

TOWARDS SCALABLE PERSISTENCE-BASED TOPOLOGICAL OPTIMIZATION

Abderrahim Bendahi *

École Polytechnique
Paris, France

abderrahim.bendahi@polytechnique.edu

Alexandre Duplessis *

ENS Ulm, PSL
Paris, France

alexandre.duplessis@ens.psl.eu

Arnaud Fickinger

UC Berkeley
Berkeley, CA, USA

arnaud.fickinger@berkeley.edu

ABSTRACT

Persistence-based topological optimization deforms a point cloud $X \subset \mathbb{R}^d$ by minimizing objectives of the form $L(X) = \ell(\text{Dgm}(X))$, where $\text{Dgm}(X)$ is a persistence diagram. In practice, optimization is limited by two coupled issues: persistent homology is typically computed on subsamples, and the resulting topological gradients are highly sparse, with only a few *anchor* points receiving nonzero updates. Motivated by diffeomorphic interpolation, which extends sparse gradients to smooth ambient vector fields via Reproducing Kernel Hilbert Space (RKHS) interpolation, we propose a more scalable pipeline that improves both subsampling and gradient extension. We introduce subsampling via random slicing, a lightweight scheme that promotes iteration-wise geometric coverage and mitigates density bias. We further replace the costly kernel solve with a fast *Nadaraya-Watson* (NW) Gaussian convolution, producing a globally defined smooth update field at a fraction of the computational cost, while being more suited for topological optimization tasks. We provide theoretical guarantees for NW smoothing, including anchor approximation bounds and global Lipschitz estimates. Experiments in 2D and 3D show that combining random slicing with NW smoothing yields consistent speedups and improved objective values over other baselines on common persistence losses.

1 INTRODUCTION

Persistent homology (PH) provides a multiscale summary of the topology of a dataset by tracking the birth and death of homological features along a filtration, producing a persistence diagram $\text{Dgm}(X)$, and enjoys well-quantified stability properties in bottleneck and Wasserstein distances as established in Chazal et al. (2014); Skraba & Turner (2020). This representation has proved expressive and robust enough to serve as a building block for learning pipelines, where topological summaries enter either as features or as loss terms. A particularly direct use is *persistence-based topological optimization*, where one updates the data itself by minimizing

$$\min_{X=\{x_i\}_{i=1}^n \subset \mathbb{R}^d} L(X) := \ell(\text{Dgm}(X)), \quad (1)$$

with losses ℓ encoding tasks such as simplification (remove features), augmentation (create features), or registration to a target diagram, in the spirit of the topological loss design used in Liu et al. (2022).

Despite favorable stability and almost-everywhere differentiability properties, the practical performance of equation 1 is often limited by *gradient sparsity*: at a given iterate, only a small set of points participating in critical birth and death simplices receive nonzero gradients, while most points have zero update, a phenomenon documented for point-cloud optimization in Carrière et al. (2024).

*Equal contribution.

A recent remedy is *diffeomorphic interpolation*, introduced in Carrière et al. (2024), which constructs a smooth ambient vector field matching the sparse topological gradient on a set of anchor points and then transports the full point set along the induced flow. This approach turns sparse gradients into global deformations with controlled Lipschitz constants and combines well with subsampling, since the diffeomorphism learned on a subsample can be applied to the entire point cloud. However, in its RKHS formulation, diffeomorphic interpolation requires solving a kernel linear system at each iteration, which can dominate runtime in large-scale settings, and when combined with uniform subsampling—a common choice for scalability—it inherits the high-variance behavior and density bias of the anchors,

Our Contributions.

- We propose random slicing to improve subsampling geometric coverage.
- We replace the kernel solve by NW Gaussian smoothing to smooth the gradient.
- We provide theoretical guarantees and validate speed and quality gains on 2D and 3D tasks.

2 RELATED WORKS

Persistence diagrams: algorithms and stability. The practical use of persistence diagrams relies on their robustness: small perturbations of the input induce small changes in the diagram, formalized in Cohen-Steiner et al. (2007). On the algorithmic side, Zomorodian & Carlsson (2005); Edelsbrunner et al. (2002) established the algebraic and computational foundations that underlie most modern PH pipelines.

Scaling persistent homology. Exact Vietoris–Rips PH becomes prohibitive at scale; efficient implementations such as Bauer (2021) reduce memory and time constants dramatically, while approximation schemes such as Sheehy (2013) justify sparsifying the filtration itself. These developments motivate subsampling and approximation inside optimization loops, but they also highlight that *how* we subsample can materially affect the optimization trajectory.

Topological losses in machine learning (ML). Topological losses have been deployed broadly in ML, for instance as regularizers that enforce desired topological structure in predictions (Chen et al., 2019; Clough et al., 2020), as constraints for geometry-aware representation learning such as *Topological Autoencoders* Moor et al. (2020), and as objectives in shape and point-cloud optimization (Carrière et al., 2021). Beyond regularization, persistence-based losses have been used for segmentation in medical imaging (Clough et al., 2020) and for uncertainty and shift detection via activation-graph topology (Lacombe et al., 2021). In contrast, we focus on directly optimizing point clouds under persistence-based losses, where gradients are typically sparse.

From sparse topological gradients to smooth ambient updates. A recent approach to address sparsity is *Diffeomorphic interpolation* (Carrière et al., 2024), which constructs a minimum-norm RKHS vector field matching the sparse anchor gradients and updates points through the induced diffeomorphic flow. Our work follows the same high-level principle (extend sparse gradients to ambient smooth vector fields) but targets scalability: we replace exact kernel interpolation (requiring repeated kernel linear solves) with a fast normalized Gaussian convolution, i.e. a Nadaraya-Watson type estimator originating in *On Estimating Regression* (Nadaraya, 1964) and *Smooth Regression Analysis* (Watson, 1964).

3 BACKGROUND

3.1 POINT CLOUDS, FILTRATIONS, AND PERSISTENCE DIAGRAMS

Let $X = \{x_i\}_{i=1}^n \subset \mathbb{R}^d$ be a point cloud. We consider the Vietoris–Rips filtration $\{K_t(X)\}_{t \geq 0}$ defined by

$$K_t(X) := \left\{ \sigma \subseteq [n] : \|x_p - x_q\|_2 \leq t, \quad \forall p, q \in \sigma \right\}.$$

For a fixed homological dimension $k \geq 0$, persistent homology yields a multiset of pairs

$$\text{Dgm}_k(X) = \{(b_\alpha, d_\alpha)\}_{\alpha \in \mathcal{I}_k}, \quad 0 \leq b_\alpha \leq d_\alpha \leq +\infty,$$

where (b_α, d_α) records the birth and death scales of a k -dimensional topological feature. We denote the persistence (distance to the diagonal) by $\text{pers}_\alpha := d_\alpha - b_\alpha$.

3.2 PERSISTENCE-BASED OBJECTIVES

We study objectives that factor through persistence diagrams:

$$\min_{X \subset \mathbb{R}^d} L(X) \quad \text{where} \quad L(X) := \ell(\text{Dgm}_k(X)), \quad (2)$$

for some loss ℓ defined on persistence diagrams. Common examples include:

$$(\text{simplification}) \quad \ell(\text{Dgm}) = \sum_{\alpha \in \tilde{\mathcal{I}}} \text{pers}_\alpha^2, \quad (\text{augmentation}) \quad \ell(\text{Dgm}) = - \sum_{\alpha \in \tilde{\mathcal{I}}} \text{pers}_\alpha^2,$$

and *diagram registration* losses $\ell(\text{Dgm}) = W_p(\text{Dgm}, \text{Dgm}_{\text{target}})$. In practice, $\tilde{\mathcal{I}}$ can select the k most persistent points or a task-specific subset.

3.3 SPARSITY OF TOPOLOGICAL GRADIENTS

The map $X \mapsto \text{Dgm}_k(X)$ is stable (Cohen-Steiner et al., 2007; 2010; Skraba & Turner, 2020), and for standard choices of ℓ , the objective L is Lipschitz and thus differentiable almost everywhere (*Rademacher theorem*, Morgan (2016)). At points of differentiability, the chain rule takes the schematic form

$$\frac{\partial L}{\partial x_i}(X) = \sum_{\alpha: x_i \rightsquigarrow (b_\alpha, d_\alpha)} \left(\frac{\partial \ell}{\partial b_\alpha} \frac{\partial b_\alpha}{\partial x_i} + \frac{\partial \ell}{\partial d_\alpha} \frac{\partial d_\alpha}{\partial x_i} \right), \quad (3)$$

where the sum runs over diagram points whose birth and death critical simplices involve x_i . For most indices $i \in [n]$, no participation occurs in equation 3, hence $\nabla L(X)_i = 0$. We call *anchors* the indices with nonzero gradient:

$$I(X) := \{i \in [n] : \nabla L(X)_i \neq 0\}.$$

Gradient sparsity slows optimization and amplifies variance when PH is computed on subsamples.

3.4 DIFFEOMORPHIC INTERPOLATION AS AMBIENT SMOOTHING

A principled way to mitigate sparsity is to extend anchor gradients to a smooth ambient vector field $v : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and update all points using this field. In particular, diffeomorphic interpolation constructs the minimum-norm RKHS vector field satisfying the interpolation constraints

$$v(x_i) = g_i \quad \text{for all } i \in I, \quad \text{where } g_i := \nabla L(X)_i, \quad (4)$$

by solving

$$v^* \in \arg \min_{v \in \mathcal{H}} \|v\|_{\mathcal{H}} \quad \text{under the constraint 4.} \quad (5)$$

With a Gaussian kernel $K(x, y) = \exp(-\|x - y\|^2 / (2\sigma^2)) I_d$, the *representer theorem* yields

$$v^*(x) = \sum_{i \in I} K(x, x_i) a_i, \quad \text{with coefficients } a = (a_i)_{i \in I} \text{ solving } \mathbb{K}a = g,$$

where $\mathbb{K} = (K(x_i, x_j))_{i, j \in I}$. This produces a smooth field that matches anchors exactly, but requires solving a kernel system (typically cubic in $|I|$) at every iteration.

4 TOWARDS SCALABLE DIFFEOMORPHIC TOPOLOGICAL OPTIMIZATION

4.1 SUBSAMPLING VIA RANDOM SLICING

Motivation and First Guarantees. Computing persistent homology on the full cloud is often prohibitive, so we rely on subsampling. However, in an optimization loop, the subsampling distribution directly affects (i) the variance of the induced anchor gradients and (ii) the geometric regions that can influence topological events. Uniform sampling (used in Carrière et al. (2024)) is cheap but tends to oversample dense regions and miss sparse boundary areas (Figure 1) where long-range simplices often form. We propose a simple structured sampler that enforces global coverage *without* requiring clustering, farthest-point sampling, or additional geometry preprocessing.

Our design is inspired by *sliced* methods in optimal transport (Rabin et al., 2012; Bonneel et al., 2015), which approximate high-dimensional geometry by repeatedly probing random one-dimensional projections where computations reduce to sorting in \mathbb{R} . We adopt the same principle for subsampling: each iteration uses one random *slice* to construct a globally spread subset.

More formally, let $X = \{x_i\}_{i=1}^n \subset \mathbb{R}^d$ and fix a target subsample size $s \in \{2, \dots, n\}$. Draw a random direction $u \sim \mathcal{N}(0, I_d)$ and normalize $u \leftarrow u/\|u\|_2$. Compute the 1D projections

$$t_i := \langle x_i, u \rangle, \quad i \in [n],$$

and let π be a permutation that sorts these values increasingly: $t_{\pi(0)} \leq t_{\pi(1)} \leq \dots \leq t_{\pi(n-1)}$. We then select indices at evenly spaced ranks:

$$S(u) := \left\{ \pi(r_k) : r_k = \lfloor k(n-1)/(s-1) \rfloor, k = 0, \dots, s-1 \right\}. \quad (6)$$

The resulting subset $S(u)$ is used for PH computation at that iteration.

Intuitively, in one dimension, selecting evenly spaced ranks enforces a strong notion of coverage. Thus, for a fixed slice, the sampler avoids the density bias of uniform sampling: extremely dense regions do not *steal* all sample budget, since each interval of the sorted list receives roughly the same number of points.

This notion of coverage can be formalized by the next lemma.

Lemma 4.1 (Maximal rank gap). *Let $r_k = \lfloor k(n-1)/(s-1) \rfloor$ be the selected ranks.*

$$\max_{k \in \{0, \dots, s-2\}} (r_{k+1} - r_k) \leq \left\lceil \frac{n-1}{s-1} \right\rceil.$$

Lemma 4.1 implies that, in the sorted 1D order induced by u , no large contiguous block of ranks is skipped: every rank is within $\mathcal{O}(n/s)$ positions of a selected one. This does not guarantee Euclidean covering in \mathbb{R}^d , but it provides an efficient, iteration-wise notion of global coverage aligned with the slice geometry that empirically stabilizes PH-based optimization. The proof is deferred to Appendix B. Computationally, the sampler requires one projection pass ($\mathcal{O}(nd)$) and a sort ($\mathcal{O}(n \log n)$), typically negligible compared to repeated PH computations.

Geometric Intuition. Figure 1 illustrates how projection-stratified (random slicing) subsampling mitigates the density bias of uniform sampling. The cloud is a mixture of two separated clusters with highly imbalanced masses (85% vs. 15%). Uniform subsampling concentrates almost entirely on the dominant cluster, so the minority cluster may be represented by very few points for small budgets. Random slicing instead selects evenly spaced ranks in a random 1D projection; the separation between clusters induces a gap in projected values, so rank stratification tends to allocate points on both sides of the gap, consistently capturing the low-mass cluster. This iteration-wise notion of global coverage is a simple mechanism to reduce sampling bias in PH computations.

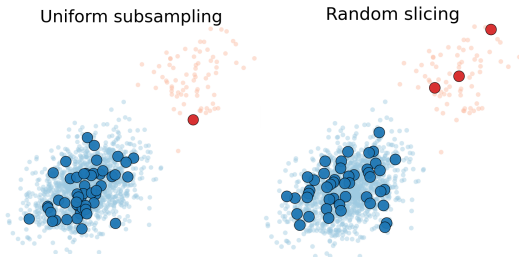


Figure 1: Uniform subsampling vs. random slicing on an imbalanced two-cluster cloud. Light points: full point set; highlighted points: a subsample ($s = 50$).

4.2 NW-CONVOLUTION SMOOTHING

Exact RKHS interpolation constructs the minimum-norm field matching anchors $v(x_i) = g_i$, but requires solving a kernel system at every step, which can dominate runtime when repeated in an optimization loop. We propose a fast alternative that preserves smoothness, yields explicit regularity control, and empirically induces update fields that are often better aligned with the geometric needs of persistence-driven optimization.

Definition 4.2 (Gaussian weights and NW field). Fix a bandwidth $\sigma > 0$ and define the unnormalized Gaussian affinities

$$\phi_i(x) := \exp\left(-\frac{\|x - x_i\|_2^2}{2\sigma^2}\right), \quad i \in I.$$

Let $Z(x) := \sum_{j \in I} \phi_j(x)$ and define normalized weights

$$w_i(x) := \frac{\phi_i(x)}{Z(x)}, \quad \sum_{i \in I} w_i(x) = 1, \quad w_i(x) \in (0, 1). \quad (7)$$

The *Nadaraya-Watson (NW) convolution field* is

$$\bar{v}(x) := \sum_{i \in I} w_i(x) g_i \in \mathbb{R}^d. \quad (8)$$

First, we state the following characterization of the NW convolution field.

Lemma 4.3 (Local kernel-regression characterization). *For each $x \in \mathbb{R}^d$, $\bar{v}(x)$ is the unique minimizer of the strictly convex problem*

$$\bar{v}(x) = \arg \min_{v \in \mathbb{R}^d} \sum_{i \in I} \phi_i(x) \|v - g_i\|_2^2. \quad (9)$$

Lemma 4.3 shows that NW is the *best local constant fit* to anchor gradients under Gaussian weights. This makes the role of σ explicit: it controls the locality of the regression and hence the trade-off between (i) propagation of sparse information to non-anchors and (ii) fidelity to anchor updates. The proof is deferred to Section B.

Note that for every $x \in \mathbb{R}^d$, $\bar{v}(x)$ lies in the convex hull of $\{g_i : i \in I\}$, and by convexity of $\|\cdot\|_2$, $\|\bar{v}(x)\|_2 \leq \max_{i \in I} \|g_i\|_2$. Besides, the following lemma holds.

Lemma 4.4 (Stability w.r.t. anchor gradients). *Fix anchor locations $\{x_i\}_{i \in I}$. Consider two anchor-gradient families $\{g_i\}$ and $\{g'_i\}$ and let \bar{v}, \bar{v}' be the corresponding NW fields. Then*

$$\sup_{x \in \mathbb{R}^d} \|\bar{v}(x) - \bar{v}'(x)\|_2 \leq \max_{i \in I} \|g_i - g'_i\|_2.$$

Lemma 4.4 implies that NW smoothing is *non-expansive* in the anchor gradients: noise in sparse gradients is not amplified by the extension step. The proof is deferred to Section B.

In the following, for each anchor $i \in I$, define

$$\rho_i(\sigma) := \sum_{j \in I \setminus \{i\}} \exp\left(-\frac{\|x_i - x_j\|_2^2}{2\sigma^2}\right).$$

The following theorem studies when NW *nearly* interpolates, i.e., when the exact weight at anchors are close to the corresponding gradients.

Theorem 4.5 (Anchor error bound). *Assume $\|g_i\|_2 \leq M$ for all $i \in I$. Then for each $i \in I$,*

$$\|\bar{v}(x_i) - g_i\|_2 \leq 2M \frac{\rho_i(\sigma)}{1 + \rho_i(\sigma)}. \quad (10)$$

Theorem 4.5 formalizes the *fidelity-propagation trade-off*: small σ reduces cross-anchor mixing (small $\rho_i(\sigma)$) and yields near-interpolation, whereas larger σ increases global propagation but mixes anchor directions.

In the following, we provide some regularity properties of \bar{v} . First, notice that \bar{v} is C^∞ on \mathbb{R}^d . Indeed, each ϕ_i is C^∞ and $Z(x) > 0$ for all x since all $\phi_i(x) > 0$. Thus $w_i = \frac{\phi_i}{Z}$ is C^∞ and $\bar{v} = \sum_i w_i g_i$ is C^∞ . Furthermore, we have the following Lipschitz regularity property.

Theorem 4.6 (Global Lipschitz constant of the NW field). *Assume $\|g_i\|_2 \leq M$ for all $i \in I$. Then \bar{v} is globally Lipschitz and*

$$\|D\bar{v}(x)\|_{\text{op}} \leq \frac{\text{diam}(P)}{\sigma^2} M \quad \forall x \in \mathbb{R}^d. \tag{11}$$

In particular, $\|\bar{v}(x) - \bar{v}(y)\|_2 \leq L \|x - y\|_2$ for all x, y with $L := \frac{\text{diam}(P)M}{\sigma^2}$.

The proof of Theorem 4.6 is deferred to Section B. The bound scales as $1/\sigma^2$: larger bandwidth yields smoother (more stable) deformations. The factor $\text{diam}(P)$ captures the geometric spread of anchors: if anchors are confined to a small region, the induced field cannot vary too rapidly.

Geometric intuitions on NW smoothing.

To build intuition for the behavioral gap between exact kernel-based diffeomorphic interpolation and our NW-Convolution smoothing, we consider a deliberately simple anchor configuration (Fig. 2) that isolates the key mechanism: *normalization*. We place anchors in two tight clusters with opposite vectors, and compare the induced ambient fields for a fixed bandwidth $\sigma = 0.35$.




Figure 2: Kernel vs. NW smoothing on the same anchors. Two spatial clusters of anchors carry opposite directions. Left: exact Gaussian-RKHS interpolation (KERNEL). Right: NW smoothing (NW). Colors show the norm of the smoothed vector field.

In the RKHS construction (used by Carrière et al. (2024)), the vector field takes the form $v_{\text{ker}}(x) = \sum_j k_\sigma(x, x_j) \alpha_j$, where the coefficients α are chosen so that $v_{\text{ker}}(x_i) = g_i$ at anchors. Because the expansion uses *unnormalized* Gaussians, the overall magnitude typically *decays away from the anchor set*: when x is far from every x_j , all kernels $k_\sigma(x, x_j)$ become small simultaneously. This induces a field that is strongly localized around the regions that contain anchors (and is visually reflected by the attenuation of $\|v(x)\|$ away from the clusters).

In contrast, NW defines a convex combination of anchor vectors $v_{\text{NW}}(x) = \sum_{i \in I} w_i(x) g_i$ (Definition 4.2), where the normalization enforces $\sum_j w_j(x) = 1$ for all x , so $v_{\text{NW}}(x)$ behaves as a *local barycentric average* of the anchor gradients. Two qualitative consequences follow. First, the field does not vanish simply because x is far from anchors: when all kernels are simultaneously small, the ratio defining $w_j(x)$ remains well-defined and tends to a near-uniform weighting across anchors, so the far-field approaches a (typically nonzero) average direction. Second, regions dominated by one cluster naturally inherit that cluster’s direction, while the transition zone between clusters exhibits a smooth cancellation pattern (the averaging mediates the conflict rather than enforcing exact matching).

In topological optimization, topological gradients are sparse and highly local: only a few points (anchors) receive nonzero updates. Exact interpolation propagates these updates but can produce highly localized magnitude profiles and requires solving a kernel system. NW smoothing instead provides a cheap, stable propagation mechanism whose geometry is easy to interpret: it transports the anchor signal via normalized similarity weights, yielding a globally defined field with controlled mixing across competing anchors.

In the two-cluster setting of Figure 2, a natural desired smoothing is to propagate the influence of the opposing anchor groups across the domain so that points on each side are consistently driven in the direction prescribed by the nearest cluster, rather than receiving a near-zero update away from anchors. This effect is precisely what the normalized NW averaging promotes, and it helps explain its

empirical effectiveness (see Section 5) as a faster surrogate to exact kernel interpolation. Additional configurations illustrating the same phenomena are provided and discussed in Appendix C.

Practical considerations: sensitivity to the bandwidth σ . While NW-Convolution smoothing provides a fast and principled extension of sparse topological gradients, its behavior is strongly controlled by the bandwidth parameter σ , which sets the locality of averaging. To quantify this sensitivity, we run NW-smoothing with diffeomorphic updates on a noisy circle in \mathbb{R}^2 (Gaussian noise added to a unit circle) using the *collapser* loss, i.e. a persistence objective that penalizes large birth and death times and thus encourages rapid topological simplification (see Appendix D.2 for a formal definition of the loss and visualizations for this task). Figure 3 reports the loss trajectories for several fixed values of σ . Beyond the existence of a clearly favorable range (here around $\sigma = 0.1$), the curves exhibit substantial variability in both convergence speed and attained objective value, underscoring that σ is a critical hyper-parameter rather than a minor adjustment. This observation motivates *learning* (or adaptively scheduling) σ during optimization, we provide more details about this point in Section 5.

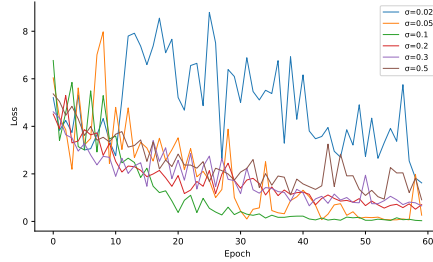


Figure 3: Loss trajectories (collapser loss) on a noisy circle in \mathbb{R}^2 for different fixed σ values, using the same diffeomorphic update scheme.

5 EXPERIMENTAL RESULTS

We evaluate our pipeline *subsampling via random slicing and NW-Convolution smoothing*¹ on a challenging 3D topological augmentation task. We consider the Stanford Bunny point cloud (Turk & Levoy, 1994), i.e. the mesh vertices, yielding an initial point set $X_0 = \{x_i\}_{i=1}^n \subset \mathbb{R}^3$, where $n = 35,947$. We optimize a persistence-based objective that promotes prominent two-dimensional topological features, i.e. we encourage the formation (or strengthening) of a cavity by increasing its persistence.

Loss. We use the *bunny loss*, a topological augmentation objective defined by

$$\mathcal{L}_{\text{bunny}}(X) = - \sum_{\alpha \in \text{Dgm}_2(X)} \left(\frac{1}{2}(d_\alpha - b_\alpha) \right)^2 + \sum_{i=1}^n \max(\|x_i\|_\infty - 1, 0), \quad (12)$$

i.e. a (negative) squared total persistence term in dimension 2, together with a soft box penalty that discourages points from leaving $[-1, 1]^3$.

Learning the bandwidth σ . The bandwidth σ controls the locality of NW-Convolution smoothing and strongly impacts the resulting deformation (Figure 3). Rather than treating σ as a fixed hyper-parameter, we can *learn* it online by backpropagating through the NW weights. Concretely, we parameterize a positive bandwidth as

$$\sigma(\theta) = \sigma_0 \exp(\theta), \quad (13)$$

with a scalar $\theta \in \mathbb{R}$. At each iteration, we first compute the raw sparse PH gradient $g_{\text{raw}}(X)$ (without internal smoothing), build the NW movement field $\bar{v}_\theta(X)$ using equation 13 inside the Gaussian weights, and form a one-step look-ahead point cloud

$$X^+(\theta) = X - \eta_X \bar{v}_\theta(X), \quad (14)$$

where η_X is the point-update step size. We then update θ by minimizing the *look-ahead* loss

$$\theta \leftarrow \theta - \eta_\sigma \nabla_\theta \mathcal{L}(X^+(\theta)), \quad (15)$$

¹In our plots we label our approach as NW-CONVOLUTION, this always refers to the *combination* of NW smoothing *and* random slicing subsampling.




Figure 4: 3D bunny augmentation task. Top: initial point cloud (Stanford Bunny) and point clouds after 400 epochs for vanilla sparse updates (VANILLA), exact kernel-based diffeomorphic interpolation (KERNEL), and our approach (NW-CONVOLUTION, i.e., random slicing subsampling and NW smoothing). Bottom: loss curves versus epochs for the three methods.

Table 1: **Bunny augmentation: best loss and runtime.** Best objective value (lower is better) and average time per epoch, averaged over the full run.

Method	Best Loss ↓	Avg Time/Epoch (s) ↓
Vanilla	-0.00026	2.044
Kernel	-0.00151	19.979
NW-Convolution (ours)	-0.01869	2.573

with η_σ a separate learning rate. In implementation, we stop gradients through X and g_{raw} when differentiating equation 15 (i.e., treat them as constants during the θ update), which yields a stable and inexpensive meta-gradient while still adapting σ to the local optimization regime. This procedure can be interpreted as selecting the bandwidth that makes the *next* deformation step most beneficial for the topological objective, and we use it as a practical alternative to manual tuning when the optimal σ varies substantially across tasks and datasets. Pseudo-codes of the algorithms for learning σ and our NW-Convolution are given in Section A.

Compared methods. We compare: (i) **Vanilla** topological gradient descent, which updates only anchor points (sparse gradient); (ii) **Kernel** diffeomorphic interpolation Carrière et al. (2024), using exact Gaussian-RKHS interpolation; and (iii) **NW-Convolution (ours)**, which replaces the kernel solve with NW smoothing and uses random slicing for subsampling.

Hyper-parameters. For all methods we run $T = 400$ optimization steps with subsampling size $s = 100$ and compute persistent homology in dimension 2. Vanilla and kernel-based diffeomorphic interpolation use SGD with learning rate $\eta_X = 0.1$, for the kernel baseline we set the Gaussian bandwidth to $\sigma = 0.05$. For our approach, we use the same subsampling budget but replace uniform subsampling by projection-stratified sampling and apply the NW-Convolution smoother with base bandwidth $\sigma_0 = 5 \times 10^{-3}$ and step size $\eta_X = 0.18$.

Results. Figure 4 shows the initial bunny and the final point clouds after 400 epochs for the three optimizers, together with the loss evolution. Table 1 reports the best loss achieved over the run and the average runtime per epoch. Our method achieves substantially lower objective values than both baselines while remaining close to the runtime of the vanilla sparse update. In particular, NW smoothing introduces only a small overhead compared to no smoothing, whereas exact kernel interpolation is an order of magnitude slower in this experiment.

Further experiments in 2D can be found in Appendix D.

6 CONCLUSION

We revisited a central practical bottleneck in persistence-based topological optimization: the sparsity and instability of topological gradients when PH is computed on subsamples. Building on the ambient-smoothing viewpoint of diffeomorphic updates, we proposed a lightweight alternative to exact kernel interpolation: a Nadaraya-Watson Gaussian convolution that extends sparse anchor gradients without solving a kernel system, together with a simple projection-stratified (random slicing) subsampling scheme to improve geometric coverage. Across our benchmarks in 2D and 3D, this combination yields substantially faster iterations than kernel-based diffeomorphic interpolation while achieving improved objective values. A key limitation of the present work is that we focus on controlled synthetic and standard geometry benchmarks, we do not yet demonstrate end-to-end impact on downstream real-world tasks. We view our contribution primarily as a step towards making persistence-driven optimization more scalable and, ultimately, more practical in application settings where large point sets and repeated PH computations are unavoidable.

REFERENCES

- Ulrich Bauer. Ripser: efficient computation of Vietoris–Rips persistence barcodes. *Journal of Applied and Computational Topology*, 2021. doi: 10.1007/s41468-021-00071-5.
- Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51(1):22–45, Jan 2015. ISSN 1573-7683. doi: 10.1007/s10851-014-0506-3. URL <https://doi.org/10.1007/s10851-014-0506-3>.
- Mathieu Carrière, Marc Theveneau, and Théo Lacombe. Diffeomorphic interpolation for efficient persistence-based topological optimization, 2024.
- Mathieu Carrière, Marc Theveneau, and Théo Lacombe. Diffeomorphic interpolation for efficient persistence-based topological optimization. In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS '24*, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9798331314385.
- Mathieu Carrière, Frédéric Chazal, Marc Glisse, Yuichi Ike, Hariprasad Kannan, and Yuhei Umeda. Optimizing persistent homology based functions. 01 2021.
- Frédéric Chazal, Marc Glisse, Catherine Labruère, and Bertrand Michel. Convergence rates for persistence diagram estimation in topological data analysis. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14*, pp. I–163–I–171. JMLR.org, 2014.
- Chao Chen, Xiuyan Ni, Qinxun Bai, and Yusu Wang. A topological regularizer for classifiers via persistent homology. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2573–2582. PMLR, 16–18 Apr 2019. URL <https://proceedings.mlr.press/v89/chen19g.html>.
- James R. Clough, Nicholas Byrne, Ilkay Oksuz, Veronika A. Zimmer, Julia A. Schnabel, and Andrew P. King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. doi: 10.1109/TPAMI.2020.3013679.
- David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007. doi: 10.1007/s00454-006-1276-5.
- David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have l_p -stable persistence. *Found. Comput. Math.*, 10(2):127–139, February 2010. ISSN 1615-3375.
- Edelsbrunner, Letscher, and Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, Nov 2002. ISSN 1432-0444. doi: 10.1007/s00454-002-2885-2. URL <https://doi.org/10.1007/s00454-002-2885-2>.
- Théo Lacombe, Yuichi Ike, Mathieu Carrière, Frédéric Chazal, Marc Glisse, and Yuhei Umeda. Topological uncertainty: Monitoring trained neural networks through persistence of activation graphs, 2021. URL <https://arxiv.org/abs/2105.04404>.
- Weiquan Liu, Hanyun Guo, Weini Zhang, Yu Zang, Cheng Wang, and Jonathan Li. Toposeg: Topology-aware segmentation for point clouds. In Lud De Raedt (ed.), *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pp. 1201–1208. International Joint Conferences on Artificial Intelligence Organization, 7 2022. doi: 10.24963/ijcai.2022/168. URL <https://doi.org/10.24963/ijcai.2022/168>. Main Track.
- 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, 2020.

- Frank Morgan. 3 - lipschitz functions and rectifiable sets. In Frank Morgan (ed.), *Geometric Measure Theory (Fifth Edition)*, pp. 25–38. Academic Press, fifth edition edition, 2016. ISBN 978-0-12-804489-6. doi: <https://doi.org/10.1016/B978-0-12-804489-6.50003-3>. URL <https://www.sciencedirect.com/science/article/pii/B9780128044896500033>.
- E. A. Nadaraya. On estimating regression. *Theory of Probability and Its Applications*, 9(1):141–142, 1964. doi: 10.1137/1109020.
- Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernot. Wasserstein barycenter and its application to texture mixing. In Alfred M. Bruckstein, Bart M. ter Haar Romeny, Alexander M. Bronstein, and Michael M. Bronstein (eds.), *Scale Space and Variational Methods in Computer Vision*, pp. 435–446, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-24785-9.
- Donald R. Sheehy. Linear-size approximations to the victoris–rips filtration. *Discrete & Computational Geometry*, 49(4):778–796, 2013. doi: 10.1007/s00454-013-9513-1.
- Primoz Skraba and Katharine Turner. Wasserstein stability for persistence diagrams. *arXiv: Algebraic Topology*, 2020. URL <https://api.semanticscholar.org/CorpusID:220265887>.
- Greg Turk and Marc Levoy. Zippered polygon meshes from range images. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '94*, pp. 311–318, New York, NY, USA, 1994. Association for Computing Machinery. ISBN 0897916670. doi: 10.1145/192161.192241. URL <https://doi.org/10.1145/192161.192241>.
- Geoffrey S. Watson. Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, 26:359–372, 1964.
- Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005. doi: 10.1007/s00454-004-1146-y.

A ALGORITHMS

A.1 PROJECTION-STRATIFIED SUBSAMPLING (RANDOM SLICING)

Algorithm 1 Projection-stratified subsampling (random slicing)

Require: Point cloud $X = \{x_i\}_{i=1}^n \subset \mathbb{R}^d$, subsample size s

- 1: Sample $u \sim \mathcal{N}(0, I_d)$ and normalize $u \leftarrow u/\|u\|_2$
- 2: Compute projections $t_i \leftarrow \langle x_i, u \rangle$ for $i = 1, \dots, n$
- 3: Let π be the permutation sorting (t_i) increasingly
- 4: **for** $k = 0$ **to** $s - 1$ **do**
- 5: $r_k \leftarrow \lfloor k(n-1)/(s-1) \rfloor$ /* evenly spaced ranks */
- 6: Add index $\pi(r_k)$ to S
- 7: **end for**
- 8: **return** S

A.2 LEARNING THE BANDWIDTH σ (ONE-STEP LOOK-AHEAD)

Algorithm 2 LEARNSIGMASTEP: one-step look-ahead update for σ

Require: Current cloud $X \in \mathbb{R}^{n \times d}$, sparse gradient $g \in \mathbb{R}^{n \times d}$, anchors I

Require: Base bandwidth σ_0 , current θ /* so that $\sigma = \sigma_0 \exp(\theta)$ */

Require: Step sizes η_X (for X) and η_σ (for θ), subsample S

- 1: $\sigma \leftarrow \sigma_0 \exp(\theta)$ /* enforce positivity */
- 2: $X^{\text{sg}} \leftarrow \text{stopgrad}(X)$, $g^{\text{sg}} \leftarrow \text{stopgrad}(g)$
- 3: Build $\bar{v}_\theta(\cdot)$ from $(x_i, g_i^{\text{sg}})_{i \in I}$ using NW weights with bandwidth σ
- 4: $\tilde{X} \leftarrow X^{\text{sg}} - \eta_X \bar{v}_\theta(X^{\text{sg}})$ /* look-ahead point cloud */
- 5: $\theta \leftarrow \theta - \eta_\sigma \nabla_\theta \ell(\text{Dgm}(\tilde{X}_S))$ /* meta-gradient step */
- 6: **return** $\sigma_0 \exp(\theta)$

A.3 NW-CONVOLUTION UPDATE WITH RANDOM SLICING (OURS)

Algorithm 3 NW-Convolution smoothing with random slicing (ours)

Require: Initial $X_0 \in \mathbb{R}^{n \times d}$, loss ℓ , bandwidth σ /* fixed, or learned via Alg. 2 */

Require: Step size η_X , subsample size s , epochs T

- 1: **for** $t = 0$ **to** $T - 1$ **do**
- 2: $S_t \leftarrow \text{RANDOMSLICING}(X_t, s)$ /* Alg. 1 */
- 3: Compute PH on $X_t[S_t]$ and sparse gradient $g^{(t)} = \nabla \ell(\text{Dgm}(X_t[S_t]))$
- 4: $I_t \leftarrow \{i \in S_t : g_i^{(t)} \neq 0\}$ /* anchor indices */
- 5: **if learning** σ **then**
- 6: $\sigma \leftarrow \text{LEARNSIGMASTEP}(X_t, g^{(t)}, I_t, \sigma)$ /* Alg. 2 */
- 7: **end if**
- 8: **for all** $x \in X_t$ **do** /* evaluate NW field on all points */
- 9: $\phi_i \leftarrow \exp(-\|x - x_i\|_2^2 / (2\sigma^2))$ for all $i \in I_t$
- 10: $Z \leftarrow \sum_{j \in I_t} \phi_j$
- 11: $\bar{v}_t(x) \leftarrow \sum_{i \in I_t} (\phi_i / Z) g_i^{(t)}$
- 12: $x \leftarrow x - \eta_X \bar{v}_t(x)$
- 13: **end for**
- 14: **end for**
- 15: **return** X_T

B OMITTED PROOFS

Proof of Lemma 4.1. Let $\Delta = \frac{n-1}{s-1}$. Since $r_k = \lfloor k\Delta \rfloor$,

$$r_{k+1} \leq (k+1)\Delta,$$

and

$$-r_k < -(k\Delta - 1).$$

Hence

$$r_{k+1} - r_k \leq \Delta + 1,$$

hence

$$r_{k+1} - r_k \leq \lceil \Delta \rceil,$$

which yields the bound. \square

Proof of Lemma 4.3. Let

$$F_x(v) := \sum_{i \in I} \phi_i(x) \|v - g_i\|_2^2.$$

Expanding,

$$F_x(v) = \left(\sum_i \phi_i(x) \right) \|v\|_2^2 - 2 \langle v, \sum_i \phi_i(x) g_i \rangle + C(x),$$

hence F_x is strictly convex and differentiable with

$$\nabla_v F_x(v) = 2Z(x)v - 2 \sum_i \phi_i(x) g_i.$$

Setting $\nabla_v F_x(v) = 0$ yields

$$v = \sum_i \left(\frac{\phi_i(x)}{Z(x)} \right) g_i = \sum_i w_i(x) g_i = \bar{v}(x).$$

\square

Proof of Lemma 4.4. For any x ,

$$\bar{v}(x) - \bar{v}'(x) = \sum_i w_i(x) (g_i - g'_i).$$

Taking norms and using $\sum_i w_i(x) = 1$,

$$\begin{aligned} \|\bar{v}(x) - \bar{v}'(x)\|_2 &\leq \sum_i w_i(x) \|g_i - g'_i\|_2 \\ &\leq \max_i \|g_i - g'_i\|_2. \end{aligned}$$

\square

Proof of Theorem 4.5. For each $i \in I$, we have,

$$w_i(x_i) = \frac{1}{1 + \rho_i(\sigma)},$$

and,

$$\sum_{j \neq i} w_j(x_i) = \frac{\rho_i(\sigma)}{1 + \rho_i(\sigma)}$$

Write $\bar{v}(x_i) = w_i(x_i)g_i + \sum_{j \neq i} w_j(x_i)g_j$. Then

$$\bar{v}(x_i) - g_i = (w_i(x_i) - 1)g_i + \sum_{j \neq i} w_j(x_i)g_j.$$

Taking norms and using $\|g_j\|_2 \leq M$,

$$\begin{aligned} \|\bar{v}(x_i) - g_i\|_2 &\leq (1 - w_i(x_i)) \|g_i\|_2 + \sum_{j \neq i} w_j(x_i) \|g_j\|_2 \\ &\leq M(1 - w_i(x_i)) + M \sum_{j \neq i} w_j(x_i). \end{aligned}$$

Furthermore,

$$1 - w_i(x_i) = \frac{\rho_i(\sigma)}{1 + \rho_i(\sigma)} \text{ and } \sum_{j \neq i} w_j(x_i) = \frac{\rho_i(\sigma)}{1 + \rho_i(\sigma)},$$

yielding equation 10. \square

Proof of Theorem 4.6. We start by stating and proving the following lemma which provides an expression of the gradient of the weights and bounds its norm .

Lemma B.1 (Gradient of weights). *Let $\bar{x}(x) := \sum_{j \in I} w_j(x)x_j$ be the weighted barycenter of anchors at x . Then for all $i \in I$ and $x \in \mathbb{R}^d$,*

$$\nabla w_i(x) = \frac{w_i(x)}{\sigma^2} (x_i - \bar{x}(x)). \quad (16)$$

Moreover, for all x and i ,

$$\|\nabla w_i(x)\|_2 \leq \frac{\text{diam}(P)}{\sigma^2} w_i(x),$$

where $\text{diam}(P) := \max_{p,q \in I} \|x_p - x_q\|_2$

Proof of Lemma B.1. Using

$$\phi_i(x) = \exp\left(-\frac{\|x - x_i\|^2}{2\sigma^2}\right) \quad \text{and} \quad w_i = \frac{\phi_i}{Z},$$

we have,

$$\nabla \phi_i(x) = -(x - x_i) \frac{\phi_i(x)}{\sigma^2},$$

and

$$\nabla Z(x) = \sum_j \nabla \phi_j(x) = -\frac{1}{\sigma^2} \sum_j (x - x_j) \phi_j(x).$$

By the quotient rule,

$$\begin{aligned} \nabla w_i(x) &= \frac{\nabla \phi_i(x) Z(x) - \phi_i(x) \nabla Z(x)}{Z(x)^2} \\ &= -\frac{\phi_i(x)}{\sigma^2 Z(x)} \left[(x - x_i) - \sum_j w_j(x) (x - x_j) \right]. \end{aligned}$$

Since

$$\begin{aligned} \sum_j w_j(x) (x - x_j) &= x - \sum_j w_j(x) x_j \\ &= x - \bar{x}(x), \end{aligned}$$

we obtain equation 16.

Hence,

$$\|\nabla w_i(x)\|_2 = \frac{w_i(x)}{\sigma^2} \|x_i - \bar{x}(x)\|_2.$$

Since $\bar{x}(x)$ is a convex combination of points in P , it lies in $\text{Conv}(P)$, hence

$$\|x_i - \bar{x}(x)\|_2 \leq \text{diam}(P).$$

\square

Differentiate $\bar{v}(x) = \sum_i w_i(x)g_i$:

$$D\bar{v}(x) = \sum_{i \in I} (\nabla w_i(x)) \otimes g_i.$$

Using $\|a \otimes b\|_{\text{op}} = \|a\|_2 \|b\|_2$,

$$\begin{aligned} \|D\bar{v}(x)\|_{\text{op}} &\leq \sum_i \|\nabla w_i(x)\|_2 \|g_i\|_2 \\ &\leq M \sum_i \|\nabla w_i(x)\|_2. \end{aligned}$$

Apply Lemma B.1:

$$\begin{aligned} \sum_i \|\nabla w_i(x)\|_2 &\leq \frac{\text{diam}(P)}{\sigma^2} \sum_i w_i(x) \\ &= \frac{\text{diam}(P)}{\sigma^2}, \end{aligned}$$

giving equation 11. The Lipschitz inequality follows from the mean value theorem. \square

C ADDITIONAL KERNEL VS. NW VISUALIZATIONS

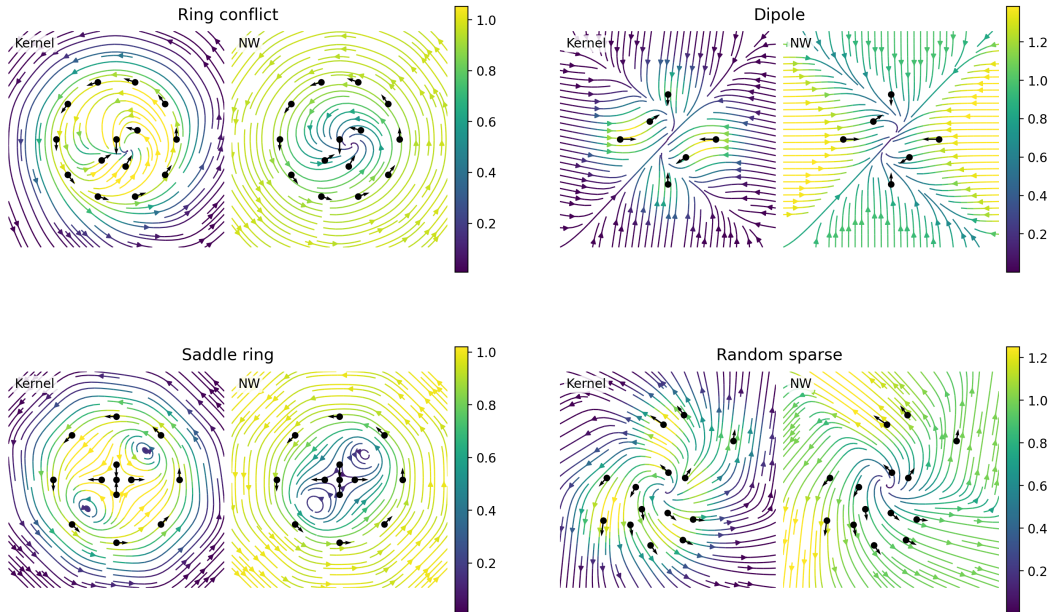


Figure 5: **Kernel interpolation vs. NW smoothing under four anchor configurations.** In each panel, the left sub-plot is exact Gaussian-RKHS interpolation (KERNEL) and the right sub-plot is Nadaraya-Watson smoothing (NW). Colors represent $\|v(x)\|$. All cases use the same anchor positions, anchor vectors, and bandwidth σ within a panel.

Fig. 5 highlights systematic geometric differences between exact kernel interpolation and NW smoothing. The most robust distinction is that KERNEL produces an *unnormalized* Gaussian expansion, so $\|v(x)\|$ tends to *decay away from anchors*, whereas NW computes a *normalized* convex combination of anchor vectors, yielding a field that behaves like a *local barycentric average* and therefore need not vanish in the far-field.

Ring conflict. This configuration resembles the qualitative behavior of persistence-driven updates: multiple anchors distributed on a ring induce a swirl-like motion, while a few interior anchors introduce competing directions. KERNEL tends to concentrate magnitude near the anchor set (visible attenuation toward the boundary of the window), and the flow lines exhibit sharper steering to satisfy exact anchor constraints. NW yields a more uniformly energized outer region and a smoother transition between the ring-induced swirl and the interior conflict, reflecting local averaging rather than exact matching.

Dipole. A small number of strong, antagonistic anchors produce a dipole-like geometry. KERNEL emphasizes the dominance of strong anchors by producing localized high-norm lobes near them, with rapid attenuation elsewhere. NW spreads influence more evenly: large regions inherit the prevailing direction under normalized weighting, resulting in wider basins of attraction/repulsion and smoother separatrices.

Saddle ring. Here, a ring component combines with a central saddle structure. KERNEL produces several localized vortical/saddle patterns whose intensities decay away from anchors. NW preserves the qualitative topology of the flow but exhibits a more “global” strength profile and smoother interactions between the saddle and ring components, consistent with a convex combination mechanism.

Random sparse. With irregular anchor placement, KERNEL can generate pronounced local features around isolated anchors, with relatively weak far-field magnitude. NW is notably robust:

because each evaluation is a normalized average over nearby anchors, the field remains well-defined and smooth across the domain, often producing globally coherent motion even when anchors are unevenly distributed.

Overall, these visualizations support two practical messages relevant to persistence-based optimization: (i) NW smoothing behaves like a geometry-aware averaging operator that propagates sparse topological gradients beyond anchors with minimal computational overhead, and (ii) normalization changes the far-field behavior compared to RKHS interpolation, which can be beneficial for spreading topological updates across the ambient point set.

D FURTHER EXPERIMENTS IN 2D

We report two additional 2D experiments that complement the main text and further illustrate the empirical behavior of our scalable pipeline (NW-CONVOLUTION, i.e. projection-stratified subsampling with NW smoothing) against (i) vanilla topological gradient descent (sparse updates) and (ii) kernel-based diffeomorphic interpolation. For both experiments, we use $N = 1000$ points, subsample size $s = 100$, and optimize with step size $\eta_X = 0.05$. For our method we use a base bandwidth $\sigma_0 = 0.1$ and, when enabled, a one-step look-ahead update of σ with learning rate $\eta_\sigma = 10^{-3}$ (cf. Section A.2), while for diffeomorphic interpolation, we use a bandwidth $\sigma = 0.1$.

D.1 MAXPERS ON A UNIFORM SQUARE

We first consider a point cloud initialized uniformly on $[-1, 1]^2$ and optimize a standard topological *augmentation* objective that encourages the emergence of persistent H_1 features while keeping points within the box. Concretely, given $\text{Dgm}_1(X) = \{(b_\alpha, d_\alpha)\}$, we use

$$\mathcal{L}_{\text{maxpers}}(X) = - \sum_{\alpha \in \text{Dgm}_1(X)} \left(\frac{1}{2}(d_\alpha - b_\alpha) \right)^2 + \sum_{i=1}^N \max(\|x_i\|_\infty - 1, 0), \quad (17)$$

i.e. a negative squared total persistence term in dimension 1 (encouraging long-lived cycles) with a soft constraint that discourages leaving $[-1, 1]^2$. We optimize for $T = 300$ epochs.

Figure 6 shows the initial and final configurations together with the loss curves. The vanilla method remains limited by the sparsity of the topological gradient, whereas both smoothing-based methods propagate the sparse signal to non-anchor points and more reliably shape the cloud toward configurations with stronger H_1 persistence. In this setting, NW-CONVOLUTION achieves the lowest objective value while being markedly faster per epoch than kernel-based interpolation (Table 2).




Figure 6: **2D maxpers from a uniform square.** Top: initial cloud sampled uniformly in $[-1, 1]^2$ and final configurations after $T = 300$ epochs for vanilla updates, kernel-based diffeomorphic interpolation (KERNEL), and our method (NW-CONVOLUTION = random slicing + NW smoothing). Bottom: loss trajectories for equation 17 (lower is better).

D.2 COLLAPSER ON A NOISY CIRCLE

We next consider a complementary *simplification* task: the initial cloud is a noisy circle in \mathbb{R}^2 , and we optimize the `collapser` loss, which penalizes both births and deaths and therefore pushes topological features toward short scales. In our implementation this reads

$$\mathcal{L}_{\text{coll}}(X) = \sum_{\alpha \in \text{Dgm}_1(X)} (d_\alpha + b_\alpha)^2, \quad (18)$$

Table 2: **2D maxpers: best loss and runtime.** Best objective value (lower is better) and average time per epoch.

Method	Best Loss ↓	Avg Time/Epoch (s) ↓
Vanilla	-0.0758	0.5615
Kernel	-0.1053	0.6895
NW-Convolution (ours)	-0.2066	0.3475

so that driving $\mathcal{L}_{\text{coll}}$ down encourages both early death and early birth, i.e. rapid topological collapse.

Figure 7 shows that NW-CONVOLUTION leads to a substantially faster and more decisive decrease of equation 18, and qualitatively produces a sharper collapse of the noisy ring than both baselines. Quantitatively, Table 3 confirms that our method achieves the lowest loss while remaining close to the runtime of the vanilla sparse updates, and significantly faster than kernel-based interpolation.




Figure 7: **2D collapser from a noisy circle.** Top: initial noisy circle and final point clouds after $T = 50$ optimization epochs using vanilla updates, kernel-based diffeomorphic interpolation (KERNEL), and our method (NW-CONVOLUTION = random slicing + NW smoothing). Bottom: loss curves for the collapser objective equation 18 (lower is better).

Table 3: **2D collapser: best loss and runtime.** Best objective value (lower is better) and average time per epoch.

Method	Best Loss ↓	Avg Time/Epoch (s) ↓
Vanilla	3.057	0.5233
Kernel	2.615	0.8800
NW-Convolution (ours)	0.064	0.5406