

Scalable, self-verifying variational quantum eigensolver using adiabatic warm starts

Bojan Žunkovič^{*1,2}, Marco Ballarín³, Lewis Wright³, and Michael Lubasch³

¹University of Ljubljana, Faculty of Computer and Information Science, Večna pot 113, 1000 Ljubljana, Slovenia

²Rudolfovo, Science and Technology Centre, Novo mesto, Slovenia

³Quantinuum, Partnership House, Carlisle Place, London SW1P 1BX, United Kingdom

Abstract

Variational quantum eigensolvers (VQEs) are among the most studied hybrid quantum–classical algorithms for near-term devices, yet they face three persistent obstacles: barren plateaus, convergence to spurious local optima, and the absence of an intrinsic certificate that the optimized state is the desired ground state. We introduce and analyze an adiabatic VQE in which the variational parameters are optimized sequentially along a discretized path of Hamiltonians, each optimization warm-started from the solution of the previous one. We prove that, under explicit spectral and geometric conditions, gradient-based optimization tracks the instantaneous ground state along the entire path, avoiding barren plateaus and local optima in the relevant region of parameter space. A central feature is runtime verification: measurements of the energy standard deviation certify that the variational state is close to an eigenstate and, given a lower bound on the spectral gap, that it is the ground state. The algorithm is thus not only variational and adiabatic but also self-verifying, adapting its step size and rejecting unreliable updates from quantities measurable during execution—providing a rigorous foundation for scalable variational ground-state preparation and reliable hybrid quantum algorithms for optimization, simulation, and quantum machine learning.

Motivation

The variational quantum eigensolver (VQE) is a central algorithmic primitive for quantum simulation, optimization, and quantum machine learning on noisy intermediate-scale quantum devices. Standard VQE, however, suffers from fundamental limitations. The optimization landscape may contain barren plateaus, where gradients vanish exponentially with system size, and many local minima that trap local optimizers. The VQE also does not provide an certificate that the final state is the desired ground state rather than an excited state or a spurious variational optimum.

Adiabatic quantum computation offers a conceptually different route. If a system is initialized in the ground state of a simple Hamiltonian H_0 and evolved slowly enough to a target Hamiltonian H_1 , the adiabatic theorem guarantees ground-state tracking as long as the spectral gap stays open. Its direct circuit-model implementation, however, can require deep circuits, limiting near-term applicability. Can one combine the reliability of adiabatic continuation with the hardware efficiency and flexibility of variational quantum algorithms?

We answer affirmatively with a scalable, self-verifying adiabatic VQE (AVQE): rather than optimizing the target Hamiltonian directly, we replace coherent real-time adiabatic evolution by a sequence of variational problems along an adiabatic path, using each optimized parameter vector as the warm

^{*}Presenting author.

start for the next. This keeps the optimizer in a locally well-conditioned region and provides a natural mechanism for avoiding the hard global features of the landscape. A closely related adiabatic warm-start strategy has been explored concurrently [2]; we add rigorous ground-state tracking guarantees and an intrinsic runtime certificate of correctness.

Adiabatic warm-start VQE

We consider a smooth Hamiltonian path

$$H(s) = (1 - s)H_0 + sH_1, \quad s \in [0, 1], \quad (1)$$

where H_0 has an efficiently preparable ground state and H_1 encodes the target problem. Let $|\psi_0(s)\rangle$ denote the instantaneous ground state and $\Delta(s) = E_1(s) - E_0(s)$ the spectral gap between the ground and first excited eigenvalues. The algorithm discretizes the path into points $0 = s_0 < s_1 < \dots < s_M = 1$ and, at each slice, minimizes the variational energy $E_s(\theta) = \langle \psi(\theta) | H(s) | \psi(\theta) \rangle$ over a parameterized state $|\psi(\theta)\rangle$. The optimizer at s_k is initialized with the parameters obtained at s_{k-1} , $\theta_{k,0} = \theta_{k-1}$, and iterated as $\theta_{k,t+1} = \theta_{k,t} - \eta d_{k,t}$, where $d_{k,t}$ may be a standard gradient, a natural-gradient direction, or another descent direction calculated at $\theta_{k,t}$. After N steps we set $\theta_k = \theta_{k,N}$ and proceed to the next Hamiltonian. The key advantage of this procedure is locality. Standard VQE attempts to reach the final solution from an initial point that may lie far away; adiabatic warm-start VQE only asks the optimizer to move from one ground state to a nearby one. If the path is finely discretized and the variational manifold represents the relevant ground states with nondegenerate local geometry, the optimizer remains inside a well-conditioned basin.

Main theoretical results

Our first result is an adiabatic tracking theorem for gradient-based variational optimization. Under three assumptions—a nonzero instantaneous spectral gap, exact or approximate representability of the instantaneous ground states by the ansatz, and nondegeneracy of the variational geometry—gradient descent tracks the instantaneous ground state along the full path: for a sufficiently small adiabatic step size and enough gradient steps per slice, the optimized state satisfies

$$|\langle \psi(\theta_k) | \psi_0(s_k) \rangle|^2 \geq 1 - \varepsilon \quad (2)$$

for every $k = 0, \dots, M$. The number of optimization updates scales polynomially in the relevant inverse gap and in the number of variational parameters, depending on system size only through operator norms and ansatz-dependent geometric quantities.

The proof admits two complementary readings. First, optimization-theoretically, the spectral gap induces a Polyak–Łojasiewicz-type condition and a locally convex basin around the ground-state minimum. Second, dynamically, the variational updates approximate a discretized imaginary-time flow whose stability is controlled by the same gap and geometric data. Both views explain why adiabatic warm starts keep the optimizer out of regions where barren plateaus and unrelated local minima dominate the landscape. Barren-plateau avoidance is therefore not an independent assumption: gradients remain nonvanishing in the tracking region because the algorithm stays close to a physically meaningful eigenstate branch, so the hard global landscape may still exist but is never explored.

Runtime verification and self-certification

The tracking theorem provides sufficient conditions for success, which may or may not be satisfied. We therefore introduce a runtime verification criterion based on the energy standard deviation

$$\sigma_s(\theta) = \sqrt{\langle \psi(\theta) | H(s)^2 | \psi(\theta) \rangle - E_s(\theta)^2}. \quad (3)$$

Unlike the energy alone, $\sigma_s(\theta)$ vanishes exactly on eigenstates, so measuring it during optimization tests whether the prepared state is close to some eigenstate of the current Hamiltonian.

Our second main result shows that if $\sigma_s(\theta)$ is small compared with a lower bound on the spectral gap, the variational state is uniquely associated with the ground-state eigenbranch and satisfies an explicit fidelity guarantee, with infidelity of order $O(\sigma_s(\theta)^2/\Delta(s)^2)$, provided the adiabatic step size is chosen so that the verified eigenbranch cannot switch between consecutive slices. This transforms energy-standard-deviation measurements into a practical certificate of correctness: the algorithm can verify that it is tracking an eigenstate and, conditioned on a gap lower bound, certify that the eigenstate is the ground state.

The criterion also enables adaptivity: if it fails at a given slice, the algorithm can reduce the adiabatic step size, perform additional optimization steps, or update the assumed gap lower bound, replacing heuristic stopping rules by physically meaningful runtime checks.

Algorithmic structure and robustness

The complete protocol prepares the ground state of H_0 and, at each adiabatic slice, performs a fixed or adaptive number of gradient-based updates, measures both the energy and its standard deviation, and either advances along the path or refines the discretization until the certificate is satisfied; at $s = 1$ the same certificate verifies that the returned state is close to the target ground state. Although our explicit scaling results are derived for Pauli-rotation ansatz circuits under suitable geometric assumptions, the construction extends to natural-gradient and momentum-based optimizers and other descent schemes, and persists under approximate representability, where the variational manifold contains states close to—rather than exactly equal to—the instantaneous ground states. Finite measurement statistics preserve this structure: stochastic gradient estimates remain unbiased, concentration bounds keep the optimization trajectory in the tracking region with high probability for an adequate shot budget, and the standard-deviation certificate stays valid when fluctuations fall below the gap threshold. Shot noise therefore rescales resource requirements without altering the logic of tracking, verification, and certification.

Significance for quantum machine learning and optimization

Variational quantum algorithms are often treated as heuristic learning procedures over quantum circuit parameter spaces; our results identify a regime in which these dynamics become controlled, certifiable, and scalable. The method applies wherever ground-state preparation or energy minimization is a subroutine—quantum optimization, variational circuits, and quantum-enhanced learning pipelines. In the talk, I will present the protocol, explain the optimization-geometric mechanism behind ground-state tracking, show how energy-standard-deviation measurements make AVQE self-verifying, and discuss near-term implementation, noise robustness, and extensions to excited states and other variational quantum algorithms.

References

- [1] B. Žunkovič, M. Ballarin, L. Wright, and M. Lubasch. Scalable, self-verifying variational quantum eigensolver using adiabatic warm starts. [arXiv:2602.17612](#), 2026.
- [2] R. Puig, B. Casas, A. Cervera-Lierta, Z. Holmes, and A. Pérez-Salinas. Warm starts, cold states: Exploiting adiabaticity for variational ground-states. [arXiv:2602.06137](#), 2026.
- [3] 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*, 5:4213, 2014. [doi:10.1038/ncomms5213](#).
- [4] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles. Variational quantum algorithms. *Nature Reviews Physics*, 3:625–644, 2021. [doi:10.1038/s42254-021-00348-9](#).
- [5] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9:4812, 2018. [doi:10.1038/s41467-018-07090-4](#).
- [6] E. R. Anschuetz and B. T. Kiani. Quantum variational algorithms are swamped with traps. *Nature Communications*, 13:7760, 2022. [doi:10.1038/s41467-022-35364-5](#).
- [7] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. Quantum computation by adiabatic evolution. [arXiv:quant-ph/0001106](#), 2000.
- [8] T. Albash and D. A. Lidar. Adiabatic quantum computation. *Reviews of Modern Physics*, 90:015002, 2018. [doi:10.1103/RevModPhys.90.015002](#).
- [9] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller. Self-verifying variational quantum simulation of lattice models. *Nature*, 569:355–360, 2019. [doi:10.1038/s41586-019-1177-4](#).
- [10] J. Stokes, J. Izaac, N. Killoran, and G. Carleo. Quantum natural gradient. *Quantum*, 4:269, 2020. [doi:10.22331/q-2020-05-25-269](#).