

Warm Starts, Cold States: Exploiting Adiabaticity for Variational Ground-States

Ricard Puig*, Berta Casas*, Alba Cervera-Lierta, Zoë Holmes, Adrián Pérez-Salinas

This work is based on Ref. [1].

Introduction Learning ground states and their energies remains among the most compelling uses of quantum computing, with applications in electronic-structure calculations [2], large-scale optimization [3], high energy physics [4], and quantum sensing [5, 6]. Yet a persistent bottleneck underlies essentially all known approaches to ground-state preparation and ground-state energy estimation [7–18]: they typically require initialization with a state that already has nontrivial overlap with the true ground state. How to prepare such “good” initial states is currently the subject of active debate. Classical experience provides insights into how one might approach approximate state preparation. One natural strategy is to optimize within a structured variational family. This perspective has motivated variational quantum methods such as the Variational Quantum Eigensolver (VQE), a hybrid quantum–classical algorithm that minimizes an energy objective of a parameterized quantum state preparation circuit. A second common strategy is to exploit adiabaticity—slowly deforming the ground state of an easy-to-solve Hamiltonian into that of a harder one [19–23]. This principle underlies quantum adiabatic algorithms for ground-state preparation (in both analogue [24] and digital settings [25]) where one evolves under a slowly varying Hamiltonian and leverages a spectral gap to suppress unwanted excitations.

Main results In this work, we analyze an iterative training strategy that integrates adiabatic state preparation [21, 24] into VQE [26]. Starting from a trivial Hamiltonian whose ground state can be prepared using known circuit parameters, we gradually deform the Hamiltonian along an adiabatic interpolation. At each step, the variational parameters are initialized around the optimal solution from the preceding Hamiltonian in the path, and then optimized as seen in Fig. 1.

More formally, consider a general loss function $\mathcal{L}(\alpha)$. An iterative training strategy defines a family of loss functions $\{\mathcal{L}(\alpha, x_k)\}_{k=1}^K$, and uses the solution at step $k - 1$

$$\alpha_{x_{k-1}}^* = \arg \min_{\alpha} \mathcal{L}(\alpha, x_{k-1}), \quad (1)$$

as an initialization bias for optimizing $\mathcal{L}(\alpha, x_k)$. Concretely, the initial parameters for $\mathcal{L}(\alpha, x_k)$ are chosen according to a distribution $\mathcal{D}$ peaked at $\alpha_{x_{k-1}}^*$. For a VQE loss function of the form of $\mathcal{L}(\alpha) = \text{Tr}[U^\dagger(\alpha)\rho U(\alpha)H]$, the family of loss functions is defined as

$$\mathcal{L}_{\text{VQE}}(\alpha, x_k) = \langle \psi | U^\dagger(\alpha) H(x_k) U(\alpha) | \psi \rangle, \quad (2)$$

with $H(x_k) = H_0 + x_k H_1$. Importantly, $\mathcal{L}(\alpha, x_K) = \text{Tr}[U^\dagger(\alpha)\rho U(\alpha)H]$ corresponds to the original problem we wanted to solve $\mathcal{L}(\alpha)$.

We further generalize this approach from vanilla VQE to meta-VQE [27], in which a single parameterized model is trained to output near-optimal parameters across a family of Hamiltonians.

To formalize the benefits of iterative methods in warm starts, we derive a theorem showing that at each step of the iterative procedure both VQE and Meta-VQE maintain trainable gradients. Concretely, the theorem establishes a non-exponentially vanishing lower bound on the loss variance within a hypercube centered at the previous optimum, provided that: (i) consecutive Hamiltonians are sufficiently close, (ii) the previous step achieved adequate accuracy, and (iii) the initial gate induces a nontrivial local state variation. We find that the permissible warm-start neighborhood shrinks as the number of variational parameters increases, but expands with a larger spectral gap. Within this neighborhood, the loss landscape retains sufficient local variation to provide useful gradients, thereby avoiding barren plateaus. We formalize this in form of Theorem 1 in the paper




Figure 1: **Schematic overview of the iterative training mechanism and its failure mode when the gap closes.** **Left:** The loss landscape $\mathcal{L}(\alpha)$ as a function of the variational parameters α at successive iterations (colors). When the ground state varies smoothly and the spectral gap remains open, the location of the minimum drifts continuously, so initializing each step with the previous optimum keeps the optimizer in a high-gradient region that reliably tracks the moving solution. As the gap closes (yellow dashed box), a competing global minimum can emerge far from the followed basin, causing the iterative procedure to remain trapped in the “wrong” minimum past the crossing. **Right:** Visual representation of the iterative warm-start strategy, where the optimal parameters at iteration k are used to initialize the step $k + 1$.

Theorem 1 (informal): Consider a loss function of the form $\mathcal{L}_{\text{VQE}}$ or the meta-VQE counterpart with an ansätze with M variational parameters. Suppose we use an iterative training strategy, and assume that the following conditions are fulfilled:

1. The Hamiltonian change rate, $x_k - x_{k-1}$ is bounded by

$$|x_k - x_{k-1}| \in \mathcal{O}\left(\frac{|\Delta_{\text{gap}}|}{\|H_1\|_s + M\|H(x_k)\|_s}\right). \quad (3)$$

2. The infidelity between the true ground state and the one learned variationally in the previous iteration is bounded by a constant.
3. The generator of the first gate acts non-trivially on the initial state, i.e., $\langle \psi | P_1 | \psi \rangle = 0$.

Consider sampling the parameters α_k from a hypercube of width $2r$ around the previous solution α_{k-1}^* , then we can guarantee that

$$r^2 \in \mathcal{O}\left(\frac{|\Delta_{\text{gap}}|}{M(\|H(x_k)\|_s + |\Delta_{\text{gap}}|)}\right) \Rightarrow \text{Var}_{\alpha \sim \mathcal{D}(\alpha^*, r)}[\mathcal{L}(\alpha, x_k)] \in \Omega(r^4). \quad (4)$$

Here we used $\|\cdot\|_s = \lambda_{\max}^{(\cdot)} - \lambda_{\min}^{(\cdot)}$, where $\lambda_{\max}^{(\cdot)}$ and $\lambda_{\min}^{(\cdot)}$ denote the largest and smallest eigenvalues of the corresponding Hamiltonian.

Theorem 1 provides a step-wise guarantee that, around each warm-start point, the loss landscape does not become exponentially flat. Consequently, as long as the gap $|\Delta_{\text{gap}}|$ does not close along the path, and as long as $\|H\|_s$ and M do not scale exponentially with the number of qubits, the variance decays polynomially, ensuring a local region with provable gradients along the iterative path.

The rate at which the Hamiltonian is iteratively updated, i.e., $|x_k - x_{k-1}|$, also needs to be sufficiently gradual to guarantee a non-exponentially vanishing variance (and hence a non-flat loss landscape). That said, it is only required to be polynomially small in the number of qubits, and this condition is obtained from a loose, worst-case bound and may be improvable in practice.

The second and third conditions capture the intuitive fact that the training at each iteration needs to be sufficiently successful in order for the next one to be close to the solution, and that the first gate in our circuit acts non-trivially on the initial state. Note that this latter condition is a convenient sufficient condition to simplify the proof, but this requirement might not be strictly necessary.

To investigate the practicality of our method, we supplement our theoretical results with numerical simulations of the full training process, accounting for realistic shot noise for system sizes up to $n = 16$. Our results confirm that the method consistently converges when the gap in the ground state energy does not vanish, providing a promising path for ground-state preparation on near-term devices in such cases. However, we see that when the gap closes the method can break down in line with our theoretical predictions and the QMA-completeness of generic Hamiltonians. This is illustrated in Fig. 2.

The behavior around the crossing can be understood from the locality of the optimization. Before the crossing, the tracked state coincides with the global minimum; after the crossing, the same continuous trajectory persists but corresponds only to a *local* minimum. Since the optimizer explores the landscape locally, nearby parameter updates typically increase the energy, making it difficult to transition to a different branch that becomes energetically favorable after the crossing. Escaping this would require an initialization already aligned with the post-crossing ground state, which cannot be guaranteed without prior knowledge of the solution.

Consequently, the issue at hand is not due to barren plateau phenomena or shot noise, but rather a fundamental limitation of local, iterative optimization strategy in the presence of spectral crossings. In this sense, the phenomenon is closely related to the bottlenecks encountered in adiabatic (or adiabatic-like) ground-state preparation strategies when the gap closes.




Figure 2: **Energy spectra learned with warm-start VQE.** Left: XY model $H_{XY}(x) = -J[H_{\sigma_x} + H_{\sigma_y} + x(H_{\sigma_x} - H_{\sigma_y})]$, with $H_{\sigma_j} = \sum_{i=1}^{n-1} \sigma_j^{(i)} \sigma_j^{(i+1)}$, $\sigma_j \in \{\sigma_x, \sigma_y\}$. Right: transverse-field Ising $H_I(x) = -J \sum_{i=1}^n Z_i Z_{i+1} - x \sum_{i=1}^n X_i - Y_1 X_2 \cdots X_{n-1} Y_n$. For each value of the field parameter x the optimization is initialized with the parameters obtained at the previous point. Solid lines show the exact ground-state and first-excited energies as functions of x , and markers show the corresponding estimates obtained under synthetic shot noise with `nshots` = 10^4 . In the Ising row, the vertical dashed line indicates the critical point at $x = 1$. All data are for $n = 16$ spins and circuit depth $L = 16$.

Conclusions This work investigates iterative warm-start training for variational optimization, where each instance is initialized using parameters optimized for a neighboring instance. This iterative framework offers distinct advantages. Unlike standard adiabatic evolution, which often requires significant evolution time (or circuit depth, if the implementation is digitalized) to satisfy the adiabatic theorem, this variational approach can search for a “shortcut to adiabaticity” within a constrained, lower-depth ansatz. Furthermore, the ability to reliably produce approximate ground states makes this method highly valuable for fault-tolerant algorithms by serving as high-overlap inputs, thereby reducing the total runtime required to achieve high-precision energy estimates. Taken together, these results position iterative warm-start variational methods as a potential route to overcoming trainability barriers in variational quantum algorithms.

References

- [1] Ricard Puig, Berta Casas, Alba Cervera-Lierta, Zoë Holmes, and Adrián Pérez-Salinas. Warm starts, cold states: Exploiting adiabaticity for variational ground-states. *arXiv preprint arXiv:2602.06137*, 2026. URL: <https://arxiv.org/abs/2602.06137>, doi:10.48550/arXiv.2602.06137.
- [2] Yudong Cao, Jonathan Romero, Jonathan P Olson, Matthias Degroote, Peter D Johnson, Mária Kieferová, Ian D Kivlichan, Tim Menke, Borja Peropadre, Nicolas PD Sawaya, et al. Quantum chemistry in the age of quantum computing. *Chemical reviews*, 119(19):10856–10915, 2019. URL: <https://pubs.acs.org/doi/10.1021/acs.chemrev.8b00803>, doi:10.1021/acs.chemrev.8b00803.
- [3] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014. URL: <https://arxiv.org/abs/1411.4028>, doi:10.48550/arXiv.1411.4028.
- [4] Alberto Di Meglio, Karl Jansen, Ivano Tavernelli, Constantia Alexandrou, Srinivasan Arunachalam, Christian W Bauer, Kerstin Borrás, Stefano Carrazza, Arianna Crippa, Vincent Croft, et al. Quantum computing for high-energy physics: State of the art and challenges. *Prx quantum*, 5(3):037001, 2024. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.5.037001>, doi:10.1103/PRXQuantum.5.037001.
- [5] Majid Hassani, Mengyao Hu, Guillem Müller-Rigat, Matteo Fadel, and Jordi Tura. Many-body k -local ground states as probes for unitary quantum metrology. *arXiv preprint arXiv:2512.02976*, 2025. URL: <https://arxiv.org/abs/2512.02976>, doi:10.48550/arXiv.2512.02976.
- [6] Ricard Puig, Pavel Sekatski, Paolo Andrea Erdman, Paolo Abiuso, John Calsamiglia, and Martí Perarnau-Llobet. From dynamical to steady-state many-body metrology: Precision limits and their attainability with two-body interactions. *PRX Quantum*, 6(3):030309, 2025. URL: <https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.6.030309>, doi:10.1103/PRXQuantum.6.030309.
- [7] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.
- [8] Kristan Temme, Tobias J Osborne, Karl G Vollbrecht, David Poulin, and Frank Verstraete. Quantum metropolis sampling. *Nature*, 471(7336):87–90, 2011. URL: <https://www.nature.com/articles/nature09770>, doi:10.1038/nature09770.
- [9] A Yu Kitaev. Quantum measurements and the abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, 1995. URL: <https://arxiv.org/abs/quant-ph/9511026>, doi:10.48550/arXiv.quant-ph/9511026.
- [10] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *arXiv preprint quant-ph/0005055*, 2000. doi:10.48550/arXiv.quant-ph/0005055.
- [11] Kok Chuan Tan, Dhiman Bowmick, and Pinaki Sengupta. Quantum stochastic series expansion methods. *arXiv preprint arXiv:2010.00949*, 2020.
- [12] Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019. doi:10.22331/q-2019-07-12-163.
- [13] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204, 2019. URL: <https://dl.acm.org/doi/abs/10.1145/3313276.3316366>, doi:10.1145/3313276.3316366.
- [14] Yimin Ge, Jordi Tura, and J Ignacio Cirac. Faster ground state preparation and high-precision ground energy estimation with fewer qubits. *Journal of Mathematical Physics*, 60(2), 2019. URL: <https://pubs.aip.org/aip/jmp/article-abstract/60/2/022202/918222/Faster-ground-state-preparation-and-high-precision?redirectedFrom=fulltext>, doi:10.1063/1.5027484.

- [15] Mario Motta, Chong Sun, Adrian T.K. Tan, Matthew J. O’Rourke, Erika Ye, Austin J. Minnich, Fernando G.S.L Brandao, and Garnet Kin-Lic Chan. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nature Physics*, 16(2):205–210, 2020. URL: <https://www.nature.com/articles/s41567-019-0704-4>, doi:10.1038/s41567-019-0704-4.
- [16] Danial Motlagh, Modjtaba Shokrian Zini, Juan Miguel Arrazola, and Nathan Wiebe. Ground state preparation via dynamical cooling. *arXiv preprint arXiv:2404.05810*, 2024. doi:10.48550/arXiv.2404.05810.
- [17] Marek Gluza, Jeongrak Son, Bi Hong Tiang, Yudai Suzuki, Zoë Holmes, and Nelly HY Ng. Double-bracket quantum algorithms for quantum imaginary-time evolution. *arXiv preprint arXiv:2412.04554*, 2024. URL: <https://arxiv.org/abs/2412.04554>, doi:10.48550/arXiv.2412.04554.
- [18] Yudai Suzuki, Bi Hong Tiang, Jeongrak Son, Nelly HY Ng, Zoë Holmes, and Marek Gluza. Double-bracket algorithm for quantum signal processing without post-selection. *arXiv preprint arXiv:2504.01077*, 2025. URL: <https://arxiv.org/abs/2504.01077>, doi:10.48550/arXiv.2504.01077.
- [19] A Garcia-Saez and JI Latorre. Addressing hard classical problems with adiabatically assisted variational quantum eigensolvers. *arXiv preprint arXiv:1806.02287*, 2018. URL: <https://arxiv.org/abs/1806.02287v1>, doi:10.48550/arXiv.1806.02287.
- [20] Stuart M Harwood, Dimitar Trenev, Spencer T Stober, Panagiotis Barkoutsos, Tanvi P Gujarati, Sarah Mostame, and Donny Greenberg. Improving the variational quantum eigensolver using variational adiabatic quantum computing. *ACM Transactions on Quantum Computing*, 3(1):1–20, 2022. URL: <https://dl.acm.org/doi/full/10.1145/3479197>, doi:10.1145/3479197.
- [21] Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Rev. Mod. Phys.*, 90:015002, Jan 2018. URL: <https://link.aps.org/doi/10.1103/RevModPhys.90.015002>, doi:10.1103/RevModPhys.90.015002.
- [22] Mohamed Hibat-Allah, Estelle M Inack, Roeland Wiersema, Roger G Melko, and Juan Carrasquilla. Variational neural annealing. *Nature Machine Intelligence*, 3(11):952–961, 2021. doi:10.1038/s42256-021-00401-3.
- [23] Bojan Žunkovič, Pietro Torta, Giovanni Pecci, Guglielmo Lami, and Mario Collura. Variational ground-state quantum adiabatic theorem. *Physical Review Letters*, 134(13):130601, 2025. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.134.130601>, doi:10.1103/PhysRevLett.134.130601.
- [24] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *arXiv preprint quant-ph/0001106*, 2000. doi:10.48550/arXiv.quant-ph/0001106.
- [25] Ioannis Kolotouros, Ioannis Petrongonas, Miloš Prokop, and Petros Wallden. Simulating adiabatic quantum computing with parameterized quantum circuits. *Quantum Science and Technology*, 10(1):015003, 2024. URL: <https://iopscience.iop.org/article/10.1088/2058-9565/ad80c0/meta>, doi:10.1088/2058-9565/ad80c0.
- [26] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1):1–7, 2014. URL: <https://www.nature.com/articles/ncomms5213#citeas>, doi:doi.org/10.1038/ncomms5213.
- [27] Alba Cervera-Lierta, Jakob S Kottmann, and Alán Aspuru-Guzik. The meta-variational quantum eigensolver (meta-vqe): Learning energy profiles of parameterized Hamiltonians for quantum simulation. *PRX Quantum*, 2(2):020329, 2021. URL: <https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.2.020329>, doi:10.1103/PRXQuantum.2.020329.