

Automated near-term quantum algorithm discovery for molecular ground states [arxiv:2603.26359]

Fabian Finger, Frederic Rapp, Pranav Kalidindi, Alexander Koziell-Pipe, David Zsolt Manrique, Gabriel Greene-Diniz, Stephen Clark, Konstantinos Meichanetzidis

Quantinuum

Kerry He, Kante Yin, Hamza Fawzi, Bernardino Romera Paredes, Alhussein Fawzi

Hiverge

1 Introduction

Realising the potential of quantum computing depends on efficient quantum algorithms. Current hardware is constrained by limited qubit numbers, gate fidelities, and quantum error correction overheads. Optimising resources is essential to bridge the gap to practical utility. In parallel, AI systems have shown promise in autonomously discovering algorithms by combining evolutionary approaches with large language models (LLMs). We focus on the electronic structure problem in quantum chemistry, targeting molecular ground-state energies to chemical precision (1.6 mHa). We investigate adaptive protocols within the variational quantum eigensolver (VQE) framework, benchmarking against ADAPT-VQE and QEB-ADAPT-VQE. These iteratively build ansätze via energy gradients but require extremely large numbers of circuit evaluations. We demonstrate that the Hive, an LLM-driven evolutionary engine for algorithm discovery, autonomously discovers VQE-style algorithms that maintain chemical precision while substantially reducing two-qubit gates and circuit evaluations.

2 Discovery Workflow

We formulate the ground-state search as a constrained program-synthesis problem. The Hive evolves adaptive ansatz-construction and optimisation logic inside a fixed VQE scaffold. For a given molecule, bond length, and basis, InQuanto provides molecular data and predefined excitation pools (e.g., UCCSD). Candidate algorithms start from the Hartree-Fock state and assemble a parameterised ansatz. The evolved object is not a single-instance circuit, but a reusable routine for ansatz growth and parameter updates. This common scaffold isolates performance differences directly to the discovered algorithms.

3 Results

3.1 Chemical Precision and Generalisation

We evaluate dissociation curves of LiH (12 qubits), H₂O (14 qubits), and F₂ (20 qubits) in the STO-3G basis. Initialised from the single-reference Hartree-Fock state, the dissociation regime poses a strong multi-reference challenge. The Hive algorithm maintains chemical precision across all bond lengths for the three molecules. For LiH, robustness stems from chemistry-informed operator prioritisation combined with late-stage global amplitude rebalancing. For H₂O, it avoids baseline plateauing by steering growth with bond-length-aware operator prioritisation and explicit




Figure 1: **Discovery workflow.** The *Hive* is provided with information (via InQuanto) about a specific molecule and a basic code skeleton. Within this fixed quantum chemistry workflow, the *Hive* evolves the operator assembly and optimisation routine to generate new quantum algorithms. Candidate quantum algorithms are evaluated through a common backend that returns the variational energy used to score each program. Each evolved quantum algorithm constructs and updates a parameterised ansatz from an excitation pool starting from the Hartree–Fock reference state.

global parameter refinement. Evolution utilized a sparse subset of bond lengths, yet achieved high precision on unseen test points, demonstrating strong generalisation.

3.2 Reduction in Quantum Resources

We report total circuit evaluations (exact statevector energy evaluations) and two-qubit gate counts compiled to the Quantinuum native gate set. For LiH, the Hive algorithm maintains a compact sequence of a single-digit number of operators and achieves chemical precision with only a few hundred evaluations per bond length. In contrast, ADAPT-VQE and QEB-ADAPT-VQE require *orders of magnitude* more evaluations. For H₂O, the Hive routine retains chemical precision with tens of operators and a few hundred compiled two-qubit gates, keeping the evaluation overhead in the $10^3 - 10^4$ range—typically an order of magnitude below ADAPT-VQE and two orders below QEB-ADAPT-VQE. For F₂, the Hive collapses evaluation overhead to the few-hundred-call regime. Toward dissociation, the algorithm becomes markedly leaner, compressing the ansatz to only a handful of operators and tens of compiled two-qubit gates.

3.3 Algorithmic Mechanisms

The AI-evolved solvers return explicit, human-readable algorithms that we can decompose into types of functional mechanisms, which we can interpret and give names to; for this decomposition, we may also employ an LLM. Ablation studies demonstrate a strict separation of these mechanisms. At the base, *Operator Scoring* prioritises chemically relevant, low-cost excitations using MP2-amplitude-guided heuristics and a soft locality bias, driving early precision gains near equilibrium. *Growth Control* regulates expansion via adaptive, bond-length-aware acceptance thresholds that relax as the ansatz grows. *Optimisation* replaces zero-initialisation with warm-started,




Figure 2: **Bond dissociation curves.** Bond dissociation curves (top) and energy errors relative to FCI (bottom) for LiH (a), H₂O (b), and F₂ (c). *Hive* results are red circles; ADAPT-VQE and QEB-ADAPT-VQE are blue squares and green pentagons. Hollow markers indicate bond lengths used in *Hive* ansatz evolution; solid markers represent unseen test points. The dashed grey line marks the chemical precision threshold (1.6×10^{-3} Ha).

constant-budget 1D coordinate updates, yielding the dominant reduction in circuit evaluations. Post-growth, *Global Refinement* performs coordinate-wise re-optimisation of all angles; this is critical for recovering chemical precision at stretched bond lengths where static correlation strengthens inter-parameter coupling. Finally, *Ansatz Compression* snaps parameters to discrete sets and explicitly prunes negligible rotations, drastically reducing the two-qubit gate count without sacrificing precision.

3.4 Hardware Implementation

Benchmarking *Hive* ansätze on the Quantinuum H2-1 emulator for LiH with Partition-Measurement Symmetry Verification (PMSV) yielded $\Delta E_{\text{FCI}} \approx 6 - 12$ mHa, outperforming baselines. Executing the *Hive* circuit on the H2-1 processor at $R = 2.0$ Å yielded an energy deviation of 3.1 mHa with 3.4 mHa statistical uncertainty, making the hardware result statistically consistent with chemical precision. We finally, tweaking the noise model of the emulator, estimate the amount of noise that would be acceptable to regain chemical precision on a quantum computer; this threshold is higher for the *Hive* algorithms vs the baselines.

4 Outlook

AI program search discovers functions that outperform human-designed methods in efficiency when designing quantum heuristics. We have shown results for VQE-style algorithms for molecular ground states. In follow-up work, we apply this method to *combinatorial optimisation*, both for evolving more efficient quantum algorithms (in the style of non-variational QAOA) as well as for evolving *problem instances* that exhibit *time-to-solution quantum advantage* against state-of-the-art classical solvers.