TOPOFORCE: LEARNING TOPOLOGY-REVEALING REPRESENTATIONS ACROSS MULTIPLE SCALES

Anonymous authors
Paper under double-blind review

ABSTRACT

Topology provides a fundamental description of high-dimensional data, yet the underlying topological organization is typically latent and can contain structures emerging at heterogeneous intrinsic scales. Conventional dimensionality-reduction methods optimize geometric or neighborhood fidelity, which need not preserve collective topology. Persistent-homology (PH)-guided methods use persistent structure inferred from the observed metric as a target for topology preservation, but this structure can become unreliable in high-dimensional noisy observations. We introduce TopoForce, a PH-free graph-based method for learning multiscale topology-revealing representations. TopoForce uses density-normalized multiscale diffusion to infer a pairwise topology scale that characterizes the structural scale relevant to each relationship, and translates these scales into adaptive force-directed interactions. Across controlled simulations, TopoForce recovers topological features spanning heterogeneous scales and remains robust to nonuniform sampling and high-dimensional nuisance variation. Across real-world applications in single-cell genomics, image analysis, and neural population activity, TopoForce reveals cyclic and toroidal organization consistent with known underlying structure. Together, these results demonstrate that TopoForce can reveal latent topological structure across multiple intrinsic scales in high-dimensional data. Our code is available at TopoForce.

1 INTRODUCTION

Topology provides a fundamental description of how data are organized, capturing collective structures such as connected components, cycles, voids, and higher-order topological features that emerge from relationships among many observations (Carlsson, 2009). In high-dimensional data, however, this underlying topology is typically latent: we observe only finite, noisy, and often nonuniform samples from which the intrinsic organization must be inferred (Carlsson, 2009; Damrich et al., 2024). For representation learning, topology is informative when the goal is to expose this underlying organization rather than faithfully reproduce the precise geometry of the observed data (Doraiswamy et al., 2021; Wagner et al., 2021). Importantly, topological structures can coexist across heterogeneous intrinsic scales within the same dataset (Edelsbrunner et al., 2002; Carlsson, 2009).

Revealing such multiscale topology poses requirements that differ from conventional geometric preservation. Methods emphasizing local geometry or neighborhood fidelity (van der Maaten & Hinton, 2008; McInnes et al., 2018) can preserve most local relationships while still destroying topology (Wagner et al., 2021; Liu et al., 2025): breaking only a small number of structurally critical relationships can open a cycle even when overall neighborhood preservation remains high (Figure 1a). Conversely, objectives with stronger global organization can expose large-scale structure at the cost of compressing or distorting finer topological patterns (Wang et al., 2021). Thus, geometric fidelity at any single scale does not determine how the observed geometry should be transformed to simultaneously expose topological structures across heterogeneous scales.

Persistent homology (PH) provides a principled description of topology across scales (Edelsbrunner et al., 2002) and has consequently motivated a growing family of PH-guided topology-preserving dimensionality-reduction methods (Moor et al., 2020; Doraiswamy et al., 2021; Wagner et al., 2021; Nelson & Luo, 2022; Trofimov et al., 2023; Clémot et al., 2025). Despite substantial differences in their formulations, these methods use persistent structure derived from a metric on the

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

observed high-dimensional data as a topological reference for the low-dimensional representation. This paradigm is powerful when the topology induced by the input metric faithfully reflects the intrinsic structure of interest. In real high-dimensional data, however, the observed geometry can be strongly affected by noise, nuisance variation, and sampling heterogeneity (Figure 1b). When these effects distort the persistence structure (Damrich et al., 2024), directly preserving it can propagate this distortion into the learned representation.

We introduce **TopoForce**, a PH-free, graph-based method for learning multiscale topology-revealing representations. Rather than prescribing the persistent homology of the observed data as a target, TopoForce uses multiscale diffusion on a neighborhood graph to infer a pairwise topology scale, capturing the intrinsic scales at which different structural relationships emerge. These scales directly modulate force-based interactions, allowing the representation to adapt its geometry to structures across heterogeneous scales while accounting for nonuniform sampling. Because PH is not used to guide the optimization, it can serve as an independent measure of the topology recovered by the representation.

Contributions. **(1) Pairwise topology scales.** From a neighborhood graph, we define pairwise topology scales that reflect which diffusion modes contribute most to each pair's separation and how slowly these modes decay. **(2) Scale-aware repulsion.** We introduce TopoForce, a PH-free method that uses these scales to modulate pairwise repulsion, together with a crossover schedule designed to encourage coarse-to-fine organization. **(3) Systematic evaluation.** We demonstrate multiscale topology recovery under geometric distortion, nonuniform sampling, and high-dimensional nuisance variation, supported by component ablations, non-cyclic controls, and real-data applications.

2 Related Work

Geometry- and neighborhood-preserving representations. PCA (Pearson, 1901) and MDS (Torgerson, 1952) emphasize global variance or pairwise distances; Isomap (Tenenbaum et al., 2000) approximates geodesic distances on a neighborhood graph; and Laplacian Eigenmaps (Belkin & Niyogi, 2003) and Diffusion Maps (Coifman & Lafon, 2006) exploit spectral graph structure. PHATE (Moon et al., 2019) constructs representations from diffusion-potential distances evaluated at a selected diffusion time. Neighborhood-preserving methods such as t-SNE (van der Maaten & Hinton, 2008) and UMAP (McInnes et al., 2018) prioritize local similarities, while force-directed layouts (Jacomy et al., 2014) balance graph attraction with repulsion. These methods primarily optimize geometric or neighborhood relationships rather than topological structure explicitly.

PH-guided topology-preserving representations. Persistent homology (PH) has motivated a complementary family of methods that explicitly use topological information derived from the input to guide the learned representation. TopoMap (Doraiswamy et al., 2021) constructively preserves 0-dimensional persistence through the minimum spanning tree, while Topological Autoencoders (TopoAE) (Moor et al., 2020) regularize topologically critical distances between input and latent spaces. DIPOLE (Wagner et al., 2021) combines local metric preservation with distributed persistence over corresponding subsets of the input and embedding. Interleaving DR (Nelson & Luo, 2022) optimizes interleavings between persistent structures, while RTD-AE (Trofimov et al., 2023) uses R-Cross-Barcodes to characterize discrepancies in topological features and their localization. Topological Autoencoders++ (TopoAE++) (Clémot et al., 2025) further targets persistent H_1 cy-

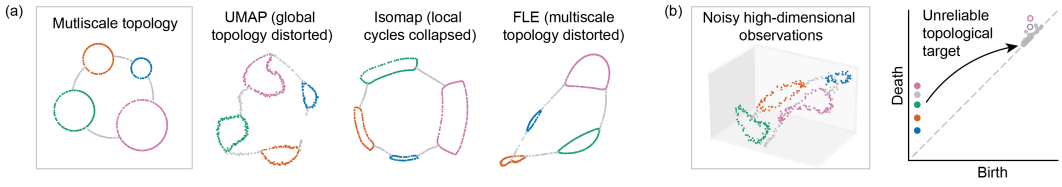

Figure 1: **Motivation.** (a) Topological structures coexist across scales, but standard representations can distort global topology, collapse local cycles, or fail to resolve multiple scales simultaneously. (b) Nuisance dimensions and measurement noise can obscure latent topology, as illustrated by the observed-space persistence diagram (right).

cles through cascade distortion, constraining their low-dimensional realization beyond persistence alone. Despite their different constructions, these methods share a *topology-preservation* paradigm in which persistent structure induced by a metric on the observed data provides a topological reference for the learned representation.

3 PRELIMINARIES

Graph spectrum and multiscale diffusion. Let $G = (V, E, W)$ be a weighted neighborhood graph with symmetric affinity matrix W and degree matrix $D = \mathrm{diag}(W\mathbf{1})$. The random-walk operator $P = D^{-1}W$ describes one-step diffusion over the graph. Let (λ_k, ψ_k) denote its eigenpairs, $P\psi_k = \lambda_k\psi_k$, ordered such that $1 = \lambda_0 > |\lambda_1| \geq |\lambda_2| \geq \cdots$, with $|\lambda_k| < 1$ for nontrivial modes. The eigenvectors ψ_k describe characteristic modes of variation over the graph, while the eigenvalues determine how these modes evolve under diffusion: after t steps, $P^t\psi_k = \lambda_k^t\psi_k$. **Modes with smaller $|\lambda_k|$ therefore decay rapidly and characterize finer-scale variation, whereas modes with $|\lambda_k|$ close to one persist longer and reflect larger-scale organization.** Diffusion thus turns the graph spectrum into a natural multiscale description of graph structure (Coifman & Lafon, 2006).

Importantly, the relevance of a spectral mode depends on the observations being compared. Mode k distinguishes observations i and j through the spectral contrast $\psi_k(i) - \psi_k(j)$, so different pairs can depend on different combinations of fast- and slow-relaxing modes. This pair-dependent spectral structure motivates the topology-scale construction introduced in Section 4.2.

Force-directed graph layout. Force-directed methods construct a low-dimensional representation by treating observations as interacting particles. Attractive forces along graph edges maintain local connectivity, while pairwise repulsion prevents collapse and promotes spatial separation. A central gravitational force is often added to limit drift and excessive expansion (Jacomy et al., 2014). The relative strength and spatial range of attraction and repulsion determine how local connectivity is balanced against larger-scale separation, providing a flexible mechanism for transforming graph structure into low-dimensional geometry.

4 TOPOFORCE

A central challenge in revealing multiscale topology is that repulsion that expands one structure can distort another. Rapidly decaying forces expand small cycles but lack long-range separation, whereas slowly decaying forces support large structures but can compress smaller ones through cumulative interactions. A shared decay profile therefore couples local expansion to long-range interactions, making different scales difficult to resolve simultaneously.

TopoForce addresses this trade-off by conditioning repulsion on embedding distance and pairwise structural scale (Figure 2). Given observations $\{\mathbf{x}_i\}_{i=1}^n$, it constructs a density-normalized diffusion graph, infers pairwise scales S_{ij} from relaxation times of distinguishing diffusion modes, and uses them to modulate force-directed repulsion. A crossover schedule encourages coarse-to-fine organization.

4.1 NEIGHBORHOOD GRAPH AND DENSITY-NORMALIZED DIFFUSION

TopoForce begins from a sparse graph that captures local connectivity while limiting density-dependent variation in graph support. Let $d_i^{(1)}$ be the nearest-neighbor distance of observation i and $r = \max_i d_i^{(1)}$. We construct the symmetrized neighborhood

$$\mathcal{N}(i) = \mathrm{kNN}_k(i) \cap \{j : d_{ij} \leq r\}, \tag{1}$$

where the radius limits neighborhood extent and the k-cap prevents excessive connectivity in dense regions.

On this fixed edge set, we assign $K_{ij} = \exp(-d_{ij}^2/\sigma^2)$ and apply $\alpha = 1$ diffusion density normalization (Coifman & Lafon, 2006; Haghverdi et al., 2015):

$$q_i = \sum_j K_{ij}, \qquad \widetilde{K}_{ij} = \frac{K_{ij}}{q_i q_j}, \qquad P = \widetilde{D}^{-1}\widetilde{K}, \quad \widetilde{D}_{ii} = \sum_j \widetilde{K}_{ij}. \tag{2}$$

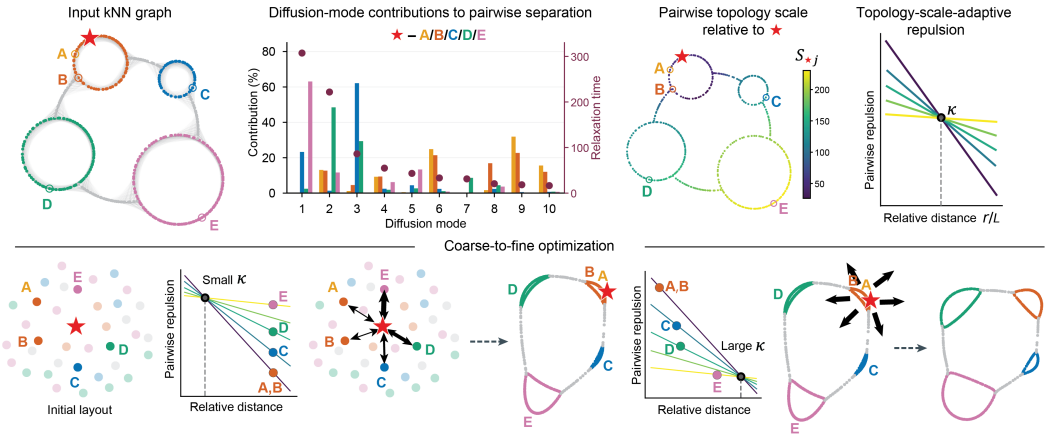

Figure 2: **Overview of TopoForce.** Given a neighborhood graph, pairwise topology scales reflect which diffusion modes contribute most to each pair's separation and how slowly these modes decay, guiding adaptive repulsion (top). Increasing κ encourages coarse-to-fine organization (bottom).

This separates graph support from diffusion weights and reduces sampling-density effects on spectral analysis. Following PHATE (Moon et al., 2019), we select the reference time t_{VNE} at the knee of the diffusion spectrum's Von Neumann entropy curve.

4.2 Pairwise Topology Scale

From a global diffusion time to pairwise structural scales. Although t_{VNE} provides a global reference, different pairwise relationships need not occur at the same structural scale. Our key idea is therefore to ask, for each pair (i, j), *which diffusion modes distinguish the pair, and how slowly do those modes relax?* We call the resulting quantity a *pairwise topology scale*: it characterizes the intrinsic scale of a pairwise structural relationship, while topological features emerge collectively from such relationships across the graph.

At t_{VNE}, the contribution and normalized weight of mode k are

$$c_k^{(ij)} = \left[\lambda_k^{t_{\mathrm{VNE}}}\left(\psi_k(i) - \psi_k(j)\right)\right]^2, \qquad q_k^{(ij)} = \frac{c_k^{(ij)}}{\sum_{\ell \geq 1} c_\ell^{(ij)}}. \tag{3}$$

Here $q_k^{(ij)}$ measures how strongly mode k distinguishes the pair (Coifman & Lafon, 2006). Since a mode with eigenvalue magnitude $|\lambda_k|$ has relaxation time $\tau_k = -1/\log|\lambda_k|$, we define

$$S_{ij} = \sum_{k \geq 1} q_k^{(ij)} \tau_k. \tag{4}$$

Thus S_{ij} is large when slowly relaxing modes distinguish a pair and small when fast modes dominate. All pairs share t_{VNE}; pair specificity arises from the spectral contrasts $\psi_k(i) - \psi_k(j)$.

4.3 Scale-Adaptive Force-Directed Embedding

To reveal topology across scales simultaneously, we use locally weaker but longer-range repulsion for coarser-scale relationships, and stronger short-range repulsion with faster decay for finer-scale ones. We encode relative scale by the symmetric rank score

$$\widehat{g}_{ij} = \frac{\mathrm{rank}_i(S_{ij}) + \mathrm{rank}_j(S_{ji}) - 2}{2(n-1)}, \tag{5}$$

where ranks include the zero diagonal and average ties.

Let $r_{ij} = \|\mathbf{z}_i - \mathbf{z}_j\|_2$, $\mathbf{u}_{ij} = (\mathbf{z}_i - \mathbf{z}_j)/r_{ij}$, and let L be the median embedding length of graph edges, recomputed each iteration. Relative density ρ_i is estimated from Gaussian row sums before

neighborhood pruning, including self-affinities, and normalized to unit mean. Repulsion is

$$\mathbf{F}_{ij}^{\text{rep}} = \frac{1}{L\sqrt{\rho_i \rho_j}} \left(\frac{\kappa}{r_{ij}/L} \right)^{1-\widehat{g}_{ij}} \mathbf{u}_{ij}. \tag{6}$$

Larger $\widehat{g}_{ij}$ gives slower distance decay, while the density factor strengthens repulsion in sparsely sampled regions. The parameter κ controls the relative-distance crossover of these scale-dependent interactions (Section 4.4).

With the same edge attraction and central gravity terms as FLE, the total force is

$$\mathbf{F}_i = c_{\text{attr}} \sum_{j:(i,j)\in E} W_{ij}(\mathbf{z}_j - \mathbf{z}_i) + c_{\text{rep}} \sum_{j \neq i} \mathbf{F}_{ij}^{\text{rep}} - c_{\text{grav}} m_i \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|_2}, \tag{7}$$

where $m_i = 1 + \deg(i)$ and gravity is zero at the origin. The coefficients c_{attr}, c_{rep}, and c_{grav} weight the three forces. Optimization builds on the ForceAtlas2-based FLE implementation in Pegasus (Jacomy et al., 2014; Li et al., 2020), with coefficients reimplemented and updated in Python.

4.4 Coarse-to-Fine Optimization

To expose the role of κ, write $\widehat{r}_{ij} = r_{ij}/L$ and suppress pair indices:

$$F^{\text{rep}}(\widehat{r}) \propto \left(\frac{\kappa}{\widehat{r}} \right)^{1-\widehat{g}}. \tag{8}$$

At $\widehat{g} = 0$, $F_{\text{fine}} \propto \kappa/\widehat{r}$, whereas at $\widehat{g} = 1$, $F_{\text{coarse}} \propto 1$. The two limits intersect at $\widehat{r}_{\times} = \kappa$, or $r_{\times} = \kappa L$. Thus, κ defines the crossover distance, relative to the current median graph-edge length, at which fine- and coarse-scale interaction laws exchange dominance.

Increasing κ preferentially strengthens finer-scale interactions, since at a fixed layout their magnitudes scale as $\kappa^{1-\widehat{g}_{ij}}$. This motivates a schedule that promotes coarse organization before finer-scale relationships, while keeping the pairwise scores $\widehat{g}_{ij}$ fixed throughout.

We use an data-adaptive schedule that start at $\kappa = 1$ and advance through pairwise-scale quantiles to $\kappa_{\max} = \max_{i<j} S_{ij}$. Since each diffusion step corresponds to one graph transition, $\max S_{ij}$ reflects the largest inferred relaxation scale; interpreting κ in median-edge-length units sets the largest crossover range reached during optimization. For disconnected graphs, quantiles and the shared endpoint use only within-component scales. Each stage continues from the preceding layout until movement plateaus or the iteration limit is reached, and the final-stage embedding is returned.

5 Experiments

We evaluate whether TopoForce (i) recovers topology across heterogeneous intrinsic scales, (ii) remains robust to geometric distortion, sampling heterogeneity, and high-dimensional nuisance variation, (iii) reveals structure consistent with independent evidence in real data, and (iv) yields interpretable component contributions and computational scaling via ablation and runtime analyses.

5.1 Experimental Setup

Benchmarks and observations. To separate structural from observational effects, we use circle, ring, bracelet, three unequal loops, torus, and sphere with known H_1 or H_2 targets, together with clusters and branching trajectories as non-cyclic controls shown in the Appendix A(Figure A1 and A3). Coordinates are planar except for the torus and sphere in $\mathbb{R}^3$.

Baselines. We compare TopoForce with UMAP, t-SNE, Isomap, PHATE, FLE, DIPOLE, RTD-AE, and TopoAE++. All baselines use their default hyperparameters, with only the output dimensionality set to match the benchmark: 2D for planar data and 3D for the torus and sphere. FLE uses ForceAtlas2 and Fruchterman–Reingold, respectively. TopoAE++ did not complete all native 3D runs within 12 hours; we therefore omit its torus and sphere results.

Topology recovery score. TRS measures target recovery and separation from spurious features. We median-normalize pairwise distances and compute barcodes on shared subsets: up to 1000 points

Dataset	Setting	TopoForce	UMAP	t-SNE	FLE	PHATE	Isomap	DIPOLE	TopoAE++	RTD-AE
Circle (H_1)	U	**0.995**	0.952	0.684	0.886	0.962	0.990	0.986	0.775	0.969
	UD	0.986	0.937	0.646	0.975	0.952	**0.988**	0.977	0.730	0.716
	NU	**0.991**	0.964	0.000	0.624	0.964	0.986	0.974	0.972	0.973
Ring (H_1)	U	**0.994**	0.940	0.487	0.981	0.941	0.828	0.475	0.735	0.942
	UD	**0.986**	0.959	0.960	0.963	0.258	0.911	0.473	0.792	0.777
	NU	**0.977**	0.891	0.479	0.966	0.913	0.927	0.485	0.951	0.949
Bracelet (H_1)	U	**0.970**	0.937	0.947	0.606	0.663	0.919	0.641	0.875	0.876
	UD	**0.877**	0.817	0.445	0.431	0.448	0.629	0.686	0.659	0.704
	NU	**0.891**	0.591	0.465	0.393	0.507	0.851	0.659	0.889	0.842
Three loops (H_1)	U	**0.975**	0.427	0.548	0.567	0.511	0.923	0.961	0.911	0.911
	UD	**0.963**	0.705	0.454	0.411	0.512	0.329	0.608	0.765	0.708
	NU	**0.976**	0.608	0.413	0.566	0.373	0.936	0.955	0.923	0.914
Torus (H_1)	U	0.633	0.672	**0.737**	0.672	0.677	0.689	0.649	—	0.696
	NU	0.661	0.655	0.627	0.644	0.628	**0.680**	0.662	—	0.641
Torus (H_2)	U	0.548	0.508	**0.700**	0.000	0.000	0.399	0.611	—	0.439
	NU	**0.790**	0.000	0.000	0.469	0.000	0.603	0.722	—	0.376
Sphere (H_2)	U	0.972	0.987	0.980	**1.000**	0.981	**1.000**	0.994	—	**1.000**
	NU	**1.000**	0.991	0.987	0.995	0.990	**1.000**	0.992	—	**1.000**

Table 1: Synthetic TRS under uniform (U), distorted uniform (UD), and nonuniform (NU) sampling. Red/orange: best/second-best distinct scores. —: omitted result.

for H_1 and 256 for H_2. Cycle matching identifies corresponding features through image persistence in a common filtration on shared point identities (García-Redondo et al., 2024). For the generator-specified count $K > 0$, targets are the longest positive reference intervals. With recovered targets $\mathcal{M}$ and matched embedding lifetimes ℓ_i, we define

$$\mathrm{TRS} = \frac{1}{K} \sum_{i \in \mathcal{M}} \frac{\ell_i}{\ell_i + b}, \qquad b = \max_{j \in \mathcal{U}} \ell_j, \tag{9}$$

where $\mathcal{U}$ contains features unmatched to targets, with $b = 0$ if empty. Missing targets contribute zero; equal weighting prevents long-lived targets from dominating. TRS lies in $[0, 1]$; larger values indicate better recovery and separation. Torus H_1 and H_2 are scored separately. For controls ($K = 0$), TRS is undefined; their persistence diagnostics are reported in the Appendix A(Figure A2).

5.2 Recovering Multiscale Topology under Geometric Distortion

To test recovery across heterogeneous topological scales, we compare uniform samples (U) with geometrically distorted counterparts (UD) (Figure 3; Table 1). The bracelet couples local cycles to a larger central cycle, whereas the three-loop benchmark contains unequal cycles. TopoForce preserves these structures in both settings, while deformation degrades several competing methods. Separation is strongest on the distorted multiscale benchmarks, where observed geometry is least informative about latent topology.

Beyond planar multiscale cycles, most methods recover the sphere's dominant H_2 feature, while t-SNE performs best on the uniform torus. These results suggest that TopoForce's primary advantage is not universal topological reconstruction, but recovery when multiple structural scales must be resolved under distorted observations. Complete embedding and persistence grids for all benchmarks and sampling settings are given in Appendix A.

5.3 Robustness to Sampling Heterogeneity

To isolate sampling heterogeneity, we compare nonuniform (NU) with uniform (U) sampling while holding structure, dimension, and observation noise fixed (Figure 4; Table 1). Three-loop recovery remains essentially unchanged under the density shift, whereas the bracelet shows a larger but still partial degradation. The nonuniform torus further separates recovery across homology dimensions, with TopoForce strongest on H_2 and Isomap on H_1. These results indicate that sampling heterogeneity affects structures differently rather than producing a uniform loss of topology. Section 5.6 isolates the contribution of density-aware repulsion.

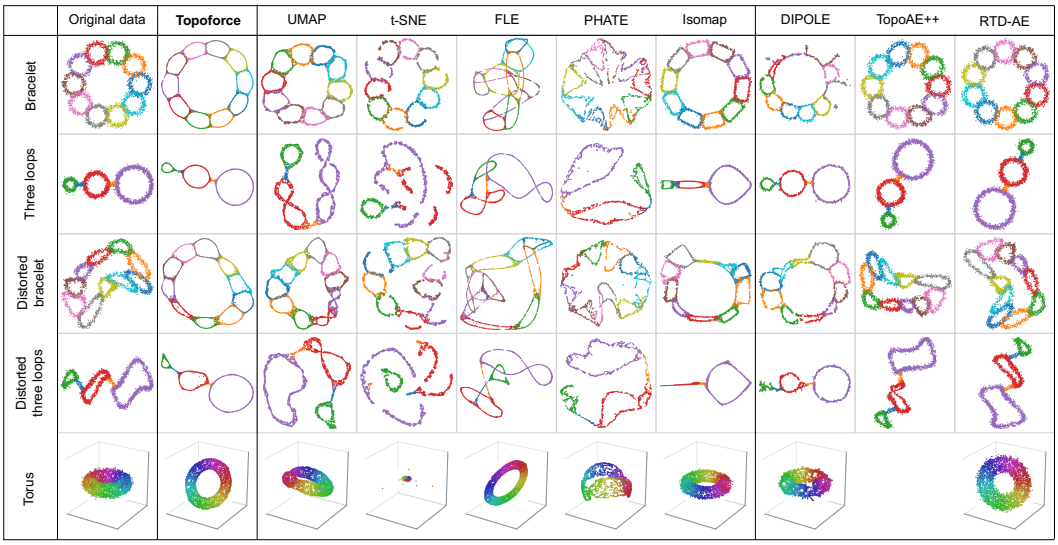

Figure 3: TopoForce preserves multiscale topology under geometric distortion, resolving both local cycles and their larger-scale organization.

5.4 Topology Recovery from High-Dimensional Observations

To test robustness to irrelevant high-dimensional variation, we append up to 2048 Gaussian nuisance dimensions ($\sigma = 0.6$) to the same 1000 observations while keeping the latent topology fixed. Methods embed the full noisy vectors, while evaluation uses the unchanged latent structure (Figure 5).

As nuisance dimensions increase, input-space persistence collapses while recovery degrades across methods (Figure 5a,c). TopoForce nevertheless retains matched target structure into the highest-nuisance regime, where several alternatives lose most or all target contribution (Figure 5b,c).

Recovery deteriorates as nuisance variation increasingly dominates the observed geometry, with rankings varying at intermediate dimensions. Input-space target persistence progressively collapses toward the diagonal (Figure 5a), indicating increasingly uninformative observed geometry. TopoForce nevertheless retains matched target structure into the highest-nuisance regime, where several alternatives lose most or all target contribution.

This separation between input-space persistence and embedding recovery is important: TopoForce can recover target topology even after it is strongly obscured in the observed geometry, rather than merely preserving an already visible loop. Its persistence diagrams further distinguish genuine recovery from embeddings that only resemble loops visually. Across the full sweep, recovery eventually degrades under sufficient nuisance variation. Full circle and torus sweeps, including persistence diagrams, are given in Appendix B.

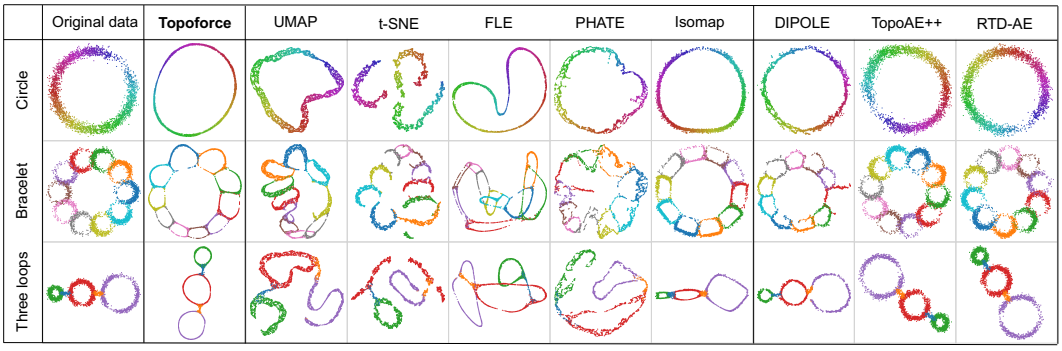

Figure 4: TopoForce remains robust to uneven sampling, preserving multiscale topology despite strong density heterogeneity.

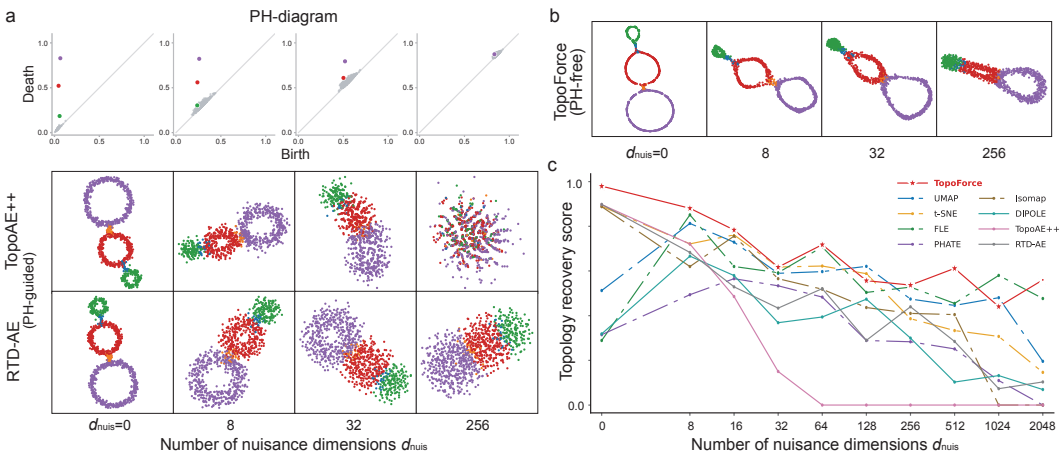

Figure 5: **High-dimensional nuisance robustness. (a)** Input-space persistence approaches the diagonal as nuisance dimensions d_{nuis} increase, indicating loss of the target topological signal. **(b)** TopoForce retains target structure after the observed geometry is strongly corrupted. **(c)** Recovery across methods over the full nuisance sweep.

5.5 Real Data with Independently Supported Topology

We next test whether TopoForce reveals organization consistent with independently supported structure in rotational images (Nene et al., 1996), neural population activity (Gardner et al., 2022), and cell-cycle transcriptomes (Mahdessian et al., 2021)(Figure 6).

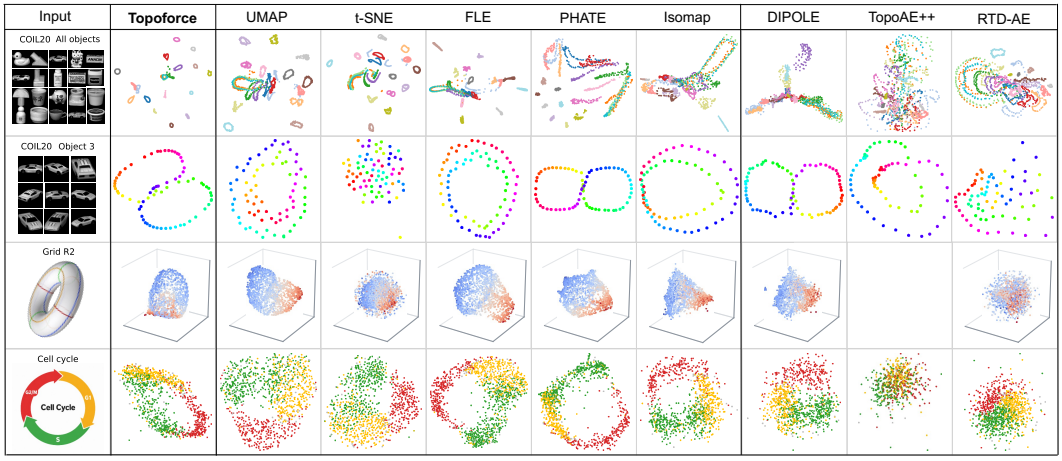

Figure 6: TopoForce reveals structure consistent with independently supported cyclic or toroidal organization across rotational images, grid-cell activity, and cell-cycle transcriptomes.

Evaluation. COIL-20 circular order is the absolute pairwise-sine correlation between centroid-based embedding angles and viewing angles, each converted to rank angles $2\pi(r_i - 1)/n$ with average ranks for ties, and averaged across objects. Cell-cycle loop dominance is $\text{LD} = (\ell_1 - \ell_2)/\ell_1$, for the longest H_1 lifetimes $\ell_1 \geq \ell_2$; missing lifetimes and LD when $\ell_1 = 0$ are zero. Both scores lie in $[0, 1]$; LD measures loop prominence, not phase agreement. Cell-cycle axes are rescaled to equal marginal standard deviations without labels before display and scoring.

Dataset	TopoForce	UMAP	t-SNE	FLE	PHATE	Isomap	DIPOLE	TopoAE++	RTD-AE
COIL-20 (all)	**0.488**	0.307	0.374	0.249	0.132	0.183	0.123	0.185	0.201
COIL-20 (Obj. 3)	0.035	0.015	0.020	0.018	0.023	0.022	**0.054**	0.012	0.004
Cell cycle	0.695	0.642	0.755	**0.834**	0.661	0.785	0.386	0.042	0.165

Table 2: COIL-20 circular order and cell-cycle loop dominance. Colors follow Table 1.

COIL-20. To test rotational agreement, we compare embedding order against viewing angle using 1440 images of 20 objects in 50-PC space. TopoForce achieves the highest circular order (Table 2), while Object 3 remains challenging across all methods.

Grid cells. To examine neural population geometry, we embed 1500 Grid R2 activity vectors from six PCs into 3D. TopoForce reveals a continuous curved organization consistent with a low-dimensional manifold, though the embedding alone does not establish toroidal topology.

Cell cycle. To assess cyclic organization, we embed 990 cells from noisy, high-dimensional transcriptomes using 40 PCs, colored by experimentally measured FUCCI phase. TopoForce reveals a cycle with graded phase colors and greater loop dominance than PH-guided baselines, although FLE and Isomap score higher.

5.6 Ablation and Scalability

To isolate the mechanisms underlying multiscale topology recovery under sampling heterogeneity, we use the nonuniform bracelet, which combines cycles at different scales with uneven sampling (Figure 7). We vary one component at a time while holding the input, graph, initialization, optimization budget, and evaluation set fixed.

Pair-specific topology scales. Averaging $\widehat{g}_{ij}$ across pairs collapses several cycles and causes the largest TRS drop, showing that pairwise scale heterogeneity is critical for resolving topology across scales.
Density-aware repulsion. Removing $1/\sqrt{\rho_i \rho_j}$ from Eq. 6 preserves all target matches but lowers TRS, indicating that density correction primarily improves target–background separation rather than creating topology.
Adaptive κ schedule. Fixing κ at its terminal value loses one target and lowers TRS despite using the same topology scales and density correction, showing that recovery depends on the coarse-to-fine optimization path, not only the terminal force law.

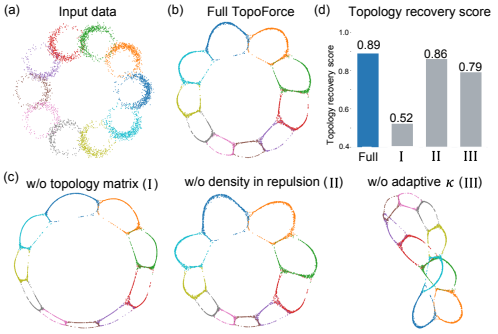

Figure 7: TopoForce ablations on the nonuniform bracelet. Pair-specific scales contribute most; density-aware repulsion improves separation, and adaptive κ preserves another target.

Scalability. We assess computational cost by varying circle sample size (Table 3), using two CPU threads on an AMD EPYC 75F3, one run per setting, and four TopoForce κ stages. The approximately quadratic runtime reflects dense pairwise scale computation and direct force evaluation.

n	TopoForce	UMAP	t-SNE	FLE	PHATE	Isomap	DIPOLE	TopoAE++	RTD-AE
500	61.79	17.51	1.11	16.08	2.01	0.41	13.36	5.39	176.21
2000	133.26	18.25	6.07	22.18	7.84	2.35	16.81	86.61	428.44
8000	2045.57	28.22	30.02	56.81	12.98	41.55	71.88	—	1409.14

Table 3: Runtime in seconds on the circle benchmark. —: failed run.

6 Conclusion

We introduced TopoForce, a PH-free method for learning topology-revealing representations by adapting repulsion to pairwise structural scales. Controlled experiments demonstrate multiscale topology recovery under geometric distortion, nonuniform sampling, and high-dimensional nuisance variation. Ablations support the contributions of pair-specific scales, density-aware repulsion, and the optimization schedule, while real-data applications demonstrate utility in complex, noisy observations. These results suggest that, when neighborhood connectivity remains informative, scale-aware repulsion can reorganize geometry to reveal latent topology. Recovery still degrades under sufficiently strong nuisance variation, and pairwise computations limit scalability. Future work will address these limitations and characterize conditions for reliable recovery.

AI USE STATEMENT

We used generative AI tools (GPT-6 Astra) for two tasks with recommended disclosure: (i) drafting portions of the algorithmic code and the simulation code used to generate synthetic data, in both cases subsequently reviewed, tested, and in part rewritten by the authors to ensure correctness and conformity with our design; and (ii) checking the manuscript for formatting and typographical issues, such as LaTeX layout, cross-references, and citation style consistency.

We have not used generative AI tools for any task that requires mandatory disclosure: the research idea, problem formulation, methodology, the design of the simulation setup (including the data-generating process and evaluation protocol), the analysis, and the writing of the manuscript were produced by the authors. The remaining mandatory-disclosure categories are not applicable to this work.

All AI-assisted output was reviewed by the authors, who take full responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.

Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009. doi: 10.1090/S0273-0979-09-01249-X.

Mattéo Clémot, Julie Digne, and Julien Tierny. Topological autoencoders++: Fast and accurate cycle-aware dimensionality reduction. *IEEE Transactions on Visualization and Computer Graphics*, 2025.

Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.

Sebastian Damrich, Philipp Berens, and Dmitry Kobak. Persistent homology for high-dimensional data based on spectral methods. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/4a32a646254d2e37fc74a38d65796552-Abstract-Conference.html.

Harish Doraiswamy, Julien Tierny, Paulo J. S. Silva, Luis Gustavo Nonato, and Claudio T. Silva. TopoMap: A 0-dimensional homology preserving projection of high-dimensional data. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):561–571, 2021. doi: 10.1109/TVCG.2020.3030441.

Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002. doi: 10.1007/s00454-002-2885-2.

Inés García-Redondo, Anthea Monod, and Anna Song. Fast topological signal identification and persistent cohomological cycle matching: I. garcía-redondo et al. *Journal of Applied and Computational Topology*, 8(3):695–726, 2024.

Richard J. Gardner, Erik Hermansen, Marius Pachitariu, Yoram Burak, Nils A. Baas, Benjamin A. Dunn, May-Britt Moser, and Edvard I. Moser. Toroidal topology of population activity in grid cells. *Nature*, 602(7895):123–128, 2022. doi: 10.1038/s41586-021-04268-7.

Laleh Haghverdi, Florian Buettner, and Fabian J. Theis. Diffusion maps for high-dimensional single-cell analysis of differentiation data. *Bioinformatics*, 31(18):2989–2998, 2015. doi: 10.1093/bioinformatics/btv325.

Mathieu Jacomy, Tommaso Venturini, Sebastien Heymann, and Mathieu Bastian. ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. *PLOS ONE*, 9(6):e98679, 2014.

Bo Li, Joshua Gould, Yiming Yang, Siranush Sarkizova, Marcin Tabaka, Orr Ashenberg, Yanay Rosen, Michal Slyper, Monika S. Kowalczyk, Alexandra-Chloé Villani, Timothy Tickle, Nir Hacohen, Orit Rozenblatt-Rosen, and Aviv Regev. Cumulus provides cloud-based data analysis for large-scale single-cell and single-nucleus rna-seq. *Nature Methods*, 17(8):793–798, 2020. doi: 10.1038/s41592-020-0905-x.

Zhexuan Liu, Rong Ma, and Yiqiao Zhong. Assessing and improving reliability of neighbor embedding methods: a map-continuity perspective. *Nature Communications*, 16:5037, 2025. doi: 10.1038/s41467-025-60434-9.

Diana Mahdessian, Anthony J. Cesnik, Christian Gnann, Frida Danielsson, Lovisa Stenström, Muhammad Arif, Cheng Zhang, Trang Le, Fredric Johansson, Rutger Schutten, Anna Bäckström, Ulrika Axelsson, Peter Thul, Nathan H. Cho, Oana Carja, Mathias Uhlén, Adil Mardinoglu, Charlotte Stadler, Cecilia Lindskog, Burcu Ayoglu, Manuel D. Leonetti, Fredrik Pontén, Devin P. Sullivan, and Emma Lundberg. Spatiotemporal dissection of the cell cycle with single-cell proteogenomics. *Nature*, 590(7847):649–654, 2021. doi: 10.1038/s41586-021-03232-9.

Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

Kevin R. Moon, David van Dijk, Zheng Wang, Scott Gigante, Daniel B. Burkhardt, William S. Chen, Kristina Yim, Antonia van den Elzen, Matthew J. Hirn, Ronald R. Coifman, Natalia B. Ivanova, Guy Wolf, and Smita Krishnaswamy. Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology*, 37(12):1482–1492, 2019. doi: 10.1038/s41587-019-0336-3.

Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7045–7054. PMLR, 2020.

Bradley J. Nelson and Yuan Luo. Topology-preserving dimensionality reduction via interleaving optimization. *arXiv preprint arXiv:2201.13012*, 2022.

S. A. Nene, S. K. Nayar, and H. Murase. Columbia object image library (COIL-20). Technical Report CUCS-005-96, Department of Computer Science, Columbia University, 1996.

Karl Pearson. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.

Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

Warren S. Torgerson. Multidimensional scaling: I. theory and method. *Psychometrika*, 17(4):401–419, 1952.

Ilya Trofimov, Daniil Cherniavskii, Eduard Tulchinskii, Nikita Balabin, Serguei Barannikov, and Evgeny Burnaev. Learning topology-preserving data representations. In *International Conference on Learning Representations*, 2023.

Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

Alexander Wagner, Elchanan Solomon, and Paul Bendich. Improving metric dimensionality reduction with distributed topology. *arXiv preprint arXiv:2106.07613*, 2021.

Yingfan Wang, Haiyang Huang, Cynthia Rudin, and Yaron Shaposhnik. Understanding how dimension reduction tools work: An empirical approach to deciphering t-SNE, UMAP, TriMap, and PaCMAP for data visualization. *Journal of Machine Learning Research*, 22(201):1–73, 2021. URL `https://www.jmlr.org/papers/v22/20-1061.html`.

APPENDIX

This appendix provides the complete visual and quantitative evidence behind the controlled experiments of the main text. Appendix A gives the full embedding and persistence-diagram grids for all eight synthetic benchmarks under uniform, distorted-uniform, and nonuniform sampling, including the non-cyclic controls used to assess spurious topology but not displayed in the main text. Appendix B extends the nuisance-dimension sweep of Section 5.4 to the circle and torus benchmarks, with embeddings and persistence diagrams at all ten noise levels. Appendix C lists the graph, diffusion, optimization, and evaluation settings, and Appendix D reports extended label-based diagnostics on the full COIL-20 set.

A SYNTHETIC EMBEDDING AND PERSISTENCE GRIDS

The main text reports topology recovery scores for all benchmarks and sampling settings (Table 1) but shows embedding grids only for a subset of rows (Figures 3 and 4). This section gives the complete grids: for each sampling setting, one figure of embeddings for all eight benchmarks and all methods, and, where available, one figure of the corresponding birth–death persistence diagrams. Together, they provide the visual evidence behind the per-benchmark scores, including the clusters and branching non-cyclic controls, for which TRS is undefined.

In every grid, rows are benchmarks and columns are methods, in the order TopoForce, UMAP, t-SNE, FLE, PHATE, Isomap, DIPOLE, TopoAE++, RTD-AE, preceded by the input (embeddings) or the clean reference (persistence diagrams). Colors encode ground-truth position or component identity on the benchmark, so a faithful representation shows coherent color regions; torus and sphere rows are three-dimensional. TopoAE++ results are absent for the torus and sphere, whose native 3D runs did not complete within the runtime budget.

A.1 UNIFORM SAMPLING

Under uniform sampling (Figure A1), all methods recover the circle's dominant H_1 feature and the sphere's dominant H_2 feature, while the multi-cycle benchmarks already separate them: t-SNE and PHATE collapse the bracelet and three-loop structures, and several methods merge the three unequal loops into a single ring. The clusters and branching controls contain no cyclic structure by construction; the corresponding persistence diagrams (Figure A2) show intervals concentrated along the diagonal for TopoForce on both controls, supporting that the cycles it recovers on the cyclic benchmarks are not sampling artifacts of the force-directed layout.

A.2 GEOMETRIC DISTORTION

Distortion is applied only to the planar cyclic coordinates, so the clusters, branching, torus, and sphere rows of Figure A3 are identical to those of Figure A1; the circle, ring, bracelet, and three-loop rows are the distorted counterparts. Distortion breaks the correspondence between observed geometry and latent topology: methods whose uniform embeddings resolve these cycles lose, merge, or open them once the coordinates are warped, while TopoForce retains the local cycles and their larger-scale organization. This is the regime in which the score separation in Table 1 is largest.

A.3 NONUNIFORM SAMPLING

Under nonuniform sampling (Figure A4), density heterogeneity degrades methods unevenly across benchmarks rather than uniformly: UMAP collapses the circle to a single blob, t-SNE loses the ring and torus, and the bracelet degrades for most methods including TopoForce, albeit partially. The persistence diagrams (Figure A2 vs. Figure A5) show that the matched target intervals survive the density shift for TopoForce on the circle, ring, and three-loop benchmarks, consistent with the essentially unchanged scores in Table 1.

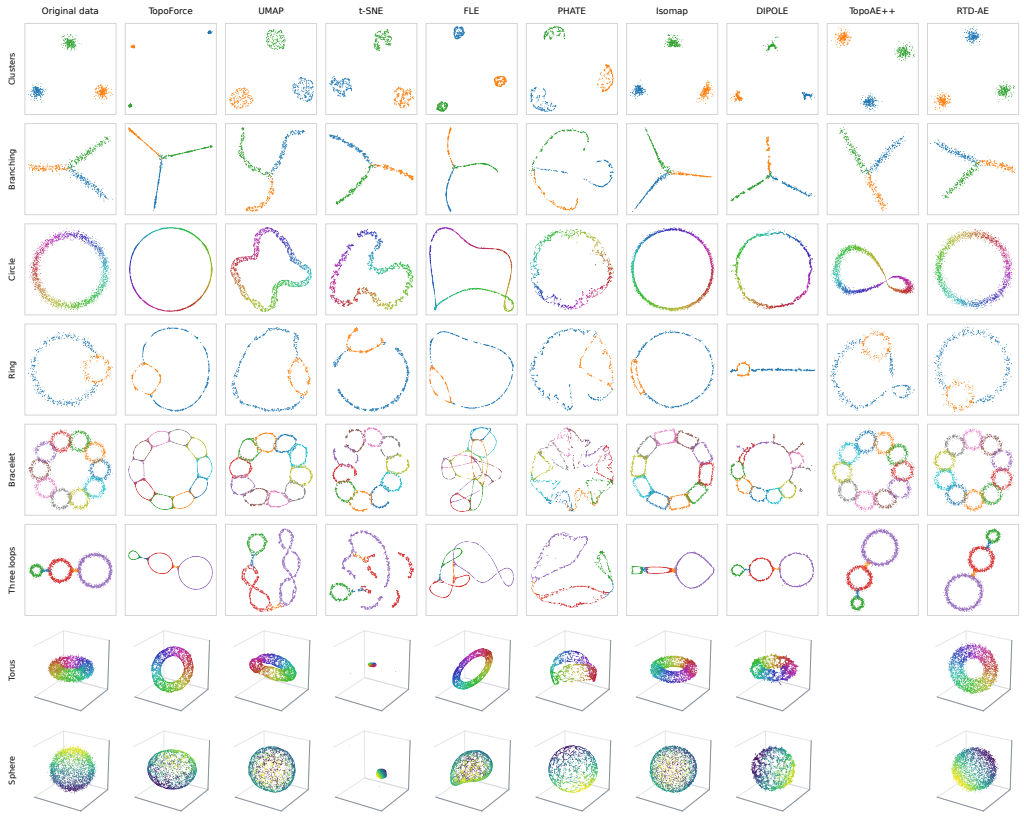

Figure A1: **Complete embedding grids under uniform sampling.** Embeddings of all eight synthetic benchmarks under uniform (U) sampling, extending Figure 3 to the benchmarks not shown in the main text. Colors encode ground-truth position or component identity; torus and sphere rows are three-dimensional. The clusters and branching rows are non-cyclic controls. TopoAE++ is omitted for the torus and sphere.

B High-Dimensional Nuisance Variation

Section 5.4 reports the nuisance-dimension sweep for the three-loop benchmark only. This section gives the full sweep for all three swept benchmarks (circle, three loops, torus): for each, the embeddings of every method at ten noise levels (Figures A6–A8) and the corresponding persistence diagrams (Figures A9–A11).

In each embedding grid, rows are noise levels, from $+0$ to $+2048$ appended Gaussian dimensions (standard deviation 0.6), and the first column is a two-dimensional PCA projection of the noisy input. Recovery deteriorates for all methods as nuisance variation increasingly dominates the observed geometry, and rankings vary at intermediate dimensions; TopoForce nevertheless retains a coherent cyclic or toroidal organization into the highest-nuisance regime, where several alternatives lose most or all target contribution. The persistence-diagram grids distinguish this genuine target recovery from embeddings that only visually resemble loops.

C Implementation Details

Graph and diffusion. TopoForce constructs the neighborhood graph exactly as in Section 4.1. Let $d_i^{(1)}$ be the nearest-neighbor distance of point i and $r = \max_i d_i^{(1)}$. We retain at most $k = 15$ neighbors per point whose distance does not exceed r, and symmetrize the resulting edge set. On this fixed support, affinities are $K_{ij} = \exp(-d_{ij}^2/h^2)$ with bandwidth $h = 0.3$ in all reported experiments; the automatic alternative uses the median distance to the tenth non-self neighbor. We

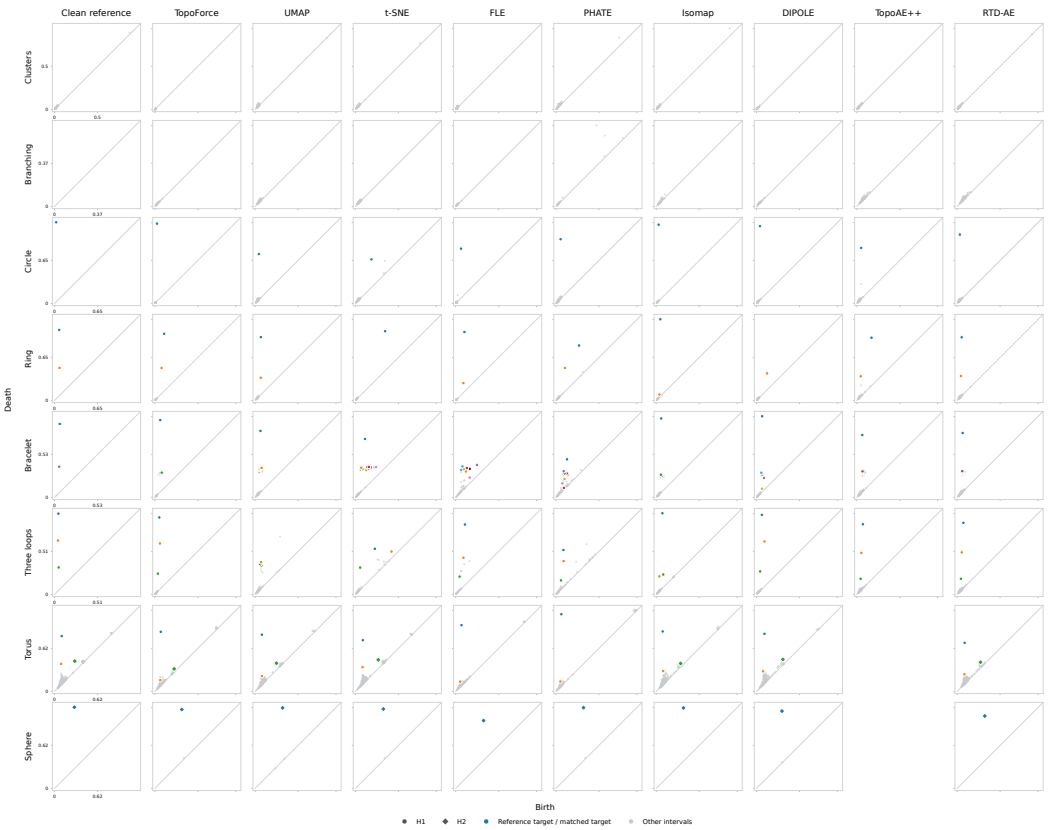

Figure A2: **Persistence diagnostics under uniform sampling.** Birth–death diagrams of the embeddings in Figure A1, computed on shared point subsets with the same protocol as the topology recovery score. Marker shape distinguishes H_1 from H_2 intervals; intervals matched to reference targets are colored by the reference feature they match (first column), while unmatched intervals are shown in light gray. On the cyclic benchmarks the matched intervals stay well separated from the diagonal cloud for TopoForce, whereas on the clusters and branching controls its intervals concentrate along the diagonal.

then apply $\alpha = 1$ diffusion normalization, $\widetilde{K}_{ij} = K_{ij}/(q_i q_j)$ with $q_i = \sum_j K_{ij}$, followed by row normalization to obtain the random-walk operator P.

For the density factor used in Eq. 6, we separately compute Gaussian row sums before neighborhood pruning, including self-affinities, and normalize them to unit mean to obtain ρ_i. Following PHATE, the reference time t_{VNE} is taken at the knee of the Von Neumann entropy curve of the diffusion spectrum, located by a two-segment straight-line fit over $t \in \{1, \ldots, 99\}$ using 30 spectral components.

Optimization. Pairwise scales S_{ij} are converted to repulsion exponents by a row-rank normalization followed by symmetrization, and the density correction $1/\sqrt{\rho_i \rho_j}$ multiplies the repulsion term only. The crossover schedule pools the within-component off-diagonal values of S, takes 16 equally spaced quantile levels, anchors the first at $\kappa = 1$ and the last at $\kappa_{\max} = \max_{i<j} S_{ij}$, and drops duplicate levels; cross-component pairs do not enter the schedule, so the realized number of stages can be smaller than 16. Each stage inherits the previous coordinates, uses the median current graph-edge length as the length scale L, and stops at the first plateau or at the iteration cap. Force constants are $c_{\mathrm{attr}} = 1$, $c_{\mathrm{rep}} = 16$ and $c_{\mathrm{grav}} = 1$, and the per-stage cap is 1500 iterations; all of these values are held fixed across experiments. The plateau rule averages per-node displacement over blocks of 25 steps after a warm-up of 50, sets its reference scale from the first 4 block means, and stops a stage once 4 consecutive blocks improve by less than 0.15 of that scale, never before step 200. Initialization is seeded, and all reported runs use seed 0. The released reference implementation realizes the full 16 stages on the 1000-point bracelet example and completes it in 127 s with eight CPU threads;

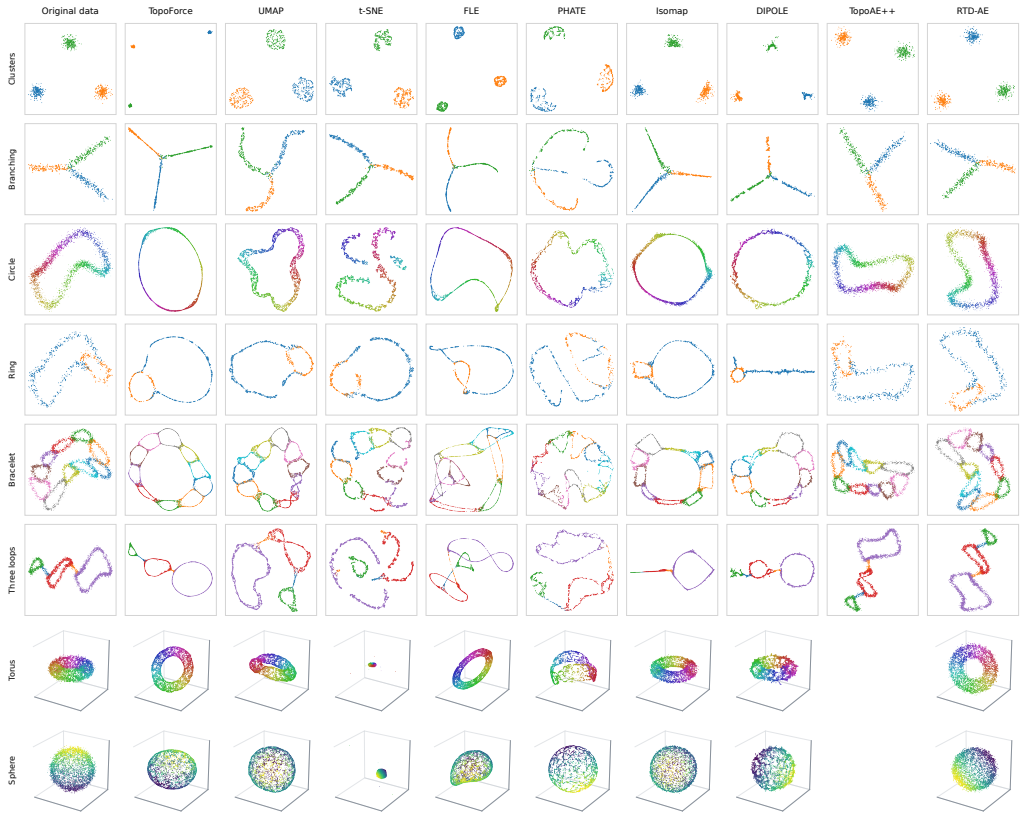

Figure A3: **Complete embedding grids under geometric distortion.** The benchmarks of Figure A1 under distorted uniform (UD) sampling for the four planar cyclic benchmarks; the remaining rows are unchanged because distortion is applied only to planar cyclic coordinates. Several methods that resolve the uniform circle, ring, bracelet, and three-loop benchmarks lose or distort these cycles under distortion, while TopoForce retains both the local cycles and their larger-scale organization.

benchmark runs in this paper use two threads on an AMD EPYC 75F3, which is the setting behind the timings of Table 3.

Evaluation. Topology recovery scores use median-normalized pairwise distances and Vietoris–Rips barcodes on shared subsets of up to 1000 points for H_1 and 256 for H_2, with cycle matching by image persistence in a common filtration. The nuisance sweep appends Gaussian dimensions of standard deviation 0.6 to the same 1000 observations. All reported numbers are single-run descriptive scores, not significance claims.

Baseline autoencoders. TopoAE++ and RTD-AE use their default training settings and published objectives; RTD-AE runs use 400 epochs, batch size 64, learning rate 5×10^{-4}, a two-hidden-layer MLP of width 32, latent dimension 2 (3 for the grid-cell data), and seed 42, trained on a single NVIDIA RTX 4090.

D Extended Real-Data Evaluation

Table 2 reports, for the real-data benchmarks, the circular rank correlation used as the primary circular metric together with the loop-dominance score. Table A1 adds, for the full COIL-20 set, the supplementary label-based diagnostics computed from the same saved embeddings: absolute Fisher–Lee circular correlation (phase agreement), angular mean absolute error under the best global orientation, neighbor recall at 2, object purity at 5, and the per-object H_1 loop-strength gap. These are in-sample, shape-sensitive diagnostics: as noted in the main text, centroid-based angular agreement depends on loop shape and sampling and does not by itself establish topology.

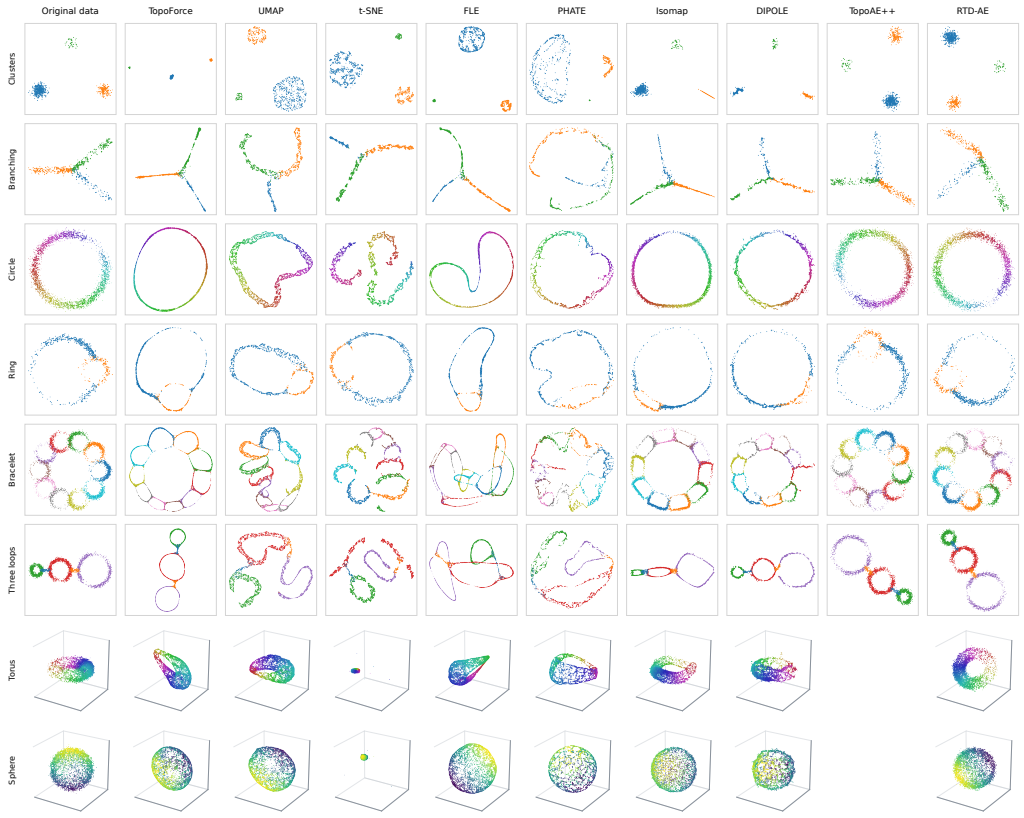

Figure A4: **Complete embedding grids under nonuniform sampling.** Embeddings of all eight benchmarks under nonuniform (NU) sampling, extending Figure 4 to the benchmarks not shown in the main text. Density heterogeneity degrades methods unevenly across benchmarks rather than producing a uniform loss of topology; TopoAE++ is omitted for the torus and sphere.

TopoForce leads the circular rank correlation, the phase agreement, and the loop-strength gap, and its angular error is the lowest among all methods, indicating that the recovered circular order is not an artifact of the primary metric choice. The two columns on which TopoForce does not lead are local rather than circular: t-SNE attains the highest neighbor recall and object purity, yet its circular rank correlation remains below TopoForce's, i.e. t-SNE keeps objects and neighboring views cohesive without arranging them in viewing order. Read together, the columns support the main-text observation that a visibly cohesive embedding need not preserve the rotational order of the underlying manifold.

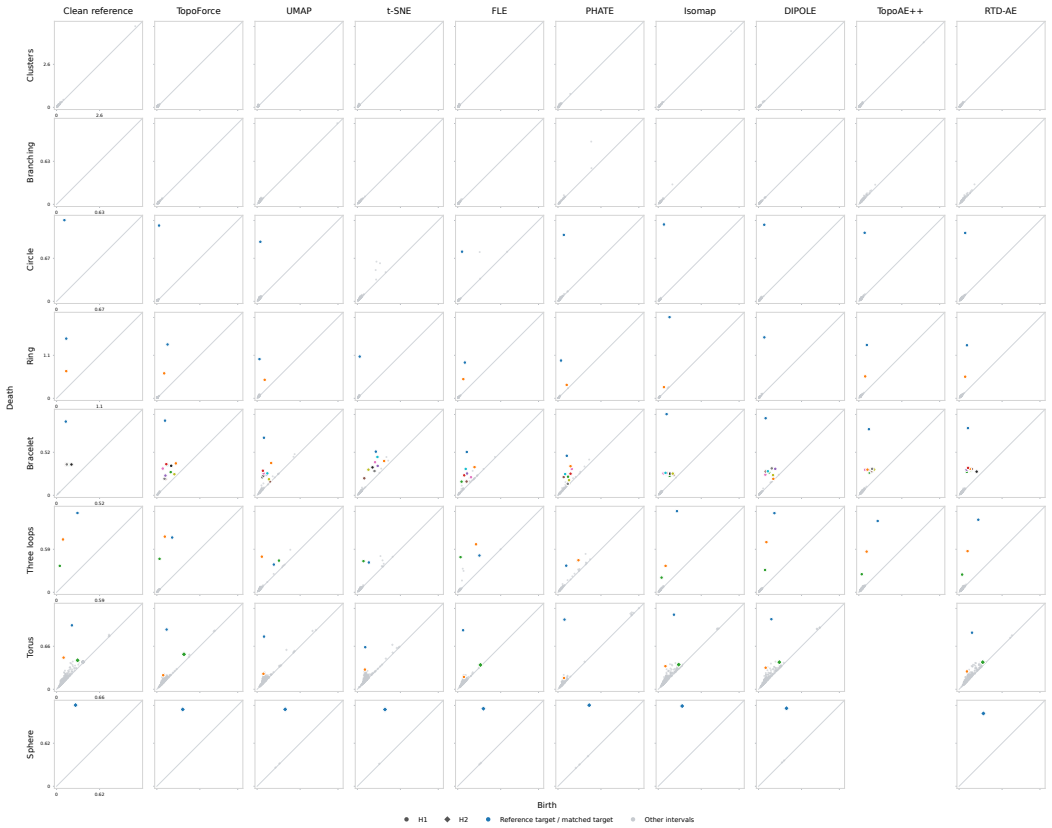

Figure A5: **Persistence diagnostics under nonuniform sampling.** Birth–death diagrams of the embeddings in Figure A4, with the same encoding as Figure A2. Matched target intervals remain separated from the diagonal for TopoForce on the circle, ring, and three-loop benchmarks despite strong density heterogeneity, while the torus separates recovery across homology dimensions.

Method	Circ. rank corr.	Phase agree.	Ang. MAE (°)	Nbr. recall@2	Purity@5	H_1 gap
TopoForce	**0.488**	**0.589**	**42.70**	0.625	**0.935**	**0.313**
UMAP	0.307	0.441	56.23	0.422	0.859	0.162
t-SNE	0.374	0.495	51.42	**0.700**	0.959	0.164
FLE	0.249	0.396	60.84	0.482	0.824	0.056
PHATE	0.132	0.368	62.90	0.227	0.867	0.020
Isomap	0.183	0.339	64.98	0.357	0.681	0.048
DIPOLE	0.123	0.386	61.57	0.209	0.592	0.010
TopoAE++	0.185	0.394	61.06	0.473	0.581	0.033
RTD-AE	0.201	0.388	60.78	0.500	0.771	0.127

Table A1: Extended label-based diagnostics on the full COIL-20 set (20 objects, 72 views each; per-object scores macro-averaged), computed from the same saved embeddings as Table 2. Circ. rank corr.: absolute pairwise-sine correlation on rank angles; Phase agree.: absolute Fisher–Lee circular correlation; Ang. MAE: angular mean absolute error under the best global orientation (lower is better); Nbr. recall@2: fraction of the two nearest embedding neighbors that are the true $\pm 5^\circ$ views; Purity@5: fraction of the five nearest neighbors in the full embedding belonging to the same object; H_1 gap: normalized largest-minus-second H_1 persistence. Red/orange: best/second-best distinct scores per column.

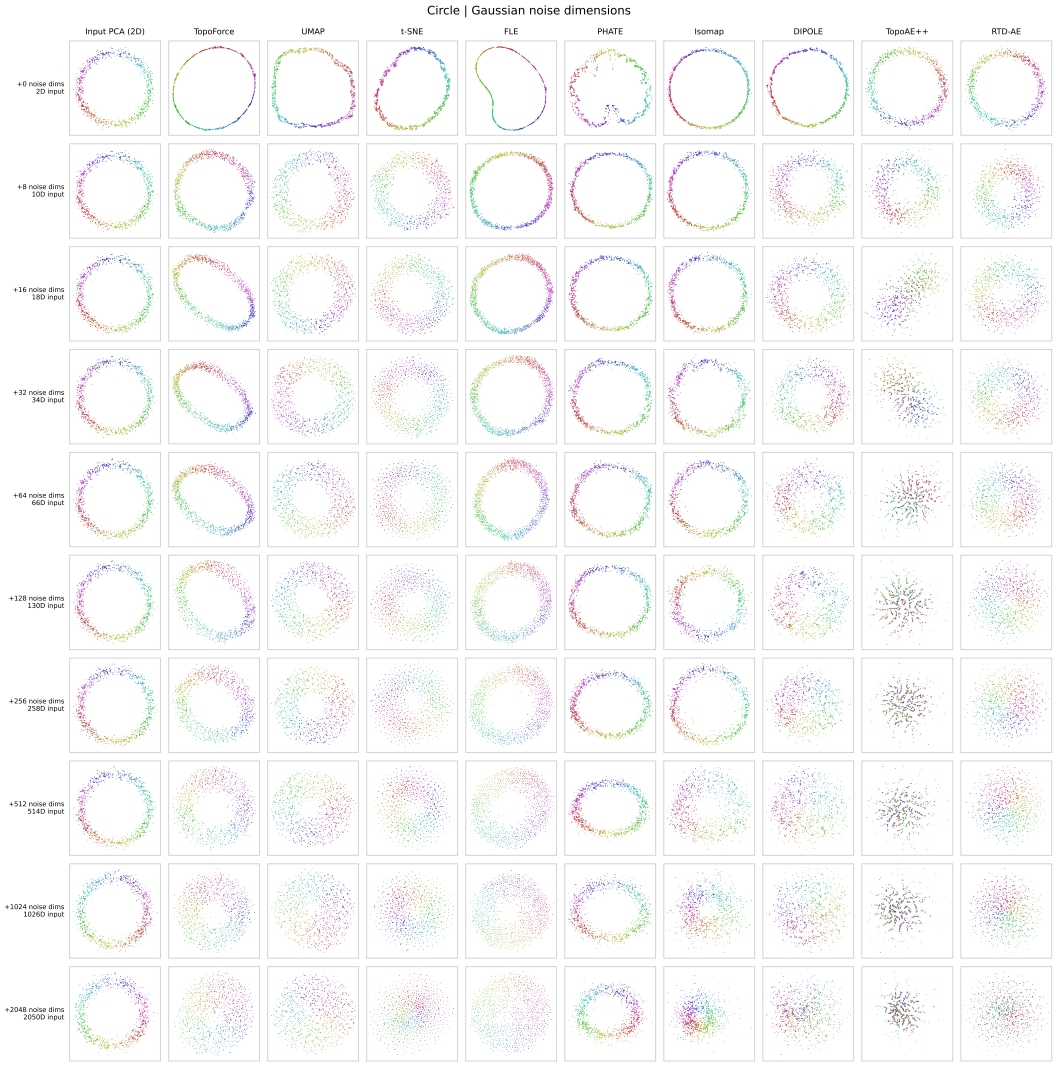

Figure A6: **Circle embeddings under increasing Gaussian nuisance dimensions.** Two-dimensional PCA projections of the noisy input (first column) and the embeddings of all methods (remaining columns) as the number of appended Gaussian nuisance dimensions increases from 0 to 2048. The input projection loses its circular organization early, whereas TopoForce retains a closed loop into the highest-nuisance regime; recovery nevertheless degrades for all methods as nuisance variation dominates the observed geometry.

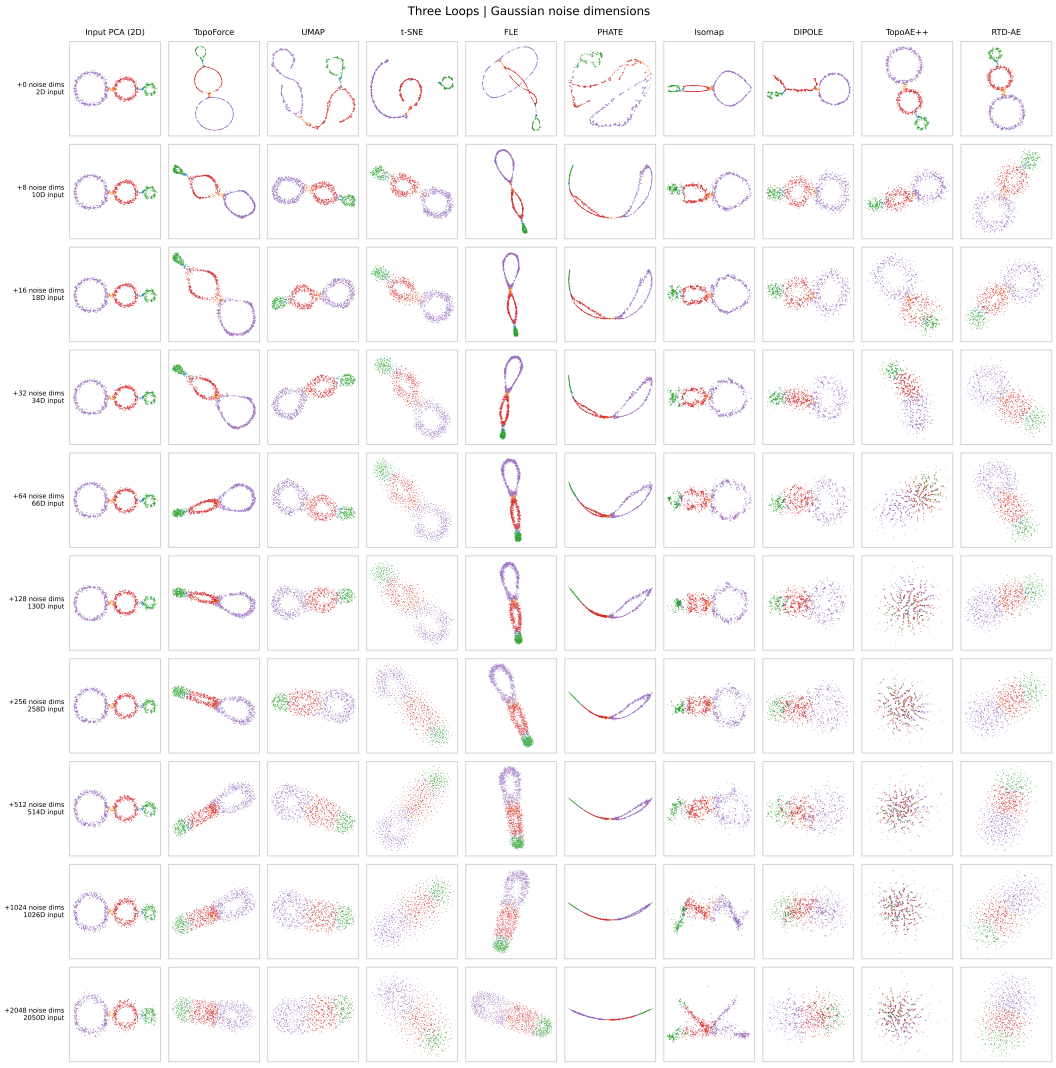

Figure A7: **Three-loop embeddings under increasing Gaussian nuisance dimensions.** The same sweep as Figure A6 on the three unequal loops, with colors encoding loop identity. The three loops remain separated for TopoForce across all noise levels shown, while several methods collapse to starburst-like clouds once the nuisance dimensions dominate; this benchmark is the one summarized quantitatively in Figure 5.



Figure A8: **Torus embeddings under increasing Gaussian nuisance dimensions.** The same sweep as Figure A6 on the torus, with three-dimensional input and embedding panels. TopoAE++ is omitted because its native 3D runs did not complete within the runtime budget. Toroidal organization is harder to read off visually than cyclic organization, which is why the persistence diagrams in Figure A11 are needed to separate genuine recovery from apparent structure.

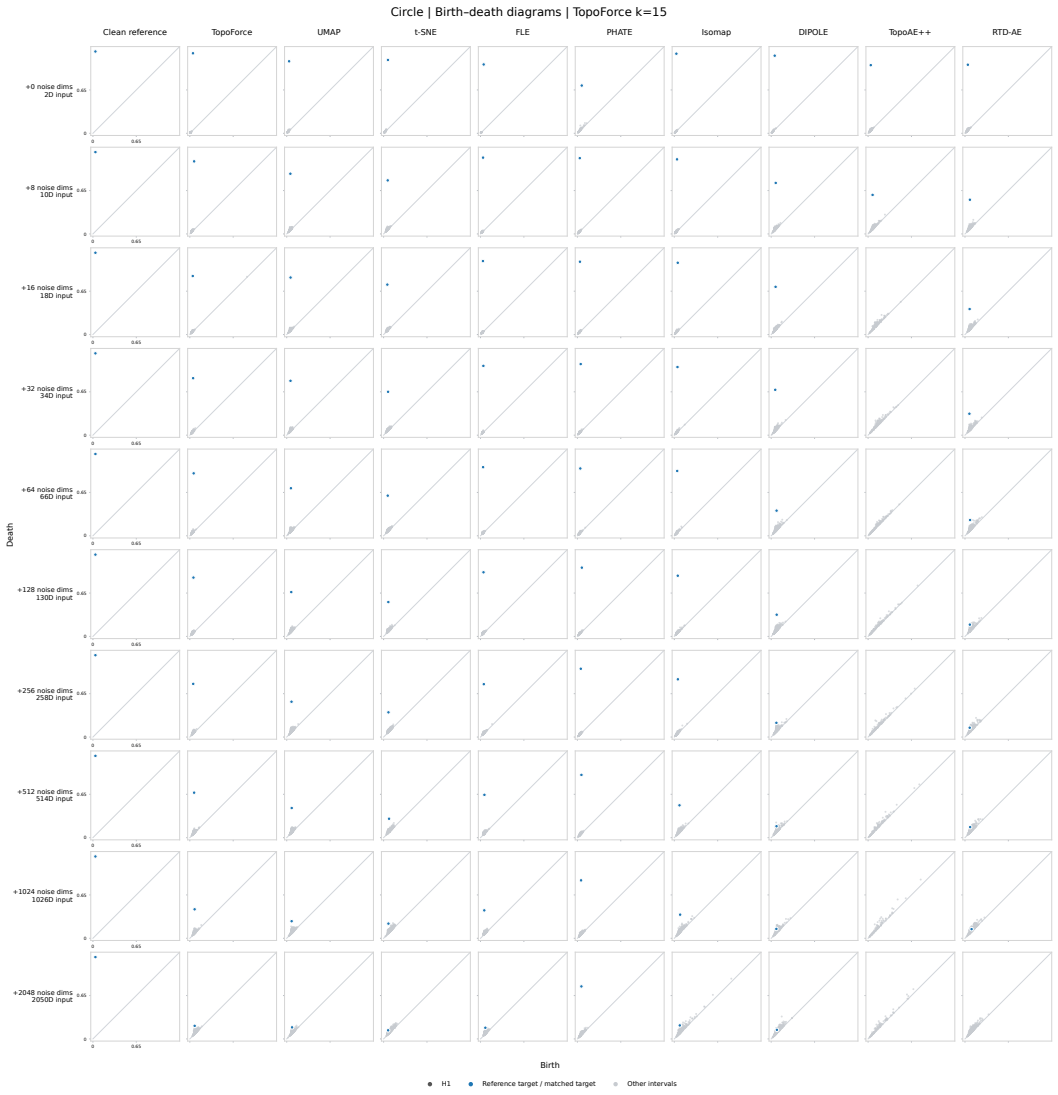

Figure A9: **Circle persistence diagrams across the nuisance sweep.** Birth–death diagrams of the embeddings in Figure A6, one row per noise level and one column per method, with the clean reference in the first column. Matched target intervals are colored and unmatched intervals are light gray; TopoForce keeps the target interval separated from the diagonal at higher noise levels than most alternatives, although the separation narrows as noise increases.

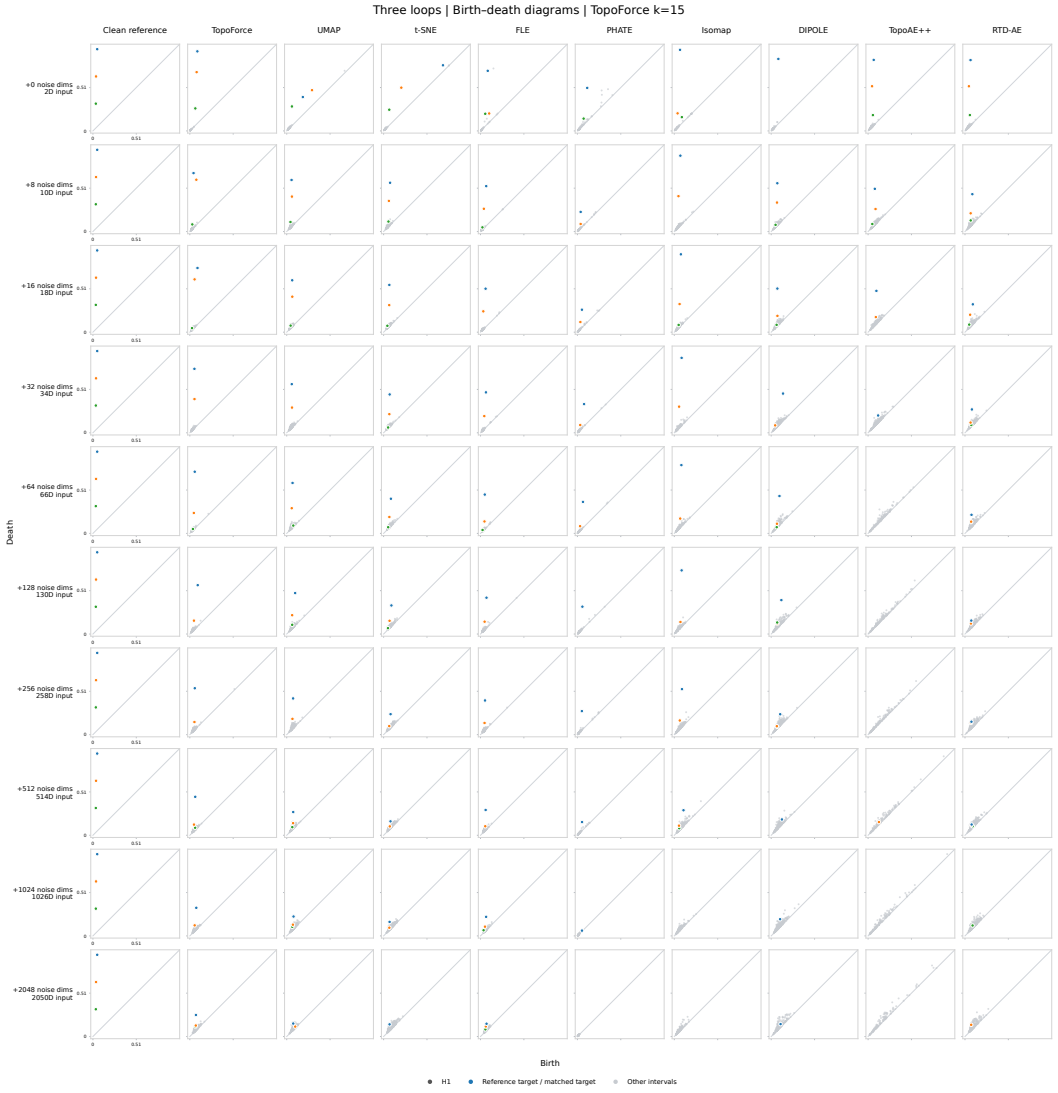

Figure A10: **Three-loop persistence diagrams across the nuisance sweep.** Birth–death diagrams of the embeddings in Figure A7, with the same layout and encoding as Figure A9 and one colored interval per reference loop. The diagrams confirm that TopoForce's visually apparent loops at high noise correspond to matched target intervals rather than to noise-induced persistence.

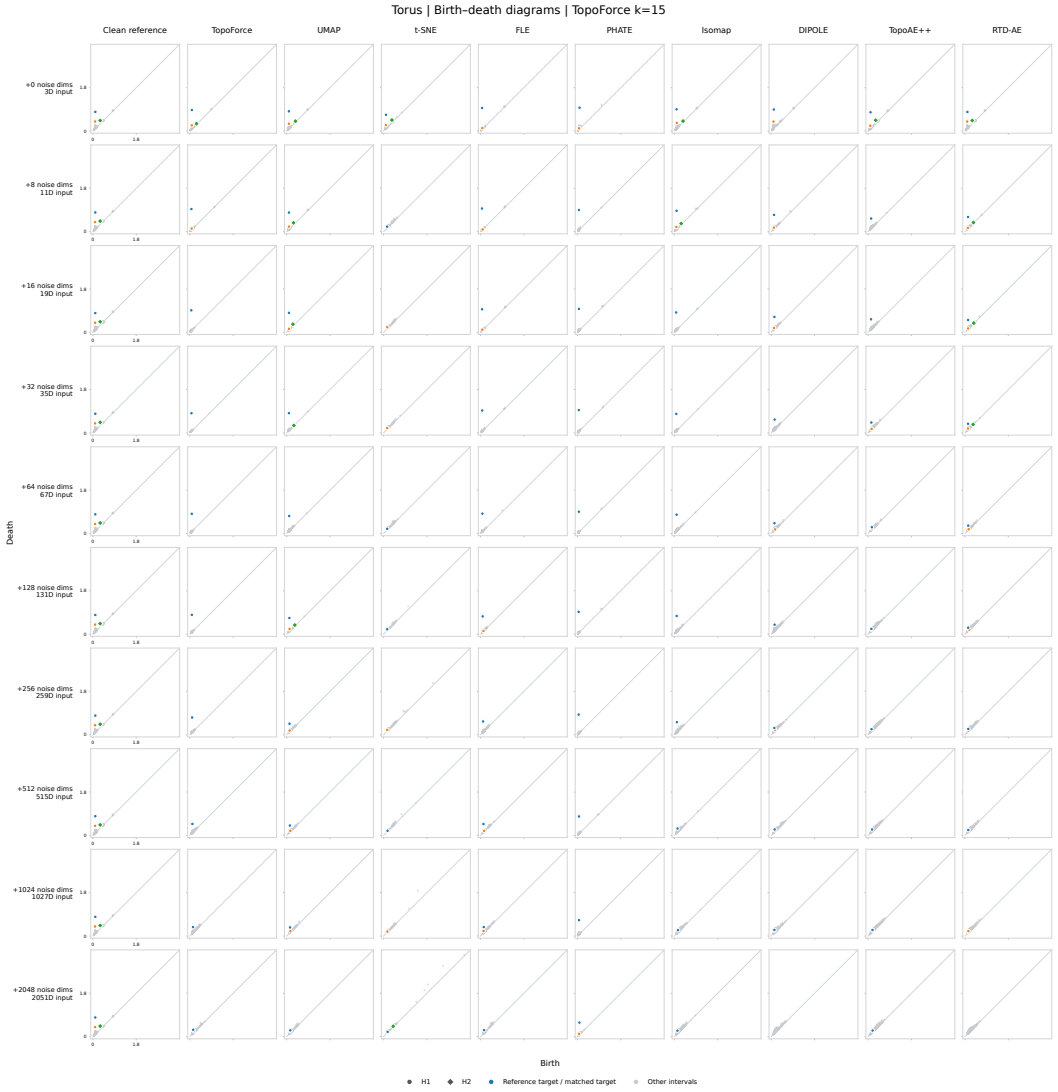

Figure A11: **Torus persistence diagrams across the nuisance sweep.** Birth–death diagrams of the embeddings in Figure A8, with the same layout as Figure A9 and marker shape additionally distinguishing H_1 from H_2 intervals. The torus is the benchmark on which recovery is weakest overall, and the diagrams make explicit which homology dimension each method retains as noise increases.