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

Samuel Graepler SAMUEL.GRAEPLER@OUTLOOK.DE
Faculty of Mathematics and Computer Science, Leipzig University, Leipzig, Germany

Diaaeldin Taha DIAAELDIN.TAHA@MIS.MPG.DE
Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

Nico Scherf
Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

Anna Wienhard
Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

Editors: List of editors' names

Abstract

We present a comparative empirical study on the impact of topological regularization on autoencoders (AEs) and variational autoencoders (VAEs) across six synthetic datasets with known topology and curvature. Particularly, we probe the alignment of the topology and geometry of the dimensionality-reduced latent representation with that of the data. To quantify geometrical alignment, we estimate the mean extrinsic curvature of the latent embedding by fitting local quadrics. We find that topological regularization can significantly improve the geometrical alignment of latent and data, even when the training objective emphasizes topological alignment alone, without regard for reconstruction quality.

Keywords: Representation Learning, Manifold Learning, Geometric Deep Learning, Autoencoder

1. Introduction

Learning meaningful representations is central to understanding and processing high - dimensional data. The *manifold hypothesis* (Bengio et al., 2013) states that real-world data often lies near a low-dimensional manifold, a view supported in computer vision (Carlsson et al., 2008), neuroscience (Chaudhuri et al., 2019), and machine learning (Naitzat et al., 2020). From a *manifold learning* perspective, an ideal representation respects both topology and geometry: local neighborhoods, distances, and angles are preserved while enabling operations like interpolation and distance computations (Lee et al., 2022; Hauberg, 2019).

However, in practice, most methods fail to satisfy these properties reliably. State-of-the-art manifold learning methods like t-SNE and UMAP struggle to represent global structure faithfully and frequently do not preserve topology (Moor et al., 2020; Nazari et al., 2023). In geometric deep learning, autoencoders provide a popular tool to face this task. A standard autoencoder (AE) can be considered as an embedding method that learns a dimensionality-reduced latent representation of the data in the latent space $\mathcal{Z}$. A probabilistic variant is the variational autoencoder (VAE) (Kingma and Welling, 2014), which instead learns a joint probability distribution over the input and latent space. For each input x, the encoder

produces a distribution $\text{enc}(x) = q(z|x)$ on $\mathcal{Z}$, such that a sample $z^* \sim q(z \mid x)$ can be decoded with a high probability to $\text{dec}(z^*) \approx x$. Thus, every forward pass of the data yields a probabilistic embedding in the latent space. Although various extensions alter the latent geometry, the standard choice is $\mathcal{Z} = \mathbb{R}^d$, a flat Euclidean space.

In this work We present preliminary results from an empirical comparison of deterministic AEs and Gaussian VAEs with Euclidean latent spaces. We investigate their ability to learn latent representations that align topologically and geometrically with the data, using six synthetic datasets lying near one and two-dimensional embedded manifolds. Topological alignment is assessed visually, while geometry preservation is assessed through estimates of the *mean extrinsic curvature*, a measure of how the manifold bends within the ambient space. We estimate curvature directly from the latent embedding via a local quadric fit, similar to Gilbert and O'Neill (2025); Yang and Lee (1999). In particular, we investigate the impact of *topological regularization* (Moor et al., 2020) on topological and geometric representational alignment.

2. Models and training

2.1. Datasets and architectures

To obtain controlled ground-truth data, we construct six synthetic datasets by embedding smoothly deformed one- and two-dimensional manifolds into $\mathcal{X} = \mathbb{R}^{10}$. Each manifold is homeomorphic to either a circle $\mathcal{S}^1$, a sphere $\mathcal{S}^2$, or a torus $\mathcal{T}^2$. For every topology, we generate two variants. In the *low*-deformation datasets ($\text{Circle}_{\text{low}}, \text{Sphere}_{\text{low}}, \text{Torus}_{\text{low}}$, adapted from Acosta et al. (2023)), the d-dimensional manifold remains in a $(d+1)$-dimensional subspace of $\mathbb{R}^{10}$. In the *high*-deformation datasets, the manifold is bent across more ambient dimensions (9 for the $\text{Circle}_{\text{high}}$, 5 for $\text{Sphere}_{\text{high}}$ and $\text{Torus}_{\text{high}}$). To all datasets we finally apply random rotations and small Gaussian noise.

We examine standard deterministic AEs and Gaussian VAEs with Euclidean latent space. Depending on the examined dataset, we set $\mathcal{Z} = \mathbb{R}^{d+1}$, where d is its intrinsic manifold dimension (e.g. $d = 1$ for $\text{Circle}_{\text{low}}$).

2.2. Topological and geometrical evaluation

We judge over topological alignment of data and latent embedding by visually inspecting the latent space focusing on the topological signature, i.e. the number of connected components, holes and voids, compared to the ground truth of the respecting dataset. This is possible since all datasets lie close to one- and two-dimensional manifolds and the latent spaces are constrained to be $\mathbb{R}^2$ or $\mathbb{R}^3$.

To evaluate the geometric alignment between the latent embedding $\{z_i\}_{i=1}^N \subset \mathcal{Z}$ and the original data $\{x_i\}_{i=1}^N \subset \mathcal{X}$, we estimate the local curvature of both manifolds. For each point in the latent space, we identify its k nearest neighbors and apply PCA to fit a quadric that locally approximates the underlying manifold structure. The mean extrinsic curvature is then computed from the Hessian of the fitted quadric, following the methods of Gilbert and O'Neill (2025); Yang and Lee (1999). See Appendix D for the algorithm. Geometric similarity is measured by MSE and SMAPE between data and latent curvature.

2.3. Intervention and Training

To encourage topological alignment, we add the persistent homology-based loss term $\mathcal{L}_t$ (see Appendix A) to the training objective, aligning latent and input topology by matching lengths of topologically relevant edges, which give birth and death to topological features. The fixed hyperparameter $\dim_t$ specifies the highest *topological feature dimension* considered (e.g., $\dim_t = 2$ captures connected components, loops and voids). The total loss is

$$\mathcal{L} = \alpha\mathcal{L}_{\text{recon}} + \gamma\mathcal{L}_{\text{t}} \quad \text{(for AEs)}, \qquad \mathcal{L} = \alpha\mathcal{L}_{\text{recon}} + \beta\mathcal{L}_{\text{KL}} + \gamma\mathcal{L}_{\text{t}} \quad \text{(for VAEs)}.$$

Each dataset–architecture pair is trained over low, medium, and high weight settings $\alpha \in \{0, 1\}$, $\beta \in \{0, 0.08, 1\}$, $\gamma \in \{0, 1, 100\}$.

3. Results

Autoencoders trained without topological regularization show topological but not geometrical alignment between data and the latent embedding. Including topological regularization can substantially improve geometric alignment (see Figure 1), even when only regularizing on topology, without including the reconstruction loss in the training objective ($\alpha = 0$), leaving the decoder unconstrained. For datasets lying in a low-dimensional subspace, topological regularization yields near-perfect embeddings. However, the effect of topological regularization depends strongly on the choice of the considered topological features ($\dim_t$), and inappropriate settings can even disrupt topological alignment.

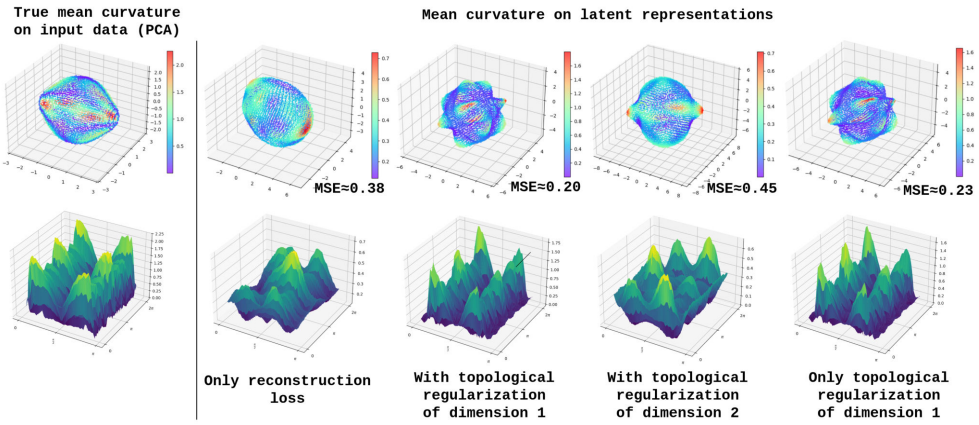

Figure 1: Impact of topological regularization on latent geometry. AE trained on $\text{Sphere}_{\text{high}}$ dataset. Top row: estimated curvature as heatmap. Bottom row: estimated curvature plotted over ground-truth angle. Scales vary across subplots.

For VAEs without topological regularization, the KL term ($\beta > 0$) pulls the latent embedding toward a Gaussian blob, preventing both topology and geometry preservation. Topological regularization can partly counteract this effect for low weights on the KL-loss term, but the latent geometry remains disrupted. Omitting the KL term makes VAEs behave similarly to AEs (see Figure 2).

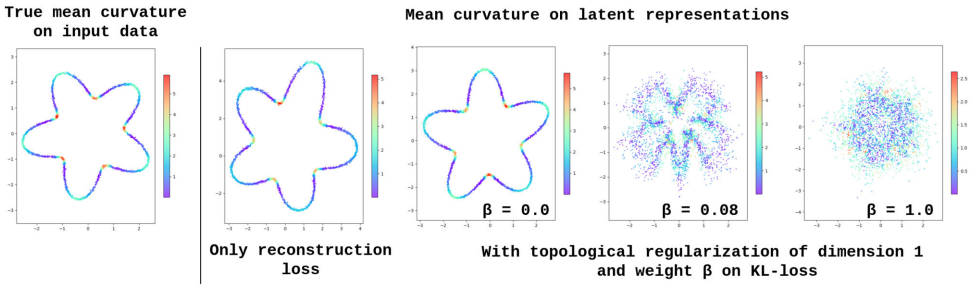

Figure 2: Impact of the KL term and topological regularization on the latent geometry of a VAE trained on the $\text{Circle}_{\text{low}}$ dataset. The KL-loss term ($\beta > 0$) disrupts topological and geometrical alignment. Scales vary across subplots.

4. Discussion and conclusions

We conducted an empirical study of autoencoders and variational autoencoders, evaluating their ability to preserve data topology and geometry in latent embeddings across six synthetic datasets. Local geometry was quantified via mean extrinsic curvature, estimated on the learned representation $\{z_i\}_{i=1}^N$ by fitting a local quadric.

For AEs, topological regularization $\mathcal{L}_t$ significantly improved geometric alignment between latent representations and data, even in the absence of the reconstruction objective. For VAEs, topological regularization acted as counterforce to the origin gravity created by the KL regularization term, but the resulting representations remained noisy and failed to preserve topology or geometry.

These findings suggest that (i) reliably topologically and geometrically aligned representations can potentially be learned by regulating only the pairwise distances between points, as indicated by the impact of the topological regularization term, and (ii) this may be achievable with a single neural network rather than a full autoencoder architecture, since the decoder became redundant when omitting the reconstruction objective.

4.1. Limitations and next steps

An immediate next step is to investigate whether architectures optimizing only for topological alignment without a decoder can yield reliable representations.

We evaluated local geometry using extrinsic curvature estimated via local quadrics, a naive method that restricts analysis to simple datasets. Future work should aim to broaden the scope toward more diverse and real-world datasets.

Topological alignment was only indirectly assessed. A similarity score based on persistent homology would provide a stronger measure. Future work should also examine intrinsic curvature, which is independent of ambient geometry. Finally, repeated experiments are needed to test whether topological regularization also leads to representational alignment across runs.

References

Francisco Acosta, Sophia Sanborn, Khanh Dao Duc, Manu Madhav, and Nina Miolane. Quantifying Extrinsic Curvature in Neural Manifolds. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 610–619, Vancouver, BC, Canada, June 2023. IEEE. ISBN 979-8-3503-0249-3. doi: 10.1109/CVPRW59228.2023.00068.

Y. Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35: 1798–1828, 08 2013. doi: 10.1109/TPAMI.2013.50.

Gunnar Carlsson, Tigran Ishkhanov, Vin de Silva, and Afra Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76(1): 1–12, 2008. ISSN 1573-1405. doi: 10.1007/s11263-007-0056-x. URL https://doi.org/10.1007/s11263-007-0056-x.

Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520, 2019. ISSN 1546-1726. doi: 10.1038/s41593-019-0460-x. URL https://doi.org/10.1038/s41593-019-0460-x.

Anna C Gilbert and Kevin O'Neill. Ca-pca: Manifold dimension estimation, adapted for curvature. *SIAM Journal on Mathematics of Data Science*, 7(1):355–383, 2025.

Søren Hauberg. Only Bayes should learn a manifold (on the estimation of differential geometric structure from data), 2019. URL http://arxiv.org/abs/1806.04994.

Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.

Yonghyeon Lee, Sangwoong Yoon, MinJun Son, and Frank C. Park. Regularized autoencoders for isometric representation learning. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=mQxt8l7JL04.

Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7045–7054. PMLR, 7 2020. URL https://proceedings.mlr.press/v119/moor20a.html.

Gregory Naitzat, Andrey Zhitnikov, and Lek-Heng Lim. Topology of deep neural networks. *Journal of Machine Learning Research*, 21(184):1–40, 2020. URL http://jmlr.org/papers/v21/20-345.html.

Philipp Nazari, Sebastian Damrich, and Fred A. Hamprecht. Geometric autoencoders: what you see is what you decode. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.

M. Yang and E. Lee. Segmentation of measured point data using a parametric quadric surface approximation. *Computer-Aided Design*, 31(7):449–457, 1999. ISSN 0010-4485. doi: https://doi.org/10.1016/S0010-4485(99)00042-1. URL https://www.sciencedirect.com/science/article/pii/S0010448599000421.

Appendix A. Persistent Homology and Topological Regularization

Persistent homology provides a way to estimate the topology of the underlying manifold of a point cloud. We can extract topological properties by tracking the birth and death of connected components, loops, and higher-dimensional cycles across a series of simplicial complexes with increasing connectivity. The Topological Autoencoder Moor et al. (2020) integrates this into training via a differentiable regularization term $\mathcal{L}_t$, aligning latent and input topology by matching lengths of *topologically relevant* edges. Let $A^{\mathcal{X}}$, $A^{\mathcal{Z}}$ be pairwise distance matrices and $\pi^{\mathcal{X}}$, $\pi^{\mathcal{Z}}$ persistence pairings in input/latent space. Then

$$\mathcal{L}_t = \frac{1}{2}\|A^{\mathcal{X}}[\pi^{\mathcal{X}}] - A^{\mathcal{Z}}[\pi^{\mathcal{X}}]\|^2 + \frac{1}{2}\|A^{\mathcal{Z}}[\pi^{\mathcal{Z}}] - A^{\mathcal{X}}[\pi^{\mathcal{Z}}]\|^2.$$

Here $A[\pi]$ extracts the lengths of topologically relevant edges. Minimizing $\mathcal{L}_t$ encourages topologically faithful embeddings.

Appendix B. Implementation details

Encoders and decoders use three hidden layers of width 100 with softplus activations, ensuring smoothness and making the decoder an immersion. Models are trained for 100 epochs with Adam optimizer (batch size 64). We implement models in `PyTorch`.

We measure the geometrical similarity of data and latent representation using both absolute and relative error metrics on curvature estimates κ:

$$\text{MSE} = \frac{1}{N}\sum_i (\kappa(x_i) - \kappa(z_i))^2, \quad \text{SMAPE} = \frac{1}{N}\sum_i \frac{|\kappa(x_i) - \kappa(z_i)|}{|\kappa(x_i)| + |\kappa(z_i)|}.$$

Appendix C. Parametrizations of the datasets

C.1. Low-deformed datasets

C.1.1. Circle

$$\phi_{\text{Circle}_{\text{low}}}(\theta) = A(\theta)\begin{pmatrix}\cos\theta\\ \sin\theta\end{pmatrix}, \quad \text{where } A(\theta) = r\,(1 + \alpha\cos(k\theta))$$

C.1.2. Sphere

$$\phi_{\text{Sphere}_{\text{low}}}(\theta,\varphi) = A(\theta)\begin{pmatrix}\sin\theta\cos\varphi\\ \sin\theta\sin\varphi\\ \cos\theta\end{pmatrix}, \quad \text{where } A(\theta) = 1 + \alpha e^{-5\theta^2} + \alpha e^{-5(\theta-\pi)^2}$$

C.1.3. Torus

$$\phi_{\text{Torus}_{\text{low}}}(\theta,\varphi) = A(\theta,\varphi)\begin{pmatrix}(R-r\cos\theta)\cos\varphi\\(R-r\cos\theta)\sin\varphi\\ r\sin\theta\end{pmatrix},$$
$$\text{where } A(\theta,\varphi) = 1+\alpha e^{-2(\varphi-\frac{\pi}{2})^2}e^{-2(\theta-\pi)^2}+\alpha e^{-2(\varphi-\frac{3\pi}{2})^2}e^{-2(\theta-\pi)^2}$$

C.2. High-deformed datasets

C.2.1. Circle

$$\phi_{\text{Circle}_{\text{high}}}(\theta) = (\sin\theta,\cos\theta,\sin(2\theta),$$
$$\alpha\cos(2\theta),\alpha\sin(3\theta),\alpha\cos(3\theta),\alpha\sin(4\theta),\alpha\cos(4\theta),\alpha\sin(5\theta),\alpha\cos(5\theta))^\top$$

C.2.2. Sphere

$$\phi_{\text{Sphere}_{\text{high}}}(\theta,\varphi) = \begin{pmatrix}A(\theta)r\sin\theta\cos\varphi\\A(\theta)r\sin\theta\sin\varphi\\A(\theta)r\cos\theta\\\sin\theta\cos(2\varphi)\\\sin(2\theta)\cos(3\varphi)\end{pmatrix}, \text{where } A(\theta) = 1+\tfrac{\alpha}{2}\left(e^{-5\theta^2}+e^{-5(\theta-\pi)^2}\right)$$

C.2.3. Torus

$$\phi_{\text{Torus}_{\text{high}}}(\theta,\varphi) = \begin{pmatrix}A_1(\theta,\varphi)(R-r\cos\theta)\cos\varphi\\A_1(\theta,\varphi)(R-r\cos\theta)\sin\varphi\\A_1(\theta,\varphi)r\sin\theta\\A_2(\theta,\varphi)\\\sin\theta\cos(2\varphi)\\\sin(2\theta)\cos(3\varphi)\end{pmatrix},$$
$$\text{where } A_1(\theta,\varphi) = 1+\tfrac{\alpha}{2}\Big(e^{-2(\varphi-\frac{\pi}{2})^2}e^{-2(\theta-\pi)^2}+e^{-2(\varphi-\frac{3\pi}{2})^2}e^{-2(\theta-\pi)^2}\Big),$$
$$\text{and } A_2(\theta,\varphi) = \alpha\Big(e^{-2(\varphi-\frac{\pi}{2})^2}e^{-2(\theta-\frac{\pi}{2})^2}+e^{-2(\varphi-\frac{3\pi}{2})^2}e^{-2(\theta-\frac{3\pi}{2})^2}\Big)$$

Appendix D. Curvature estimation algorithms

Algorithm 1: Estimate 1D Curvature via Quadratic Fitting

Input: Point cloud $P = \{p_1, \ldots, p_N\} \subset \mathbb{R}^n$, number of neighbors $k \in \mathbb{N}$
Output: Curvature estimates $\kappa = \{\kappa_1, \ldots, \kappa_N\} \subset \mathbb{R}$
for $i = 1$ **to** N **do**
  Find k nearest neighbors of p_i: $\mathcal{N}_i \subset P$;
  Compute local centroid: $c_i \leftarrow \frac{1}{k} \sum_{x \in \mathcal{N}_i} x$;
  Center neighborhood: $X_i \leftarrow \{x - c_i \mid x \in \mathcal{N}_i\}$;
  Perform PCA on X_i; obtain tangent t_i and normal $n_i \in \mathbb{R}^n$;
  Project neighbors: $x_j \leftarrow \langle X_j, t_i \rangle$, $z_j \leftarrow \langle X_j, n_i \rangle$;
  Fit parabola: $z_j \approx a x_j^2 + b x_j + c$ via least squares;
  Set $\kappa_i \leftarrow |2a|$;
end
return κ

Algorithm 2: Estimate 2D Mean Curvature via Quadratic Fitting

Input: Point cloud $P = \{p_1, \ldots, p_N\} \subset \mathbb{R}^n$, number of neighbors $k \in \mathbb{N}$
Output: Curvature estimates $\kappa = \{\kappa_1, \ldots, \kappa_N\} \subset \mathbb{R}$
for $i = 1$ **to** N **do**
  Find k nearest neighbors of p_i: $\mathcal{N}_i \subset P$;
  Compute local centroid: $c_i \leftarrow \frac{1}{k} \sum_{x \in \mathcal{N}_i} x$;
  Center neighborhood: $X_i \leftarrow \{x - c_i \mid x \in \mathcal{N}_i\}$;
  Perform PCA on X_i; obtain tangent basis $T_i \in \mathbb{R}^{2 \times n}$, normal $n_i \in \mathbb{R}^n$;
  Project neighbors onto local frame: $[x_j, y_j] \leftarrow T_i X_j^\top \in \mathbb{R}^2$, $z_j \leftarrow \langle X_j, n_i \rangle \in \mathbb{R}$;
  Fit quadric: $z_j \approx a x_j^2 + b y_j^2 + c x_j y_j + d x_j + e y_j + f$ via least squares;
  Set $\kappa_i \leftarrow |a + b|$;
end
return κ

Appendix E. Similarity measuring results

MSE and SMAPE error of the estimated curvature on the input $\{x_i\}_{i=1}^N$ data and the representation $\{z_i\}_{i=1}^N$ learned by AEs for different weight settings and datasets. Here α is the weight of the reconstruction loss $\mathcal{L}_{\text{recon}}$, γ the weight of the topological regularization term $\mathcal{L}_{\text{t}}$, β the weight of the KL loss term $\mathcal{L}_{\text{KL}}$ for VAEs. The value of $\dim_t$ sets the dimension up to which topological features in $\mathcal{L}_{\text{t}}$ are considered.

Table 1: MSE and SMAPE error of the estimated curvature on the input $\{x_i\}_{i=1}^N$ data and the representation $\{z_i\}_{i=1}^N$ learned by AEs for different weight settings and datasets.

α	γ	$\dim_t$	$\text{Circle}_{\text{low}}$		$\text{Circle}_{\text{high}}$		$\text{Sphere}_{\text{low}}$		$\text{Sphere}_{\text{high}}$		$\text{Torus}_{\text{low}}$		$\text{Torus}_{\text{high}}$	
			MSE	SMAPE	MSE	SMAPE	MSE	SMAPE	MSE	SMAPE	MSE	SMAPE	MSE	SMAPE
1	0	-	0.4509	16.53%	2.7776	69.40%	0.4102	28.45%	0.3836	42.10%	0.0329	21.44%	0.8026	50.19%
1	1	0	**0.0630**	14.44%	10.8412	**33.13%**	0.0504	8.27%	0.3024	41.06%	0.0009	4.16%	**0.2754**	**34.01%**
1	100	0	0.5671	20.39%	12.6125	34.08%	0.2032	15.00%	0.2784	39.54%	**0.0004**	**2.81%**	0.3150	35.37%
0	1	0	0.5587	20.30%	12.8751	34.51%	0.2081	15.15%	0.3049	38.33%	0.0006	3.14%	0.3171	35.61%
1	1	1	0.0820	**12.01%**	3.6430	52.55%	**0.0003**	**0.45%**	**0.1981**	**37.59%**	0.0158	13.23%	0.7181	52.58%
1	100	1	1.9507	44.60%	2.7692	64.42%	0.0092	3.56%	0.2468	41.49%	0.0203	14.30%	0.9048	65.42%
0	1	1	1.7502	41.90%	2.6157	65.67%	0.0077	3.19%	0.2301	40.79%	0.0184	12.70%	0.8989	64.69%
1	1	2	109.6815	65.88%	**2.1280**	49.61%	0.7751	50.25%	0.4535	53.31%	0.0483	26.38%	0.8652	55.09%
1	100	2	2342.8889	84.12%	2.4720	53.55%	0.8100	52.06%	0.4580	54.00%	0.1506	52.26%	1.0173	67.17%
0	1	2	64217.8867	84.00%	2.3382	53.52%	0.7440	48.28%	0.4682	54.47%	0.1554	51.75%	1.0158	65.96%

Appendix F. Supplementary figures

F.1. Autoencoders

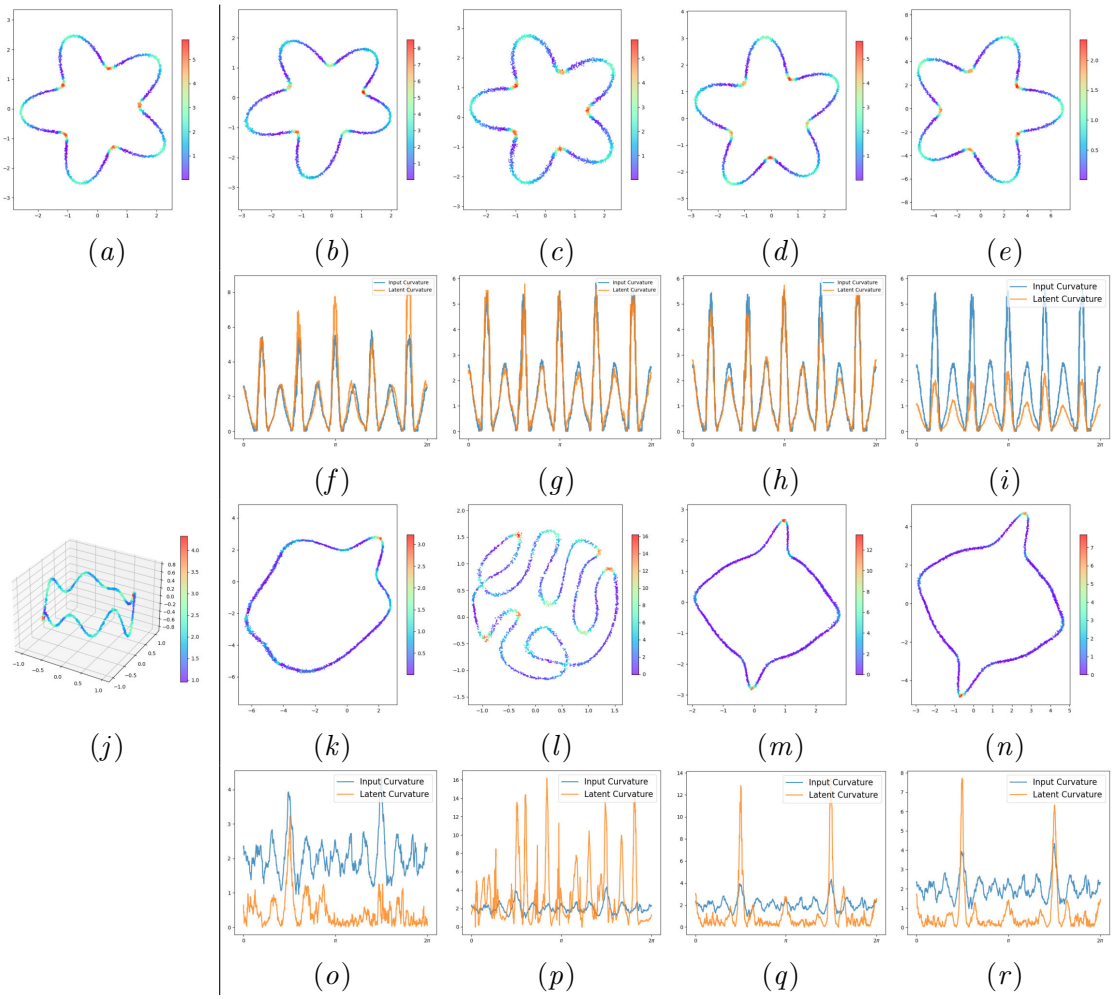

Figure 3: Euclidean AE trained on $\text{Circle}_{\text{low}}$ and $\text{Circle}_{\text{high}}$ datasets. Here $(\alpha, \gamma, \dim_t) =$ (reconstruction weight, topological weight, topological dimension). **Top two rows**: $\text{Circle}_{\text{low}}$. Topological regularization yields an almost perfect representation. (a) True curvature (PCA). (b)+(f) Only reconstruction, $(1, 0, -)$. (c)+(g) With topo. loss, $(1, 1, 0)$. (d)+(h) With topo. loss, $(1, 1, 1)$. (e)+(i) Only topo. loss $(0, 1, 1)$. **Bottom two rows**: $\text{Circle}_{\text{high}}$. Topological regularization can destroy topological alignment and focuses on high curvature areas. (j) True curvature (PCA). For the same weights as above: (k)+(o) Only reconstruction. (l)+(p), (m)+(q): With topo. loss. (n)+(r) Only topo. loss.

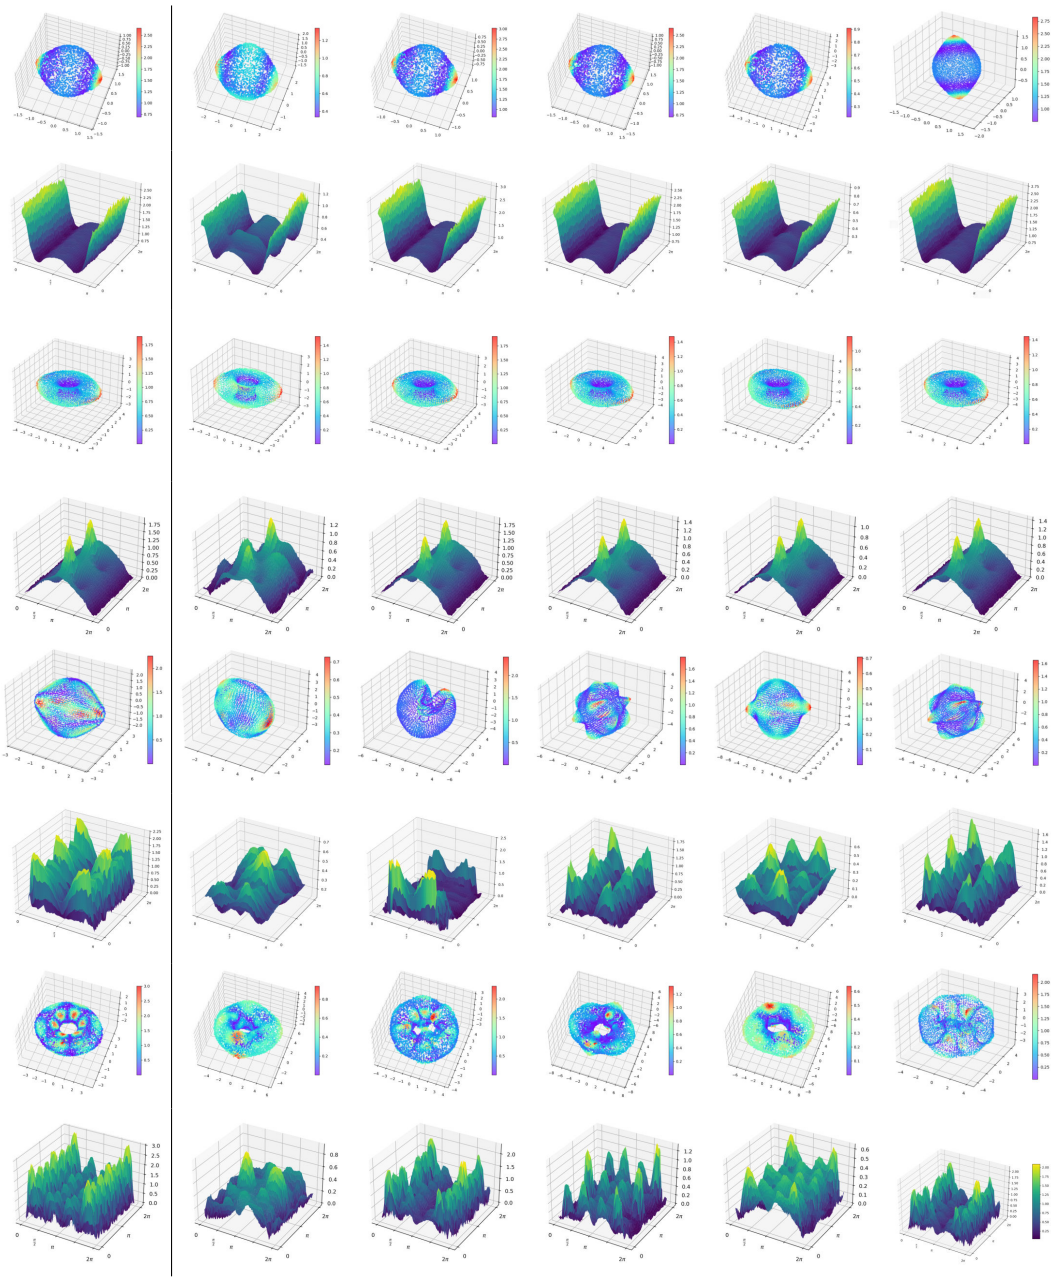

Figure 4: Autoencoder trained on $\text{Sphere}_{\text{low}}$, $\text{Torus}_{\text{low}}$, $\text{Sphere}_{\text{high}}$, $\text{Torus}_{\text{high}}$. Topological regularization gives near-perfect alignment on low-datasets and clear improvements on high-datasets depending on $\dim_t$. Each pair of rows shows curvature heatmaps (top) and curvature plots (bottom) over ground truth angle. Columns: Input (PCA), latent curvature with reconstruction loss $(1, 0, -)$, and with topological regularization $(1, 1, 0)$, $(1, 1, 1)$, $(1, 1, 2)$, $(0, 1, 1)$. Last column: only topological regularization for $\dim_t$ with best result. Tuples denote weights $(\alpha, \gamma, \dim_t)$.

F.2. Variational autoencoder

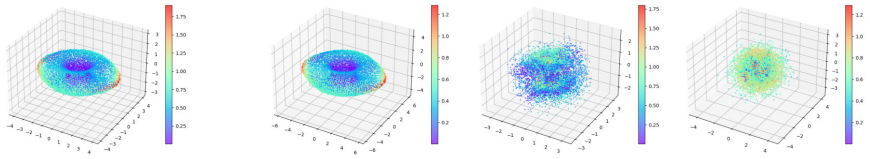

(a) (b) (c) (d)

Figure 5: Effect of topological regularization on latent geometry for a VAE trained on $\text{Torus}_{\text{low}}$. The KL term ($\beta > 0$) disrupts topological and geometrical alignment. (a) True curvature of input data (PCA). (b–d) Latent curvature with topological regularization ($\dim_t = 2$) and increasing β: (c) $\beta = 0$, (d) $\beta = 0.08$, (e) $\beta = 1$.