

Neural network assisted fermionic compression encoding: A lossy quantum-selected configuration interaction framework for resource-efficient ground-state simulations

Yu-Cheng Chen^{1,*}, Ronin Wu^{2,†}, M. H. Cheng^{3,4} and Min-Hsiu Hsieh¹

¹*Hon Hai Research Institute, Taipei, Taiwan*

²*QunaSys Europe, Copenhagen, Denmark*

³*Department of Physics Blackett Laboratory, Imperial College London, London SW7 2AZ, United Kingdom*

⁴*Fraunhofer Institute for Industrial Mathematics, Fraunhofer-Platz 1, 67663 Kaiserslautern, Germany*



(Received 26 May 2025; accepted 15 September 2025; published 4 February 2026)

Quantum computing promises to revolutionize many-body simulations for quantum chemistry, but its potential is constrained by limited qubits and noise in current devices. In this work, we introduce the lossy quantum-selected configuration interaction (Lossy-QSCI) framework, which combines a lossy subspace Hamiltonian preparation pipeline with a generic QSCI selection process. This framework integrates a chemistry-inspired lossy random linear encoder (Chemical-RLE) with a neural network assisted fermionic expectation decoder (NN-FED). The RLE leverages fermionic number conservation to compress quantum states, reducing qubit requirements to $O(N \log_2 M)$ for M spin orbitals and N electrons while preserving crucial ground-state information and enabling self-consistent configuration recovery. NN-FED, powered by a neural network trained with minimal data, efficiently decodes these compressed states, overcoming the measurement challenges common in the approaches of the traditional QSCI and its variants. Through iterative quantum sampling and classical postprocessing, our hybrid method refines ground-state estimates with high efficiency. This framework offers a resource-efficient pathway for ground-state simulations on near-term noisy hardware and could inspire resource-efficient extensions to future devices by minimizing qubit overhead.

DOI: [10.1103/9x8b-4j41](https://doi.org/10.1103/9x8b-4j41)

I. INTRODUCTION

Quantum computing enables precise emulation of large-scale electronic structure problems, delivering promises to advance quantum chemistry [1]. This potential stems from the natural compatibility of quantum computers in representing quantum many-body systems, a task intractable for classical computers in many cases [2,3]. However, in the presence of quantum noise, near-term quantum devices face tremendous obstacles in the quest to realize this promise. Hybrid algorithms such as the variational quantum algorithm (VQA) [4,5] and quantum-selected configuration interaction (QSCI) [6] subsequently emerged as counter strategies, harnessing the power of quantum sampling to approximate molecular ground states.

VQA, for example, exploits parameterized quantum circuits [7] combined with classical optimization to converge to the ground-state energy iteratively. Recent advances have demonstrated improved ansatz design [8–11], enhanced error mitigation [12], and reduced measurement overhead [13], thereby bolstering VQA's robustness against noise and its

scalability for simulating complex electronic correlations. These developments reinforce the promise of hybrid algorithms in overcoming hardware limitations and paving the way for more efficient quantum chemistry simulations. Although VQA shows promise, it faces serious challenges. Its variational optimization is highly sensitive to quantum noise, prone to barren plateaus, and saddled with a measurement overhead that grows rapidly with system size. These limitations call for alternatives.

In contrast, QSCI leverages quantum sampling to construct effective subspaces for diagonalization and uses chemical priors to enhance state selection and energy estimation. It thereby addresses the noise sensitivity and high measurement cost issues inherent in variational quantum eigensolver (VQE), and offers a pathway for accurate electronic structure simulations. However, traditional QSCI does not compress the qubit space; it requires the full Hilbert space representation, leading to high resource demands and substantial classical postprocessing overhead that remains sensitive to noise. To overcome this, it is demonstrated in a scaled-up QSCI experiment (77 qubits) that by incorporating self-consistent configuration recovery and employing chemical priors via the localized unitary cluster Jastrow (LUCJ) ansatz [14], one can improve noise robustness; this yields higher state quality but adds complexity and further strains classical resources [15].

On the other hand, Adaptive Derivative-Assembled Pseudo-Trotter ansatz quantum-selected configuration interaction (ADAPT-QSCI) [16] iteratively grows the input state by ADAPT-VQE-style operator selection, thereby removing the need for a fixed chemical ansatz but at the cost of many

*Contact author: kesson.yc.chen@foxconn.com

†Contact author: ronin@qunasys.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

TABLE I. Comparison of original QSCI, its variants, and Lossy-QSCI frameworks. Criteria are defined as follows: Qubit efficiency: Denotes optimization of qubit usage. Bit-flip error mitigation: How the method mitigates bit-flip errors with particle number conservation. Chemical prior: Specifies how prior chemical knowledge constrains the solution subspace.

Method	QSCI [6]	Scaled-up QSCI [15]	TE-QSCI [17]	Lossy-QSCI (this work)
Qubit efficiency	$O(M)$	$O(M)$	$O(M)$	$O(N \log(M))$
Bit-flip error mitigation	Postselection	Configuration recovery	Postselection	Compression + postselection
Chemical prior	Customizable	LUCJ ansatz	Time-evolved-HF/UCCSD	Chemical-RLE

optimization rounds and repeated QSCI executions. More recently, time-evolved QSCI (TE-QSCI) [17] and concurrent work by Sugisaki *et al.* [18] have been proposed to prepare the QSCI input state via time evolution instead of variational optimization, naturally generating a compact subspace rich in electronically excited configurations. Later, Yu *et al.* [19] independently adopted the same “time-evolved input-state” motif, signifying growing interest in this direction. By reducing measurement overhead through effective selections of important configurations, TE-QSCI offers a more resource-efficient alternative that complements and, in some aspects, outperforms the scaled-up QSCI presented in Robledo-Moreno *et al.* [15] in balancing qubit efficiency and overall computational cost.

Nonetheless, even with these advancements, QSCI and its variants still suffer from sampling inefficiency. The full Hilbert space representation results in less compact wave functions and higher qubit resource demands [20]. Moreover, the inherent complex ansatz design and the variational optimization in standard QSCI or even TE-QSCI introduce additional overheads and challenges that strain both classical and quantum resources. These issues strongly motivate our work in advanced fermionic encoding methods, which aim to directly compress the qubit space while preserving the essential physics of the electronic Hamiltonian. By addressing both the qubit resource inefficiency and the burden of complex ansatz design, our approach offers a more resource-efficient pathway for ground-state simulations.

Traditional fermionic encodings, such as the Jordan-Wigner (JW) transformation [21], capture the full Fock space of electronic Hamiltonians, resulting in linear qubit scaling with system size. Recent number-conserving encodings [22] perform a log-qubit scaling by restricting the simulation to the number-conserving subspace or even compact configuration interaction (CI) subspace [23]. Specifically, focusing on states with a fixed electron number N from the M available spin orbitals, the effective number of qubits needed is roughly given by $\log_2 \binom{M}{N}$. For cases where $N \ll M$, this quantity scales approximately as the $O(N \log_2 M)$ scaling. However, they incur an increased energy measurement overhead, typically scaling as the number of states [e.g., $O(N^M)$ for number conserving], due to the need to project onto these conserved subspaces, which restricts the scaling efficiency of solving the subspace Hamiltonian. The fermionic expectation decoder (FED) has offered a partial solution by decoupling measurement scaling from electron number theoretically but has yet to be realized through implementation [24]. A related first-quantized approach also attains the nominal $O(N \log_2 M)$ qubit scaling by encoding an antisymmetrized N -electron wave function in a register of $N \log_2 M$ qubits [25]. However, because the wave

function lives in first quantization, Hamiltonian application and the method are not directly compatible with selected-CI workflows. These limitations highlight the need for innovative approaches that optimize both qubit usage and measurement efficiency.

Motivated by these limitations, we present a lossy quantum-selected configuration interaction (Lossy-QSCI) framework that integrates quantum encoding and classical decoding techniques. This approach compresses the qubit space while mitigating measurement overhead and preserving essential chemical information. Central to this approach is a chemistry-inspired lossy randomized linear encoding (Chemical-RLE) modified from the original lossless RLE [24], which lossy compresses quantum states while preserving essential chemical properties, substantially reducing qubit cost and scaling. In complement, a neural network assisted FED (NN-FED) is introduced to rapidly and accurately decode compressed states using minimal training data, surpassing traditional classical decoding methods. The framework operates iteratively, sampling compressed quantum states and refining ground-state estimates through classical postprocessing. Our method achieves chemical accuracy with fewer resources on C_2 and LiH molecules, demonstrating its efficacy and accuracy.

Table I presents a comparative analysis of the original QSCI, its variants, and the Lossy-QSCI framework. For qubit efficiency, Lossy-QSCI excels in reducing the qubit needs to $O(N \log_2 M)$ or fewer via Chemical-RLE, while others require the M qubits. For error mitigation, all approaches address bit-flip errors that violate the conservation of particle number. QSCI and TE-QSCI use postselection to reduce errors, and the scaled-up QSCI employs self-consistent configuration recovery to further mitigate them, whereas Lossy-QSCI leverages qubit compression to preclude the inclusion of nonconserving configurations. Incorporating chemical priors, QSCI and its scaled-up version utilize a chemical ansatz to constrain the solution subspace and boost sampling efficiency, while TE-QSCI relies on a carefully chosen ansatz (e.g., Hartree-Fock or UCCSD) combined with time evolution to generate a configuration-rich subspace naturally. On the other hand, Lossy-QSCI relies on its chemistry-informed encoding strategy, Chemical-RLE.

In practice, the advantages of Lossy-QSCI rely on two empirical steps: a biased, chemistry-guided injectivity check and a neural decoder trained on a limited set of sample determinants. Both procedures scale favorably for the benchmarked medium-sized systems studied here, but they are not yet universal tools that guarantee success for arbitrary, much larger molecules. Consequently, the present framework should be viewed as a pragmatic extension of existing

hybrid schemes (VQE, QSCI, etc.), whose own demonstrations have so far been confined to $\lesssim 10^2$ physical qubits and well-characterized Hamiltonians. Within these limitations, Lossy-QSCI still delivers a worthwhile trade-off between qubit count, measurement overhead, and final accuracy. It provides the first proof-of-concept algorithm that demonstrates improved configuration sampling not only via ansatz choice, but also encoding choice. Such a classical-quantum interface not only improves the expressivity of variational algorithms by using the classical encoding degrees of freedom, but it also significantly reduces the hardware burdens. By identifying encoding as a quantum resource for chemistry computation, we extend the potential quantum advantage for variational quantum algorithms.

II. PREVIOUS WORKS

A. Fermionic encoding

Let us consider a second quantized quantum chemistry Hamiltonian, comprising the fermionic creation and annihilation operators, $\hat{a}_{p\sigma}^\dagger$ and $\hat{a}_{p\sigma}$, and the one and two electronic integrals h_{pr} , h_{pqrs} associated [26]:

$$\hat{H} = \sum_{p,r,\sigma} h_{pr} \hat{a}_{p\sigma}^\dagger \hat{a}_{r\sigma} + \frac{1}{2} \sum_{\substack{p,r,q,s \\ \sigma,\tau}} h_{pqrs} \hat{a}_{p\sigma}^\dagger \hat{a}_{q\tau}^\dagger \hat{a}_{s\tau} \hat{a}_{r\sigma}, \quad (1)$$

where h_{pq} and h_{pqrs} are one- and two-electron integrals over spin orbitals (indices p, q, r, s), regarding to quantum chemistry standards; σ denotes spin indices (up or down). The Hamiltonian finds a natural encoding onto the qubit system via JW transformation, resulting in a linear combination of unitary $\hat{H} = \sum_i c_i U_i$ for the expectation estimation [21], where the index i refers to the i th unitary operator in the ansatz decomposition. Given the system with M spin orbitals and N electrons, it has been shown that the number-conserving encoding reduces the qubit resource requirements up to $O(N \log_2 M)$ scaling in both linear and nonlinear compression schemes [22,24]. These schemes can be collectively denoted as an encoder $\mathcal{E}(\cdot)$ that maps the JW basis $\{\vec{b}_i\} \in \{0, 1\}^M$ of $\hat{H}$ to $\{\mathcal{E}(\vec{b}_i)\}$ that preserves the N electron number-conserving basis states $\mathbf{b}_N$. Linear encoding further preserves the linear bitwise addition property to crucial for describing fermionic transitions. $\mathcal{E}(b_i \oplus b_j) = \mathcal{E}(b_i) \oplus \mathcal{E}(b_j)$. Although qubit-efficient type nonlinear compression achieves optimal qubit compression rate [22], the number of measurement bases will scale up to $O(M^N)$.

With linear encoding, operations can be avoided by performing all computations in the compressed subspace. The existence of an efficient decoder guarantee that map-encoded quantum states $\mathcal{E}(b_i)$ back to b_i at least for all number-conserving states $b_i \in \mathbf{b}_N$, where $\mathbf{b}_N$ is a bitstring representation for N -electron configurations. That is,

$$\mathcal{D}(\mathcal{E}(s_i)) = s_i. \quad (2)$$

This observation leads to an expectation estimation method—the FED algorithm [24]. This method also guarantees at most $O(M^4)$ measurement bases required for sampling each element of fermionic reduced density matrices. These expectation values can be measured, classically postprocessed, and computed in polynomial time. Meanwhile, operation in the

compressed subspace reduces the impact of quantum noise since (1) number-conserved bitstrings are more likely mapped to number-conserved bitstrings and (2) the qubit cost significantly reduces.

However, there are still two key challenges for the optimal linear encoding framework. First, encoding requires N -electron states to be injectively mapped to the encoded states. Utilizing random linear code, encoder generation will require the worst-case $O(M^N)$ checks. Although there are deterministic decoding methods, they either require access to a sparse oracle [27] or require a great number of qubits (Polylog) [28] to achieve an advantage. Second, decoding this problem is generally hard, the naive way is to store the lookup table to get optimal time complexity while trading-off with space complexity $O(M^N)$ to restore the data. While there exist efficient methods widely used in quantum error correction (QEC) [29–31], they are not specifically designed for this problem, since we only want to ensure the decoding of $\mathbf{b}_N$ is correct.

B. QSCI framework

QSCI algorithms—including the 77-qubit demonstration and the more recent TE-QSCI—show how near-term hardware can tackle quantum chemistry. They sample a trial state $|\Psi\rangle$ on a quantum device, retain the most probable bitstrings $\mathbf{x}_i$, and use these determinants to define a reduced configuration-interaction subspace in which the electronic Hamiltonian $\hat{H}$ is classically diagonalized to obtain ground-state energies.

To address the sensitivity of QSCI to the *a priori* choice of trial state (unclear), an adaptive variant has been introduced. ADAPT-QSCI [16] mirrors the ADAPT-VQE strategy. At each macroiteration it ranks all single- and double-excitation operators by the size of their QSCI energy gradient, appends the best candidate to the circuit, then runs a brief classical reoptimization before moving to the next iteration. Because the ansatz is grown on-the-fly, ADAPT-QSCI alleviates the “good initial state” problem, but this comes at the cost of many repeated QSCI and gradient evaluations, and no saving in qubit count since the full JW register is still required. Beside the adaptive approaches, LUCJ ansatz has also been proposed to tackle hardware noise by concentrating probability of $|\Psi\rangle$ in configurations near the ground-state support, enabling efficient sampling even on noisy devices [15].

Beyond the variational approaches, TE-QSCI represents an alternative to initial-state preparation problems. Time-evolution operator $e^{-i\hat{H}t}$ is employed to an initial state, such as Hartree-Fock or UCCSD, naturally generating a subspace rich in electronically excited configurations. This method avoids overheads from iterative optimization and the recurring use of quantum circuits that plagues the variational methods, while ensuring sampling in the physical effective subspace necessary for QSCI diagonalization.

Overall, the computational cost of all QSCI variants critically depends on the overlap between trial state $|\Psi\rangle$ with the ground-state wave function. Let P_G represent the ground-state projector, which indicates the probability of the operator projecting onto the low-energy subspace. If $|\Psi\rangle$ preferentially samples configurations P_G , the resulting subspace




FIG. 1. This diagram illustrates the key concept of Lossy-QSCI. RLE is preformed around the ground state to perform VQE. The resulting minima are sampled to gauge the region of the QSCI subspace.

diagonalization yields accurate energy estimates with samples that scale polynomially with system size.

Despite notable progress, all existing QSCI variants exhibit two fundamental limitations. First, they are qubit inefficient: Current methods encode the full Fock space using one qubit per fermionic mode, rapidly depleting quantum hardware resources as the system size increases. Second, they depend on complex ansatz constructions or iterative optimization routines—such as the multiterm LUCJ ansatz or repeated state preparation in TE-QSCI—which lead to high circuit depth and significant classical overhead. These challenges motivate our Lossy-QSCI approach. By leveraging Chemical-RLE, we compress the qubit register to $O(N \log_2 M)$ and employ a lightweight NN-FED decoder to recover and measure only the most relevant configurations. This direct compression in qubit space, coupled with efficient decoding, addresses both the qubit-scaling and ansatz-complexity bottlenecks of conventional QSCI frameworks.

III. LOSSY-QSCI

In this work, we introduce the Lossy-QSCI framework. It enhances quantum chemistry simulations by integrating a chemistry-inspired Chemical-RLE with a neural network FED, building on the strengths of fermionic encodings. Leveraging the number-conserving basis $\mathbf{b}_N$ from prior sections—where qubit requirements scale as $O(N \log_2 M)$ for M spin orbitals and N electrons—RLE compresses basis states $\{\vec{b}_i\}$ into a qubit-efficient representation $\{\mathcal{E}(\vec{b}_i)\}$, preserving a bijective mapping on important chemical subspace and avoiding the $O(M^N)$ measurement overhead of nonlinear encoding. To decode these compressed states efficiently, we propose an NN decoder trained on the structured patterns of electronic configurations, sidestepping the combinatorial space complexity of naive methods, and improving the performance of the FED. The use of Chemical-RLE naturally integrates QSCI with particle-conserving postselection by using the Q -qubit compressed states as trial wave functions $|\Psi\rangle \in F_2^Q$, sampling configurations to build subspaces for diagonalization. It can further enhance accuracy by adopting a lossy subspace as demonstrated in Fig. 1, concentrating probability on configurations near the ground state with even fewer

qubits, ensuring significant overlap with the true solution, and enabling precise energy estimation with polynomially many samples.

A. Chemistry-inspired Chemical-RLE

To integrate bit-flip error mitigation and ground-state concentration in QSCI with RLE, we provide a Chemical-RLE design that can work with QSCI to approach the full space ground state by sampling the subspace approximate ground state to improve usability. The original RLE begins its random search by choosing a Q within the bounded region $N \log_2 M < Q < 2N \log_2 M$. Then, the algorithm initializes the parity check matrix $\mathbf{G} = [I_Q | D]$ in the standard basis, where I_Q is a $Q \times Q$ identity matrix. In previous work [24], the $Q \times (M - Q)$ matrix D is randomly generated such that each column has an even Hamming weight equal to the value $Q/2$, denoted as $\text{even}(Q/2)$. However, such rules restrict the compressibility of qubit numbers, and the original RLE also requires a heavy subroutine check to ensure lossless configurations during compression. To address this, we propose a lossy compression scheme without a bijective checking process or with a biased bijective checking, ensuring the compressed Hamiltonian retains sufficient information. This approach is particularly well-suited for quantum chemistry, where ground states are often sparse in a well-defined basis and conform to specific symmetries or chemical rules (e.g., classical selected CI), making them ideal for lossy compression with minimal loss of critical information.

For the bit-flip error mitigation, encoding particle-number symmetry at the RLE stage removes much of the Hilbert space redundancy that the original QSCI exploited for postselection. In the uncompressed mapping, a single bit flip typically drives the register outside the selected N -electron sector, so postselection suppresses logical errors from $O(p)$ to $O(p^2)$ (see Appendix B of Ref. [6]). After compression, however, a single bit flip can map one valid encoded determinant to another, restoring an $O(p)$ error rate. We compensate for this loss by (1) discarding any decoded string that falls outside the complete active space (CAS)-biased domain ignored by the lossy map and (2) rejecting shots whose decoder likelihood or estimated energy lies beyond predefined thresholds.

For the preparation of a compact wave function with chemical knowledge, one of the selections is to use a classical heuristic [32] for chemical-inspired lossy RLE. By biased bijective checking, RLE will be enforced to generate a compressed Hamiltonian span by the classical heuristic. One can then further use QSCI to filter out more configurations for preparing even compact CI expansions. Another design is backed by the molecular orbitals (MOs) theory [26]. Since the chemical Hamiltonian's ground state normally has a biased distribution, configurations with lower energy molecular orbitals will contribute most to the ground-state wave function. With this, the CAS method [33] has been introduced to remove the high-energy MOs for better resource usage. Inspired by this fact, we can easily construct the Chemical-RLE that preserves the lower energy configurations while randomly losing the configurations that might have higher energy. By setting the computing basis with MOs order from lowest




FIG. 2. Result of different RLE strategies performance. The x axis represents the case number sorted by the final converged energy of QSCI, and the y axis denotes the energy difference from the exact ground state.

energy to highest energy as

$$|1_u 1_d 2_u 2_d \cdots M_u M_d\rangle,$$

where number is energy ordered MOs, u is spin-up orbitals, and d is spin-down orbitals. Combined with the original RLE design that can prepare a lossless map of N -electron configurations whose electrons are all in the front Q orbitals as

$$|1_u 1_d 2_u 2_d \cdots Q_u Q_d 000 \cdots 000\rangle,$$

the I_Q part in $\mathbf{G}$ can guarantee the lossless while electron belongs in $\mathbf{D}$ part. Then, we can further reduce the checking process by bias sampling the one-electron ($\{ \cdot \}_1$), two-electron ($\{ \cdot \}_2$), ..., N -electron excited state with different probabilities,

$$\begin{aligned} & \{|1_u \cdots Q_d 100 \cdots 000\rangle, \dots, |1_u \cdots Q_d 000 \cdots 001\rangle\}_1 \\ & \{|1_u \cdots Q_d 110 \cdots 000\rangle, \dots, |1_u \cdots Q_d 000 \cdots 011\rangle\}_2 \\ & \vdots \\ & \{|1_u \cdots Q_d 111 \cdots 100\rangle, \dots, |1_u \cdots Q_d 001 \cdots 111\rangle\}_N. \end{aligned}$$

This basis selection provides two advantages if the ground state is *sparse*. First, it can further reduce the number of qubits lower than the theoretical low bounds $\log_2 \binom{M}{N}$ while still being able to recover the ground-state support. Second, it can skip the costly injectivity checking process by not checking or only checking the important states from the predefined prior subspace.

To test the performance of Chemical-RLE, we compared three different strategies for solving 16 spin orbitals and six electrons C_2H_4 system with 12 qubits ($\log_2 \binom{16}{6}$); see Fig. 2. The first strategy, random encoding, involves randomly selecting the computing basis without considering the energy-sorted order of MOs and without any injectivity checking. The second strategy, chemical encoding, uses the energy-sorted basis order of MOs without injectivity checking. The third strategy, biased chemical encoding, employs the energy-sorted basis order of MOs, incorporating preselected configurations, which are considered more likely to be ground-state support for injectivity checking. In this test, we select the bias

injectivity checking states from the top 200 most frequent configurations of ground states. In a real case simulation scenario, preselected configurations could be obtained from prior knowledge, e.g., derived from classical approximations like Hartree-Fock, selected configuration interaction, or complete active space self-consistent field (CASSCF) methods, which identify chemically relevant subspaces and symmetries of the system. A random uniform noise is added to the exact ground state to prepare an approximate ground state with about 0.9 fidelity. This process is repeated 50 times for Lossy-QSCI, each performed 20 times, and the results are collected to prepare the effective Hamiltonian for classical diagonalization. This result shows that the RLE that guarantees bias injectivity could provide better lossy compression quality. It also shows that more prior knowledge of the given chemical system could lead to a better lossy compression rate since the bias injectivity can be more constrained to a smaller subspace.

B. Efficient neural network decoder, NN-FED

Beyond the injectivity check required by RLE, the main bottleneck in Lossy-QSCI is the decoding stage of the FED. Reconstructing an M -qubit occupation string from its Q -qubit codeword is difficult because the number of admissible states grows as $\mathcal{O}(M^N)$. A naïve lookup table demands exponential memory, whereas heuristic optimizers (e.g., simulated annealing or genetic algorithms) yield only approximate solutions and scale poorly.

To provide an efficient alternative, we introduce an NN-FED, trained with the workflow in Algorithm 1. We first fix a random linear encoder $\mathcal{E}(\cdot)$ and initialize a feed-forward network $\mathcal{D}_{\bar{\theta}}(\cdot)$ that maps a Q -qubit codeword to an M -qubit bitstring; the network contains $\mathcal{O}(MN)$ parameters in all test cases. During each training iteration, we draw k random N electron strings $\{s_i\}$, encode them, and let the network predict $\{\hat{s}_i\}$. The parameters $\bar{\theta}$ are updated by minimizing the binary cross-entropy (BCE) [34]:

$$L_{\text{BCE}} = -\frac{1}{k} \sum_{i=1}^k [s_i \cdot \log_2 \hat{s}_i + (1 - s_i) \cdot \log_2 (1 - \hat{s}_i)].$$

After a few epochs, the converged network $\mathcal{D}_{\bar{\theta}}$ becomes the fast decoder $\mathcal{D}_{\text{NN}}$ employed by the FED in Eq. (2).

NN-FED learns to decode a number-conserving state with only a few randomly generated samples in each iteration. For each update of parameters, the model will get the gradient of loss from the randomly generated dataset. Via gradient descent over the cross-entropy, it learns to decode the whole number-conserving states. This workflow bridges the original scalability guarantee of FED and enables the Parameterized Quantum Circuit (PQC) to be efficiently updated compared to the measurement methods that lacks a decoder. To benchmark the trained NN-FED's performance, we compare its decoding time for a single measurement shot to the genetic algorithm and simulated annealing. We also compare their average success rate over 1000 shots. The results are shown as Table II, where NN-FED surpasses both classical algorithms in terms of decoding time and accuracy on the scaling test cases.

ALGORITHM 1. NN-FED.

Require: Random linear encoder $\mathcal{E}(\cdot)$, random N -electron M -qubit bit string generator $\mathbf{S}_N^M \in F_2^M$;
Ensure: $\mathcal{D}_{\mathcal{NN}}(\mathcal{E}(s_i)) = s_i, \forall s_i \in \mathbf{S}_N^M$;
 Initialize a neural network
 $\mathcal{D}_{\tilde{\theta}}(\cdot) : F_2^Q \rightarrow F_2^M$;
for $t < t_{\max}$ **do**
 Randomly sample
 $\{s_i^{(t)}\} \in \mathbf{S}_N^M$;
 Prepare training data
 $\{b_i^{(t)}\} \leftarrow \mathcal{E}(\{s_i^{(t)}\}) \in F_2^Q$;
 $\{\hat{s}_i^{(t)}\} \leftarrow \mathcal{D}_{\tilde{\theta}}(b_i^{(t)}) \in F_2^M$;
 $\Delta\tilde{\theta} \leftarrow \nabla_{\tilde{\theta}^{(t)}} L_{\text{BCE}}(\{s_i\}, \{\hat{s}_i\})$
 $\tilde{\theta}^{(t+1)} \leftarrow \tilde{\theta}^{(t)} + \Delta\tilde{\theta}$;
end for
 $\mathcal{NN} \leftarrow \tilde{\theta}^{(t_{\max})}$;
return $\mathcal{D}_{\mathcal{NN}}$

C. Revisit the design of QSCI

Although both (1) compressing the register below $\log_2[M]$ and (2) relaxing injectivity checks inevitably discard certain basis states, a carefully designed Chemical-RLE eliminates primarily nonphysical configurations that do not contribute to the electronic Hamiltonian. If the surviving subspace still contains all determinants that dominate the ground state, the associated accuracy loss is negligible. Assuming this ground state is indeed sparse—an assumption supported by Fig. 1—we embed Chemical-RLE into the QSCI workflow to construct Lossy-QSCI.

The resulting algorithm (Fig. 3) introduces an encode-decode layer that (1) samples configurations in the compressed Hilbert space, (2) decodes them on the fly, and (3) iteratively updates the CI subspace used for classical diagonalization. This strategy simultaneously reduces the qubit count and enhances sampling efficiency, all while retaining chemical accuracy.

In the Lossy-QSCI workflow, we iteratively refine a compact CI subspace until the ground-state energy converges (Algorithm 2). Each outer iteration comprises five steps:

Compression. Generate a random Chemical-RLE that maps the original M -orbital register to a Q -qubit subspace, thereby significantly reducing qubit count.

Decoder training. Fit a problem-specific NN-FED that inverts the encoding and enables fast evaluation of observables in the full Fock space.

State preparation. Use a hybrid routine (e.g., VQE) to approximate the ground state for energy evaluations with the compressed Hamiltonian and the trained NN-FED.

Sampling and decoding. Draw the R most probable bit-strings from the prepared state and decode them to N -electron configurations in the original basis.

Collection. Augment the candidate set S_R with the decoded configurations to form S_{new} ; diagonalize $\hat{H}$ in this enlarged subspace. If the resulting energy is lower, set $S_R \leftarrow S_{\text{new}}$ and repeat.

The iterative “collect-and-refine” loop compensates for the information discarded by the lossy Chemical-RLE compression. After a sufficient number of cycles, all accepted configurations are merged into a final candidate set; the Hamiltonian is then diagonalized classically in this subspace to produce our best ground-state estimate. By fusing qubit-efficient encoding with targeted quantum sampling and lightweight classical postprocessing, the Lossy-QSCI workflow offers a resource-efficient path to chemically accurate energies on today’s noise-limited quantum hardware.

In summary, Lossy-QSCI incorporates QSCI with chemically aware Lossy-RLE to enhance the sample efficiency and NN-FED to provide efficient expectation measurement. As pointed out by Reinholdt *et al.* [20], (1) subspace methods such as QSCI require an exponential number of samples if the trial state has significant weight outside the optimal configuration manifold, which further results in (2) the classical diagonalization cost will grow prohibitively when the number of redundant selected determinants increases. Lossy-QSCI provides mitigation for both concerns by the following: First, the chemistry-inspired Chemical-RLE concentrates the support of the trial state in an $O(N \log_2 M)$ -qubit register; empirically, the required shots thus drop compared with uncompressed QSCI at the same chemical accuracy. Second, by design the lossy map and the $O(MN)$ -parameter NN-FED decoder, it can enable efficient VQE on classical heuristics subspace that can remove redundant selected determinants, so the exact diagonalization (ED) remains a negligible fraction of the total run-time. Hence, the two scalability bottlenecks highlighted in Reinholdt *et al.* [20] can thus be mitigated when a symmetry-aware lossy encoding and a lightweight decoder are combined with QSCI.

IV. NUMERICAL RESULT OF LOSSY-QSCI

Here, we test the lossy compression QSCI on different molecules with the active space selections labeled as




FIG. 3. This figure illustrates the Lossy-QSCI workflow. RLE compression is initialized for PQC optimization. The resulting minimum is sampled for contributing configurations. This set of configurations is embedded in the candidate set S_{new} if the minimum shows better convergence to the previous optimization trial.

TABLE II. This figure illustrates the NN-FED benchmarks. M is the number of spin orbitals and Q is the number of qubits.

Algorithms	4-Electron	(M, Q)	(30, 23)	(40, 27)	(50, 30)	(60, 32)	(70, 34)
NN-FED		Training time (min)	2.56	5.09	6.97	11.50	14.22
		Decoding times (s)	0.0022	0.0022	0.0032	0.0076	0.0072
		Accuracy	100.0%	99.9%	100.0%	99.8%	99.3%
Genetic algorithm		Decoding times (s)	3.6693	3.9568	3.6679	3.6423	3.7163
		Accuracy	90.4%	83.9%	76.0%	71.8%	65.0%
Simulated annealing		Decoding times (s)	0.0845	0.1017	0.1112	0.1187	0.1206
		Accuracy	40.3%	31.5%	24.3%	18.3%	16.1%

($N_{\text{Spin Orbitals}}$, $N_{\text{Electrons}}$) to benchmark the performance of Lossy-QSCI, where $N_{\text{Spin Orbitals}}$ is the number of spin orbitals and $N_{\text{Electrons}}$ the number of electrons. The ground-state support for all cases is constructed by iteratively sampling the R most frequent configurations from iteratively prepared approximate ground states (see Appendix for details). We first examine the C_2 molecule with different lossy compression rates represented to study how the number of qubits scales with the energy in Lossy-QSCI. Then, we test the framework on LiH with the hardware efficient ansatz (HEA) and varying QSCI basis size. Finally, we perform a noisy simulation with a bit-flip error channel on the H_2 system with both the typical QSCI and Lossy-QSCI frameworks to investigate the quality of convergence in realistic circumstances.

A. C_2 molecule

We begin by classically diagonalizing the full 6–31G Hamiltonian for C_2 [35] at a bond length of 0.9–3 Å. The resulting eigenvector is taken as the exact ground state, $|\psi_g\rangle$. To emulate the imperfect trial states that a realistic, noise-affected VQE optimization would deliver, we form a surrogate

variational state

$$|\psi_n\rangle = \frac{|\psi_g\rangle + |\psi_u\rangle}{\| |\psi_g\rangle + |\psi_u\rangle \|_2},$$

where $|\psi_u\rangle$ is a random vector drawn from the same N -electron, M -orbital sector. This construction yields an overlap $|\langle \psi_n | \psi_g \rangle|^2 \simeq 0.85$, in line with fidelities reported for converged VQE circuits of comparable size. Hence, $|\psi_n\rangle$ is adopted as the common input state for all QSCI variants discussed below, providing a VQE-level baseline from which the benefits of Lossy-QSCI can be assessed. Finally, we select the top 50 most probable states (based on their highest occurrence probabilities) to define the compression ratio R .

Figure 4 illustrates the performance of Lossy-QSCI as a function of the number of qubits. We benchmark our results against reference energies obtained via ED in smaller active spaces, shifted by 1 kcal/mol (0.001 59 hartree) to denote chemical accuracy thresholds: (1) ED in the (16, 4) active space, marked as “chemical accuracy (16, 4)” (solid blue line); and (2) ED in the (20, 4) active space, marked as “chemical accuracy (20, 4)” (solid gray line). Discrete markers of

ALGORITHM 2. Lossy-QSCI workflow.

Require: Hamiltonian H , initial state $|\Psi\rangle$
Ensure: Ground-state energy E_{best} and configuration subspace S_R

- 1: **repeat**
- 2: Generate lossy RLE $E(\cdot)$ to obtain compressed Hamiltonian H_{comp}
- 3: Prepare trial wave function $|\Psi\rangle$ using a PQC under H_{comp}
- 4: Train NN-FED to decode compressed bitstrings to N -electron configurations
- 5: **repeat**
- 6: Sample bitstrings from $|\Psi\rangle$
- 7: Decode samples using NN-FED and get energy
- 8: Update PQC parameters if required
- 9: **until** convergence criteria met
- 10: Select R -most frequent bitstrings to extend subspace S_{new}
- 11: Diagonalize H in $S_R \leftarrow S_R + S_{\text{new}}$ to get energy E_{new}
- 12: **if** $E_{\text{new}} < E_{\text{best}}$ **then**
- 13: $S_R \leftarrow S_R + S_{\text{new}}$
- 14: $E_{\text{best}} \leftarrow E_{\text{new}}$
- 15: **end if**
- 16: **until** convergence criteria met
- 17: **return** E_{best} , S




FIG. 4. The C_2 molecule Lossy-QSCI results with approximate ground state prepared by adding random noise to the exact ground state. The x axis represents the bond length in angstroms, ranging from 0.9 to 3.0, while the y axis shows the energy in Hartree units, spanning from approximately -75.050 to -75.050 . The subfigure is the energy difference from the exact ground-state energy of the active space (20, 4).




FIG. 5. The LiH molecule Lossy-QSCI results with VQE. The x axis represents the number of QSCI basis states (R) and the y axis shows the computed energy in hartree (Ha).

different colors and shapes denote the Lossy-QSCI results for qubit counts ranging from 10 to 16: red triangles for 10 qubits, green crosses for 11 qubits, purple squares for 12 qubits, orange pentagon for 15 qubits, and black circles for 16 qubits.

For lower qubit counts (10 and 11), Lossy-QSCI yields energies below the (16,4) chemical accuracy threshold at shorter bond lengths (indicating superior performance in compact regimes), but exceeds it as the bond length increases and the system begins to dissociate. This trend highlights the varying computational complexity across the dissociation curve and underscores the potential benefits of adaptively selecting qubit counts for compression in different molecular configurations. As the qubit count increases to 12, 15, and 16, the Lossy-QSCI energies progressively approach and often surpass the (20,4) chemical accuracy benchmark, particularly at intermediate and longer bond lengths.

Overall, these results demonstrate that Lossy-QSCI, by compressing a larger active space into a reduced qubit representation, can achieve greater accuracy than ED in equivalently sized or even larger active spaces using the same or fewer qubits. This capability positions Lossy-QSCI as a resource-efficient strategy for quantum simulations, enabling closer approximations to exact QSCI results while mitigating constraints on qubit resources and circuit depth.

B. LiH molecule

In the study of LiH molecule at a bond length of 2.5 Å with the STO-3G basis set [36], the candidate configuration sets, essential for a single round Chemical-RLE, is constructed from the five most frequently sampled configurations of the approximate ground states prepared via VQE. For the optimization, we employ L-BFGS-B optimizer [37].

Figure 5 illustrates the result of the Lossy-QSCI method encoded in five qubits. The two dashed lines, colored in black and blue, provide references for the chemical accuracy in the (6, 2) and (10, 2) active spaces, respectively. The red line depicts the energy convergence of the Lossy-QSCI method as




FIG. 6. This figure illustrates the H_2 molecule noisy simulation. A lossy compressed QSCI simulation not only significantly reduces the qubit cost, but also enhances the convergence with respect to increasing QSCI basis.

we increase R , starting at slightly above the chemical accuracy of (6, 2) active space for $R = 4$, converges quickly to the chemical accuracy of (10, 2) active space, and stabilizes as $R \geq 12$. The green line represents the collective optimized energy of VQE on randomly compressed Hamiltonians for five qubits.

The results demonstrate that iterative sampling of CI states from different compressed subspaces (Fig. 3) can significantly enhance the performance of Lossy-QSCI. Specifically, the Lossy-QSCI method achieves the (10, 2) chemical accuracy with as few as 12 basis states, outperforming all collective VQE results that lie above -7.8 . Improved convergence indicates a potential of Lossy-QSCI to deliver more accurate ground-state energies with fewer basis states. This approach achieves the required accuracy with fewer qubit counts compared to conventional QSCI approaches through iterative VQE, thereby improving resource efficiency in quantum simulations, and boosts the simulation power of noisy hardware.

C. H_2 molecule

We investigate the performance of noisy VQE optimization on the H_2 molecule with the 6-31G basis set at bond length 4 Å. We employed a bit-flip error model with single-qubit error probability $p_{\text{gate}1} = 0.1$, two-qubit error composed as the tensor product of two single-qubit errors, qubit resetting error—erroneous preparation of $|1\rangle$ instead of $|0\rangle$ — $p_{\text{reset}} = 0.1$, and measurements readout error probability $p_{\text{meas}} = 0.1$. The noise model is implemented as Pauli X bit-flip channel via the Qiskit NoiseModel object. Within this setting, we benchmark the Lossy-QSCI algorithm implemented with four qubits and a HEA ansatz comprised of 12 CNOT gates and 20 Ry gates against QSCI methods implemented with 8 qubits and a HEA ansatz made up of 14 CNOT gates and 24 Ry gates.

Figure 6 illustrates the performance of the Lossy-QSCI and QSCI under the influence of bit-flip error. For the Lossy-QSCI method, the candidate configuration set was constructed by

sampling the 10 most frequent configurations from 20 subspaces of ground states (due to random compression) prepared by VQE. On the other hand, the QSCI method's candidate configuration set was constructed by sampling the 10 most frequent configurations from 50 subspaces of ground states prepared by VQE. Since Lossy-QSCI with Chemical-RLE is only optimized and sampled on the number-conserving subspace, we postprocess the QSCI candidate configuration set only containing the number-conserving state to improve the construction of the effective Hamiltonian.

Both results demonstrate that enforcing number-conserving constraints can provide error robustness for both QSCI and noisy QSCI and enable them to recover from noisy VQE results, as depicted by the red and blue lines. Meanwhile, the Lossy-QSCI method (blue) outperforms QSCI (red), achieving energy below chemical accuracy with as few as 12 basis states compared to the 15 basis states in QSCI. The total sampled configurations from Lossy-QSCI are 200 samples of four-qubit configurations compared to the 500 samples of eight-qubit configurations for QSCI. This approach highlights the potential of Lossy-QSCI to deliver better searching efficiency. Robustness against bit-flip error shows the potential for making it far more suitable than QSCI in near-term devices.

V. CONCLUSION

In this work, we have introduced the Lossy-QSCI framework, enhanced by the NN-FED and the Chemical-RLE, to address the challenges of qubit-efficient quantum chemistry simulations on near-term quantum devices. By leveraging the number-conserving properties of fermionic encodings, our framework achieves significant qubit reduction scaling as $O(N \log_2 M)$ for M spin orbitals and N electrons, while mitigating the measurement overhead typically associated with compressed Hamiltonians. Furthermore, the comparison of RLE strategies for the C_2H_4 system underscores the advantage of biased chemical encoding, which leverages prior knowledge to improve compression quality. This approach not only allows qubit compression below the theoretical lower bound of $\log_2(\binom{M}{N})$ but also minimizes the computational overhead of qubit encoding generation by focusing only on chemically relevant configurations. Inspired by MO theory and the CAS method, our lossy encoding ensures that critical ground-state information is preserved, making it well suited for quantum chemistry simulations. On the other hand, the integration of NN decoder, as detailed in Algorithm 1, enables fast high-precision bitstring decoding of compressed quantum states as shown in Table II, outperforming traditional classical decoders like genetic algorithms and simulated annealing.

Our numerical experiments on the C_2 , LiH, and H_2 molecules demonstrate the efficacy of Lossy-QSCI. For the C_2 molecule, Fig. 4 shows that Lossy-QSCI achieves energies closer to the exact diagonalization energy curve results of the (20, 4) active space as the number of qubits increases from 10 to 16, where 16 qubits results surpass the exact diagonalization benchmark for the (16, 4) active space throughout the length of the bond. This highlights the ability of Lossy-QSCI

to map larger active spaces to smaller qubit representations without greatly sacrificing accuracy. Similarly, the LiH molecule illustrates that Lossy-QSCI with five qubits reaches the chemical accuracy of the (10, 2) active space with as few as 12 basis states, outperforming the best collected VQE result. The iterative CI states from compressed subspaces, combined with statistical ranking, enable this enhanced performance, as detailed in the improved workflow. Finally, the H_2 molecule case shows that the number-conserving subspace prepared by Chemical-RLE naturally protects the sample of nonphysical state and therefore provides error robustness compared to the original QSCI with number-conserving post-processing.

Overall, our framework demonstrates that Lossy-QSCI can deliver accurate ground-state energies with fewer basis states and highlights the potential of near-term devices with fewer qubits and circuit depth to achieve the same accuracy as larger devices through iterative VQE on different random subspaces with classical postprocessing. This balance between computational efficiency and accuracy paves the way for resource-efficient ground-state simulations on near-term hardware, offering a promising pathway to explore quantum advantages in solving complex many-body problems. While this work demonstrates an improvement of time and space resources for QSCI simulation via efficient encoding-decoding to a smaller subspace, it remains open to characterize the qubit-to-gate trade-offs in the encoded subspace, which will allow us to encode arbitrary QSCI algorithms that leverage both the gate advantages from the unencoded space and subspace simulation. Beyond it, further research is required to devise provably scalable, symmetry-aware encoders and replace the data-driven decoder algorithms with formal performance bounds. In this work, we present Lossy-QSCI as an intermediate step that enhances near-term resource usage and informs the design of future, fault-tolerant workflows rather than a universal, black-box solution. In future work, we will explore the possibility of an encoder-decoder framework that generalizes the framework of Lossy-QSCI.

ACKNOWLEDGMENT

The authors thank Dr. Mathias Mikkelsen, Dr. Masaya Kohda, Dr. Keita Kanno, and Dr. Ming-Zhi Chung (QunaSys Inc.) for their insightful discussions and constructive feedback, which substantially improved the manuscript.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX: NUMERICAL DETAILS

For all numerical simulations, the ground-state supports are constructed by iteratively sampling from iteratively prepared approximate ground states using either VQE or an approximate ground state that aims to mimic the VQE solution. In each sampling, we will go through the same steps for constructing the ground-state support: (1) having the sample




FIG. 7. Configuration distribution for the ground state of the H_2 molecule using Lossy-QSCI. Red bars represent the support of the exact ground-state probabilities [thresholded at 10^{-3} , heights proportional to $3.5 \times \log_2(p_i \times 2000)$]. Colored stacked bars show frequency counts of configurations from 20 subspaces (200 total samples, 10 per subspace), with each color indicating a subspace (samples 0–19). The x axis lists all 28 basis states. The y axis is the sampling frequency.

probabilities of all sampled configurations, (2) picking the R most frequent configurations (R -largest sample probabilities) as candidates of ground-state support, (3) postselecting the number-conserving states, and (4) add configurations to the existing ground-state support and do the exact diagonalization to see if the energy will decrease; if not, keep the previous ground-state support.

For illustrating the sampling process, in Fig. 7, we take the Lossy-QSCI on the H_2 molecule as an example. The computational basis consists of 28 bitstring states, each of which represents a possible configuration in the number-conserving computing basis of the H_2 molecule. These states are listed on the x axis, ordered by their binary representation (from ‘00000110’ to ‘11000000’). The red bars represent the support of the exact ground-state probability distribution, derived from the squared amplitudes of the wave function ($|\psi_i|^2$). To focus on significant contributions, probabilities below 10^{-3} are thresholded to zero, and the heights of the red bars are scaled as $3.5 \times \log_2(p_i \times 2000)$, where p_i are the probabili-

ties for scaling to align with the data frequency scale. This logarithmic scaling enhances visibility of the ground-state support. The colored stacked bars represent the frequency counts of configurations sampled from 20 subspaces, with each subspace contributing 10 samples, resulting in 200 total samples. Each subspace is assigned a unique color from the “viridis” colormap, labeled as “Sample 0” to “Sample 19” in the legend. The counts are computed by mapping each sampled bitstring to its index in the basis, excluding any bitstrings not present in the 28-state basis to ensure number conserving. The bars for each subspace are stacked over the red bars.

The complete code, including the dataset, basis states, and plotting parameters, is available in a public repository at Ref. [38]. This repository includes the exact bitstring data, the ground-state wave function amplitudes, ensuring full reproducibility. The Lossy-QSCI parameters, such as the subspace definitions and sampling procedure, are detailed in the repository’s documentation, aligned with the methodology described in the main text.

- [1] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, *Science* **309**, 1704 (2005).
- [2] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [3] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [4] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [5] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
- [6] K. Kanno, M. Kohda, R. Imai, S. Koh, K. Mitarai, W. Mizukami, and Y. O. Nakagawa, Quantum-selected configuration interaction: Classical diagonalization of Hamiltonians in subspaces selected by quantum computers, *arXiv:2302.11320*.
- [7] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, Expressive power of parametrized quantum circuits, *Phys. Rev. Res.* **2**, 033125 (2020).

- [8] S. Kanasugi, S. Tsutsui, Y. O. Nakagawa, K. Maruyama, H. Oshima, and S. Sato, Computation of Green's function by local variational quantum compilation, *Phys. Rev. Res.* **5**, 033070 (2023).
- [9] K. Mizuta, Y. O. Nakagawa, K. Mitarai, and K. Fujii, Local variational quantum compilation of a large-scale Hamiltonian dynamics, *PRX Quantum* **3**, 040302 (2022).
- [10] C.-T. Li, T. T. Ong, L. Wang, M.-C. Hsu, H. Lin, and M.-H. Hsieh, A Gell-Mann and low theorem perspective on quantum computing: New paradigm for designing quantum algorithm, [arXiv:2311.18271](https://arxiv.org/abs/2311.18271).
- [11] Y. Du, T. Huang, S. You, M.-H. Hsieh, and D. Tao, Quantum circuit architecture search for variational quantum algorithms, *npj Quantum Inf.* **8**, 62 (2022).
- [12] W. Inoue, K. Aoyama, Y. Teranishi, K. Kanno, Y. O. Nakagawa, and K. Mitarai, Almost optimal measurement scheduling of molecular Hamiltonian via finite projective plane, *Phys. Rev. Res.* **6**, 013096 (2024).
- [13] K. Heya, K. M. Nakanishi, K. Mitarai, Z. Yan, K. Zuo, Y. Suzuki, T. Sugiyama, S. Tamate, Y. Tabuchi, K. Fujii, and Y. Nakamura, Subspace variational quantum simulator, *Phys. Rev. Res.* **5**, 023078 (2023).
- [14] M. Motta, K. J. Sung, K. B. Whaley, M. Head-Gordon, and J. Shee, Bridging physical intuition and hardware efficiency for correlated electronic states: The local unitary cluster Jastrow ansatz for electronic structure, *Chem. Sci.* **14**, 11213 (2023).
- [15] J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, I. Sitdikov, R.-Y. Sun, K. J. Sung, M. Takita, M. C. Tran, S. Yunoki, and A. Mezzacapo, Chemistry beyond exact solutions on a quantum-centric supercomputer, *Sci. Adv.* **11**, eadu9991 (2025).
- [16] Y. O. Nakagawa, M. Kamoshita, W. Mizukami, S. Sudo, and Y.-Y. Ohnishi, Adapt-QSCI: Adaptive construction of an input state for quantum-selected configuration interaction, *J. Chem. Theory Comput.* **20**, 10817 (2024).
- [17] M. Mikkelsen and Y. O. Nakagawa, Quantum-selected configuration interaction with time-evolved state, *Phys. Rev. Res.* **7**, 043043 (2025).
- [18] K. Sugisaki, S. Kanno, T. Itoko, R. Sakuma, and N. Yamamoto, Hamiltonian simulation-based quantum-selected configuration interaction for large-scale electronic structure calculations with a quantum computer, *Phys. Chem. Chem. Phys.* **27**, 20869 (2025).
- [19] J. Yu, J. R. Moreno, J. T. Iosue, L. Bertels, D. Claudino, B. Fuller, P. Groszkowski, T. S. Humble, P. Jurcevic, W. Kirby, T. A. Maier, M. Motta, B. Pokharel, A. Seif, A. Shehata, K. J. Sung, M. C. Tran, V. Tripathi, A. Mezzacapo, and K. Sharma, Quantum-centric algorithm for sample-based Krylov diagonalization, [arXiv:2501.09702](https://arxiv.org/abs/2501.09702).
- [20] P. Reinholdt, K. M. Ziems, E. R. Kjellgren, S. Coriani, S. P. A. Sauer, and J. Kongsted, Fundamental limitations in sample-based quantum diagonalization methods, *J. Chem. Theory Comput.* **21**, 6811 (2025).
- [21] P. Jordan and E. Wigner, Über das paulische Äquivalenzverbot, *Z. Phys.* **47**, 631 (1928).
- [22] Y. Shee, P.-K. Tsai, C.-L. Hong, H.-C. Cheng, and H.-S. Goan, Qubit-efficient encoding scheme for quantum simulations of electronic structure, *Phys. Rev. Res.* **4**, 023154 (2022).
- [23] D. Yoffe, N. Entin, A. Natan, and A. Makmal, A qubit-efficient variational selected configuration-interaction method, *Quantum Sci. Technol.* **10**, 015020 (2025).
- [24] M. H. Cheng, Y.-C. Chen, Q. Wang, V. Bartsch, M. S. Kim, A. Hu, and M.-H. Hsieh, Optimal particle-conserved linear encoding for practical fermionic simulation, *Phys. Rev. Res.* **7**, 033280 (2025).
- [25] R. Babbush, D. W. Berry, Y. R. Sanders, I. D. Kivlichan, A. Scherer, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, Exponentially more precise quantum simulation of fermions in the configuration interaction representation, *Quantum Sci. Technol.* **3**, 015006 (2018).
- [26] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, Mineola, NY, 1996).
- [27] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [28] A. M. Childs, R. Kothari, and R. D. Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, *SIAM J. Comput.* **46**, 1920 (2017).
- [29] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Phys. Rev. A* **52**, R2493 (1995).
- [30] A. M. Steane, Error correcting codes in quantum theory, *Phys. Rev. Lett.* **77**, 793 (1996).
- [31] D. Gottesman, Stabilizer codes and quantum error correction, [arXiv:quant-ph/9705052](https://arxiv.org/abs/quant-ph/9705052).
- [32] N. M. Tubman, C. D. Freeman, D. S. Levine, D. Hait, M. Head-Gordon, and K. B. Whaley, Modern approaches to exact diagonalization and selected configuration interaction with the adaptive sampling CI method, *J. Chem. Theory Comput.* **16**, 2139 (2020).
- [33] B. O. Roos, P. R. Taylor, and P. E. Sigbahn, A complete active space SCF method (CASSCF) using a density matrix formulated super-CI approach, *Chem. Phys.* **48**, 157 (1980).
- [34] C. M. Bishop, *Pattern Recognition and Machine Learning, Information Science and Statistics* (Springer, New York, 2006).
- [35] W. J. Hehre, R. Ditchfield, and J. A. Pople, Self-consistent molecular orbital methods. XII. Further extensions of Gaussian-type basis sets for use in molecular orbital studies of organic molecules, *J. Chem. Phys.* **56**, 2257 (1972).
- [36] W. J. Hehre, R. F. Stewart, and J. A. Pople, Self-consistent molecular-orbital methods. I. Use of Gaussian expansions of Slater-type atomic orbitals, *J. Chem. Phys.* **51**, 2657 (1969).
- [37] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, A limited memory algorithm for bound constrained optimization, *SIAM J. Sci. Comput.* **16**, 1190 (1995).
- [38] Lossy-QSCI: Lossy Quantum Selected Configuration, available at Interaction <https://github.com/ALS15204/lossy-qsci>.