

Simulating the Ground-State Energy of the Hydrogen Molecule Using the Variational Quantum Eigensolver

Tanveersingh Rajpurohit, Dhruv Baldha, Bhargav Dobariya,
Hetav Chaudhari

Asha M. Tarsadia Institute of Computer Science and Technology, Uka
Tarsadia University, Maliba Campus, Bardoli-Mahuva Road, Tarsadi,
Bardoli, Surat, 394350, Gujarat, India .

Abstract

Estimating the ground-state energy of a molecule is an important problem in quantum chemistry and a common benchmark for near-term quantum computing. This paper presents a simulation study of the hydrogen molecule (H_2) using the Variational Quantum Eigensolver (VQE), a hybrid quantum-classical algorithm in which a parameterized quantum circuit is optimized by a classical routine. The electronic H_2 Hamiltonian is represented as a weighted sum of four-qubit Pauli operators using the Jordan-Wigner mapping. A hardware-efficient parameterized circuit is used as the variational ansatz, and the Constrained Optimization BY Linear Approximation (COBYLA) optimizer minimizes the corresponding energy expectation value. The experiment is performed with Qiskit's noise-free statevector simulator, with the energy recorded throughout the optimization process. The final computed molecular energy is approximately -1.117 Hartree, compared with a reference value of approximately -1.13727 Hartree, giving an absolute difference of about 0.0202 Hartree. The results show that the VQE procedure moves the estimated energy toward the ground state, while the remaining gap is associated primarily with the selected ansatz and optimization budget. The study documents the complete implementation, the practical development constraints, and the factors limiting the final accuracy.

Keywords: Variational Quantum Eigensolver, Quantum Computing, Quantum Chemistry, Hydrogen Molecule, Qiskit, Hardware-Efficient Ansatz, Jordan-Wigner Transformation

1 Introduction

Quantum chemistry is concerned with predicting molecular properties by solving the equations of quantum mechanics for electrons and nuclei. A central task is determining a molecule’s ground-state energy, the lowest energy accessible to the system. This quantity is important for understanding properties such as molecular stability, bond strength, and reaction energies. For very small molecules, high-accuracy classical calculations are possible; however, the computational cost of solving the underlying many-electron problem becomes increasingly impractical as system size grows.

This scaling problem is one motivation for quantum computing in chemistry. Quantum systems can naturally represent quantum states, although current quantum processors remain limited by qubit counts, coherence times, and gate errors. These devices are commonly described as Noisy Intermediate-Scale Quantum (NISQ) systems. Algorithms intended for this regime therefore aim to use relatively shallow circuits while relying on classical computation where appropriate.

The Variational Quantum Eigensolver (VQE) is one such hybrid algorithm. Instead of solving the complete eigenvalue problem on a quantum processor, VQE prepares a parameterized trial state and evaluates its energy. A classical optimizer then updates the circuit parameters to reduce the energy. This makes VQE particularly suitable for studying near-term quantum algorithms.

The hydrogen molecule, H_2 , is a useful benchmark because it is small enough to have a well-established reference energy while still containing the essential steps of a molecular quantum-computing workflow: Hamiltonian construction, fermion-to-qubit mapping, variational-state preparation, and classical optimization.

This project addresses the following research question: How closely can a small, hardware-efficient VQE simulation approach the known ground-state energy of H_2 , and what factors are responsible for the remaining gap?

To investigate this, the project constructs a four-qubit H_2 Hamiltonian, builds a parameterized hardware-efficient circuit, optimizes its parameters using COBYLA, and compares the resulting energy with a reference value.

The specific objectives are:

1. to represent the electronic structure of H_2 as a qubit Hamiltonian;
2. to implement VQE using a parameterized quantum circuit and a classical optimizer;
3. to record and analyze the energy convergence behaviour during optimization;
4. to compare the final VQE energy against the known reference ground-state energy; and
5. to identify and explain the specific factors that prevent the simulation from matching the reference value exactly.

The remainder of this paper reviews relevant work, defines the problem, describes the proposed solution and methodology, presents the implementation and results, discusses the observed energy gap, identifies limitations, and concludes with directions for future improvement.

2 Literature Review

2.1 Quantum Simulation of Molecular Systems

The use of quantum computers for molecular simulation has been investigated for many years. Early experiments demonstrated that small quantum processors could represent electronic-structure problems and extract molecular information [1]. These studies also highlighted an important computational step: molecular Hamiltonians are naturally expressed using fermionic creation and annihilation operators, whereas quantum circuits operate on qubits. A fermion-to-qubit transformation is therefore required.

The Jordan-Wigner transformation is one of the standard mappings used for this purpose. It converts fermionic operators into combinations of Pauli operators acting on qubits while preserving the required fermionic anticommutation structure.

2.2 The Variational Quantum Eigensolver

Peruzzo et al. introduced VQE as a hybrid quantum-classical method for estimating eigenvalues of quantum systems [2]. The key idea is based on the variational principle: the expected energy of any normalized trial state is greater than or equal to the true ground-state energy. A classical optimizer can therefore search for circuit parameters that minimize this expected energy.

McClean et al. later developed the theoretical framework for variational hybrid quantum-classical algorithms and emphasized the importance of ansatz design and optimization in determining the quality of the solution [3].

2.3 Hardware-Efficient Ansätze

Kandala et al. demonstrated VQE on superconducting quantum hardware using hardware-efficient ansätze [4]. Such circuits generally combine parameterized single-qubit rotations with entangling operations selected to match available hardware connectivity.

This approach is attractive for small and introductory VQE experiments because it is straightforward to construct. Its general-purpose nature, however, can make it less efficient than a chemistry-specific ansatz for representing a particular molecular ground state.

2.4 Quantum Chemistry Software

Modern frameworks provide integrated workflows for molecular quantum simulation. Qiskit Nature, for example, supports construction of the H_2 electronic-structure problem, fermion-to-qubit mappings, and ground-state solvers [5]. The standard H_2 configuration used in this project provides an independent reference for assessing the VQE result.

The present study uses a more manual implementation. The qubit Hamiltonian coefficients are entered directly into the code rather than generated at runtime through

a classical quantum-chemistry package. This allows the VQE optimization process to be studied independently of the Hamiltonian-generation software.

3 Problem Statement

For a molecular Hamiltonian $\hat{H}$, the ground-state energy E_0 is the smallest eigenvalue satisfying

$$\hat{H}|\psi_0\rangle = E_0|\psi_0\rangle. \quad (1)$$

Directly solving the eigenvalue problem becomes increasingly difficult for larger quantum systems. VQE instead uses a parameterized trial state $|\psi(\boldsymbol{\theta})\rangle$ and minimizes its expected energy:

$$E(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|\hat{H}|\psi(\boldsymbol{\theta})\rangle. \quad (2)$$

The variational principle guarantees

$$E(\boldsymbol{\theta}) \geq E_0, \quad (3)$$

so minimizing $E(\boldsymbol{\theta})$ produces an upper-bound estimate of the ground-state energy.

The practical question addressed in this work is how close a small hardware-efficient ansatz and fixed classical optimization budget can come to the reference H_2 ground-state energy, and which aspects of the implementation limit that accuracy.

4 Proposed Solution

The project follows the standard VQE pipeline:

$$\text{H}_2 \rightarrow \text{Qubit Hamiltonian} \rightarrow \text{Parameterized Ansatz} \rightarrow \text{VQE Optimization} \rightarrow \text{Estimated Ground State Energy} \quad (4)$$

Five components make up the implementation:

1. a four-qubit Pauli Hamiltonian representing H_2 ;
2. a hardware-efficient parameterized ansatz;
3. a statevector-based expectation-value estimator;
4. the COBYLA classical optimizer; and
5. convergence and reference-energy analysis.

The implementation is deliberately modular so that each stage can be inspected and reproduced independently.

5 Methodology

5.1 Molecular Configuration

The study considers H_2 at a fixed nuclear separation of approximately 0.735 Å using a minimal STO-3G basis. This is a standard benchmark configuration for introductory quantum chemistry simulations [5].

5.2 Qubit Hamiltonian

The electronic H_2 Hamiltonian is represented as a weighted sum of Pauli strings acting on four qubits:

$$\hat{H}_{\text{elec}} = \sum_{i=1}^{15} c_i P_i. \quad (5)$$

where c_i is a real coefficient and P_i is a tensor product of the single-qubit Pauli operators I , X , Y , and Z .

The implemented Hamiltonian contains fifteen Pauli terms, including single-qubit terms such as $IIIZ$ and $ZIII$, two-qubit terms such as $IIZZ$, and four-qubit terms such as $XXXX$ and $YYYY$.

The electronic Hamiltonian alone does not represent the complete molecular energy. Under the Born-Oppenheimer approximation, the total energy is

$$E_{\text{total}} = E_{\text{electronic}} + E_{\text{nuclear}}. \quad (6)$$

For two hydrogen nuclei separated by distance R in atomic units, the nuclear repulsion contribution is

$$E_{\text{nuclear}} = \frac{1}{R}. \quad (7)$$

At the selected separation, the corresponding nuclear contribution is approximately 0.7198 Hartree.

The Hamiltonian coefficients are entered directly into the implementation because the required PySCF workflow could not be installed in the local Windows development environment. The coefficients correspond to the same H_2 /STO-3G configuration used for the reference calculation.

5.3 Variational Ansatz

The trial state is prepared using a hardware-efficient ansatz based on parameterized single-qubit rotations and two-qubit entangling gates. Each qubit is acted on by R_Y and R_Z rotations, followed by controlled- Z gates connecting neighboring qubits in a linear pattern.

The rotation-entanglement block is repeated twice, resulting in a 24-parameter circuit:

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle^{\otimes 4}. \quad (8)$$

The entangling operations allow the circuit to represent correlated states that cannot be represented by independent single-qubit rotations alone.

5.4 VQE Optimization

The expectation value is evaluated with Qiskit’s StatevectorEstimator, which performs exact statevector simulation without finite-shot measurement noise.

The classical optimizer is COBYLA (Constrained Optimization BY Linear Approximation), a derivative-free optimizer. The optimization is run for a maximum of 300 iterations.

The initial parameter vector is sampled uniformly from $[-\pi, \pi]$ using a fixed random seed to ensure reproducibility. During optimization, the energy at each evaluation is recorded for subsequent convergence analysis.

5.5 Software Environment

Table 1 Software components used in the experiment.

Component	Purpose
Python	Implementation
Qiskit	Quantum circuits and operators
Qiskit Algorithms	VQE
StatevectorEstimator	Exact expectation values
COBYLA	Classical optimization
NumPy	Numerical operations
Matplotlib	Visualization

6 Implementation

The implementation is divided into independently runnable Python scripts.

The Hamiltonian is constructed in `01_hamiltonian.py` using Qiskit’s `SparsePauliOp` class. The parameterized circuit is constructed in `02_ansatz.py` using R_Y/R_Z rotations and linear CZ entanglement.

The main VQE calculation is implemented in `03_vqe_run.py`. It initializes the Hamiltonian, ansatz, statevector estimator, COBYLA optimizer, and reproducible initial parameter vector. The resulting electronic energy is combined with the nuclear repulsion energy to obtain the total molecular energy.

A fourth script records the optimization trajectory and produces the convergence plot and VQE-versus-reference comparison chart.

The complete source code and generated figures are available in the public GitHub repository:

<https://github.com/Tanveer-rajpurohit/vqe-h2-project>

7 Results

7.1 Convergence Behaviour

Figure 1 shows the energy returned during successive COBYLA evaluations.

The energy decreases rapidly during the early part of the optimization and then approaches a stable region. The visible fluctuations are not quantum hardware noise




Fig. 1 VQE convergence for the H_2 molecular ground-state energy. The dashed line indicates the reference ground-state energy.

because the experiment uses an exact statevector estimator. Instead, they reflect the exploratory behavior of the classical optimizer while searching the parameter space.

7.2 Final Energy

The final VQE calculation produced

$$E_{\text{VQE}} \approx -1.117 \text{ Ha.} \quad (9)$$

The reference value used for comparison is

$$E_{\text{reference}} = -1.13727 \text{ Ha.} \quad (10)$$

The absolute error is

$$\Delta E = |E_{\text{VQE}} - E_{\text{reference}}| \approx 0.0202 \text{ Ha.} \quad (11)$$

The VQE energy remains above the reference value, which is consistent with the variational principle in Eq. 3.

7.3 Energy Comparison

Figure 2 compares the final VQE energy with the reference value.

The two values are close but distinct. The VQE result approaches the reference energy but does not reproduce it exactly, leaving an absolute difference of approximately 0.0202 Hartree.

8 Discussion

The results provide several useful observations.




Fig. 2 Comparison of the final VQE energy with the reference ground-state energy.

First, the optimization clearly lowers the energy from the random initial parameter configuration. This indicates that the selected ansatz provides a search space containing states substantially closer to the molecular ground state than the initial state.

Second, optimizer convergence should not be confused with exact ground-state convergence. COBYLA can settle in a locally optimal region within the states reachable by the chosen ansatz. The remaining energy gap can therefore result from both the expressiveness of the ansatz and the optimization procedure.

The hardware-efficient ansatz used here is not specifically designed around the electronic structure of H_2 . A chemistry-informed ansatz, such as UCCSD, could potentially represent the target state more efficiently. In addition, the fixed 300-iteration limit may prevent the optimizer from reaching the lowest energy available within the chosen parameterized circuit.

Third, the use of an exact statevector simulator means that the observed error is not caused by hardware noise, decoherence, gate errors, or finite-shot measurement uncertainty. A real quantum device would introduce additional sources of error beyond this algorithmic baseline.

Finally, manually specifying the Hamiltonian coefficients allows the VQE stage to remain reproducible despite the local PySCF installation problem. However, it also means that the present implementation is not directly reusable for other molecular geometries without supplying new Hamiltonian coefficients.

9 Limitations

The current study has the following limitations:

- **Fixed Hamiltonian.** The qubit Hamiltonian coefficients are entered directly rather than generated automatically from molecular geometry.
- **Statevector simulation only.** The experiment does not include hardware noise, decoherence, or finite-shot measurement effects.
- **Single ansatz configuration.** Only one hardware-efficient ansatz was evaluated.
- **Single optimizer.** Only COBYLA was used.
- **Single molecular geometry.** The study uses only the equilibrium bond length and does not construct a complete potential-energy curve.

10 Future Work

Future work can address these limitations by generating the Hamiltonian automatically through Qiskit Nature and PySCF, comparing hardware-efficient and chemistry-informed ansatzes such as UCCSD, evaluating alternative classical optimizers, and introducing realistic noise models.

Another useful extension would be to vary the H–H bond length and calculate the molecular energy at multiple geometries. This would produce a potential-energy curve and provide a more complete evaluation of the VQE approach.

11 Conclusion

This paper presented a complete simulation of the Variational Quantum Eigensolver for estimating the ground-state energy of the hydrogen molecule. A four-qubit electronic Hamiltonian was constructed from Pauli operators, a 24-parameter hardware-efficient ansatz was used to prepare trial states, and the COBYLA optimizer minimized the energy over a maximum of 300 iterations using a noise-free statevector simulator.

The final computed molecular energy was approximately -1.117 Hartree, compared with the reference value of -1.13727 Hartree. The resulting absolute difference of approximately 0.0202 Hartree demonstrates that the implemented VQE pipeline successfully approaches the molecular ground-state energy but does not reproduce it exactly.

The remaining gap is best explained by the combination of a generic hardware-efficient ansatz and a fixed optimization budget rather than by a failure of the VQE method itself. The experiment therefore demonstrates both the usefulness of VQE for small molecular simulations and the importance of ansatz design and classical optimization in determining final accuracy.

The implementation also provides a reproducible foundation for further experiments involving automatic Hamiltonian generation, chemistry-informed ansatzes, alternative optimizers, noisy simulation, and molecular bond-length scans.

Source Code

The complete implementation and generated results are publicly available at:

<https://github.com/Tanveer-rajpurohit/vqe-h2-project>

References

- [1] Lanyon, B.P., Whitfield, J.D., Gillett, G.G., Goggin, M.P., Almeida, M.P., Kasal, I., Biamonte, J.D., Mohseni, M., Powell, B.J., Barbieri, M., Aspuru-Guzik, A., White, A.G.: Towards quantum chemistry on a quantum computer. *Nature Chemistry* **2**, 106–111 (2010) <https://doi.org/10.1038/nchem.483>
- [2] Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.-H., Zhou, X., Love, P.J., Aspuru-Guzik, A., O’Brien, J.L.: A variational eigenvalue solver on a photonic quantum processor. *Nature Communications* **5**, 4213 (2014) <https://doi.org/10.1038/ncomms5213>
- [3] McClean, J.R., Romero, J., Babbush, R., Aspuru-Guzik, A.: The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics* **18**(2), 023023 (2016) <https://doi.org/10.1088/1367-2630/18/2/023023>
- [4] Kandala, A., Mezzacapo, A., Temme, K., Takita, M., Brink, M., Chow, J.M., Gambetta, J.M.: Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017) <https://doi.org/10.1038/nature23879>
- [5] Qiskit Nature Development Team: Qiskit Nature Documentation. <https://qiskit-community.github.io/qiskit-nature/>. Accessed: Aug. 13, 2026 (2026)