Variational Manifold Engineering in Quantum Chemistry: Structural Challenges in Ansatz Design

Abhoy Kole[Δ], Julie Maria Raju[Ш], Kamalika Datta[Δ,Ш], Rolf Drechsler[Δ,Ш]
[Δ]*Cyber-Physical Systems, DFKI, Bremen, Germany*
[Ш]*Institute of Computer Science, University of Bremen, Bremen, Germany*
abhoy.kole@dfki.de, {jraju, datta, drechsler}@uni-bremen.de

***Abstract*—Variational quantum algorithms have emerged as a leading approach for molecular ground-state energy estimation on near-term quantum devices. Central to their performance is the choice of ansatz, which defines the accessible variational manifold and directly influences expressibility, trainability, symmetry preservation, and resource requirements. Despite extensive development of fixed chemistry-inspired ansatz, hardware-efficient constructions, and adaptive Hamiltonian-driven approaches, a unified structural understanding of their design principles and limitations remains incomplete. This work presents a comparative structural analysis of major ansatz paradigms from the perspective of constrained variational manifold engineering. We examine how each construction strategy encodes physical symmetries, balances expressibility and trainability, scales in parameter and measurement complexity, and responds to noise and hardware-induced constraints.**

***Index Terms*—VQE, UCC, Hardware Efficient Ansatz, Adaptive Ansatz**

I. Introduction

Quantum computing has evolved from a theoretical concept to a physically realizable paradigm with the potential to solve computationally intractable problems. The development of quantum algorithms, such as Shor's factoring algorithm [1], and their proof-of-concept demonstrations on emerging quantum hardware have significantly accelerated research in both quantum devices and software ecosystems. These advances aim to position quantum computing as a potential accelerator for high-performance computing (HPC) [2]. While substantial progress has been made in scaling quantum hardware [3], [4], the realization of fault-tolerant quantum computation for practically relevant problems remains an open challenge [5].

In this context, variational quantum algorithms (VQAs) have emerged as a promising approach for leveraging noisy intermediate-scale quantum (NISQ) devices. VQAs utilize a hybrid quantum-classical (HQC) framework, where a parameterized quantum circuit prepares trial states and a classical optimizer iteratively updates its parameters to minimize a cost function [6]–[8]. This hybrid formulation reduces circuit depth requirements and enables computations within the limited coherence times of current quantum hardware, thereby avoiding the immediate need for full error correction. Furthermore, recent studies have highlighted the rapid growth of entanglement in quantum systems and the resulting limitations of classical simulation methods, reinforcing the need for efficient variational representations of quantum states [9].

Among VQAs, the Variational Quantum Eigensolver (VQE) has been widely studied for estimating molecular ground-state energies [7]. Despite significant progress in classical quantum chemistry, accurately describing strongly correlated systems, such as transition-metal complexes and reaction dynamics, remains challenging [10]–[12]. VQE addresses this challenge by partitioning the computational workload between quantum and classical resources, where the quantum processor prepares parameterized trial states (ansätze), and the classical optimizer minimizes the expectation value of the molecular Hamiltonian.

A central component of VQE is the design of the ansatz, which determines the accessible variational manifold and directly impacts expressibility, trainability, symmetry preservation, and resource efficiency. Existing ansatz construction approaches can be broadly categorized into: (i) chemistry-inspired ansätze based on traditional methods such as unitary coupled cluster [13], (ii) hardware-efficient ansätze designed to align with device constraints [14], and (iii) adaptive or Hamiltonian-driven ansätze that iteratively build problem-specific circuits [15]. While each of these approaches offers distinct advantages, they also exhibit inherent limitations related to optimization landscape complexity, sensitivity to noise, scalability, measurement overhead, and the preservation of physical symmetries.

In this work, we present a structural analysis of ansatz design for VQE from the perspective of constrained variational manifold engineering. We systematically examine the construction principles, advantages, and limitations of major ansatz paradigms, and highlight key research challenges associated with their practical deployment on NISQ hardware.

The organization of the paper is as follows. Section II introduces the necessary background on quantum computing, the VQE framework, and existing approaches for efficient quantum simulation of molecular systems. Section III presents the variational manifold perspective of ansatz design and provides a comparative analysis of different ansatz families. Finally, Section IV concludes the paper.

This work was partly funded by the Federal Ministry of Research, Technology and Space (BMFTR) through the EASEPROFIT project (grant no. 16KIS2127) and the EQeS project (grant no. 13N17378), and by the German Research Foundation (DFG) through the CONAD-QC project (grant no. 559888852). This research was conducted within the scope of the DFG Priority Programme 2514: Quantum Software, Algorithms and Systems.

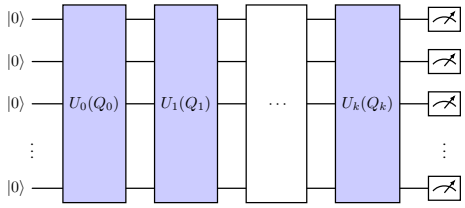

Fig. 1: Quantum circuit diagram illustrating the sequence of gate operations on qubits.

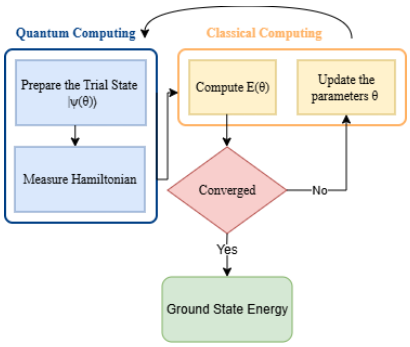

Fig. 2: Workflow of VQE.

II. Background

A. Basics of Quantum Computing

Quantum information is stored in the states of qubits, which are two-level quantum systems. In contrast to classical bits, a qubit can exist in a coherent superposition of basis states, and is generally written as $\alpha|0\rangle + \beta|1\rangle$, where $\alpha, \beta \in \mathbb{C}$ satisfy the normalization condition $|\alpha|^2 + |\beta|^2 = 1$. For a system of n qubits, the joint state is described by a unit vector in the Hilbert space $\mathcal{H} = \mathbb{C}^{2^n}$, formed via tensor products of individual qubit states.

Information is extracted from a quantum system through measurement, which yields probabilistic outcomes according to Born's rule [16]. The computational basis consists of tensor products of $|0\rangle$ and $|1\rangle$ states, corresponding to the eigenbasis of the Pauli-Z operator. Other Pauli operators, such as X and Y, define alternative measurement bases that are frequently used in quantum algorithms.

A distinctive feature of quantum systems is entanglement, wherein the state of a composite system cannot be decomposed into independent single-qubit states. For instance, the Bell state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ exhibits maximal entanglement between two qubits. Such correlations are central to the potential computational advantages of quantum devices, particularly in applications involving strongly correlated systems in condensed matter physics and quantum chemistry.

Quantum circuits are implemented through sequences of unitary transformations, commonly referred to as quantum gates, acting on subsets of qubits. Starting from an initial state $|0\rangle^{\otimes n}$, a computation is described by

$$U_k(Q_k)\, U_{k-1}(Q_{k-1}) \cdots U_0(Q_0)\, |0\rangle^{\otimes n}\,, \tag{1}$$

where each $U_k(Q_k)$ denotes a unitary operation applied to a specified subset of qubits Q_k at step k. Fig. 1 shows a diagrammatic representation of the quantum circuit.

B. Variational Quantum Eigensolver (VQE)

The Variational Quantum Eigensolver (VQE) is a hybrid quantum–classical algorithm designed to estimate the ground state energy of a Hamiltonian H, such as those arising in molecular systems [7]. The ground state energy is a fundamental quantity in quantum chemistry, as it enables the computation of molecular properties such as reaction rates, binding energies, and reaction pathways. The method is based on the variational principle, which guarantees that the expectation value of H with respect to a trial state provides an upper bound to the true ground state energy E_g.

In VQE, a parameterized quantum circuit prepares a trial state of the form

$$|\psi(\theta)\rangle = U(\theta)\,|\phi_0\rangle\,, \tag{2}$$

where $|\phi_0\rangle$ is a reference state and $U(\theta)$ is a parameterized unitary operator, commonly referred to as the ansatz. In practice, the Hartree–Fock (HF) state is commonly used as the reference state. The corresponding energy expectation value is given by

$$E(\theta) = \langle\phi_0|\, U^\dagger(\theta)\, H\, U(\theta)\,|\phi_0\rangle \geqslant E_g, \tag{3}$$

where the Hamiltonian is typically decomposed as $H = \sum_i H_i$, allowing each term to be estimated individually on a quantum device.

A classical optimizer is then employed to iteratively update the parameters θ so as to minimize $E(\theta)$. This hybrid quantum–classical loop continues until convergence as shown in Fig. 2. Since the computation is iterative, inaccuracies in early stages can propagate through subsequent updates; therefore, the choice of an appropriate initial state and parameterization is critical for obtaining accurate results. VQE significantly reduces coherence requirements by offloading the optimization procedure to a classical processor, making it well-suited for near-term quantum devices.

C. Prior Works

The use of quantum systems for simulating physical processes was first formalized by Lloyd [17], establishing the theoretical foundation for quantum simulation. Early quantum algorithms, including Trotter-based methods and phase estimation, provided rigorous approaches for solving quantum chemistry problems [18], [19], but required deep circuits and fault-tolerant hardware, limiting their applicability on near-term devices [20]. This foundational result further motivated extensive research in quantum algorithms for simulating physical and chemical systems, particularly where classical methods face exponential scaling limitations.

To this end, the emergence of VQE [6], [7] marked a shift toward hybrid quantum-classical approaches suitable for NISQ systems. Subsequent developments have focused extensively

on ansatz construction, which plays a critical role in determining the efficiency and accuracy of VQE. Existing ansatz design strategies can be broadly categorized into chemistry-inspired methods, hardware-efficient constructions, and adaptive or Hamiltonian-driven approaches.

Chemistry-inspired ansätze, such as the Unitary Coupled Cluster (UCC) family [21]–[23], retain the structure of electronic excitations and offer systematically improvable accuracy, but typically lead to deep quantum circuits due to the large number of excitation operators and their decomposition into elementary gates. In contrast, hardware-efficient ansätze (HEA) [24]–[26] are tailored to the native gate set and connectivity of quantum processors, enabling shallow circuits suitable for NISQ devices; however, they may suffer from reduced physical interpretability, symmetry breaking, and trainability issues. To bridge this gap, adaptive strategies [27]–[29] construct ansätze iteratively based on the problem Hamiltonian, yielding compact and accurate representations. However, these approaches introduce additional measurement and optimization overhead due to repeated operator selection and gradient evaluation.

Recent work has demonstrated large-scale integrated quantum–classical simulation pipelines [30]. In parallel, quantum-centric workflows extending beyond classical exact diagonalization limits have been demonstrated [31], including simulations of N_2 dissociation and the properties of [2Fe–2S] and [4Fe–4S] clusters using circuits with up to 77 qubits and 10,570 gates. More recently, such workflows have been extended to complex biological systems [32]. Despite these advances, the design of efficient and scalable ansätze remains a fundamental challenge due to issues such as barren plateaus [33] and limited expressibility [34], motivating the need for a systematic structural analysis of existing ansatz construction approaches.

III. Ansatz Design Approach

A. Unifying Variational Manifold Perspective

The set of states $\{|\psi(\boldsymbol{\theta})\rangle\}$, generated by a parameterized ansatz $U(\boldsymbol{\theta})$ acting on a reference state $|\phi_0\rangle$, defines a parameterized submanifold $\mathcal{M}$ embedded in the full Hilbert space $\mathcal{H}$, i.e.,

$$\mathcal{M} = \{U(\boldsymbol{\theta})\,|\phi_0\rangle \mid \boldsymbol{\theta} \in \mathbb{R}^p\} \subseteq \mathcal{H}, \tag{4}$$

where p indicates the dimension of the parameter space, i.e., $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_p)$. The effectiveness of the variational ansatz is determined by the geometric and structural properties of this manifold.

Expressibility refers to the ability of the ansatz manifold $\mathcal{M}$ to approximate arbitrary states in the Hilbert space. Formally, for any target state $|\psi\rangle \in \mathcal{H}$, there exists a parameter vector $\boldsymbol{\theta} \in \mathbb{R}^p$ such that

$$\min_{\boldsymbol{\theta}\in\mathbb{R}^p} \| \,|\psi\rangle - U(\boldsymbol{\theta})\,|\phi_0\rangle\,\| \leq \varepsilon, \tag{5}$$

for some $\varepsilon > 0$. A more expressive ansatz corresponds to a smaller achievable approximation error. While a larger manifold generally increases expressibility, excessively large or unstructured manifolds may resemble random states and hinder optimization [34].

Trainability is governed by the geometry of the cost landscape defined over the manifold $\mathcal{M}$. For a given Hamiltonian H, the cost function is defined as $E(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|\,H\,|\psi(\boldsymbol{\theta})\rangle$, and its optimization depends on the behavior of the gradient

$$\nabla_{\boldsymbol{\theta}} E(\boldsymbol{\theta}) = \left(\frac{\partial E}{\partial \boldsymbol{\theta}_1}, \ldots, \frac{\partial E}{\partial \boldsymbol{\theta}_p}\right). \tag{6}$$

Poorly structured ansätze may induce concentration of measure in the parameter space, leading to exponentially vanishing gradients as the system size increases, a phenomenon commonly referred to as barren plateaus [33].

Symmetry compliance [15] is characterized by whether the ansatz manifold $\mathcal{M}$ lies within an invariant subspace of $\mathcal{H}$ defined by conserved quantities (e.g., particle number, spin, or parity). This is ensured when the parameterized unitary $U(\boldsymbol{\theta})$ commutes with the symmetry operator S, i.e.,

$$[U(\boldsymbol{\theta}), S] = 0, \tag{7}$$

and the reference state $|\phi_0\rangle$ is chosen as an eigenstate of S. In this case, all states $|\psi(\boldsymbol{\theta})\rangle$ remain within the corresponding invariant subspace of $\mathcal{H}$, ensuring that the optimization is restricted to physically relevant states.

Noise robustness of a variational ansatz characterizes the stability of states in the manifold $\mathcal{M}$ under the action of noisy quantum operations, modeled as completely positive trace-preserving (CPTP) maps [16]. For a given parameter vector $\boldsymbol{\theta} \in \mathbb{R}^p$, the robustness of the ansatz under a quantum channel $\mathcal{E}$ can be assessed by the deviation between the ideal state $\rho(\boldsymbol{\theta}) = |\psi(\boldsymbol{\theta})\rangle\,\langle\psi(\boldsymbol{\theta})|$ and the noisy state $\mathcal{E}(\rho(\boldsymbol{\theta}))$. For instance, using the trace distance[1],

$$\delta_{\mathrm{tr}}(\boldsymbol{\theta}) = \frac{1}{2}\,\|\rho(\boldsymbol{\theta}) - \mathcal{E}(\rho(\boldsymbol{\theta}))\|_1\,, \tag{8}$$

where $\|\cdot\|_1$ denotes the trace norm. An ansatz is said to be noise-robust if $\delta_{\mathrm{tr}}(\boldsymbol{\theta})$ remains uniformly small over the parameter manifold. In practice, ansätze with shallow circuit depth and structured operator forms tend to exhibit greater resilience to noise, as they accumulate fewer errors under the action of $\mathcal{E}$.

Measurement cost of a variational ansatz characterizes the resources required to estimate expectation values of observables over the manifold $\mathcal{M}$. By decomposing the given Hamiltonian H using Pauli operators as $H = \sum_i c_i P_i$ the corresponding measurement cost [6] is determined by the number of samples required to estimate $E(\boldsymbol{\theta}) = \sum_i c_i\,\langle\psi(\boldsymbol{\theta})|\,P_i\,|\psi(\boldsymbol{\theta})\rangle$. The accuracy of this estimation depends on the variance

$$\mathrm{Var}[E(\boldsymbol{\theta})] = \sum_i c_i^2\,\mathrm{Var}[P_i]_{\boldsymbol{\theta}}, \tag{9}$$

where $\mathrm{Var}[P_i]_{\boldsymbol{\theta}} = 1 - \langle\psi(\boldsymbol{\theta})|\,P_i\,|\psi(\boldsymbol{\theta})\rangle^2$, as well as the number of Pauli terms and their commutation structure. Grouping commuting operators can reduce the number of distinct

[1]Equivalently, fidelity-based measures may be used to characterize the preservation of the prepared state under noisy evolution.

measurement settings, thereby lowering the overall cost. The total cost therefore depends not only on the number of Pauli terms, but also on their commutation structure and the chosen measurement strategy.

Different ansatz families can be understood as defining distinct variational manifolds with specific geometric and structural properties. Chemistry-inspired ansätze constrain the manifold to physically motivated subspaces, hardware-efficient ansätze aim to maximize coverage of the Hilbert space with shallow circuits, while adaptive approaches dynamically construct problem-specific manifolds.

B. Fixed Chemistry Ansätze (UCC-type)

Chemistry-inspired ansätze, particularly those based on the Unitary Coupled Cluster (UCC) framework, form the foundation of many variational quantum algorithms for molecular simulation [35]. The UCC wavefunction remains a central construct in these approaches, owing to its direct connection to established methods in quantum chemistry. In this formulation, the variational state is expressed as

$$|\psi(\boldsymbol{\theta})\rangle = e^{T(\boldsymbol{\theta})-T^\dagger(\boldsymbol{\theta})}\,|\phi_0\rangle\,, \tag{10}$$

where $|\phi_0\rangle$ is a reference state, typically the Hartree–Fock state, and $T(\boldsymbol{\theta})$ is the excitation operator composed of fermionic creation and annihilation operators [7]. In the commonly used UCCSD approximation [13], it is decomposed into single and double excitations as

$$T(\boldsymbol{\theta}) = \sum_{i,a} \theta_i^a\, a_a^\dagger a_i + \sum_{i<j,a<b} \theta_{ij}^{ab}\, a_a^\dagger a_b^\dagger a_j a_i. \tag{11}$$

To construct a quantum circuit implementation of the UCC ansatz, the fermionic anti-Hermitian generator $T(\boldsymbol{\theta}) - T^\dagger(\boldsymbol{\theta})$ is first mapped to a qubit operator using fermion-to-qubit transformations such as the Jordan–Wigner [36] or Bravyi–Kitaev [37] mappings. This yields a representation of the form

$$T(\boldsymbol{\theta}) - T^\dagger(\boldsymbol{\theta}) = \sum_j \alpha_j(\boldsymbol{\theta}) P_j, \tag{12}$$

where P_j are Pauli strings and $\alpha_j(\boldsymbol{\theta}) \in \mathbb{R}$. The corresponding unitary operator $U_{\mathrm{UCC}}(\boldsymbol{\theta})$ is then implemented via a Trotter–Suzuki decomposition [38], whereby

$$U_{\mathrm{UCC}}(\boldsymbol{\theta}) = e^{\sum_j \alpha_j P_j} \approx \prod_j e^{\alpha_j P_j}. \tag{13}$$

Each exponential term $e^{\alpha_j P_j}$ is realized as a Pauli rotation, implemented using a sequence of basis-change gates, entangling CNOT chains, and a single-qubit R_z rotation acting on an appropriate qubit basis [13].

From a geometric perspective, the UCC ansatz is generated via the exponential map of an anti-Hermitian operator algebra formed by fermionic excitation operators, endowing the variational manifold with a structured Lie-algebraic organization [23]. The corresponding unitary $U_{\mathrm{UCC}}(\boldsymbol{\theta})$ defines a variational manifold of the form given in Eq. (4), which we denote by $\mathcal{M}_{\mathrm{UCC}}$.

By construction, UCC ansätze satisfy the symmetry condition in Eq. (7), ensuring that $\mathcal{M}_{\mathrm{UCC}}$ remains confined to an invariant subspace of $\mathcal{H}$ and thereby enforcing symmetry compliance [14]. However, from the perspective of expressibility in Eq. (5), this physically motivated restriction limits the ability of $\mathcal{M}_{\mathrm{UCC}}$ to approximate arbitrary states compared to more flexible ansätze [39].

Regarding trainability, as characterized in Eq. (6), the structured nature of UCC ansätze can lead to more favorable optimization landscapes and reduced susceptibility to barren plateaus [33]. In contrast, noise robustness (cf. Eq. (8)) is often adversely affected in practice due to circuit depth arising from non-local Pauli strings and Trotterized implementations [40].

The measurement cost in Eq. (9) is not directly determined by the ansatz but by the Pauli decomposition of the Hamiltonian. However, UCC ansätze can indirectly increase sampling overhead due to increased circuit depth and circuit complexity, which may lead to increased variance in measured observables and slower convergence [6].

C. Hardware-Efficient Ansätze (HEA)

Hardware-efficient ansätze (HEA) provide a practical alternative to chemistry-inspired constructions, particularly for NISQ devices. In contrast to UCC-based approaches, HEA are defined directly at the circuit level using hardware-native gate sets and connectivity, without relying on an underlying fermionic or problem-specific operator structure. The variational unitary is typically constructed as a sequence of parameterized single-qubit rotations interleaved with entangling layers [14],

$$U(\boldsymbol{\theta}) = \prod_{\ell=1}^{L} U_{\mathrm{ent}}^{\ell}\, U_{\mathrm{rot}}^{\ell}(\boldsymbol{\theta}^\ell), \tag{14}$$

where $U_{\mathrm{rot}}^{\ell}(\boldsymbol{\theta}^\ell)$ denotes a layer of single-qubit rotations with parameter vector $\boldsymbol{\theta}^\ell \in \mathbb{R}^{p_\ell}$ (where p_ℓ is the number of rotation parameters in layer ℓ), and U_{ent}^{ℓ} represents an entangling operation compatible with the device topology. This construction defines a variational manifold $\mathcal{M}_{\mathrm{HEA}}$ as an instance of the general variational form in Eq. (4).

From the perspective of expressibility in Eq. (5), HEA typically induce large and flexible manifolds capable of approximating a wide range of states in $\mathcal{H}$. However, as discussed in Eq. (6), this increased expressibility may lead to concentration of measure in the parameter space, resulting in barren plateaus [33] and vanishing gradients [39] for sufficiently deep circuits.

With respect to symmetry compliance, as defined in Eq. (7), HEA generally do not preserve conserved quantities, as their construction is not constrained by the underlying Hamiltonian. This may lead to symmetry leakage during optimization unless additional constraints or penalty terms are introduced [24], [25].

From the standpoint of noise robustness (cf. Eq. (8)), the shallow circuit structure and compatibility with hardware-native gates can enhance resilience to noise and reduce compilation overhead [14]. However, increasing circuit depth to

improve expressibility may counteract these advantages by amplifying noise effects.

Similarly, the measurement cost described in Eq. (9) is influenced by the expressibility of the ansatz and the resulting variance of observables. While HEA do not inherently increase the number of Hamiltonian terms, highly expressive states may lead to larger variances [6], thereby increasing the number of measurements required for accurate estimation.

D. Adaptive Ansätze (Hamiltonian-Driven)

Adaptive ansätze, exemplified by approaches such as ADAPT-VQE [15], construct problem-specific variational circuits by iteratively expanding the ansatz using operators selected based on their contribution to the energy gradient. Starting from a reference state $|\phi_0\rangle$, the ansatz is built as

$$U(\boldsymbol{\theta}) = \prod_{k=1}^{K} e^{\theta_k G_k}, \tag{15}$$

where G_k are generators drawn from a predefined anti-Hermitian operator pool, e.g., $\{a_a^\dagger a_i - a_i^\dagger a_a,\ \ a_a^\dagger a_b^\dagger a_j a_i - a_i^\dagger a_j^\dagger a_b a_a\}$. Each operator is selected iteratively according to a Hamiltonian-driven criterion, typically the magnitude of the gradient

$$\frac{\partial E}{\partial \theta_k} = \langle \psi_\ell | \left[H, G_k\right] | \psi_\ell \rangle , \tag{16}$$

where $|\psi_\ell\rangle = \prod_{j=1}^{\ell} e^{\theta_j G_j} |\phi_0\rangle$ denotes the current variational state, and the operators $\{G_j\}_{j=1}^{\ell}$ have been selected in previous iterations. This defines a variational manifold $\mathcal{M}_{\mathrm{ADAPT}}$ as a dynamically constructed instance of Eq. (4).

From the perspective of expressibility in Eq. (5), adaptive ansätze achieve efficient coverage of the relevant subspace of $\mathcal{H}$ by selectively including only the most important operators. This typically yields compact representations with high accuracy compared to fixed ansatz constructions [27].

Regarding trainability, as characterized in Eq. (6), the iterative, gradient-informed construction can mitigate barren plateau behavior by restricting the ansatz to directions that significantly impact the cost function [33]. This often results in improved optimization performance compared to highly expressive, unstructured ansätze.

In terms of symmetry compliance (cf. Eq. (7)), the properties of the ansatz depend on the choice of operator pool [28]. When symmetry-adapted operators are employed, the resulting manifold can remain confined to invariant subspaces of $\mathcal{H}$; otherwise, symmetry violations may occur during the adaptive growth process.

The noise robustness (cf. Eq. (8)) of adaptive ansätze is influenced by both circuit depth and the iterative construction procedure [27]. While the resulting circuits are often shallow due to their compactness, the repeated gradient evaluations required during ansatz construction increase sensitivity to noise and shot-noise-induced errors.

Finally, the measurement cost described in Eq. (9) is significantly affected by the operator selection procedure. At each iteration, gradients must be evaluated for all operators in the pool [27], leading to substantial measurement overhead that scales with both the pool size and system dimension.

E. Ansatz Comparison and Analysis

The different ansatz families discussed above can be understood through the geometric properties of the variational manifolds they define. Table I summarizes their key characteristics. Each ansatz family reflects a distinct trade-off between expressibility, trainability, and physical structure.

Chemistry-inspired ansätze such as UCC define structured manifolds constrained by physical symmetries, resulting in improved interpretability and stability of optimization, but at the cost of increased circuit depth and measurement overhead.

In contrast, hardware-efficient ansätze define large and relatively unstructured manifolds that enable efficient implementation on NISQ devices. While this enhances expressibility and hardware compatibility, it often leads to poor trainability and symmetry violations due to the absence of problem-specific structure.

TABLE I: Comparison of ansatz families from the perspective of variational manifold properties.

Property	**UCC**	**HEA**	**Adaptive**
Construction	Fermionic	Layered	Gradient-driven
Manifold	Structured	Unstructured	Dynamic
Expressibility	Moderate	High	Adaptive
Trainability	Moderate	Low	Improved
Symmetry Comp.	Preserved	Weak	Pool-dependent
Noise Robustness	Low	Moderate–High	Moderate
Measurement Cost	High	Moderate	Very High

Adaptive ansätze provide a middle ground by dynamically constructing the variational manifold based on Hamiltonian-informed operator selection. This allows efficient exploration of problem-relevant subspaces and results in compact circuit representations. However, this advantage comes at the expense of increased measurement and optimization cost during the ansatz construction phase.

IV. Conclusion

Within HPC platforms supporting hybrid quantum-classical workflows, VQE plays a central role in molecular simulation, while its success largely depends on efficient ansatz preparation on NISQ devices. The foregoing analysis frames ansatz design from a variational manifold perspective. Within this framework, different ansatz families reflect a fundamental trade-off between structure and flexibility. Highly structured manifolds, as in UCC, restrict the search space to physically meaningful regions, improving optimization behavior but limiting expressibility. In contrast, unstructured manifolds, as in HEA, enable broader exploration of the Hilbert space but often suffer from unfavorable optimization landscapes. Adaptive approaches attempt to balance these extremes by incrementally expanding the manifold along directions informed by the problem Hamiltonian. These observations demonstrate that efficient

ansatz design fundamentally requires balancing expressibility, trainability, symmetry preservation, and hardware constraints.

REFERENCES

[1] P. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 1994, pp. 124–134.

[2] S. Seelam, J. M. Chow, A. Córcoles, S. Sheldon, T. Mittal, A. Kandala, S. Dague, I. Hincks, H. Horii, B. Johnson, M. Le, H. Jamjoom, and J. M. Gambetta, "Reference architecture of a quantum-centric supercomputer," *arXiv preprint arXiv:2603.10970 [cs.ET]*, 2026.

[3] H. J. Manetsch, G. Nomura, E. Bataille, X. Lv, K. H. Leung, and M. Endres, "A tweezer array with 6,100 highly coherent atomic qubits," *Nature*, vol. 647, no. 8088, pp. 60–67, Nov 2025.

[4] D. Bluvstein, A. A. Geim, S. H. Li *et al.*, "A fault-tolerant neutral-atom architecture for universal quantum computation," *Nature*, vol. 649, no. 8095, pp. 39–46, Jan 2026.

[5] M. Cain, Q. Xu, R. King, L. R. B. Picard, H. Levine, M. Endres, J. Preskill, H.-Y. Huang, and D. Bluvstein, "Shor's algorithm is possible with as few as 10,000 reconfigurable atomic qubits," *arXiv preprint arXiv:2603.28627 [quant-ph]*, 2026.

[6] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," *New Journal of Physics*, vol. 18, no. 2, p. 023023, Feb. 2016.

[7] 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," *Nature Communications*, vol. 5, no. 1, Jul. 2014.

[8] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028 [quant-ph]*, 2014.

[9] F. Alam, M. Crichigno, E. Crosson, S. T. Flammia, F. M. Gambetta, M. H. Gordon, M. Kreshchuk, A. Montanaro, A. Nocera, and R. A. Santos, "Onset of ergodicity across scales on a digital quantum processor," *arXiv preprint arXiv:2603.12236 [quant-ph]*, 2026.

[10] F. Jensen, *Introduction to Computational Chemistry*. John Wiley & Sons, 2017.

[11] M. Troyer and U.-J. Wiese, "Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations," *Phys. Rev. Lett.*, vol. 94, p. 170201, May 2005.

[12] G. K.-L. Chan and S. Sharma, "The density matrix renormalization group in quantum chemistry," *Annual Review of Physical Chemistry*, vol. 62, no. Volume 62, 2011, pp. 465–481, 2011.

[13] J. Romero, R. Babbush, J. R. McClean, C. Hempel, P. Love, and A. Aspuru-Guzik, "Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz," *arXiv preprint arXiv:1701.02691 [quant-ph]*, 2018.

[14] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, Sep 2017.

[15] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, "An adaptive variational algorithm for exact molecular simulations on a quantum computer," *Nature Communications*, vol. 10, no. 1, p. 3007, Jul 2019.

[16] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*. Cambridge Univ. Press, Oct. 2000.

[17] S. Lloyd, "Universal quantum simulators," *Science*, vol. 273, no. 5278, pp. 1073–1078, 1996.

[18] D. S. Abrams and S. Lloyd, "Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors," *Phys. Rev. Lett.*, vol. 83, pp. 5162–5165, Dec 1999.

[19] A. Y. Kitaev, "Quantum measurements and the abelian stabilizer problem," *arXiv preprint quant-ph/9511026 [quant-ph]*, 1995.

[20] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, H. Neven, and R. Babbush, "Improved Fault-Tolerant Quantum Simulation of Condensed-Phase Correlated Electrons via Trotterization," *Quantum*, vol. 4, p. 296, Jul. 2020.

[21] J. Lee, W. J. Huggins, M. Head-Gordon, and K. B. Whaley, "Generalized unitary coupled cluster wave functions for quantum computation," *Journal of Chemical Theory and Computation*, vol. 15, no. 1, pp. 311–324, Jan 2019.

[22] D. A. Fedorov, Y. Alexeev, S. K. Gray, and M. Otten, "Unitary Selective Coupled-Cluster Method," *Quantum*, vol. 6, p. 703, May 2022.

[23] M. Haidar, O. Adjoua, S. Badreddine, A. Peruzzo, and J.-P. Piquemal, "Non-iterative disentangled unitary coupled-cluster based on lie-algebraic structure," *Quantum Science and Technology*, vol. 10, no. 2, p. 025031, feb 2025.

[24] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, "Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm," *npj Quantum Information*, vol. 6, no. 1, p. 10, Jan 2020.

[25] G.-L. R. Anselmetti, D. Wierichs, C. Gogolin, and R. M. Parrish, "Local, expressive, quantum-number-preserving VQE ansätze for fermionic systems," *New Journal of Physics*, vol. 23, no. 11, p. 113010, nov 2021.

[26] H. G. A. Burton, "Accurate and gate-efficient quantum ansätze for electronic states without adaptive optimization," *Phys. Rev. Res.*, vol. 6, p. 023300, Jun 2024.

[27] H. L. Tang, V. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou, "Qubit-ADAPT-VQE: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor," *PRX Quantum*, vol. 2, p. 020310, Apr 2021.

[28] A. Anand, P. Schleich, S. Alperin-Lea, P. W. Jensen, S. Sim, M. Díaz-Tinoco, J. S. Kottmann, M. Degroote, A. F. Izmaylov, and A. Aspuru-Guzik, "A quantum computing view on unitary coupled cluster theory," *Chemical Society Reviews*, vol. 51, no. 5, pp. 1659–1684, 2022.

[29] W. Li, Y. Ge, S.-X. Zhang, Y.-Q. Chen, and S. Zhang, "Efficient and robust parameter optimization of the unitary coupled-cluster ansatz," *Journal of Chemical Theory and Computation*, vol. 20, no. 9, pp. 3683–3696, May 2024.

[30] T. Shirakawa, J. Robledo-Moreno, T. Itoko, V. Tripathi, K. Ueda, Y. Kawashima, L. Broers, W. Kirby, H. Pathak, H. Paik, M. Tsuji, Y. Kodama, M. Sato, C. Evangelinos, S. Seelam, R. Walkup, S. Yunoki, M. Motta, P. Jurcevic, H. Horii, and A. Mezzacapo, "Closed-loop calculations of electronic structure on a quantum processor and a classical supercomputer at full scale," *arXiv preprint arXiv:2511.00224 [quant-ph]*, 2025.

[31] 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 the scale of exact diagonalization on a quantum-centric supercomputer," *Science Advances*, vol. 11, no. 25, p. eadu9991, 2025.

[32] A. Shajan, D. Kaliakin, F. Liang, T. Pellegrini, H. Doga, S. Bhowmik, S. Das, A. Mezzacapo, M. Motta, and K. M. M. Jr, "Molecular quantum computations on a protein," *arXiv preprint arXiv:2512.17130 [quant-ph]*, 2026.

[33] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *Nature Communications*, vol. 9, no. 1, p. 4812, Nov 2018.

[34] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms," *Advanced Quantum Technologies*, vol. 2, no. 12, p. 1900070, 2019.

[35] R. J. Bartlett and M. Musiał, "Coupled-cluster theory in quantum chemistry," *Rev. Mod. Phys.*, vol. 79, pp. 291–352, Feb 2007.

[36] P. Jordan and E. Wigner, "About the pauli exclusion principle," *Zeitschrift für Physik*, vol. 47, no. 9, pp. 631–651, Sep 1928.

[37] S. B. Bravyi and A. Y. Kitaev, "Fermionic quantum computation," *Annals of Physics*, vol. 298, no. 1, pp. 210–226, 2002.

[38] M. Suzuki, "Fractal decomposition of exponential operators with applications to many-body theories and monte carlo simulations," *Physics Letters A*, vol. 146, no. 6, pp. 319–323, 1990.

[39] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, "Connecting ansatz expressibility to gradient magnitudes and barren plateaus," *PRX Quantum*, vol. 3, p. 010313, Jan 2022.

[40] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, "Chemical basis of trotter-suzuki errors in quantum chemistry simulation," *Phys. Rev. A*, vol. 91, p. 022311, Feb 2015.