

# FibeRed: Fiberwise Dimensionality Reduction of Topologically Complex Data with Vector Bundles

Luis Scoccola   

Department of Mathematics, Northeastern University, Boston, MA, USA

Jose A. Perea   

Department of Mathematics and Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA

---

## Abstract

Datasets with non-trivial large scale topology can be hard to embed in low-dimensional Euclidean space with existing dimensionality reduction algorithms. We propose to model topologically complex datasets using vector bundles, in such a way that the base space accounts for the large scale topology, while the fibers account for the local geometry. This allows one to reduce the dimensionality of the fibers, while preserving the large scale topology. We formalize this point of view and, as an application, we describe a dimensionality reduction algorithm based on topological inference for vector bundles. The algorithm takes as input a dataset together with an initial representation in Euclidean space, assumed to recover part of its large scale topology, and outputs a new representation that integrates local representations obtained through local linear dimensionality reduction. We demonstrate this algorithm on examples coming from dynamical systems and chemistry. In these examples, our algorithm is able to learn topologically faithful embeddings of the data in lower target dimension than various well known metric-based dimensionality reduction algorithms.

**2012 ACM Subject Classification** Mathematics of computing → Algebraic topology

**Keywords and phrases** topological inference, dimensionality reduction, vector bundle, cocycle

**Digital Object Identifier** 10.4230/LIPIcs.SoCG.2023.56

**Related Version** *Full Version*: <https://arxiv.org/abs/2206.06513>

**Supplementary Material** Proof-of-concept implementation [44]

**Funding** *Luis Scoccola*: supported by the NSF through grants CCF-2006661 and CAREER award DMS-1943758.

*Jose A. Perea*: supported by the NSF through grants CCF-2006661 and CAREER award DMS-1943758.

**Acknowledgements** The authors thank Matt Pickenbrock for various fruitful conversations.

## 1 Introduction

**Motivation.** We take the manifold hypothesis at face value and consider data consisting of a finite sample of a Riemannian manifold. We take the goal of *dimensionality reduction* to be that of learning an embedding of the input data in low dimension, in such a way that the differentiable structure of the underlying manifold is preserved. This is different from *charting*, whose objective we take to be that of producing local parametrizations of the data that, together, cover the entire manifold.

We refer to dimensionality reduction algorithms which aim to preserve metric relationships and do not explicitly incorporate large scale topology in their objective function as *metric-based*. Metric-based algorithms work best when the Riemannian manifold underlying the data can be isometrically embedded in the target dimension. For example, algorithms such as Isomap [53], Local Tangent Space Alignment (LTSA) [59], and Hessian Eigenmaps (HLE)



© Luis Scoccola and Jose A. Perea;

licensed under Creative Commons License CC-BY 4.0

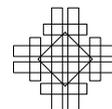
39th International Symposium on Computational Geometry (SoCG 2023).

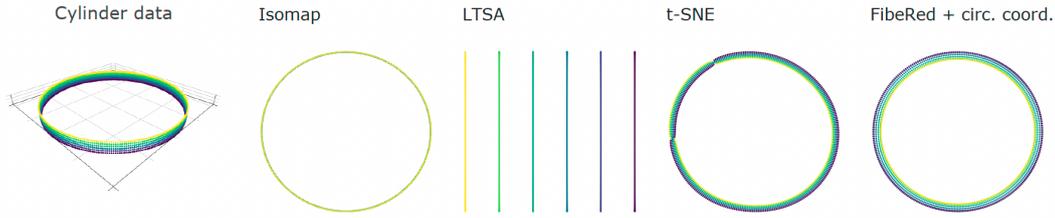
Editors: Erin W. Chambers and Joachim Gudmundsson; Article No. 56; pp. 56:1–56:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany





■ **Figure 1** *Left*: A sample from a cylinder with height equal to 0.15 times its radius, colored by height. The cylinder is developable, since it is diffeomorphic to an annulus in  $\mathbb{R}^2$ , and is also flat, but it is not isometric to the annulus, which also has a flat, yet distinct, Riemannian metric. *Center three*: Well known dimensionality reduction algorithms run on the cylinder data. The outputs are representative of other parameter choices and of Laplacian Eigenmaps (LE) [8], Diffusion Maps (DM) [13], LLE [42], HLLE, t-SNE [55], and UMAP [28]. Some algorithms only capture the circularity, others only the local 2D structure, while others capture both, but they are not able to consistently align the local 2D structure. *Right*: The output of fiberwise dimensionality reduction.

[15] assume that the manifold  $\mathcal{X}$  underlying the data is *isometrically developable*, in the sense that  $\mathcal{X}$  is a  $d$ -dimensional Riemannian manifold for which there exists an embedded  $d$ -dimensional manifold  $\mathcal{X}' \subseteq \mathbb{R}^d$  and a diffeomorphism  $\mathcal{X}' \rightarrow \mathcal{X}$  which is a Riemannian isometry. An isometrically developable manifold  $\mathcal{X}$  is necessarily *flat* (i.e., locally isometric to Euclidean space) and *developable* (i.e., diffeomorphic to an embedded  $d$ -dimensional manifold  $\mathcal{X}' \subseteq \mathbb{R}^d$ ). But a manifold can be flat and developable without it being isometrically developable: a simple example is that of a straight cylinder in  $\mathbb{R}^3$  (Figure 1). As observed in [24], and shown in Figure 1, already in the setting of a flat and developable  $d$ -dimensional manifold, metric-based dimensionality reduction algorithms can fail to find an embedding of the data in  $\mathbb{R}^d$ . On the mathematical side, while Whitney’s embedding theorem [57] guarantees that any closed  $d$ -dimensional manifold admits a smooth ( $C^\infty$ ) embedding in  $2d$  dimensions, a smooth, Riemannian isometric embedding of a closed  $d$ -dimensional Riemannian manifold can require in the order of  $d^2$  dimensions [11]. Thus, the preservation of distances requires more complicated embeddings than the preservation of topology.

If we remove a small portion of the cylinder of Figure 1, in order to make it a curved rectangle, most metric-based dimensionality reduction algorithms have no problem finding an embedding in  $\mathbb{R}^2$ . It is thus the non-trivial topology of the cylinder – its circularity – that causes difficulties. This suggests that embeddings of topologically non-trivial manifolds can be built by gluing local representations along a representation of the global topological structure: in the case of the cylinder, one would try to glue 2D patches around a circle in a globally consistent manner. This leads to the following problem, formalized as the vector bundle embedding problem (Problem 1):

*Given a dataset  $X$  and an initial map  $X \rightarrow \mathbb{R}^D$  capturing the large scale topology of  $X$ , find a new representation  $X \rightarrow \mathbb{R}^D$  that captures the large scale topology as well as the local geometry.*

We call our approach to the above problem *fiberwise dimensionality reduction* (FIBERED). In the examples of Section 4, we focus on manifolds with an essential loop, and, as initial map, we use circular coordinates based on persistent cohomology [14, 37], a technique from Topological Data Analysis [32, 18]. Nevertheless, the approach is not restricted to the case of a circular initial embedding and one could use as initial map one constructed by, e.g., other cohomological coordinates [35, 38, 43], standard non-linear dimensionality reduction methods [23, 13], or lens functions as in [50, Section 4].

**Contributions.** We show that the theory of vector bundles is useful in abstracting (Section 2.1), devising solutions to (Section 3), and computing obstructions to solving (Section 2.3) the problem of extending an initial coarse representation of data to a new, more descriptive representation. We demonstrate with computational examples (Section 4) that topological inference for vector bundles can be carried out in practice. In particular, we show that efficient embeddings and chartings of topologically non-trivial data can be learned with this approach and give examples supporting the claim that metric-based dimensionality reduction algorithms are often not able to find such representations. We implement our main algorithm in [44].

**Related work.** Various dimensionality reduction schemes [52, 41, 10] learn a global alignment of local linear models from the local interactions of the models, which can be challenging in the presence of non-trivial topology. In contrast, our approach assumes a global topological representation is given and builds and aligns the local linear models along this representation.

There has been recent interest in designing topology-preserving dimensionality reduction schemes [26, 58, 31, 56]. Our approach is different from previous approaches we are aware of, as it builds a new representation around an initial topological representation, instead of using topology to regularize an essentially metric objective.

Our cut-unfold technique of [45, Appendix C.1] has a similar goal to that of [24, 58], which propose to tear a data manifold in order to find efficient representations of it. A main difference is that our technique allows the user to select a specific hole to cut and to use topological persistence to guide this choice.

## 2 The vector bundle embedding problem

For background, please refer to [45, Appendix A.1]. In Section 2.1 we describe the Vector Bundle Embedding problem; in Section 2.2, we recall the notion of discrete vector bundle that we use to estimate vector bundles from finite samples; and in Section 2.3 we explain how characteristic classes of vector bundles give computable obstructions to solving the vector bundle embedding problem and can thus be used for parameter selection.

### 2.1 Main problem

Let  $\mathcal{B}$  be a closed differentiable manifold and let  $\pi : \mathcal{X} \rightarrow \mathcal{B}$  be a rank  $r$  Euclidean vector bundle with zero-section  $s_0 : \mathcal{B} \rightarrow \mathcal{X}$ , where by *Euclidean* we mean that  $\pi$  is endowed with a scalar product on each fiber  $\pi^{-1}(b) \subseteq \mathcal{X}$ , which varies smoothly with  $b \in \mathcal{B}$ . The main problem we seek to solve is that of extending an embedding  $\mathcal{B} \rightarrow \mathbb{R}^D$  to a fiberwise isometric embedding of  $\mathcal{X}$ , as follows:

► **Problem 1.** *Given an embedding  $\iota : \mathcal{B} \rightarrow \mathbb{R}^D$ , find a fiberwise isometric embedding  $\bar{\iota} : \mathcal{X} \rightarrow \mathbb{R}^D$  that extends  $\iota$  in the sense that  $\bar{\iota} \circ s_0 = \iota$ , and that is orthogonal to  $\mathcal{B}$ , in the sense that  $\bar{\iota}(\pi^{-1}(b)) \perp \iota(T_b\mathcal{B})$  for all  $b \in \mathcal{B}$ .*

By *fiberwise isometric embedding*  $\mathcal{X} \rightarrow \mathbb{R}^D$  we mean a map that is a linear isometry when restricted to each fiber  $\pi^{-1}(b) \subseteq \mathcal{X}$ , where  $b \in \mathcal{B}$ .

Let  $\nu : N \rightarrow \mathcal{B}$  be the normal bundle of the embedding  $\iota : \mathcal{B} \rightarrow \mathbb{R}^D$  and endow  $\nu$  with the Euclidean structure inherited from  $\mathbb{R}^D$ . The following result reduces Problem 1 to a problem only involving vector bundles.

► **Lemma 2.** *Problem 1 admits a solution if and only if there exists a morphism  $\mathcal{X} \rightarrow N$  of vector bundles over  $\mathcal{B}$  that is an isometry in each fiber.*

In order to do this, we trivialize the bundles  $\mathcal{X}$  and  $N$  over a common cover of the base  $\mathcal{B}$  and construct the embedding  $\mathcal{X} \rightarrow N$  by restricting to each element of the cover. Formally, we proceed as follows.

Let  $e$  be the dimension of  $\mathcal{B}$ , so that the rank of  $\nu$  is  $D - e$ . Let  $\mathcal{U} = \{U_i\}$  be a cover of  $\mathcal{B}$  such that both  $\pi$  and  $\nu$  can be trivialized over  $\mathcal{U}$  and let  $\mathcal{X}_i := \pi^{-1}(U_i)$ . Recall that  $\mathbb{V}(n, m)$  denotes the Stiefel manifold, which consist of  $m$ -by- $n$  matrices with orthonormal columns and that  $O(n) = \mathbb{V}(n, n)$  denotes the orthogonal group. Let  $\alpha = \{\alpha_i : U_i \rightarrow \mathbb{V}(D - e, D)\}$  be local bases for  $N$ , and let  $\Theta = \{\Theta_{ij} : U_i \cap U_j \rightarrow O(D - e)\}$  be defined by  $\Theta_{ij}(b) = \alpha_i(b)\alpha_j(b)^T$  for all  $b \in U_i \cap U_j$ , so that  $\Theta$  is a cocycle with associated vector bundle  $\nu$ . Finally, let  $\{(\pi|_{\mathcal{X}_i}, f_i) : \mathcal{X}_i \rightarrow U_i \times \mathbb{R}^r\}$  be a metric trivialization of  $\mathcal{X}$  over  $\mathcal{U}$  and let  $\Omega = \{\Omega_{ij} : U_i \cap U_j \rightarrow O(r)\}$  be defined as the unique set of maps satisfying

$$\Omega_{ij}(\pi(x)) f_j(x) = f_i(x), \text{ for all } x \in \mathcal{X}_i \cap \mathcal{X}_j, \quad (1)$$

so that  $\Omega$  is a cocycle with associated vector bundle  $\pi$ . We refer to the maps  $\{f_i : \mathcal{X}_i \rightarrow \mathbb{R}^r\}$  as the *fiber coordinates*. With these definitions, one can use Lemma 2 to prove the following.

► **Proposition 3.** *There exists a fiberwise isometric embedding  $\mathcal{X} \rightarrow N$  if and only if there exist maps  $\Phi = \{\Phi_i : U_i \rightarrow \mathbb{V}(r, D - e)\}$  such that*

$$\Phi_i(b)\Omega_{ij}(b) = \Theta_{ij}(b)\Phi_j(b), \text{ for all } i \text{ and } j \text{ and } b \in U_i \cap U_j. \quad (2)$$

Given the maps  $\Phi = \{\Phi_i : U_i \rightarrow \mathbb{V}(r, D - e)\}$  of Proposition 3, one obtains the fiberwise isometric embedding  $\bar{\iota} : \mathcal{X} \rightarrow \mathbb{R}^D$  by  $\bar{\iota}(x) = \alpha_i(b)\Phi_i(b)f_i(x) + \iota(b)$ , where  $b = \pi(x)$ .

In general, the fiberwise isometric embedding  $\bar{\iota} : \mathcal{X} \rightarrow \mathbb{R}^D$  is not an embedding of  $\mathcal{X}$  as a manifold, since different fibers may intersect. Nonetheless, if  $\tau > 0$  is the *reach* [1, Definition 2.1] of  $\iota(\mathcal{B}) \subseteq \mathbb{R}^D$ , i.e., the largest possible radius of a uniform tubular neighborhood around  $\iota(\mathcal{B})$ , one can find an embedding of a full-dimensional compact subset of  $\mathcal{X}$  by scaling the fibers by a fraction of  $\tau$ , as follows. Let  $\text{disk}(\pi) \subseteq \mathcal{X}$  be the unit disk of the bundle  $\pi$ , namely, the subspace of points  $x \in \mathcal{X}$  such that  $\|x - s_0(\pi(x))\| \leq 1$ , where  $\|\cdot\|$  denotes the norm of the fiber  $\pi^{-1}(\pi(x))$  induced by the Euclidean structure of  $\pi$ . Then, the following formula gives an embedding  $\text{disk}(\pi) \rightarrow \mathbb{R}^D$ :

$$x \mapsto c\tau \cdot \alpha_i(\pi(x))\Phi_i(\pi(x))f_i(x) + \iota(\pi(x)), \text{ for } \pi(x) \in U_i, \quad (3)$$

where  $0 < c < 1$  is any fixed constant.

## 2.2 Vector bundles from finite samples

In practice, continuous maps to a Stiefel manifold or orthogonal group – such as the maps  $\{\alpha_i : U_i \rightarrow \mathbb{V}(D - e, D)\}$  or the cocycle  $\{\Omega_{ij} : U_i \cap U_j \rightarrow O(r)\}$  of Section 2.1 – are hard to work with, as they are potentially determined by an infinite amount of data. One of the main takeaways of [46] is that one can work with Euclidean vector bundles in practice by considering only constant maps into Stiefel manifolds or orthogonal groups. In order to accomplish this, one relaxes the notion of Euclidean vector bundle as follows.

Given a simplicial complex  $S$ , a rank  $r$  *discrete approximate* cocycle on  $S$  ([46, Definition 5.1]) consists of a family of matrices  $\{\Omega_{ij} \in O(r)\}$  indexed by the oriented 1-simplices of  $S$ , which satisfies  $\Omega_{ij} = \Omega_{ji}^T$ . There is a similar way of discretizing maps into a Stiefel manifold ([46, Definition 5.4]). These discretizations can be used to represent usual vector bundles [46, Theorem A] and any vector bundle can be represented in this way [46, Proposition 5.7].

This justifies the fact that, in Section 3, we discretize the base  $\mathcal{B}$  by considering the simplicial complex given by the nerve of a cover  $\mathcal{U} = \{U_i\}$  and we consider constant maps from  $U_i$  into a Stiefel manifold and from  $U_i \cap U_j$  into an orthogonal group.

## 2.3 Computable obstructions to vector bundle embedding

The theory of vector bundles provides us with algebraic obstructions to solving Problem 1, namely, characteristic classes. We now give a few details about the subject; we refer the reader to [30] for a detailed account of the theory of characteristic classes.

To a vector bundle  $\pi : \mathcal{X} \rightarrow \mathcal{B}$  and number  $i \in \mathbb{N}$ , one can associate an element  $w_i(\pi) \in H^i(\mathcal{B}; \mathbb{Z}/2)$  of the  $i$ th cohomology group of  $\mathcal{B}$  with coefficients in the group  $\mathbb{Z}/2$ , called the  $i$ th *Stiefel–Whitney class* of  $\pi$ . This procedure is such that, if  $\pi$  and  $\pi'$  are isomorphic vector bundles over the same base  $\mathcal{B}$ , then  $w_i(\pi) = w_i(\pi')$ .

If Problem 1 admits a solution, then there exists a complement of  $\pi$  in  $\nu$ , that is there exists a vector bundle  $\kappa$  over  $\mathcal{B}$  such that  $\pi \oplus \kappa \cong \nu$ , where  $\oplus$  denotes the direct sum of vector bundles. It follows from the Whitney product formula [30, Section 4, Axiom 3] that  $w(\pi) \smile w(\kappa) = w(\nu)$ , where  $\smile$  denotes the cup-product in cohomology [21, Section 3.2]. In particular, when Problem 1 admits a solution, we have the following:

- If  $D = r + e$ , then  $w_1(\pi) = w_1(\nu) \in H^1(\mathcal{B}; \mathbb{Z}/2)$ .
- If  $D = r + e + 1$ , then  $w_2(\pi) - w_1(\pi)^2 + w_1(\pi) \smile w_1(\nu) = w_2(\nu) \in H^2(\mathcal{B}; \mathbb{Z}/2)$ .

Thus, if any of these equalities is not satisfied, then Problem 1 does not admit a solution. These obstructions can be computed from finite samples using [46, Theorem C].

## 3 The fiberwise dimensionality reduction scheme

We describe the FIBERED algorithm in Sections 3.1–3.4. In Section 3.5 we justify a main subroutine of the algorithm. In Section 3.6, we explain how we choose parameters. We represent vector bundles using discrete approximate cocycles as in [46] (see Section 2.2).

To facilitate the interpretation of the different steps of the algorithm, the notation is kept as in Section 2.1, except for the spaces  $\mathcal{X}$  and  $\mathcal{B}$ , which we denote here by  $X$  and  $B$  to emphasize the fact that we are working with finite samples  $X \subseteq \mathcal{X}$  and  $B \subseteq \mathcal{B}$ . See also Figure 2 for a schematic representation of some of the main steps (4,5,6) of the algorithm.

Precise assumptions about the input of the algorithm are in [45, Appendix B.1]. Our algorithm can be efficiently implemented; we give more details in [45, Appendix B.2].

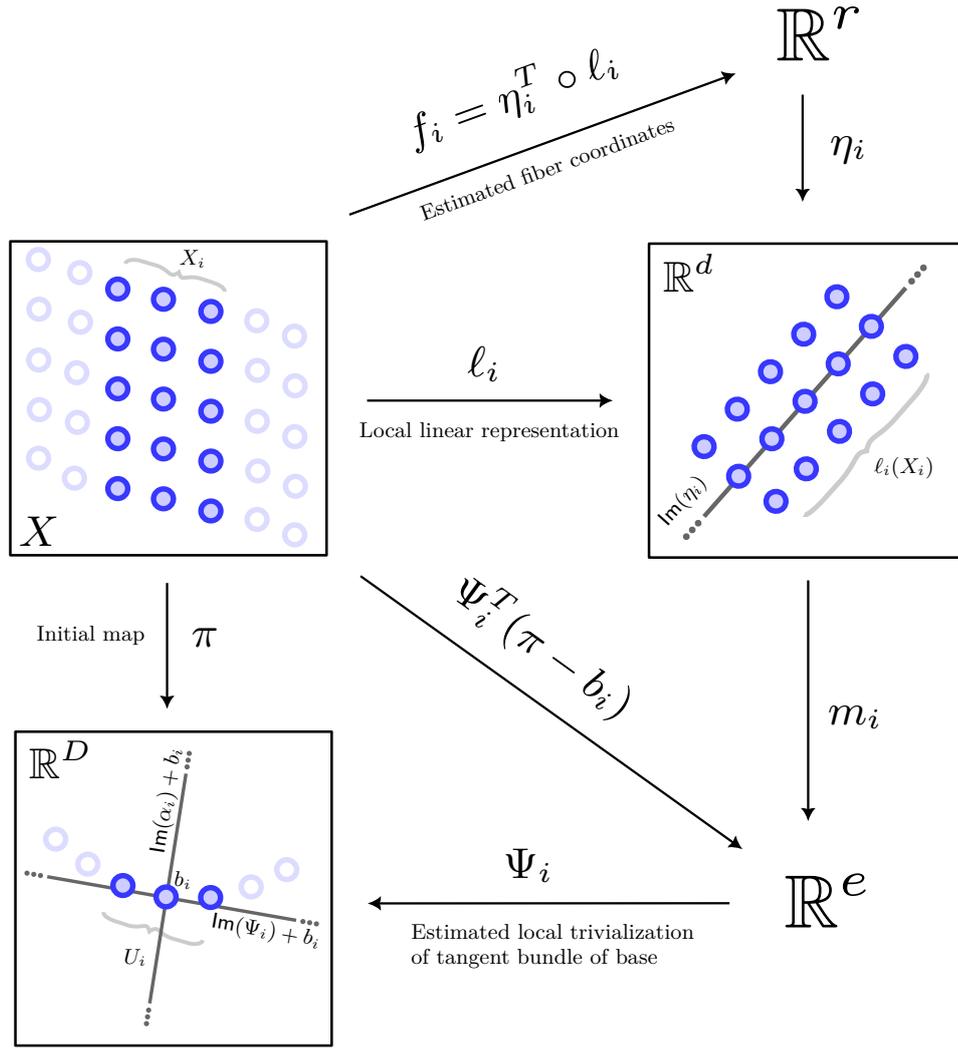
### 3.1 Main routine

**Inputs.** A dataset represented by a finite set  $X$  together with a distance matrix  $\partial : X \times X \rightarrow \mathbb{R}$ ; and a function  $\pi : X \rightarrow \mathbb{R}^D$ . We let  $B := \pi(X) \subseteq \mathbb{R}^D$ .

**Parameters.** A number  $k \in \mathbb{N}$ , the number of sets we use to construct a cover of  $B$ ; a number  $\mathbf{n\_iter} \in \mathbb{N}$  used in the ALIGNFIBERS subroutine; an estimate  $e \in \mathbb{N}$  of the intrinsic dimension of  $\mathcal{B}$ ; an estimate  $d \in \mathbb{N}$  of the intrinsic dimension of  $\mathcal{X}$ ; a fiber scale  $0 < c < 1$ .

**Output.** A map  $X \rightarrow \mathbb{R}^D$ .

The pseudocode is in Algorithm 1. With this notation, the map  $\iota$  of Section 2.1 corresponds to the inclusion  $B = \pi(X) \subseteq \mathbb{R}^D$  and the rescaling of  $\bar{\iota}$  of Equation (3) corresponds to the output of the algorithm.



■ **Figure 2** Schematic representation of the main constructions in the FIBERED pipeline.

■ **Algorithm 1** FIBERED( $X, \pi, k, e, d, c, n\_iter$ ).

- 
- 1:  $\mathcal{U}, \rho \leftarrow \text{COVERANDPARTITIONUNITY}(k, B)$
  - 2:  $\mathcal{N} \leftarrow \text{NERVE}(B, \mathcal{U})$
  - 3: **for**  $1 \leq i \leq k$  **do**
  - 4:      $l_i \leftarrow \text{LOCALLINEARREPRESENTATION}(X, \mathcal{U}, d, i)$
  - 5:      $\Psi_i, \alpha_i \leftarrow \text{ESTTANGANDNORMBUN}(B, \mathcal{U}, e, i)$
  - 6:      $\bar{f}_i \leftarrow \text{ESTNORMFIBERCOORDINATES}(B, \Psi_i, l_i)$
  - 7: **end for**
  - 8:  $\tau \leftarrow \text{ESTREACH}(B, \mathcal{U}, \Psi)$
  - 9: **for**  $(ij) \in \mathcal{N}$  **do**
  - 10:      $\Omega_{ij}, \Theta_{ij} \leftarrow \text{ESTCOCYCLES}(\bar{f}_i, \bar{f}_j, \alpha_i, \alpha_j)$
  - 11: **end for**
  - 12:  $\Phi \leftarrow \text{ALIGNFIBERS}(\mathcal{N}, \Omega, \Theta, n\_iter)$
  - 13: **return** ASSEMBLE( $\rho, \tau, c, \alpha, \Phi, \bar{f}, \pi$ )
-

### 3.2 Subroutines

**Compute cover and partition of unity (COVERANDPARTITIONUNITY).** We compute a cover  $\mathcal{U} = \{U_i \subseteq B\}_{1 \leq i \leq k}$  of  $B$  as follows. We first run on  $B$  an approximate algorithm for the  $k$ -center problem. We use a simple, greedy approach, but more sophisticated options are available (see, e.g., [17] for a survey). This results in  $k$  points  $\{b_1, \dots, b_k\} \subseteq B$  and in a radius  $c > 0$  such that any point of  $B$  is at distance at most  $c$  from some  $b_i$ . We then let  $U_i = \{b \in B : \|b - b_i\| < 3c\}$ . The factor of 3 is arbitrary; we choose it to ensure that elements of the cover have sufficiently large intersections.

We compute a partition of unity  $\rho = \{\rho_i : U_i \rightarrow \mathbb{R}\}$  subordinate to  $\mathcal{U}$  by first defining  $p_i(x) = \exp(-1/(1 - (\|x - b_i\|/(3c))^2))$  for  $x \in U_i$  and  $p_i(x) = 0$  for  $x \notin U_i$ , and then normalizing as follows  $\rho_i(x) = p_i(x) / \sum_j p_j(x)$ .

**Compute nerve of cover (NERVE).** We let  $\mathcal{N}$  be the undirected graph with vertices  $1 \leq i \leq k$  and an edge  $(ij)$  with weight  $s_{ij} = |U_i \cap U_j|$  when  $U_i \cap U_j \neq \emptyset$ .

**Compute local linear representation (LOCALLINEARREPRESENTATION).** Given  $1 \leq i \leq k$ , we let  $X_i := \pi^{-1}(U_i)$  and apply a linear dimensionality reduction algorithm to each  $X_i$ , resulting in a function  $\ell'_i : X_i \rightarrow \mathbb{R}^d$ . In our implementation, we use classical multidimensional scaling (see, e.g., [9]). We then mean-center  $\ell'_i$  to get a function  $\ell_i : X_i \rightarrow \mathbb{R}^d$ .

**Estimate local trivialization of tangent and normal bundle (ESTTANGANDNORMBUN).** Given  $1 \leq i \leq k$ , we compute an orthonormal frame  $\Psi_i \in \mathcal{V}(e, D)$  by applying PCA with target dimension  $e$  to  $U_i \subseteq \mathbb{R}^D$ . We then compute an orthonormal frame  $\alpha_i \in \mathcal{V}(D - e, D)$  such that  $\alpha_i \perp \Psi_i$ .

**Estimate normalized fiber coordinates (ESTNORMFIBERCOORDINATES).** Given  $1 \leq i \leq k$ , we define  $t : X_i \rightarrow \mathbb{R}^e$  by  $t(x) = \Psi_i^T(\pi(x) - b_i)$ . We find a linear transformation  $m_i : \mathbb{R}^d \rightarrow \mathbb{R}^e$ , which has minimal Frobenius norm and minimizes

$$\sum_{x \in X_i} \|t(x) - m_i(\ell_i(x))\|^2, \quad (4)$$

and compute an orthonormal frame  $\eta_i \in \mathcal{V}(r, d)$  with image in the kernel of  $m_i$ . We let  $f_i := \eta_i^T \circ \ell_i : X_i \rightarrow \mathbb{R}^r$ , and obtain a normalized fiber coordinate  $\bar{f}_i : X_i \rightarrow \mathbb{R}^r$  with image contained in the unit ball by normalizing  $f_i$ . We justify these choices in Section 3.5.

**Estimate reach (ESTREACH).** If  $\{b_1, \dots, b_k\} \subseteq B$  are the centers of the  $k$  balls used to construct the cover  $\mathcal{U}$  in COVERANDPARTITIONUNITY, we compute an estimate of the reach of  $B$  by

$$\tau = \inf_{i \neq j} \frac{\|b_j - b_i\|^2}{2\sqrt{\|b_j - b_i\|^2 - \|\Psi_i^T(b_j - b_i)\|^2}}.$$

This formula is equivalent to [1, Equation 6.1], where it is proven that, under suitable assumption, it yields a consistent estimator of the reach.

**Estimate cocycles for  $\nu$  and  $\pi$  (ESTCOCYCLES).** Based on Equation (1), given  $(ij) \in \mathcal{N}$ , we compute an orthogonal matrix  $\Omega_{ij} \in O(r)$  which minimizes

$$\sum_{x \in X_i \cap X_j} \|\Omega_{ij} f_j(x) - f_i(x)\|^2.$$

We also compute an orthogonal matrix  $\Theta_{ij} \in O(D - e)$  which minimizes  $\|\Theta_{ij} - \alpha_i^T \alpha_j\|_F$ , where  $\|\cdot\|_F$  denotes the Frobenius norm. Both minimizations are instances of the orthogonal Procrustes problem, which can be solved using SVD (see, e.g., [22, Section 7.4]).

**Align fibers (ALIGNFIBERS).** Based on Equation (2), we compute orthonormal frames  $\{\Phi_i \in \mathcal{V}(r, D - e)\}$  minimizing the following expression; we describe the minimization procedure in Section 3.3:

$$\sum_{(ij) \in \mathcal{N}} s_{ij} \|\Phi_i \Omega_{ij} - \Theta_{ij} \Phi_j\|_F. \quad (5)$$

**Compute final representation (ASSEMBLE).** Based on Eq. 3, we represent  $x \in X$  by

$$\sum_{1 \leq i \leq k} \rho_i(x) (c\tau \cdot \alpha_i \Phi_i \bar{f}_i(x) + \pi(x)).$$

### 3.3 Minimizing Equation (5)

The minimization problem in ALIGNFIBERS is non-convex, so a possible solution is to do gradient descent in a product Stiefel manifold. This is the approach we take, except that we avoid explicitly computing a gradient, and take a sampling based approach, as done in, e.g., LargeVis [51]. Before describing the approach, we note that, in the case  $D = r + e$ , the Stiefel manifold  $\mathcal{V}(r, r)$  is equal to the orthogonal group  $O(r)$ , which is disconnected. Thus, in this case, any local optimization approach to minimizing Equation (5), such a gradient descent, is bound to fail. In Section 3.4 we describe a procedure based on the notion of synchronization (see, e.g., [48]) that reduces the problem from having to align using matrices in  $O(r)$  to using matrices in  $SO(r)$ , which is connected.

**Iterative procedure.** We start by initializing  $\{\Phi_i \in \mathcal{V}(r, D - e)\}$  at random and setting  $a = 1$ . For  $1 \leq n \leq \text{n\_iter}$ , we proceed as follows. We sample an edge  $(ij) \in \mathcal{N}$  with probability proportional to its weight  $s_{ij}$ , let  $M$  be an orthonormal frame minimizing  $\|M \Omega_{ij} - \Theta_{ij} \Phi_j\|_F$ , and replace  $\Phi_i$  with a closest orthonormal frame to the convex combination  $(1 - a)\Phi_i + aM$ . Finally, we replace  $a$  with  $1 - n/\text{n\_iter}$ .

### 3.4 Preprocessing in the case $D = r + e$

In this case, the matrices  $\Phi$ ,  $\Omega$ , and  $\Theta$  are in  $O(r)$ . The preprocessing consists of replacing the matrices  $\{\Theta_{ij}\}$  by matrices that induce an equivalent problem to the one of minimizing Equation (5), but for which the matrices  $\{\Phi_i\}$  we look for can be taken to be in the special orthogonal group  $SO(r)$ , which is connected.

Note that, if we want  $\Phi_i \Omega_{ij}$  and  $\Theta_{ij} \Phi_j$  to belong to the same connected component of  $O(r)$ , then we must have  $\det(\Omega_{ij}) \det(\Theta_{ij}) = \det(\Phi_i) \det(\Phi_j) \in O(1) = \{-1, +1\}$ . This

suggests that we can let  $\omega_{ij} = \det(\Omega_{ij}) \det(\Theta_{ij}) \in O(1)$  and consider first the problem of finding  $\{\lambda_i \in O(1)\}$  such that  $\lambda_i \lambda_j = \omega_{ij}$ , which leads to minimizing the objective function

$$\sum_{(ij) \in \mathcal{N}} s_{ij} |\omega_{ij} - \lambda_i \lambda_j|^2.$$

This is a well known synchronization problem, for which an approximate solution can be found effectively and efficiently with spectral methods [49, 4]. Here, we use [4, Algorithm 2.3], with  $d = 1$ , which yields an approximate solution  $\{\lambda_i \in O(1)\}$ .

Given  $\lambda \in O(1) = \{-1, +1\}$  let  $M(\lambda) \in O(r)$  be the diagonal matrix with all diagonal entries equal to 1, except for the first one, which is equal to  $\lambda$ . With this in mind, we can replace  $\Theta_{ij}$  by  $M(\lambda_i) \Theta_{ij} M(\lambda_j)$ . Having done this, we can now restrict the matrices  $\{\Phi_i\}$  to belong to the connected component of  $O(r)$  of orthogonal matrices with  $+1$  as determinant. More specifically, we now can carry out the optimization procedure described above, but restricting the matrices  $\{\Phi_i\}$  to be in  $SO(r) \subseteq O(r) = V(r, r)$ .

### 3.5 Justification of estimate of fiber coordinates

Let  $x_i := s_0(b_i) \in \mathcal{X}$ . We interpret the local model  $\ell_i : X_i \rightarrow \mathbb{R}^d$  as a projection  $\ell_i : X_i \rightarrow T_x \mathcal{X} \cong \mathbb{R}^d$  of  $X_i$  onto the tangent space at the origin of the fiber  $\pi^{-1}(b_i)$ . In the idealized case ([45, Appendix B.1]), the fiber coordinate  $f_i : \mathcal{X}_i \rightarrow \mathbb{R}^r$  is given by any map fitting into a fiberwise isometric diffeomorphism  $(\pi|_{\mathcal{X}_i}, f_i) : \mathcal{X}_i \rightarrow U_i \times \mathbb{R}^r$ . When dealing with finite samples, we use the following composite as a proxy for  $f_i$ :

$$X_i \xrightarrow{\ell_i} T_{x_i} \mathcal{X} \xrightarrow{(df_i)_{x_i}} T_{f_i(x_i)} \mathbb{R}^r \cong \mathbb{R}^r.$$

Note that, by assumption,  $(df_i)_{x_i}$  is the second component of an isometric isomorphism of Euclidean vector spaces  $d(\pi, f_i)_{x_i} : T_{x_i} \mathcal{X} \rightarrow T_{b_i} \mathcal{B} \oplus \mathbb{R}^r$ , in which the two direct summands are orthogonal. It is thus sufficient to estimate the first component  $d\pi_{x_i} : T_{x_i} \mathcal{X} \rightarrow T_{b_i} \mathcal{B}$  and to then compose  $\ell_i$  with the orthogonal projection onto the orthogonal complement of  $d\pi_{x_i}$ . We do have an estimate for the composite

$$X_i \xrightarrow{\ell_i} T_{x_i} \mathcal{X} \xrightarrow{d\pi_{x_i}} T_{b_i} \mathcal{B} \cong \mathbb{R}^e,$$

namely  $t = \Psi_i^T \circ \pi|_{X_i} : X_i \rightarrow \mathbb{R}^e$ , but, since the embedding  $\mathcal{B} \subseteq \mathbb{R}^D$  is not required to preserve the Riemannian structure of  $\mathcal{B}$  inherited from that of  $\mathcal{X}$ , the map  $t$  is an approximation of  $d\pi_{x_i} \circ \ell_i$  up to a linear map  $m_i : \mathbb{R}^d \rightarrow \mathbb{R}^e$ . This justifies finding  $m_i$  by minimizing Equation (4), and getting the approximate fiber coordinate  $f_i$  by composing  $\ell_i$  with the orthogonal projection onto the kernel of  $m_i$ .

### 3.6 Choosing input and parameters

We discuss some guiding principles to choose parameters for our pipeline. We focus mostly on parameter selection for the examples of Section 4.

**Parameters.** An estimate of the dimensions  $e$  of  $B$  and  $d$  of  $X$  can be obtained by analyzing the explained variance of PCA applied to each of the sets  $U_i$  and  $X_i$  with a range of target dimensions, but more sophisticated algorithms are available; see, e.g., [25]. The parameter  $k$  is chosen to be large enough so that the cover  $\mathcal{U}$  captures the topology of the base space  $B$ , and such that each open ball of the cover is sufficiently small so that it can be approximated

reasonable well by a linear space. Admittedly, this is in general a difficult choice and producing good covers of data is an interesting problem in its own right. In our case, when the base space is the circle, we use  $k = 16$ ; see also [45, Appendix C.5] for a parameter sensitivity analysis. The algorithm is robust to the choice of parameter `n_iter`, which we choose to be 1000 in all of our examples.

**Choosing base map and  $D$ .** We construct the initial map  $\pi : X \rightarrow \mathbb{R}^D$  in two ways.

The first way is to use the persistent cohomology of the initial data  $X$  to construct circular coordinates  $X \rightarrow S^1$  and then embed the circle  $S^1$  as the unit circle in the plane spanned by the first two coordinates of  $\mathbb{R}^D$ ,  $D \geq 2$ , which gives us the initial map  $\mathcal{X} \rightarrow \mathbb{R}^D$ . In order to choose the embedding dimension  $D$ , we compute the Stiefel–Whitney obstructions, as in Section 2.3. Since the base space is the circle, which is 1-dimensional, only the first Stiefel–Whitney class provides an obstruction. The Stiefel–Whitney class of the normal bundle of the embedding  $S^1 \subseteq \mathbb{R}^D$  is trivial. Thus, if the first Stiefel–Whitney class of the estimated cocycle  $\Omega$  is trivial, we set  $D = r + 1$ , and if it is non-trivial, we set  $D = r + 2$ .

The second way is to use the cut-unfold technique, explained in [45, Appendix C.1], with the circular coordinates and map  $\mathcal{X} \rightarrow \mathbb{R}^D$  by embedding the interval  $[0, 1)$  as the unit interval of the line spanned by the first coordinate of  $\mathbb{R}^D$ . In this case, since the interval is topologically trivial (contractible), the Stiefel–Whitney classes give no obstructions, and thus we set  $D = r + 1$ .

## 4 Examples

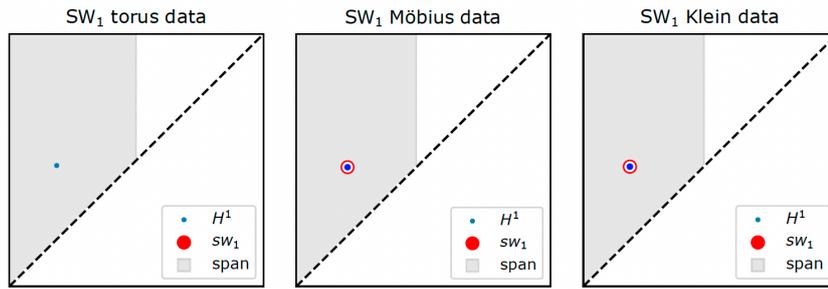
We apply FIBERED to three examples. We reproduce a dynamical system simulation from [12] and reconstruct an attractor – a torus. We reconstruct the conformation space – a Möbius band – and energy landscape of the pentane molecule from a simulation using RDKit [39]; this is inspired by an analysis in [29]. Finally, we reconstruct the conformation space of the cyclooctane molecule – a Klein bottle glued to a 2-sphere – using the data of [27].

We compare FIBERED to various well known dimensionality reduction algorithms (see [45, Appendix C.4] for more results). Given that we consider topologically non-trivial data, we follow [40, 33] and evaluate the output of algorithms using persistent homology and persistence diagrams (PDs) to quantify the preservation of large scale topology (see [45, Appendix A.1] for background and references). When we do not clarify the field of coefficients used to compute a PD, the PD is independent of this choice. See Table 1 for a summary.

For the initial map  $\pi : X \rightarrow B$  we use the implementation of circular coordinates in [54]. The parameters for FIBERED are chosen as in Section 3.6 and the computed Stiefel–Whitney obstructions are in Figure 3. For persistent homology computations, we use `ripser` [5] on geodesic distance, estimated as shortest path distance in a 15-nearest neighbor graph. For other dimensionality reduction algorithms, we use their `scikit-learn` implementation [34].

An implementation and Jupyter notebooks to reproduce the examples is in [44].

**Torus from attractor of double-gyre dynamical system.** Dynamical systems can be analyzed by studying the topology of their attractors [2, 36]. Given a real-valued time series coming from measurements of a given particle on which a dynamical system acts, one can obtain a pointcloud by constructing a delay embedding of the time series, which, under certain conditions, is concentrated around a diffeomorphic copy of the attractor the particle is converging to [36]. Using the delay embedding method with target dimension 4, it was shown in [12, Section 4.1] that a certain attractor of the double-gyre dynamical system [47]



**Figure 3** We proceed as in [45, Appendix C.3]. The persistence diagram of  $\{\mathcal{N}_r\}_{r \in [0,1]}$  for each of the three examples, with the span of  $\Omega$  shaded in grey, and the classes summing to  $w_1(\Omega)$  circled in red. Recall that in the three examples, the nerve  $\mathcal{N}$  is a circle and thus the persistence diagram consists of just one prominent 1-dimensional cohomology class. In the case of the torus, the first Stiefel–Whitney class is zero and thus there is no obstruction to choosing  $D = 1 + 2$  (1 being the dimension of the circle and 2 the rank of the vector bundle). In the case of the Möbius band, the first Stiefel–Whitney class coincides with the only point in the persistence diagram and is thus non-trivial, which gives an obstruction to selecting  $D = 1 + 1$ , which reflects the fact that the Möbius band cannot be embedded in the plane. Similarly, in the case of the Klein bottle, the Stiefel–Whitney computation gives an obstruction to selecting  $D = 1 + 2$ , which reflects the fact that the Klein bottle cannot be embedded in  $\mathbb{R}^3$ .

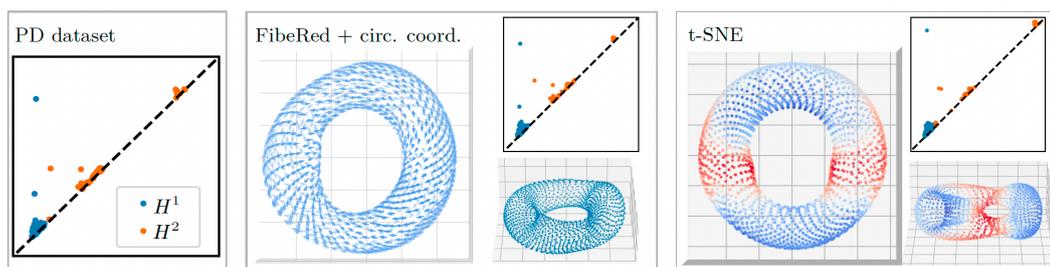
	Optimal	Isomap	t-SNE	LE/DM	LLE	HLLE	LTSA	UMAP	FibeRed
Cyl.	2	3	3	3	3	3	3	3	<b>2</b>
Torus	3	4	4	4	4	4	4	4	<b>3</b>
Möb.	3	4	4	<b>3*</b>	N/A	N/A	N/A	<b>3*</b>	<b>3</b>
Klein	4	5	5	7	7	5	5	<b>4</b>	<b>4</b>

**Table 1** The minimal target dimension that can be chosen for each of the algorithms considered in this section, so that there exist parameters that return a topologically faithful embedding of the data. “Optimal” refers to the theoretical minimal embedding dimension. “Cylinder” refers to the dataset of Figure 1. “Torus”, “Möbius band”, and “Klein bottle” refer to the three datasets considered in this section. Since the Möbius band data is not Euclidean, some algorithms cannot be run on these data; we denote this with “N/A”. Asterisks indicate that the data had to be preprocessed with MDS and target dimension 20 in order to get a topologically faithful embedding with the corresponding algorithm and dimension.

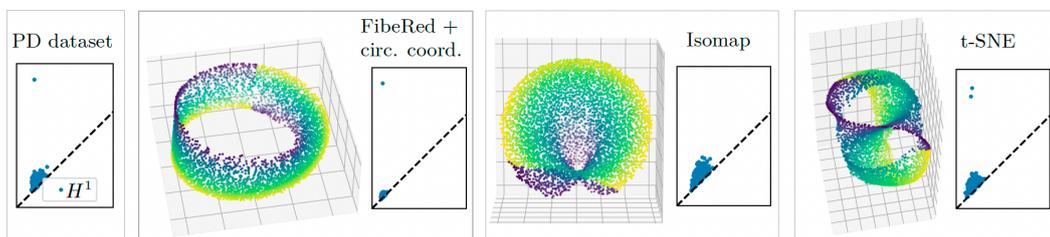
is orientable and has the homology of a torus. Here, we reproduce the simulation of [12] using the code from [16] and apply dimensionality reduction to this 4D pointcloud, with the goal of embedding the attractor and its dynamics in  $\mathbb{R}^3$ .

In Figure 4, we show the results of FIBERED and t-SNE. In order to highlight self-intersections in low-dimensional representations, we use the following function: given a dataset  $X$  and a representation of it  $f : X \rightarrow Y$  let  $\kappa : X \rightarrow \mathbb{R}$  be defined by  $\kappa(x) = \min_{y \in X} d_Y(f(x), f(y)) / d_X(x, y)$ . The output of t-SNE in Figure 4 is representative of the output with other parameter choices and other dimensionality reduction algorithms we have tried on this data (LE, DM, LLE, HLLE, Isomap, UMAP): if the target dimension is 3, there are always self-intersections or tears. The difficulty faced by metric-based algorithms in this example is that the input torus in 4D has an approximately flat metric and thus it does not admit a smooth isometric embedding in  $\mathbb{R}^3$ .

**Möbius band from conformation space of pentane.** Any fixed molecule admits different realizations, or *conformations*, in three-dimensional space. In, e.g., molecular dynamics [19],



■ **Figure 4** The PD of the original pointcloud (two prominent 1-dimensional classes, and one prominent 2-dimensional class); the output of FIBERED with the reconstructed dynamics and side view, and the PD of the output (which matches the PD of the original pointcloud well); the output of t-SNE on the same data and side view, colored by  $\kappa$  (red is smaller), there seem to be two self-intersections, and the PD of the output of t-SNE, which has one prominent 1-dimensional hole and two 2-dimensional voids, confirming that the red regions have been pinched.

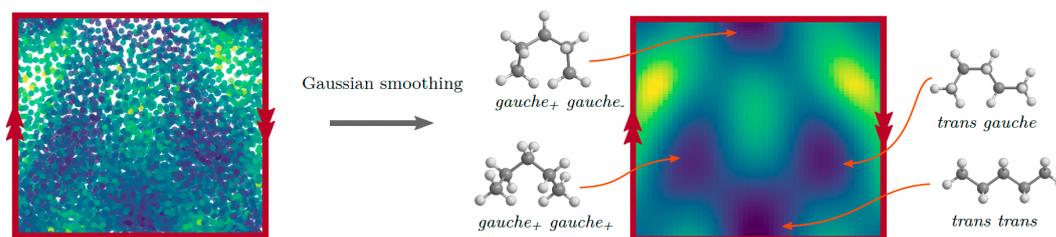


■ **Figure 5** The PD of the original pointcloud, which has one prominent 1-dimensional class; the output of FIBERED and its PD; the output of Isomap and its PD (regardless of the parameter for Isomap, the algorithm is unable to capture the circularity of the data, and thus its PD has no prominent features); the output of t-SNE and its PD (regardless of the parameters for t-SNE, the algorithm is unable to capture the circularity and non-orientability of the data without tears, which cause the output to have two holes). Outputs are colored by the (aligned) fiber coordinates estimated by FIBERED.

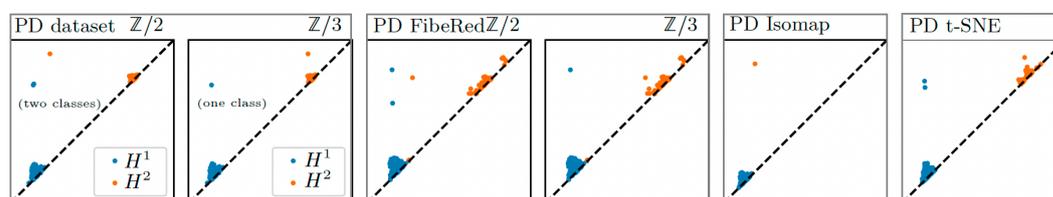
one is interested in understanding all possible conformations of a molecule. The collection of conformations up to rotations and translations is known as the *conformation space* of the molecule. Each conformation has an associated energy and the conformation space together with the energy function is known as the *energy landscape* of the molecule.

We reconstruct the conformation space and energy landscape of the pentane molecule from a simulation (see [45, Appendix C.2] for details). The pentane molecule has two rotational degrees of freedom (modelled as a torus  $S^1 \times S^1$ ) but also has a symmetry which interchanges the two angles of rotation. For this reason, the (unlabeled) conformation space consists of a quotient of the torus, which can be seen to be a Möbius band. In Figure 5, we embed the conformation space of pentane in  $\mathbb{R}^3$  and compare the output of FIBERED to that of Isomap and t-SNE. LE and DM are able to recover a Möbius band in  $\mathbb{R}^3$ ; since UMAP uses LE as initialization, it is also able to recover the Möbius band in  $\mathbb{R}^3$ . In Figure 6, we use the cut-unfold technique to find a fundamental domain of the conformation space and estimate the energy landscape.

The difficulty faced by some of the metric-based dimensionality reduction algorithms in this example is that, with respect to the intrinsic metric, the ratio between the height of the Möbius band and its circumference is approximately  $2/3$  and thus there is no isometric embedding in  $\mathbb{R}^3$  [20, Theorem 15.1].



■ **Figure 6** *Left*: the 2D representation of the conformation space of pentane using FIBERED with the cut-unfold technique, colored by energy. Arrows indicate how the data must be glued in order to recover its global topology; this information can be extracted from the cocycle  $\Omega$  of ESTCOCYCLES. *Right*: a 2D representation of the energy landscape of pentane, where the energy is estimated using the representation on the left and Gaussian smoothing. We see that there are four local minima of the energy function. By going back to the molecule simulation, we confirm that these four minima correspond to the four well known conformations of pentane [3].

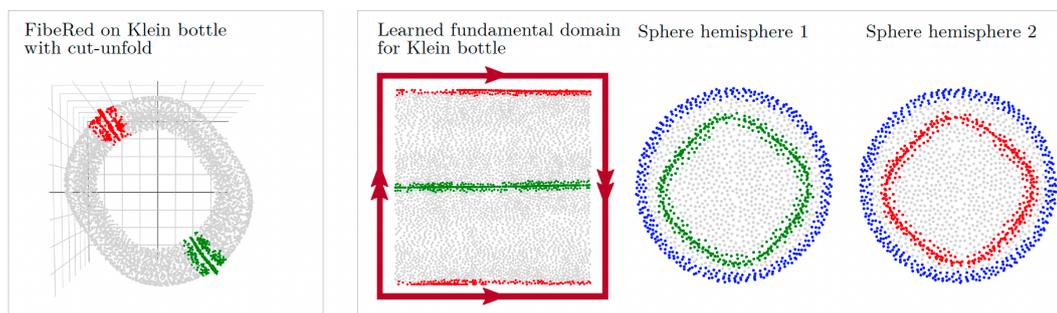


■ **Figure 7** The PD of the original data with  $\mathbb{Z}/2$  (two prominent 1-dimensional and one 2-dimensional classes) and  $\mathbb{Z}/3$  coefficients (one prominent 1-dimensional class), which suggests the data is a Klein bottle; the PD of the representation obtained using FIBERED, which matches the original topology well; the PD of a representation using Isomap; the PD of a representation using t-SNE. For Isomap and t-SNE, the PD is the same regardless of the field of coefficients.

**Klein bottle from conformation space of cyclooctane.** In this example, we reconstruct the conformation space and energy landscape of the cyclooctane molecule using the dataset of [27]. In [27], it is shown that the conformation space of cyclooctane consists of a 2-sphere glued to a Klein bottle along two disjoint circles and a parametrization of the dataset is given using Isomap and knowledge about how the data was generated.

By estimating the local dimension of the data, we first separate the Klein bottle part of the dataset from the 2-sphere. In Figure 7, we embed the Klein bottle part of the data in 4D. We were not able to recover the right topology in  $\mathbb{R}^4$  using any of LE, DM, LLE, HLLE, LTSA, Isomap, or t-SNE. Meanwhile, UMAP is able to recover the right topology in  $\mathbb{R}^4$ . In order to evaluate the 4D embeddings, we use the following distinguishing feature of the Klein bottle  $K$ : with  $\mathbb{Z}/2$  coefficients we have  $\dim(H^1(K; \mathbb{Z}/2)) = 2$  and  $\dim(H^2(K; \mathbb{Z}/2)) = 1$ , while with  $\mathbb{Z}/3$  coefficients we have  $\dim(H^1(K; \mathbb{Z}/3)) = 1$  and  $\dim(H^2(K; \mathbb{Z}/3)) = 0$ . In Figure 8, we produce an efficient 2D parametrization of the conformation space of cyclooctane without using a priori knowledge of how the data was generated.

The difficulty faced by metric-based algorithms in this example is that the Klein bottle in high dimensional space has aspect ratio close to 1 (i.e., an isometric representation by a fundamental domain such as the one Figure 8 (left) has commensurable height and width), and thus it does not admit a simple isometric embedding in  $\mathbb{R}^4$ .



■ **Figure 8** *Left*: The output of FIBERED with the cut-unfold technique on the portion of the conformation space belonging to the Klein bottle. Colored in red and green are the two circles that glue the Klein bottle to the 2-sphere. Using this representation – a cylinder – we compute a new circular coordinate, which we combine with the initial circular coordinate to get a fundamental domain for the Klein bottle. *Right*: A 2D model of the conformation space of cyclooctane. The two circles are two hemispheres of the 2D sphere and were obtained using Isomap. Points not colored in grey indicate the gluings that have to be performed to recover the conformation space.

## 5 Discussion

We have presented a procedure to learn vector bundles from data and demonstrated that it can be used to decouple the global topology from the local geometry in topologically complex data. We showed with examples that this can be helpful for embedding topologically complex data in low dimension, as well as for charting such data. We have also developed a mathematical foundation for this point of view.

**Limitations.** The theory and methods presented in this paper assume that the data lives in the total space of a vector bundle. There are two main ways in which real data can deviate from these assumptions: (1) There are singularities in the data manifold and thus the base map is not a vector bundle since fibers may have different dimensions; (2) the data contains outliers and only a core subset of the data satisfies the assumptions. Two other important caveats are that (3) the procedure assumes that a base map is given and that (4) success depends on the first step of the procedure finding a good cover of the data. We comment on these remarks below.

**Future work.** With respect to (1), the situation in which the fibers of the base map can have different dimensions can be abstracted using the theory of stratified vector bundles [7, 6]. We believe that the main algorithm of Section 3 can be adjusted to account for different local dimensions by allowing the cocycle  $\Omega$  between patches with different dimension to be a matrix in a Stiefel manifold instead of an orthogonal matrix. With respect to (2), our procedures are robust with respect to limited amount of noise and the problem of devising extensions robust to outliers is left as future work.

With respect to (3), there are several ways to obtain non-linear initial representations. First, other cohomological coordinates besides circular coordinates have been developed [35, 38]. Second, one could use standard non-linear representations, such as the ones learned by Diffusion Maps [23, 13]. Third, one could use any of the lens functions [50, Section 4] Mapper uses. Another interesting avenue for constructing coarse topological representations is to build a graph on the data, simplify it while preserving part of its large scale topology, and use a graph layout algorithm. Regarding (4), finding good covers of noisy data is an

interesting problem in itself; we believe the approach presented in this paper can be made more robust by developing a more nuanced subroutine for computing a cover.

Our approach depends on several constructions, some of which are known to be consistent estimators. Addressing the consistency of the entire pipeline is left for future work.

Finally, FIBERED can be interpreted as principal component analysis relative to an initial representation, as it works by linearly embedding the local coordinates of  $X$  that are not already accounted by the initial map, in a way that is globally consistent and orthogonal to the coordinates already accounted by the initial map. This suggests considering versions of other popular dimensionality reduction algorithms relative to an initial representation.

---

## References

- 1 Eddie Aamari, Jisu Kim, Frédéric Chazal, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Estimating the reach of a manifold. *Electronic journal of statistics*, 13(1):1359–1399, 2019.
- 2 Henry D. I. Abarbanel. *Analysis of observed chaotic data*. Institute for Nonlinear Science. Springer-Verlag, New York, 1996. doi:10.1007/978-1-4612-0763-4.
- 3 Roman M Balabin. Enthalpy difference between conformations of normal alkanes: Raman spectroscopy study of n-pentane and n-butane. *The Journal of Physical Chemistry A*, 113(6):1012–1019, 2009.
- 4 Afonso S. Bandeira, Amit Singer, and Daniel A. Spielman. A Cheeger inequality for the graph connection Laplacian. *SIAM J. Matrix Anal. Appl.*, 34(4):1611–1630, 2013. doi:10.1137/120875338.
- 5 Ulrich Bauer. Ripser: efficient computation of vietoris-rips persistence barcodes. *Journal of Applied and Computational Topology*, 2021. doi:10.1007/s41468-021-00071-5.
- 6 Hans-Joachim Baues and Davide L. Ferrario.  $K$ -theory of stratified vector bundles. *K-Theory*, 28(3):259–284, 2003. doi:10.1023/A:1026215632002.
- 7 Hans-Joachim Baues and Davide L. Ferrario. Stratified fibre bundles. *Forum Math.*, 16(6):865–902, 2004. doi:10.1515/form.2004.16.6.865.
- 8 Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- 9 Ingwer Borg and Patrick J. F. Groenen. *Modern multidimensional scaling*. Springer Series in Statistics. Springer, New York, second edition, 2005. Theory and applications.
- 10 Matthew Brand. Charting a manifold. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems*, volume 15. MIT Press, 2002. URL: <https://proceedings.neurips.cc/paper/2002/file/8929c70f8d710e412d38da624b21c3c8-Paper.pdf>.
- 11 Elie Joseph Cartan. Sur la possibilité de plonger un espace riemannien donné dans un espace euclidien. *Annales de la Société Polonaise de Mathématique*, 1928.
- 12 Gisela D. Charó, Guillermo Artana, and Denisse Sciamarella. Topology of dynamical reconstructions from Lagrangian data. *Phys. D*, 405:132371, 12, 2020. doi:10.1016/j.physd.2020.132371.
- 13 Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Appl. Comput. Harmon. Anal.*, 21(1):5–30, 2006. doi:10.1016/j.acha.2006.04.006.
- 14 Vin de Silva, Dmitry Morozov, and Mikael Vejdemo-Johansson. Persistent cohomology and circular coordinates. *Discrete Comput. Geom.*, 45(4):737–759, 2011. doi:10.1007/s00454-011-9344-x.
- 15 David L. Donoho and Carrie Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596, 2003. doi:10.1073/pnas.1031596100.
- 16 Ximena Fernández. Topology of fluid flows. [https://github.com/ximenafernandez/topology\\_fluids](https://github.com/ximenafernandez/topology_fluids), 2022.

- 17 Jesus Garcia-Diaz, Rolando Menchaca-Mendez, Ricardo Menchaca-Mendez, Saúl Pomares Hernández, Julio César Pérez-Sansalvador, and Nouredine Lakouari. Approximation algorithms for the vertex k-center problem: Survey and experimental evaluation. *IEEE Access*, 7:109228–109245, 2019. doi:10.1109/ACCESS.2019.2933875.
- 18 Robert Ghrist. Barcodes: the persistent topology of data. *Bull. Amer. Math. Soc. (N.S.)*, 45(1):61–75, 2008. doi:10.1090/S0273-0979-07-01191-3.
- 19 James M Haile. *Molecular dynamics simulation: elementary methods*. John Wiley & Sons, Inc., 1992.
- 20 B. Halpern and C. Weaver. Inverting a cylinder through isometric immersions and isometric embeddings. *Trans. Amer. Math. Soc.*, 230:41–70, 1977. doi:10.2307/1997711.
- 21 Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- 22 Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, second edition, 2013.
- 23 S. Lafon and A.B. Lee. Diffusion maps and coarse-graining: a unified framework for dimensionality reduction, graph partitioning, and data set parameterization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(9):1393–1403, 2006. doi:10.1109/TPAMI.2006.184.
- 24 John Aldo Lee and Michel Verleysen. Nonlinear dimensionality reduction of data manifolds with essential loops. *Neurocomput.*, 67:29–53, August 2005. doi:10.1016/j.neucom.2004.11.042.
- 25 Anna V Little, Jason Lee, Yoon-Mo Jung, and Mauro Maggioni. Estimation of intrinsic dimensionality of samples from noisy low-dimensional manifolds in high dimensions with multiscale svd. In *2009 IEEE/SP 15th Workshop on Statistical Signal Processing*, pages 85–88. IEEE, 2009.
- 26 Zixiang Luo, Chenyu Xu, Zhen Zhang, and Wenfei Jin. A topology-preserving dimensionality reduction method for single-cell rna-seq data using graph autoencoder. *Scientific reports*, 11(1):1–8, 2021.
- 27 Shawn Martin, Aidan Thompson, Evangelos A Coutsias, and Jean-Paul Watson. Topology of cyclo-octane energy landscape. *The journal of chemical physics*, 132(23):234115, 2010.
- 28 Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction, 2018. doi:10.48550/ARXIV.1802.03426.
- 29 Ingrid Membrillo-Solis, Mariam Pirashvili, Lee Steinberg, Jacek Brodzki, and Jeremy G. Frey. Topology and geometry of molecular conformational spaces and energy landscapes, 2019. doi:10.48550/ARXIV.1907.07770.
- 30 John W. Milnor and James D. Stasheff. *Characteristic classes*. Annals of Mathematics Studies, No. 76. Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1974.
- 31 Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. In *International conference on machine learning*, pages 7045–7054. PMLR, 2020.
- 32 Steve Y. Oudot. *Persistence theory: from quiver representations to data analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2015. doi:10.1090/surv/209.
- 33 Rahul Paul and Stephan K Chalup. A study on validating non-linear dimensionality reduction using persistent homology. *Pattern Recognition Letters*, 100:160–166, 2017.
- 34 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- 35 Jose A. Perea. Multiscale projective coordinates via persistent cohomology of sparse filtrations. *Discrete Comput. Geom.*, 59(1):175–225, 2018. doi:10.1007/s00454-017-9927-2.
- 36 Jose A. Perea. Topological times series analysis. *Notices Amer. Math. Soc.*, 66(5):686–694, 2019.
- 37 Jose A. Perea. Sparse circular coordinates via principal  $\mathbb{Z}$ -bundles. In *Topological data analysis—the Abel Symposium 2018*, volume 15 of *Abel Symp.*, pages 435–458. Springer, Cham, 2020.

- 38 Luis Polanco and Jose A. Perea. Coordinatizing data with lens spaces and persistent cohomology. In Zachary Friggstad and Jean-Lou De Carufel, editors, *Proceedings of the 31st Canadian Conference on Computational Geometry, CCCG 2019, August 8-10, 2019, University of Alberta, Edmonton, Alberta, Canada*, pages 49–58, 2019.
- 39 RDKit developers. RDKit: Open-source cheminformatics. <http://www.rdkit.org>, 2022.
- 40 Bastian Rieck and Heike Leitte. Persistent homology for the evaluation of dimensionality reduction schemes. In *Computer Graphics Forum*, volume 34, pages 431–440. Wiley Online Library, 2015.
- 41 Sam Roweis, Lawrence Saul, and Geoffrey E Hinton. Global coordination of local linear models. In T. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001. URL: <https://proceedings.neurips.cc/paper/2001/file/850af92f8d9903e7a4e0559a98ecc857-Paper.pdf>.
- 42 Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.
- 43 Luis Scoccola, Hitesh Gakhar, Johnathan Bush, Nikolas Schonsheck, Tatum Rask, Ling Zhou, and Jose A. Perea. Toroidal coordinates: Decorrelating circular coordinates with lattice reduction, 2022. doi:10.48550/ARXIV.2212.07201.
- 44 Luis Scoccola and Jose A. Perea. FibeRed implementation. <https://github.com/LuisScoccola/fibered.git>, 2022.
- 45 Luis Scoccola and Jose A. Perea. Fiberwise dimensionality reduction of topologically complex data with vector bundles, 2022. doi:10.48550/ARXIV.2206.06513.
- 46 Luis Scoccola and Jose A. Perea. Approximate and discrete Euclidean vector bundles. *Forum of Mathematics, Sigma*, 11, March 2023. doi:10.1017/fms.2023.16.
- 47 Shawn C. Shadden, Francois Lekien, and Jerrold E. Marsden. Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Phys. D*, 212(3-4):271–304, 2005. doi:10.1016/j.physd.2005.10.007.
- 48 A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.*, 30(1):20–36, 2011. doi:10.1016/j.acha.2010.02.001.
- 49 Amit Singer and Hau-tieng Wu. Orientability and diffusion maps. *Appl. Comput. Harmon. Anal.*, 31(1):44–58, 2011. doi:10.1016/j.acha.2010.10.001.
- 50 Gurjeet Singh, Facundo Mémoli, Gunnar E Carlsson, et al. Topological methods for the analysis of high dimensional data sets and 3d object recognition. *PBG Eurographics*, 2, 2007.
- 51 Jian Tang, Jingzhou Liu, Ming Zhang, and Qiaozhu Mei. Visualizing large-scale and high-dimensional data. In *Proceedings of the 25th International Conference on World Wide Web, WWW '16*, pages 287–297, Republic and Canton of Geneva, CHE, 2016. International World Wide Web Conferences Steering Committee. doi:10.1145/2872427.2883041.
- 52 Yee Teh and Sam Roweis. Automatic alignment of local representations. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems*, volume 15. MIT Press, 2002. URL: <https://proceedings.neurips.cc/paper/2002/file/3a1dd98341fafc1dfe9bcf36360e6b84-Paper.pdf>.
- 53 Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- 54 Chris Tralie, Tom Mease, and Jose Perea. Dreimac. <https://github.com/ctralie/DREiMac>, 2017.
- 55 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008.
- 56 Alexander Wagner, Elchanan Solomon, and Paul Bendich. Improving metric dimensionality reduction with distributed topology, 2021. doi:10.48550/ARXIV.2106.07613.
- 57 Hassler Whitney. The self-intersections of a smooth  $n$ -manifold in  $2n$ -space. *Ann. of Math. (2)*, 45:220–246, 1944. doi:10.2307/1969265.

## 56:18 Fiberwise Dimensionality Reduction

- 58 Lin Yan, Yaodong Zhao, Paul Rosen, Carlos Scheidegger, and Bei Wang. Homology-preserving dimensionality reduction via manifold landmarking and tearing. *arXiv preprint arXiv:1806.08460*, 2018.
- 59 Zhenyue Zhang and Hongyuan Zha. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM Journal on Scientific Computing*, pages 313–338, 2004.