projection transitions from the approximate truncation of the CS-VQE to an exact mapping. This shift eliminates the truncation errors inherent in the CS-VQE method, allowing the subspace to retain the core operators of the excitation pool and UCCSD. Hence, we observe a sudden discontinuity in the energy error at this point. We found that LiH, HF, and H₂O all show a clear jump in accuracy, but HCl does not show this within the tested subspace range. This is mainly because the tapered Hamiltonian of HCl




Figure 2: Energy error scaling of the CS-VQE and CS-QSE compared with the CS-VQD using a UCCSD ansatz as a function of the number of qubits.

contains 17 qubits, while the present calculations are limited to $u = 10$; therefore, the tested range remains below the full tapered-Hamiltonian dimension. Furthermore, the CS-QSE method typically needs more qubits for excited states than for the ground state to reach the chemical-precision threshold. This is because accurate results for excited states depend on including the key orbital degrees of freedom.

The CS-VQD data in Fig. 2 are restricted to the first excited state, as the algorithm necessitates a sequential optimization with overlap penalty terms for each successive level. This serial approach incurs prohibitive measurement and execution overhead as the number of target states grows, which often undermines numerical stability. In contrast, the CS-QSE resolves multiple excited states simultaneously via the classical post-processing. By solving the subspace generalized eigenvalue problem in a single execution, the CS-QSE eliminates the need for repetitive, high-cost variational searches for each target state. Furthermore, CS-QSE avoids the repeated overlap measurements required by the sequential CS-VQD procedure and obtains multiple excited

states through diagonalization of the projected subspace matrices, thereby reducing the measurement overhead for excited-state calculation and improving its practicality for near-term quantum simulations.

To validate the extension of the QSE within the CS-VQE framework, we assessed its capability to track ground and excited state potential energy curves across a range of nuclear geometries in Figs. 3 and 4. For each molecular system, potential energy surface (PES) scans are performed over a bond length range of $R = 0.7\text{--}3.5$ Å.

To assess the impact of qubit scaling on precision, the PES scans are performed at two simulation-qubit sizes, with $u = 6$ and $u = 8$ for LiH and HF, and $u = 8$ and $u = 10$ for H₂O and HCl. A chemical-precision threshold of 1.6×10^{-3} Ha relative to the exact FCI values for the ground and low-lying excited states is attainable at the small-scale qubit systems. However, the pursuit of higher precision through increased the qubit counts significantly inflates the computational costs due to the escalating demands of circuit depth and measurement overhead. Therefore, minimizing the qubit requirements




Figure 3: Energy error scaling of the CS-VQE and CS-QSE using a UCCSD ansatz as a function of 6 qubits and 8 qubits.

is of higher practical significance within the resource-constrained NISQ regime.

As illustrated in Figs. 3 and 4, the $E_0 \sim E_4$ energies obtained via the CS-QSE align closely with the FCI benchmarks across the entire bond-length range, forming smooth and continuous PES profiles. This demonstrates that the CS-QSE method can accurately evaluate the ground-state energy while reliably capturing the relative energy relations and tracking multiple low-lying excited-state curves during configuration scans. Note that the FCI benchmarks are plotted as smooth curves using fifth-order polynomial fitting for visual clarity, whereas the CS-QSE results are presented as discrete sampling points. The ‘CS-VQE’ markers represent the ground-state energies obtained by the CS-VQE method at each bond length, which serve as the reference for constructing the CS-QSE subspace.

The symbol (T) denotes that the residual has reached the predefined numerical threshold. Tables 2–5 summarize the absolute energy errors calculated relative to the FCI benchmarks. The tables detail the ground-state reference errors from

the CS-VQE and the $E_0 \sim E_4$ errors obtained via the CS-QSE. Numerical results are presented in the form of $a(b)$, representing the n -qubit and N -qubit calculations, respectively. Notably, for the 8-qubit (LiH/HF) and 10-qubit (H_2O) cases, the simulation qubit count matches the tapered Hamiltonian size.

As a result, the CS-VQE procedure is not performed, and the reference state is generated through the VQE on the tapered Hamiltonian, thereby avoiding the additional approximation introduced by the CS projection. Across the scanned bond-length range, the calculated energies for both the ground and the first few excited states consistently reach the chemical-precision threshold when tightly converged. In certain configurations, the excited-state errors are comparable to or even lower than those of the ground state, indicating that the CS-QSE subspace expansion effectively captures the underlying physics of the target states relative to the reference.

In summary, despite the gains in accuracy, the precision and the computational cost should be traded off subtly with the associated increases in measurement overhead and subspace dimension. Our ap-




Figure 4: Energy error scaling of the CS-VQE and CS-QSE using a UCCSD ansatz as a function of 8 qubits and 10 qubits.

Table 2: Energy errors of LiH versus the number of qubits.

Bond length (Å)	6 qubits (8 qubits)					
	CS-VQE GS	CS-QSE E_0	CS-QSE E_1	CS-QSE E_2	CS-QSE E_3	CS-QSE E_4
0.700	1.05e-3 (1.01e-7)	1.05e-3 (9.77e-15)	1.53e-3 (9.77e-15)	1.72e-3 (T)	5.89e-3 (T)	6.69e-4 (1.95e-14)
0.875	6.34e-4 (1.74e-8)	6.34e-4 (8.97e-14)	7.82e-4 (5.06e-14)	9.23e-4 (6.04e-14)	6.54e-4 (1.09e-13)	4.05e-3 (3.02e-14)
1.050	3.81e-4 (7.63e-9)	3.81e-4 (1.07e-14)	5.55e-4 (4.00e-14)	6.40e-4 (5.06e-14)	5.14e-4 (1.07e-14)	2.68e-3 (9.77e-15)
1.225	2.59e-4 (1.35e-8)	2.59e-4 (9.77e-15)	5.45e-4 (T)	5.95e-4 (1.07e-14)	4.02e-4 (9.77e-15)	1.95e-3 (9.77e-15)
1.400	2.23e-4 (5.78e-8)	2.23e-4 (3.02e-14)	6.06e-4 (T)	6.42e-4 (1.07e-14)	3.48e-4 (3.02e-14)	1.55e-3 (3.02e-14)
1.575	3.45e-4 (7.86e-8)	3.45e-4 (5.33e-14)	6.15e-4 (1.78e-14)	7.53e-4 (3.55e-15)	3.44e-4 (1.07e-14)	1.09e-3 (1.95e-14)
1.750	4.03e-4 (9.62e-8)	4.03e-4 (3.55e-15)	6.24e-4 (2.84e-14)	8.22e-4 (1.07e-14)	3.30e-4 (4.97e-14)	7.76e-4 (4.97e-14)
1.925	4.38e-4 (1.11e-7)	4.38e-4 (3.55e-15)	6.33e-4 (2.84e-14)	8.71e-4 (1.78e-14)	3.16e-4 (5.68e-14)	5.56e-4 (2.13e-14)
2.100	4.58e-4 (1.24e-7)	4.58e-4 (5.33e-14)	6.41e-4 (3.55e-14)	9.10e-4 (2.13e-14)	3.05e-4 (5.33e-14)	4.18e-4 (3.55e-14)
2.275	4.70e-4 (1.35e-7)	4.70e-4 (4.97e-14)	6.46e-4 (6.04e-14)	9.41e-4 (3.55e-14)	2.96e-4 (2.13e-14)	3.19e-4 (3.20e-14)
2.450	4.77e-4 (1.43e-7)	4.77e-4 (4.97e-14)	6.49e-4 (4.97e-14)	9.65e-4 (4.97e-14)	2.90e-4 (4.97e-14)	2.53e-4 (T)
2.625	4.81e-4 (1.47e-7)	4.81e-4 (5.33e-14)	6.51e-4 (4.97e-14)	9.83e-4 (3.55e-14)	2.86e-4 (4.97e-14)	2.11e-4 (3.55e-14)
2.800	4.84e-4 (1.49e-7)	4.84e-4 (5.33e-14)	6.52e-4 (3.55e-14)	9.97e-4 (3.55e-14)	2.84e-4 (4.97e-14)	1.83e-4 (4.97e-14)
2.975	4.85e-4 (1.50e-7)	4.85e-4 (5.33e-14)	6.52e-4 (3.55e-14)	1.01e-3 (3.55e-14)	2.84e-4 (3.55e-14)	1.64e-4 (T)
3.150	4.85e-4 (1.50e-7)	4.85e-4 (5.33e-14)	6.52e-4 (3.55e-14)	1.01e-3 (3.55e-14)	2.86e-4 (4.97e-14)	1.52e-4 (4.97e-14)
3.325	4.85e-4 (1.50e-7)	4.85e-4 (5.33e-14)	6.52e-4 (3.55e-14)	1.02e-3 (3.55e-14)	2.89e-4 (3.55e-14)	1.42e-4 (3.55e-14)
3.500	4.85e-4 (1.50e-7)	4.85e-4 (5.33e-14)	6.51e-4 (3.55e-14)	1.03e-3 (3.55e-14)	2.93e-4 (3.55e-14)	1.35e-4 (4.97e-14)

proach reconciles chemical precision with minimal qubit overhead, maintaining the feasibility of the algorithm within the constraints of NISQ hardware. In the VQE calculations, the parameters are initialized to zero at the first scanned bond length, corresponding to the Hartree–Fock reference state. For each subsequent scanned bond length, the converged pa-

rameters from the preceding point are used as the initial guess for optimization. This warm-start strategy accelerates VQE convergence and facilitates robust state tracking, resulting in physically consistent and continuous potential energy profiles across the scanned coordinate space.

Table 3: Energy errors of HF versus the number of qubits.

Bond length (Å)	6 qubits (8 qubits)					
	CS-VQE GS	CS-QSE E_0	CS-QSE E_1	CS-QSE E_2	CS-QSE E_3	CS-QSE E_4
0.700	2.32e-5 (9.63e-12)	2.32e-5 (1.56e-13)	4.35e-5 (7.11e-14)	1.28e-4 (2.70e-13)	4.59e-5 (1.85e-13)	1.24e-4 (3.13e-13)
0.875	2.12e-5 (9.63e-12)	2.12e-5 (1.42e-13)	2.78e-5 (2.84e-13)	1.06e-4 (1.14e-13)	7.11e-14 (8.53e-14)	1.38e-4 (2.13e-13)
1.050	1.93e-5 (9.71e-12)	1.93e-5 (6.39e-13)	1.64e-5 (5.83e-13)	7.89e-5 (5.68e-13)	2.84e-14 (1.56e-13)	1.31e-4 (1.14e-13)
1.225	1.67e-5 (9.75e-12)	1.67e-5 (3.69e-13)	9.51e-6 (3.69e-13)	5.55e-5 (6.82e-13)	1.14e-13 (1.28e-13)	9.87e-5 (9.95e-14)
1.400	1.33e-5 (9.32e-12)	1.33e-5 (4.12e-13)	5.52e-6 (2.42e-13)	3.68e-5 (5.40e-13)	3.69e-13 (4.69e-13)	4.80e-5 (3.13e-13)
1.575	9.70e-6 (9.56e-12)	9.70e-6 (9.95e-14)	3.18e-6 (6.25e-13)	2.26e-5 (3.41e-13)	5.68e-14 (1.14e-13)	1.50e-5 (2.42e-13)
1.750	6.42e-6 (9.35e-12)	6.42e-6 (2.84e-14)	1.82e-6 (8.53e-14)	1.27e-5 (5.68e-14)	5.58e-6 (1.85e-13)	9.95e-14 (1.56e-13)
1.925	3.93e-6 (9.79e-12)	3.93e-6 (T)	1.04e-6 (3.98e-13)	6.58e-6 (1.42e-14)	3.84e-6 (1.85e-13)	1.56e-13 (1.42e-14)
2.100	2.27e-6 (1.01e-11)	2.27e-6 (3.41e-13)	6.02e-7 (1.14e-13)	3.25e-6 (2.98e-13)	3.21e-6 (4.26e-14)	9.95e-14 (T)
2.275	1.27e-6 (1.17e-11)	1.27e-6 (5.68e-14)	3.52e-7 (3.13e-13)	1.64e-6 (8.24e-13)	2.53e-6 (1.99e-13)	1.28e-13 (7.11e-14)
2.450	6.95e-7 (1.13e-11)	6.95e-7 (1.28e-13)	2.13e-7 (1.42e-13)	8.79e-7 (3.41e-13)	1.87e-6 (4.26e-13)	8.53e-14 (5.68e-14)
2.625	3.81e-7 (9.68e-12)	3.81e-7 (4.26e-13)	1.34e-7 (1.28e-13)	4.98e-7 (8.53e-14)	1.35e-6 (5.26e-13)	2.56e-13 (2.84e-13)
2.800	2.12e-7 (1.05e-11)	2.12e-7 (3.13e-13)	8.82e-8 (5.83e-13)	2.94e-7 (1.14e-13)	9.59e-7 (1.71e-13)	2.84e-14 (4.26e-14)
2.975	1.21e-7 (1.15e-11)	1.21e-7 (4.26e-13)	6.08e-8 (2.84e-14)	1.79e-7 (2.98e-13)	6.85e-7 (2.27e-13)	4.26e-14 (8.53e-14)
3.150	7.31e-8 (9.72e-12)	7.31e-8 (1.71e-13)	4.38e-8 (4.26e-13)	1.13e-7 (2.84e-13)	4.95e-7 (1.85e-13)	6.39e-13 (1.42e-13)
3.325	4.68e-8 (1.04e-11)	4.68e-8 (7.11e-14)	3.29e-8 (2.13e-13)	7.42e-8 (3.27e-13)	3.62e-7 (5.68e-14)	8.53e-14 (5.54e-13)
3.500	3.21e-8 (1.40e-11)	3.21e-8 (3.69e-13)	2.55e-8 (2.98e-13)	5.06e-8 (3.84e-13)	2.69e-7 (4.55e-13)	1.99e-13 (1.42e-13)

Table 4: Energy errors of H₂O versus the number of qubits.

Bond length (Å)	8 qubits (10 qubits)					
	CS-VQE GS	CS-QSE E_0	CS-QSE E_1	CS-QSE E_2	CS-QSE E_3	CS-QSE E_4
0.700	1.21e-4 (3.81e-6)	1.18e-4 (4.80e-13)	1.35e-4 (4.50e-13)	1.44e-4 (1.00e-13)	6.12e-5 (2.00e-13)	2.44e-4 (1.00e-13)
0.875	9.42e-5 (4.99e-6)	8.93e-5 (2.00e-13)	1.15e-4 (4.00e-13)	1.41e-4 (1.00e-13)	4.02e-5 (1.00e-13)	1.33e-4 (2.00e-13)
1.050	6.94e-5 (3.11e-6)	6.76e-5 (2.00e-13)	8.66e-5 (3.00e-13)	1.30e-4 (1.00e-13)	2.78e-5 (1.00e-13)	6.76e-5 (1.00e-13)
1.225	5.64e-5 (2.22e-6)	5.28e-5 (1.00e-13)	5.96e-5 (1.00e-13)	1.18e-4 (2.00e-13)	2.20e-5 (1.00e-13)	8.08e-5 (1.00e-13)
1.400	4.57e-5 (1.38e-4)	4.12e-5 (2.00e-13)	4.06e-5 (1.00e-13)	8.41e-5 (1.00e-13)	1.69e-5 (1.00e-13)	1.51e-5 (1.00e-13)
1.575	3.63e-5 (1.39e-4)	3.03e-5 (1.00e-13)	1.91e-5 (1.00e-13)	5.71e-5 (4.00e-13)	1.13e-5 (1.00e-13)	3.57e-5 (1.00e-13)
1.750	2.88e-5 (1.31e-4)	2.16e-5 (1.00e-13)	1.14e-5 (1.00e-13)	3.82e-5 (2.00e-13)	7.80e-6 (2.00e-13)	3.95e-6 (1.00e-13)
1.925	2.22e-5 (1.17e-4)	1.45e-5 (1.00e-13)	4.70e-6 (1.00e-13)	2.44e-5 (1.00e-13)	4.70e-6 (1.00e-13)	3.24e-6 (1.00e-13)
2.100	1.63e-5 (8.93e-5)	7.82e-6 (1.00e-13)	1.95e-6 (1.00e-13)	1.51e-5 (1.00e-13)	3.88e-6 (1.00e-13)	6.27e-6 (1.00e-13)
2.275	1.10e-5 (6.14e-5)	2.91e-6 (1.00e-13)	1.21e-6 (1.00e-13)	8.85e-6 (2.00e-13)	2.33e-6 (2.00e-13)	1.19e-6 (1.00e-13)
2.450	6.81e-6 (3.66e-5)	1.99e-6 (1.00e-13)	1.03e-6 (1.00e-13)	4.75e-6 (4.00e-13)	2.45e-6 (1.00e-13)	1.35e-6 (2.00e-13)
2.625	3.85e-6 (1.98e-5)	1.89e-6 (1.00e-13)	8.47e-7 (1.00e-13)	2.37e-6 (1.00e-13)	2.68e-6 (1.00e-13)	1.68e-6 (2.00e-13)
2.800	1.82e-6 (1.03e-5)	1.89e-6 (1.00e-13)	7.37e-7 (1.00e-13)	1.54e-6 (1.00e-13)	2.33e-6 (2.00e-13)	2.11e-6 (2.00e-13)
2.975	4.78e-7 (5.59e-9)	1.86e-6 (1.20e-13)	8.26e-7 (1.00e-13)	1.32e-6 (5.00e-13)	2.96e-7 (2.00e-13)	4.44e-5 (2.00e-13)
3.150	1.60e-7 (2.68e-8)	1.89e-6 (1.00e-13)	9.04e-7 (1.00e-13)	1.13e-6 (2.00e-13)	1.71e-7 (1.00e-13)	1.99e-5 (4.00e-13)
3.325	7.99e-8 (7.93e-9)	1.92e-6 (1.00e-13)	1.55e-7 (T)	1.03e-6 (2.00e-13)	2.14e-7 (2.00e-13)	9.44e-6 (1.00e-13)
3.500	3.92e-8 (4.10e-9)	1.85e-6 (1.00e-13)	1.01e-6 (1.00e-13)	9.46e-7 (1.00e-13)	2.61e-7 (2.00e-13)	3.69e-6 (3.00e-13)

Table 5: Energy errors of HCl versus the number of qubits.

Bond length (Å)	8 qubits (10 qubits)					
	CS-VQE GS	CS-QSE E_0	CS-QSE E_1	CS-QSE E_2	CS-QSE E_3	CS-QSE E_4
0.700	1.57e-5 (1.04e-5)	1.57e-5 (1.04e-5)	8.17e-6 (8.17e-6)	3.60e-5 (3.37e-5)	2.32e-5 (2.25e-5)	5.04e-5 (4.39e-5)
0.875	1.14e-5 (8.77e-6)	1.14e-5 (8.77e-6)	7.27e-6 (7.27e-6)	2.09e-5 (2.03e-5)	1.86e-5 (1.81e-5)	3.51e-5 (3.45e-5)
1.050	8.05e-6 (5.74e-6)	8.07e-6 (5.74e-6)	5.60e-6 (5.60e-6)	1.08e-5 (1.05e-5)	1.39e-5 (1.36e-5)	2.26e-5 (2.27e-5)
1.225	5.08e-6 (3.43e-6)	5.09e-6 (3.43e-6)	3.44e-6 (3.44e-6)	5.66e-6 (5.58e-6)	1.05e-5 (1.03e-5)	1.43e-5 (1.43e-5)
1.400	2.93e-6 (2.20e-6)	2.93e-6 (2.20e-6)	1.95e-6 (1.95e-6)	3.02e-6 (2.93e-6)	7.63e-6 (7.45e-6)	5.09e-6 (4.64e-6)
1.575	1.73e-6 (1.76e-6)	1.73e-6 (1.76e-6)	1.12e-6 (1.12e-6)	3.25e-6 (3.09e-6)	5.18e-6 (5.03e-6)	4.25e-6 (4.11e-6)
1.750	1.01e-6 (1.48e-6)	1.01e-6 (1.48e-6)	6.37e-7 (6.35e-7)	4.61e-6 (4.35e-6)	4.03e-6 (3.90e-6)	4.40e-6 (4.26e-6)
1.925	6.00e-7 (1.24e-6)	5.99e-7 (1.24e-6)	3.66e-7 (3.65e-7)	6.01e-6 (5.57e-6)	3.45e-6 (3.34e-6)	5.24e-6 (5.04e-6)
2.100	3.68e-7 (1.00e-6)	3.67e-7 (1.00e-6)	2.16e-7 (2.16e-7)	7.00e-6 (6.25e-6)	3.12e-6 (3.02e-6)	5.86e-6 (5.61e-6)
2.275	2.52e-7 (7.53e-7)	2.52e-7 (7.53e-7)	1.42e-7 (1.42e-7)	7.67e-6 (6.57e-6)	3.00e-6 (2.90e-6)	6.30e-6 (6.00e-6)
2.450	1.93e-7 (5.16e-7)	1.93e-7 (5.16e-7)	1.04e-7 (1.04e-7)	7.97e-6 (6.51e-6)	2.99e-6 (2.90e-6)	6.56e-6 (6.26e-6)
2.625	1.56e-7 (3.22e-7)	1.56e-7 (3.22e-7)	7.96e-8 (7.96e-8)	7.83e-6 (6.05e-6)	3.04e-6 (2.96e-6)	6.74e-6 (6.46e-6)
2.800	1.33e-7 (2.07e-7)	1.32e-7 (2.06e-7)	6.46e-8 (6.45e-8)	7.30e-6 (5.32e-6)	3.14e-6 (3.07e-6)	6.74e-6 (6.49e-6)
2.975	1.20e-7 (1.26e-7)	1.20e-7 (1.26e-7)	5.52e-8 (5.51e-8)	6.43e-6 (4.48e-6)	3.32e-6 (3.25e-6)	6.33e-6 (6.12e-6)
3.150	1.15e-7 (1.01e-7)	1.15e-7 (1.01e-7)	4.91e-8 (4.91e-8)	5.28e-6 (3.34e-6)	3.57e-6 (3.51e-6)	5.23e-6 (5.03e-6)
3.325	1.16e-7 (9.64e-8)	1.16e-7 (9.64e-8)	4.46e-8 (4.46e-8)	3.91e-6 (2.07e-6)	3.93e-6 (3.87e-6)	3.96e-6 (3.78e-6)
3.500	1.19e-7 (9.77e-8)	1.19e-7 (9.77e-8)	4.14e-8 (4.13e-8)	2.43e-6 (7.33e-7)	4.42e-6 (4.36e-6)	2.92e-6 (2.75e-6)

4.2 Performance and Resource Efficiency Analysis

In the traditional QSE, the measurement cost for constructing subspace matrices is typically dictated by an $O(N^8)$ scaling, stemming from the $O(N^4)$ growth of the excitation operator pool. This frame-

work achieves a significant efficiency gain through a dual compression mechanism: First, it integrates $\mathbb{Z}_2$ symmetry tapering with the CS-VQE to reduce the effective simulation space from N qubits to a compact u -qubit manifold; Second, a projection mapping π_ν^u is introduced to automatically filter out excitation

terms that violate stabilizer-sector constraints.

Figs. 5 and 6 quantify the reduction in the excitation-operator pool induced by the tapering and contextual-subspace projection. As shown in Fig. 5, the full excitation-operator pool O_{full} contains $10^4 - 10^5$ Pauli terms (e.g., 10,903 for LiH/HF, 20,854 for H₂O, and 92,171 for HCl), whereas symmetry tapering reduces the raw excitation pool to 730 Pauli strings for LiH/HF, 1,541 for H₂O, and 11,723 for HCl. Applying π_v^u further prunes the pool substantially: for LiH the projected pool contains only 12, 40, 148, 203, and 279 Pauli strings at $u = 2 - 6$, and reaches the tapered-pool limit (730) only when u coincides with the number of tapered qubits; Similar trends hold for HF (12, 40, 64, 108, 279 at $u = 2 - 6$) and for the larger systems, where the size of the projected pool increases with u but remains far below the tapered pool across the compact subspaces (e.g., 748 at $u = 8$ for H₂O and 1,641 at $u = 10$ for HCl). Moreover, Fig. 6 shows that the number of Pauli strings retained after applying π_v^u remains nearly unchanged along the bond-length scan when u is fixed, indicating that the stabilizer-sector constraints largely fix which Pauli components are retained by the projection, with nuclear-geometry changes reflected mainly in their coefficients.

These reductions imply that the projected excitation operators act entirely within the chosen stabilizer sector. In practice, we select u according to the minimum contextual register size that achieves chemical precision in Fig. 2, and use the corresponding projected pool size N_{ops} as the effective operator budget for QSE.

As illustrated in Fig. 2, the minimal simulation qubits (u_{sim}) required to reach this precision threshold for LiH, HF, H₂O, and HCl are 4, 6, 8, and 6, respectively. Correlating these results with Table 6, the combination of qubit-tapering and contextual projection leads to a substantial reduction in the operator pool size across all considered systems.

The LiH operator pool, originally comprising 10,903 terms, is pared down to 730 via qubit-tapering and further distilled to 148 effective operators (N_{ops}) via CS projection, marking a significant 98.6% reduction in total operator count. This compression efficiency extends to HF and H₂O, where the operator pool sizes are pruned down to 279 and 748 terms, yielding substantial reduction rates of 97.4% and 96.4%, respectively.

The efficacy of this framework is most striking for HCl, where a massive 92,171 terms are compressed to a mere 279, corresponding to an aggregate reduction of 99.7%. This sparsity directly alleviates the measurement overhead dictated by $O(N_{ops}^2 \cdot N_H)$ scaling

Table 6: Molecular resource thresholds and operator reduction statistics.

Molecule	u_{sim}	N_{ops}	O_{full}	O_{full}^{taper}	R_{cs} (%)	R_{total} (%)
LiH	4	148	10,903	730	79.7	98.6
HF	6	279	10,903	730	61.8	97.4
H ₂ O	8	748	20,854	1,541	51.5	96.4
HCl	6	279	92,171	11,723	97.6	99.7

Note: u_{sim} denotes the minimum simulation qubits required to reach the chemical-precision threshold. O_{full} and O_{full}^{taper} are pool sizes before and after qubit-tapering. R_{cs} is the reduction from CS-projection ($O_{full}^{taper} \rightarrow N_{ops}$), and R_{total} is the cumulative reduction ($O_{full} \rightarrow N_{ops}$).

before measurement grouping, and improves the feasibility of implementing such workflows on resource-constrained NISQ devices.

Compared to the standard CS-VQD algorithm, the CS-QSE facilitates higher parallel scalability and numerical resilience, addressing key bottlenecks in excited-state simulations. The CS-VQD relies on a serial optimization protocol involving k successive variational searches to access k states, leading to a linear scaling of measurement complexity and often necessitating a $2n$ -qubit register.

In contrast, the CS-QSE requires only a single optimization run. By performing subspace diagonalization within the original n -qubit register, it simultaneously extracts the low-lying spectrum (e.g., $E_0 \sim E_4$). This linear subspace architecture precludes cumulative error propagation, facilitating a more robust representation of physical correlations across the energy spectrum.

The results in Table 7 demonstrate that the CS-QSE consistently achieves high-fidelity excited-state energies for all molecular systems studied. For the HCl system, the first excited-state error remains as low as 1.94×10^{-5} Ha. Despite the drastic sparsification of the operator pool, this result is consistent with CS-VQD benchmarks. Notably, the second excited-state error (E_2) is held to 7.41×10^{-5} Ha, attained concurrently without the need for state-specific optimization. These results demonstrate that the pruned operator space retains the necessary many-body correlations to support high-precision excited-state simulations.

For LiH, the CS-QSE achieves an E_2 error of 1.16×10^{-3} Ha, on par with the 1.11×10^{-3} Ha precision of the CS-VQD, while obviating the three successive optimization runs required by the latter. For the more complex HF and H₂O systems, the parallel extraction scheme attains accuracies of 8.71×10^{-5} Ha and 1.36×10^{-4} Ha, respectively.

The N_{ops} scaling demonstrates that the CS-QSE effectively offloads the computational burden from




Figure 5: Scaling of Pauli operator cardinality with qubit count.

Table 7: Comparison between the CS-VQD and the proposed CS-QSE framework.

Metric Type	Item / Molecule	CS-VQD	CS-QSE
Algorithm Resources	Qubits Required	$2n$ (with ancilla)	n (Direct)
	VQE Optimization	k times (Serial)	1 time
	Complexity Scaling	$O(k \cdot N_H)$	$O(N_{ops}^2 \cdot N_H)$
	Numerical Stability	Sequence error risk	Linear subspace
Accuracy (E_1 Error)	LiH ($u = 4$)	6.76×10^{-4}	6.59×10^{-4}
	HF ($u = 6$)	1.94×10^{-5}	1.94×10^{-5}
	H ₂ O ($u = 8$)	1.05×10^{-4}	1.02×10^{-4}
	HCl ($u = 6$)	1.94×10^{-5}	1.94×10^{-5}
High-state (E_2 Error)	LiH	1.11×10^{-3}	1.16×10^{-3}
	HF	—	8.71×10^{-5}
	H ₂ O	—	1.36×10^{-4}
	HCl	—	7.41×10^{-5}

Note: Energy error values are in Hartree. N_{ops} denotes the number of effective excitation operators in the projected pool, N_H represents the number of Pauli terms in the Hamiltonian, and k is the number of targeted states. “—” indicates calculations that are not performed.

quantum hardware to classical eigensolvers. By leveraging linear subspace diagonalization, the framework obviates the convergence instabilities endemic to CS-VQD, ensuring numerical robustness and physical self-consistency even under a significantly reduced measurement budget. The aforementioned opera-

tor pool reduction is paralleled by a contraction of the Hamiltonian Pauli one-norm, $\|\hat{H}\|_1 = \sum_a |w_a|$. The contextual-subspace projection π_ν^u reduces the observed Hamiltonian Pauli one-norm through two mechanisms: the projector P_ν exactly annihilates any terms in $\hat{H}$ that anticommute with the chosen sta-




Figure 6: Scaling of Pauli string counts for excitation operators along bond-length scan.

Table 8: Variation of the total Hamiltonian Pauli 1-norm.

Molecule	$\ \hat{H}\ _1$	$\ \hat{H}_{\text{taper}}\ _1$	$\ \hat{H}_{\text{CS}}\ _1$ (Ha) for u qubits in the contextual subspace						
			$u=4$	$u=5$	$u=6$	$u=7$	$u=8$	$u=9$	$u=10$
LiH	16.47	16.12	10.03	11.10	14.27	14.45	16.12	—	—
HF	138.87	137.87	103.77	113.54	118.26	125.77	137.87	—	—
H ₂ O	118.42	117.80	79.38	81.12	84.33	91.25	97.17	105.31	117.80
HCl	607.98	607.98	458.32	464.43	466.67	474.39	476.87	484.49	494.48

Note: All values are in Hartree. “—” indicates that the contextual-subspace size exceeds the tapered register size for that molecule, in which case CS projection is not applied. $\|\hat{H}\|_1$, $\|\hat{H}_{\text{taper}}\|_1$, and $\|\hat{H}_{\text{CS}}\|_1$ denote the Pauli 1-norms of the original full Hamiltonian, the Hamiltonian after qubit tapering, and the Hamiltonian after contextual-subspace projection, respectively. u is the number of the qubit in the contextual subspace.

bilizers, and surviving Pauli operators that map to identical reduced Pauli strings are algebraically combined, which can produce coefficient cancellation and thereby lower the resulting Pauli one-norm. As shown in Table 8, $\|\hat{H}_{\text{CS}}\|_1$ is notably reduced compared to $\|\hat{H}\|_1$ for the systems studied here. This reduction lowers a conservative shot-noise bound for Hamiltonian energy estimation; under an optimal allocation weighted by $|w_a|$, the variance of the energy estimator satisfies $\text{Var}(\hat{E}) \leq \|\hat{H}\|_1^2 / N_{\text{shots}}$ [39, 49]. Thus, the lower Pauli one-norm implies a meaningful decrease in the theoretical upper bound of the shot count re-

quired for Hamiltonian expectation-value estimation at a target accuracy. However, the CS-QSE sampling cost also depends on the number, grouping, and conditioning of the QSE matrix-element observables, as assessed in Fig. 7 and Appendix A.

5 Conclusion

A resource-efficient framework for molecular excited-state calculations is developed by integrating the CS method with the QSE. Through a unified algebraic

mapping, the excitation operator pool is projected into the same contextual subspace as the reduced Hamiltonian. This structural alignment circumvents the exponential basis scaling inherent in the traditional QSE, enabling the simulation of complex systems under a restricted quantum resource budget.

It is demonstrated that this unified scheme curtails the effective operator count by an order of magnitude while maintaining errors within the targeted tolerance relative to the exact FCI baseline. Specifically, the generalized eigenvalue problem formulated within the reduced subspace affords a refined linear-response correction to the ground state, while resolving the low-lying excited-state manifold to within the chemical-precision threshold.

While hardware noise remains a pertinent challenge in the NISQ era, future developments will focus on enhancing the robustness of the algorithm by reasonable hardware-adaptive optimizations. We anticipate that the synergy between this framework and some advanced error mitigation techniques such as Pauli grouping, classical shadows, and Zero-Noise Extrapolation will provide a scalable and practical pathway for high-precision quantum chemistry simulations on the near-term quantum devices.

CRedit authorship contribution statement

Chao Liu: Conceptualization, Methodology, Investigation, Data curation, Software, Writing-original draft. **Yuxin Deng:** Investigation, Supervision, Funding acquisition, Project administration, Writing-review and editing.

Declaration of competing interest

The authors declare that there are no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

The work was supported by the National Key R&D Program of China under Grant No. 2023YFA1009403, the National Natural Science Foundation of China under Grant No. 62472175, CCF-QuantumCtek Superconducting Quantum Computing Special Cooperation Program under

Grant No. CCF-QC2025004, and the Fundamental Research Funds for the Central Universities.

Data availability

Data will be made available on request.

Appendix

A. Finite-Shot Benchmarks and Energy Errors

To assess the practical measurement cost of the ground- and excited-state calculations, we use CS-VQE to obtain the ground-state reference and compare CS-VQD with CS-QSE for the low-lying excited states of LiH in the STO-3G basis at bond length 1.57473 Å. We use $u = 4$ and $u = 6$ as the number of qubits retained in the contextual correction and compare how the finite-shot energy error changes with the total measurement cost.

In the finite-shot benchmarks, the ansatz parameters are fixed at the optima obtained from noiseless statevector simulations, and shot-based sampling is used only to estimate the post-optimization energies. For CS-VQE and CS-VQD, finite-shot sampling is used to estimate the Hamiltonian expectation value of the optimized ansatz state. In CS-VQD, the k -th excited state is obtained by minimizing a modified variational cost function that augments the Hamiltonian expectation value with penalty terms proportional to the overlaps with the previously obtained lower-energy states. Consequently, a finite-shot implementation of CS-VQD requires sampling both the Hamiltonian expectation value and the overlap terms entering these deflation penalties. For CS-QSE, finite-shot sampling is used to estimate the projected Hamiltonian and the QSE metric matrix defined in Eq. 20. Solving the corresponding generalized eigenvalue problem yields all low-lying QSE eigenvalues from the same sampled matrices. Therefore, the QSE measurement budget is shared across the ground, first excited, and second excited states reported in Fig. 7.

For each shot count, we use $R = 20$ independent repeats to estimate the mean absolute error and its statistical uncertainty, and the horizontal-axis cost in Fig. 7 includes all measurements from these repeats. The shot counts are varied over 500, 1000, 1500, 2000, 3000, 4000, 5000, 7000, 10000, 20000, 50000, 100000, and 200000.

For each observable or matrix element to be sampled, the corresponding Pauli terms are partitioned

into mutually qubit-wise commuting (QWC) sets. In practice, we construct this partition by coloring the QWC-incompatibility graph, whose vertices are Pauli terms and whose edges connect pairs of terms that are not QWC. Each color class defines one tensor-product-basis measurement setting, allowing all Pauli terms in that class to be estimated from the same shot data after the appropriate single-qubit basis rotations.

For each independent repeat, the measurement cost is defined as

$$N_{\text{meas}}^{\text{repeat}} = N_{\text{shots}} G_{\text{QWC}}, \quad (23)$$

where G_{QWC} is the number of qubit-wise commuting measurement groups. The horizontal-axis cost in Fig. 7 is the repeat-aggregated cost used to estimate the plotted mean error and error bar.

For the LiH data in Fig. 7, the CS-VQE ground-state Hamiltonian sampling uses $G_H(4) = 25$ and $G_H(6) = 47$ QWC groups. The CS-QSE global operator pool uses $G_{\text{QSE}}(4) = 49$ and $G_{\text{QSE}}(6) = 721$ QWC groups. For CS-VQD, the first excited state uses $G_{\text{VQD}}^{(1)}(4) = 26$ and $G_{\text{VQD}}^{(1)}(6) = 48$, while the second excited state uses $G_{\text{VQD}}^{(2)}(4) = 27$ and $G_{\text{VQD}}^{(2)}(6) = 49$. These CS-VQD effective group counts include the additional overlap measurements required by the deflation penalty terms.


This QWC partition reduces the number of distinct measurement settings, but it does not change the Pauli 1-norm that controls the shot-noise scaling [49, 50]. For an observable $\hat{O} = \sum_k c_k P_k$, the standard finite-shot scaling [39] is expressed as

$$N_{\text{shots}} = \mathcal{O}\left(\frac{\|\hat{O}\|_1^2}{\epsilon^2}\right), \quad \|\hat{O}\|_1 = \sum_k |c_k|. \quad (24)$$


The energy error in Fig. 7 is measured relative to the corresponding FCI reference energy. The dotted gray guide line indicates the ideal shot-noise scaling proportional to $N_{\text{meas}}^{-1/2}$ and is anchored at the first data point. The horizontal gray dash-dotted line marks chemical precision, 1.6×10^{-3} Ha. The results show that CS-QSE can obtain multiple low-lying eigenstates from one shared matrix-element sampling budget. In contrast, CS-VQD evaluates excited states sequentially and requires additional overlap measurements for the deflation penalties.

B. Numerical Experiments with the 6-31G Basis Set


To further assess the robustness of the proposed approach, we performed additional numerical experiments on LiH with the 6-31G basis set. In this case,



(a) Ground State



(b) the first Excited state



(c) the second Excited state

Figure 7: Energy-error scaling of CS-QSE and CS-VQD versus total measurement cost for the (a) ground state, (b) the first excited state, and (c) the second excited state within the symmetry sector of LiH.

the full electronic Hamiltonian acts on 22 qubits, and qubit tapering reduces the problem to 19 qubits. The excitation pool contains 136,998 Pauli terms after Jordan–Wigner mapping; reaching the targeted precision relative to the FCI baseline requires retaining only 7,029 of these after contextual-subspace projection, i.e., about 5.1% of the original pool. Compared




Figure 8: Energy error scaling of the CS-VQE and CS-QSE for LiH using a UCCSD ansatz as a function of the number of qubits.

with the STO-3G calculations, the compression in the 6-31G basis is less aggressive, reflecting the much larger excitation space.

Figure 8 reports the energy error of CS-VQE and CS-QSE as a function of the number of qubits retained after contextual-subspace compression. For the ground state, CS-QSE (E_0) reaches the chemical-precision threshold at $N = 8$, indicating that QSE post-processing effectively recovers the correlation energy lost during contextual-subspace compression. For the low-lying excited states, CS-QSE (E_1) and CS-QSE (E_2) reach this same threshold at $N = 16$, showing that the same strategy remains effective for excited-state calculations in a larger basis set. These results indicate that qubit tapering alone provides only a limited reduction in qubit count, since it is constrained by the underlying symmetries of the Hamiltonian.

By contrast, the combined contextual-subspace and QSE procedure yields a substantially stronger reduction while maintaining the chemical-precision threshold. As the basis set grows, the excitation space expands rapidly and the compression ratio becomes smaller than in the STO-3G case, making QSE post-processing increasingly important for maintaining accuracy under stronger compression. Overall, the LiH/6-31G results indicate that the proposed CS-QSE framework remains effective beyond the minimal basis for this test case.

References

[1] Raffaele Santagati, Alan Aspuru-Guzik, Ryan Babbush, Matthias Degroote, Leticia González, Elica Kyoseva, Nikolaj Moll, Markus Oppel,

Robert M Parrish, Nicholas C Rubin, Michael Streif, Christofer S Tautermann, Horst Weiss, Nathan Wiebe, and Clemens Utschig-Utschig. Drug design on quantum computers. *Nat. Phys.*, 20(4):549–557, 2024.

- [2] Jiabin Yu, B Andrei Bernevig, Raquel Queiroz, Enrico Rossi, Päivi Törmä, and Bohm Jung Yang. Quantum geometry in quantum materials. *NPJ Quantum Mater.*, 10(1):101, 2025.
- [3] Shubin Sun, Yudi Zhang, Xin Shi, Wen Sun, Claudia Felser, Wei Li, and Guowei Li. From charge to spin: an in-depth exploration of electron transfer in energy electrocatalysis. *Adv. Mater.*, 36(37):2312524, 2024.
- [4] Paul Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *J. Stat. Phys.*, 22(5):563–591, 1980.
- [5] Richard P Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21(6/7):553–555, 1982.
- [6] Yifan Ding, Zi Li, and Nanrun Zhou. Quantum generative adversarial network based on the quantum Born machine. *Adv. Eng. Inform.*, 68:103622, 2025.
- [7] Xiaochang He, Lihua Gong, and Nanrun Zhou. Image denoising with hybrid classical-quantum convolutional neural network. *Memet. Comput.*, 17(2):23, 2025.
- [8] Lihua Gong, Yuqi Chen, Shun Zhou, and Qingwei Zeng. Dual discriminators quantum generation adversarial network based on quantum convolutional neural network. *Adv. Quantum Technol.*, 8(10):e2500224, 2025.
- [9] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.*, 5:4213, 2014.
- [10] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.*, 18(2):023023, 2016.
- [11] Shengbin Wang, Peng Wang, Guihui Li, Shubin Zhao, Dongyi Zhao, Jing Wang, Yuan Fang, Menghan Dou, Yongjian Gu, Yuchun Wu,

- et al. Variational quantum eigensolver with linear depth problem-inspired ansatz for solving portfolio optimization in finance. *Sci. China Inf. Sci.*, 68(8):1–11, 2025.
- [12] Harper R Grimsley, Sophia E Economou, Edwin Barnes, and Nicholas J Mayhall. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.*, 10(1):3007, 2019.
- [13] Panagiotis G Anastasiou, Yanzhu Chen, Nicholas J Mayhall, Edwin Barnes, and Sophia E Economou. TETRIS-ADAPT-VQE: An adaptive algorithm that yields shallower, denser circuit ansätze. *Phys. Rev. Res.*, 6(1):013254, 2024.
- [14] Nonia Vaquero-Sabater, Abel Carreras, and David Casanova. Pruned-adapt-vqe: compacting molecular ansätze by removing irrelevant operators. *J. Chem. Theory Comput.*, 21(18):8720–8728, 2025.
- [15] Gerald Knizia and Garnet Kin-Lic Chan. Density matrix embedding: A strong-coupling quantum embedding theory. *J. Chem. Theory Comput.*, 9(3):1428–1432, 2013.
- [16] Qiming Sun and Garnet Kin-Lic Chan. Quantum embedding theories. *Acc. Chem. Res.*, 49(12):2705–2712, 2016.
- [17] Akhil Shajan, Danil Kaliakin, Abhishek Mitra, Javier Robledo Moreno, Zhen Li, Mario Motta, Caleb Johnson, Abdullah Ash Saki, Susanta Das, Iskandar Sitdikov, et al. Toward quantum-centric simulations of extended molecules: Sample-based quantum diagonalization enhanced with density matrix embedding theory. *J. Chem. Theory Comput.*, 21(14):6801–6810, 2025.
- [18] William M Kirby and Peter J Love. Contextuality test of the nonclassicality of variational quantum eigensolvers. *Phys. Rev. Lett.*, 123(20):200501, 2019.
- [19] William M Kirby and Peter J Love. Classical simulation of noncontextual Pauli Hamiltonians. *Phys. Rev. A*, 102(3):032418, 2020.
- [20] William M Kirby, Andrew Tranter, and Peter J Love. Contextual subspace variational quantum eigensolver. *Quantum*, 5:456, 2021.
- [21] Tim Weaving, Alexis Ralli, William M Kirby, Andrew Tranter, Peter J Love, and Peter V Coveney. A stabilizer framework for the contextual subspace variational quantum eigensolver and the noncontextual projection ansatz. *J. Chem. Theory Comput.*, 19(3):808–821, 2023.
- [22] Alexis Ralli, Tim Weaving, Andrew Tranter, William M Kirby, Peter J Love, and Peter V Coveney. Unitary partitioning and the contextual subspace variational quantum eigensolver. *Phys. Rev. Res.*, 5(1):013095, 2023.
- [23] Zhiding Liang, Zhixin Song, Jinglei Cheng, Hang Ren, Tianyi Hao, Rui Yang, Yiyu Shi, and Tongyang Li. Combining parameterized pulses and contextual subspace for more practical VQE. In *Proceedings of the 61st ACM/IEEE Design Automation Conference, DAC ’24*, New York, NY, USA, 2024. Association for Computing Machinery.
- [24] Tim Weaving, Alexis Ralli, Peter J Love, Sauro Succi, and Peter V Coveney. Contextual subspace variational quantum eigensolver calculation of the dissociation curve of molecular nitrogen on a superconducting quantum computer. *NPJ Quantum Inf.*, 11(1):25, 2025.
- [25] Oscar Higgott, Daochen Wang, and Stephen Brierley. Variational quantum computation of excited states. *Quantum*, 3:156, 2019.
- [26] Qianjun Yao and He Li. Efficient excited-state calculations for molecules based on contextual subspace method and symmetry optimizations. *New J. Phys.*, 27(9):094508, 2025.
- [27] Michal Ďuriška, Ivana Miháliková, and Martin Friák. Quantum computing of the electronic structure of crystals by the variational quantum deflation algorithm. *Phys. Scr.*, 100(4):045105, 2025.
- [28] Jarrod R McClean, Mollie E Kimchi-Schwartz, Jonathan Carter, and Wibe A De Jong. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys. Rev. A*, 95(4):042308, 2017.
- [29] Seonghoon Choi and Artur F. Izmaylov. Measurement optimization techniques for excited electronic states in near-term quantum computing algorithms. *J. Chem. Theory Comput.*, 19(11):3184–3193, 2023.

- [30] Smik Patel, Praveen Jayakumar, Rick Huang, Tao Zeng, and Artur F. Izmaylov. Quantum seniority-based subspace expansion: Linear combinations of short-circuit unitary transformations for the electronic structure problem. *J. Chem. Theory Comput.*, 22(8):3937–3949, 2026.
- [31] João C Getelina, Prachi Sharma, Thomas Iadecola, Peter P Orth, and Yong-Xin Yao. Quantum subspace expansion in the presence of hardware noise. *APL Quantum*, 1(3), 2024.
- [32] Pauline J. Ollitrault, Abhinav Kandala, Chun-Fu Chen, Panagiotis Kl. Barkoutsos, Antonio Mezzacapo, Marco Pistoia, Sarah Sheldon, Stefan Woerner, Jay M. Gambetta, and Ivano Tavernelli. Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor. *Phys. Rev. Res.*, 2:043140, 2020.
- [33] Simon Kochen and Ernst P. Specker. The problem of hidden variables in quantum mechanics. *J. Math. Mech.*, 17:59–87, 1967.
- [34] RW Spekkens. Contextuality for preparations, transformations, and unsharp measurements. *Phys. Rev. A*, 71(5):052108, 2005.
- [35] N David Mermin. Hidden variables and the two theorems of John Bell. *Rev. Mod. Phys.*, 65(3):803–815, 1993.
- [36] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 10th anniversary edition, 2010.
- [37] Attila Szabo and Neil S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. McGraw-Hill, New York, 1989.
- [38] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.
- [39] Dave Wecker, Matthew B Hastings, and Matthias Troyer. Progress towards practical quantum variational algorithms. *Phys. Rev. A*, 92(4):042303, 2015.
- [40] Alexis Ralli and Tim Weaving. Symmer. <https://github.com/qmatter-labs/symmer>, 2022.
- [41] Tzu Ching Yen, Robert A Lang, and Artur F Izmaylov. Exact and approximate symmetry projectors for the electronic structure problem on a quantum computer. *J. Chem. Phys.*, 151(16), 2019.
- [42] Kanav Setia, Richard Chen, Julia E Rice, Antonio Mezzacapo, Marco Pistoia, and James D Whitfield. Reducing qubit requirements for quantum simulations using molecular point group symmetries. *J. Chem. Theory Comput.*, 16(10):6091–6097, 2020.
- [43] Qiming Sun, Timothy C Berkelbach, Nick S Blunt, George H Booth, Sheng Guo, Zhendong Li, Junzi Liu, James D McClain, Elvira R Sayfutyarova, Sandeep Sharma, Sebastian Wouters, and Garnet Kin-Lic Chan. PySCF: the Python-based simulations of chemistry framework. *WIREs Comput. Mol. Sci.*, 8(1):e1340, 2018.
- [44] Ali Javadi Abhari, Matthew Treinish, Kevin Krsulich, Christopher J Wood, Jake Lishman, Julien Gacon, Simon Martiel, Paul D Nation, Lev S Bishop, Andrew W Cross, Blake R Johnson, and Jay M Gambetta. Quantum computing with Qiskit. *arXiv preprint arXiv:2405.08810*, 2024.
- [45] The OpenFermion Developers. OpenFermion, 2025.
- [46] Jonas Jäger, Thierry N. Kaldenbach, Max Haas, and Erik Schultheis. Fast gradient-free optimization of excitations in variational quantum eigensolvers. *Commun. Phys.*, 8(418), 2025.
- [47] A. R. Hoy and P. R. Bunker. A precise solution of the rotation bending schrödinger equation for a triatomic molecule with application to the water molecule. *J. Mol. Spectrosc.*, 74(1):1–8, 1979.
- [48] Attila G. Császár, Gábor Czako, Tibor Furtenbacher, Jonathan Tennyson, Viktor Szalay, Sergei V. Shirin, Nikolai F. Zobov, and Oleg L. Polyansky. On equilibrium structures of the water molecule. *J. Chem. Phys.*, 122:214305, 2005.
- [49] Ophelia Crawford, Barnaby van Straaten, Daochen Wang, Thomas Parks, Earl Campbell, and Stephen Brierley. Efficient quantum measurement of Pauli operators in the presence of finite sampling error. *Quantum*, 5:385, 2021.
- [50] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F Izmaylov. Measurement optimization in the variational quantum eigensolver using a minimum clique cover approach. *J. Chem. Phys.*, 152:124114, 2020.

TOC Graphic

