On the Impact of Topological Regularization on Geometrical and Topological Alignment in Autoencoders: An Empirical Study

Samuel Graepler[1] Diaaeldin Taha[2] Nico Scherf[3] Anna Wienhard[2]

[1] Faculty of Mathematics and Computer Science, Leipzig University, Leipzig, Germany
[2] Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany
[3] Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

UNIVERSITÄT LEIPZIG

MAX PLANCK INSTITUTE FOR MATHEMATICS IN THE SCIENCES

MAX PLANCK INSTITUTE FOR HUMAN COGNITIVE AND BRAIN SCIENCES

Motivation

Manifold learning aims to construct meaningful low-dimensional representations of underlying data manifolds. Ideally, a representation should preserve both topological structure (components, loops, voids) and geometric properties (distances, curvature). However, common manifold-learning methods such as t-SNE and UMAP often distort global structure, while standard autoencoders provide no guarantees on topology or geometry. Topologically regularized autoencoders aim to preserve the data topology, but what happens to the latent geometry in these models?

Datasets & Models

- **Six synthetic datasets:** circles, spheres, and tori, deformed and embedded in $\mathbb{R}^{10}$.
- Each with a low-deformation version confined to a small subspace and a high-deformation version that bends across more ambient dimensions.
- Known ground-truth topology and curvature.

- **Two models:** AEs and Gaussian VAEs with euclidean latent spaces and topological regularization.
- Latent dimension $d+1$, for d the intrinsic dimension of the data manifold.

Intervention and Training

We use the persistent-homology-based topological regularization loss $\mathcal{L}_t$ from [1]. It aims to match the lengths of topology-relevant edges, i.e. edges from the persistence pairings that give birth or death to topological features on a mini-batch scale.

$$\mathcal{L}_{\text{AE}} = \alpha \mathcal{L}_{\text{recon}} + \gamma \mathcal{L}_t,$$

$$\mathcal{L}_{\text{VAE}} = \alpha \mathcal{L}_{\text{recon}} + \beta \mathcal{L}_{\text{KL}} + \gamma \mathcal{L}_t$$

We explore $\alpha \in \{0, 1\}$, $\beta \in \{0, 0.08, 1\}$, $\gamma \in \{0, 1, 100\}$. The parameter $\dim_t$ defines the highest feature dimension considered (0: components, 1: loops, 2: voids).

Curvature Estimation

We estimate mean extrinsic curvature by fitting a local quadric to each point's k-nearest neighbors. The curvature is extracted from the quadric's Hessian. We compare input vs. latent curvature using MSE and SMAPE.

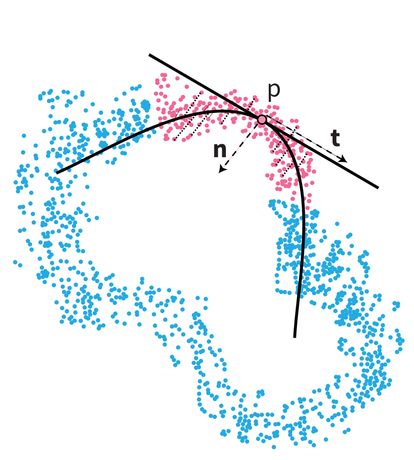

Objective

In an empirical study, we investigate the impact of topological regularization in autoencoders (AEs) and Variational Autoencoders (VAEs) on topological and geometrical alignment between data and latent representation, focusing on extrinsic curvature.

Take-Away Message

Key findings in AEs. The use of the persistent homology-based topological regularization can lead to improved geometrical alignment of data and latent representation in autoencoders.

Key findings in VAEs. Topological regularization can only partially mitigate the disruption of the latent geometry caused by the KL term.

Results: Autoencoders

Without topological regularization, AEs roughly preserve topology but distort geometry in the latent space. Adding topological regularization substantially improves curvature alignment between data and latent representations across datasets. Geometry can improve even when training uses only the topological loss, with no reconstruction term. The effect depends strongly on the chosen topological feature dimension $\dim_t$ in a non-monotonic way.

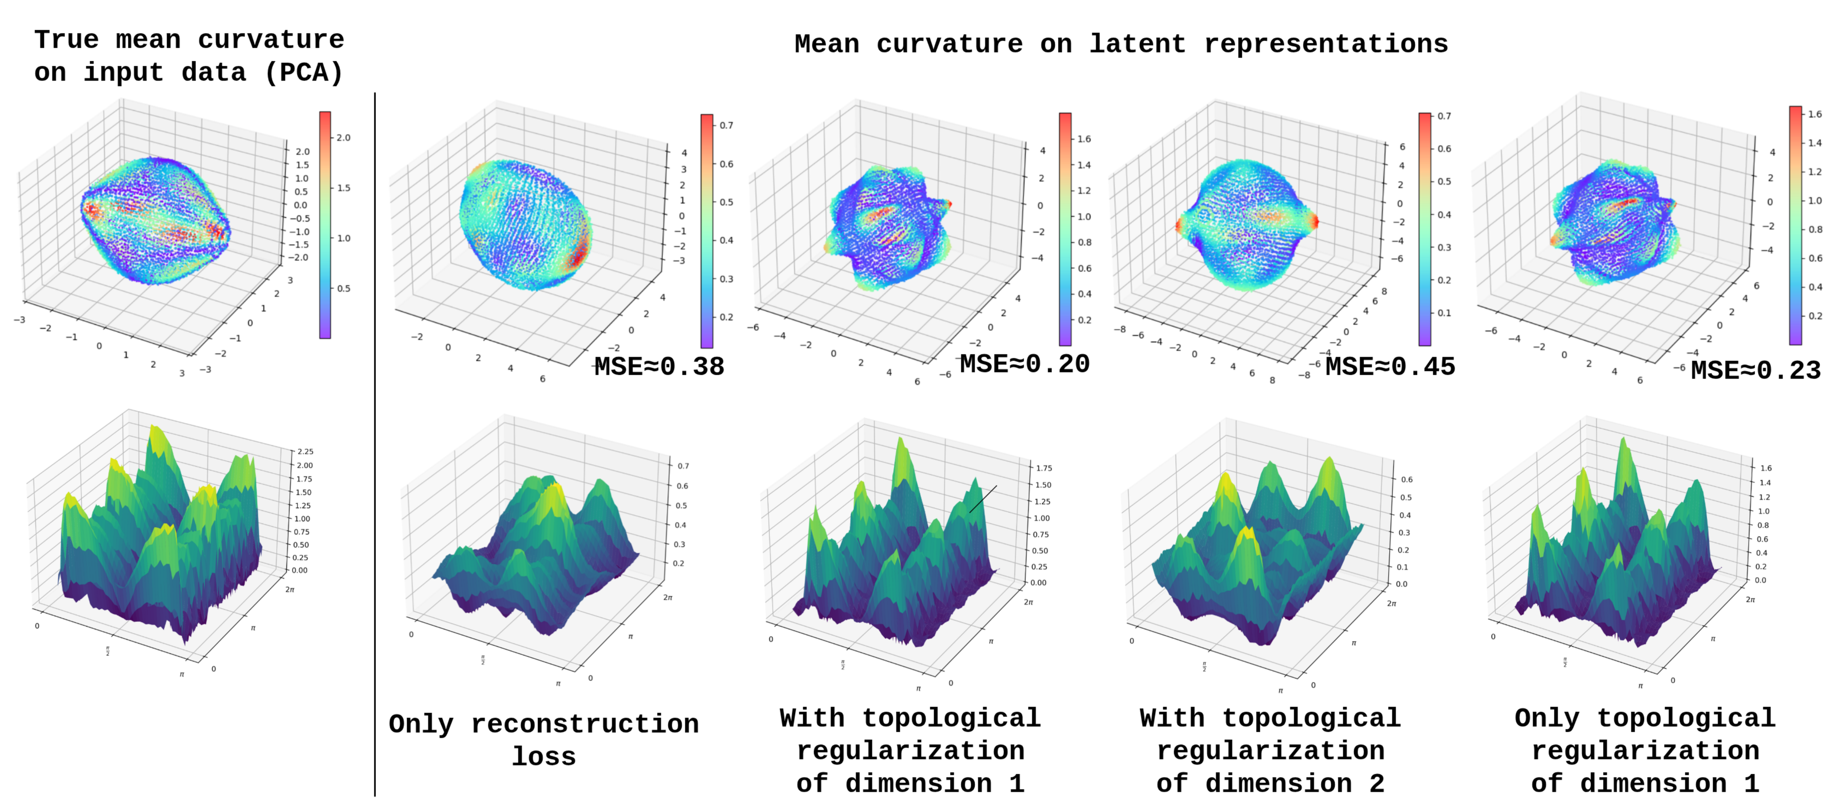

Figure: Impact of topological regularization on the latent geometry of an AE trained on the $\mathrm{Sphere}_{\mathrm{high}}$ dataset. The upper row shows the curvature of the input data and the learned representation as heatmap. The lower row shows the latent curvature plotted over the ground truth angles. Color indicates curvature. Scales vary across plots.

Results: Variational Autoencoders

The KL term pushes the latent space toward a Gaussian blob, disrupting topology and geometry. Topological regularization mitigates this effect for small KL weights β. Removing the KL term makes VAEs behave like AEs.

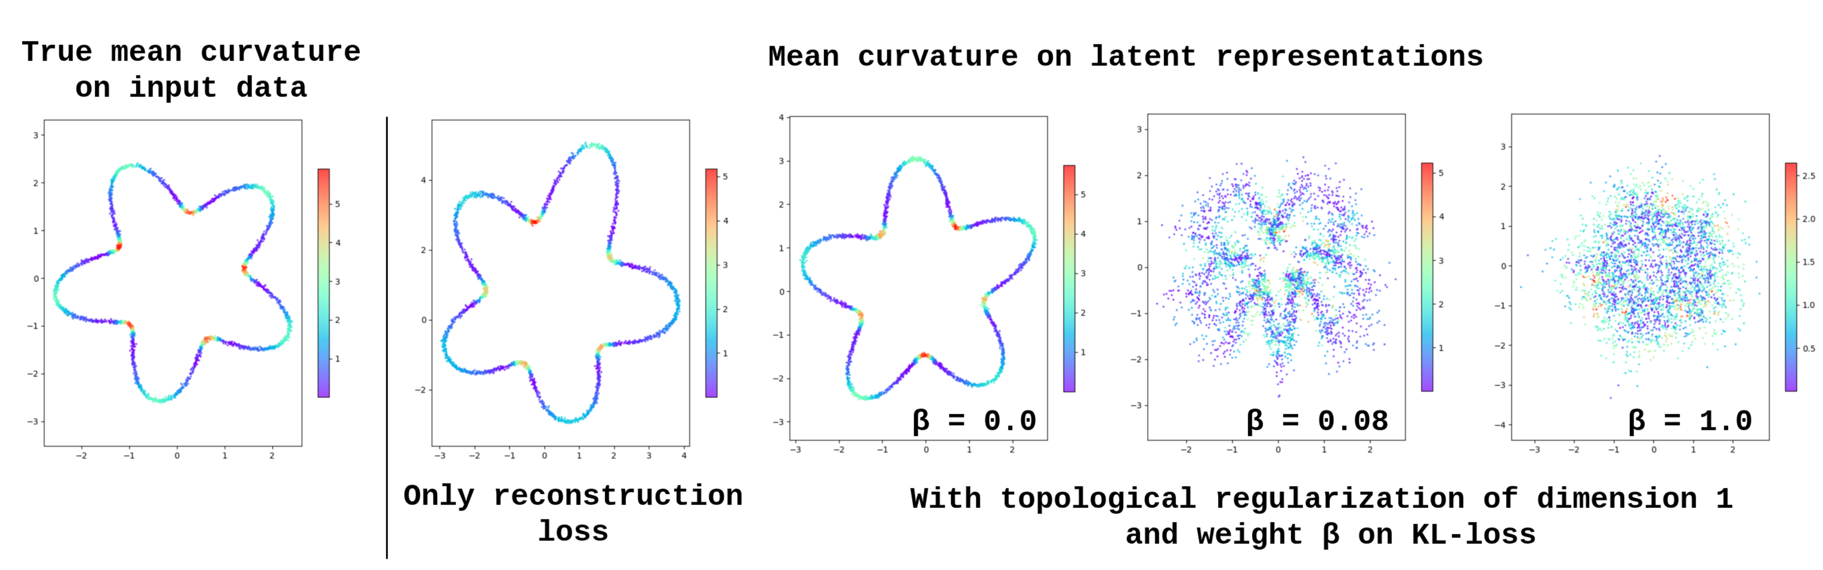

Figure: Impact of the KL term and topological regularization on the latent geometry of a VAE trained on the $\mathrm{Circle}_{\mathrm{low}}$ dataset. Color indicates curvature.

Discussion

- Topological regularization improves geometric alignment in AEs, even without reconstruction.
- Controlling topology-relevant distances already induces useful geometric structure.
- This suggests potential for models without decoder.
- The KL term makes VAEs unsuitable for this approach.

Future Work

- Quantify topological alignment rather than relying solely on visual inspection.
- Benchmark against other manifold-learning methods.

References

[1] Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt.
Topological autoencoders.
In *International conference on machine learning*, pages 7045--7054. PMLR, 2020.