

A Compass on the Quantum State Sphere: The Hopf Ansatz for Arbitrary Pure-State Optimization

Ruge Lin¹, Guangxi Li², and Xin Wang¹

¹Thrust of Artificial Intelligence, Information Hub, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou 511453, China

²Quantum Science Center of Guangdong–Hong Kong–Macao Greater Bay Area, Shenzhen 518045, China

Abstract

Universal state-preparation circuits answer which pure states can be reached; optimization also requires coordinates, geometry, and derivative access. We introduce the Hopf ansatz, a binary-tree circuit chart for arbitrary normalized real and complex states. One fixed skeleton provides state preparation, a linear-time inverse from amplitudes to physical angles, a closed-form diagonal pullback metric, and exact preparation of normalized coordinate tangents. The tangents convert objective derivatives into measurable transition moments and admit layer- and leaf-indexed batching, requiring $1 + 2n$ real or $1 + 2(n + 1)$ complex compiled gradient-access configurations for n qubits. Deterministic real-state VQE and metrology-inspired stress tests with known global optima show numerical-precision median gaps and substantially more threshold crossings for geometry-aware Hopf optimizers than for coordinate Adam baselines. For quantum machine learning, the construction supplies an executable full-state layer whose initialization, geometry, and derivative interface are exposed within one architecture.

1. Introduction

Parameterized quantum circuits act as models, encoders, and coordinate systems, so their usefulness depends on expressivity, initialization, geometry, and derivative access [1–3]. Full-state preparation pushes expressivity to its limit: arbitrary-state preparation solves reachability [4], but optimization also requires navigability through a physical chart that can be inverted and differentiated. The Hopf ansatz provides such a chart by combining an explicit inverse, analytic geometry, and exactly preparable tangent directions within one state-preparation architecture. While quantum natural-gradient methods obtain geometry from a quantum geometric tensor [5], the Hopf metric is available in closed form, avoiding estimation cost.

2. The Hopf chart

For an n -qubit system, let $N = 2^n$ denote the Hilbert-space dimension, equivalently the number of leaves in the Hopf tree. Figure 1 summarizes the construction. The ansatz represents a normalized state as a binary tree: $N - 1$ internal nodes split probability mass; real signs are encoded by final-layer split angles, and the complex chart attaches one phase to each leaf. Each split uses $\cos \theta_j$ and $\sin \theta_j$, so every root-to-leaf path produces one computational-basis amplitude. For example, the leaf $13 = (1101)_2$ follows a right–right–left–right path,

$$x_{13} = \sin \theta_1 \sin \theta_3 \cos \theta_7 \sin \theta_{14}. \quad (1)$$

On a quantum circuit, the tree maps directly to gates: internal nodes become controlled R_y gates, with complex final-layer gates carrying the sibling phases. Given an arbitrary target state, a bottom-up pass computes subtree masses, recovers split angles by an atan2 rule, and reads real signs or complex phases at the leaves, giving, up to usual zero-subtree nonuniqueness, an $O(N)$ -time and $O(N)$ -memory inverse on the same physical tree.


Metric. The tree induces a diagonal pullback metric; for complex states we use the normalized state-vector sphere convention. For a magnitude node j , let $\mathcal{L}(j)$ be its descendant leaves. The metric weights are

$$g_{jj} = \sum_{\ell \in \mathcal{L}(j)} |x_\ell|^2, \quad g_{N+\ell, N+\ell} = |x_\ell|^2 \quad (\text{complex phase}). \quad (2)$$


The magnitude weight is the probability mass entering node j , while a phase weight is the corresponding leaf occupation. Disjoint mass partition between child nodes makes the metric diagonal, so natural-gradient rescaling is componentwise without estimating or inverting a dense metric [6], and the diagonal form enables efficient state-sphere geodesics and vector transport.

Tangent states. The metric yields explicit tangent vectors at each node. Let $|\psi_L\rangle$ and $|\psi_R\rangle$ be the normalized states of the two child subtrees of node j . Its normalized coordinate tangent is

$$|e_j\rangle = -\sin \theta_j |\psi_L\rangle + \cos \theta_j |\psi_R\rangle. \quad (3)$$



(a) Four-qubit Hopf tree. A root-to-leaf path multiplies sine and cosine factors; the leaf labels $\theta_{16}, \dots, \theta_{31}$ are complex phases. Circled numbers are gates implementation order.



(b) Fixed circuit skeleton. Complex final-layer gates also carry the two sibling phases.

Figure 1: The same binary tree defines amplitudes, physical gates, inverse coordinates, metric weights, and tangent-state assignments.

It is prepared on the same skeleton by clamping the ancestor path, shifting $\theta_j \mapsto \theta_j + \pi/2$, retaining the descendant subtree, and zeroing irrelevant branches. Complex phase tangents are single-leaf states with the inherited phase shifted by $\pi/2$.

Gradient. Tangents convert the objective landscape into measurable quantities. For $E(\boldsymbol{\theta}) = \langle \psi | H | \psi \rangle$, the coordinate derivative is the transition moment $\partial_j E = 2\sqrt{g_{jj}} \operatorname{Re} \langle e_j | H | \psi \rangle$. Expressing this as a difference of expectation values avoids a Hadamard-test circuit: for $s \in \{+1, -1\}$, define $|\phi_j^{(s)}\rangle = (|\psi\rangle + s|e_j\rangle)/\sqrt{2}$, which gives

$$\partial_j E = 2\sqrt{g_{jj}} s \left(\langle H \rangle_{\phi_j^{(s)}} - \frac{\langle H \rangle_{\psi} + \langle H \rangle_{e_j}}{2} \right). \quad (4)$$

Thus the gradient is expressed through expectation values on prepared states; fixed Hopf assignments prepare derivative directions rather than shifted evaluations of the original objective.

Compiled-circuit families. Tangent preparations are structured, enabling batched compilation. Tangents at one magnitude-tree depth share one indexed tangent preparation and one indexed signed-branch preparation; all complex phase tangents and branches form one additional leaf-indexed pair. Together with one baseline setting, this gives

$$N_{\text{set}}^{\mathbb{R}} = 1 + 2n, \quad N_{\text{set}}^{\mathbb{C}} = 1 + 2(n + 1) = O(\log N), \quad (5)$$

where a setting is a fixed compiled gate pattern and indexed preparation choice. Under the paper’s no-clean-ancilla CNOT-counting model for decomposing multi-controlled gates, each setting has the same $O(nN)$ compiled CNOT cost as forward preparation. Batching does not change per-setting compiled cost; sampling precision remains componentwise.

3. Numerical validation

To test whether the geometric structure translates into optimization performance, we compare three Hopf-geometric methods: exact-geodesic-transport conjugate gradient (EGT-CG) [6], Riemannian Barzilai–Borwein (R-BB) [10], and Riemannian L-BFGS (R-LBFGS) [11] with Hopf-Adam (coordinate Adam on the Hopf chart) and a real Möttönen physical-angle Adam baseline [4, 12]. The deterministic state-vector




Figure 2: Aggregate final gaps and threshold hitting over 150 traces per optimizer and application class. Top: final optimality gaps after 200 accepted updates. Bottom: first update below 10^{-8} ; parenthetical labels report crossing counts. Lower is better.

benchmark uses the real Hopf ansatz at $n = 6, \dots, 10$, corresponding to $N = 64, \dots, 1024$, three scrambled VQE tasks, three metrology-inspired tasks (MET), ten deterministic initial states per task, and 200 accepted updates, giving 150 traces per optimizer and application class. All six tasks have known global optima, so convergence is measured directly by optimality gaps.

Figure 2 shows that all three geometric methods reach numerical-precision median gaps in both application classes. The advantage is clearest in VQE: R-BB obtains a mean final gap of 1.61×10^{-12} , compared with 3.47×10^{-3} for the stronger coordinate baseline, and crosses 10^{-8} in all 150 traces. In the metrology-inspired aggregate, the Möttönen baseline has a mean comparable to the best geometric method but retains a larger median; the geometric methods cross 10^{-8} in 117–146 of 150 traces, versus 11–12 for the coordinate baselines. Hopf-Adam isolates the geometric contribution: using the same chart and coordinate gradients without the induced-metric and state-sphere updates does not reproduce the geometric per-update behavior. Within the geometric methods, R-BB has the strongest VQE aggregate. In MET no method is uniformly dominant: EGT-CG has the smallest mean gap and most best-gap wins, while R-BB has the strongest rank statistics. Per-task curves at $n = 10$ in the paper show that the easiest instances are solved rapidly; harder objectives are distinguished mainly by slow tails and threshold failures.

4. Applications and outlook

Hopf is an executable full-state variational layer rather than only a loader. Three properties make it suitable for optimization: the inverse map closes the Hopf chart, the analytic metric replaces estimated natural-gradient quantities, and the tangent states give derivative access through prepared circuits. VQE is the most direct application: for $E(\theta) = \langle H \rangle_\psi$, the Hamiltonian H is already the transition-moment observable. With exact initialization, analytic natural-gradient weights, and an inverse map returning state-space updates to circuit angles, Hopf becomes a geometry-aware full-state ansatz for ground-state searches where a sufficiently expressive problem-specific ansatz is not known a priori [3]. When the relevant sector is known, it complements structured encoders such as fixed-Hamming-weight constructions [13].

The interface also extends to any probe-state differential expressible as a Hermitian chain-rule observable, such as pure-state quantum Fisher information (QFI) and local fixed-readout Ramsey classical Fisher information (CFI). The Hopf chart, tangent preparations, and compiled-circuit families remain unchanged; the application enters through the chain-rule observable.

Hopf inherits the exponential scaling of state-vector preparation; its advantage lies in optimization geometry, not dimensionality reduction. For full-state variational problems, Hopf turns universal state preparation into an optimizer-ready circuit chart by exposing initialization, inversion, geometry, and derivative access within one architecture.

References

- [1] V. Havlíček et al., “Supervised learning with quantum-enhanced feature spaces,” *Nature* **567**, 209–212 (2019).
- [2] M. Schuld, R. Sweke, and J. J. Meyer, “Effect of data encoding on the expressive power of variational quantum-machine-learning models,” *Physical Review X* **11**, 041011 (2021).
- [3] M. Cerezo et al., “Variational quantum algorithms,” *Nature Reviews Physics* **3**, 625–644 (2021).
- [4] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, “Transformation of quantum states using uniformly controlled rotations,” arXiv:quant-ph/0407010 (2004).
- [5] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, “Quantum natural gradient,” *Quantum* **4**, 269 (2020).
- [6] A. J. Ferreira-Martins et al., “Quantum optimization with exact geodesic transport,” arXiv:2506.17395 (2025).
- [7] B. Koczor et al., “Variational-state quantum metrology,” *New Journal of Physics* **22**, 083038 (2020).
- [8] J. J. Meyer, J. Borregaard, and J. Eisert, “A variational toolbox for quantum multi-parameter estimation,” *npj Quantum Information* **7**, 89 (2021).
- [9] R. Kaubruegger et al., “Quantum variational optimization of Ramsey interferometry and atomic clocks,” *Physical Review X* **11**, 041045 (2021).
- [10] J. Barzilai and J. M. Borwein, “Two-point step size gradient methods,” *IMA Journal of Numerical Analysis* **8**, 141–148 (1988).
- [11] D. C. Liu and J. Nocedal, “On the limited memory BFGS method for large scale optimization,” *Mathematical Programming* **45**, 503–528 (1989).
- [12] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” arXiv:1412.6980 (2014).
- [13] R. M. S. Farias et al., “Quantum encoder for fixed-Hamming-weight subspaces,” *Physical Review Applied* **23**, 044014 (2025).
- [14] R. Lin and G. Li, “Hopf ansatz GitHub repository” (2026).