Variational Quantum Eigensolver

Use a one-qubit VQE example to separate energy estimation, optimisation and ansatz error.

Undergraduate · 22 min lesson · Dr. Charith Jayasekara

By the end, you'll be able to

- Derive an energy curve for a trial state.
- Estimate Pauli terms in separate measurement settings.
- Explain the limits of the variational principle for sampled estimates.

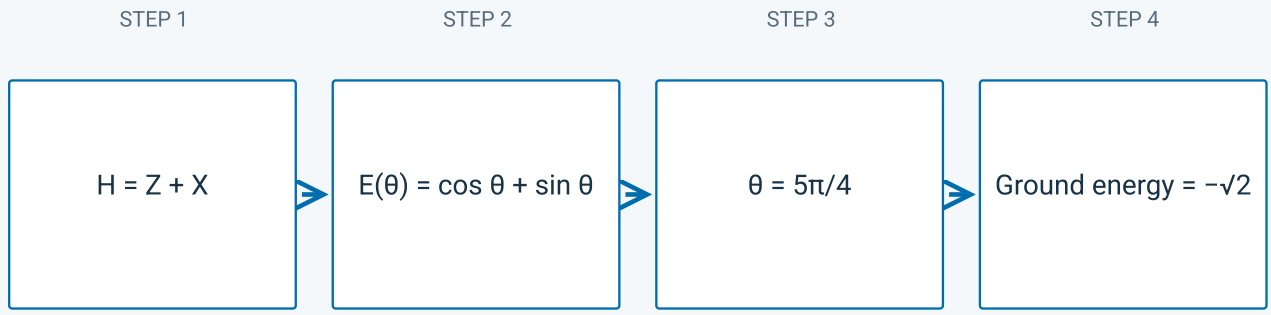
Search for a low-energy state

VQE combines a quantum state-preparation circuit with a classical parameter search. The quantum part estimates the energy of a trial state. The classical part chooses another trial. The aim is usually to approximate a Hamiltonian's lowest eigenvalue and a state near its ground state. The energy is not measured by reading every amplitude; it is reconstructed from observable measurements.

Prerequisites: bras and inner products, Pauli operators and sampling uncertainty. The original experimental work demonstrates the hybrid idea on a small chemical example. Peruzzo and colleagues.

WORKED DIAGRAM · A19

One-qubit VQE



The Y-rotation ansatz reaches the exact ground energy −√2 at θ = 5π/4; restricting to basis states only reaches −1.

The setup and the variational principle

The input includes a Hermitian Hamiltonian, a parameterised state-preparation family and an energy-estimation rule. The output is a selected parameter vector and an estimated energy, optionally accompanied by observables of the prepared state. It is not automatically a complete classical description of the ground-state wavefunction.

For an exactly evaluated normalised state, its energy expectation is at least the ground energy. Expand the state in energy eigenvectors: the expectation is a probability-weighted average of their eigenvalues, so it cannot be below the smallest one. This is the variational principle. It gives a lower target for minimisation, but not a guarantee that a chosen ansatz or optimiser reaches it.

A trial state's ideal expectation is bounded below by the true ground energy.

$$E(\theta) \geq E_0, \qquad E_0 = -\sqrt{2}, \qquad E\left(\frac{5\pi}{4}\right) = -\sqrt{2}.$$

E_0 denotes the lowest eigenvalue. The worked ansatz contains a ground state, so it attains this lower bound at the stated angle; a finite-shot estimate still has statistical uncertainty.

Worked example: one qubit

Take the toy Hamiltonian H equal to Z plus X. Its matrix has diagonal entries one and negative one, and off-diagonal entries one. Squaring it gives twice the identity, so its eigenvalues are positive and negative square root of two. The ground energy is therefore approximately negative 1.414214.

Prepare a real state using a Y rotation of angle theta on zero. Its Z expectation is cosine theta and its X expectation is sine theta. Its energy is their sum. At theta zero, the energy is one. At theta pi, it is negative one. At theta five pi over four, it reaches negative square root of two, the exact ground value.

This example separates expressivity from optimisation. The one-angle family contains the desired ground state, so a perfect search can reach it. A restricted family containing only zero and one cannot do better than negative one. A poor optimisation run might also stop at negative one even when the continuous family is available. The same observed energy can therefore reflect an ansatz limitation or an optimisation failure.

The one-angle trial family gives an energy function that can be minimised exactly.

$$H = X + Z, \qquad |\psi(\theta)\rangle = R_y(\theta)|0\rangle, \qquad E(\theta) = \langle\psi(\theta)|H|\psi(\theta)\rangle = \sin\theta + \cos\theta.$$

$R_y(\theta)$ is a rotation about the Bloch y axis. Its state has X expectation $\sin\theta$ and Z expectation $\cos\theta$.

Estimate the energy through measurements

Measure Z on one set of fresh preparations. Assign positive one to outcome zero and negative one to outcome one, then average. For X, apply H before the Z detector on another set of fresh preparations and average the signed results. Add the two estimates. One physical shot cannot reveal both incompatible observables as exact values.

At the optimum in this example, both expectations are negative one over square root of two. Each signed measurement has variance one half. With an equal independent budget of M shots for each term, the energy estimate has variance one over M. A noisy estimate can lie below the true ground energy by statistical fluctuation even though the exact variational expectation cannot. Report uncertainty rather than declaring that a below-ground sample violates the theorem.

For a many-qubit Hamiltonian, write it as a weighted sum of Pauli strings. Some terms can be measured in compatible groups; others require different settings. Shot allocation should reflect coefficients and variances. The number of terms and desired precision can make measurement cost a dominant part of VQE even when state-preparation depth is modest.

The complete hybrid loop

Choose an ansatz and initial parameters. Prepare the state repeatedly, estimate all required observables and combine them into an energy with uncertainty. Supply that estimate to a classical optimiser. Repeat until a predetermined stopping rule or budget is reached. Freeze the selected parameters and perform a separate evaluation with adequate precision.

Record the number of objective evaluations, total shots, circuit depths, optimiser settings and initialisations. A plot showing only the best energy after each iteration hides the work required to obtain it. Include classical preprocessing, Hamiltonian construction and any error-mitigation cost when comparing with a classical solver.

Energy accuracy and state accuracy

The toy Hamiltonian has a gap of twice the square root of two between its ground and excited energies. In a two-level system, the excited-state population equals the energy error divided by that gap. An exact energy expectation of negative 1.4 therefore implies an excited population of about 0.00503 and ground-state population about 0.99497.

This conclusion uses the known spectrum and an exact expectation. If negative 1.4 is only a noisy estimate, propagate its uncertainty before claiming a fidelity bound. In a system with a tiny spectral gap, the same energy error can permit a much larger excited population. A low energy by itself is not a universal certificate of a nearly correct wavefunction.

The example also distinguishes ansatz bias from measurement uncertainty. More shots can sharpen the energy estimate but cannot make a restricted ansatz contain a missing state. A better ansatz can improve the achievable energy while increasing training and measurement difficulty. Diagnose those effects separately.

Limits and useful interpretations

VQE is a heuristic workflow in general. Difficult landscapes, poor initial overlap, insufficient ansatz expressivity, barren plateaus, noise and measurement expense can prevent useful convergence. A low energy does not always imply high state fidelity unless a spectral gap and error bound are available. Different states close in energy can have different other properties.

The one-qubit calculation is intentionally classically easy. Its purpose is to make every step checkable before studying chemistry or many-body systems. Real applications need an encoding, symmetries, truncation choices, a defensible accuracy target and competitive classical baselines. A quantum chemistry label alone does not establish an advantage.

Exercises and worked answers

Exercise 1. Restrict the toy ansatz to computational-basis states. What is the best energy and its error relative to the exact ground energy?

<details>
<summary>Check the worked answer</summary>

Choose one, giving energy negative one. The energy error is square root of two minus one, approximately 0.414214. More shots cannot remove this expressivity error.
</details>

Exercise 2. An energy estimate is negative 1.43 with standard error 0.03. Does it prove the algorithm found an energy below the ground energy?

Check the worked answer

No. The estimate is statistically compatible with the exact lower bound near negative 1.414. The variational principle constrains the ideal expectation, not every finite-sample estimate. Investigate uncertainty and systematic errors before interpreting the numerical difference.

Check your understanding

PAUSE & THINK

What is the lowest energy attainable using only |0⟩ and |1⟩ for H = Z + X?

- [ ] −√2
- [ ] −1
- [ ] 0

Answer: −1. The basis-state energies are +1 and −1. The lower value −√2 requires a superposition available in a richer ansatz.

Sources

- Peruzzo and colleagues