Symmetry in Landscapes

Variational Quantum Algorithms · Module 3 — Cost Landscapes

Node 3.4

Position in Knowledge Graph

System: Variational Quantum Algorithms **Structural Domain:** Cost Landscapes **Node:** 3.4

Symmetry in a VQA cost landscape is both a **gift** and a **curse**, depending on how you look at it.

The **gift**: a symmetric landscape has equivalent critical points related by group actions. Once you find one, you've effectively found them all. Symmetries can be exploited to restrict the search space.

The **curse**: symmetric landscapes have *redundant* parameter directions. Two distinct parameter vectors can produce the same physical state, which means the parameter-space gradient has a kernel — there are directions in parameter space that don't change anything, but the optimizer still tries to move along them, wasting effort.

This node walks through both sides of the symmetry coin and explains how to think about symmetry-aware optimization.

Three Sources of Symmetry

There are three distinct kinds of symmetry that show up in VQA cost functions, and they have very different consequences.

1. Hamiltonian Symmetries

If the Hamiltonian H commutes with some unitary G, then the eigenstates of H split into invariant subspaces under G, and the ground state can be chosen to lie in a specific symmetry sector.

Examples: - **Particle number symmetry**: $[H, \hat{N}] = 0$ for fermionic Hamiltonians. - **Total spin / spin projection**: for spin-conserving Hamiltonians. - **Translation invariance**: for lattice Hamiltonians on translation-invariant lattices. - **Parity** (Z2 symmetries): for many quantum chemistry Hamiltonians after Jordan-Wigner. - **Time-reversal**: for Hamiltonians made of real combinations of Pauli operators.

Consequences for the cost landscape: - The ground state lies in one specific symmetry sector. If the ansatz prepares states outside that sector, those states have higher energy than they "should" — the cost function has local minima outside the relevant sector that the optimizer might mistakenly find. - Symmetry-preserving ansätze (constructed to commute with G) restrict the search to the right sector, removing those spurious minima. - Symmetry-violating ansätze must "discover" the symmetry through optimization, which is generally slow and unreliable.

2. Ansatz Symmetries

If the ansatz $U(\boldsymbol{\theta})$ has a structural symmetry — e.g., parameters θ_i and θ_j are interchangeable — then the cost function inherits a **discrete symmetry group action** on parameter space.

Examples: - **Permutation symmetry of parallel layers**: in an HEA where every layer is identical, swapping the parameters of layer 3 with layer 5 produces a different parameter vector with the same cost. - **Sign-flip symmetry**: $R_Y(\theta)$ and $R_Y(-\theta + 2\pi)$ prepare states differing by a phase. This can be a 1- or 2-element symmetry depending on the surrounding circuit. - **Qubit permutation symmetry**: when the device topology and ansatz are symmetric under qubit relabeling.

Consequences for the cost landscape: - Critical points come in equivalence classes. A local minimum at $\boldsymbol{\theta}^*$ is also a local minimum at every $g \cdot \boldsymbol{\theta}^*$ for $g \in G$. - The number of critical points is multiplied by $|G|$. - Symmetry-related basins are equally good — the optimizer is indifferent between them.

This is the source of the "16 equivalent minima" phenomenon described in node 3.3.

3. Parameterization Redundancy (Continuous Gauge)

Some ansatz parameterizations have **continuous redundancies** — directions in parameter space that don't change the prepared state at all.

Example: $R_X(\theta_1)R_X(\theta_2)|0\rangle$ depends only on the sum $\theta_1 + \theta_2$, not on the individual values. The direction $(\theta_1, \theta_2) \to (\theta_1 + s, \theta_2 - s)$ is a 1-parameter gauge orbit — no physical motion, just relabeling.

Consequences for the cost landscape: - The cost function has flat directions along the gauge orbits. - The Hessian has a kernel; the metric g_{ij} is degenerate (rank deficient). - Any local minimum is part of an entire submanifold of equivalent minima of dimension equal to the gauge orbit dimension. - The natural gradient is undefined in the gauge directions.

These continuous redundancies are pathological — they waste parameters and slow optimization without producing any genuine minimum structure. Good ansatz design avoids them where possible.

How To Use Symmetry

The good move is: **identify symmetries and exploit them to shrink the search space.**

For Hamiltonian symmetries:

- **Restrict the ansatz** to states obeying the symmetry. For particle number, this means using particle-number-preserving gates. For parity, this means using only even Pauli rotations.
- **Encode the symmetry into the cost function** as a penalty. Add $\mu(\langle G\rangle - g_0)^2$ to the energy expectation, where g_0 is the desired symmetry charge. Pushes the optimizer toward the correct sector.
- **Project the state** onto the symmetric subspace before measurement. Equivalent to ground-state preparation in a smaller Hilbert space.

For ansatz symmetries:

- **Quotient out the symmetry** — use a parameterization where redundant directions are explicitly removed. Sometimes possible by clever ansatz design.
- **Symmetry-aware initialization** — start the optimizer at a representative point of each equivalence class to ensure broad exploration.
- **Symmetry-aware aggregation** — when reporting "the converged parameters," report the equivalence class rather than a specific representative.

For continuous gauge:

- **Eliminate the gauge** by reparameterizing. $R_X(\theta_1)R_X(\theta_2) \to R_X(\theta_1 + \theta_2)$.
- If you can't eliminate it (gauges arising from circuit composition aren't always removable), use the **natural gradient** (Module 6) which automatically projects out gauge directions.

How Symmetry Misleads

The bad move is to **ignore symmetry**. Three failure modes:

1. Optimizer wastes effort on gauge. A vanilla gradient descent on a gauge-redundant parameterization spends real shots evaluating the gradient in gauge directions, where it always reads zero plus noise. The optimizer is "moving" in those directions but the state isn't.

2. Symmetry-broken minima look like progress. If the cost function has a Hamiltonian symmetry but the ansatz doesn't preserve it, the optimizer will find local minima at symmetry-broken states. These look like "good progress" (the cost is decreasing) but the converged state has the wrong quantum numbers — it is not the physical ground state.

3. Spurious basins from broken parameterization symmetries. If your ansatz has a discrete symmetry the cost function does not, then the symmetry-related "minima" are not actually equivalent — they have slightly different costs because of the asymmetry — and the optimizer can get confused about which one to commit to.

The right diagnostic for all of these is to **inspect the converged state**, not just the converged cost. Compute the expectation values of the symmetry generators. If they don't match the expected charges, you have a symmetry-related bug.

Symmetry and Module 5

Module 5's regularization framework has a clean interaction with symmetry:

- **Prior injection (objective-side)** can encode symmetry as a penalty term added to the cost.
- **Dynamical modification (dynamics-side)** can encode symmetry as a constraint on the optimizer's update rule (project the gradient onto the symmetric subspace).

These produce *different* converged states even when nominally enforcing the same symmetry. The thesis's argument that these two regularization styles have different semantics applies in the symmetry case as well — perhaps especially clearly there.

Structural Role

This is the **structural awareness** node of Module 3. It rounds out the topology discussion (3.2-3.3) with the observation that not all topology is independent — structural relations between minima exist and can be exploited.

Forwards: every symmetry-aware ansatz design choice in Module 2.3 (chemically-inspired ansätze) is consistent with this node.

Relations to Other Concepts

- **Group theory** — the mathematical setting for discrete symmetries.
- **Gauge theory** — the formal treatment of continuous parameterization redundancies.
- **Equivariant neural networks** — the classical ML analogue of symmetry-preserving ansätze.
- **Symmetry breaking in physics** — the process by which a symmetric Hamiltonian acquires a non-symmetric ground state. Different phenomenon, often confused.
- **Quotient space optimization** — optimizing on the quotient by a symmetry group.

Worked Example — Sign-Flip Symmetry of R_Y

Consider $|\psi(\theta)\rangle = R_Y(\theta)|0\rangle = \cos(\theta/2)|0\rangle + \sin(\theta/2)|1\rangle$.

For Hamiltonian $H = X$, the cost is

$$C(\theta) = \langle\psi|X|\psi\rangle = 2\cos(\theta/2)\sin(\theta/2) = \sin(\theta).$$

This has minima at $\theta = -\pi/2$ and $\theta = 3\pi/2$, and these differ by 2π — they are the **same point on the torus S^1**. One minimum, no symmetry.

Now consider $H = Z$. The cost is

$$C(\theta) = \cos(\theta).$$

Minimum at $\theta = \pi$. Maximum at $\theta = 0$. The function has reflection symmetry $\theta \to -\theta$, but $-\pi = \pi$ on the torus, so the symmetry has just one fixed point and no equivalence class structure. Again one minimum.

Now consider $H = X^2 + Z^2 = 2I$ (a useless Hamiltonian, but pedagogically clean): the cost is constant. Every point is equivalent. The "symmetry" is the entire group acting trivially.

These three examples show that the *manifest* symmetry of the cost function depends jointly on the ansatz and the Hamiltonian. **Symmetry isn't a property of the ansatz alone or the Hamiltonian alone; it's a property of their composition.**

Visualization

Pending. A 2D parameter space colored by cost, with discrete symmetry orbits highlighted (e.g., 4-fold permutation orbits). Multiple equivalent minima shown connected by group action arrows. Caption: "Symmetry $\times$ cost function = orbit structure."

Exercises

1. The Hamiltonian $H = Z_0 + Z_1$ is invariant under qubit swap. The ansatz $R_Y(\theta_1) \otimes R_Y(\theta_2)|00\rangle$ is also invariant under qubit swap. Show that the cost function has a $\mathbb{Z}_2$ symmetry, and characterize the orbits.
2. Design a symmetry-violating perturbation to the above Hamiltonian and show that the symmetry of the cost function is broken.
3. (VQA) Explain why a symmetry-preserving ansatz can sometimes have **fewer** local minima than a symmetry-breaking one. Is fewer minima always good?

Research Extensions

- **Equivariant VQAs** — ansätze whose action commutes with a target symmetry group by construction.
- **Symmetry-protected initial states** — restricting the reference state to lie in a specific symmetry sector.
- **Symmetry detection algorithms** — automatic discovery of symmetries in arbitrary Hamiltonians.
- **Quotient-space optimization** — optimizing on the parameter space modulo the symmetry group.

Cross-links to Other Courses

- **Group Theory** — discrete and continuous symmetries.
- **Differential Geometry** — quotient manifolds, principal bundles.
- **Equivariant ML** — neural networks invariant under group actions.
- **Module 2.3** — chemically-inspired ansätze are symmetry-aware by construction.

Variational Quantum Algorithms · Module 3 — Cost Landscapes · Node 3.4
Concepts: optimization-landscape, eigenvalue-decomposition, cost-function
Prerequisites: 3.2, 3.3

hbar.university — own.your.cognition