The Data Manifold under the Microscope

Marios Koulakis[1] Constantin Seibold[2]

[1]Independent Researcher, Germany [2]University Clinic Heidelberg, Germany

UNIVERSITÄTS KLINIKUM HEIDELBERG

A controlled setup for representation geometry

We build image-like manifolds where curvature, reach, volume, and layer-wise geometry are **measurable rather than assumed**. This turns hidden representation geometry into something we can probe directly.

Controlled image manifolds

Real datasets hide geometric quantities used in manifold-based theory and interpretability: dimension, curvature, reach, volume, and sampling density.

Idea: Expand image datasets along controlled transformation axes, giving explicit charts $u : \Theta \subset \mathbb{R}^d \to \mathbb{R}^{64\times 64}$.

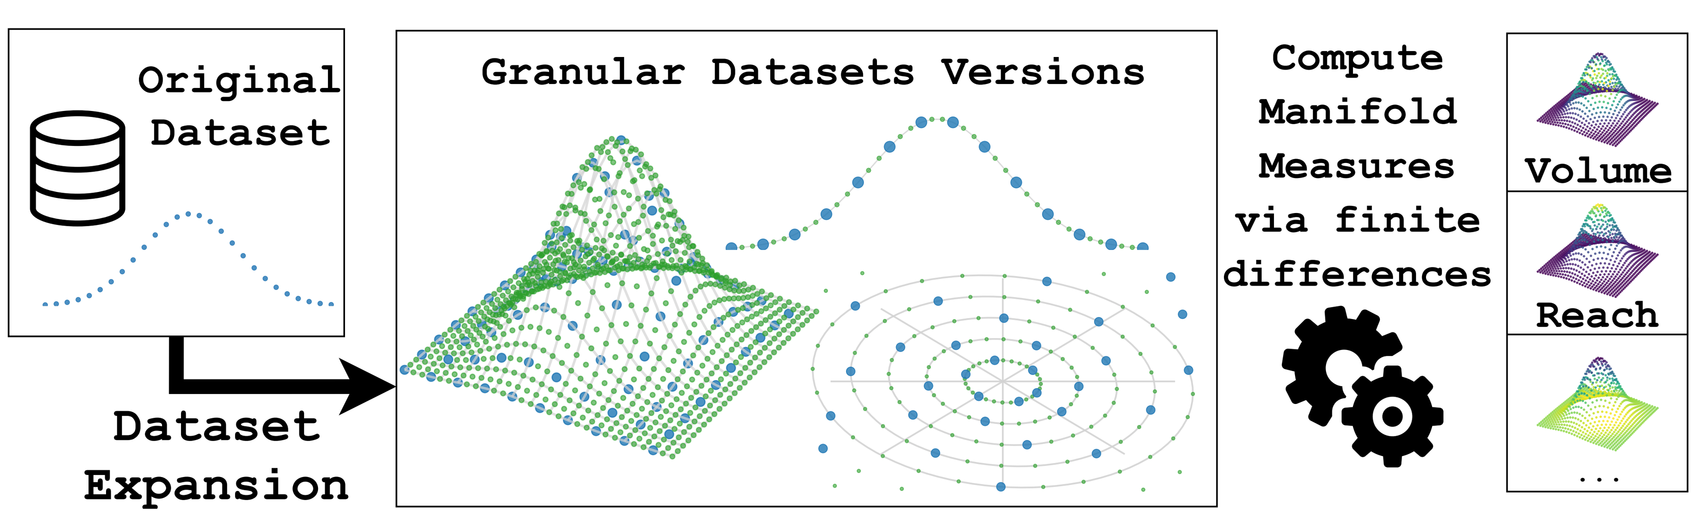

- **dSprites:** $d = 4$, scale, rotation, x, y.
- **COIL-20:** $d = 3$, view angle, rotation, scale.
- Dense grids provide train/test splits and finite-difference stencils.

Geometry from the grid

Tangent vectors and the induced metric come from finite differences:

$$u_{,i}(x) \approx \frac{u(x + he_i) - u(x - he_i)}{2h}, \qquad g_{ij} = \langle u_{,i}, u_{,j} \rangle.$$

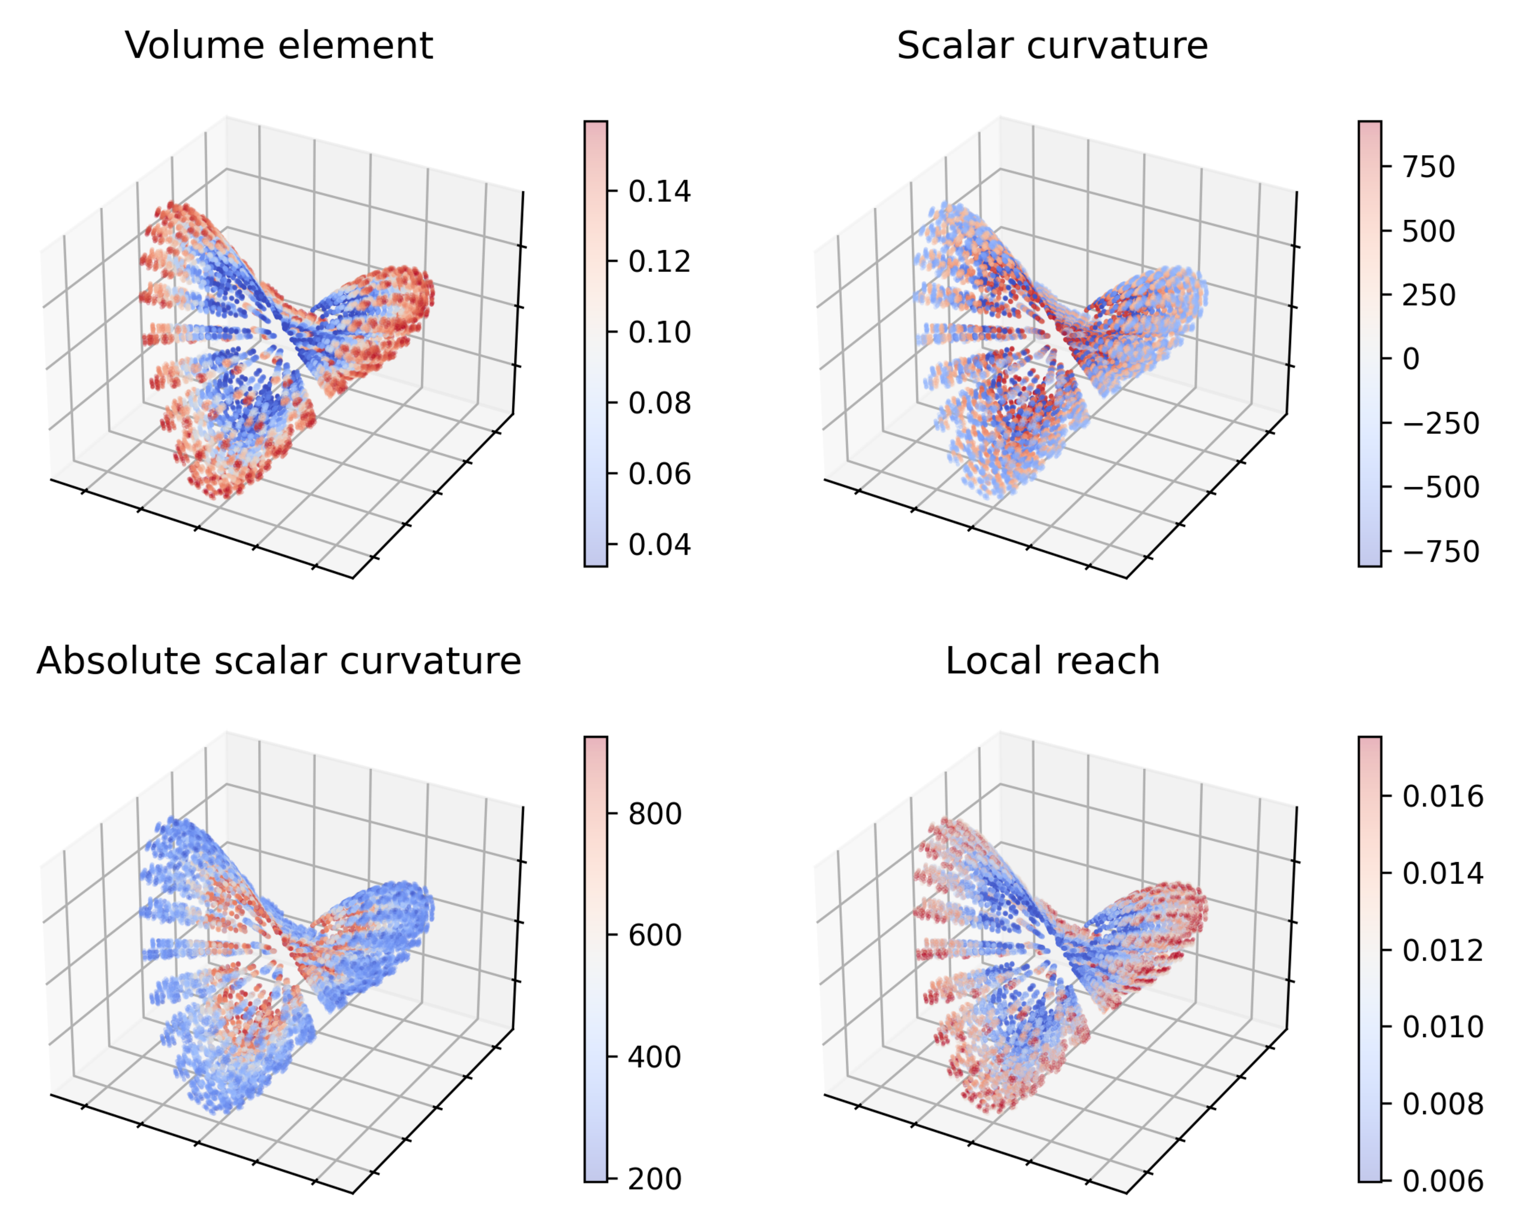

Geometric fields on the 4D dSprites square manifold.

Use case: What happens inside a β-VAE?

The same transformation grid lets us measure how hidden representations reshape the data manifold.

dSprites

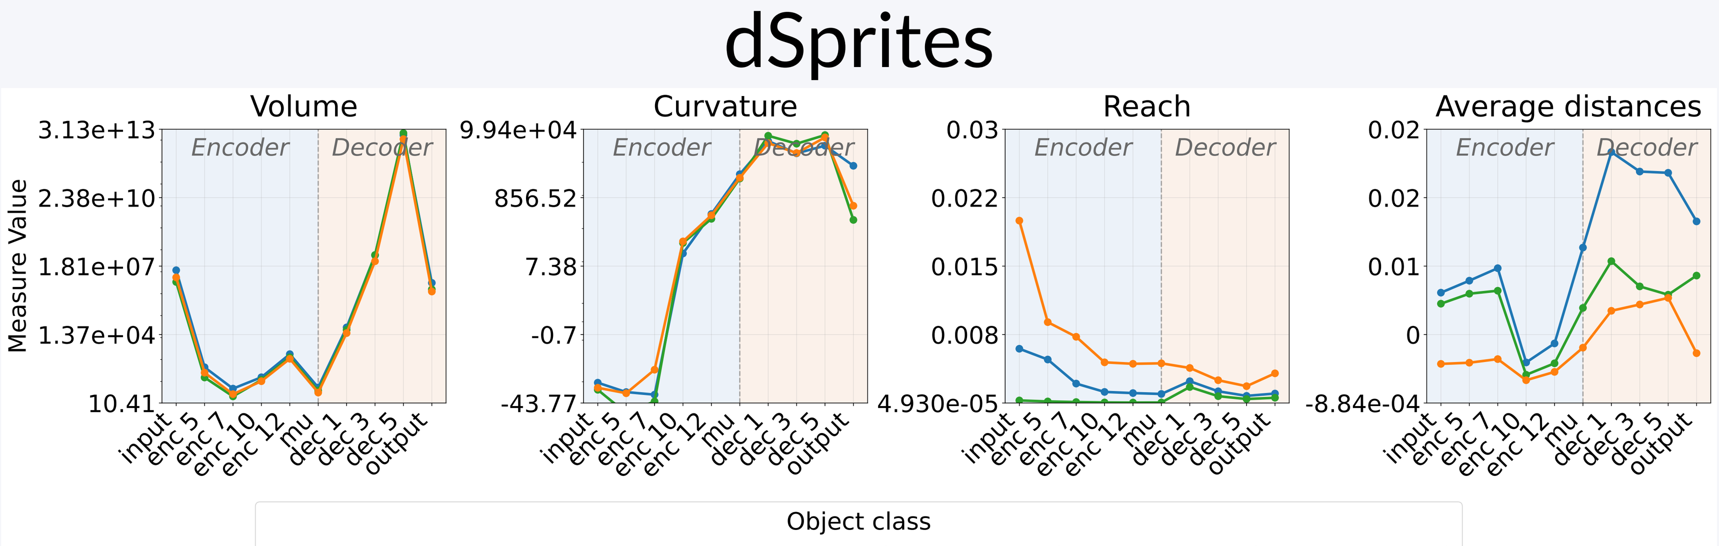

Layer-wise geometry, dSprites.

COIL-20

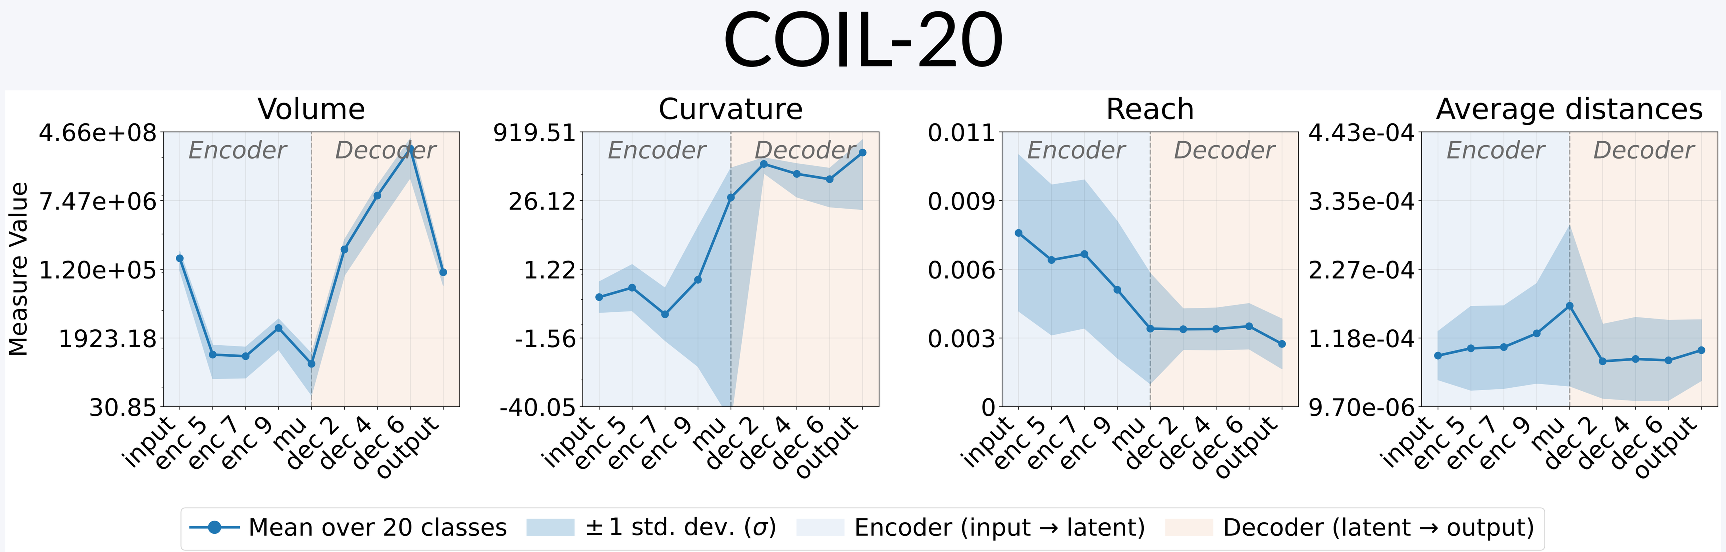

Layer-wise geometry, COIL-20.

Curvature increases, reach decreases, and class distances peak near the latent representation.

Use case: Minimax rates

Scaling of manifold fitting: $\widehat{R}_n^{\text{fit}} = \widehat{C} n^{-\widehat{A}}$.

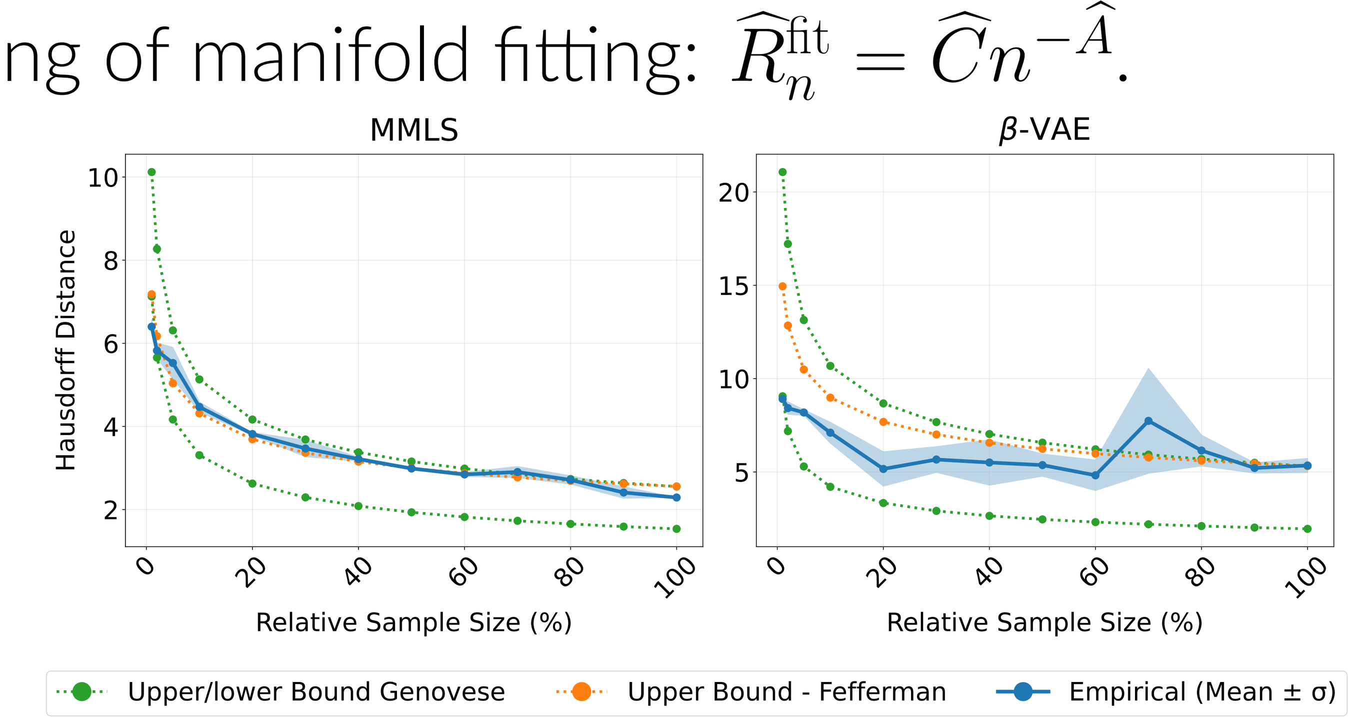

Takeaways and next steps

- Controlled image manifolds make geometric assumptions testable.
- Explicit grids provide finite-difference geometry.
- Layer-wise measurements probe representation geometry.
- **Next:** Richer transforms, occlusions, modalities, rasterization.