Geometric Regularization

Variational Quantum Algorithms · Module 5 — Regularization

Node 5.2

Position in Knowledge Graph

System: Variational Quantum Algorithms **Structural Domain:** Regularization **Node:** 5.2

This node specializes the prior-injection framework of node 5.1 to the case where the prior is **geometric** — that is, defined on the manifold of reachable quantum states rather than on the raw parameter vector.

Formal Definition

Let $\mathcal{M} = \{|\psi(\boldsymbol{\theta})\rangle : \boldsymbol{\theta} \in \mathbb{R}^d\}$ be the ansatz manifold (see node 2.9).

A **geometric regularizer** is a penalty function

$$R_{\text{geom}}(\boldsymbol{\theta}) = d_{\mathcal{M}}\big(|\psi(\boldsymbol{\theta})\rangle,\, |\psi_0\rangle\big)^2$$

where $d_{\mathcal{M}}$ is a distance function on $\mathcal{M}$ and $|\psi_0\rangle$ is a reference state (often the previous iterate).

The most common choice is the **Fubini-Study metric**:

$$d_{\text{FS}}(|\psi\rangle, |\phi\rangle) = \arccos|\langle\psi|\phi\rangle|.$$

The regularized cost becomes

$$C_{\text{reg}}(\boldsymbol{\theta}) = C(\boldsymbol{\theta}) + \lambda\, d_{\text{FS}}(|\psi(\boldsymbol{\theta})\rangle, |\psi(\boldsymbol{\theta}_k)\rangle)^2,$$

evaluated at iteration k.

Intuition

L2 regularization on $\boldsymbol{\theta}$ penalizes large parameter values. But two parameter vectors that differ a lot in $\boldsymbol{\theta}$-space may produce nearly identical quantum states — and conversely. **The geometry of the parameter manifold and the geometry of the state manifold do not agree.**

Geometric regularization fixes this by penalizing motion in *state space*, which is the physically meaningful place.

This is the same idea that motivates **quantum natural gradient** (node 6.2) — both replace the Euclidean parameter metric with the Fubini-Study metric on the state manifold.

Structural Role

Geometric regularization is the **manifold-aware** member of the standard regularization family. It connects Module 5 to Module 2 (the ansatz manifold) and Module 6 (natural gradient).

It is also a stepping stone toward node 5.5: once you accept that the parameter metric is the wrong unit of "distance," you are halfway to seeing regularization as something the optimizer dynamics need to consume — not a property of an idealized objective.

Relations to Other Concepts

- **Quantum Fisher Information** (node 6.2) — the metric tensor whose square root is the local Fubini-Study distance.
- **Trust regions** (node 5.6) — bounds the step size in some metric. When the metric is FS, this is geometric regularization in disguise.
- **Mirror descent** — generalized gradient methods that respect a non-Euclidean parameter geometry.

Worked Derivations

Fubini-Study to leading order

For an infinitesimal step $\delta\boldsymbol{\theta}$ from $\boldsymbol{\theta}_k$, the Fubini-Study distance squared is

$$d_{\mathrm{FS}}^2 = \delta\boldsymbol{\theta}^\top F(\boldsymbol{\theta}_k)\,\delta\boldsymbol{\theta} + O(\|\delta\boldsymbol{\theta}\|^4),$$

where F is the quantum Fisher information matrix. At first order in $\delta\boldsymbol{\theta}$, geometric regularization is equivalent to a **state-space Tikhonov penalty** with the QFI matrix as the metric.

Equivalence to natural gradient damping

A geometrically regularized step minimizes

$$C(\boldsymbol{\theta}_k + \delta\boldsymbol{\theta}) + \lambda\,\delta\boldsymbol{\theta}^\top F(\boldsymbol{\theta}_k)\,\delta\boldsymbol{\theta}.$$

Setting the gradient to zero gives

$$\delta\boldsymbol{\theta} = -(2\lambda F + H)^{-1}\nabla C,$$

where H is the Hessian of C. For small H, this approaches the **damped natural gradient step** $-(2\lambda F)^{-1}\nabla C$.

Visualization

Pending. Diagram: parameter space (Euclidean grid) on the left, state manifold (curved surface) on the right, with the same parameter step producing different state distances depending on the local curvature.

Exercises

1. For a single-qubit ansatz $|\psi(\theta)\rangle = R_y(\theta)|0\rangle$, compute the Fubini-Study distance between $\theta = 0$ and $\theta = \pi$. Compare to the Euclidean distance $|\pi - 0| = \pi$.
2. Show that geometric regularization with the FS metric is invariant under reparameterization of the ansatz, while L2 is not.
3. Implement geometric regularization for a 2-qubit hardware-efficient ansatz and compare convergence to plain L2 on a small VQE instance.

Research Extensions

- **Information-geometric optimization** — full theory in Amari's framework.
- **Riemannian VQE** — treating the optimization as a Riemannian descent on the state manifold from the start.
- **Geodesic regularization** — penalizing total geodesic length traversed, not just final distance.

Cross-links to Other Courses

- **Statistics Module 5** — Information Geometry (Fisher information as a Riemannian metric).
- **Linear Algebra** — Riemannian metrics on submanifolds.

Variational Quantum Algorithms · Module 5 — Regularization · Node 5.2
Concepts: regularization, inner-product, parameterized-circuits
Prerequisites: 2.9, 5.1
Enables: 5.5

hbar.university — own.your.cognition