

A Givens-Exchange Ansatz for Molecular Variational Eigensolvers

Azadeh Alavi^{1,*}, Fatemeh Kouchmeshki^{2,*}, Muhammad Usman^{3,4}, Yongli Ren¹, Ke Deng¹,
Hossein Akhondi², Abdolrahman Alavi²

¹School of Computing Technologies, RMIT University, Melbourne, Australia

²Pattern Recognition Pty Ltd, Australia; ³Quantum Systems, CSIRO, Clayton, Australia

⁴School of Physics, The University of Melbourne, Australia

*Equal contribution. Correspondence: azadeh.alavi@rmit.edu.au; admin@pr2aid.com

Abstract

Variational quantum eigensolvers require ansatzes that are expressive, trainable, and implementable on constrained quantum hardware. Quantum architecture search can discover compact circuits, but the learned topology is obtained only after an additional search process. We present a deterministic Givens-exchange ansatz for molecular VQE that removes this outer architecture-search loop while retaining a chemically interpretable exchange prior. The circuit initializes the computational-basis state with the lowest diagonal Hamiltonian expectation and applies three local R_Y layers interleaved with two ordered all-pair Givens-exchange blocks. Optimization uses only Hamiltonian expectation values; exact diagonalization is reserved for post-optimization errors and fidelities. Across six fixed seeds, every run on coefficient-verified LiH-6 and H₂O-8 Hamiltonians, and on a BeH₂-6 public-specification candidate, satisfies chemical accuracy, with mean errors of 0.000124, 0.128558, and 0.002152 mHa, respectively. On LiH-6 and H₂O-8 these six-seed means are lower than the published point errors of the compared quantum-architecture-search methods. The method is not claimed to be a hardware-resource optimum: Givens macros and CNOT counts are reported at distinct levels, and hardware realization requires native-gate compilation, routing, finite-shot estimation, and noisy reoptimization. The contribution is an accurate, reproducible, and search-free reference architecture for molecular VQE and quantum-architecture-search benchmarking.

1 Motivation and contribution

Ansatz design is a central bottleneck in near-term variational quantum algorithms. A molecular VQE ansatz must represent correlated electronic structure, remain trainable under a finite classical optimization budget, and compile to circuits compatible with hardware connectivity and noise. Hardware-efficient circuits can be shallow but weakly aligned with molecular excitation structure, whereas chemically inspired circuits may become expensive after compilation. Quantum architecture search (QAS) addresses this trade-off by learning gate placement and circuit structure, often through reinforcement learning or differentiable search, but introduces a second optimization problem in addition to the VQE parameter optimization itself [1, 2, 3, 4].

This work asks whether a transparent, fixed circuit can provide a strong molecular-VQE reference point without learning the architecture. The answer is affirmative on two coefficient-verified primary benchmarks. The proposed ansatz is deterministic at the topology level; only continuous angles are optimized. This makes it especially relevant to QTML because it supplies a controlled baseline against which learned QAS

methods can be compared without conflating circuit discovery, parameter refinement, Hamiltonian convention, and seed selection.

The prior literature already contains the essential primitives: Givens rotations, fermionic swap networks, particle-conserving exchange circuits, and Stiefel or Grassmann formulations of orthogonality-constrained quantum states. The claim here is not to invent those objects. The novelty is the particular search-free circuit template, the coefficient-verified molecular comparison protocol, and the insistence that accuracy, architecture-search cost, and native-gate resource cost are separate axes of evidence. This positioning is important because molecular-VQE benchmarks can otherwise obscure whether an observed gain arises from a better ansatz, a different Hamiltonian convention, a favorable seed, or an unreported architecture-search process.

The main contributions are: (i) a specified Givens-exchange circuit topology with diagonal-basis initialization; (ii) a strict separation between variational information and post-optimization benchmarking information; (iii) six-seed reporting rather than best-seed selection; (iv) same-Hamiltonian LiH-6 and H₂O-8 comparisons against recent QAS methods; and (v) level-separated resource accounting that does not identify a pre-compilation Givens macro with a native CNOT gate.

2 Ansatz and evaluation protocol

After fermion-to-qubit mapping, each molecular problem is represented by an n -qubit Hamiltonian

$$H = \sum_{\mu=1}^M h_{\mu} P_{\mu}, \quad P_{\mu} \in \{I, X, Y, Z\}^{\otimes n}, \quad H = H^{\dagger}. \quad (1)$$

If the stored representation includes a scalar shift s , the total energy of a variational state is $E(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle + s$. Initialization uses only diagonal Hamiltonian entries,

$$b_0 = \arg \min_{b \in \{0,1\}^n} \langle b | H | b \rangle. \quad (2)$$

The optimizer then minimizes the Hamiltonian expectation value. Exact diagonalization is introduced only after completion of the prescribed optimization to compute the exact ground energy E_0 , the absolute error

$$\epsilon_{\text{mHa}} = 1000 |E(\theta_{\star}) - E_0|, \quad (3)$$

and the ground-state fidelity $F = |\langle \psi_0 | \psi(\theta_{\star}) \rangle|^2$.

For qubits $i < j$, the Givens-exchange macro is the two-qubit

operation

$$\begin{aligned} G_{ij}(\beta)|00\rangle &= |00\rangle, & G_{ij}(\beta)|11\rangle &= |11\rangle, \\ G_{ij}(\beta)|01\rangle &= \cos\beta|01\rangle - \sin\beta|10\rangle, \\ G_{ij}(\beta)|10\rangle &= \sin\beta|01\rangle + \cos\beta|10\rangle. \end{aligned} \quad (4)$$

In a Jordan–Wigner occupation encoding, this operation has a direct single-excitation interpretation. In parity-mapped or tapered encodings it remains a valid variational macro but should not be interpreted as the same literal fermionic excitation. The exchange block applies G_{ij} to all qubit pairs in fixed lexicographic order. With $L = 2$ blocks, the implemented state is

$$|\psi(\theta)\rangle = R_2 G_1 R_1 G_0 R_0 |b_0\rangle, \quad (5)$$

where each R_ℓ is a layer of local R_Y rotations. The Stiefel viewpoint is used only as geometric language: the ordered unitary circuit traces a norm-preserving trajectory on $\text{St}_{\mathbb{C}}(2^n, 1)$, not a Riemannian optimization procedure [8, 5, 6, 7]. A two-qubit exchange embedded in an n -qubit Hilbert space should therefore be understood as the corresponding two-dimensional rotation in every compatible spectator-qubit sector, rather than as one isolated coordinate-plane rotation in the full 2^n -dimensional basis. This more careful interpretation preserves the useful geometric intuition while avoiding an overstatement of the manifold picture.

3 Results

All reported runs use six fixed seeds, 150 Adam steps, one restart, diagonal-basis initialization, two all-pair exchange blocks, and local R_Y layers. LiH-6 at 2.2 Å and H₂O-8 are coefficient-verified against the reference dense Hamiltonians used in the HyRLQAS comparison; the maximum matrix difference, Frobenius-norm difference, shift difference, and total exact-energy difference are zero under the distributed dense-matrix convention. BeH₂-6 is retained as a public-specification candidate because coefficient identity with an official external BeH₂ coefficient file is not claimed. The evaluation protocol is intentionally conservative: exact eigenvalues and exact eigenvectors are not used to initialize, update, or select parameters during optimization. They enter only after the retained parameter vector has been fixed, where they provide the reported error and fidelity. This separation is especially important when comparing with learned-architecture studies, because it prevents the exact solution from implicitly guiding the variational search.

Table 1: Six-seed performance. Errors are absolute mHa. The table reports the mean error and sample standard deviation for this work. QAS values are published point errors; “best QAS” is the lowest listed QAS error for that verified Hamiltonian.

Case/status	this work (mHa)	pass	comparison
LiH-6 JW, verified	0.000124	6/6	best QAS 0.290; HyRLQAS 0.590
	std 0.000113		
H ₂ O-8 JW, verified	0.128558	6/6	best QAS 0.170
	std 0.133018		
BeH ₂ -6, candidate	0.002152	6/6	label-level only
	std 0.004242		
LiH-4 parity, stress	0.555932	5/6	not primary
	std 0.898445		

Table 1 summarizes the central numerical evidence. All primary rows are below the chemical-accuracy threshold of 1.60 mHa for every seed, with mean fidelities between

0.9999527 and 0.9999996. The LiH-6 mean is over three orders of magnitude below the best listed QAS point error for that Hamiltonian, while the H₂O-8 mean improves on the HyRLQAS point error with a smaller but still positive margin. The LiH-4 stress row is deliberately not treated as the primary LiH claim: it differs from LiH-6 in geometry, qubit count, mapping, tapering, and landscape structure.




Figure 1: Energy-error comparison on the two coefficient-verified primary Hamiltonians. The Givens-exchange values are six-seed means; comparison values are published point errors.

Figure 1 illustrates the same-Hamiltonian comparison. The result should be read as an accuracy–resource trade-off, not a hardware-resource dominance claim. The fixed ansatz uses 30 and 56 pre-compilation Givens macros for LiH-6 and H₂O-8, respectively, whereas several searched circuits report fewer CNOTs. These quantities are not hardware-equivalent. The present advantage is that the circuit topology is known before optimization and does not require a learned gate-placement policy.




Figure 2: Resource counts separated by reporting level. This work reports pre-compilation Givens macros and retained one-qubit rotations; published QAS methods report CNOT and rotation counts in their own circuit representations.

A depth ablation further supports the mechanism of the result. With no exchange blocks, the diagonal initialization plus local rotations is not chemically accurate for the primary labels. One exchange block improves LiH but does not make the BeH₂ candidate or H₂O uniformly accurate. Two ordered all-pair exchange blocks are the smallest tested depth that makes all six seeds chemically accurate for LiH-6, H₂O-8, and the BeH₂ candidate. This pattern indicates that the observed performance is produced by structured exchange-induced amplitude redistribution, not by diagonal initialization alone.

4 Scope and hardware significance

The method is strongest when the encoded qubits retain an occupation-like interpretation and the molecular state is not dominated by stretched-bond correlation. The parity/tapered and stretched LiH tests identify the present boundary of the exchange prior: the ansatz can find chemically accurate LiH-4 stress solutions in several seeds, but the distribution is seed-sensitive, and stretched LiH-6 is not uniformly accurate. This limitation is useful for QTML benchmarking because it distinguishes search-free expressivity from optimization robustness and mapping dependence. It also clarifies the role of learning-based architecture design: QAS may be most valuable precisely

in regimes where a fixed exchange prior becomes brittle, such as stretched geometries, parity-tapered encodings, or hardware-constrained connectivity graphs.

The method is also reproducible at the reporting level. The public artifacts contain Hamiltonians, scalar shifts, seed-level numerical outputs, Hamiltonian-identity checks, figure data, and scripts that regenerate the reported tables and figures. The production optimization implementation is not required to audit the table and figure values, although an independent end-to-end rerun of the proprietary optimization pipeline is not claimed. This distinction follows the same principle as the resource accounting: the manuscript states exactly which level of evidence is supplied.

Hardware conclusions also require care. The reported simulations are noiseless dense-state calculations. A hardware-level comparison must compile each retained Givens macro into a native basis, account for connectivity and routing, estimate finite-shot measurement cost, and reoptimize under noise. The present result nevertheless establishes a strong reference point: a transparent all-pair exchange topology can match or exceed the noiseless accuracy of recent learned QAS circuits on verified molecular Hamiltonians without performing architecture search.

For the QTML community, the main use of this construction is comparative rather than only chemical. It can act as a fixed control architecture for studies of reinforcement-learning QAS, differentiable architecture search, and noise-aware circuit discovery: if a learned method improves on the present circuit, the improvement can be attributed more clearly to learned structure or hardware adaptation rather than to a hidden change in Hamiltonian, seed convention, or post-selection rule.

5 Conclusion

We introduced a fixed-topology Givens-exchange ansatz for molecular VQE and evaluated it with coefficient verification, six-seed statistics, and separated resource accounting. The

ansatz achieves chemical accuracy in every seed for LiH-6, H₂O-8, and a BeH₂ public-specification candidate. On the two coefficient-verified primary comparisons, the six-seed means are lower than the published point errors of the compared QAS methods. The contribution is an interpretable and reproducible benchmark architecture for molecular VQE: not the first use of Givens rotations or Stiefel geometry, and not yet a compiled hardware optimum, but a precise search-free reference for the next generation of learned and hardware-aware quantum-circuit design. Future work should replace the all-pair graph with Hamiltonian-informed pair selection, compile retained exchanges to device-native gates, and compare against learned QAS methods after common-shot, common-noise, and common-connectivity reoptimization.

References

- [1] A. Peruzzo *et al.*, “A variational eigenvalue solver on a photonic quantum processor,” *Nature Communications*, vol. 5, 4213, 2014.
- [2] 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, 023023, 2016.
- [3] M. Cerezo *et al.*, “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, pp. 625–644, 2021.
- [4] J. Niu *et al.*, “Hybrid Action Reinforcement Learning for Quantum Architecture Search,” arXiv:2511.04967v3, 2025.
- [5] I. D. Kivlichan *et al.*, “Quantum simulation of electronic structure with linear depth and connectivity,” *Physical Review Letters*, vol. 120, 110501, 2018.
- [6] J. M. Arrazola *et al.*, “Universal quantum circuits for quantum chemistry,” *Quantum*, vol. 6, 742, 2022.
- [7] M. R. Nirmal *et al.*, “Parallelized Givens ansatz for molecular ground states,” *J. Phys. Chem. A*, vol. 129, pp. 10794–10805, 2025.
- [8] A. Edelman, T. A. Arias, and S. T. Smith, “The geometry of algorithms with orthogonality constraints,” *SIAM J. Matrix Anal. Appl.*, vol. 20, pp. 303–353, 1998.