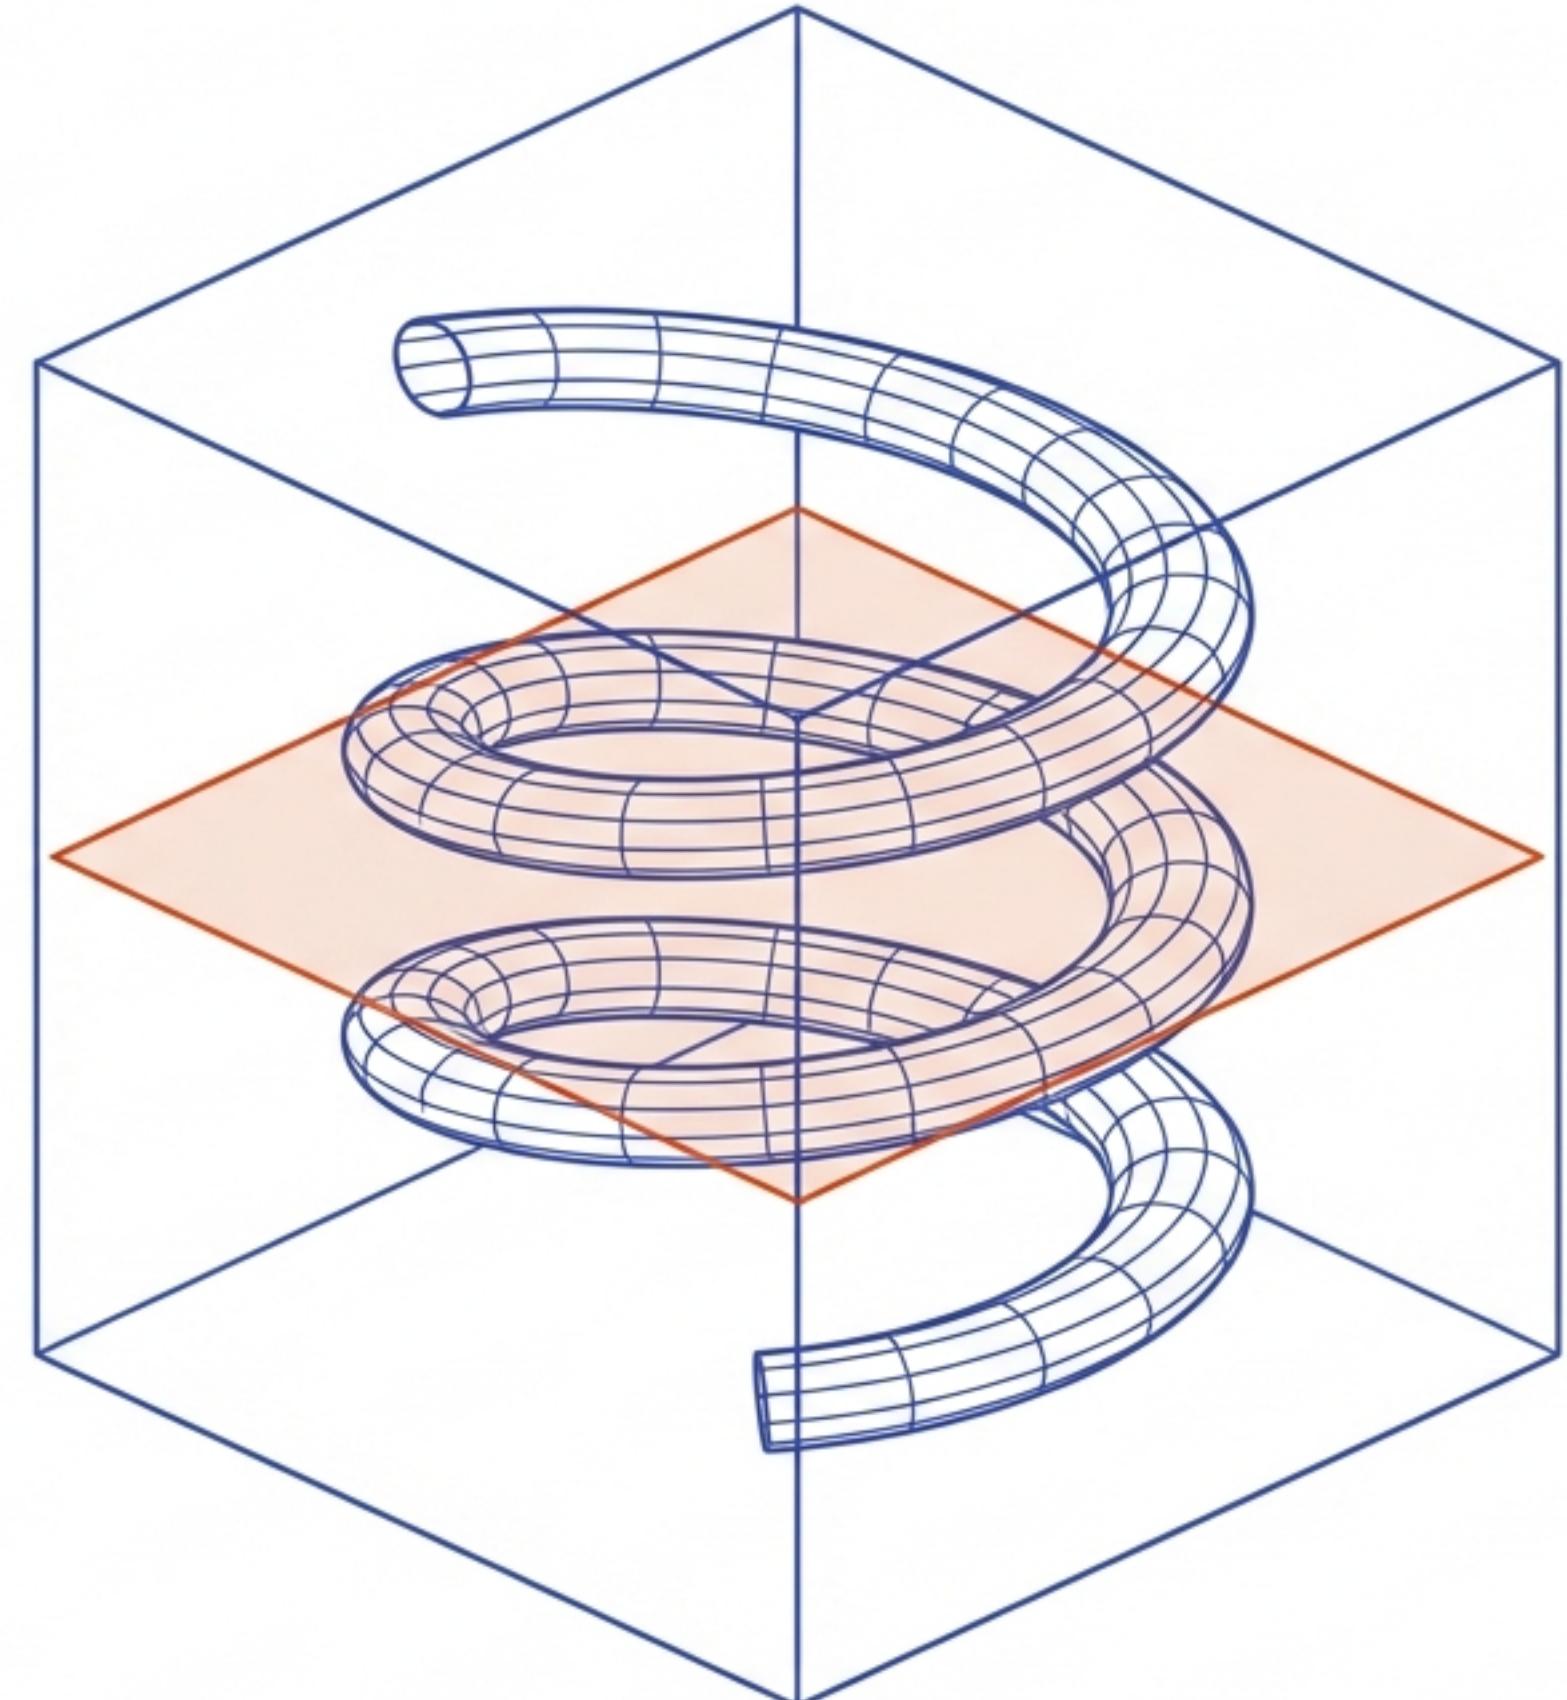


Optimization Landscape and Feasibility in Updated Riemannian AmbientFlow.

A Systematic Empirical Investigation
of the Landscape-Feasibility Trade-off.



Based on 'Optimization Landscape and Feasibility in Updated
Riemannian AmbientFlow' (Anonymous Authors).

Executive Summary: The Landscape is the Culprit.

Helvetica Now Display
THE GOAL.

Investigate if theoretical assumptions (F1, F2, F3) of Riemannian AmbientFlow hold in practice using synthetic "torture tests" (Circle, Sphere, Helix).

Helvetica Now Display
THE VERDICT.

Feasibility assumptions fail. Increasing geometric regularization (Λ) actively degrades data matching.



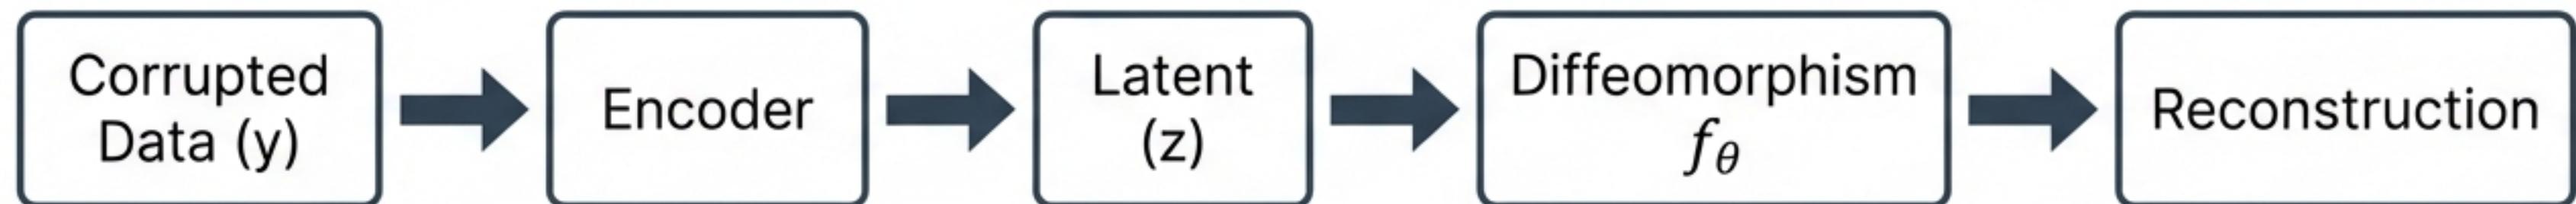
Helvetica Now Display
THE ALIBI.

It is NOT a capacity issue. Oracle experiments prove the architectures can represent the manifolds perfectly.

Helvetica Now Display
THE CAUSE.

The non-convex optimization landscape traps the model in **local minima** where the metric is systematically underestimated.

The Subject: Riemannian AmbientFlow



$$\mathcal{L}(\theta, \phi) = \text{ELBO} + \lambda \|\mathbf{J}f_{\theta\theta}(\theta)\|_F^2$$

Variational Lower Bound
(Data Fit)

Geometric Regularization
(Penalizes Stretching)

Objective: To learn low-dimensional manifolds from noisy ambient data by enforcing geometric simplicity at the origin.

The Theoretical Promise: The Recoverability Theorem

Theory guarantees recovery IF three conditions are met at the solution

Condition F1: Data Matching

Learned distribution
approx Ground Truth

$$p_{\text{theta}} = p_{\text{data}}$$

Condition F2: Posterior Matching

Learned posterior approx
True posterior

$$q_{\text{phi}} = p_{\text{theta}}(\mathbf{z}|\mathbf{y})$$

Condition F3: Geometric Constraint

Jacobian norm is bounded
by the true constant

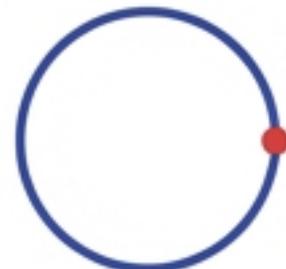
$$\|Jf_{\text{theta}}(0)\|_F^2 \leq C^*$$

The Question: Do gradient-based optimizers actually find these solutions?



The Interrogation Room: Experimental Design

The Suspects (The Manifolds)



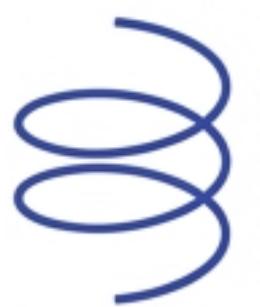
Circle (R^2)

$C^* = 1.0$. Simple topology.



Sphere (R^3)

$C^* = 8.0$. High curvature via stereographic projection.



Helix (R^3)

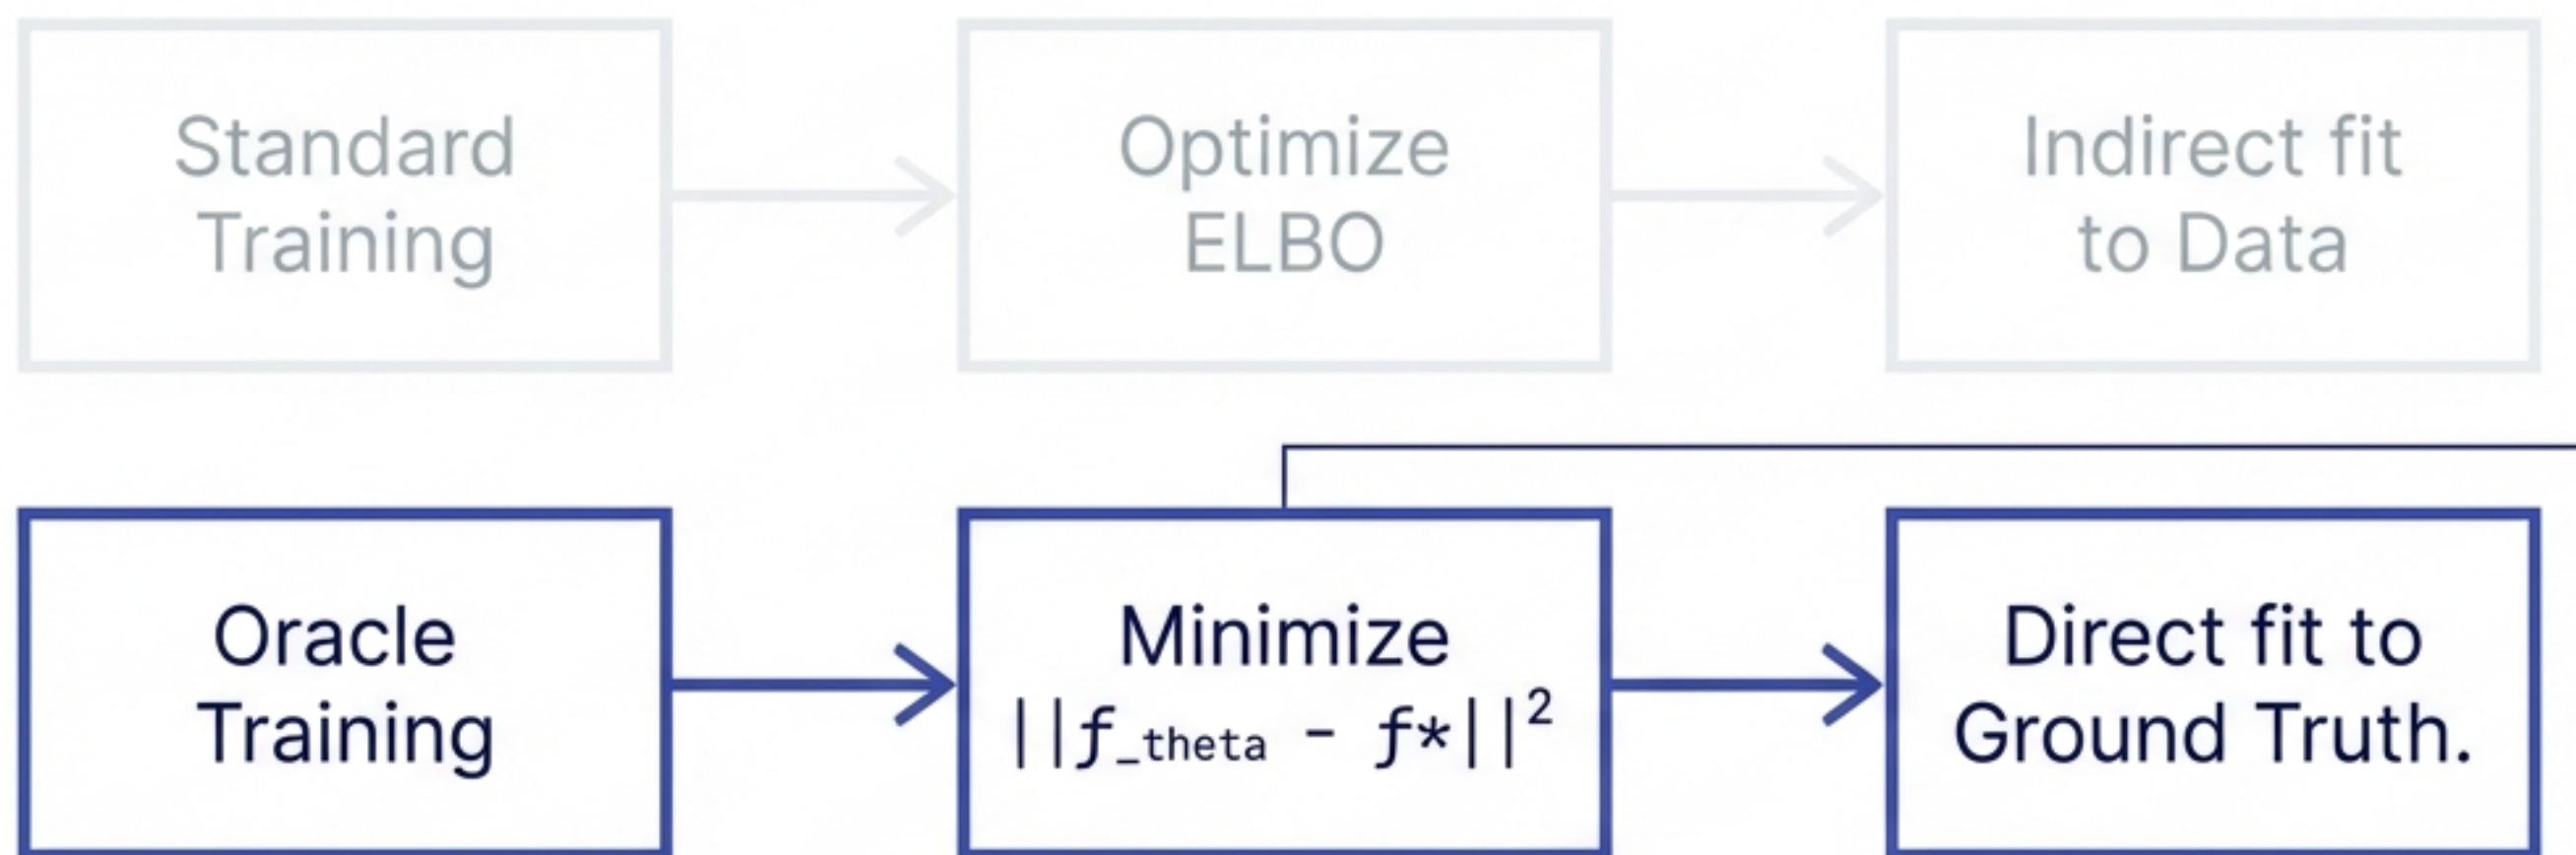
C^* approx 1.025. Non-compact geometry.

The Stress Test Parameters

Parameter Class	Value/Description
Corruption	Gaussian Noise ($\sigma = 0.1$)
Optimizer	L-BFGS-B (200 Iterations)
Architectures	Simple (Affine + tanh, 9-19 params) vs. MLP (2 hidden layers, ~1200 params)

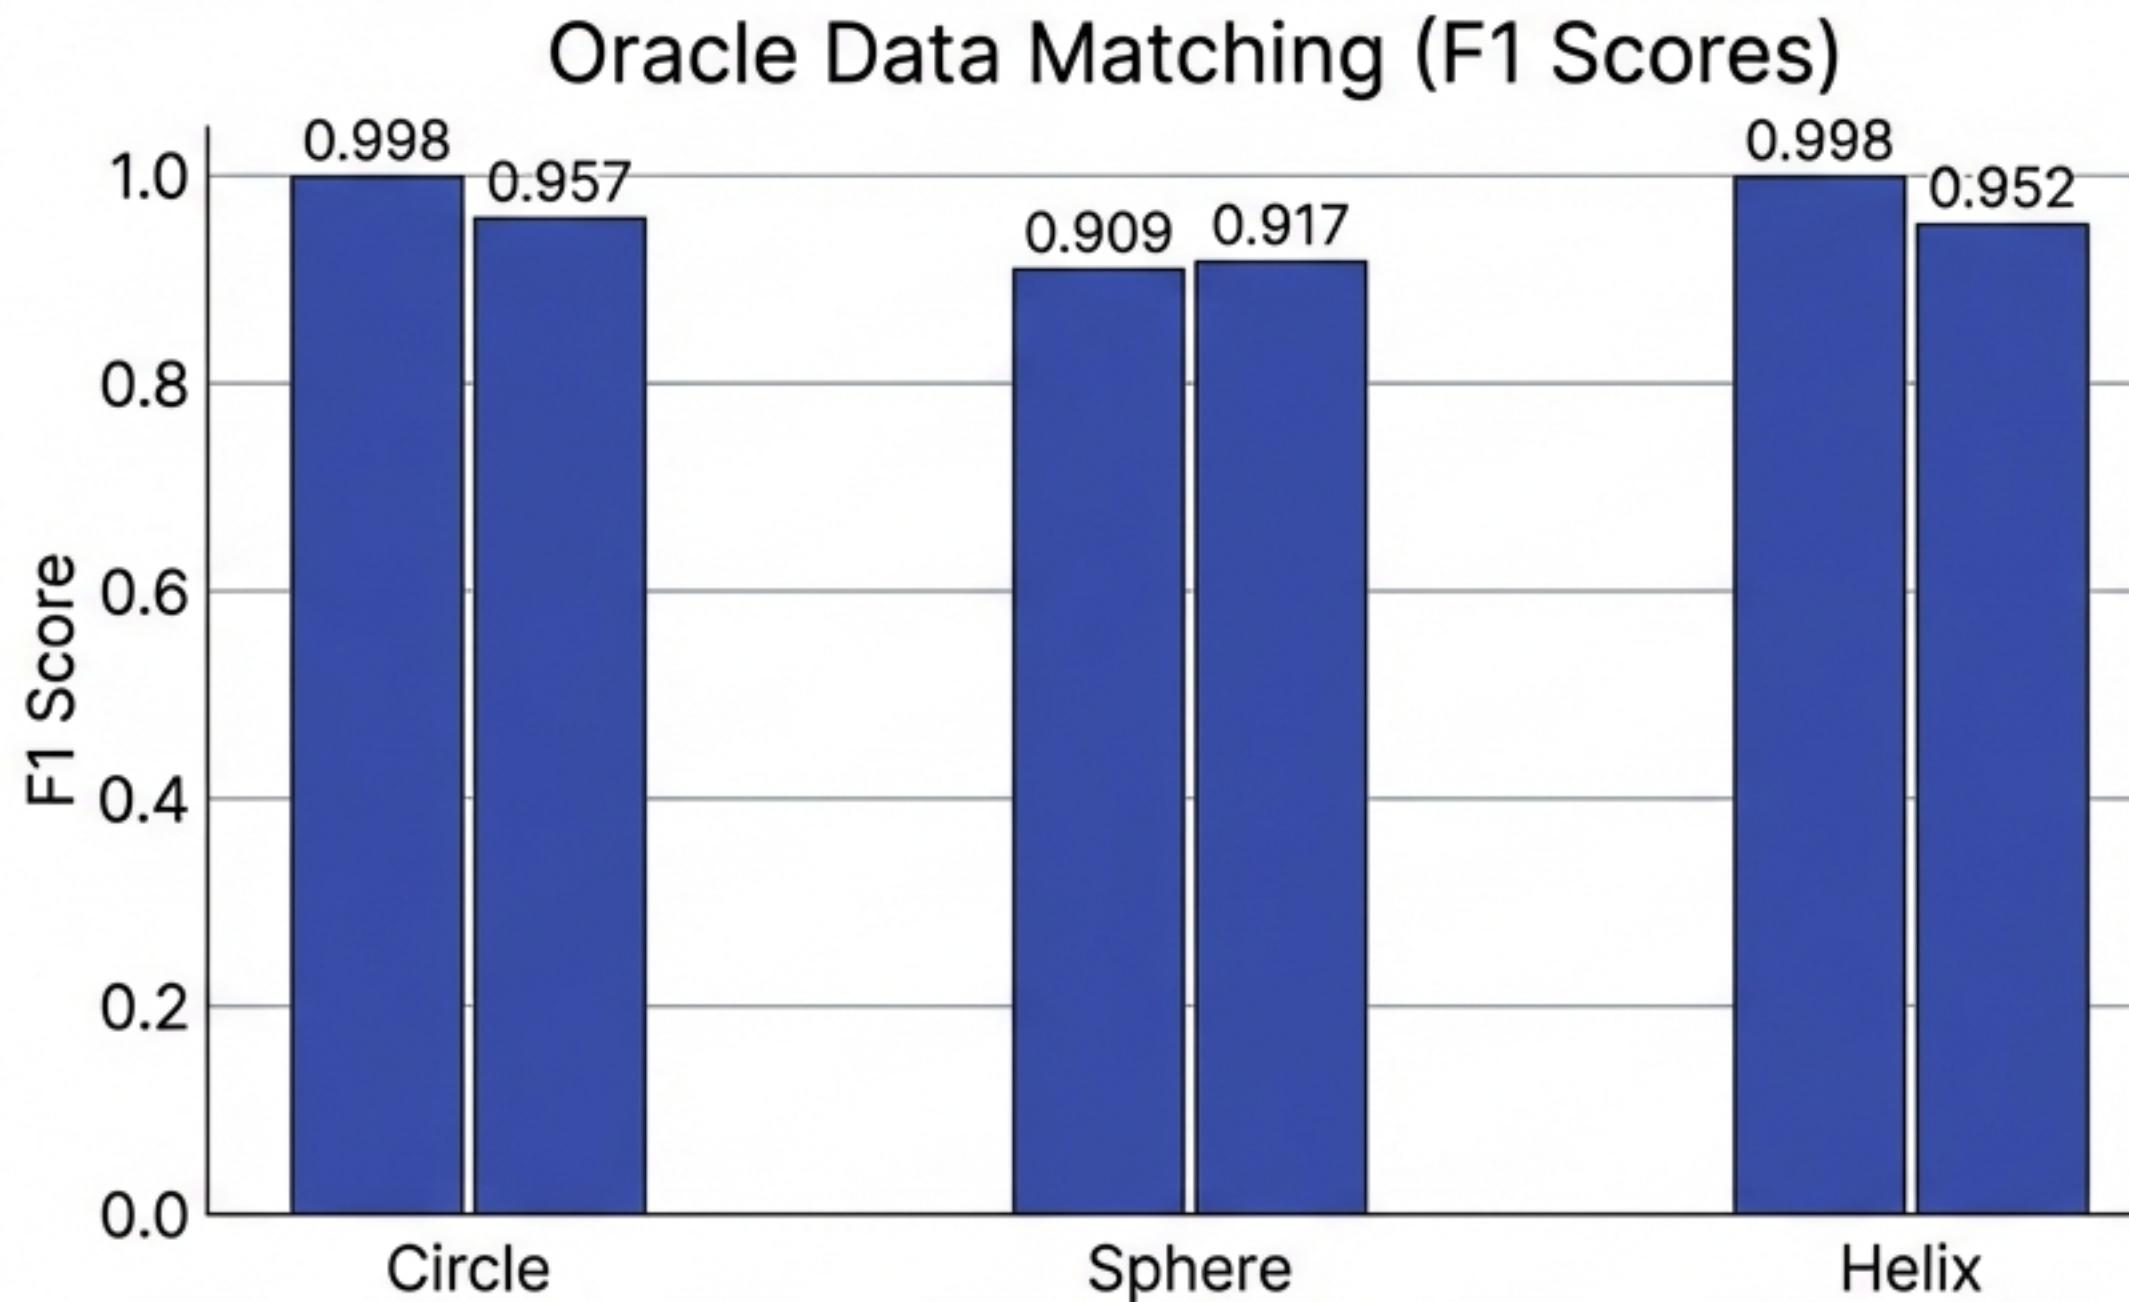
Ruling Out Incompetence: The Oracle Experiment

Before blaming the optimization landscape, we must ensure the model has the Capacity to learn the manifold.



Testing if the model CAN represent the solution, bypassing the optimization difficulty.

Verdict on Capacity: The Models Are Capable

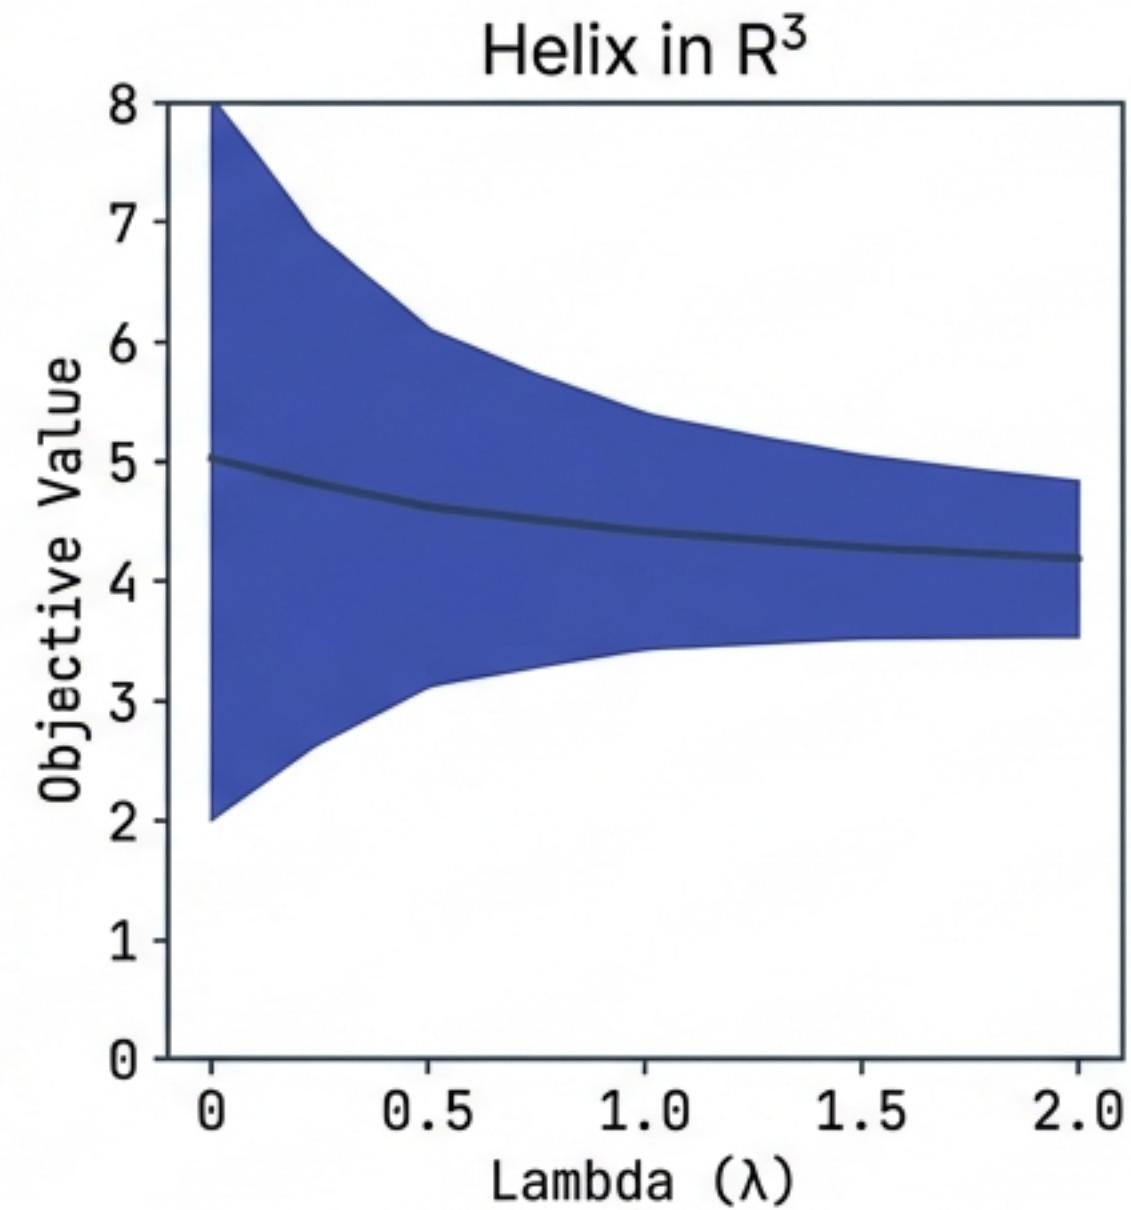
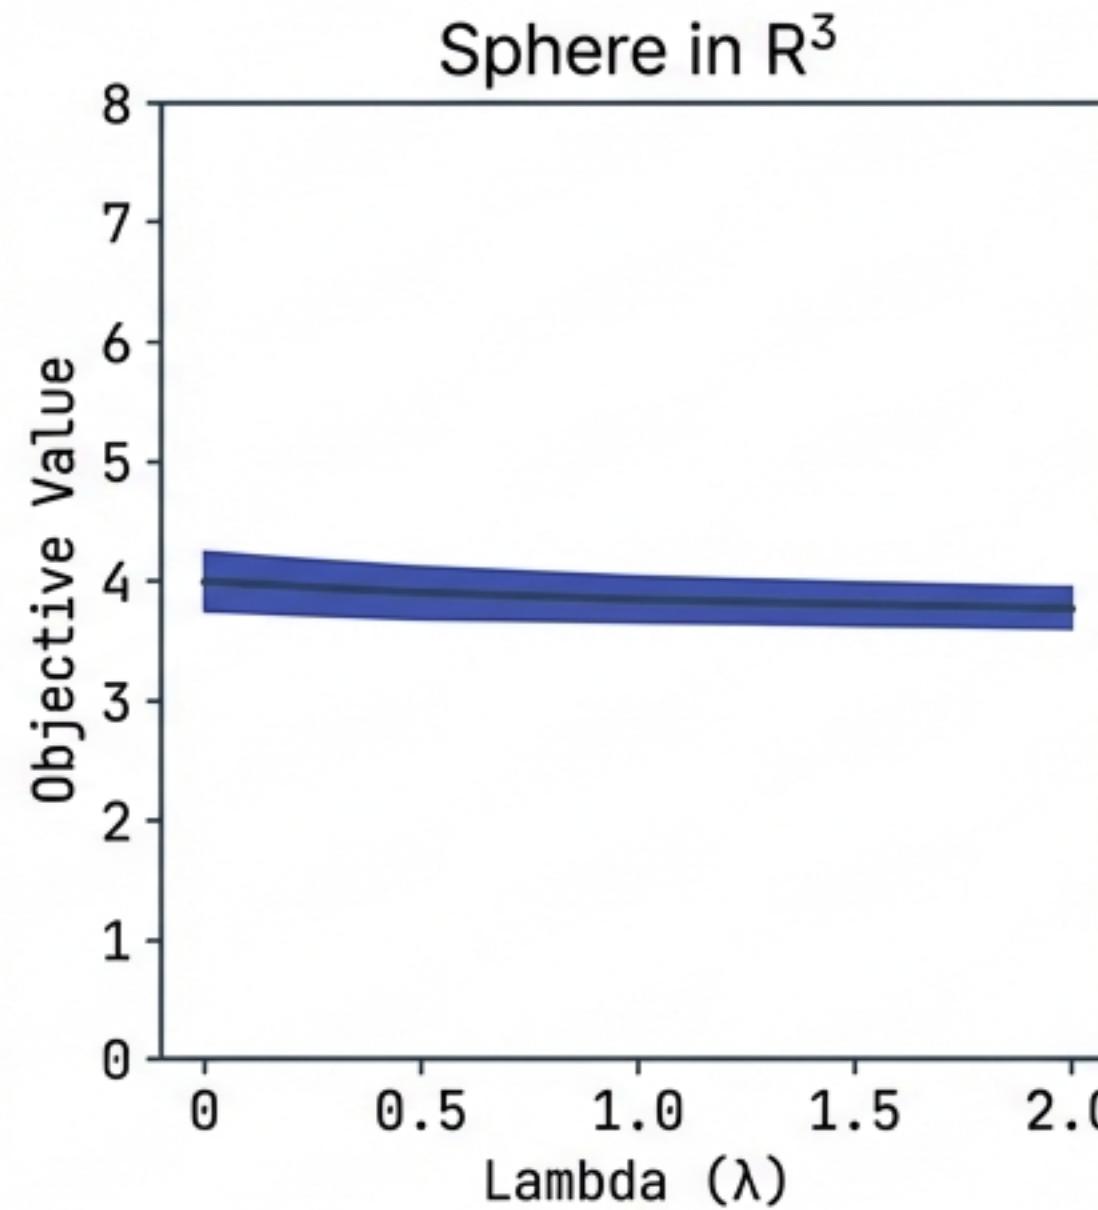
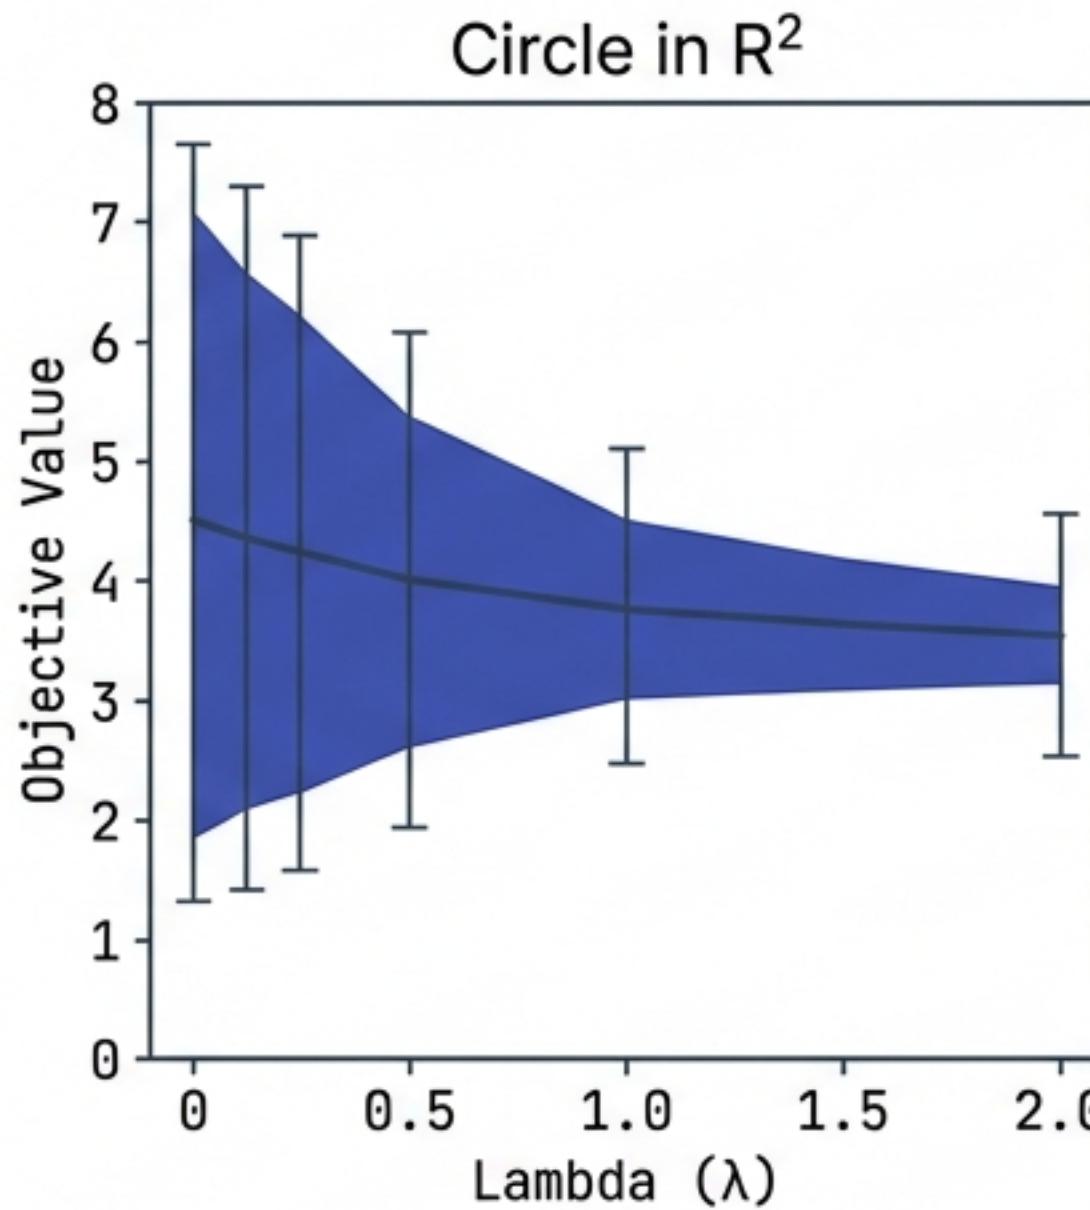


⚠️ Caveat: The Sphere model underestimates the metric trace (~3.2 vs 8.0) even at Oracle, suggesting intrinsic geometric difficulty.

Implication: If ELBO training fails later, it is a LANDSCAPE problem, not a capacity problem.

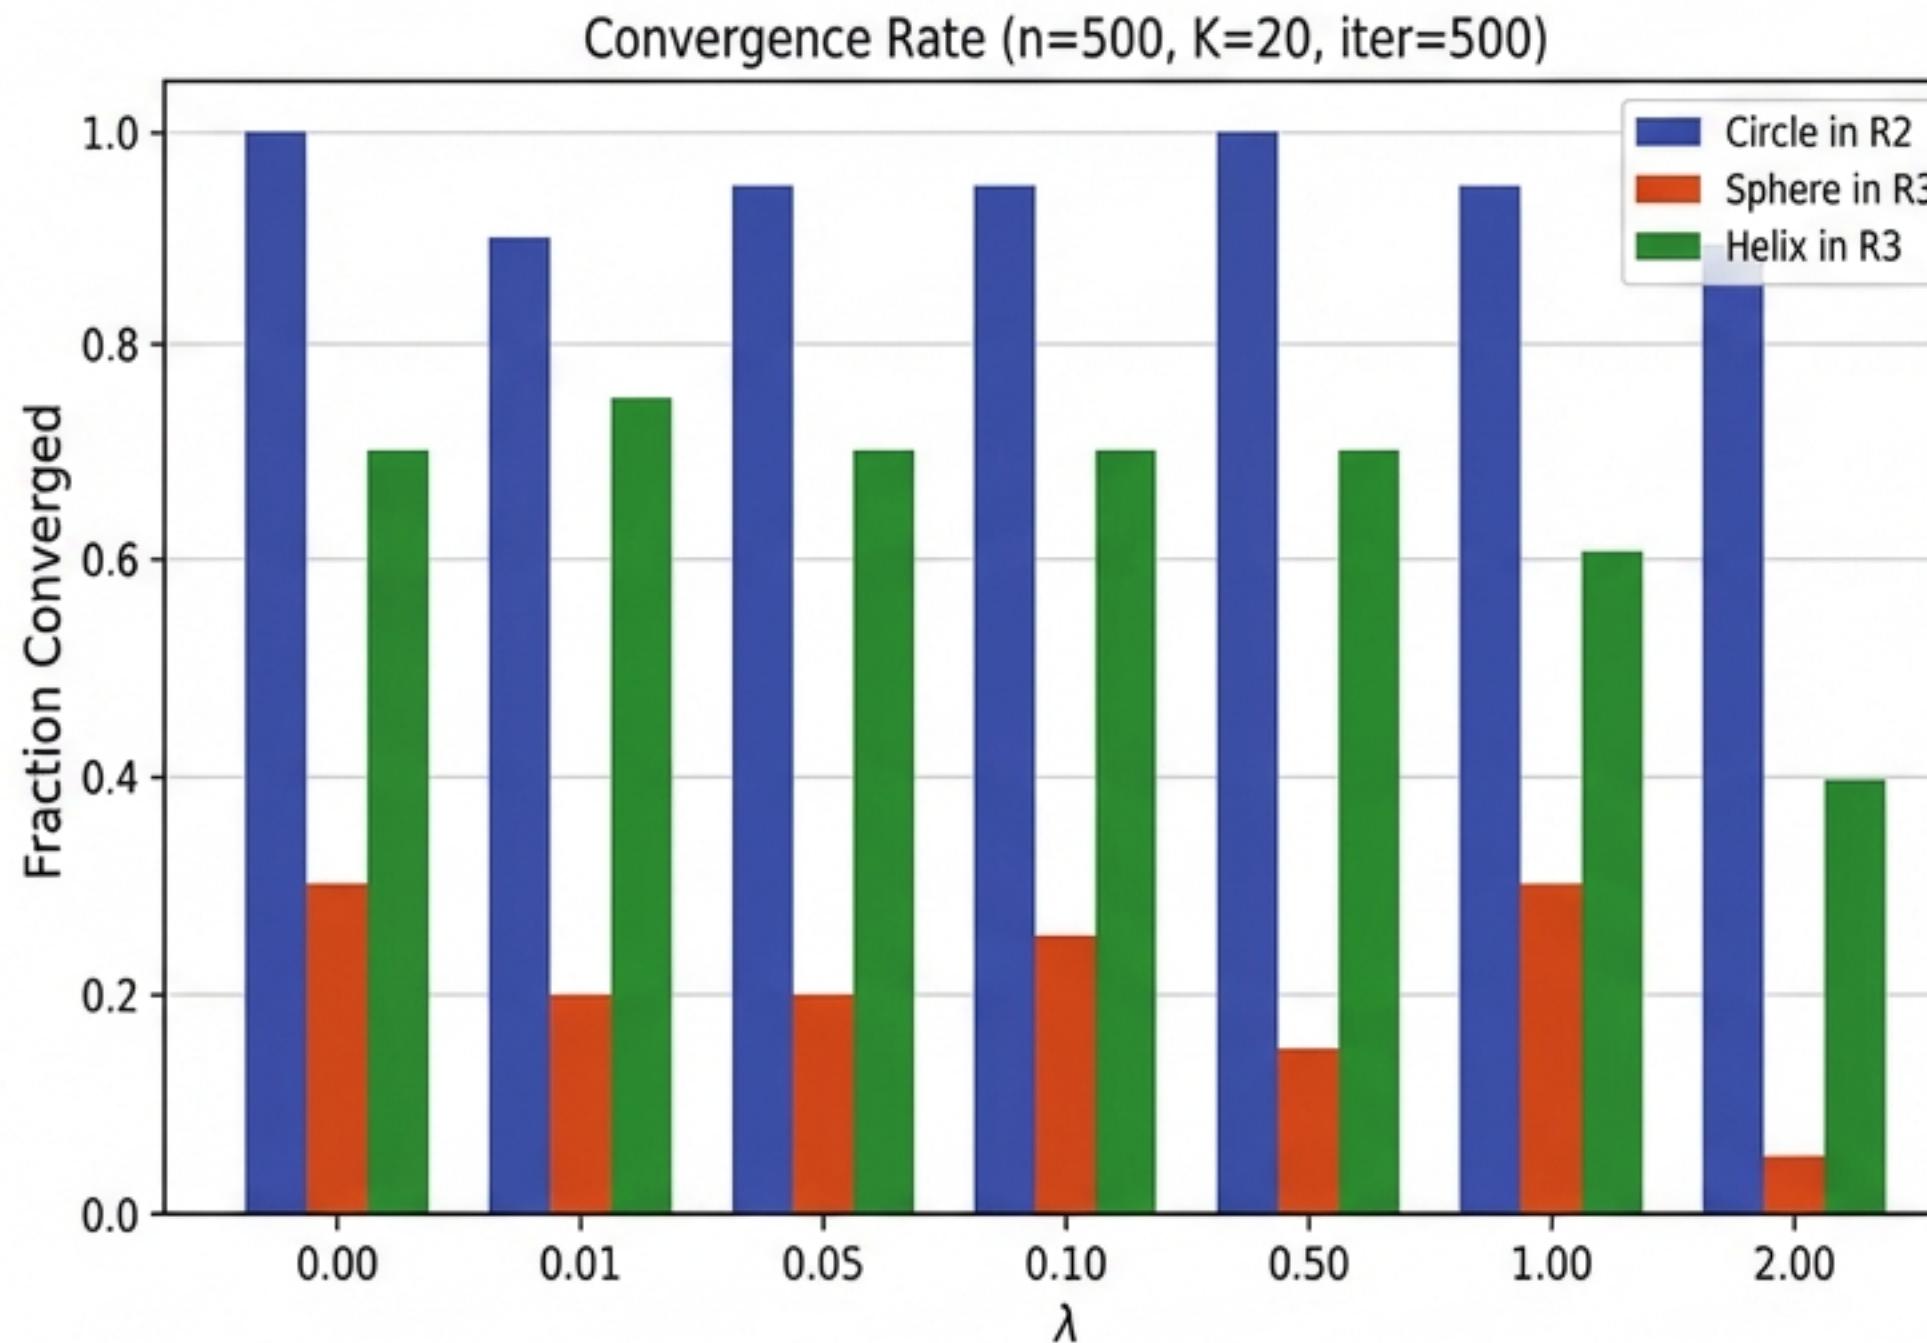
Landscape Exploration: Instability at Low Regularization

Objective Value vs. Regularization Strength



High variance at Lambda=0 proves the existence of multiple distinct basins of attraction.

Convergence Diagnostics: Stalled, Not Lost.



93%
of runs hit the
200-iteration limit.

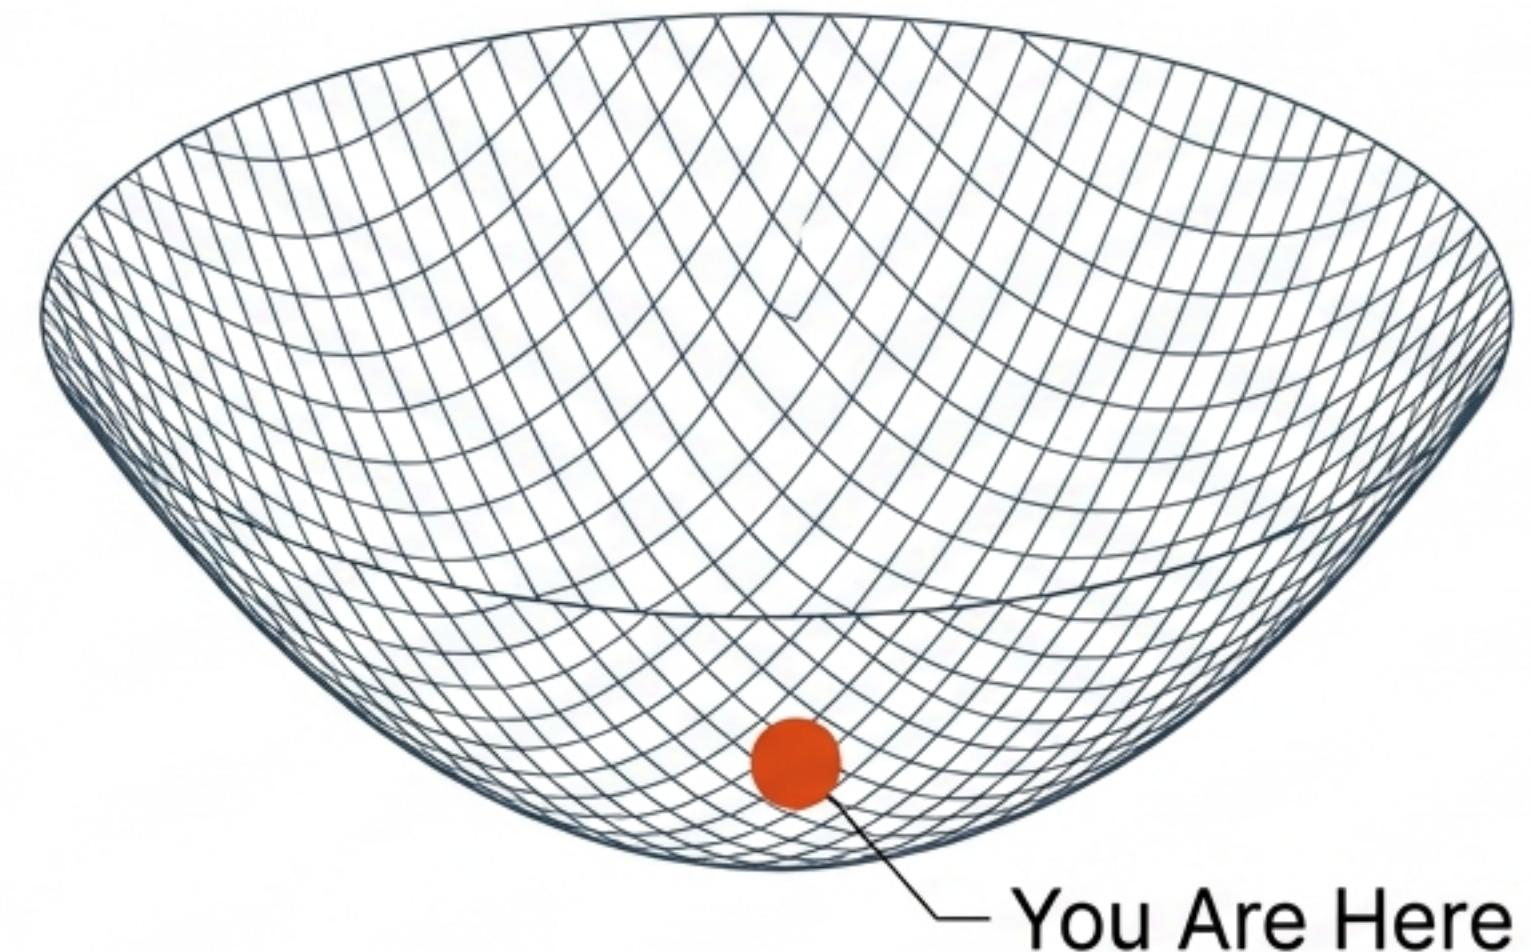
However, final gradient norms are small (median < 0.3). The optimizer reaches near-stationary points but struggles to settle completely.

Hessian Spectral Analysis: Confirmed Local Minima.

Are we stuck in saddle points or true traps?

- Method: Sampled 50 directional second derivatives ($v^T H v$) at converged solutions.

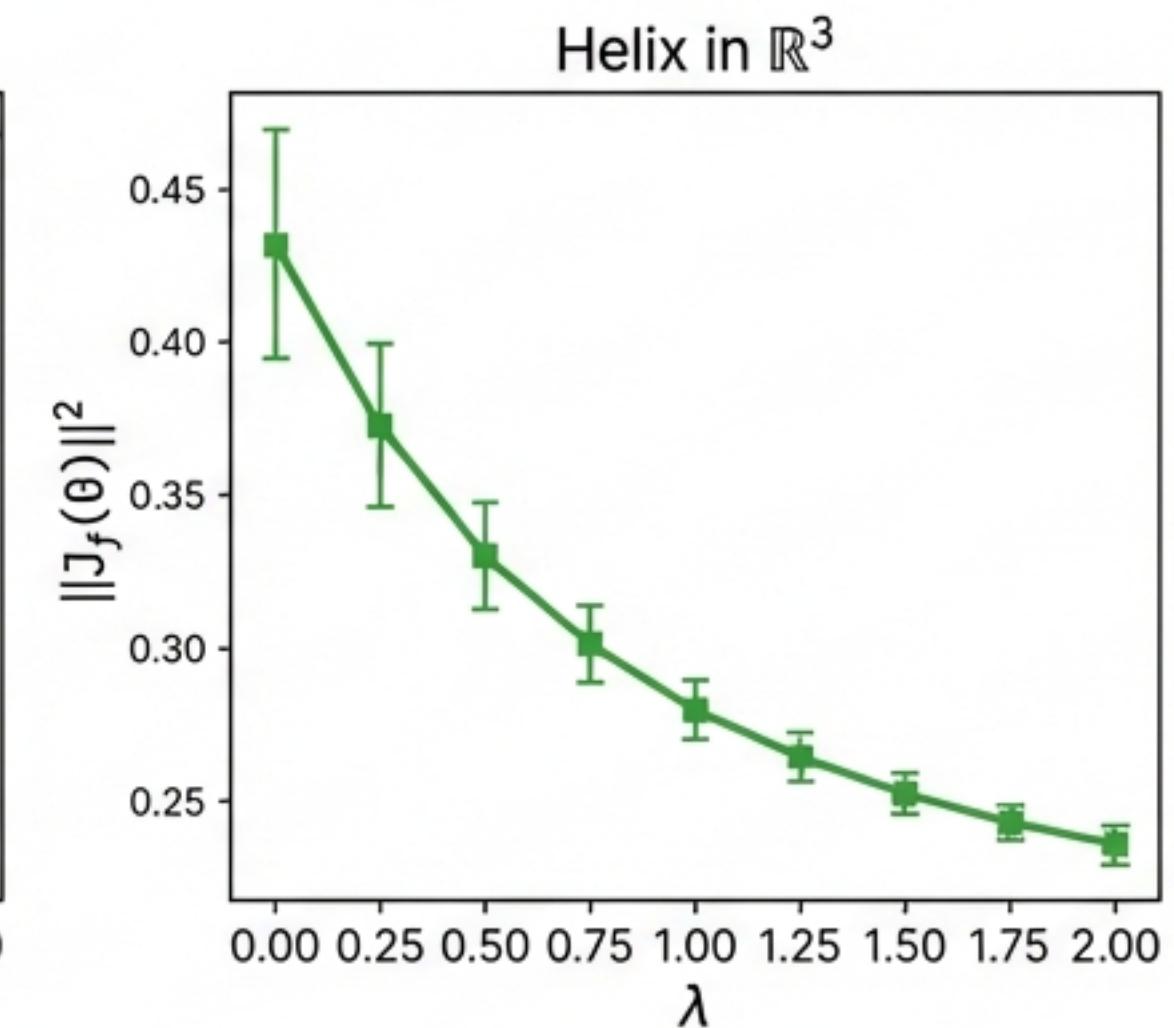
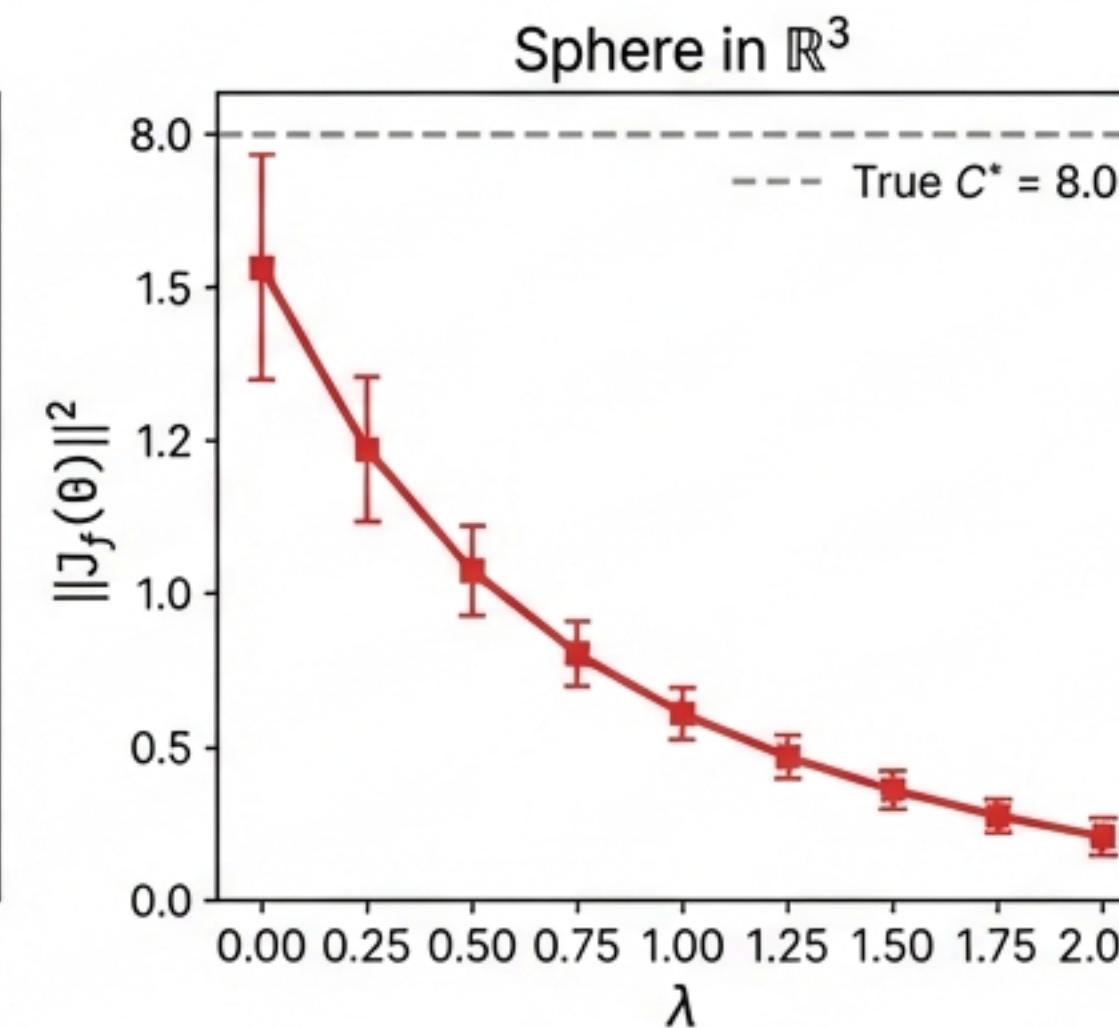
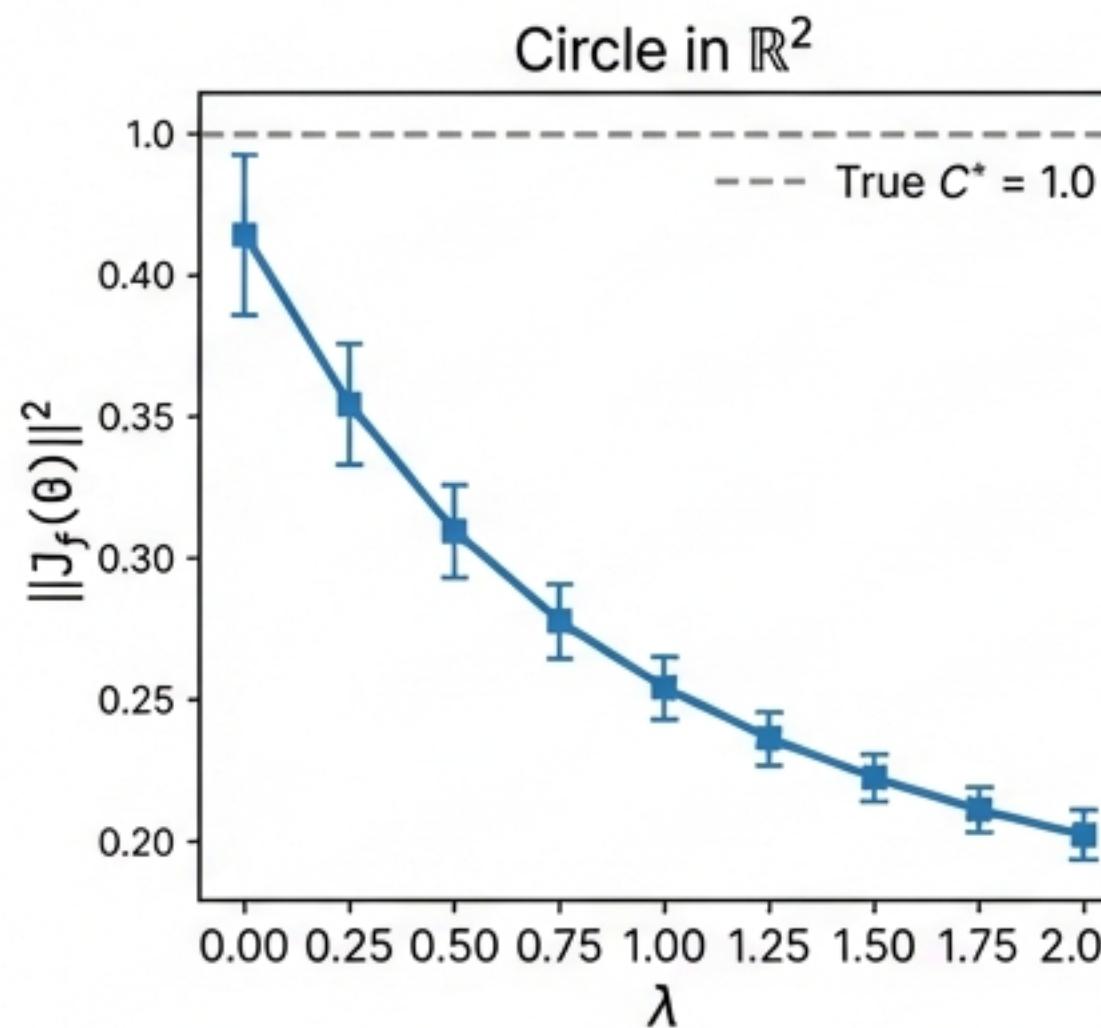
Result: 0 negative curvature directions detected across 600+ samples.



Verdict: We are stuck in genuine local minima. The infeasibility is stable.

The Regularization Works (Perhaps Too Well)

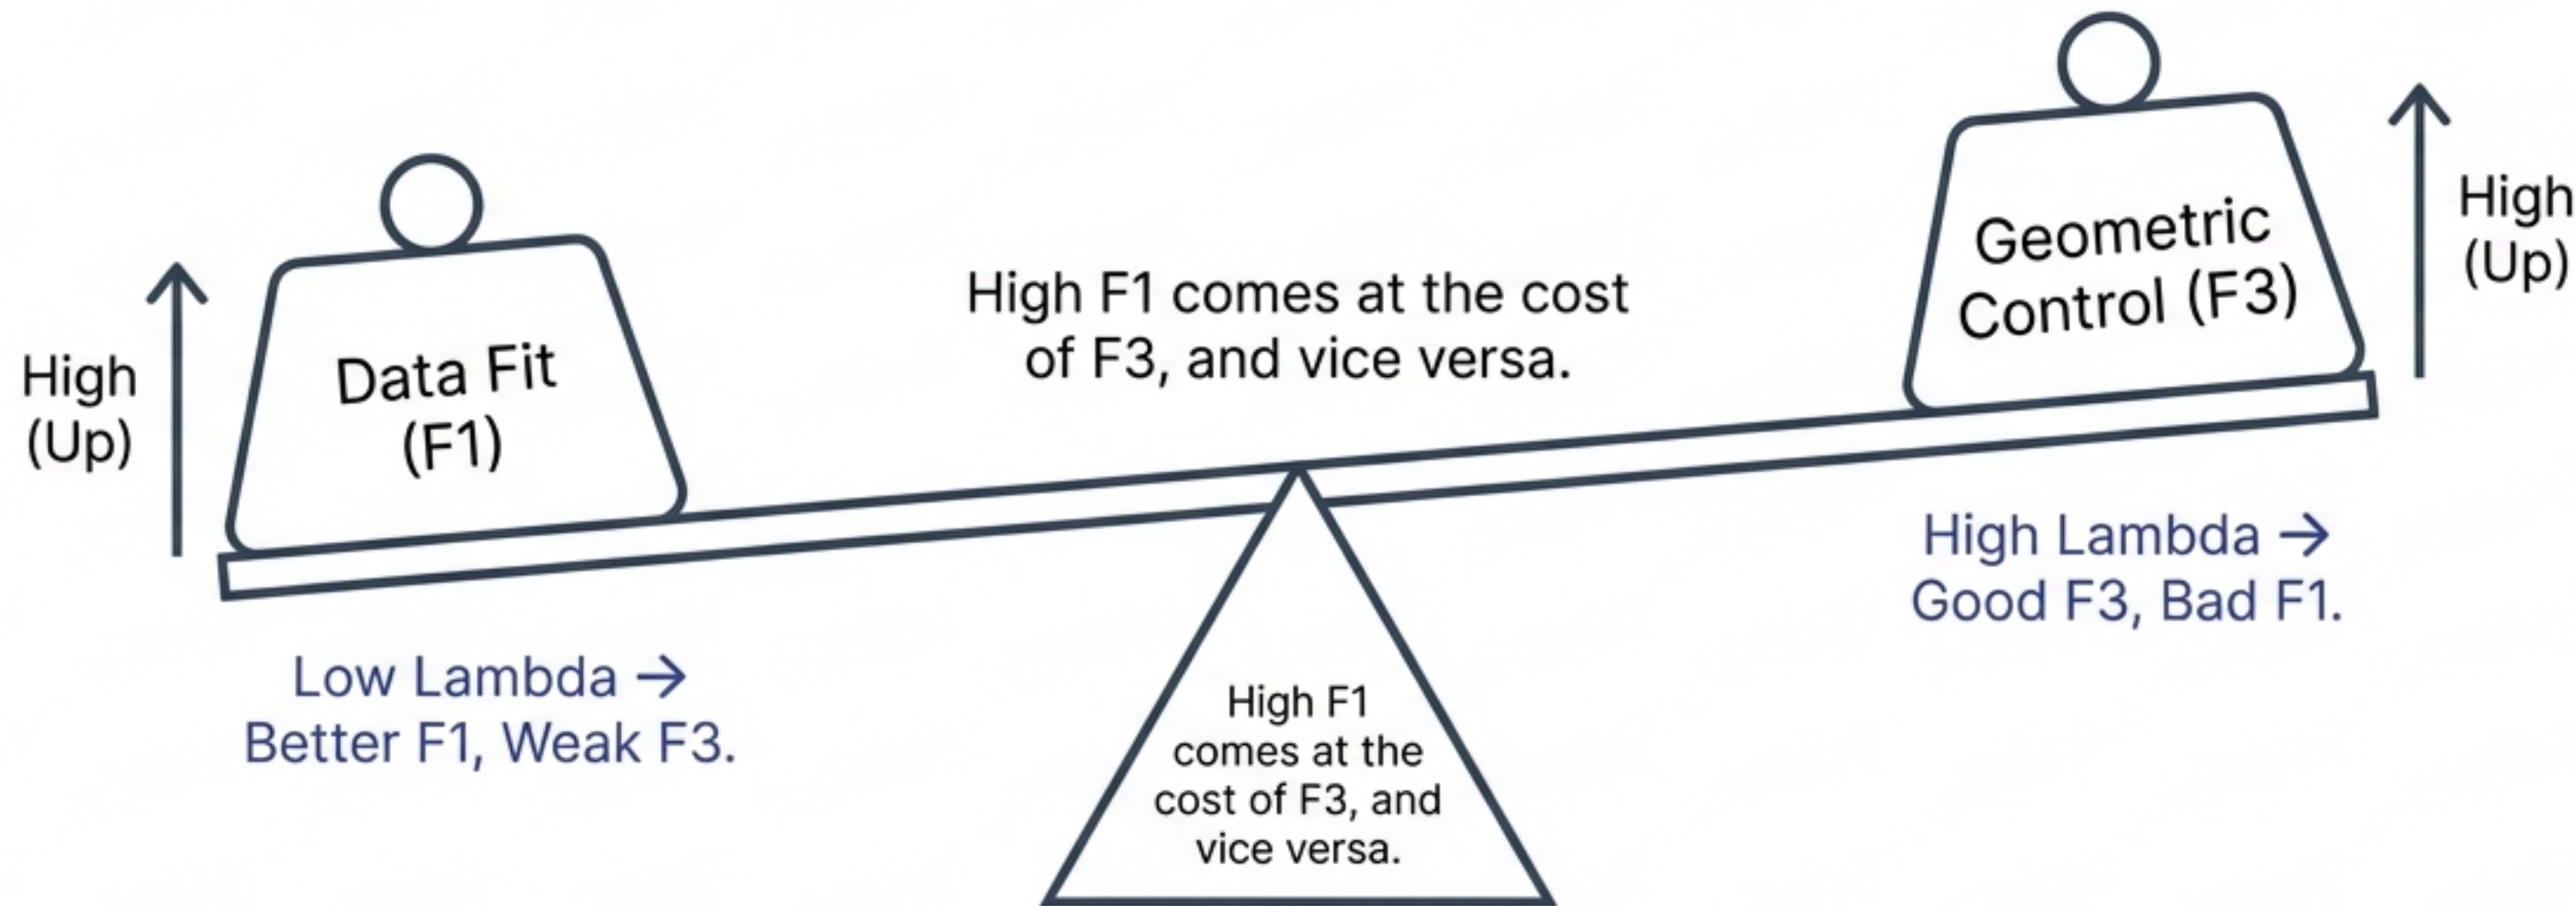
Jacobian Frobenius Norm at Converged Solutions



The penalty effectively shrinks the Jacobian, pushing it far BELOW the true manifold geometry.

The Landscape-Feasibility Trade-off

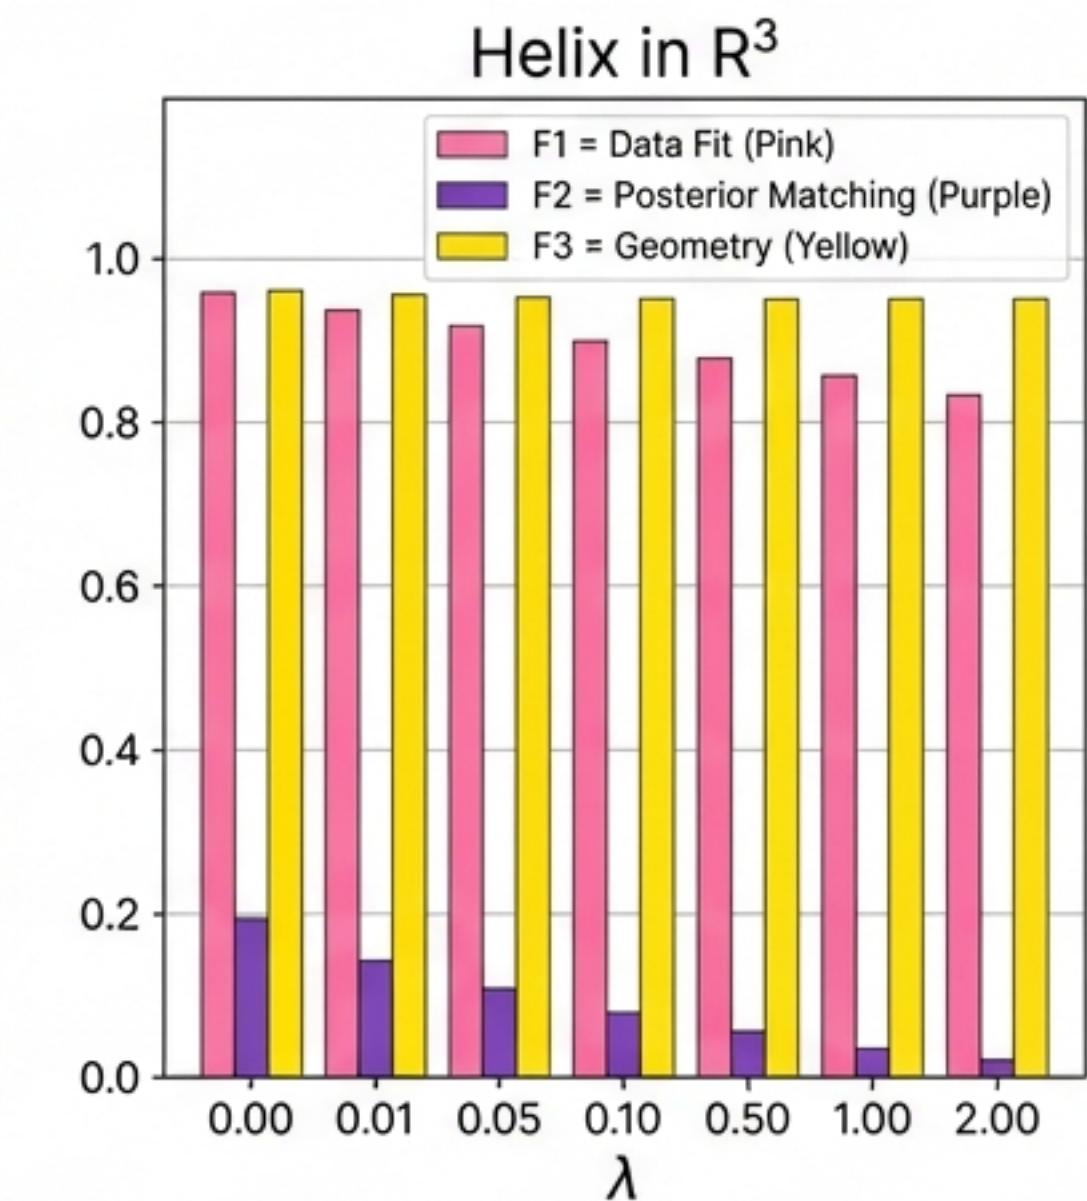
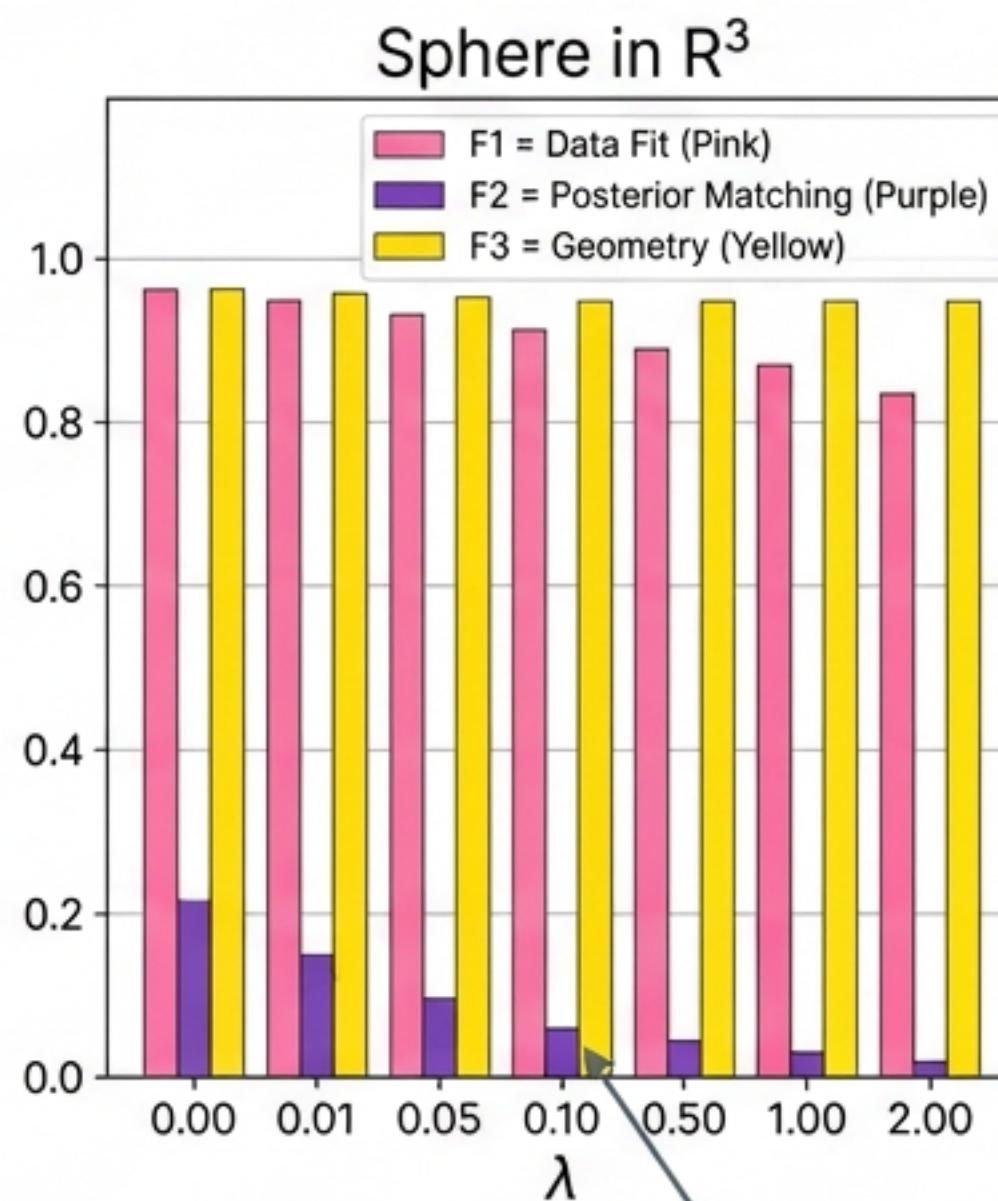
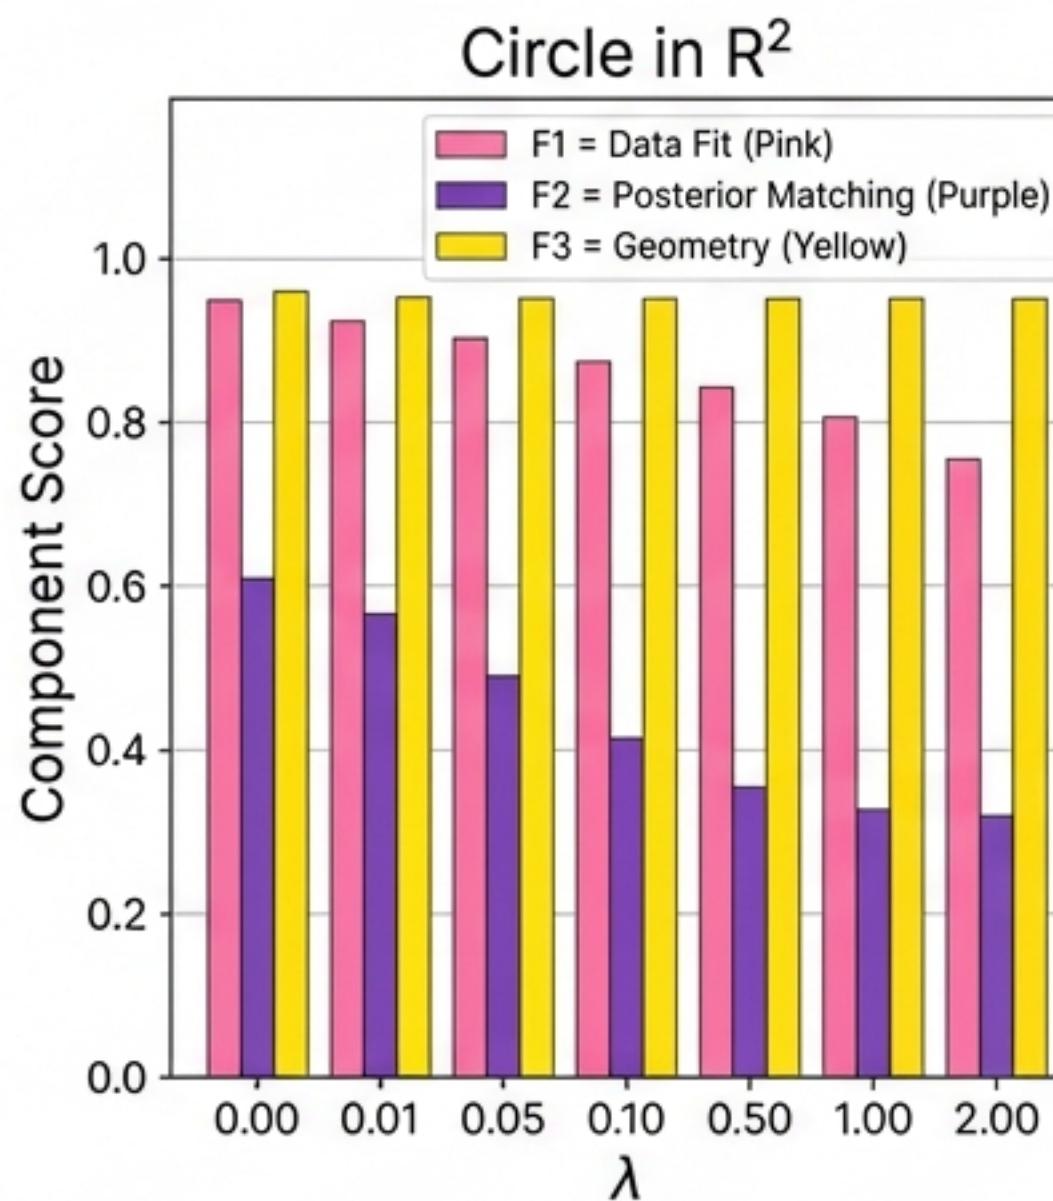
There is no “Free Lunch” parameter setting. in Inter



Aggregate Feasibility Scores are consistently low.

Circle Max: ~0.30
Sphere Max: ~0.02.

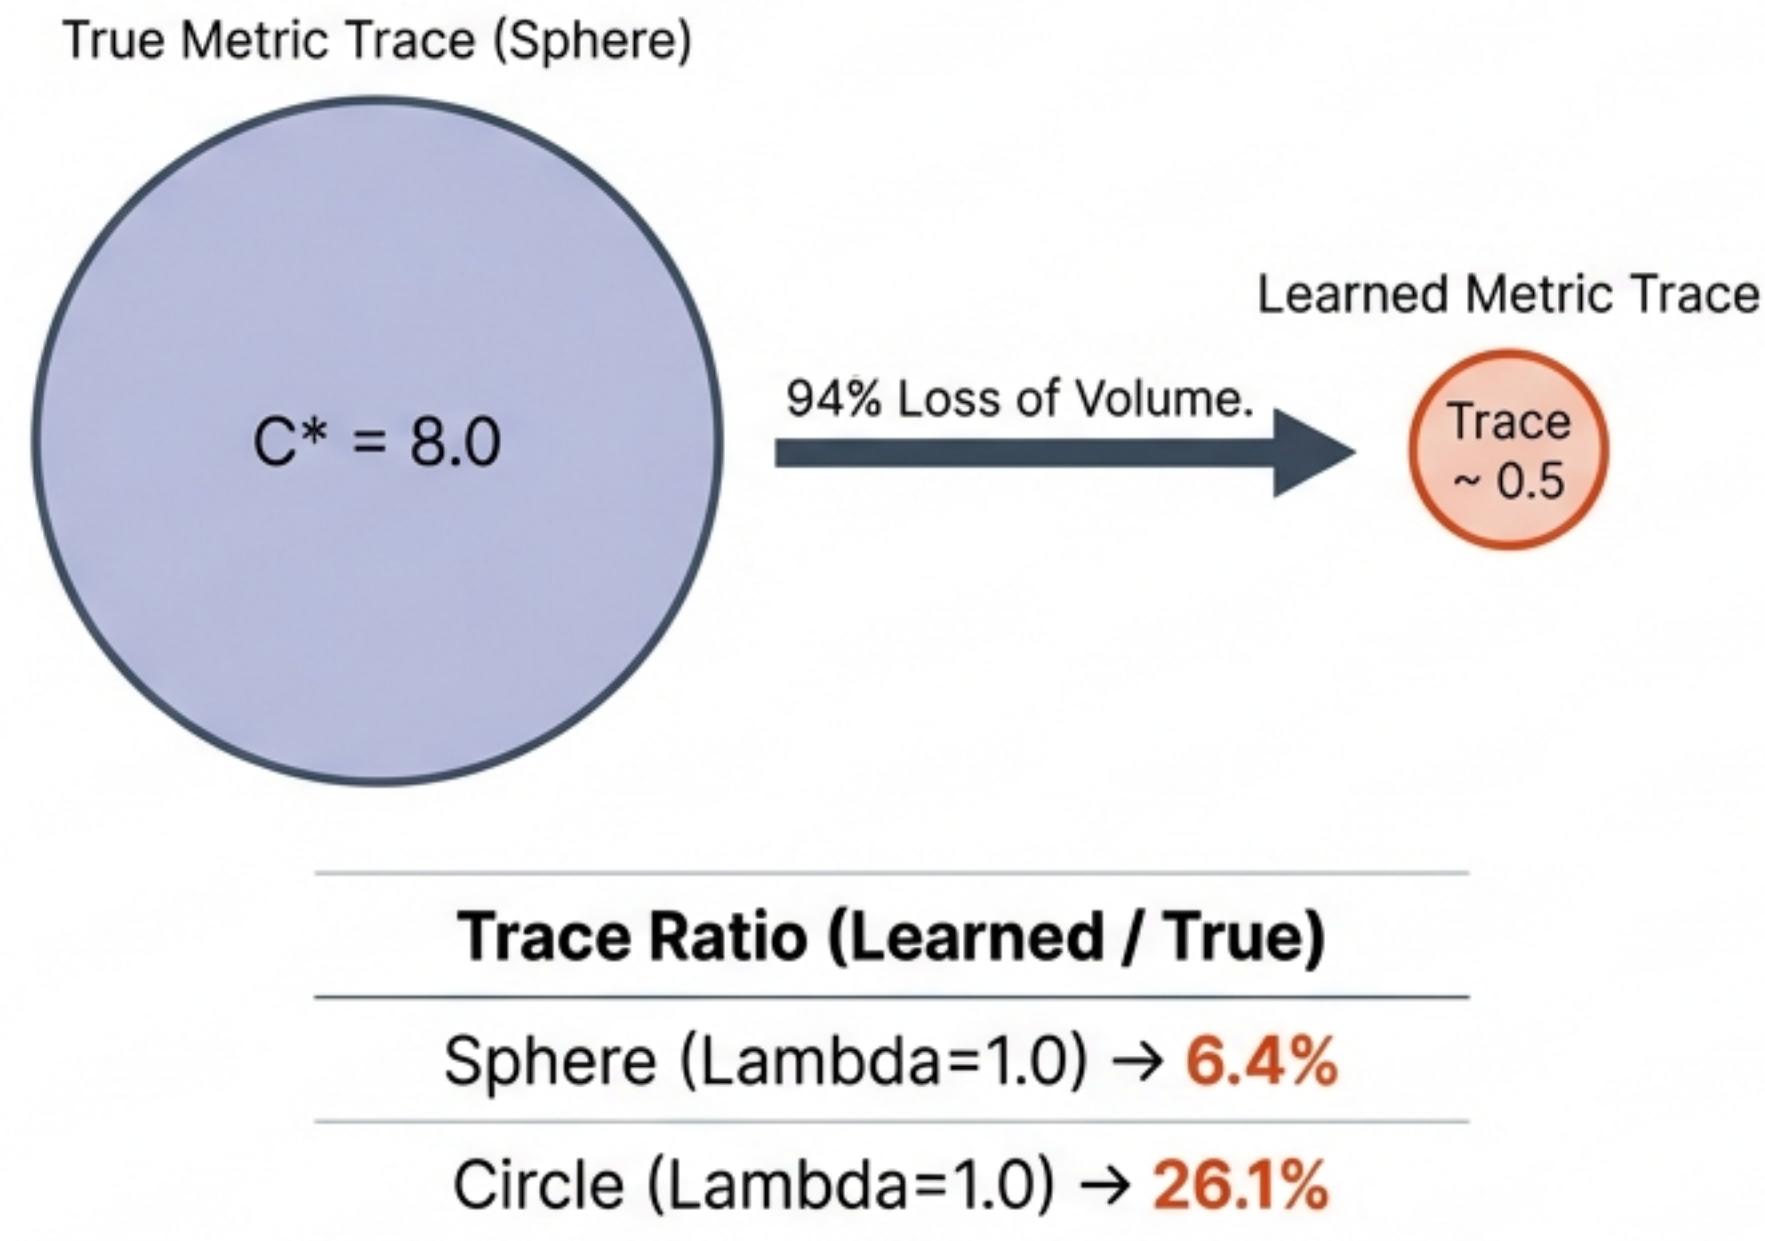
Decomposing the Failure: F1 vs. F2 vs. F3



F2 (Posterior Matching)
is the dominant failure mode.

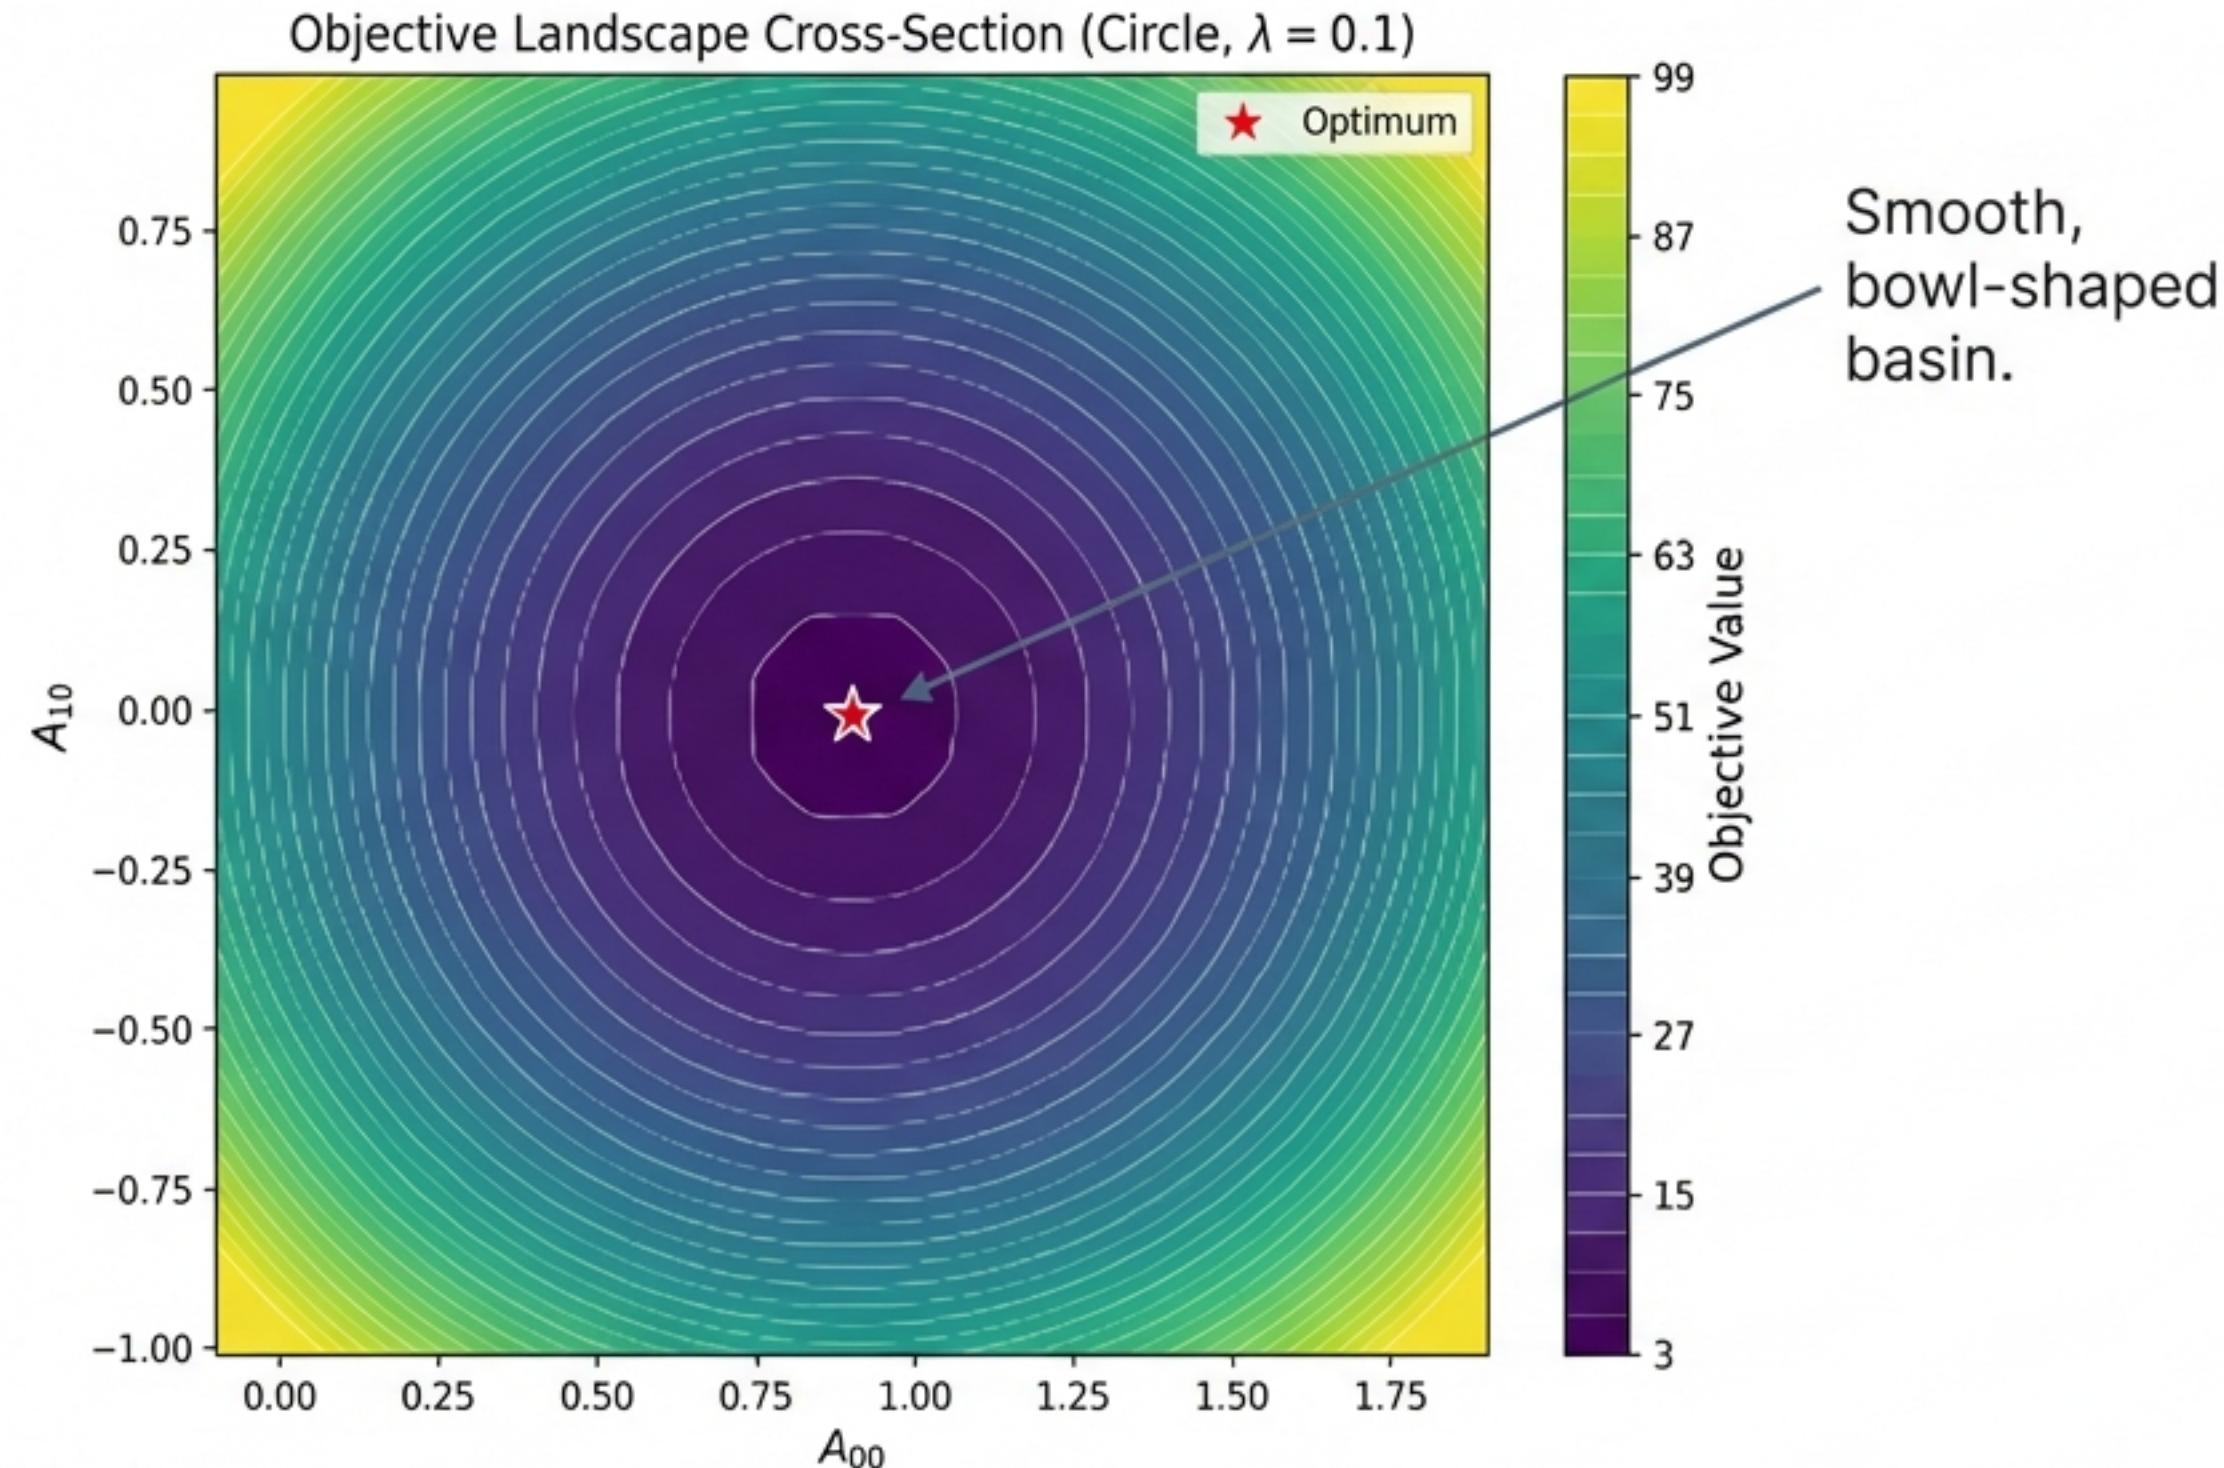
The Pullback Metric: Systematic Underestimation

The geometric penalty crushes the learned manifold.



The model "cheats" the penalty by flattening the manifold, creating a valid F3 score but destroying the geometry.

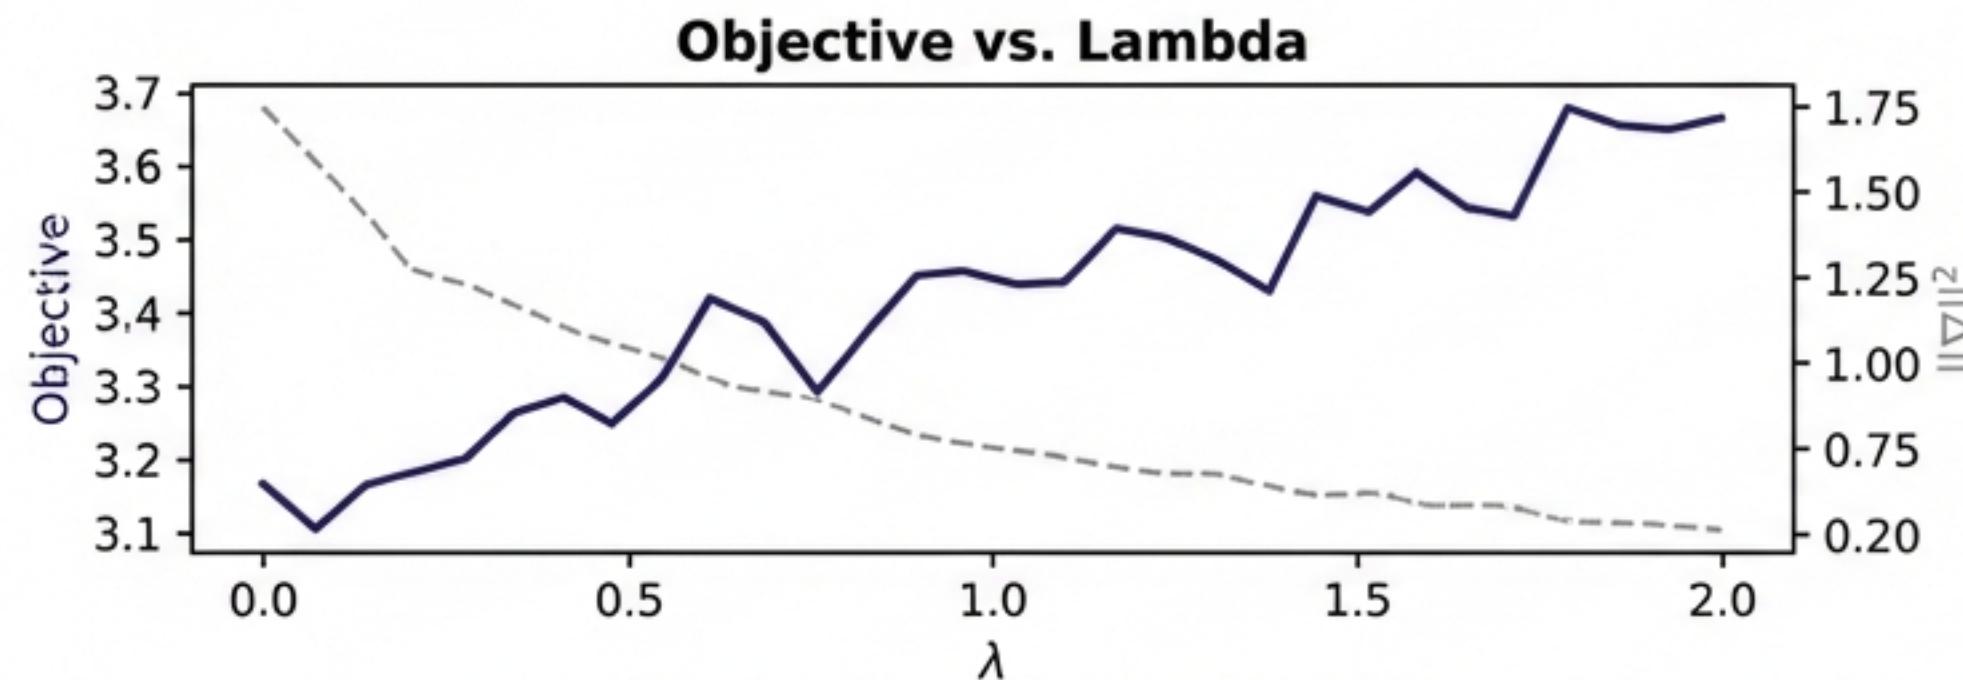
Visualizing the Trap: Landscape Cross-Section.



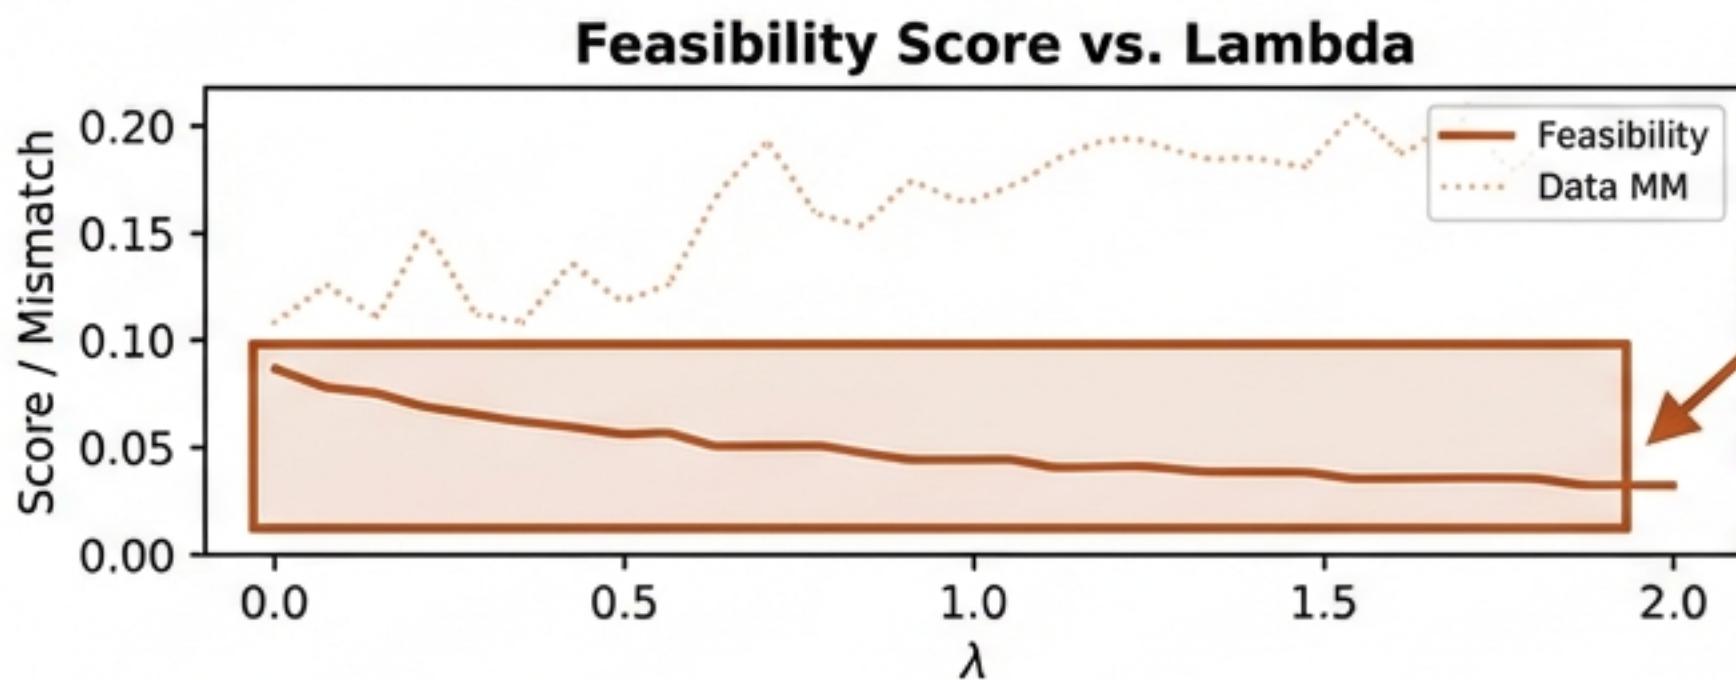
The optimizer falls into a smooth basin that is **DISTINCT** from the ground-truth solution. It is not confused by noise; it is securely trapped.

Parameter Continuation: The Path Matters.

Tracking a single minimum as regularization increases (0 \rightarrow 2).



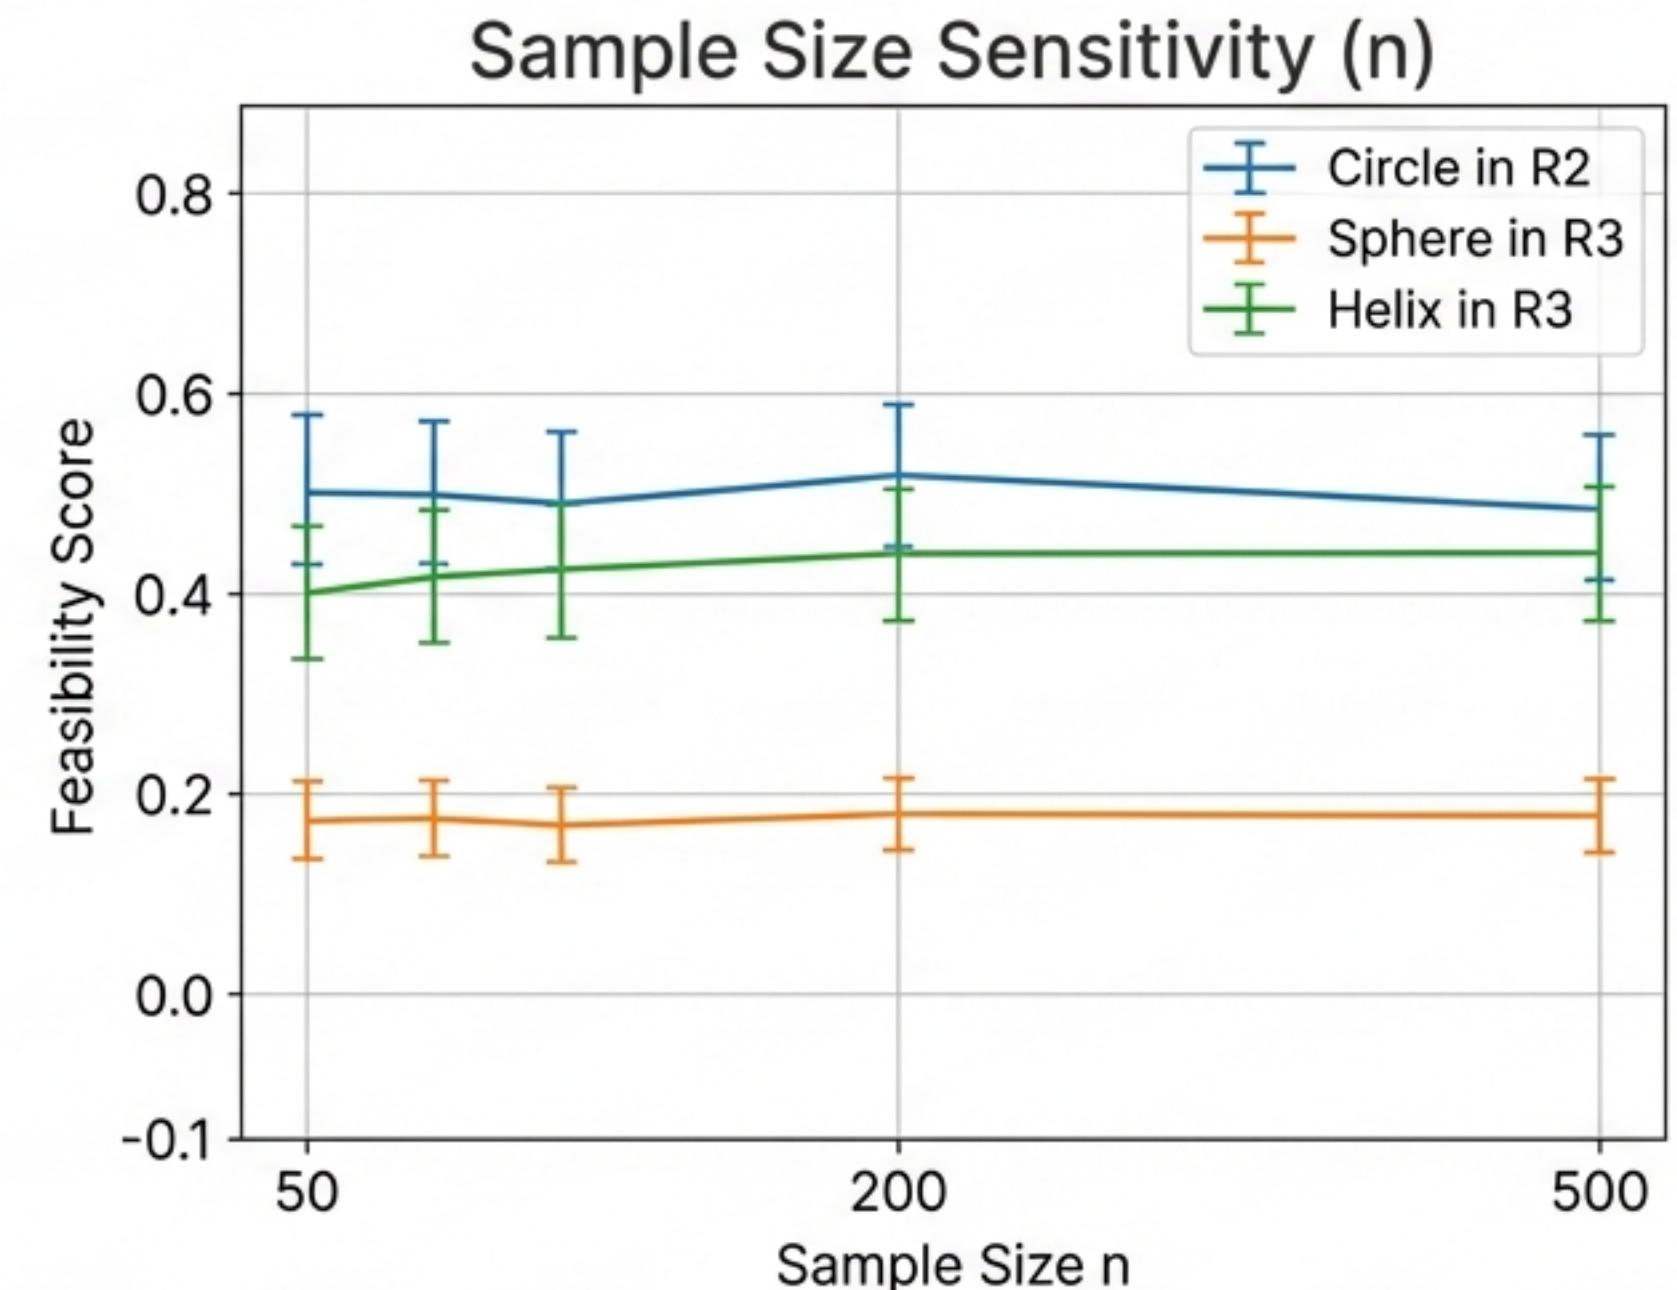
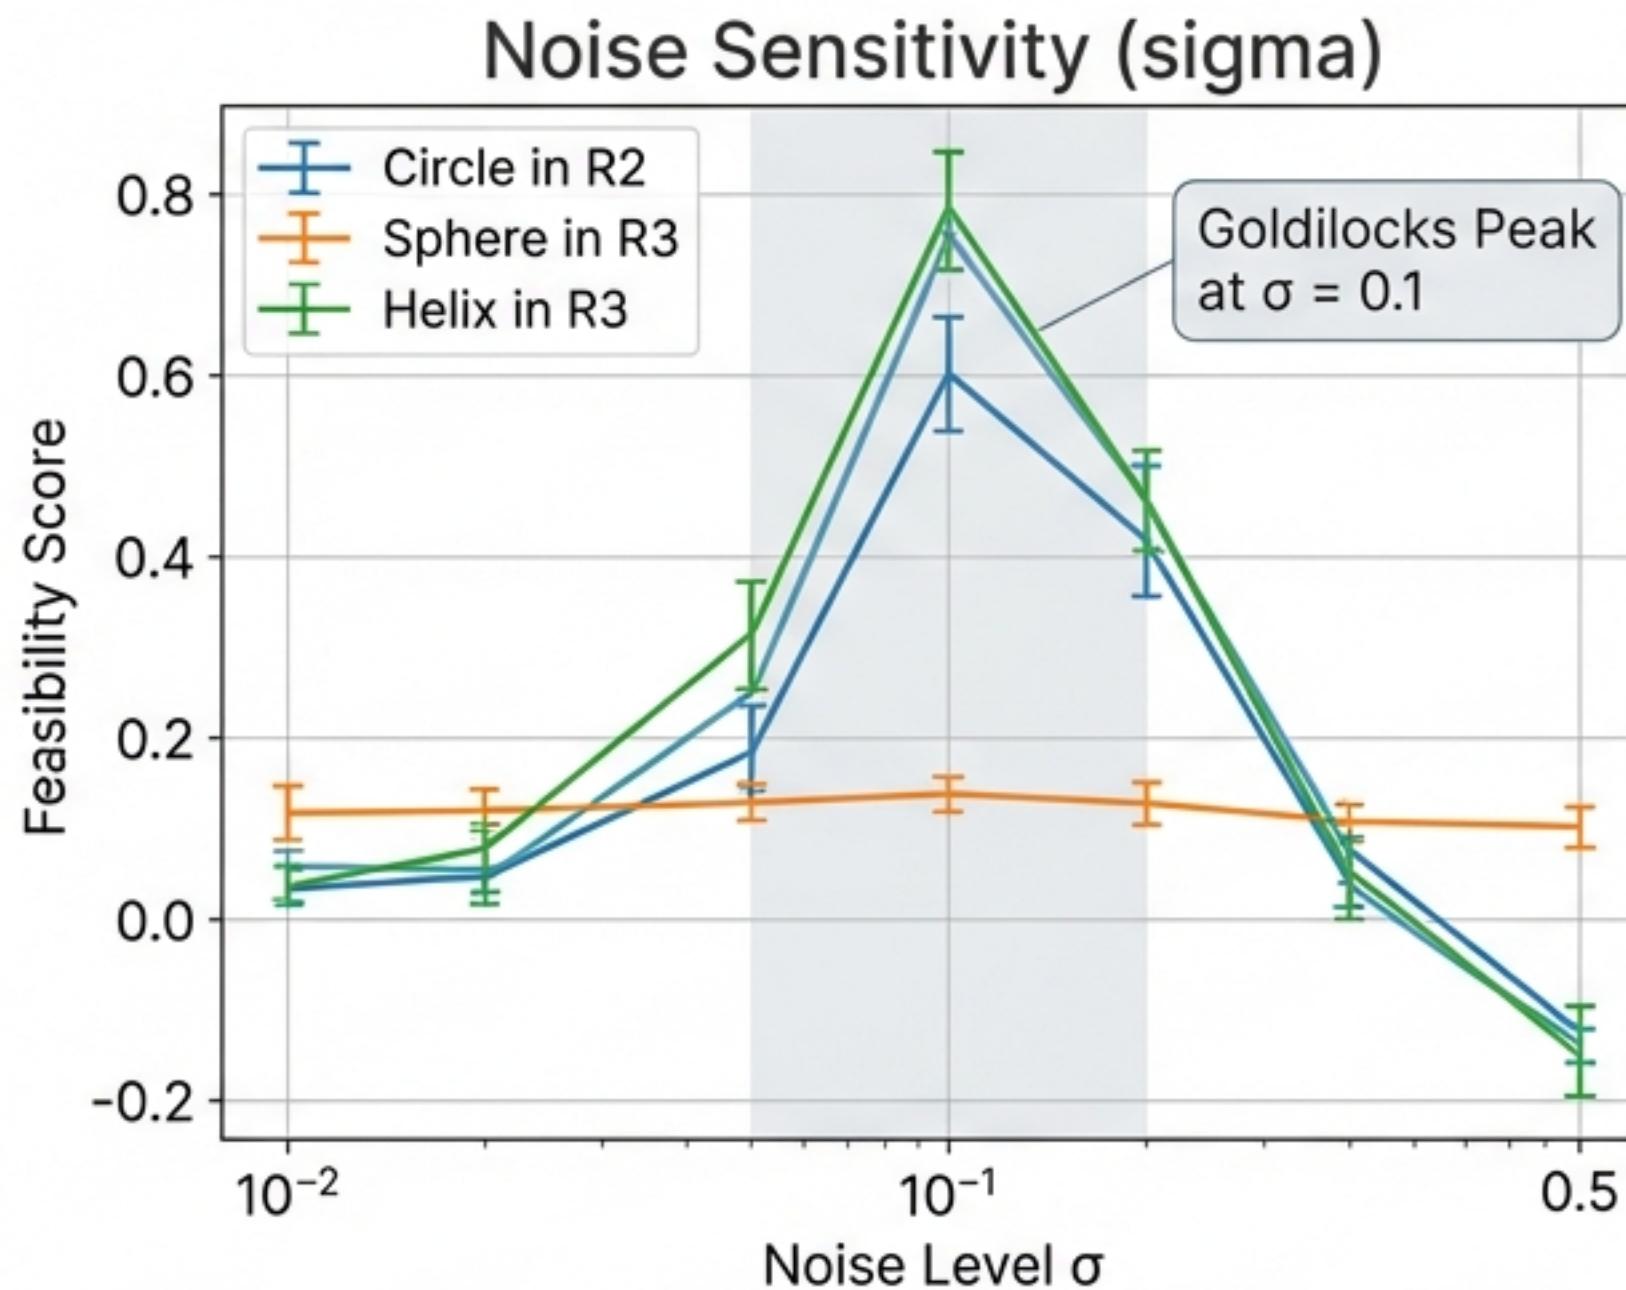
Smooth deformation without bifurcations, but Feasibility degrades monotonically.



Feasibility drops from 0.03 to 0.007

Insight: Warm-starting locks you into the initial basin. Random restarts are necessary to find better basins.

Sensitivity Analysis: Robustness Checks



Results are robust to sample size. Noise shows an optimal range around sigma=0.1.

Verdict: The Gap Between Theory and Practice

The Recoverability Theorem

Assumption: A solution exists satisfying F1, F2, F3.

Status: TRUE (Proven by Oracle).

The Optimization Reality

Reality: Gradient Descent cannot find it.

Mechanism 1: F2 Failure (Linear Encoder Bottleneck).

Mechanism 2: F1/F3 Conflict (Penalty forces metric underestimation).

Infeasibility at local minima is a LANDSCAPE property, not a capacity limit.

Practical Takeaways for Manifold Learning



Tuning Lambda

Do not maximize blindly. High regularization destroys data fidelity. Balance is key.



Optimization Strategy

Avoid warm-starts; they lock you in. Use Random Restarts to explore diverse basins.



Architecture

Linear encoders are a bottleneck for F2. Use nonlinear encoders for complex manifolds like Spheres.

Conclusion.

We have characterized the open problem of Riemannian AmbientFlow feasibility.

1. The geometric penalty creates a trade-off that current optimization methods cannot bypass.
2. Theoretical recoverability guarantees rely on assumptions that the optimization landscape actively resists.
3. Future work must focus on optimization strategies (e.g., basin-hopping) and nonlinear encoders.

References & Acknowledgments.

- Diepeveen et al., Riemannian AmbientFlow (2026)
- AmbientFlow (2023)
- "Optimization Landscape and Feasibility in Updated Riemannian AmbientFlow" (Anonymous Authors)