



RESEARCH ARTICLE

Approximate and discrete Euclidean vector bundles

Luis Scoccola ¹ and Jose A. Perea ²

¹Department of Mathematics, Northeastern University, 43 Leon St, Boston, MA 02115, USA;

E-mail: l.scoccola@northeastern.edu

²Department of Mathematics and Khoury College of Computer Sciences, Northeastern University, 43 Leon St, Boston, MA 02115, USA; E-mail: j.pereabenitez@northeastern.edu

Received: 12 May 2021; **Revised:** 9 September 2022; **Accepted:** 7 February 2023

2020 Mathematics Subject Classification: *Primary* – 55R99, 55N31, 68W05; *Secondary* – 55U99

Abstract

We introduce ε -approximate versions of the notion of a Euclidean vector bundle for $\varepsilon \geq 0$, which recover the classical notion of a Euclidean vector bundle when $\varepsilon = 0$. In particular, we study Čech cochains with coefficients in the orthogonal group that satisfy an approximate cocycle condition. We show that ε -approximate vector bundles can be used to represent classical vector bundles when $\varepsilon > 0$ is sufficiently small. We also introduce distances between approximate vector bundles and use them to prove that sufficiently similar approximate vector bundles represent the same classical vector bundle. This gives a way of specifying vector bundles over finite simplicial complexes using a finite amount of data and also allows for some tolerance to noise when working with vector bundles in an applied setting. As an example, we prove a reconstruction theorem for vector bundles from finite samples. We give algorithms for the effective computation of low-dimensional characteristic classes of vector bundles directly from discrete and approximate representations and illustrate the usage of these algorithms with computational examples.

Contents

1	Introduction	1
2	Background	8
3	Three notions of approximate vector bundle	12
4	Relationships between the notions	15
5	Discrete approximate vector bundles	23
6	Effective computation of characteristic classes	27
7	Computational examples	35
8	Future work	39
A	Technical details	40

1. Introduction

1.1. Context and problem statement

The notion of *fiber bundle* is fundamental in mathematics and physics [33, 43, 3, 6]. Informally, a fiber bundle with fiber F consists of a continuous function $p : Y \rightarrow X$ from the *total space* Y to the *base space* X , that, locally, looks like a projection $U \times F \rightarrow U$, in the following sense: X can be covered by open sets $U \subseteq X$, each equipped with a homeomorphism $i : U \times F \rightarrow p^{-1}(U)$ such that $(p \circ i)(x, f) = x$

© The Author(s), 2023. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike licence (<https://creativecommons.org/licenses/by-nc-sa/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the same Creative Commons licence is included and the original work is properly cited. The written permission of Cambridge University Press must be obtained for commercial re-use.

for every $(x, f) \in U \times F$. In particular, $p^{-1}(x)$ is homeomorphic to F for every $x \in X$. *Vector bundles* are fiber bundles for which F is a vector space, and a key example is the tangent bundle $T\mathcal{M} \rightarrow \mathcal{M}$ of a d -dimensional differentiable manifold \mathcal{M} . The fiber of this bundle is \mathbb{R}^d , as the tangent space at each point of \mathcal{M} is d -dimensional. The Möbius band $\text{Mob} \rightarrow S^1$ is another example of a vector bundle, interpreted as a collection of one-dimensional real vector spaces that change orientation as one goes around the equatorial circle S^1 . Of particular interest are fiber bundles for which the fiber F is only allowed to ‘twist’ according to a certain group G of automorphisms of F ; these correspond to *principal G -bundles*. Vector bundles can be identified with principal G -bundles with G the general linear group, while *Euclidean vector bundles* arise when G is the orthogonal group.

Many problems in mathematics and physics can be reduced to finding sections of a fiber bundle (that is, maps $s : X \rightarrow Y$ for which $p \circ s = \text{id}_X$) satisfying certain properties. For this reason, one is interested in finding computable obstructions to the existence of certain sections and, more generally, in defining computable invariants of fiber bundles that can aid in their classification up to isomorphism.

Characteristic classes are examples of such invariants [47]. Indeed, any principal bundle determines a collection of elements in the cohomology of its base space, called the characteristic classes of the bundle. This is done in such a way that isomorphic bundles have the same characteristic classes. The *Stiefel–Whitney classes* of a vector bundle are a particular type of characteristic class, which provide obstructions to solving several geometric problems. For instance, the first Stiefel–Whitney class determines whether or not the vector bundle is orientable, while, for a differentiable manifold \mathcal{M} , the Stiefel–Whitney classes of its tangent bundle provide obstructions to embedding \mathcal{M} in \mathbb{R}^n for small n . Similarly, the *Euler class* of an oriented vector bundle is yet another characteristic class, which provides an obstruction to the existence of a nowhere vanishing section.

Part of the ubiquity of principal bundles stems from the fact that they can be defined in several, a posteriori, equivalent ways. Of particular interest to us are: (1) the definition of principal G -bundles by means of G -valued Čech cocycles, which, roughly speaking, consist of local data on the base space specifying how the fibers must be glued to reconstruct the total space, and (2) the definition of principal G -bundles by means of *classifying maps*, which are continuous functions from the base space to a certain *classifying space* BG .

Principal bundles and their characteristic classes appear also in practical applications. Many synchronization problems, in which independent, local measurements need to be assembled into a global quantity, can be interpreted as the problem of trivializing a Čech cochain of pairwise alignments. Dimensionality reduction problems, where a high-dimensional point cloud concentrated around a low-dimensional manifold needs to be represented with low distortion in a low-dimensional space, can be interpreted as an embedding problem for which estimates of the tangential characteristic classes can provide obstructions. Although it is informally clear that vector bundles are relevant for these kinds of problems, the discrete and noisy nature of the input data makes it unclear whether the data actually determine a true vector bundle and whether topological information of this bundle can be extracted from the noisy and incomplete input.

We identify two main difficulties for working with vector bundles in a practical setting. One comes from a discrete aspect: Mathematically, vector bundles are continuous entities specified, for instance, by (continuous) cocycles or by classifying maps. How can one specify arbitrary vector bundles on, say, the geometric realization of a finite simplicial complex using a finite amount of data? The other difficulty comes from the fact that, in practical applications, nothing is exact (e.g., cocycle conditions from noisy pairwise alignments), so one needs a notion of vector bundle that is robust to some degree of noise. Although the results in this paper are mainly theoretical, they are motivated by problems in applied topology. In the rest of this introduction, we describe some of these problems in more detail.

1.1.1. Synchronization and cocycles

Consider a synchronization problem with input a set V of local measurements that are pairwise aligned by elements $\{g_{ij}\}_{i,j \in V}$ of a group G . One instance of this problem arises in cryogenic electron microscopy (cryo-EM), where, broadly speaking, one seeks to reconstruct the three-dimensional (3D) shape of




Figure 1. 2D projections of an unknown 3D shape.

a molecule from several two-dimensional (2D) projections taken from unknown viewing directions [18, 28]. Here, the measurements are the 2D pictures, which are pairwise aligned by elements of the rotation group $SO(2)$. Figure 1 shows examples of the input data for this kind of problem.

Letting K be the graph with vertex set V and an edge (ij) if g_{ij} aligns i and j sufficiently well, one expects the family of elements $\{g_{ij}\}_{(ij) \in K}$ to satisfy an *approximate cocycle condition*. This means that if (ij) , (jk) and (ik) are edges of V , then $\|g_{ij}g_{jk} - g_{ik}\| < \varepsilon$ for some sufficiently small $\varepsilon > 0$. This condition indicates that the measurements can be approximately locally synchronized: In the case of cryo-EM, this implies that sets of 2D projections with very similar viewing directions can be aligned simultaneously; this is key for averaging and thus denoising images with similar viewing directions. Global synchronization is a different matter: In the case of cryo-EM, we do not expect to be able to align all 2D projections simultaneously. This is justified in [58] by observing that the pairwise alignments $\{g_{ij}\}$ approximate an $SO(2)$ -cocycle representing the tangent bundle of the 2-sphere, which is nontrivial. We will demonstrate in Section 7.4 how a data-derived estimate of the associated Euler class yields a nontrivial obstruction to globally aligning the data from Figure 1.

In the general case, if the approximate cocycle $\{g_{ij}\}$ does indeed determine a true principal G -bundle, then global synchronization corresponds to the bundle being trivial. This view of synchronization is studied in-depth in [19] for the case of flat principal G -bundles (i.e., bundles determined by a group morphism from π_1 of the base to G). However, as it is observed in [69], the bundles underlying the cryo-EM problem are not flat.

To our knowledge, the problem of determining when an arbitrary approximate $O(d)$ -valued cocycle defines a true vector bundle has not yet been addressed in the literature. This is one of the goals of this paper. Possible applications of a general theory of approximate $O(d)$ -cocycles are the usage of characteristic classes of the underlying true vector bundle for model validation, detecting nonsynchronizability in data and for guiding local synchronization and averaging algorithms.

1.1.2. Local trivializations and dimensionality reduction

Consider a point cloud $X \subseteq \mathbb{R}^D$ concentrated around a d -dimensional embedded manifold \mathcal{M} and the ensuing problem of extracting information about the tangent bundle $T\mathcal{M} \rightarrow \mathcal{M}$ from the sample X . Possible applications of this include using the tangential characteristic classes as computable obstructions for local-to-global dimensionality reduction on X .

Here is an example of such an approach. Fix $k \in \mathbb{N}$, and let N_x denote the set of the k closest neighbors of x in X . The application of principal component analysis (PCA) to each N_x with target dimension d yields an ordered orthonormal basis $\Phi(x)$ of the d -dimensional linear subspace that best approximates N_x . This method, or variations thereof, is sometimes known as *local PCA* [59] and can be used, for instance, to estimate the local dimension of the data [41].

The aforementioned process defines a function $\Phi : X \rightarrow V(d, D)$ from the data to the Stiefel manifold of orthonormal d -frames in \mathbb{R}^D , which can be interpreted as an approximate vector bundle given by local trivializations. One expects this construction to approximate the tangent bundle of the manifold \mathcal{M} . Our goal here is to formalize this intuition and, in particular, to quantify the extent to which an approximate

local trivialization induces a true vector bundle and whether topological information of the bundle can be extracted from the approximation.

We apply these ideas in Section 7.2 to the problem of distinguishing nonhomeomorphic but homotopy equivalent attractors in the double-gyre dynamical system. Indeed, we do so by combining local trivializations and data-driven estimates of the first tangential Stiefel–Whitney class. Moreover, we also demonstrate in Section 7.3 how the first two Stiefel–Whitney classes can be estimated from data to yield nontrivial obstructions to dimensionality reduction and parallelization.

1.2. Contributions

In this paper, we introduce relaxations and discretizations of the notion of vector bundle, we determine the extent to which these approximate representations determine true vector bundles and we give algorithms for the effective computation of low-dimensional characteristic classes of the true vector bundle directly from approximate and discrete representations. We run these algorithms on examples to illustrate their usage.

The notions of vector bundle we focus on correspond to Euclidean vector bundles, that is, to vector bundles endowed with a compatible, fiberwise inner product, or, equivalently, vector bundles whose structure group is an orthogonal group. To avoid cumbersome nomenclature, we drop the modifier ‘Euclidean’ in our main definitions.

Approximate notions of vector bundle.

Let $\varepsilon \geq 0$. For a topological space B with an open cover \mathcal{U} , we define the set of ε -approximate $O(d)$ -valued 1-cocycles $Z^1_\varepsilon(\mathcal{U}; O(d))$ (Definition 3.3), which we use to define the ε -approximate cohomology set $\check{H}^1_\varepsilon(\mathcal{U}; O(d))$ (Definition 3.5). We also introduce the set of ε -approximate classifying maps $\text{Maps}(B, \text{Gr}(d)^\varepsilon)$ (Definition 3.9), which consists of continuous maps with codomain a thickened Grassmannian (Section 2.1.3), and the set of ε -approximate classifying maps up to homotopy $[B, \text{Gr}(d)^\varepsilon]$. In order to relate these notions, we define the set of ε -approximate local trivializations $\mathcal{T}_\varepsilon(\mathcal{U}; d)$ (Definition 3.8). We introduce metrics on these sets, the most relevant of which being the metric d_{cl} on $\check{H}^1_\varepsilon(\mathcal{U}; O(d))$, which we use to state stability results, below.

When B is a paracompact topological space and \mathcal{U} is a countable open cover, we define maps between the sets, as follows:

$$\begin{array}{ccccc}
 Z^1_\varepsilon(\mathcal{U}; O(d)) & \begin{array}{c} \xrightarrow{\text{triv}} \\ \xleftarrow{w} \end{array} & \mathcal{T}_\varepsilon(\mathcal{U}; d) & \xrightarrow{\text{av}} & \overset{\sqrt{2}}{\text{Maps}}(B, \text{Gr}(d)^\varepsilon) \\
 \downarrow \text{cl}_1 & & & & \downarrow \text{cl}_1 \\
 \check{H}^1_\varepsilon(\mathcal{U}; O(d)) & \xrightarrow{\text{cl}} & & & \overset{\sqrt{2}}{[B, \text{Gr}(d)^\varepsilon]}
 \end{array}$$

Here, a sub- or superscript with a constant c at an arrow tip indicates that the map sends one kind of ε -approximate vector bundle to another kind of $c\varepsilon$ -approximate vector bundle. The maps triv and w are defined in Section 4.1, the map av is defined in Construction 4.15 and the map cl is defined in Theorem 4.21, which, in particular, also implies that the diagram above commutes when going from top left to bottom right. In Section 4.11, we describe a way in which triv and w are approximate inverses of each other.

By a result of Tinarrage [63, Lemma 2.1], there is a bijection $\pi_* : [B, \text{Gr}(d)^\varepsilon] \rightarrow [B, \text{Gr}(d)]$ whenever $\varepsilon \leq \sqrt{2}/2$. In particular, when $\varepsilon \leq 1/2$, the composite

$$\pi_* \circ \text{cl} : \check{H}^1_\varepsilon(\mathcal{U}; O(d)) \rightarrow [B, \text{Gr}(d)]$$

assigns a true (classical) vector bundle to every ε -approximate cohomology class. In this sense, an ε -approximate cocycle determines a true vector bundle, as long as ε is sufficiently small.

The following result, which appears as Theorem 4.21 and Theorem 4.27, says in particular that the map cl is stable.

Theorem A. *For any $\varepsilon \geq 0$, the map $\text{cl} : \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow [B, \text{Gr}(d)^{\sqrt{2}\varepsilon}]$ is independent of arbitrary choices, such as a choice of a partition of unity subordinate to \mathcal{U} or a choice of enumeration of the opens of \mathcal{U} , and is such that, if $d_{\check{H}}(\Omega, \Lambda) < \delta$ in $\check{H}_\varepsilon^1(\mathcal{U}; O(d))$, then $\text{cl}(\Omega)$ and $\text{cl}(\Lambda)$ become equal in $[B, \text{Gr}(d)^{\sqrt{2}(\varepsilon+\delta)}]$. Moreover, if $\Omega \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$ and $\varepsilon \leq \sqrt{2}/4$, then there exists $\Lambda \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$ such that $\pi_*(\text{cl}(\Omega)) = \text{cl}(\Lambda) \in [B, \text{Gr}(d)]$ and $d_{\check{H}}(\Omega, \Lambda) \leq 9\varepsilon$.*

Theorem 4.21 also addresses refinements of covers and their action on approximate cohomology classes; see also Remark 4.22 for a notion of approximate Čech cohomology independent of a cover.

Discrete approximate notions of vector bundle.

For a simplicial complex K , we introduce discrete analogues of approximate cocycles, approximate cohomology and approximate local trivializations. Most importantly, we introduce the discrete approximate cohomology set $\text{DH}_\varepsilon^1(K; O(d))$ and the set of discrete approximate local trivializations $\text{DT}_\varepsilon(K; d)$. These are useful in practice, as specifying elements of these sets requires a finite amount of data when K is finite. To highlight their simplicity, we give here the notion of discrete approximate cocycle over a simplicial complex K :

$$\text{DZ}_\varepsilon^1(K; O(d)) = \left\{ \{ \Omega_{ij} \in O(d) \}_{(ij) \in K_1} : \Omega_{ij} = \Omega_{ji}^t \text{ and } \| \Omega_{ij} \Omega_{jk} - \Omega_{ik} \| < \varepsilon \text{ for all } (ijk) \in K_2 \right\}$$

so that a discrete approximate cocycle consists of a set of orthogonal matrices indexed by the ordered 1-simplices of K , which satisfies an approximate cocycle condition. The definition of discrete approximate local trivialization is equally simple.

We motivate the introduction of these constructions by showing that, when $\varepsilon \leq 1/2$, any discrete ε -approximate cocycle and any discrete ε -approximate local trivialization represent a true vector bundle on the geometric realization of K . Moreover, we prove a completeness result (Proposition 5.7), which says that any vector bundle over a compact triangulable space can be represented by a discrete approximate cocycle on a sufficiently fine triangulation of the space.

We remark here that the map w restricts to an algorithmic map from discrete approximate local trivializations to discrete approximate cocycles (Remark 5.6).

Reconstruction from finite samples.

Building on a result of Niyogi, Smale and Weinberger [49], we prove a reconstruction result for vector bundles on compact, embedded manifolds. For readability, we give here a version of the result using big- \mathcal{O} notation and an informal notion of (ε, δ) -closeness; a formal statement with precise bounds is given in Theorem 5.15. In the statement, $\text{Gr}(d, n)$ is the Grassmannian of d -planes in \mathbb{R}^n , which we metrize using the Frobenius distance (Section 2.1.1).

Theorem B. *Let $\mathcal{M} \subseteq \mathbb{R}^N$ be a smoothly embedded compact manifold with reach $\tau > 0$, and let $f : \mathcal{M} \rightarrow \text{Gr}(d, n)$ be ℓ -Lipschitz. Let $P \subseteq \mathbb{R}^N$ be a finite set, and let $g : P \rightarrow \text{Gr}(d, n)$ be a function such that (P, g) is (ε, δ) -close to (\mathcal{M}, f) . If $\alpha \in (\mathcal{O}(\varepsilon), \tau - \mathcal{O}(\varepsilon)) \cap (0, \sqrt{2}/4\ell^{-1} - 2\delta\ell^{-1} - \varepsilon)$, then there exists a homotopy commutative diagram as follows, in which the vertical maps are homotopy equivalences:*

$$\begin{array}{ccc} \mathcal{M} & \xrightarrow{f} & \text{Gr}(n, d) \\ \downarrow & & \downarrow \\ \check{C}(P)(\alpha) & \xrightarrow{\check{C}(g)} & \text{Gr}(n, d)^{2(\ell\alpha + \ell\varepsilon + \delta)}. \end{array}$$

Moreover, the map $\check{C}(g)$ can be represented by a discrete local trivialization $\Phi \in \text{DT}_{2\ell, \varepsilon+\delta}(\check{C}(P)(\varepsilon); d)$, in the sense that $\check{C}(g) = \text{av}(\Phi)$.

Applying w to the discrete local trivialization of Theorem B, we get an approximate cocycle that can be used to compute low-dimensional characteristic classes of the true vector bundle using the algorithms of Section 6, which we describe next. The extent to which the approximate cocycle $w(\Phi)$ recovers the true vector bundle f is made precise in Lemma 5.10; see also Remark 5.17.

Effective computation of characteristic classes.

Our last main contribution is the definition of algorithms for the stable and consistent computation of the one- and two-dimensional characteristic classes of a vector bundle given by an approximate $O(d)$ -cocycle.

The algorithms are based on well-known results which say that the characteristic classes we consider are obstructions to lifting the structure group of the cocycle from an orthogonal group to certain other Lie groups. In this sense, the algorithms are classical and most of our work goes into adapting them to the approximate setting and into giving precise bounds for their stability and consistency. The following theorem is a summary of the results in Section 6. In the theorem, \check{H} denotes Čech cohomology.

Theorem C. *Let \mathcal{U} be a cover of a topological space B with the property that nonempty binary intersections are locally path connected and simply connected. There are maps*

$$\begin{aligned} \text{sw}_1 &: \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow \check{H}^1(\mathcal{U}; \mathbb{Z}/2) \text{ for } \varepsilon \leq 2; \\ \text{sw}_2 &: \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow \check{H}^2(\mathcal{U}; \mathbb{Z}/2) \text{ for } \varepsilon \leq 1; \\ \text{eu} &: \check{H}_\varepsilon^1(\mathcal{U}; SO(2)) \rightarrow \check{H}^2(\mathcal{U}; \mathbb{Z}) \text{ for } \varepsilon \leq 1. \end{aligned}$$

The map sw_1 is 2-stable, in the sense that, if $\Omega, \Omega' \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$ with $\varepsilon \leq 2$, and $d_{\check{H}}(\Omega, \Omega') < 2$, then $\text{sw}_1(\Omega) = \text{sw}_1(\Omega') \in \check{H}^1(\mathcal{U}; \mathbb{Z}/2)$. In this same sense, the maps sw_2 and eu are 1-stable.

Assume that \mathcal{U} is countable and that B is paracompact and locally contractible. The map sw_1 is 2/9-consistent, in the sense that, if $\Omega \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$ with $\varepsilon < 2/9$, then $\text{sw}_1(\Omega)$ is the first Stiefel–Whitney class of the vector bundle classified by $\pi_*(\text{cl}(\Omega)) : B \rightarrow \text{Gr}(d)$. In this same sense, the map sw_2 is 1/9-consistent and computes the second Stiefel–Whitney class, and the map eu is 1/9-consistent and computes the Euler class of an oriented vector bundle of rank 2.

We show that the maps of Theorem C are algorithmic when the input approximate cocycles are discrete approximate cocycles on a finite simplicial complex (see Tables 1–3). The time and space complexity of the algorithms is polynomial in the number of vertices of the simplicial complex. For sw_1 , the complexity is also polynomial in d , while, for sw_2 , the complexity is exponential in d and depends on calculating geodesic distances on the Spin group. In Section 6.4, we explain how to perform these calculations for small values of d .

Computational examples.

We demonstrate our algorithms on data and show how they can be combined with persistent cohomology computations to obtain cohomological as well as bundle-theoretical information of the data; this is done in Section 7. A proof-of-concept implementation of our algorithms, together with code to replicate the examples, can be found at [56].

1.3. Related work

Cohomology of synchronization problems.

Cohomological aspects of synchronization problems are discussed in [69, 19, 25].

In [19], Gao, Brodzki and Mukherjee describe a general framework for studying cohomological obstructions to synchronization problems using principal G -bundles. Although useful in many practical

applications, the approach is limited by the fact that it only encompasses so-called *flat* principal G -bundles, namely, bundles over a space B classified by group homomorphisms $\pi_1(B) \rightarrow G$. To see that this is indeed a limitation, note that no nontrivial vector bundle over S^2 can be represented by a group homomorphism $\pi_1(S^2) \rightarrow O(d)$ since $\pi_1(S^2) = 0$. As mentioned previously, vector bundles over S^2 are central in the cryo-EM synchronization problem.

The fact that the cryo-EM bundles cannot be represented by discrete and exact $SO(2)$ -cocycles on a triangulation of S^2 is observed by Ye and Lim in [69], where a solution to this problem is proposed in the special case of $SO(2)$ -cocycles over a two-dimensional simplicial complex.

Vector bundles over simplicial complexes.

In [37], Knöppel and Pinkall give a method for describing arbitrary complex line bundles over finite simplicial complexes using a finite amount of data. They also describe applications to physics and computer graphics. Their theory relies on the fact that, up to isomorphism, a complex line bundle over a simplicial complex can be described exactly using a finite amount of data consisting of an angle $\eta_{ij} \in S^1$ (encoded as a complex number of absolute value 1) for each edge (ij) of the simplicial complex as well as a real number $\Omega_{ijk} \in \mathbb{R}$ for each 2-simplex (ijk) of the simplicial complex, which satisfy $\eta_{ki}\eta_{jk}\eta_{ij} = e^{i\Omega_{ijk}}$ for every 2-simplex (ijk) , where i denotes the imaginary unit. This fact does not seem to generalize in a direct way to vector bundles of higher rank.

In [25], Hansen and Ghrist discuss synchronization on networks and describe a framework based on cellular sheaves. The framework specializes to encompass flat vector bundles over a simplicial complex. This application of their theory has the same limitation as [19]: By having the vector bundles be described by a family of matrices $\{\Omega_{ij}\}$ indexed by the edges of the simplicial complex, subject to the restriction that $\Omega_{jk}\Omega_{ij} = \Omega_{ik}$ for every 2-simplex (ijk) of the simplicial complex, only flat vector bundles can be described.

These two approaches require a certain equality to hold for each 2-simplex of the simplicial complex. The main difference between these approaches and the approach described in this paper is that we do not require an equality to hold for each 2-simplex; instead, we keep track of how much a certain equality does not hold.

Computation of characteristic classes.

In [59], Singer and Wu describe an algorithm for the consistent estimation of the orientability of an embedded manifold. This can be interpreted as estimating whether $\text{sw}_1(\Omega)$ is zero or not, where Ω is the approximate cocycle given by the transition functions of the approximate local trivialization defined using a local PCA computation, as sketched in Section 1.1.2. Their approach is robust to outliers, a property not enjoyed by the approach presented in this paper. But the algorithm does not provide the user with an actual cocycle and thus, in the case $\text{sw}_1(\Omega)$ is deemed to be nonzero by the algorithm, it is not clear how one can write it in a basis of the cohomology of a simplicial complex built from the data.

In [63], Tinarrage presents a framework for the consistent estimation of characteristic classes of vector bundles over embedded compact manifolds. The input data of the framework consist of the value of a sufficiently tame classifying map on a sufficiently dense sample of the manifold. Theoretical algorithms are given for the stable and consistent computation of arbitrary characteristic classes. The characteristic classes are computed in a geometric way, as the algorithm uses explicit triangulations of the Grassmannians and pulls back the universal characteristic classes to a simplicial complex built from the sample. The practicality of the algorithms is limited by the fact that the number of simplices required to triangulate a Grassmannian $\text{Gr}(d, n)$ is exponential in both d and n [21] and the fact that the algorithm often requires iterated subdivisions of the simplicial complex built from the data.

Vector bundles from finite samples.

As mentioned above, in [63] Tinarrage presents a framework for the consistent estimation of characteristic classes of a vector bundle, given a sample.

In [53], Rieffel addresses the problem of giving a precise correspondence between vector bundles on metric spaces X and Y that are at small Gromov–Hausdorff distance. In particular, one can use Rieffel’s framework to extend a vector bundle on a sample of a manifold to the entire manifold.

1.4. Structure of the paper

In Section 2, we give basic background. In Section 3, we present our three notions of approximate vector bundle, and in Section 4, we relate them to each other and to classical vector bundles. In particular, we prove Theorem 4.21 and Theorem 4.27. In Section 5, we introduce our notions of discrete approximate vector bundles and show that they can be used to represent vector bundles over triangulable spaces and to reconstruct vector bundles from finite samples. In Section 6, we give algorithms for the stable and consistent computation of low-dimensional characteristic classes starting from a discrete and approximate cocycle. In Section 7, we run our algorithms on examples. In Section 8, we discuss open questions.

2. Background

In this section, we introduce the basic background needed to state and prove the results in this paper. Some more technical definitions and results are in Appendix A. In Section 2.1, we introduce orthogonal groups, Grassmannians and Stiefel manifolds and fix notation. In Section 2.2, we recall some of the basics of the theory of principal bundles and vector bundles. We assume familiarity with the very basics of algebraic topology, including cohomology.

2.1. Orthogonal groups, Grassmannians and Stiefel manifolds

2.1.1. Main definitions

We start by recalling the definition of the Frobenius norm. Let $l, m \in \mathbb{N}_{\geq 1}$, and let $\mathbb{R}^{l \times m}$ denote the set of $l \times m$ matrices with real coefficients. Let $A \in \mathbb{R}^{l \times m}$. The **Frobenius norm** of A is defined by

$$\|A\| = \sqrt{\operatorname{tr}(A^t A)} = \sqrt{\sum_{1 \leq i \leq l, 1 \leq j \leq m} A_{ij}^2},$$

where $A^t \in \mathbb{R}^{m \times l}$ denotes the transpose of A . The Frobenius norm, as any norm, induces a distance on $\mathbb{R}^{l \times m}$ defined by $d_{\text{Fr}}(A, B) = \|A - B\|$. We refer to this distance as the **Frobenius distance**.

We now introduce the spaces of matrices we are most interested in. Let $n \geq d \geq 1 \in \mathbb{N}$.

The **Grassmannian** $\operatorname{Gr}(d, n)$ has as elements the $n \times n$ real matrices A that satisfy $A = A^t = A^2$ and have rank equal to d and is thus a subset of $\mathbb{R}^{n \times n}$. We metrize and topologize $\operatorname{Gr}(d, n)$ using the Frobenius distance. Note that the elements of $\operatorname{Gr}(d, n)$ are canonically identified with the orthogonal projection operators $\mathbb{R}^n \rightarrow \mathbb{R}^n$ of rank d . Since these orthogonal projections are completely determined by the subspace of \mathbb{R}^n which they span, it follows that the elements of $\operatorname{Gr}(d, n)$ correspond precisely to the d -dimensional subspaces of \mathbb{R}^n .

The (compact) **Stiefel manifold** $V(d, n)$ has as elements the set of $n \times d$ real matrices with orthonormal columns. We metrize $V(d, n)$ with the Frobenius distance. Note that the elements of $V(d, n)$ can be identified with the d -dimensional subspaces of \mathbb{R}^n equipped with an ordered orthonormal basis. The elements of a Stiefel manifold are sometimes referred to as **frames**.

When $d = n$, the Stiefel manifold $V(d, n)$ coincides with the **orthogonal group** $O(d)$, which consists of real $d \times d$ matrices Ω such that $\Omega \Omega^t = \operatorname{id}$. We metrize $O(d)$ using the Frobenius distance. With this definition, $O(d)$ is a topological group, as matrix multiplication and inversion are continuous.

We remark here that, although we shall encounter other metrics for $O(d)$ and $\operatorname{Gr}(d, n)$, our main results are stated using the Frobenius distance. The relevant results about other metrics on the orthogonal group can be found in Appendix A.4, and the ones about other metrics on the Grassmannian are in Appendix A.4.

2.1.2. Infinite-dimensional Grassmannians and Stiefel manifolds

Let $n \geq d \geq 1 \in \mathbb{N}$. There is an inclusion $\text{Gr}(d, n) \subseteq \text{Gr}(d, n + 1)$ given by adding a row and a column of zeros at the bottom and right, respectively. This inclusion is norm-preserving and thus metric-preserving. We define the infinite-dimensional Grassmannian as $\text{Gr}(d) = \bigcup_{n \in \mathbb{N}} \text{Gr}(d, n)$. We topologize $\text{Gr}(d)$ using the direct limit topology; recall that the direct limit topology on a union $X = \bigcup_{n \in \mathbb{N}} X_n$ of topological spaces $\{X_n\}_{n \in \mathbb{N}}$, where X_n is a subspace of X_{n+1} for all $n \in \mathbb{N}$, is the topology where $A \subseteq X$ is open if and only if $A \cap X_n$ is open in X_n for all $n \in \mathbb{N}$. Note that this topology is *finer* (i.e., has more opens) than the topology induced by the metric inherited by $\text{Gr}(d)$ by virtue of it being an increasing union of metric spaces.

Similarly, we have an inclusion $V(d, n) \subseteq V(d, n + 1)$ given by taking the matrix representation of a d -frame in \mathbb{R}^n and adding a row of zeros at the bottom of the matrix. Again, these inclusions are metric preserving, and we define $V(d) = \bigcup_{n \in \mathbb{N}} V(d, n)$, with the direct limit topology induced by the inclusions $V(d, n) \rightarrow V(d)$.

There is a principal $O(d)$ -bundle $\text{Proj} : V(d, n) \rightarrow \text{Gr}(d, n)$, defined by mapping a d -frame M to the matrix MM^t (see Section 2.2.2 for the notion of principal bundle). Note that the maps $\text{Proj} : V(d, n) \rightarrow \text{Gr}(d, n)$ for each $n \geq 1$ assemble into a map

$$\text{Proj} : V(d) \rightarrow \text{Gr}(d).$$

It is clear that Proj is continuous, as it restricts to a continuous map $\text{Proj} : V(d, n) \rightarrow \text{Gr}(d, n)$ for each $n \geq d$.

2.1.3. Thickenings of Grassmannians

For the general notion of thickening and some basic properties, we refer the reader to Appendix A.1; here, we briefly introduce the thickenings of Grassmannians, as they play an important role in our results.

We will be interested in thickenings of Grassmannians, and for that we need to include Grassmannians into a larger metric space. A natural candidate is to let $\text{Gr}(d, n) \subseteq \mathbb{R}^{n \times n}$, which is metric preserving if we metrize $\mathbb{R}^{n \times n}$ using the Frobenius distance. Similarly, we have $V(d, n) \subseteq \mathbb{R}^{n \times d}$. Analogously to what we did for Grassmannians and Stiefel manifolds, we define $\mathbb{R}^{\infty \times \infty} := \bigcup_{n \in \mathbb{N}} \mathbb{R}^{n \times n}$ and $\mathbb{R}^{\infty \times d} := \bigcup_{n \in \mathbb{N}} \mathbb{R}^{n \times d}$.

The elements of $\mathbb{R}^{\infty \times \infty}$ thus consist of infinite matrices with rows and columns indexed by the positive natural numbers, which have finite support, meaning that they have only finitely many nonzero entries. Similarly, the elements of $\mathbb{R}^{\infty \times d}$ are matrices with finite support, with d columns and rows indexed by the positive natural numbers. Again, although $\mathbb{R}^{\infty \times \infty}$ and $\mathbb{R}^{\infty \times d}$ inherit natural metrics, we use instead the direct limit topologies induced by inclusions $\mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{\infty \times \infty}$ and $\mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{\infty \times d}$, respectively.

Let $\varepsilon > 0$. The ε -**thickening** of $\text{Gr}(d, n)$, denoted $\text{Gr}(d, n)^\varepsilon \subseteq \mathbb{R}^{n \times n}$, consists of all matrices in $\mathbb{R}^{n \times n}$ at Frobenius distance strictly less than ε from a matrix in $\text{Gr}(d, n)$. Similarly, we define $\text{Gr}(d)^\varepsilon \subseteq \mathbb{R}^{\infty \times \infty}$. Clearly, we have $\text{Gr}(d, n)^\varepsilon \subseteq \text{Gr}(d, n + 1)^\varepsilon \subseteq \text{Gr}(d)^\varepsilon$. The ε -thickenings $\text{Gr}(d, n)^\varepsilon \subseteq \mathbb{R}^{n \times n}$ and $\text{Gr}(d)^\varepsilon \subseteq \mathbb{R}^{\infty \times \infty}$ are open subsets. If $\varepsilon = 0$, it is convenient to define $\text{Gr}(d, n)^\varepsilon = \text{Gr}(d, n)$ and $\text{Gr}(d)^\varepsilon = \text{Gr}(d)$.

The main result about thickenings of Grassmannians we will need is Proposition A.12. The result follows directly from a result of Tinarrage (Lemma A.11) and says that, if $\varepsilon \leq \sqrt{2}/2$, then $\text{Gr}(d, n)^\varepsilon$ deformation-retracts onto $\text{Gr}(d, n)$.

2.2. Principal bundles and vector bundles

We recall the main notions and results that are relevant to this article. For a thorough exposition, we refer the reader to, for example, [38, Appendices A and B], [61], [34] and [33]. The classic book [47] contains all of the standard results on vector bundles we need; see also [27] for a modern exposition.

2.2.1. Covers and nerve

A **cover** $\mathcal{U} = \{U_i\}_{i \in I}$ of a topological space B consists of an indexing set I together with, for every $i \in I$, an open subset $U_i \subseteq B$ such that $B = \bigcup_{i \in I} U_i$. The **nerve** of \mathcal{U} , denoted $N(\mathcal{U})$, is the simplicial complex

with underlying set I and simplices consisting of finite, nonempty subsets $J \subseteq I$ such that $\bigcap_{j \in J} U_j \neq \emptyset$. An **ordered simplex** of $N(\mathcal{U})$ consists of a list $(i_0 i_1 \dots i_n)$ such that $U_{i_n} \cap U_{i_{n-1}} \cap \dots \cap U_{i_0} \neq \emptyset$. When quantifying over ordered 1-simplices of $N(\mathcal{U})$, we will write $(ij) \in N(\mathcal{U})$, and, similarly, when quantifying over ordered 2-simplices of $N(\mathcal{U})$ we will write $(ijk) \in N(\mathcal{U})$.

2.2.2. Principal bundles

Let B and F be topological spaces. A **fiber bundle** over B with fiber F consists of a continuous map $p : E \rightarrow B$ such that, for every $y \in B$, there exists an open neighborhood $U \subseteq B$ of y and a homeomorphism $\rho_U : U \times F \rightarrow p^{-1}(U)$ such that $p \circ \rho_U = \pi_1 : U \times F \rightarrow U$, where π_1 denotes projection onto the first factor. Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a cover of B . A collection of maps $\{\rho_i : U_i \times F \rightarrow p^{-1}(U_i)\}_{i \in I}$ with the property above is referred to as a **local trivialization** of p . By definition, any fiber bundle admits some local trivialization.

Let G be a topological group. A **principal G -bundle** over B consists of a fiber bundle $p : E \rightarrow B$ with fiber G together with a continuous, fiberwise right action $-\cdot- : E \times G \rightarrow E$ such that there exists a cover $\mathcal{U} = \{U_i\}_{i \in I}$ and a local trivialization $\{\rho_i : U_i \times G \rightarrow p^{-1}(U_i)\}_{i \in I}$ that is equivariant, meaning that, for every $g, h \in G$ and $y \in U_i$, we have $\rho_i(y, gh) = \rho_i(y, g) \cdot h \in E$.

Two principal G -bundles $p : E \rightarrow B$ and $p' : E' \rightarrow B$ are **isomorphic** if there exists a G -equivariant map $m : E \rightarrow E'$ such that $p' \circ m = p$. Denote by $\text{Prin}_G(B)$ the set of isomorphism classes of principal G -bundles over B (this is a set and not a proper class).

We remark that the definitions we have given are sometimes referred to as *locally trivial fiber bundle* and *locally trivial principal bundle*.

2.2.3. Čech cocycles

Let B be a topological space, G a topological group and \mathcal{U} a cover of B . A **Čech 1-cocycle** subordinate to \mathcal{U} with coefficients in G consists of a family of continuous maps $\{\rho_{ij} : U_j \cap U_i \rightarrow G\}_{(ij) \in N(\mathcal{U})}$ indexed by the ordered 1-simplices of $N(\mathcal{U})$, which satisfies the **cocycle condition**, meaning that for every $(ijk) \in N(\mathcal{U})$ and $y \in U_k \cap U_j \cap U_i$, we have

$$\rho_{ij}(y)\rho_{jk}(y) = \rho_{ik}(y) \in G.$$

We remark that we are using a convention for defining ρ_{ij} so that ρ_{ij} and ρ_{jk} can be composed from left to right. The opposite convention is also common. The set of 1-cocycles subordinate to \mathcal{U} with coefficients in G is denoted by $Z^1(\mathcal{U}; G)$.

The following construction associates a cocycle to any principal G -bundle and motivates the notion of cocycle. Given a principal G -bundle over B , there exists, by definition, a cover $\mathcal{U} = \{U_i\}_{i \in I}$ and a family of G -equivariant local trivializations $\{\rho_i : U_i \times G \rightarrow p^{-1}(U_i)\}_{i \in I}$. Note that, whenever $i, j \in I$ are such that $U_j \cap U_i \neq \emptyset$, the map $\rho_j^{-1} \circ \rho_i : (U_j \cap U_i) \times G \rightarrow (U_j \cap U_i) \times G$ is G -equivariant. It follows that $\rho_j^{-1} \circ \rho_i$ induces a continuous map $\rho_{ij} : U_j \cap U_i \rightarrow G$ satisfying $(\rho_j^{-1} \circ \rho_i)(y, g) = (y, g \cdot \rho_{ij}(y))$ for all $(y, g) \in (U_j \cap U_i) \times G$, and that the family $\{\rho_{ij} : U_j \cap U_i \rightarrow G\}_{(ij) \in N(\mathcal{U})}$ satisfies the cocycle condition.

2.2.4. Čech cohomology

Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a cover of B . A **Čech 0-cochain** subordinate to \mathcal{U} with values in G consists of a family of continuous maps $\Theta = \{\Theta_i : U_i \rightarrow G\}_{i \in I}$. Let $C^0(\mathcal{U}; G)$ denote the set of 0-cochains subordinate to \mathcal{U} with values in G . There is an action $C^0(\mathcal{U}; G) \curvearrowright Z^1(\mathcal{U}; G)$ with $\Theta = \{\Theta_i : U_i \rightarrow G\}_{i \in I}$ acting on a cocycle $\rho = \{\rho_{ij} : U_j \cap U_i \rightarrow G\}_{(ij) \in N(\mathcal{U})}$ by $\Theta \cdot \rho = \{\Theta_i \rho_{ij} \Theta_j^{-1} : U_j \cap U_i \rightarrow G\}_{(ij) \in N(\mathcal{U})}$. The quotient of $Z^1(\mathcal{U}; G)$ by the action of $C^0(\mathcal{U}; G)$ is denoted by $\check{H}^1(\mathcal{U}; G)$.

Let $\mathcal{U} = \{U_i\}_{i \in I}$ and $\mathcal{V} = \{V_j\}_{j \in J}$ be two covers of a common topological space B . A **refinement** $\nu : \mathcal{U} \rightarrow \mathcal{V}$ consists of a function $\nu : I \rightarrow J$ such that for every $i \in I$ we have $U_i \subseteq V_{\nu(i)}$. Given a refinement $\nu : \mathcal{U} \rightarrow \mathcal{V}$ and $\rho \in Z^1(\mathcal{V}; G)$, define a cocycle $\nu(\rho) \in Z^1(\mathcal{U}; G)$ by $\nu(\rho)_{ij} = \rho_{\nu(i)\nu(j)}$. It is

not hard to check that any two refinements $\nu, \nu' : \mathcal{U} \rightarrow \mathcal{V}$ induce the same map $\check{H}^1(\mathcal{V}; G) \rightarrow \check{H}^1(\mathcal{U}; G)$. One then defines the **Čech cohomology** of B with coefficients in G as

$$\check{H}^1(B; G) = \operatorname{colim}_{\mathcal{U}} \check{H}^1(\mathcal{U}; G),$$

where the colimit is indexed by the poset whose objects are covers of B and where $\mathcal{V} \leq \mathcal{U}$ if there exists a refinement $\mathcal{U} \rightarrow \mathcal{V}$.

As we saw previously, any principal G -bundle over B can be trivialized over some cover of B and induces a cocycle over that cover. It is well known, and easy to see, that this construction induces a bijection

$$\operatorname{Prin}_G(B) \rightarrow \check{H}^1(B; G).$$

For a description of principal G -bundles from this point of view, see [38, Appendix A].

2.2.5. Vector bundles

Let $d \in \mathbb{N}_{\geq 1}$. A **vector bundle** of rank d over B consists of a fiber bundle $p : E \rightarrow B$ with fiber \mathbb{R}^d , where each fiber comes with the structure of a real vector space of dimension d and such that p admits a local trivialization that is linear on each fiber.

It follows that $\rho_j^{-1} \circ \rho_i$ induces a well-defined continuous map $\rho_{ij} : U_j \cap U_i \rightarrow GL(d)$ satisfying $(\rho_j^{-1} \circ \rho_i)(y, g) = (y, g \cdot \rho_{ij}(y))$ for all $(y, g) \in (U_j \cap U_i) \times \mathbb{R}^d$. So every vector bundle of rank d over B that trivializes over a cover \mathcal{U} gives a cocycle in $Z^1(\mathcal{U}; GL(d))$.

An isomorphism between vector bundles $p : E \rightarrow B$ and $p' : E' \rightarrow B$ consists of a fiberwise map $m : E \rightarrow E'$ that is a linear isomorphism on each fiber. The family of isomorphism classes of rank- d vector bundles over B is denoted by $\operatorname{Vect}_d(B)$.

A partition of unity argument shows that, if B is paracompact (in the sense of [47, Section 5.8]), the trivialization of a vector bundle on B can be taken so that the associated cocycle takes values in the orthogonal group. When B is paracompact, this gives a bijection

$$\operatorname{Vect}_d(B) \rightarrow \check{H}^1(B; O(d)).$$

In particular, $\operatorname{Vect}_d(B) \cong \check{H}^1(B; O(d)) \cong \operatorname{Prin}_{O(d)}(B)$.

2.2.6. Classifying maps

For details about the claims in this section, see [45, Theorem 3.1] for the existence of classifying spaces of topological groups and [38, Appendix B] for an account of classifying spaces of Lie groups.

For every topological group G , there exists a **classifying space** $\mathcal{B}G$, which consists of a topological space with the property that, for every paracompact topological space B , there is a natural bijection

$$\operatorname{Prin}_G(B) \cong [B, \mathcal{B}G].$$

The classifying space of the orthogonal group $O(d)$ is the Grassmannian $\operatorname{Gr}(d)$, and thus

$$\operatorname{Vect}_d(B) \cong \check{H}^1(B; O(d)) \cong \operatorname{Prin}_{O(d)}(B) \cong [B, \operatorname{Gr}(d)].$$

The map $[B, \operatorname{Gr}(d)] \rightarrow \operatorname{Prin}_{O(d)}(B)$ is constructed as follows. Given $[f] \in [B, \operatorname{Gr}(d)]$, let $f : B \rightarrow \operatorname{Gr}(d)$ be a representative. Then, the pullback of $\operatorname{Proj} : \mathcal{V}(d) \rightarrow \operatorname{Gr}(d)$ along f is a principal $O(d)$ -bundle over B , whose isomorphism type is independent of the choice of representative for $[f]$.

2.2.7. Characteristic classes

Let G be a topological group, B a paracompact topological space and p be a principal G -bundle over B . Fix an abelian group A , and let $n \in \mathbb{N}$. The bundle p can be represented by a continuous map $f : B \rightarrow \mathcal{B}G$, which can then be used to pull back any cohomology class $x \in H^n(\mathcal{B}G; A)$ to a class

$f^*(x) \in H^n(B; A)$. The cohomology classes $f^*(x)$ obtained in this way are the **characteristic classes** of the principal G -bundle p , and they are invariant under isomorphism of principal G -bundles. For a presentation of characteristic classes from this point of view, see [38, Appendix B].

The cohomology ring $H^*(\text{Gr}(d); \mathbb{Z}/2)$ is isomorphic to a polynomial ring $(\mathbb{Z}/2)[\sigma_1, \dots, \sigma_d]$ with $\sigma_i \in H^i(\text{Gr}(d); \mathbb{Z}/2)$ [47, Section 7]. For $p : E \rightarrow B$ a rank- d vector bundle over a paracompact topological space B , one defines the i th **Stiefel–Whitney class** of p as the characteristic class corresponding to $\sigma_i \in H^i(\text{Gr}(d); \mathbb{Z}/2)$.

For any even $d \geq 1$, there is a distinguished element $e_d \in H^d(\text{BSO}(d); \mathbb{Z})$, the *universal Euler class* [47, Section 9],[17, Theorem 1]. For $p : E \rightarrow B$ an oriented, rank- d vector bundle over a paracompact topological space B , which corresponds, up to isomorphism, to an principal $SO(d)$ -bundle, one defines the **Euler class** of p as the characteristic class corresponding to $e_d \in H^d(\text{BSO}(d); \mathbb{Z})$. We remark that the Euler class can be defined for odd d too, but in this case it is often less useful as, for odd d , we have $2e_d = 0$ [47, Property 9.4].

3. Three notions of approximate vector bundle

In this section, we introduce relaxations of three standard definitions of vector bundle. The base space of our bundles will be denoted by B and a typical element will usually be denoted by $y \in B$.

The classical notions of vector bundle that we consider are those of a vector bundle given by an $O(d)$ -valued Čech cocycle, a vector bundle given by a family of compatible maps from opens of a cover of B to the Stiefel manifold $V(d)$, which we interpret as a local trivialization; and a vector bundle given by a continuous map from B to the Grassmannian $\text{Gr}(d)$.

The reader may be more familiar with the notion of vector bundle given by a $GL(d)$ -valued cocycle. Being able to lift the structure group from $GL(d)$ to $O(d)$ corresponds to endowing the vector bundle with a compatible, fiberwise inner product. Vector bundles endowed with this extra structure are sometimes referred to as *Euclidean vector bundles*. Any vector bundle over a paracompact space can be endowed with an inner product in an essentially unique way [47, Problems 2-C and 2-E], and thus is isomorphic to the underlying vector bundle of a Euclidean vector bundle. The main reason for working with Euclidean vector bundles is that, for many choices of distance on the orthogonal group $O(d)$, the group acts on itself by isometries. This is not the case for $GL(d)$.

3.1. Approximate cocycles

The notion of approximate cocycle makes sense for any metric group, which we define next. This extra generality makes some arguments clearer and will be of use in Section 6.1.

Definition 3.1. A **metric group** consists of a group G endowed with the structure of a metric space such that left and right multiplication by any fixed element $g \in G$, and taking inverse are all isometries $G \rightarrow G$.

The main example of metric group to keep in mind is that of the orthogonal group $O(d)$ endowed with the Frobenius distance. Other relevant examples include all connected Lie groups endowed with the geodesic distance induced by a bi-invariant Riemannian metric; recall that all compact Lie groups, such as the orthogonal groups, the unitary groups and the compact symplectic groups, admit a bi-invariant Riemannian metric (see, e.g., [46, Corollary 1.4]). It is also relevant to note that the Frobenius distance on $O(d)$ does not arise as the geodesic distance induced by a Riemannian metric; Section A.4 deals with some of the relationships between the Frobenius distance and the (usual) geodesic distance on $O(d)$.

Let G be a metric group and let us denote its distance by $d_G : G \times G \rightarrow \mathbb{R}$. Let B be a topological space and let $\mathcal{U} = \{U_i\}_{i \in I}$ be a cover of B .

Definition 3.2. A **1-cochain** on B subordinate to \mathcal{U} with values in G consists of a family of continuous maps $\Omega = \{\Omega_{ij} : U_j \cap U_i \rightarrow G\}_{(ij) \in N(\mathcal{U})}$ indexed by the ordered 1-simplices of $N(\mathcal{U})$ that is **symmetric**, that is, such that for all $(ij) \in N(\mathcal{U})$ and $y \in U_j \cap U_i$ we have $\Omega_{ij}(y) = \Omega_{ji}(y)^{-1}$.

We denote the set of all 1-cochains on B subordinate to \mathcal{U} with values in G by $C^1(\mathcal{U}; G)$. If there is no risk of confusion, we may refer to an element of $C^1(\mathcal{U}; G)$ simply as a cochain subordinate to \mathcal{U} .

Definition 3.3. Let $\varepsilon \in (0, \infty]$. A cochain Ω subordinate to \mathcal{U} is an ε -**approximate cocycle** if for every $(ijk) \in N(\mathcal{U})$ and every $y \in U_k \cap U_j \cap U_i$ we have $d_G(\Omega_{ij}(y)\Omega_{jk}(y), \Omega_{ik}(y)) < \varepsilon$, and an **exact cocycle** if we have $\Omega_{ij}(y)\Omega_{jk}(y) = \Omega_{ik}(y)$.

We denote the set of ε -approximate cocycles by $Z^1_\varepsilon(\mathcal{U}; G) \subseteq C^1(\mathcal{U}; G)$, and the set of exact cocycles by either $Z^1_0(\mathcal{U}; G)$ or simply $Z^1(\mathcal{U}; G)$. Of course, when $\varepsilon = \infty$, ε -approximate cocycles are merely cochains, so to keep notation uniform, we denote $C^1(\mathcal{U}; G)$ by $Z^1_\infty(\mathcal{U}; G)$. We endow the set $Z^1_\infty(\mathcal{U}; G)$ with the metric given by

$$d_Z(\Omega, \Lambda) := \sup_{(ij) \in N(\mathcal{U})} \sup_{y \in U_j \cap U_i} d_G(\Omega_{ij}(y), \Lambda_{ij}(y)),$$

for $\Omega, \Lambda \in Z^1_\infty(\mathcal{U}; G)$. This induces a metric on all spaces of approximate and exact cocycles. In particular, for $\varepsilon \leq \varepsilon'$, we have metric embeddings

$$Z^1(\mathcal{U}; G) \subseteq Z^1_\varepsilon(\mathcal{U}; G) \subseteq Z^1_{\varepsilon'}(\mathcal{U}; G) \subseteq Z^1_\infty(\mathcal{U}; G).$$

Definition 3.4. A **0-cochain** subordinate to \mathcal{U} with values in G consists of a family of continuous maps $\Theta = \{\Theta_i : U_i \rightarrow G\}_{i \in I}$.

We denote the set of all 0-cochains by $C^0(\mathcal{U}; G)$. The set $C^0(\mathcal{U}; G)$ forms a group, by pointwise multiplication. There is an action $C^0(\mathcal{U}; G) \curvearrowright Z^1_\infty(\mathcal{U}; G)$ with Θ acting on Ω by

$$(\Theta \cdot \Omega)_{ij}(y) = \Theta_i(y) \Omega_{ij}(y) \Theta_j^{-1}(y)$$

for every $y \in U_j \cap U_i$. Since G acts on itself by isometries, the above action restricts to an action on $Z^1_\varepsilon(\mathcal{U}; G)$ for every $\varepsilon \geq 0$.

Definition 3.5. Let $\varepsilon \in [0, \infty]$. Define the ε -**approximate cohomology set** $\check{H}^1_\varepsilon(\mathcal{U}; G)$ as the quotient of $Z^1_\varepsilon(\mathcal{U}; G)$ by the action of $C^0(\mathcal{U}; G)$.

Notation 3.6. We denote a typical element of $\check{H}^1_\varepsilon(\mathcal{U}; G)$ by Ω , or by $[\Omega]$ if we want to refer to the equivalence class of an approximate cocycle $\Omega \in Z^1_\varepsilon(\mathcal{U}; G)$. In the latter case, we say that the approximate cocycle $\Omega \in Z^1_\varepsilon(\mathcal{U}; G)$ is a **representative** of the approximate cohomology class $[\Omega] \in \check{H}^1_\varepsilon(\mathcal{U}; G)$.

Since the action $C^0(\mathcal{U}; G) \curvearrowright Z^1_\infty(\mathcal{U}; G)$ is by isometries, the set $\check{H}^1_\varepsilon(\mathcal{U}; G)$ inherits a metric $d_{\check{H}}$ from $Z^1_\varepsilon(\mathcal{U}; G)$, given as follows. For $\Omega, \Lambda \in \check{H}^1_\varepsilon(\mathcal{U}; G)$, we have

$$d_{\check{H}}(\Omega, \Lambda) := \inf_{\bar{\Omega}, \bar{\Lambda} \in Z^1_\varepsilon(\mathcal{U}; O(d))} d_Z(\bar{\Omega}, \bar{\Lambda}),$$

where $\bar{\Omega}$ and $\bar{\Lambda}$ range over all representatives of Ω and Λ , respectively.

Remark 3.7. Although we will not use this in what follows, we remark that, for \mathcal{U} a cover of a topological space and G a metric group, we have constructed a filtered (or persistent) metric space $\check{H}^1_\varepsilon(\mathcal{U}; G)$, parametrized by $\varepsilon \in [0, \infty]$, which one can interpret as a functor $\check{H}^1_\bullet(\mathcal{U}; G) : ([0, \infty], \leq) \rightarrow \mathbf{Met}$.

3.2. Approximate local trivializations

We start with some considerations about the notion of local trivialization of a vector bundle (Section 2.2.5), which motivate our notion of approximate local trivialization. Any rank- d Euclidean vector bundle $p : E \rightarrow B$ admits an isometric local trivialization, that is, a local trivialization over an open cover $\mathcal{U} = \{U_i\}_{i \in I}$ of B given by a family of homeomorphisms $\{U_i \times \mathbb{R}^d \rightarrow p^{-1}(U_i)\}_{i \in I}$ with

the property that, given $y \in U_i$, we obtain an orthonormal basis of the vector space $p^{-1}(y)$ by evaluating the map $U_i \times \mathbb{R}^d \rightarrow p^{-1}(U_i)$ on (y, e_j) for $1 \leq j \leq d$, where e_j is the j th canonical basis vector of \mathbb{R}^d . Recall that any vector bundle over a paracompact base B is classified by some continuous map $B \rightarrow \text{Gr}(d)$ (Section 2.2.6). This implies, in particular, that, up to isomorphism of vector bundles, any vector bundle $p : E \rightarrow B$ over a paracompact base is a Euclidean vector bundle whose fibers $p^{-1}(y)$ that are not just abstract vector spaces, but d -dimensional subspaces of \mathbb{R}^∞ .

With the above in mind, any isometric local trivialization of a Euclidean vector bundle over a paracompact base B gives us maps $\{\Phi_i : U_i \rightarrow V(d)\}$, where, as in Section 2.1.1, the space $V(d)$ stands for the Stiefel manifold of d -frames in \mathbb{R}^∞ . One can check that a family of maps $\{\Phi_i : U_i \rightarrow V(d)\}_{i \in I}$ comes from a vector bundle over B precisely when, for every intersection $U_j \cap U_i \neq \emptyset$ and every $y \in U_j \cap U_i$, we have that $\Phi_i(y)$ and $\Phi_j(y)$ span the same subspace of \mathbb{R}^∞ and thus, equivalently, when there exist continuous maps $\{\Omega_{ij} : U_j \cap U_i \rightarrow O(d)\}_{(ij) \in N(\mathcal{U})}$ such that $\Phi_i(y)\Omega_{ij}(y) = \Phi_j(y)$ for all $y \in U_j \cap U_i$. Our notion of approximate local trivialization is based on this last equivalent characterization of local trivializations and, specifically, on relaxing this last equality.

Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a cover of a topological space B .

Definition 3.8. Let $\varepsilon \in (0, \infty]$. An ε -**approximate local trivialization** subordinate to \mathcal{U} consists of a family of continuous maps $\Phi = \{\Phi_i : U_i \rightarrow V(d)\}_{i \in I}$ such that, for every $(ij) \in N(\mathcal{U})$, there exists a continuous map $\Omega_{ij} : U_j \cap U_i \rightarrow O(d)$ such that, for every $y \in U_j \cap U_i$, we have $\|\Phi_i(y)\Omega_{ij}(y) - \Phi_j(y)\| < \varepsilon$. An **exact local trivialization** is an approximate local trivialization for which $\Phi_i(y)\Omega_{ij}(y) = \Phi_j(y)$.

We say that the family $\Omega = \{\Omega_{ij}\}_{(ij) \in N(\mathcal{U})}$ in Definition 3.8 is a **witness** of the fact that Φ is an ε -approximate (or exact) local trivialization. We remark that this witness is not part of the data of an ε -approximate local trivialization and that we merely require that a witness exists.

An **approximate local trivialization** consists of an ε -approximate local trivialization for some $\varepsilon \in [0, \infty]$. We denote the set of ε -approximate local trivializations subordinate to \mathcal{U} by $T_\varepsilon(\mathcal{U}; d)$. We define a metric on $T_\varepsilon(\mathcal{U}; d)$ by

$$d_T(\Phi, \Psi) := \sup_{i \in I} \sup_{y \in U_i} \|\Phi_i(y) - \Psi_i(y)\|.$$

An approximate local trivialization $\{\Phi_i\}_{i \in I}$ is **nondegenerate** if, for every $(ij) \in N(\mathcal{U})$ and every $y \in U_j \cap U_i$, the $d \times d$ matrix $\Phi_i(y)^t \Phi_j(y)$ has full rank.

3.3. Approximate classifying maps

Recall from Section 2.1.3 the definition of the thickened Grassmannians, and recall, in particular, that the topology of $\text{Gr}(d)^\varepsilon$ is the direct limit topology and not the one induced by the Frobenius metric. Let B be a topological space.

Definition 3.9. An ε -**approximate classifying map** consists of a continuous map $B \rightarrow \text{Gr}(d)^\varepsilon$.

The set of ε -approximate classifying maps is denoted by $\text{Maps}(B, \text{Gr}(d)^\varepsilon)$. We define a metric on this set by

$$d_C(f, g) = \sup_{y \in B} \|f(y) - g(y)\|.$$

We are also interested in the set of classifying maps up to homotopy, which we denote by $[B, \text{Gr}(d)^\varepsilon]$. Although we will not make use of this fact, we mention that one can interpret this as a persistent set $[B, \text{Gr}(d)^\varepsilon] : ([0, \infty], \leq) \rightarrow \mathbf{Set}$, in the sense of [15, Definition 2.2]. The following result is easily proven using a linear homotopy.

Lemma 3.10. *Let $f, g \in \text{Maps}(B, \text{Gr}(d)^\varepsilon)$, and let $[f]$ and $[g]$ denote their images in $[B, \text{Gr}(d)^\varepsilon]$. Let $\delta > 0$. If $\|f(y) - g(y)\| < \delta$ for all $y \in B$, then $[f]$ and $[g]$ become equal in $[B, \text{Gr}(d)^{\varepsilon+\delta}]$. In particular, if $d_C(f, g) < \delta$, then $[f]$ and $[g]$ become equal in $[B, \text{Gr}(d)^{\varepsilon+\delta}]$.*

4. Relationships between the notions

We now consider the problem of going back and forth between the different notions of approximate vector bundle. As a consequence of this study, we relate approximate vector bundles to true (exact) vector bundles. In particular, this lets us extract a true vector bundle from an ε -approximate vector bundle when ε is sufficiently small.

There are two main results in this section. Theorem 4.21 associates an approximate classifying map to any approximate cocycle and lets us, in particular, assign a true vector bundle to any ε -approximate cocycle as long as $\varepsilon \leq 1/2$. This is done in a way that is stable and independent of arbitrary choices. Theorem 4.27 gives an upper bound for the distance from an ε -approximate cocycle to an exact cocycle representing the same true vector bundle when $\varepsilon \leq \sqrt{2}/4$. This is used in Section 6 to prove the consistency of algorithms to compute characteristic classes.

Many proofs in this section rely on various results stated and proven in Appendix A.

4.1. Cocycles and local trivializations

In this section, we relate approximate cocycles and approximate local trivializations. We give constructions (Construction 4.2 and Construction 4.8) to go back and forth between the notions, and we show that, in a sense, these constructions are approximate inverses of each other (Remark 4.11). The construction to go from approximate local trivializations to approximate cocycles is in general not canonical; we conclude the section by showing that, when $\varepsilon \leq 1$, the construction can be made canonical.

To motivate the assumptions made in the following construction, recall that a set I is **countable** if there exists an injection $\iota : I \rightarrow \mathbb{N}_{\geq 1}$. Recall also that any vector bundle on a paracompact topological space can be trivialized on a countable open cover [47, Lemma 5.9], and that every open cover of a paracompact topological space admits a subordinate partition of unity.

We start with a simplification. Given $\mathcal{V} = \{V_i\}_{i \in I}$ a countable open cover of a topological space B and an injection $\iota : I \rightarrow \mathbb{N}_{\geq 1}$, consider a new open cover $\mathcal{U} = \iota_*(\mathcal{V})$ of B indexed by $\mathbb{N}_{\geq 1}$ with $U_n = V_i$ if $\iota(i) = n$ or $U_n = \emptyset$ if there is no $i \in I$ such that $\iota(i) = n$.

Remark 4.1. Note that, using ι , one can construct a canonical bijection between the set of partitions of unity subordinate to \mathcal{V} and the set of partitions of unity subordinate to \mathcal{U} . The same is true for the sets of approximate cocycles subordinate to \mathcal{V} and \mathcal{U} and for the sets of approximate local trivializations subordinate to \mathcal{V} and \mathcal{U} .

We give the main constructions of this section for covers indexed by $\mathbb{N}_{\geq 1}$, and we will later generalize them to arbitrary countable covers, as this simplifies exposition.

Construction 4.2. Let B be a paracompact topological space and let $\mathcal{V} = \{V_i\}_{i \in \mathbb{N}_{\geq 1}}$ be a cover of B . Let φ be a partition of unity subordinate to \mathcal{V} . Given a cochain Ω subordinate to \mathcal{V} define, for each $i \in \mathbb{N}_{\geq 1}$, a map $\Phi_i : V_i \rightarrow \mathbb{V}(d)$, where the rows of $\Phi_i(y)$ from $d \times j$ to $d \times (j + 1) - 1$ are given by

$$\sqrt{\varphi_j(y)} \Omega_{ij}(y)^t. \tag{4.1}$$

Note that the maps Φ_i of Construction 4.2 are continuous and are well defined since, if $y \notin V_j$, then $\varphi_j(y) = 0$.

Lemma 4.3. *Let $\varepsilon \in [0, \infty]$ and let Ω be an ε -approximate cocycle subordinate to an open cover $\mathcal{V} = \{V_i\}_{i \in \mathbb{N}_{\geq 1}}$ of a paracompact topological space. The maps Φ_i of Construction 4.2 form an ε -approximate local trivialization.*

Proof. We give the proof for $\varepsilon \in (0, \infty]$, the case $\varepsilon = 0$ being similar. Let $(ij) \in N(\mathcal{V})$. We claim that the original ε -approximate cocycle Ω is a witness that the family Φ is an ε -approximate local trivialization. To prove this, we must show that, for all $y \in V_j \cap V_i$, the Frobenius distance between $\Phi_i(y)\Omega_{ij}(y)$ and $\Phi_j(y)$ is less than ε . Carrying out the product $\Phi_i(y)\Omega_{ij}(y)$, we get an element of $V(d)$ with rows from $d \times k$ to $d \times (k + 1) - 1$ given by

$$\sqrt{\varphi_k(y)}\Omega_{ik}(y)^t\Omega_{ij}(y).$$

So $\|\Phi_i(y)\Omega_{ij}(y) - \Phi_j(y)\| = \left(\sum_{k \geq 1} \|\sqrt{\varphi_k(y)}(\Omega_{ik}(y)^t\Omega_{ij}(y) - \Omega_{kj}(y))\|^2\right)^{1/2} < \varepsilon$, as required. \square

We have thus constructed a map

$$\text{triv}^\varphi : Z_\varepsilon^1(\mathcal{V}; O(d)) \rightarrow T_\varepsilon(\mathcal{V}; d),$$

for any cover \mathcal{V} of a paracompact topological space B that is indexed by $\mathbb{N}_{\geq 1}$. It is important to note that this map depends on the choice of partition of unity φ . Nevertheless, using two different partitions of unity gives homotopic local trivializations, in the following sense.

Lemma 4.4. *Let Ω be an ε -approximate cocycle subordinate to $\mathcal{V} = \{V_i\}_{i \in \mathbb{N}_{\geq 1}}$. If φ and φ' are two partitions of unity subordinate to \mathcal{V} , then $\text{triv}^\varphi(\Omega)$ and $\text{triv}^{\varphi'}(\Omega)$ are homotopic through a family of ε -approximate local trivializations that admit Ω as a witness.*

Proof. For any $\alpha \in [0, 1]$, the formula $\varphi_i^\alpha = \alpha\varphi_i + (1 - \alpha)\varphi'_i$ gives a partition of unity. Now observe that the family of ε -approximate local trivializations $\text{triv}^{\varphi^\alpha}(\Omega)$ admit Ω as a witness. \square

Next, we show that the construction assigning an approximate local trivialization to an approximate cocycle is stable.

Lemma 4.5. *Let $\mathcal{V} = \{V_i\}_{i \in \mathbb{N}_{\geq 1}}$ be a cover of a paracompact topological space B , and let φ be a partition of unity subordinate to \mathcal{V} . For Ω and Λ ε -approximate cocycles subordinate to \mathcal{V} ,*

$$d_T(\text{triv}^\varphi(\Omega), \text{triv}^\varphi(\Lambda)) \leq d_Z(\Omega, \Lambda).$$

Proof. Let $d_Z(\Omega, \Lambda) = \delta$, $\Phi := \text{triv}^\varphi(\Omega)$, and $\Psi := \text{triv}^\varphi(\Lambda)$. For $y \in V_i$, we have

$$\|\Phi_i(y) - \Psi_i(y)\|^2 = \sum_{k \in I} \varphi_k(y)\|\Omega_{ik}(y) - \Lambda_{ik}(y)\|^2 \leq \delta^2 \sum_{k \in I} \varphi_k(y) = \delta^2,$$

which proves the claim. \square

Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a countable open cover of a paracompact topological space B , and let $\iota : I \rightarrow \mathbb{N}_{\geq 1}$ be an injection. Using Remark 4.1, we can generalize Construction 4.2, Lemma 4.3, Lemma 4.4 and Lemma 4.5 to \mathcal{U} .

In particular, we have a map $\text{triv}^{\varphi, \iota} : Z_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow T_\varepsilon(\mathcal{U}; d)$ that now also depends on the choice of injection ι . We now prove that, up to homotopy, triv is independent of the choice of ι . In order to show this, we prove a more general lemma that will be of use later.

Lemma 4.6. *Given an injection $\iota : \mathbb{N}_{\geq 1} \rightarrow \mathbb{N}_{\geq 1}$ define $\chi^\iota : V(d) \rightarrow V(d)$ by mapping a frame $\Phi \in V(d)$ to the frame whose $\iota(k)$ th row is the k th row of Φ and whose other rows are identically 0. If $\iota, \iota' : \mathbb{N}_{\geq 1} \rightarrow \mathbb{N}_{\geq 1}$ are injections, then χ^ι and $\chi^{\iota'} : V(d) \rightarrow V(d)$ are homotopic.*

The proof of the lemma is standard, but we give it here for completeness.

Proof. Assume that ι and ι' have disjoint images. Then, there is a homotopy between χ^ι and $\chi^{\iota'}$ given by $\sqrt{\alpha}\chi^\iota + \sqrt{1 - \alpha}\chi^{\iota'}$ for $\alpha \in [0, 1]$. Let $2\iota : \mathbb{N}_{\geq 1} \rightarrow \mathbb{N}_{\geq 1}$ be given by $2\iota(k) = 2 \times \iota(k)$ and define $2\iota' + 1$ in an analogous way. Since 2ι and $2\iota' + 1$ have disjoint images, our previous reasoning reduces the

problem to showing that χ^ι and $\chi^{2\iota}$ are homotopic and that $\chi^{\iota'}$ and $\chi^{2\iota'-1}$ are homotopic. Since the two proofs are entirely analogous, we give the details only for the case of ι . Moreover, since $2\iota : \mathbb{N}_{\geq 1} \rightarrow \mathbb{N}_{\geq 1}$ is the composite of ι with multiplication by 2, it is enough to give the proof for $\iota = \text{id}$, which we now do.

We must show that the identity $V(d) \rightarrow V(d)$ is homotopic to $\chi^2 : V(d) \rightarrow V(d)$. Informally, the proof works by moving each row of Φ at a time. To simplify exposition, in the rest of this proof, the notation Φ_m will be used to refer to the m th row of a frame $\Phi \in V(d)$. Consider the family of functions $f^\alpha : V(d) \rightarrow V(d)$ indexed by $\alpha \in [0, 1]$ defined as follows. For $\alpha = 0$, let $f^\alpha = \text{id}$. For $n \in \mathbb{N}_{\geq 1}$, $\alpha \in [1/(n + 1), 1/n]$, and $m \in \mathbb{N}_{\geq 1}$, define

$$(f^\alpha(\Phi))_m = \begin{cases} \Phi_m, & m < n \\ \sqrt{\beta} \Phi_m, & m = n \\ 0, & n < m < 2n \\ \sqrt{1 - \beta} \Phi_m, & m = 2n \\ \Phi_{m/2}, & m > 2n \text{ and } m \text{ even} \\ 0, & m > 2n \text{ and } m \text{ odd,} \end{cases}$$

where $\beta = (1/n - \alpha) \times (1/n - 1/(n + 1))$. By inspection, we see that f^α gives a homotopy between the identity and χ^2 , using that $V(d)$ has the direct limit topology. □

Lemma 4.7. *Let Ω be an ε -approximate cocycle subordinate to a countable open cover $\mathcal{U} = \{U_i\}_{i \in I}$ and let $\iota, \iota' : I \rightarrow \mathbb{N}_{\geq 1}$ be injections. Let φ be a partition of unity subordinate to \mathcal{U} . Then $\text{triv}^{\varphi, \iota}(\Omega)$ and $\text{triv}^{\varphi, \iota'}(\Omega)$ are homotopic through a family of ε -approximate local trivialisations that admit Ω as a witness.*

Proof. The result follows at once from Lemma 4.6 by noticing that, for any two injections $\iota, \iota' : I \rightarrow \mathbb{N}_{\geq 1}$, there exists a bijection b of $\mathbb{N}_{\geq 1}$ such that $\iota' = b \circ \iota$. □

We now consider the problem of assigning an approximate cocycle to an approximate local trivialisaton.

Construction 4.8. Let $\Phi = \{\Phi_i\}_{i \in I}$ be an ε -approximate local trivialisaton subordinate to $\mathcal{U} = \{U_i\}_{i \in I}$. By definition, there exists a witness that Φ is an ε -approximate local trivialisaton. Choose, arbitrarily, such a witness Ω . Without loss of generality, we may assume that Ω is symmetric and thus that it is a cochain.

Lemma 4.9. *Let $\varepsilon \in [0, \infty]$. Let Φ be an ε -approximate local trivialisaton. Then, the cochain described in Construction 4.8 is a 3ε -approximate cocycle. Thus, Construction 4.8 gives a map $w : \mathbb{T}_\varepsilon(\mathcal{U}; d) \rightarrow Z^1_{3\varepsilon}(\mathcal{U}; O(d))$.*

Proof. We address the case $\varepsilon \in (0, \infty]$, the case $\varepsilon = 0$ being similar. Let $(ijk) \in N(\mathcal{U})$ and let $y \in U_k \cap U_j \cap U_i$. Since $\|\Phi_i(y)\Omega_{ik}(y) - \Phi_k(y)\| < \varepsilon$, we have that $\|\Phi_k(y)^t \Phi_i(y)\Omega_{ik}(y) - \text{id}\| < \varepsilon$, by Lemma A.5. This implies that $\|\Phi_i(y)^t \Phi_k(y) - \Omega_{ik}(y)\| < \varepsilon$.

A similar computation shows that $\|\Phi_i(y)^t \Phi_j(y)\Omega_{jk}(y) - \Phi_i(y)^t \Phi_k(y)\| < \varepsilon$. Using the triangle inequality and the first bound in the proof, we get that $\|\Phi_i(y)^t \Phi_j(y)\Omega_{jk}(y) - \Omega_{ik}(y)\| < 2\varepsilon$. By Lemma A.5 and the first bound in this proof but for i and j , we have that $\|\Phi_i(y)^t \Phi_j(y)\Omega_{jk}(y) - \Omega_{ij}(y)\Omega_{jk}(y)\| < \varepsilon$. The triangle inequality then finishes the proof. □

Although we can associate a 3ε -approximate cocycle to every ε -approximate local trivialisaton, this choice is not canonical, as an approximate local trivialisaton can have many distinct witnesses. Nonetheless, the following result says that any two witnesses cannot be too far apart.

Lemma 4.10. *Let Ω and Λ be witnesses that Φ and Ψ are, respectively, ε - and δ -approximate local trivialisatons. Then $d_Z(\Omega, \Lambda) \leq \varepsilon + \delta + \sqrt{2}d_T(\Phi, \Psi)$.*

Proof. We address the case in which $\varepsilon, \delta > 0$, the case in which any of them is 0 being similar. Let $(ij) \in N(\mathcal{U})$ and $y \in U_j \cap U_i$. We have $\|\Omega_{ij}(y) - \Phi_i(y)^t \Phi_j(y)\| < \varepsilon$ and $\|\Lambda_{ij}(y) - \Psi_i(y)^t \Psi_j(y)\| < \delta$, so it suffices to show that $\|\Phi_i(y)^t \Phi_j(y) - \Psi_i(y)^t \Psi_j(y)\| \leq \sqrt{2}\|\Phi_i(y) - \Psi_i(y)\|$, which follows from Lemma A.6. \square

Remark 4.11. From Lemma 4.10, it follows that, if Ω and Λ are witnesses that Φ is an ε -approximate local trivialization, then $d_Z(\Omega, \Lambda) \leq 2\varepsilon$, and thus w is, approximately, a left inverse of triv^φ , in the sense that $d_Z(\Omega, w(\text{triv}^\varphi(\Omega))) \leq 2\varepsilon$ for every ε -approximate cocycle Ω . The following result can be interpreted as saying that w is also a right inverse of triv^φ since it implies, in particular, that $\text{triv}^\varphi(w(\Phi))$ is homotopic to Φ through 3ε -approximate local trivializations, whenever Φ is an ε -approximate local trivialization.

Lemma 4.12. *Let $\varepsilon \geq 0$. Let Ω be a witness that Φ and Ψ are ε -approximate cocycles. Then Φ and Ψ are homotopic through ε -approximate local trivializations subordinate to $\mathcal{U} = \{U_i\}_{i \in I}$ that admit Ω as a witness.*

Proof. We use the language of Lemma 4.6. Consider the maps $\iota, \iota' : \mathbb{N}_{\geq 1} \rightarrow \mathbb{N}_{\geq 1}$ given by $\iota(k) = 2k$ and $\iota'(k) = 2k - 1$. It is clear that Ω is a witness that $\{\chi^\iota \circ \Phi_i\}_{i \in I}$ is an ε -local trivialization, and Lemma 4.6 implies that Φ is homotopic to $\{\chi^\iota \circ \Phi_i\}_{i \in I}$ through ε -approximate local trivializations that admit Ω as a witness. Similarly, we deduce that Ψ is homotopic to $\{\chi^{\iota'} \circ \Psi_i\}_{i \in I}$ through ε -approximate local trivializations that admit Ω as a witness.

Consider, for $\alpha \in [0, 1]$, the family $\{\sqrt{\alpha}(\chi^\iota \circ \Phi_i) + \sqrt{1 - \alpha}(\chi^{\iota'} \circ \Psi_i)\}_{i \in I}$. Since the images of ι and ι' are disjoint, this constitutes an ε -approximate local trivialization that admits Ω as a witness that varies continuously with α . The result follows. \square

The following result gives conditions under which there is a canonical approximate cocycle associated to an approximate local trivialization.

Lemma 4.13. *Let Φ be a nondegenerate ε -approximate local trivialization. For $(ij) \in N(\mathcal{U})$ and $y \in U_j \cap U_i$, let $\Omega_{ij}(y) \in O(d)$ minimize $\|\Phi_i(y)\Omega - \Phi_j(y)\|$, where Ω ranges over $O(d)$. Then, the matrices $\Omega_{ij}(y)$ assemble into a 3ε -approximate cocycle.*

Proof. To see that the mappings $\Omega_{ij} : U_j \cap U_i \rightarrow O(d)$ are continuous, use Corollary A.4. To see that $\Omega_{ij} = \Omega_{ji}^t$ use that the minimizers are unique. Then, Lemma 4.9 directly implies that Ω satisfies the 3ε -approximate cocycle condition, as required. \square

We now give sufficient conditions for an approximate local trivialization to be nondegenerate.

Lemma 4.14. *If $\varepsilon \leq 1$, then any ε -approximate local trivialization is nondegenerate.*

Proof. This is a consequence Lemma A.15, which says that a matrix that is at Frobenius distance less than 1 from an orthogonal matrix is invertible. \square

We conclude by remarking that Lemma 4.14 implies that, if $\varepsilon \leq 1$, we have a canonical map (that is independent of any arbitrary choices)

$$w : T_\varepsilon(\mathcal{U}; d) \rightarrow Z_{3\varepsilon}^1(\mathcal{U}; O(d))$$

given by taking the best witness, as in Lemma 4.13.

4.2. Local trivializations and classifying maps

In this section, we give a construction that, given an approximate local trivialization, returns an approximate classifying map. We also observe that this construction behaves well with respect to homotopies between approximate local trivializations.

Recall from Section 2.1.2 the definition of the map $\text{Proj} : V(d) \rightarrow \text{Gr}(d)$.

Construction 4.15. Let \mathcal{U} be a countable open cover of a paracompact topological space B and let φ be a partition of unity subordinate to \mathcal{U} . Let Φ be an approximate local trivialization subordinate to \mathcal{U} . Define a map $\text{av}^\varphi(\Phi) : B \rightarrow \mathbb{R}^{\infty \times \infty}$ by

$$\text{av}^\varphi(\Phi)(y) = \sum_{i \in I} \varphi_i(y) \text{Proj}(\Phi_i(y)). \quad \triangleleft$$

Note that the map $\text{av}^\varphi(\Phi)$ is continuous.

Lemma 4.16. Let $\varepsilon \in [0, \infty]$. Under the hypotheses of Construction 4.15, if Φ is an ε -approximate local trivialization, then $\text{av}^\varphi(\Phi)$ is a $\sqrt{2\varepsilon}$ -approximate classifying map.

Proof. We address the case $\varepsilon \in (0, \infty]$, the case $\varepsilon = 0$ being similar. By definition, there exist, for each $(ij) \in N(\mathcal{U})$, a continuous map $\Omega_{ij} : U_j \cap U_i \rightarrow O(d)$ such that $\|\Phi_i(y)\Omega_{ij}(y) - \Phi_j(y)\| < \varepsilon$ for every $y \in U_j \cap U_i$. Since $\text{Proj} : V(d) \rightarrow \text{Gr}(d)$ is $O(d)$ -invariant and $\sqrt{2}$ -Lipschitz (Lemma A.6), it follows that $\|\text{Proj}(\Phi_i(y)) - \text{Proj}(\Phi_j(y))\| < \sqrt{2\varepsilon}$. The result then follows from Lemma A.2. \square

The following is clear.

Lemma 4.17. Let Φ and Ψ be approximate local trivializations subordinate to \mathcal{U} , a countable open cover of a paracompact topological space B . Let φ be a partition of unity subordinate to \mathcal{U} .

1. We have that $d_C(\text{av}^\varphi(\Phi), \text{av}^\varphi(\Psi)) \leq \sqrt{2}d_T(\Phi, \Psi)$.
2. If Φ and Ψ are homotopic through ε -approximate local trivializations subordinate to \mathcal{U} , then $\text{av}^\varphi(\Phi)$ and $\text{av}^\varphi(\Psi)$ are homotopic as maps $B \rightarrow \text{Gr}(d)^{\sqrt{2\varepsilon}}$.

4.3. Cocycles and classifying maps

In this section, we relate approximate cocycles to approximate classifying maps in a way that is independent of any partition of unity and of any enumeration of the sets in the open cover the cocycle is subordinate to (Theorem 4.21). We also study the action of refinements on approximate cohomology.

Let B be a paracompact topological space and let $\mathcal{U} = \{U_i\}_{i \in I}$ be a countable open cover. Let φ be a partition of unity subordinate to \mathcal{U} and let $\iota : I \rightarrow \mathbb{N}_{\geq 1}$ be an injection. Using Lemma 4.3 and Lemma 4.16, we get a map $\text{av}^\varphi \circ \text{triv}^{\varphi, \iota} : Z_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow \text{Maps}(B, \text{Gr}(d)^{\sqrt{2\varepsilon}})$ which we compose with the quotient map $\text{Maps}(B, \text{Gr}(d)^{\sqrt{2\varepsilon}}) \rightarrow [B, \text{Gr}(d)^{\sqrt{2\varepsilon}}]$ to obtain a map

$$\text{cl}' : Z_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow [B, \text{Gr}(d)^{\sqrt{2\varepsilon}}].$$

Together, Lemma 4.4, Lemma 4.17, and Lemma 4.7 imply that this map is independent of the choice of partition of unity φ and of injection ι . The following result says that two approximate cocycles that differ in a 0-cochain are sent to the same approximate classifying map by cl' .

Lemma 4.18. Let B be a paracompact topological space and let \mathcal{U} be a countable open cover. The map $\text{cl}' : Z_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow [B, \text{Gr}(d)^{\sqrt{2\varepsilon}}]$ factors through $\check{H}_\varepsilon^1(\mathcal{U}; O(d))$.

Proof. By construction, we may assume that \mathcal{U} is indexed by $I = \mathbb{N}$. Consider the following open cover $\mathcal{V} = \{V_j\}_{j \in J}$ indexed by $J = \mathbb{N}$. Let $V_j = U_i$ whenever $j = 2i + z$ with $z = 0$ or $z = 1$. So the open cover \mathcal{V} consists of two copies of each open set of \mathcal{U} , where each open U_i appears with an even index as V_{2i} and with an odd index as V_{2i+1} .

Suppose that Ω and Ω' are equal in $\check{H}_\varepsilon^1(\mathcal{U}; O(d))$ so that there is a 0-cochain Θ such that $\Theta \cdot \Omega = \Omega'$. Let φ be a partition of unity subordinate to \mathcal{U} . This induces two partitions of unity φ^0 and φ^1 subordinate to \mathcal{V} , where φ_j^0 is equal to $\varphi_{j/2}$ if j is even and is identically 0 if j is odd. Similarly, φ_j^1 is equal to $\varphi_{(j-1)/2}$ if j is odd and identically 0 if j is even.

Consider the following cochain Λ subordinate to \mathcal{V} . For $(jk) \in N(\mathcal{V})$, define $\Lambda_{jk} = \Omega_{j/2k/2}$ if j and k are even, $\Lambda_{jk} = \Theta_{(j-1)/2}^t \Omega_{(j-1)/2(k-1)/2} \Theta_{(k-1)/2}$ if j and k are odd, $\Lambda_{jk} = \Theta_{(j-1)/2}^t \Omega_{(j-1)/2k/2}$ if j is odd and k is even, and $\Lambda_{jk} = \Omega_{j/2(k-1)/2} \Theta_{(k-1)/2}$ if j is even and k is odd. It is clear that Λ is an ε -approximate cocycle subordinate to \mathcal{V} .

Finally, using Lemma 4.7, if we use the partition of unity φ^0 , we see that $\text{cl}'(\Lambda) = \text{cl}'(\Omega)$, and if we use the partition of unity φ^1 , we see that $\text{cl}'(\Lambda) = \text{cl}'(\Omega')$. The result follows. \square

Recall from Section 2.2.4 the notion of refinement of a cover.

Construction 4.19. Let $\nu : \mathcal{U} \rightarrow \mathcal{V}$ be a refinement of covers of a topological space B and let $\varepsilon \in [0, \infty]$. Let $\Omega \in Z_\varepsilon^1(\mathcal{V}; O(d))$. Define $\nu(\Omega) \in Z_\varepsilon^1(\mathcal{U}; O(d))$ by letting $\nu(\Omega)_{jk} = \Omega_{\nu(j)\nu(k)}$ for all $(jk) \in N(\mathcal{U})$.

Construction 4.19 gives a map $\nu : Z_\varepsilon^1(\mathcal{V}; O(d)) \rightarrow Z_\varepsilon^1(\mathcal{U}; O(d))$ that descends to a map

$$\nu : \check{H}_\varepsilon^1(\mathcal{V}; O(d)) \rightarrow \check{H}_\varepsilon^1(\mathcal{U}; O(d)).$$

It is clear that both these maps are 1-Lipschitz with respect to d_Z and with respect to $d_{\check{H}}$.

Lemma 4.20. Let $\varepsilon \in [0, \infty]$, let $\mu, \nu : \mathcal{U} \rightarrow \mathcal{V}$ and let $\Omega \in Z_\varepsilon^1(\mathcal{V}; O(d))$. Then, for all $(jk) \in N(\mathcal{U})$ and $y \in U_k \cap U_j$, we have $\|\mu(\Omega)_{jk}(y) - \nu(\Omega)_{jk}(y)\| < 2\varepsilon$, and thus $d_{\check{H}}(\mu(\Omega), \nu(\Omega)) \leq 2\varepsilon$.

Proof. We address the case $\varepsilon \in (0, \infty]$, the case $\varepsilon = 0$ being similar. We start by defining a 0-cochain Θ subordinate to \mathcal{U} . Given $j \in N(\mathcal{U})$, let $\Theta_j = \Omega_{\mu(j)\nu(j)}$ if $\mu(j) \neq \nu(j)$ and the identity if $\mu(j) = \nu(j)$. Let $(jk) \in N(\mathcal{U})$, and let $y \in U_k \cap U_j$. To simplify notation in the rest of this proof, let us denote $\Omega_{ab}(y)$ by Ω_{ab} . We have

$$\begin{aligned} \|\mu(\Omega)_{jk}(y) - (\Theta \cdot \nu(\Omega))_{jk}(y)\| &= \|\Omega_{\mu(j)\mu(k)} - \Omega_{\mu(j)\nu(j)} \Omega_{\nu(j)\nu(k)} \Omega_{\nu(k)\mu(k)}\| \\ &= \|\Omega_{\nu(j)\mu(j)} \Omega_{\mu(j)\mu(k)} - \Omega_{\nu(j)\nu(k)} \Omega_{\nu(k)\mu(k)}\| \\ &\leq \|\Omega_{\nu(j)\mu(j)} \Omega_{\mu(j)\mu(k)} - \Omega_{\nu(j)\mu(k)}\| \\ &\quad + \|\Omega_{\nu(j)\mu(k)} - \Omega_{\nu(j)\nu(k)} \Omega_{\nu(k)\mu(k)}\| \\ &< 2\varepsilon, \end{aligned}$$

where for the second equality we used the fact that $\Omega_{\nu(j)\mu(j)} = \Omega_{\mu(j)\nu(j)}^{-1}$ combined with the fact that the Frobenius norm is invariant under multiplication by an orthogonal matrix, and for the inequalities we used the triangle inequality and the approximate cocycle condition. \square

We are now ready to state and prove the main result of this section.

Theorem 4.21. Let B be a paracompact topological space, and let \mathcal{U} be a countable cover of B . Let $\varepsilon \in [0, \infty]$. The map cl' induces a map

$$\text{cl} : \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow [B, \text{Gr}(d)^{\sqrt{2}\varepsilon}]$$

such that, if $d_{\check{H}}(\Omega, \Lambda) < \delta$ in $\check{H}_\varepsilon^1(\mathcal{U}; O(d))$, then $\text{cl}(\Omega)$ and $\text{cl}(\Lambda)$ become equal in $[B, \text{Gr}(d)^{\sqrt{2}(\varepsilon+\delta)}]$. Moreover, if $\mu, \nu : \mathcal{V} \rightarrow \mathcal{U}$ are refinements and \mathcal{V} is a countable cover of B , then $\text{cl}(\mu(\Omega))$ and $\text{cl}(\nu(\Omega))$ become equal in $[B, \text{Gr}(d)^{2\sqrt{2}\varepsilon}]$.

Proof. The map cl is well defined thanks to Lemma 4.18. For the stability of cl , note that, using Lemma 4.17, we see that if Ω and Λ are ε -approximate cocycles subordinate to a countable cover $\mathcal{U} = \{U_i\}_{i \in I}$, φ is a partition of unity subordinate to \mathcal{U} , and $\iota : I \rightarrow \mathbb{N}_{\geq 1}$ is an injection, then

$$d_C(\text{av}^\varphi \circ \text{triv}^{\varphi, \iota}(\Omega), \text{av}^\varphi \circ \text{triv}^{\varphi, \iota}(\Lambda)) \leq \sqrt{2}d_Z(\Omega, \Lambda).$$

The stability then follows from Lemma 3.10. Finally, the claim about refinements follows directly from Lemma 4.20. \square

Note that the map cl is independent of any choice of partition of unity or enumeration of the cover \mathcal{U} . We conclude with an interesting remark that is not used in the rest of the paper.

Remark 4.22. Let $\text{cov}(B)$ be the category whose objects are the countable covers of a paracompact topological space B and whose morphisms are the refinements. Let $\varepsilon \in [0, \infty]$. Construction 4.19 gives a functor $\check{H}_\varepsilon^1(-; O(d)) : \text{cov}(B)^{\text{op}} \rightarrow \mathbf{Met}$. We can then define

$$\check{H}_\varepsilon^1(B; O(d)) = \text{colim}_{\mathcal{U} \in \text{cov}(B)} \check{H}_\varepsilon^1(\mathcal{U}; O(d)),$$

with the caveat that $\check{H}_\varepsilon^1(B; O(d))$ may be a *pseudo* metric space. Theorem 4.21 implies that there is a well-defined map $\text{cl} : \check{H}_\varepsilon^1(B; O(d)) \rightarrow [B, \text{Gr}(d)^{2\sqrt{2}\varepsilon}]$, natural in $\varepsilon \in [0, \infty]$.

4.4. Relationship to classical vector bundles

We now relate approximate vector bundles to exact vector bundles, following the intuition that ε -approximate vector bundles should correspond to true vector bundles as long as ε is sufficiently small. For this, we use Theorem 4.21. In the case where an approximate cocycle represents a true vector bundle, we study the problem of constructing an exact cocycle that represents the same vector bundle. We also give upper and lower bounds for the distance from an approximate cocycle to an exact cocycle representing the same vector bundle (Theorem 4.27 and Lemma 4.30).

We start by recalling that small thickenings of the Grassmannian embedded in $\mathbb{R}^{\infty \times \infty}$ retract to the Grassmannian. More precisely, if $\varepsilon \leq \sqrt{2}/2$, there is a map $\pi : \text{Gr}(d)^\varepsilon \rightarrow \text{Gr}(d)$ which is a homotopy inverse of the inclusion $\text{Gr}(d) \subseteq \text{Gr}(d)^\varepsilon$, by Proposition A.12. Let B be a topological space. By postcomposing with π , we get an inverse for the natural map $[B, \text{Gr}(d)] \rightarrow [B, \text{Gr}(d)^\varepsilon]$ which we denote by

$$\pi_* : [B, \text{Gr}(d)^\varepsilon] \rightarrow [B, \text{Gr}(d)].$$

By an abuse of notation, we also let $\pi_* : \text{Maps}(B, \text{Gr}(d)^\varepsilon) \rightarrow \text{Maps}(B, \text{Gr}(d))$.

Recall that we constructed a map $\text{cl} : \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow [B, \text{Gr}(d)^{\sqrt{2}\varepsilon}]$, so, if $\varepsilon \leq 1/2$, then any ε -approximate cocycle Ω represents a true vector bundle, namely $\pi_*(\text{cl}(\Omega)) \in [B, \text{Gr}(d)]$. To summarize, if $\varepsilon \leq 1/2$, we have defined a map

$$\pi_* \circ \text{cl} : \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow [B, \text{Gr}(d)].$$

Upper bound.

For the rest of this section, we let Φ be an ε -approximate local trivialization subordinate to a countable cover $\mathcal{U} = \{U_i\}_{i \in I}$ of a paracompact topological space B , with $\varepsilon \in [0, \infty]$; we let φ be a partition of unity subordinate to \mathcal{U} and let $\iota : I \rightarrow \mathbb{N}_{\geq 1}$ be an injection. Recall from Lemma 4.16 that

$$\text{av}^\varphi(\Phi)(y) = \sum_{i \in I} \varphi_i(y) \text{Proj}(\Phi_i(y))$$

defines a $\sqrt{2}\varepsilon$ -approximate classifying map $B \rightarrow \text{Gr}(d)^{\sqrt{2}\varepsilon}$. We will make use of results in Appendix A.2.

Lemma 4.23. *If $\varepsilon \in (0, \infty]$, then for, $i \in I$ and $y \in U_i$, we have*

$$\|\Phi_i(y)\Phi_i(y)^t - \pi_*(\text{av}^\varphi(\Phi)(y))\| < 2\sqrt{2}\varepsilon \quad \text{and} \quad \|\Phi_i(y) - \pi_*(\text{av}^\varphi(\Phi)(y))\Phi_i(y)\| < 2\sqrt{2}\varepsilon.$$

If $\varepsilon = 0$, then $\Phi_i(y)\Phi_i(y)^t = \pi_(\text{av}^\varphi(\Phi)(y))$ and $\Phi_i(y) = \pi_*(\text{av}^\varphi(\Phi)(y))\Phi_i(y)$.*

Proof. We address the case $\varepsilon \in (0, \infty]$, the case $\varepsilon = 0$ being similar. The second inequality is a consequence of the first one and Lemma A.5. For the first inequality, use Lemma A.2 together with Lemma A.6. \square

Lemma 4.24. *Assume that $\varepsilon \leq \sqrt{2}/4$. For every $i \in I$ and $y \in U_i$, the matrix $\pi_*(\text{av}^\varphi(\Phi)(y))\Phi_i(y)$ has rank d .*

Proof. To simplify notation, let us omit from the formulas the evaluations on $y \in U_i$. It is enough to show that $(\pi_*(\text{av}^\varphi(\Phi))\Phi_i)^t \pi_*(\text{av}^\varphi(\Phi))\Phi_i$ has full rank. Note that

$$(\pi_*(\text{av}^\varphi(\Phi))\Phi_i)^t \pi_*(\text{av}^\varphi(\Phi))\Phi_i = \Phi_i^t \pi_*(\text{av}^\varphi(\Phi)) \pi_*(\text{av}^\varphi(\Phi))\Phi_i = \Phi_i^t \pi_*(\text{av}^\varphi(\Phi))\Phi_i.$$

Using Lemma 4.23 and Lemma A.5, we conclude that

$$\|\Phi_i^t \pi_*(\text{av}^\varphi(\Phi))\Phi_i - \text{id}\| = \|\Phi_i^t \pi_*(\text{av}^\varphi(\Phi))\Phi_i - \Phi_i^t \Phi_i \Phi_i^t \Phi_i\| < 2\sqrt{2}\varepsilon.$$

The result then follows from Lemma A.15, as $2\sqrt{2}\varepsilon \leq 1$ by assumption. \square

Given an ε -approximate local trivialization Φ with $\varepsilon \leq \sqrt{2}/4$, we now define an exact local trivialization Ψ that represents the same vector bundle. Given $i \in I$ and $y \in U_i$, we let

$$\Psi_i(y) := Q(\pi_*(\text{av}^\varphi(\Phi)(y))\Phi_i(y)),$$

where the map Q is the one of Corollary A.4. By Lemma 4.24 and Corollary A.4, the maps Ψ_i are well defined and continuous. To see that it is an exact local trivialization it suffices to check that, if $y \in U_j \cap U_i$, then the columns of $\Psi_i(y)$ and $\Psi_j(y)$ span the same subspace of \mathbb{R}^∞ . This is a consequence of the fact that the columns of $\Psi_i(y)$ span the image of $\pi_*(\text{av}^\varphi(\Phi)(y))$. We also deduce the following.

Lemma 4.25. *Let φ be a partition of unity subordinate to \mathcal{U} . Then $\text{av}^\varphi(\Psi) = \pi_*(\text{av}^\varphi(\Phi))$.*

We now bound the distance between Ψ and Φ .

Lemma 4.26. *Assume that $\varepsilon \leq \sqrt{2}/4$, then $d_\tau(\Phi, \Psi) \leq 4\sqrt{2}\varepsilon$.*

Proof. To simplify notation, let us omit from the formulas the evaluations on $y \in U_i$. For every $i \in I$ and $y \in U_i$, we have

$$\begin{aligned} \|\Phi_i - Q(\pi_*(\text{av}^\varphi(\Phi))\Phi_i)\| &\leq \|\Phi_i - \pi_*(\text{av}^\varphi(\Phi))\Phi_i\| + \|\pi_*(\text{av}^\varphi(\Phi))\Phi_i - Q(\pi_*(\text{av}^\varphi(\Phi))\Phi_i)\| \\ &< 2\sqrt{2}\varepsilon + 2\sqrt{2}\varepsilon = 4\sqrt{2}\varepsilon, \end{aligned}$$

where we bounded the first summand using Lemma 4.23 and the second summand using Lemma A.8. In order to satisfy the hypotheses of Lemma A.8, we use the same argument as in Lemma 4.24. \square

Theorem 4.27. *Let $\varepsilon \leq \sqrt{2}/4$ and let $\Omega \in Z_\varepsilon^1(\mathcal{U}; O(d))$. There exists $\Lambda \in Z^1(\mathcal{U}; O(d))$ such that $\text{cl}(\Lambda) = \pi_*(\text{cl}(\Omega))$ and such that $d_Z(\Omega, \Lambda) \leq 9\varepsilon$.*

Proof. Let φ be a partition of unity subordinate to \mathcal{U} . Since $\varepsilon \leq \sqrt{2}/4 \leq 1/2$, it follows that $\pi_*(\text{cl}(\Omega))$ is well defined. By construction, $\pi_*(\text{cl}(\Omega))$ is the homotopy class of $\pi_* \circ \text{av}^\varphi \circ \text{triv}^{\varphi, \iota}(\Omega)$. By Lemma 4.26, there is an exact local trivialization Ψ such that $d_\tau(\text{triv}^{\varphi, \iota}(\Omega), \Psi) \leq 4\sqrt{2}\varepsilon$ and such that $\text{av}^\varphi(\Psi) = \pi_* \circ \text{av}^\varphi \circ \text{triv}^{\varphi, \iota}(\Omega)$. Let $\Lambda = w(\Psi)$. Then Λ is a witness that Ψ is an exact cocycle, so, by Lemma 4.10, we have that $d_Z(\Omega, \Lambda) \leq \varepsilon + \sqrt{2} \times 4\sqrt{2}\varepsilon = 9\varepsilon$. To conclude, note that $\text{cl}(\Lambda) = [\text{av}^\varphi \circ \text{triv}^{\varphi, \iota} \circ w(\Psi)] = [\text{av}^\varphi(\Psi)] = [\pi_* \circ \text{av}^\varphi \circ \text{triv}^{\varphi, \iota}(\Omega)] = \pi_*(\text{cl}(\Omega))$, where in the second equality we used Lemma 4.12 to conclude that $\text{triv}^{\varphi, \iota}(w(\Psi))$ and Ψ are homotopic through 0-approximate local trivializations. \square

Lower bound.

The following definition and result are inspired by Robinson’s notion of consistency radius [55, 54]. The idea of this short section is to give a lower bound for the distance from an approximate cocycle to an exact cocycle.

Definition 4.28. Let Ω be a cochain subordinate to \mathcal{U} . The **consistency radius** of Ω , denoted by $r(\Omega)$, is the infimum over all ε such that Ω belongs to $Z^1_\varepsilon(\mathcal{U}; O(d))$.

A similar argument to the one in Lemma 4.9 proves the following.

Lemma 4.29. Let $\Lambda \in Z^1_\varepsilon(\mathcal{U}; O(d))$ and $\Omega \in C^1(\mathcal{U}; O(d))$. Let $\delta > 0$. If $d_Z(\Lambda, \Omega) < \delta$, then $\Omega \in Z^1_{\varepsilon+3\delta}(\mathcal{U}; O(d))$.

Lemma 4.30. Let Ω be a cochain, and let $\varepsilon > r(\Omega)$. Then, the distance from Ω to $\check{H}^1_\varepsilon(\mathcal{U}; O(d))$ is bounded below by $\frac{r(\Omega)-\varepsilon}{3}$. In particular, if Λ is an exact cocycle, then $d_{\check{H}}(\Omega, \Lambda) \geq r(\Omega)/3$.

Proof. Let Λ be an exact cocycle, and let $\delta > d_Z(\Omega, \Lambda)$. It is enough to show that $\delta \geq r(\Omega)/3$. Equivalently, it is enough to show that, for every $(ijk) \in N(\mathcal{U})$ and $y \in U_k \cap U_j \cap U_i$, we have $\|\Omega_{ij}(y)\Omega_{jk}(y) - \Omega_{ik}(y)\| < 3\delta$. This follows from Lemma 4.29. □

5. Discrete approximate vector bundles

In this section, we specialize the notions of approximate cocycle and approximate local trivialization to a certain open cover associated to any simplicial complex. This gives us the notions of discrete approximate cocycle and of discrete approximate local trivializations.

We show in Proposition 5.7 that any vector bundle over a compact triangulable space can be represented by a discrete approximate cocycle over a sufficiently fine triangulation of the space. In Section 5.2, we study the problem of reconstructing a vector bundle from finite samples as a discrete approximate local trivialization and as a discrete approximate cocycle. We prove in Theorem 5.15 that this is possible provided the classifying map of the vector bundle we wish to reconstruct is sufficiently regular and that we are given a sufficiently dense sample.

5.1. Discrete approximate vector bundles over simplicial complexes

We introduce two notions of discrete approximate vector bundle over a simplicial complex. These notions of discrete approximate vector bundle induce approximate vector bundles over the geometric realization of the simplicial complex.

Since this will be relevant in Section 6, we define discrete approximate cocycles with values in an arbitrary metric group G . Fix a simplicial complex K .

Definition 5.1. A **discrete ε -approximate cocycle** on K with values in G consists of, for every ordered 1-simplex $(ij) \in K$, an element $\Omega_{ij} \in G$ such that, for every ordered 2-simplex $(ijk) \in K$, we have $d_G(\Omega_{ij}\Omega_{jk}, \Omega_{ik}) < \varepsilon$, and such that $\Omega = \{\Omega_{ij}\}_{(ij) \in K}$ is **symmetric**, that is, we have $\Omega_{ij} = \Omega_{ji}$. A **discrete exact cocycle** consists of the same data but subject to $\Omega_{ij}\Omega_{jk} = \Omega_{ik}$.

We denote the set of discrete ε -approximate cocycles on a simplicial complex K with values in G by $DZ^1_\varepsilon(K; G)$.

Definition 5.2. Let $i \in K$ be a vertex. Let $st(i)$ be the geometric realization of the open star of i seen as a vertex of the geometric realization $|K|$. The **star cover** of $|K|$ consists of the family of open sets $\{st(i)\}_{i \in K}$. Denote the star cover of $|K|$ by st_K .

Note that there is a canonical isomorphism of simplicial complexes $N(st_K) \cong K$ that maps a vertex $i \in K$ to itself.

Construction 5.3. Let Ω be a discrete ε -approximate cocycle on a simplicial complex K with values in G . Define, for each $(ij) \in K$, a continuous map $\text{st}(j) \cap \text{st}(i) \rightarrow G$ that is constantly Ω_{ij} . This defines a natural map $\text{DZ}_\varepsilon^1(K; G) \rightarrow Z_\varepsilon^1(\text{st}_K; G)$.

Note that the map $\text{DZ}_\varepsilon^1(K; G) \rightarrow Z_\varepsilon^1(\text{st}_K; G)$ is injective. With this in mind, we endow $\text{DZ}_\varepsilon^1(K; G)$ with the metric d_Z , and interpret the map $\text{DZ}_\varepsilon^1(K; G) \rightarrow Z_\varepsilon^1(\text{st}_K; G)$ as an embedding of metric spaces.

Definition 5.4. A **discrete ε -approximate local trivialization** on a simplicial complex K consists of a frame $\Phi_i \in V(d)$ for every $(i) \in K$ such that, for every $(ij) \in K$, there exists $\Omega_{ij} \in O(d)$ such that $\|\Phi_i \Omega_{ij} - \Phi_j\| < \varepsilon$. A **discrete exact local trivialization** consists of the same data but subject to $\Phi_i \Omega_{ij} = \Phi_j$.

Denote the set of discrete ε -approximate local trivializations on a simplicial complex K by $\text{DT}_\varepsilon(K; d)$. In this discrete case too, the witness Ω that Φ is a discrete ε -approximate local trivialization is not part of the data of the approximate local trivialization.

Construction 5.5. Let Φ be a discrete ε -approximate local trivialization on K . Define, for each $(i) \in K$, a map $\text{st}(i) \rightarrow V(d)$ that is constantly Φ_i . This defines a natural map $\text{DT}_\varepsilon(K; d) \rightarrow \tau_\varepsilon(\text{st}_K; d)$.

Remark 5.6. Using Construction 4.8 we obtain a map $\text{DT}_\varepsilon(K; d) \rightarrow \text{DZ}_{3\varepsilon}^1(K; O(d))$. If K is finite, this map is algorithmic since the minimization problem

$$\min_{\Omega \in O(d)} \|\Phi_i \Omega - \Phi_j\|$$

can be solved by using the polar decomposition (Appendix A.2).

The next result guarantees that any vector bundle on a compact triangulable space can be encoded as a discrete approximate cocycle on a sufficiently fine triangulation of the space.

Proposition 5.7. *Let $E \rightarrow B$ be a vector bundle over a compact triangulable space B , and let $\varepsilon \leq 3/8$. There exists a triangulation K of B and a discrete ε -approximate cocycle $\Omega \in \text{DZ}_\varepsilon^1(K; O(d))$ such that $\pi_*(\text{cl}(\Omega))$ represents the vector bundle $E \rightarrow B$.*

Proof. Let S be a finite simplicial complex such that $|S| \cong B$. Without loss of generality, we assume $|S| = B$. Since the star cover of S consists of contractible sets, the vector bundle $E \rightarrow |S|$ trivializes over st_S and thus is represented by an exact cocycle $\Lambda \in Z^1(\text{st}_S; O(d))$. Moreover, since the closed stars $\text{st}(i)$ are also contractible, for each $(ij) \in S$, the continuous map $\Lambda_{ij} : \text{st}(j) \cap \text{st}(i) \rightarrow O(d)$ can be taken such that it extends to $\overline{\text{st}(j) \cap \text{st}(i)}$, which is a closed set. Pick a metric that metrizes $|S|$, which must exist since S is a finite simplicial complex. It follows that the maps Λ_{ij} are uniformly continuous, and thus there exists $\delta > 0$ such that for every $(ij) \in S$ and every $T \subseteq \text{st}(j) \cap \text{st}(i)$ of diameter less than δ , the diameter of $\Lambda_{ij}(T) \subseteq O(d)$ is less than $\varepsilon/3$.

For $n \in \mathbb{N}$, let $S^{(n)}$ denote the n th barycentric subdivision of S . Note that, for every $n \in \mathbb{N}$, the star cover $\text{st}_{S^{(n+1)}}$ refines the star cover $\text{st}_{S^{(n)}}$. Choose a refinement map $r^n : \text{st}_{S^{(n+1)}} \rightarrow \text{st}_{S^{(n)}}$ for each n , and let $\Lambda^n \in Z_\varepsilon^1(\text{st}_{S^{(n)}}; O(d))$ denote the restriction of Λ along these refinements, obtained by using Construction 4.19. Note that, by Lemma 4.20, Λ still represents the original vector bundle $E \rightarrow |S|$.

As n goes to ∞ , the maximum of the diameters $\max_{i \in S^{(n)}} \text{diam}(\text{st}(i))$ goes to 0, since the diameter of simplices goes to 0 uniformly, as S has finitely many simplices. In particular, there exists n_0 such that the diameter of the image of $\Lambda_{ij}^{n_0}$ is less than $\varepsilon/3$ for every $(ij) \in S^{(n_0)}$ since the original Λ consists of uniformly continuous maps. Let $K = S^{n_0}$.

For every $(ij) \in K$, pick a matrix Ω_{ij} in the image of $\Lambda_{ij}^{n_0}$ in such a way that $\Omega_{ij} = \Omega_{ji}^t$. The matrices Ω_{ij} assemble into a cochain $\Omega \in C^1(\text{st}_K; O(d))$ such that $d_Z(\Omega, \Lambda^{n_0}) < \varepsilon/3$. It follows from Lemma 4.29 that Ω is a ε -approximate cocycle, and since it is constant on each intersection, we have $\Omega \in \text{DZ}_\varepsilon^1(K; O(d))$.

To conclude, note that, by Theorem 4.21, we have that $\text{cl}(\Omega)$ and $\text{cl}(\Lambda^{n_0}) = \text{cl}(\Lambda)$ become equal in $\left[B, \text{Gr}(d)^{\sqrt{2}(\varepsilon+\varepsilon/3)} \right]$. Since $\sqrt{2}(\varepsilon + \varepsilon/3) \leq \sqrt{2}/2$, by assumption, we have that $\pi_*(\text{cl}(\Omega)) = \text{cl}(\Lambda)$, as required. \square

5.2. Reconstruction of vector bundles from finite samples

Let $\delta \geq 0$ and $\ell > 0$. A function $f : X \rightarrow Y$ between metric spaces is a δ -approximate ℓ -Lipschitz map if, for every $x, x' \in X$, we have $\ell d_X(x, x') + \delta \geq d_Y(f(x), f(x'))$.

Definition 5.8. Let $X \subseteq \mathbb{R}^N$. The Čech complex of X at distance scale $\varepsilon > 0$, denoted $\check{C}(X)(\varepsilon)$, consists of the simplicial complex given by the nerve of the cover $\{B(x, \varepsilon)\}_{x \in X}$ of X^ε .

Note that the cover $\{B(x, \varepsilon)\}_{x \in X}$ is indexed by the elements of X , so the 0-simplices of $\check{C}(X)(\varepsilon)$ consist of the elements of X . As a set, $|\check{C}(X)(\varepsilon)|$ consists of formal linear combinations

$$p = \sum_{x \in X} c_x [x]$$

such that $S_p = \{x \in X : c_x > 0\}$ is a finite set with the property that the intersection $\bigcap_{x \in S_p} B(x, \varepsilon) \subseteq \mathbb{R}^N$ is nonempty. We will often write $\check{C}(X)(\varepsilon)$ for the geometric realization $|\check{C}(X)(\varepsilon)|$.

Construction 5.9. Let $X \subseteq \mathbb{R}^N$ and let $\varepsilon > 0$, $\delta \geq 0$ and $\ell > 0$. Assume that X is finite. Let $Y \subseteq \mathbb{R}^M$, and let $f : X \rightarrow Y$ be a δ -approximate ℓ -Lipschitz map with respect to the distances induced by the Euclidean norm $\| \cdot \|_2$. Define a continuous map

$$\begin{aligned} \check{C}(f)(\varepsilon) : \check{C}(X)(\varepsilon) &\rightarrow Y^{2\ell\varepsilon+\delta} \subseteq \mathbb{R}^M \\ \sum_{x \in X} c_x [x] &\mapsto \sum_{x \in X} c_x f(x). \end{aligned} \quad \triangleleft$$

The map $\check{C}(f)$ is well defined since, if $p = \sum_{x \in X} c_x [x]$ is such that $c_{x_0} > 0$, then $\|f(z) - f(x_0)\|_2 < 2\ell\varepsilon + \delta$ for every $z \in X$ such that $c_z > 0$, since, in that case, $\|x_0 - z\| < 2\varepsilon$, as the balls $B(x_0, \varepsilon)$ and $B(z, \varepsilon)$ must intersect.

Lemma 5.10. Let $X \subseteq \mathbb{R}^N$ and let $\varepsilon > 0$, $\delta \geq 0$ and $\ell > 0$. Assume that X is finite. Let $f : X \rightarrow \text{Gr}(n, d) \subseteq \mathbb{R}^{n \times n}$ be a δ -approximate ℓ -Lipschitz map with respect to the Euclidean norm and the Frobenius norm. The $(2\ell\varepsilon + \delta)$ -approximate classifying map $\check{C}(f)(\varepsilon) : \check{C}(X)(\varepsilon) \rightarrow \text{Gr}(d, n)^{2\ell\varepsilon+\delta}$ of Construction 5.9 can be represented by a discrete approximate local trivialization, in the sense that there exists a partition of unity φ of the star cover of $\check{C}(X)(\varepsilon)$ and $\Phi \in \text{DT}_{2\ell\varepsilon+\delta}(\check{C}(X)(\varepsilon); d)$ such that $\text{av}^\varphi(\Phi) = \check{C}(f)(\varepsilon)$.

Moreover, there is a discrete $(3(2\ell\varepsilon + \delta))$ -approximate cocycle $w(\Phi)$ such that $\text{cl}(\Omega)$ is equal to $\check{C}(f)(\varepsilon)$ in $\left[\check{C}(X)(\varepsilon), \text{Gr}(d, n)^{3\sqrt{2}(2\ell\varepsilon+\delta)} \right]$.

Proof. The second claim is a consequence of the first one and Lemma 4.12, where w is the map defined in Lemma 4.9.

For the first claim, for each $x \in X$, let $\Phi_x \in V(d, n)$ be an orthonormal basis of the subspace of \mathbb{R}^n spanned by $f(x) \in \text{Gr}(d, n)$. Since f is a δ -approximate ℓ -Lipschitz map, the family $\{\Phi_x\}_{x \in X}$ constitutes a discrete $(2\ell\varepsilon + \delta)$ -approximate local trivialization on the simplicial complex $\check{C}(X)(\varepsilon)$, by Lemma A.9. By taking the partition of unity φ subordinate to the star cover of $\check{C}(X)(\varepsilon)$ given by $\varphi_x(\sum_{x \in X} c_x [x]) = c_x$, we see that $\text{av}^\varphi(\Phi) = \check{C}(f)(\varepsilon)$. \square

The following result is well known; see, for example, [26, Corollary 4G.3].

Lemma 5.11 (Nerve lemma). Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of a paracompact topological space B , and let φ be a partition of unity subordinate to B . If \mathcal{U} has the property that any finite intersection of

its elements is either contractible or empty, then the map $B \rightarrow |N(\mathcal{U})|$ that sends y to $\sum_{i \in I} \varphi_i(y)[i]$ is well defined, continuous and a homotopy equivalence.

Corollary 5.12. *Let $X \subseteq \mathbb{R}^N$ be a finite subset, let $\varepsilon > 0$ and let $\varphi = \{\varphi_x\}_{x \in X}$ be a partition of unity subordinate to $\{B(x, \varepsilon)\}_{x \in X}$. Then, the following map is a homotopy equivalence:*

$$R^\varphi : X^\varepsilon \rightarrow \check{C}(X)(\varepsilon)$$

$$z \mapsto \sum_{x \in X} \varphi_x(z)[x]. \quad \square$$

Our reconstruction theorem for vector bundles builds on the following result by Niyogi, Smale and Weinberger, which allows one to recover the homotopy type of a compact manifold smoothly embedded into \mathbb{R}^N from a sufficiently close and dense sample. In the result, d_H denotes the Hausdorff distance.

Proposition 5.13 [49, Proposition 7.1]. *Let $\mathcal{M} \subseteq \mathbb{R}^N$ be a smoothly embedded compact manifold with $\tau = \text{reach}(\mathcal{M}) > 0$. Let $P \subseteq \mathbb{R}^N$ such that $d_H(P, \mathcal{M}) < \varepsilon < (3 - \sqrt{8})\tau$ and let*

$$\alpha \in \left(\frac{(\varepsilon + \tau) - \sqrt{\varepsilon^2 + \tau^2 - 6\tau\varepsilon}}{2}, \frac{(\varepsilon + \tau) + \sqrt{\varepsilon^2 + \tau^2 - 6\tau\varepsilon}}{2} \right),$$

which is a nonempty open interval. Then $\mathcal{M} \subseteq P^\alpha$ and the inclusion $\mathcal{M} \rightarrow P^\alpha$ is a homotopy equivalence.

We are now ready to prove the reconstruction theorem for vector bundles. Before doing so, we give a short remark about representing vector bundles by Lipschitz maps.

Remark 5.14. In [53, Proposition 3.1] it is shown that any rank- d vector bundle on a compact metric space X can be represented by a Lipschitz map $X \rightarrow \text{Gr}(d, n)$ for some n , where $\text{Gr}(d, n)$ is seen as a subspace of the space of square matrices $\mathbb{R}^{n \times n}$ with the operator norm. It follows that a vector bundle on X can also be represented by a Lipschitz map $X \rightarrow \text{Gr}(d, n)$ where now we use the Frobenius distance on $\text{Gr}(d, n)$, as we do in this paper. This motivates the assumptions made in the following result.

Theorem 5.15. *Let $\mathcal{M} \subseteq \mathbb{R}^N$ be a smoothly embedded compact manifold and let $f : \mathcal{M} \rightarrow \text{Gr}(d, n)$ be an ℓ -Lipschitz map with respect to the Euclidean distance on \mathcal{M} and the Frobenius distance on $\text{Gr}(d, n)$. Assume that $\text{reach}(\mathcal{M}) = \tau > 0$. Let $P \subseteq \mathbb{R}^N$ be a finite set and let $g : P \rightarrow \text{Gr}(d, n)$ be a function. Let $\varepsilon, \delta > 0$ be such that*

- for every $x \in \mathcal{M}$ there exists $p \in P$ such that $\|p - x\|_2 < \varepsilon$;
- for every $p \in P$ there exists $x \in \mathcal{M}$ such that $\|p - x\|_2 < \varepsilon$ and $\|g(p) - f(x)\| < \delta$

so that g is a $2(\delta + \ell\varepsilon)$ -approximate ℓ -Lipschitz map. If $\varepsilon < (3 - \sqrt{8})\tau$, then, for every

$$\alpha \in \left(\frac{(\varepsilon + \tau) - \sqrt{\varepsilon^2 + \tau^2 - 6\tau\varepsilon}}{2}, \frac{(\varepsilon + \tau) + \sqrt{\varepsilon^2 + \tau^2 - 6\tau\varepsilon}}{2} \right) \cap \left(0, \frac{\sqrt{2}/2 - 2\delta - 2\ell\varepsilon}{2\ell} \right),$$

there is a homotopy commutative diagram as follows, in which the vertical maps are homotopy equivalences:

$$\begin{array}{ccc} \mathcal{M} & \xrightarrow{f} & \text{Gr}(n, d) \\ \downarrow & & \downarrow \\ \check{C}(P)(\alpha) & \xrightarrow{\check{C}(g)} & \text{Gr}(n, d)^{2(\ell\alpha + \ell\varepsilon + \delta)}. \end{array}$$

Note that the interval to which α must belong to is nonempty as long as ε and δ are sufficiently small, a condition that depends only on τ and ℓ .

Proof. We start by showing that g is indeed a $2(\delta + \ell\varepsilon)$ -approximate ℓ -Lipschitz map so that the bottom map of the diagram in the statement is well defined. To do this, note that, if $p, q \in P$, then there exist $x, y \in \mathcal{M}$ with $\|p - x\|_2, \|q - y\|_2 < \varepsilon$ and $\|g(p) - f(x)\|, \|g(q) - f(y)\| < \delta$. Since f is ℓ -Lipschitz, we have $\|g(p) - g(q)\| \leq 2\delta + \ell\|x - y\| \leq 2\delta + \ell 2\varepsilon + \ell\|p - q\|$, as required.

The homotopy equivalence $\mathcal{M} \rightarrow \check{C}(P)(\alpha)$ is given by composing the inclusion $\mathcal{M} \subseteq P^\alpha$ with the map $R^\varphi : P^\alpha \rightarrow \check{C}(P)(\alpha)$ for a choice of partition of unity φ subordinate to $\{B(p, \alpha)\}_{p \in P}$. The map R^φ is a homotopy equivalence by Corollary 5.12. The fact that the inclusion $\mathcal{M} \rightarrow P^\alpha$ is well defined and a homotopy equivalence is the content of Proposition 5.13, whose hypotheses are satisfied since our conditions imply that $d_H(P, \mathcal{M}) < \varepsilon$.

The map $\text{Gr}(n, d) \rightarrow \text{Gr}(n, d)^{2(\ell\alpha + \ell\varepsilon + \delta)}$ is simply the inclusion and it is a homotopy equivalence by Proposition A.12 since $\alpha < (\sqrt{2}/2 - 2\delta - 2\ell\varepsilon)/(2\ell)$.

To conclude the proof, we must show that the diagram in the statement commutes up to homotopy. For this, let $z \in \mathcal{M}$. We have $\check{C}(g)(R^\varphi(z)) = \check{C}(g)\left(\sum_{p \in P} \varphi_p(z)[p]\right) = \sum_{p \in P} \varphi_p(z)g(p)$. Consider the linear path $\beta f(z) + (1 - \beta) \sum_{p \in P} \varphi_p(z)g(p)$ for $\beta \in [0, 1]$. It suffices to show that it is included in $\text{Gr}(n, d)^{2\ell\alpha + 2\ell\varepsilon + 2\delta}$ for all $\beta \in [0, 1]$, and we will show that it is at distance less than $\ell\alpha + \ell\varepsilon + \delta$ from $f(z)$. We compute

$$\begin{aligned} & \left\| f(z) - \left(\beta f(z) + (1 - \beta) \sum_{p \in P} \varphi_p(z)g(p) \right) \right\|_2 \\ &= \left\| \beta f(z) + (1 - \beta)f(x) - \left(\beta f(z) + (1 - \beta) \sum_{p \in P} \varphi_p(z)g(p) \right) \right\|_2 \\ &\leq (1 - \beta) \sum_{p \in P} \varphi_p(z) \|f(z) - g(p)\|_2 \\ &< (1 - \beta) \sum_{x \in X} \varphi_x(z) (\ell\alpha + \ell\varepsilon + \delta) \leq \ell\alpha + \ell\varepsilon + \delta, \end{aligned}$$

where the strict inequality comes from the fact that, if $\varphi_p(z)$ is nonzero, then $z \in B(p, \alpha)$, and thus there exists $x \in \mathcal{M}$ such that $\|f(z) - g(p)\|_2 \leq \|f(z) - f(x)\| + \|f(x) - g(p)\| < \ell\|z - x\|_2 + \delta \leq \ell(\|z - p\|_2 + \|p - x\|_2) + \delta < \ell\alpha + \ell\varepsilon + \delta$. \square

We conclude this section with a few remarks.

Remark 5.16. For simplicity, we have proven Theorem 5.15 using the Čech complex. A similar result can be obtained for the Vietoris–Rips complex, using, for instance, the results of [2].

Remark 5.17. As proven in Lemma 5.10, the map reconstructed in Theorem 5.15 has a combinatorial description using a discrete approximate local trivialization that in turn induces a discrete approximate cocycle. This discrete approximate cocycle can be used to compute characteristic classes combinatorially, which is the subject of the next section.

6. Effective computation of characteristic classes

In this section, we present three algorithms to compute, respectively, the first two Stiefel–Whitney classes of an approximate vector bundle given by an approximate $O(d)$ -cocycle and the Euler class of an oriented approximate vector bundle of rank 2 given by an approximate $SO(2)$ -cocycle. The algorithms are based on well-known results which say that the characteristic classes we consider are obstructions to lifting the structure group of the cocycle to certain other Lie groups. The difficulty is in showing that these algorithms can be extended in a stable and consistent way to ε -approximate cocycles, provided ε is sufficiently small. Throughout the section, we will make use of basic Riemannian geometry; a reference for this topic is [39].

In Section 6.1, we recall two standard constructions used to change the coefficient group of a Čech cocycle and we extend them to approximate cocycles. In Section 6.2, we give the algorithm for the first Stiefel–Whitney class, in Section 6.3 we give the algorithm for the Euler class and in Section 6.4 we give the algorithm for the second Stiefel–Whitney class.

6.1. Change of coefficients

In this section, we will make use of basic Čech cohomology with coefficients in a sheaf of abelian groups. We recall the essential components now; for an introduction to the subject, see, for example, [66, Chapter 5].

Let \mathcal{U} be a cover of a topological space B . For A an abelian group and $n \in \mathbb{N}$, we let $\check{H}^n(\mathcal{U}; A)$ denote the n th Čech cohomology group of (B, \mathcal{U}) with coefficients in the sheaf of locally constant functions with values in A [66, p. 201]. This cohomology group is a quotient of the subgroup of cocycles of $C^n(\mathcal{U}; A)$, which is the abelian group of locally constant functions defined on all $(n + 1)$ -fold intersections $U_{i_1} \cap \dots \cap U_{i_n} \rightarrow A$. As usual, the Čech cohomology of B with values in A is defined as $\check{H}^n(B; A) = \text{colim}_{\mathcal{U} \text{ cover}} \check{H}^n(\mathcal{U}; A)$. It is well known that, when B is paracompact and locally contractible, there is a natural isomorphism $\check{H}^n(B; A) \cong H^n(B; A)$, where the right-hand side denotes singular cohomology with coefficients in A .

Since we use these constructions only for $n = 1, 2$, we elaborate on these two cases. For $n = 1$, the Čech cohomology is precisely the one we introduced in Section 2.2.4, where the group A is endowed with the discrete topology. For $n = 2$, the 2-cocycles $Z^2(\mathcal{U}; A)$ consist of families of locally constant functions $\{\Gamma_{ijk} : U_k \cap U_j \cap U_i \rightarrow A\}_{(ijk) \in N(\mathcal{U})}$ such that, for every $(ijkl) \in N(\mathcal{U})$ and every $y \in U_l \cap U_k \cap U_j \cap U_i$, we have

$$\Gamma_{ijk}(y)\Gamma_{ijl}(y)^{-1}\Gamma_{ikl}(y)\Gamma_{jkl}(y)^{-1} = 1_A,$$

where we are writing the operations of the abelian group A multiplicatively. The operation on $Z^2(\mathcal{U}; A)$ is pointwise multiplication. The cohomology group $\check{H}^2(\mathcal{U}; A)$ is the quotient of $Z^2(\mathcal{U}; A)$ by the subgroup of 2-cocycles of the form $y \mapsto \Omega_{ij}(y)\Omega_{jk}(y)\Omega_{ki}(y)$, for $\Omega \in C^1(\mathcal{U}; A)$.

Let G and H be topological groups, and let $\zeta : G \rightarrow H$ be a continuous group morphism. The map ζ induces a map $Z^1(\mathcal{U}; G) \rightarrow Z^1(\mathcal{U}; H)$, simply by applying ζ pointwise to a cocycle. This map induces a well-defined map $\check{H}^1(\mathcal{U}; G) \rightarrow \check{H}^1(\mathcal{U}; H)$. A bit more interestingly, given a central extension of groups $1 \rightarrow F \rightarrow G \rightarrow H \rightarrow 1$, there is a well-defined so-called *connecting morphism* $\check{H}^1(\mathcal{U}; H) \rightarrow \check{H}^2(\mathcal{U}; F)$. For a short introduction to these concepts, see [38, Appendix A].

In this section, we generalize these two constructions to the case of approximate cocycles when F, G and H are well-behaved metric groups. We start by generalizing the first construction.

Construction 6.1. Let $\zeta : G \rightarrow H$ be a continuous group morphism between topological groups. Given $\Omega \in C^1(\mathcal{U}; G)$, define a cochain $\zeta(\Omega)$ with values in H by $\zeta(\Omega)_{(ij)}(y) = \zeta(\Omega_{(ij)}(y))$ for every $(ij) \in N(\mathcal{U})$ and $y \in U_j \cap U_i$.

Lemma 6.2. Let G and H be metric groups and let $\zeta : G \rightarrow H$ be an ℓ -Lipschitz group morphism. Let $\varepsilon \in [0, \infty]$. Construction 6.1 induces maps $\zeta : Z^1_{\varepsilon}(\mathcal{U}; G) \rightarrow Z^1_{\ell\varepsilon}(\mathcal{U}; H)$ and $\zeta : \check{H}^1_{\varepsilon}(\mathcal{U}; G) \rightarrow \check{H}^1_{\ell\varepsilon}(\mathcal{U}; H)$ and these maps are ℓ -Lipschitz with respect to the distances d_Z and $d_{\check{H}}$, respectively.

In particular, if the infimum over all distances between distinct elements of H is bounded below by δ and $\ell\varepsilon \leq \delta$, then ζ induces a map $\zeta : \check{H}^1_{\varepsilon}(\mathcal{U}; G) \rightarrow \check{H}^1(\mathcal{U}; H)$, such that, if $\Omega, \Omega' \in \check{H}^1_{\varepsilon}(\mathcal{U}; G)$ satisfy $d_Z(\Omega, \Omega') < \delta/\ell$, then $\zeta(\Omega) = \zeta(\Omega') \in \check{H}^1(\mathcal{U}; H)$.

Proof. We start by checking that, if $\Omega \in Z^1_{\varepsilon}(\mathcal{U}; G)$, then $\zeta(\Omega)$ is an $\ell\varepsilon$ -approximate cocycle. For this, let $(ijk) \in N(\mathcal{U})$ and let $y \in U_k \cap U_j \cap U_i$. We have $d_H(\zeta(\Omega)_{ij}(y)\zeta(\Omega)_{jk}(y), \zeta(\Omega)_{ik}(y)) = d_H(\zeta(\Omega_{ij}(y)\Omega_{jk}(y)), \zeta(\Omega_{ik}(y))) \leq \ell d_G(\Omega_{ij}(y)\Omega_{jk}(y), \Omega_{ik}(y)) < \ell\varepsilon$, using the fact that ζ is an ℓ -Lipschitz group morphism. The fact that the maps are ℓ -Lipschitz is clear.

To see that ζ descends to approximate cohomology, note that, if $\Theta \in C^0(\mathcal{U}; G)$, then we can define $\zeta(\Theta) \in C^0(\mathcal{U}; H)$ by $\zeta(\Theta)_i(y) = \zeta(\Theta_i(y))$ for all $i \in N(\mathcal{U})$ and $y \in U_i$ and that, with this definition, we have $\zeta(\Theta) \cdot \zeta(\Omega) = \zeta(\Theta \cdot \Omega)$, for every $\Omega \in Z^1_\varepsilon(\mathcal{U}; G)$.

The second claim is a consequence of the first one. □

We now generalize the connecting morphism construction.

Construction 6.3. Let $\zeta : G \rightarrow H$ be a continuous and surjective group morphism between metric groups and let $\varepsilon \geq 0$. Let F be the kernel of ζ and assume that F is locally compact and discrete in G so that, in particular, $G \rightarrow H$ is a covering map. Suppose that \mathcal{U} is a cover of a topological space with the property that nonempty binary intersections of elements of \mathcal{U} are locally path connected and simply connected. Let $\Omega \in C^1(\mathcal{U}; H)$. Given $(ij) \in N(\mathcal{U})$, choose a continuous lift $\Lambda_{ij} : U_j \cap U_i \rightarrow G$ of $\Omega_{ij} : U_j \cap U_i \rightarrow H$ such that $\Lambda_{ij} = (\Lambda_{ji})^{-1}$. Finally, for $(ijk) \in N(\mathcal{U})$ and $y \in U_k \cap U_j \cap U_i$, let $\Gamma_{ijk}(y) \in F$ be a closest point of F to $\Lambda_{ij}(y)\Lambda_{jk}(y)\Lambda_{ki}(y)$.

A priori, the maps Γ_{ijk} of Construction 6.3 do not necessarily assemble into a 2-cocycle with values in F since the maps may not even be continuous. We now give conditions under which the maps Γ_{ijk} constitute a 2-cocycle. In order to do this, we need to introduce some definitions.

Definition 6.4. Let M be a metric space. The **systole** of M , denoted $\text{sys}(M)$, is the infimum of the lengths of all nonnullhomotopic loops of M .

Recall from, for example, [39, Chapter 2], that any connected Riemannian manifold M can be endowed with the **geodesic distance** where the distance between two points is taken to be the infimum of the lengths of all piecewise regular (i.e., with nonzero velocity) paths between the two points. If the manifold is not connected, then the same construction gives an **extended distance**, meaning a distance that can also take the value ∞ . Whenever we endow a Riemannian manifold with the geodesic distance, we are referring to this extended distance. Finally, recall that, if the manifold is complete (as are all the examples we consider here), then the geodesic distance between two points can be calculated as the infimum of the length of all geodesics between the two points [39, Corollary 6.21].

Before being able to give conditions under which the connecting morphism is well behaved, we need to prove the following technical lemma.

Lemma 6.5. *Let $F \subseteq G$ be an isometric inclusion of a discrete subgroup into a metric group. Let $2r$ be the infimum of the distances between distinct elements of F . Let $\langle - \rangle : F^r \rightarrow F$ denote the projection to the closest element of F , which is well-defined and continuous.*

1. *If $a \in F^s, b \in F^t, c \in F^u$ with $s + t + u \leq r$ then $\langle a \rangle \langle b \rangle \langle c \rangle = \langle abc \rangle$.*
2. *If $a \in F^r$ and $b \in F$, then $\langle ab \rangle = \langle a \rangle b$.*
3. *If $a, c \in G, b \in F^s$, and $a \langle b \rangle c \in F^t$ with $s + t \leq r$, then $\langle a \langle b \rangle c \rangle = \langle abc \rangle$.*

Proof. As we will use this for each claim, start by noting that, if $g \in G, h \in F$ and $d_G(g, h) < r$, then $\langle g \rangle = h$. Since multiplication is an isometry, we have $d_G(abc, \langle a \rangle bc) < s$. Using this idea two more times and the triangle inequality, we deduce $d_G(abc, \langle a \rangle \langle b \rangle \langle c \rangle) < s + t + u \leq r$ and the first claim follows. For the second claim, note that $d_G(ab, \langle a \rangle b) < r$. Since $\langle a \rangle b \in F$, the second claim follows. Finally, $d_G(\langle a \rangle \langle b \rangle c, a \langle b \rangle c) < s$ and $d_G(a \langle b \rangle c, abc) < t$, so $d_G(\langle a \rangle \langle b \rangle c, abc) < s + t \leq r$ and the third claim follows. □

Theorem 6.6. *Let $1 \rightarrow F \rightarrow G \rightarrow H \rightarrow 1$ be a central extension of Lie groups with H compact and F discrete. Fix a bi-invariant Riemannian metric on H and use it to metrize H and G with the geodesic distance and F with the distance inherited from G . Let \mathcal{U} be a cover of a topological space B such that each set and each nonempty binary intersection is locally path connected and simply connected. Let $\varepsilon \leq \text{sys}(H)/8$. Then, Construction 6.3 induces maps $\partial : Z^1_\varepsilon(\mathcal{U}; H) \rightarrow Z^2(\mathcal{U}; F)$ and $\partial : \check{H}^1_\varepsilon(\mathcal{U}; H) \rightarrow \check{H}^2(\mathcal{U}; F)$.*

The second map is independent of any choice of lift and is stable in the sense that if $\Omega, \Omega' \in \check{H}^1_\varepsilon(\mathcal{U}; H)$ are such that $d_{\check{H}}(\Omega, \Omega') < \text{sys}(H)/8$, then $\partial(\Omega) = \partial(\Omega') \in \check{H}^2(\mathcal{U}; F)$.

Proof. Let $2r$ denote the infimum over the distances between distinct elements of F . By Lemma A.19, we have $\text{sys}(H)/2 \leq r$ and thus $\varepsilon \leq r/4$. For $x \in F^r$, let $\langle x \rangle \in F$ denote its closest point in F , which is well defined and continuous.

We start by showing that, for $(ijk) \in N(\mathcal{U})$, the map $\Gamma_{ijk} : U_k \cap U_j \cap U_i \rightarrow F$ is continuous. It suffices to show that $\Lambda_{ij}(y)\Lambda_{jk}(y)\Lambda_{ki}(y) \in F^r$, as $\Gamma_{ijk}(y) = \langle \Lambda_{ij}(y)\Lambda_{jk}(y)\Lambda_{ki}(y) \rangle$. By definition of the metrics on G and H , the distance from $\Lambda_{ij}(y)\Lambda_{jk}(y)\Lambda_{ki}(y)$ to F is equal to the distance from $\Omega_{ij}(y)\Omega_{jk}(y)\Omega_{ki}(y)$ to the identity of H , which is less than $\varepsilon \leq r/4$, by assumption.

We now show that Γ is an H -valued 2-cocycle. Fix $y \in U_l \cap U_k \cap U_j \cap U_i$ and let us write (ij) for $\Lambda_{ij}(y)$ and likewise for the other elements of G . We must show that

$$\langle (ij)(jk)(ki) \rangle \langle (ij)(jl)(li) \rangle^{-1} \langle (ik)(kl)(li) \rangle \langle (jk)(kl)(lj) \rangle^{-1} = 1_F.$$

Note that, although F is abelian, we are using multiplicative notation to be consistent with the fact that G need not be abelian. We now compute:

$$\begin{aligned} & \langle (ij)(jk)(ki) \rangle \langle (ij)(jl)(li) \rangle^{-1} \langle (ik)(kl)(li) \rangle \langle (jk)(kl)(lj) \rangle^{-1} \\ &= \langle (ij)(jk)(ki) \rangle \langle (il)(lj)(ji) \rangle \langle (ik)(kl)(li) \rangle \langle (jl)(lk)(kj) \rangle \quad (\text{taking inverse is an isometry}) \\ &= \langle (ij)(jk)(ki) \rangle \langle (ik)(kl)(li) \rangle \langle (il)(lj)(ji) \rangle \langle (jl)(lk)(kj) \rangle \quad (F \text{ is central}) \\ &= \langle (ij)(jk)(ki)(ik)(kl)(li)(il)(lj)(ji) \rangle \langle (jl)(lk)(kj) \rangle \quad (\text{Lem. 6.5(1), with } s, t, u = \varepsilon \leq r/4) \\ &= \langle (ij)(jk)(kl)(lj)(ji) \rangle \langle (jl)(lk)(kj) \rangle \quad (\text{cancellations}) \\ &= \langle (ij)(jk)(kl)(lj)(ji) \rangle \langle (jl)(lk)(kj) \rangle \quad (\text{Lem. 6.5(2), with } s = 3\varepsilon \leq 3/4r) \\ &= \langle (ij)(jk)(kl)(lj) \rangle \langle (jl)(lk)(kj) \rangle \langle (ji) \rangle \quad (F \text{ is central}) \\ &= \langle (ij)(jk)(kl)(lj)(jl)(lk)(kj)(ji) \rangle \quad (\text{Lem. 6.5(3), with } s = \varepsilon, t = 3\varepsilon) \\ &= 1_F \quad (\text{cancellations}). \end{aligned}$$

We now prove that the composite $Z_\varepsilon^1(\mathcal{U}; H) \rightarrow Z^2(\mathcal{U}; F) \rightarrow \check{H}^2(\mathcal{U}; F)$ is independent of the choice of lifts Λ . In order to see this note that, if $\Lambda_{ij}, \Lambda'_{ij} : U_j \cap U_i \rightarrow G$ are two lifts of $\Omega_{ij} : U_j \cap U_i \rightarrow H$, then they differ by a function $f_{ij} : U_j \cap U_i \rightarrow F$. Since F is central, it follows that Γ_{ijk} and Γ'_{ijk} differ by $f_{ij}f_{jk}f_{ki}$, using Lemma 6.5(2).

Now, let $\Theta \in C^0(\mathcal{U}; H)$ and let $\Pi \in C^0(\mathcal{U}; G)$ be a lift of it, which exists since the elements of \mathcal{U} are simply connected. Let us write (i) for $\Pi_i(y)$. A lift for $\Theta \cdot \Omega$ is given by $\Pi \cdot \Lambda$ and the 2-cochain in this case is

$$\begin{aligned} \langle (i)(ij)(j)^{-1}(j)(jk)(k)^{-1}(k)(ki)(i)^{-1} \rangle &= \langle (i)(ij)(jk)(ki)(i)^{-1} \rangle = \langle (i) \langle (ij)(jk)(ki) \rangle (i)^{-1} \rangle \\ &= \langle (i)(i)^{-1} \langle (ij)(jk)(ki) \rangle \rangle = \Gamma_{ijk}(y), \end{aligned}$$

where in the second equality we used Lemma 6.5(3) and in the third one we used that F is central. This shows that the map $\partial : \check{H}_\varepsilon^1(\mathcal{U}; H) \rightarrow \check{H}^2(\mathcal{U}; F)$ is well defined.

We conclude the proof by proving the stability claim. Let $[\Omega], [\Omega'] \in \check{H}_\varepsilon^1(\mathcal{U}; H)$ be such that $d_{\check{H}}([\Omega], [\Omega']) < \text{sys}(H)/8$ and choose representatives $\Omega, \Omega' \in Z_\varepsilon^1(\mathcal{U}; H)$ such that $d_Z(\Omega, \Omega') < \text{sys}(H)/8$. By Lemma A.20(2), we can pick lifts $\Lambda, \Lambda' \in C^1(\mathcal{U}; G)$ of Ω and Ω' , respectively, such that $d_Z(\Lambda, \Lambda') < \text{sys}(H)/8$. Now, let $y \in U_k \cap U_j \cap U_i$, and let us write $(ij)'$ for $\Lambda'_{ij}(y)$. We have

$$\begin{aligned} d_G((ij)'(jk)'(ki)', \langle (ij)(jk)(ki) \rangle) &\leq d_G((ij)'(jk)'(ki)', (ij)(jk)(ki)) \\ &\quad + d_G((ij)(jk)(ki), \langle (ij)(jk)(ki) \rangle) \\ &< 3 \text{sys}(H)/8 + \varepsilon = \text{sys}(H)/2, \end{aligned}$$

so $\langle \Lambda'_{ij}(y)\Lambda'_{jk}(y)\Lambda'_{ki}(y) \rangle = \langle \Lambda_{ij}(y)\Lambda_{jk}(y)\Lambda_{ki}(y) \rangle$, and the result follows. □

Table 1. Pseudocode for the algorithm sw_1 .

Input: A simplicial complex K with vertices $\{1, \dots, k\}$ and a discrete approximate cocycle $\Omega \in \text{DZ}_\varepsilon^1(K; O(d))$ with $\varepsilon \leq 2$.

Output: A simplicial cocycle $\text{sw}_1(\Omega) \in Z^1(K; \mathbb{Z}/2)$.

For each (ij) -simplex of K with $i < j$, let $\text{sw}_1(\Omega)_{ij} := \det(\Omega_{ji})$.

6.2. First Stiefel–Whitney class

Consider the map $\det : O(d) \rightarrow \{\pm 1\} \cong \mathbb{Z}/2$, and metrize $O(d)$ using the Frobenius distance and $\{\pm 1\}$ by $d(1, -1) = 2$. It follows from Lemma A.13 that \det is 1-Lipschitz.

Algorithm.

Let \mathcal{U} be a cover of a topological space. Let $\varepsilon \leq 2$. By applying \det to a representative cocycle, we get a map $\text{sw}_1 : \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow \check{H}^1(\mathcal{U}; \mathbb{Z}/2)$ such that, if $\Omega, \Omega' \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$ satisfy $d_Z(\Omega, \Omega') < 2$, then $\text{sw}_1(\Omega) = \text{sw}_1(\Omega') \in \check{H}^1(\mathcal{U}; \mathbb{Z}/2)$. This follows directly from Lemma 6.2.

In particular, if K is a simplicial complex, we get an algorithm $\text{sw}_1 : \text{DH}_\varepsilon^1(K; O(d)) \rightarrow H^1(K; \mathbb{Z}/2)$ that is polynomial in the number of vertices of K ; see Table 1 for the pseudocode of this algorithm. From Lemma 6.2, it follows that the algorithm is stable, in the sense that if $\Omega, \Omega' \in \text{DH}_\varepsilon^1(K; O(d))$ satisfy $d_Z(\Omega, \Omega') < 2$, then $\text{sw}_1(\Omega) = \text{sw}_1(\Omega') \in H^1(K; \mathbb{Z}/2)$.

Oriented approximate vector bundles.

The following observation is relevant for the computation of Euler classes in the next section. Let \mathcal{U} be a cover of a topological space with the property that nonempty binary intersections are locally path connected and simply connected. If $\varepsilon \leq 2$ and $\Omega \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$ is such that $\text{sw}_1(\Omega) = 0$, then $\Omega = [\Lambda]$ for some $\Lambda \in Z_\varepsilon^1(\mathcal{U}; SO(d))$. To prove this, let $[\Gamma] = \text{sw}_1(\Omega)$, let $\Theta \in C^0(\mathcal{U}; \mathbb{Z}/2)$ be such that $\Theta \cdot \Gamma = \text{id}$ and lift Θ to $\Pi \in C^0(\mathcal{U}; O(d))$. Then $\Pi \cdot \Omega \in Z_\varepsilon^1(\mathcal{U}; SO(d))$.

Consistency of the algorithm.

The following well-known result can be found in, for example, [38, Example A.3].

Lemma 6.7. *Let \mathcal{U} be a cover of a locally contractible topological space. If $\varepsilon = 0$, the construction $\text{sw}_1 : \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow \check{H}^1(\mathcal{U}; \mathbb{Z}/2)$ computes the first Stiefel–Whitney class of the vector bundle represented by the cocycle.*

Proposition 6.8. *Let \mathcal{U} be a countable cover of a locally contractible, paracompact space B . Let $\varepsilon < 2/9$, and let $\Omega \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$. Then $\text{sw}_1(\Omega) \in \check{H}^1(\mathcal{U}; \mathbb{Z}/2)$ coincides with the first Stiefel–Whitney class of the vector bundle classified by $\pi_*(\text{cl}(\Omega)) : B \rightarrow \text{Gr}(d)$.*

Proof. Since $2/9 \leq \sqrt{2}/4$, there exists, by Theorem 4.27, an exact cocycle $\Lambda \in \check{H}^1(\mathcal{U}; O(d))$ such that $d_{\check{H}^1}(\Omega, \Lambda) < 2/9 \times 9 = 2$ and such that $\pi_*(\text{cl}(\Omega)) = \text{cl}(\Lambda)$. From the stability of sw_1 , it follows that $\text{sw}_1(\Omega) = \text{sw}_1(\Lambda)$. Finally, by Lemma 6.7, $\text{sw}_1(\Lambda)$ is the first Stiefel–Whitney class of the vector bundle classified by $\text{cl}(\Lambda)$, so the result follows. \square

6.3. Euler class of oriented, rank-2 vector bundles

Consider the group $SO(2)$ of orthogonal 2-by-2 matrices with positive determinant. There is a short exact sequence of Lie groups $1 \rightarrow \mathbb{Z} \rightarrow \mathbb{R} \rightarrow SO(2) \rightarrow 1$ given by mapping a real number r to the rotation matrix

$$\begin{pmatrix} \cos(2\pi r) & -\sin(2\pi r) \\ \sin(2\pi r) & \cos(2\pi r) \end{pmatrix}.$$

Table 2. Pseudocode for the algorithm eu.

Input: A simplicial complex K with vertices $\{1, \dots, k\}$ and a discrete approximate cocycle $\Omega \in \text{DZ}_\varepsilon^1(K; SO(2))$ with $\varepsilon \leq 1$.

Output: A simplicial cocycle $\text{eu}(\Omega) \in Z^2(K; \mathbb{Z})$.

1. For each (ij) 1-simplex of K with $i < j$, let $\Lambda_{ij} \in \mathbb{R}$ be a lift of $\Omega_{ij} \in SO(2)$.
2. For each (ijk) 2-simplex of K with $i < j < k$, let $\text{eu}(\Omega)_{ijk} \in \mathbb{Z} \subseteq \mathbb{R}$ the closest element to $\Lambda_{ij} + \Lambda_{jk} - \Lambda_{ik} \in \mathbb{R}$, using the usual distance of \mathbb{R} .

Algorithm.

Let \mathcal{U} be a cover of a topological space with the property that nonempty binary intersections are locally path connected and simply connected. As usual, we endow $SO(2)$ with the Frobenius distance, and we let $SO(2)_g$ denote the same group but endowed with the geodesic distance, with Riemannian structure inherited from the inclusion $SO(2) \subseteq \mathbb{R}^{2 \times 2} = \mathbb{R}^4$.

By Lemma A.18, we have $\text{sys}(SO(2)) = 2\sqrt{2}\pi$, so, applying Theorem 6.6 to the extension $1 \rightarrow \mathbb{Z} \rightarrow \mathbb{R} \rightarrow SO(2)_g \rightarrow 1$, we see that Construction 6.3 gives a map $\text{eu} : \check{H}_\varepsilon^1(\mathcal{U}; SO(2)_g) \rightarrow \check{H}^2(\mathcal{U}; \mathbb{Z})$, as long as $\varepsilon \leq \sqrt{2}\pi/4$. In order to give an algorithm using Frobenius distances, we note that, if $\varepsilon \leq 1$, then, by Corollary A.17, any ε -approximate $SO(2)$ -cocycle is a $(\sqrt{2}\pi/4)$ -approximate $SO(2)_g$ -cocycle. So Construction 6.3 gives a map $\text{eu} : \check{H}_\varepsilon^1(\mathcal{U}; SO(2)) \rightarrow \check{H}^2(\mathcal{U}; \mathbb{Z})$, as long as $\varepsilon \leq 1$. By the stability statement of Theorem 6.6, we see that, if $\Omega, \Omega' \in \check{H}_\varepsilon^1(\mathcal{U}; SO(2))$ are such that $d_{\check{H}}(\Omega, \Omega') < 1$, then $\text{eu}(\Omega) = \text{eu}(\Omega')$.

In particular, if K is a simplicial complex, we get a 1-stable algorithm $\text{eu} : \text{DH}_\varepsilon^1(K; SO(2)) \rightarrow H^2(K; \mathbb{Z})$; see Table 2 for the pseudocode of this algorithm.

Consistency of the algorithm.

We need the following well known result.

Lemma 6.9. *Let \mathcal{U} be a cover of locally contractible space with the property that nonempty binary intersections are locally path connected and simply connected. If $\varepsilon = 0$, the construction $\text{eu} : \check{H}_\varepsilon^1(\mathcal{U}; SO(2)) \rightarrow \check{H}^2(\mathcal{U}; \mathbb{Z})$ computes the Euler class of the vector bundle represented by the cocycle.*

Proof. In this proof, we freely use the language of complex vector bundles. The Lie group $SO(2)$ is isomorphic to the group of unit complex numbers $U(1)$ and thus, up to isomorphism, real, oriented, rank-2 vector bundles are exactly the same as complex rank-1 vector bundles. By [7, (20.10.6)], the top Chern class of a complex vector bundle coincides with the Euler class of the bundle, seen as an oriented real vector bundle in view of the inclusions $U(n) \subseteq SO(2n)$. Finally, when $\varepsilon = 0$, the construction eu computes the first Chern class [38, Example A.5]. □

Proposition 6.10. *Let \mathcal{U} be a countable cover of a locally contractible, paracompact space B , with the property that nonempty binary intersections are locally path connected and simply connected. Let $\varepsilon < 1/9$ and let $\Omega \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$. Then $\text{eu}(\Omega) \in \check{H}^2(\mathcal{U}; \mathbb{Z})$ coincides with the Euler class of the vector bundle classified by $\pi_*(\text{cl}(\Omega)) : B \rightarrow \text{Gr}(d)$.*

Proof. Since $1/9 \leq \sqrt{2}/4$, there exists, by Theorem 4.27, an exact cocycle $\Lambda \in \check{H}^1(\mathcal{U}; O(d))$ such that $d_{\check{H}}(\Omega, \Lambda) < 1$ and such that $\pi_*(\text{cl}(\Omega)) = \text{cl}(\Lambda)$. From the stability of eu, it follows that $\text{eu}(\Omega) = \text{eu}(\Lambda)$. Finally, by Lemma 6.9, eu(Λ) is the Euler class of the vector bundle classified by $\text{cl}(\Lambda)$. □

6.4. Second Stiefel–Whitney class

In order to compute the second Stiefel–Whitney class, we will use a certain short exact sequence $1 \rightarrow \mathbb{Z}/2 \rightarrow \text{Pin}(d) \rightarrow O(d) \rightarrow 1$. To describe this sequence, we introduce the Clifford algebra

Table 3. Pseudocode for the algorithm sw_2 .

Input: A simplicial complex K with vertices $\{1, \dots, k\}$ and a discrete approximate cocycle $\Omega \in DZ_\varepsilon^1(K; O(d))$ with $\varepsilon \leq 1$.

Output: A simplicial cocycle $sw_2(\Omega) \in Z^2(K; \mathbb{Z}/2)$.

1. For each (ij) -simplex of K with $i < j$, let $\Lambda_{ij} \in \text{Pin}(d)$ be a lift of $\Omega_{ij} \in O(d)$.
2. For each (ijk) -simplex of K with $i < j < k$, let $sw_2(\Omega)_{ijk} \in \{\pm \mathbf{1}\} \subseteq \text{Pin}(d)$ the closest element to $\Lambda_{ij}\Lambda_{jk}\Lambda_{ik}^{-1} \in \text{Pin}(d)$, using the geodesic distance of $\text{Pin}(d)$.

associated to the standard inner product of \mathbb{R}^d . We state some well-known results whose proofs can be found in [36], [38, Chapter 1] and [8, Chapter 1, Section 6].

The group Pin .

Fix $d \in \mathbb{N}_{\geq 1}$. The **Clifford algebra** corresponding to the inner product space $(\mathbb{R}^d, \langle -, - \rangle)$, which we denote by $\text{Cl}(d)$, is the quotient of the tensor algebra $T(\mathbb{R}^d)$ by the two-sided ideal generated by elements of the form $v \otimes w + w \otimes v - 2\langle v, w \rangle \mathbf{1}$, for $v, w \in \mathbb{R}^d$, where $\mathbf{1}$ is the unit of $T(\mathbb{R}^d)$. We denote the product of two elements $x, y \in \text{Cl}(d)$ by $x \cdot y \in \text{Cl}(d)$.

Here, we are using the ‘positive convention’ for Clifford algebras, as we want the Pin group to be the group Pin^+ discussed in, for example, [36]. We refer the interested reader to [36, Section 1] for further details about this choice.

Remark 6.11. A more concrete description of $\text{Cl}(d)$ is given by the free noncommutative \mathbb{R} -algebra generated by elements $\{e_1, \dots, e_d\}$ subject to the relations $e_i \cdot e_j = -e_j \cdot e_i$ if $i > j$, and $e_i^2 = \mathbf{1}$. This description makes it evident that $\dim(\text{Cl}(d)) = 2^d$ since the elements of the form $e_{i_1} \cdots e_{i_k}$ for $i_1 < \dots < i_k$ form a basis of $\text{Cl}(d)$ [8, Chapter 1, Corollary 6.7].

If $v \in \mathbb{R}^d$ is of unit length, then it is invertible when seen as an element of $\text{Cl}(d)$ since $v \cdot v = \mathbf{1}$. Let $\text{Pin}(d) \subseteq \text{Cl}(d)^\times$ be the subgroup of the group of units of $\text{Cl}(d)$ generated by elements $v \in \mathbb{R}^d$ of unit length. There is a group morphism $\rho : \text{Pin}(d) \rightarrow O(d)$ that is defined on generators by mapping $v \in \mathbb{R}^d$ with $\|v\| = 1$ to the orthogonal transformation given by reflection about the hyperplane orthogonal to v . This morphism is surjective and its kernel consists of $\{\pm \mathbf{1}\}$ [36, Section 1]. This gives a central sequence of Lie groups $\mathbb{Z}/2 \cong \{\pm \mathbf{1}\} \rightarrow \text{Pin}(d) \rightarrow O(d)$.

Algorithm.

Let \mathcal{U} be a cover of a topological space with the property that nonempty binary intersections are locally path connected and simply connected. Let $\varepsilon \leq 1$. An analogous analysis to the one made for the Euler class shows that, applying Theorem 6.6 to the extension $1 \rightarrow \mathbb{Z}/2 \rightarrow \text{Pin}(d) \rightarrow O(d) \rightarrow 1$, we get a map $sw_2 : \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow \check{H}^2(\mathcal{U}; \mathbb{Z}/2)$. And that sw_2 is such that, if $\Omega, \Omega' \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$ satisfy $d_Z(\Omega, \Omega') < 1$, then $sw_2(\Omega) = sw_2(\Omega') \in \check{H}^2(\mathcal{U}; \mathbb{Z}/2)$.

In particular, if K is a simplicial complex, we get a 1-stable algorithm $sw_2 : \text{DH}_\varepsilon^1(K; O(d)) \rightarrow H^2(K; \mathbb{Z}/2)$; see Table 3 for the pseudocode of this algorithm.

Technicalities about the algorithm.

In order to see that sw_2 is really an algorithm, we must explain how to lift elements from $O(d)$ to $\text{Pin}(d)$, how to multiply elements in $\text{Pin}(d)$ and how to decide, given an element $x \in \text{Pin}(d)$, which one of $\mathbf{1}$ or $-\mathbf{1}$ is closest to it in the geodesic distance.

We see $\text{Pin}(d)$ as a subset of the Clifford algebra, which we represent as in Remark 6.11. By definition of the map $\rho : \text{Pin}(d) \rightarrow O(d)$, to lift a matrix $M \in O(d)$ to an element of $\text{Pin}(d)$, we must write M as a product of reflection matrices, which is exactly what the QR factorization of an orthogonal matrix by means of Householder reflections does [31, Theorem 2.1.14].

In order to multiply elements of $\text{Pin}(d)$, we again use the representation of Remark 6.11. We can thus multiply elements of $\text{Pin}(d)$ in $\mathcal{O}(4^d)$ time.

The problem of deciding which of $\mathbf{1}$ or $-\mathbf{1}$ is closest to an arbitrary element of $\text{Pin}(d)$ in the geodesic distance is more involved, as it depends on the geometry of $\text{Pin}(d)$. We explain how this can be done explicitly for the cases $d = 2, 3, 4$ using extraordinary isomorphisms. The same idea works for $d = 5, 6$ but one must be able to compute geodesic distances in $Sp(4)$ and $SU(4)$.

We first remark that, since $O(d)$ has two connected components, so must $\text{Pin}(d)$, and that $+\mathbf{1}$ and $-\mathbf{1}$ belong to the same connected component of $\text{Pin}(d)$, as they are both preimages of the identity matrix under the covering map $\text{Pin}(d) \rightarrow O(d)$. Let the **spin group** $\text{Spin}(d) \subseteq \text{Pin}(d)$ be the connected component of $+\mathbf{1}$. Checking if $x \in \text{Pin}(d)$ belongs to $\text{Spin}(d)$ is easy since it amounts to checking if its image in $O(d)$ has positive determinant.

If $x \in \text{Pin}(d)$ does not belong to $\text{Spin}(d)$, then its distance to $+\mathbf{1}$ is the same as the distance to $-\mathbf{1}$, namely infinity. With this in mind, we may assume that $x \in \text{Spin}(d)$ and that we want to know if x is closest to $+\mathbf{1}$ or to $-\mathbf{1}$ in the geodesic distance of $\text{Spin}(d)$.

For explicit formulas for the extraordinary isomorphisms mentioned below, see the example after Lemma 1.8.3 in [35] or [42]. One should keep in mind that [35] is working with the ‘negative convention’ for the Clifford algebra, and thus the isomorphisms are given for $\text{Spin}(d) \subseteq \text{Cl}^-(d)$, but it is not hard to modify their formulas to get the isomorphisms when $\text{Spin}(d) \subseteq \text{Cl}(d)$.

- $(d = 2)$ In this case, $\text{Spin}(d)$ is isomorphic to the circle as Lie groups, so the geodesic distance on $\text{Spin}(2)$ can be computed, up to a multiplicative constant, using arclengths on the circle. Note that the multiplicative constant is inconsequential for the purposes of determining if an element is closer to $+\mathbf{1}$ or to $-\mathbf{1}$.
- $(d = 3)$ In this case, $\text{Spin}(d)$ is isomorphic to $SU(2)$, which in turn is diffeomorphic to the 3-sphere S^3 . The geodesic distance is again computed, up to a multiplicative constant, using arcs.
- $(d = 4)$ In this case, $\text{Spin}(d)$ is isomorphic to $SU(2) \times SU(2)$, which in turn is diffeomorphic to $S^3 \times S^3$. Moreover, the Riemannian metric on $\text{Spin}(4)$ induced by the double cover $\text{Spin}(4) \rightarrow SO(4)$ coincides, up to a multiplicative constant, with the product metric on $S^3 \times S^3$, where the two copies of S^3 are endowed with the Riemannian metric inherited from the inclusion $S^3 \subseteq \mathbb{R}^4$. To compute the geodesic distance on $S^3 \times S^3$, recall that a geodesic in a product Riemannian manifold corresponds to a pair of geodesics, one on each factor [39, Problem 5-7]. This implies that, in a product of complete Riemannian manifolds $X \times Y$, the geodesic distance satisfies the Pythagorean theorem, meaning that the geodesic distance from (x, y) to (x', y') is equal to the square root of the sum of the squares of the geodesic distances from x to x' and from y to y' .

So, when $d = 2, 3, 4$, this results in a numerical algorithm sw_2 that is polynomial in the number of vertices of K .

Consistency of the algorithm.

We need a well-known result; see, for example, [36, Lemma 1.3].

Lemma 6.12. *Let \mathcal{U} be a cover of a locally contractible space with the property that nonempty binary intersections are locally path connected and simply connected. If $\varepsilon = 0$, the construction $\text{sw}_2 : \check{H}_\varepsilon^1(\mathcal{U}; O(d)) \rightarrow \check{H}^2(\mathcal{U}; \mathbb{Z}/2)$ computes the second Stiefel–Whitney class of the vector bundle represented by the cocycle.*

Note that, also by [36, Lemma 1.3], if we would have used the negative convention for the Clifford algebra, the algorithm would be computing $\text{sw}_1^2 + \text{sw}_2$ instead of sw_2 .

The proof of the following proposition is analogous to the proof of Proposition 6.10 but uses Lemma 6.12 instead of Lemma 6.9.

Proposition 6.13. *Let \mathcal{U} be a countable cover of a locally contractible, paracompact space B , with the property that nonempty binary intersections are locally path connected and simply connected. Let $\varepsilon < 1/9$ and let $\Omega \in \check{H}_\varepsilon^1(\mathcal{U}; O(d))$. Then $\text{sw}_2(\Omega) \in \check{H}^2(\mathcal{U}; \mathbb{Z}/2)$ is the second Stiefel–Whitney class of $\pi_*(\text{cl}(\Omega))$.*

7. Computational examples

In this section, we run the algorithms of Section 6 on data. We describe the basic pipeline used in the examples in Lemma 7.1.

In Section 7.2, we use sw_1 to study the topology and, in particular, the orientability of attractors in the time-variant double-gyre dynamical system. This illustrates how characteristic classes can be used to classify such attractors. The trajectory of particles in the dynamical system was computed numerically using the code in [16].

In Section 7.3, we use sw_1 and sw_2 to provide experimental evidence that confirms that a certain dataset of lines in \mathbb{R}^2 can be parametrized by a projective plane. The theoretical explanation of this phenomenon appears in [52, Section 2.4]. The code used to generate the dataset is from [64].

In Section 7.4, we show how eu can be used to detect nonsynchronizability of data. The example is a version of the cryo-EM problem mentioned in Section 1.1.1.

In Table 4, we summarize the runtime of our algorithms. Code to replicate these examples can be found at [56].

7.1. Pipeline

We make free use of basic tools from persistence theory such as Vietoris–Rips complexes, persistent cohomology and persistence diagrams; see, for example, [20, 50]. The persistent cohomology computations are done using the Python interface [65] for Ripser [4].

Let $\{K_r\}_{r \in \mathbb{R}_{\geq 0}}$ be a filtration of a finite simplicial complex K by subcomplexes so that we have $K_s \subseteq K_r \subseteq K$ for $s \leq r \in \mathbb{R}$. For $\Omega \in \check{H}^1_\infty(K; \mathcal{O}(d))$ and $\varepsilon \geq 0$, define the ε -**death** of Ω as

$$\delta = \sup(\{r \geq 0 : \Omega \text{ is an } \varepsilon\text{-approximate cohomology class on } K_r\} \cup \{0\}).$$

The ε -**span** of Ω is the subset of the persistence diagram of K of classes whose (homological) birth is at most δ and whose (homological) death is at least δ .

Fix a basis for the persistent cohomology of K , that is, cocycles $B = \{\Lambda_1, \dots, \Lambda_n\}$ representing each point of the persistence diagram of the $\mathbb{Z}/2$ -cohomology of K . Let $\varepsilon = 2$. For any $r < \delta$, we can write $sw_1(\Omega)$ in the basis B . To represent $sw_1(\Omega)$ in the persistence diagram of K , we decorate the diagram by circling the classes of B that appear with a nonzero coefficient. Note that any such class must live inside the 2-span of Ω . An analogous analysis, replacing 2 with 1, applies for sw_2 . For eu , we use $\mathbb{Z}/3$ -cohomology and the mod 3 reduction of the Euler class.

7.2. Orientability of attractors

The **double-gyre** is the dynamical system characterized by the differential equations $\dot{x} = \partial\psi/\partial y$ and $\dot{y} = -\partial\psi/\partial x$, where

$$\psi(x, y, t) = A \sin(\pi f(x, t)) \sin(\pi y), \quad f(x, t) = a(t) x^2 + (1 - 2a(t)) x, \quad a(t) = \varepsilon \sin(\omega t),$$

and A, ε and ω are positive parameters. The system is defined over the domain $(x, y) \in [0, 2] \times [0, 1]$ with $t \in \mathbb{R}$ representing time. It was introduced in [57] as a simplified model of the double-gyre pattern

Table 4. Runtime of the algorithms on laptop with a 2.20 GHz Intel Core i7 and 16GB of RAM.

Dataset X	VR(X) size	Algorithms	Time
ATTR1	22921-simplices at 2-death of Ω	w and sw_1	180ms
ATTR2	78341-simplices at 2-death of Ω	w and sw_1	450ms
LINES	31741-simplices at 2-death of Ω	w and sw_1	270ms
LINES	266832-simplices at 1-death of Ω	w and sw_2	3950ms
PROJS	455732-simplices at 1-death of Ω	eu	180ms

observed in geophysical flows. Geometrically, the double-gyre system consists of a pair of vortices oscillating back and forth, horizontally, in time; see, for example, [57, Figure 5].

Since the vector field characterizing the system varies periodically with time, the flow line of a particle initially at (x_0, y_0) depends on the time t_0 at which the particle is at that spot. Because of this, the *phase space* of the system (the space parametrizing the possible states of a particle) is $[0, 2] \times [0, 1] \times S^1$, where the last coordinate represents time \mathbb{R} modulo the period π/ω .

Dynamical systems can be analyzed by studying the topology of their attractors [62, 1, 51]. Informally, an *attractor* \mathcal{M} of a dynamical system consists of a subset of the phase space that is invariant under the action of the system and such that any point sufficiently close to the attractor gets arbitrarily close to it as the system evolves. We refer the interested reader to [60, 67, 62] for formal notions of attractor.


Usually, one only has access to partial information about the trajectory of a particle, which one wants to use to study the topology of the attractor \mathcal{M} the particle is converging to. For example, one may be given a real-valued time series $\{x_n\}_{n=1, \dots, N}$, which comes from applying a differentiable map F defined on the phase space to a finite sample of the trajectory of the particle. If the attractor \mathcal{M} is a smooth manifold and certain other technical conditions are satisfied, a theorem of Takens [62] implies that the *delay embedding* of the time series

$$X = \{(x_i, x_{i+\tau}, x_{i+2\tau}, \dots, x_{i+(d-1)\tau})\}_{i=1, \dots, N-(d-1)\tau} \subseteq \mathbb{R}^d$$


is concentrated around a diffeomorphic copy of \mathcal{M} , and is sufficiently dense in \mathcal{M} so that the Vietoris–Rips complex of X can be used to estimate the homology of \mathcal{M} , and local PCA can be used to estimate the tangent bundle of \mathcal{M} . Here, d is the *target dimension* and τ is the *delay*. We refer the reader to [62, 1, 51, 68] for more information about recovering the geometry of attractors from a delay embedding.

In [13, Section 4.1], four attractors of the double-gyre dynamical system are studied. The four attractors can be distinguished using their homology, except for two of them which, topologically, look like a cylinder and a Möbius strip, respectively, and thus have the same homology groups. In order to deal with such examples, the authors of [13] develop an algorithm for detecting orientability of this kind of data. Here, we show that our general purpose algorithms readily apply to this kind of data and confirm, using sw_1 , that [13, Section 4.1, Example 4] is parametrized by a Möbius band and [13, Section 4.1, Example 1] by a cylinder.

Fixing the initial condition $(x_0, y_0, t_0) = (0.55, 0.5, 0) \in [0, 2] \times [0, 1] \times S^1$, for a double-gyre with $A = 0.1$, $\varepsilon = 0.1$, and $\omega = \pi/5$, as in [13], we sample the trajectory of the particle at 2000 equally spaced times from $t = 0$ to $t = 1000$, obtaining a time series of positions $\{(x_n, y_n)\}_{n=1, \dots, 2000} \subseteq [0, 2] \times [0, 1]$. We take the function F to be projection onto the x -coordinate, as in [13], the delay $\tau = 5$ and the target dimension $d = 5$. A 2D projection of X is depicted in Figure 2a. The dataset ATTR1 consists of a subsample of 1000 points of X .



(a) 2D projection of the ATTR1 dataset.



(b) 2D projection of the ATTR2 dataset.

Figure 2. 2D projections of samples approximating two attractors of the double-gyre dynamical system. Consecutive samples are joined by an edge.

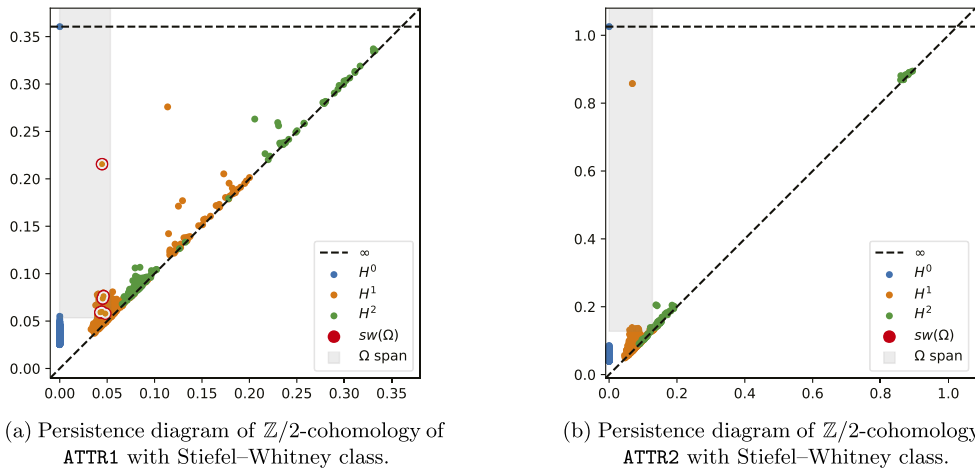


Figure 3. Persistence diagrams of the Vietoris–Rips cohomology of finite samples approximating two attractors of the double-gyre dynamical system, with Stiefel–Whitney classes highlighted.

We approximate the tangent bundle of \mathcal{M} using local PCA as sketched in Section 1.1.2. This gives us a discrete approximate local trivialization Φ , which we use to compute an approximate cocycle $\Omega = w(\Phi)$ over the Vietoris–Rips complex of X . We choose $k = 58$ for the local PCA computations as this value maximizes the 2-death of Ω .

Finally, we apply the algorithm sw_1 to Ω and write the cohomology class thus obtained in the basis of the persistent cohomology provided by Ripser. The classes circled in red in Figure 3a are the classes that sum to $sw_1(\Omega)$. The fact that the most persistent class of the persistent diagram is colored in red tells us that, as one goes around the one-dimensional hole this class represents, the approximate vector bundle given by local PCA is changing orientation. This, coupled with the fact that the dataset is locally two-dimensional, suggests that the dataset is parametrized by a Möbius band.


We run the same pipeline but with initial condition $(x_0, y_0, t_0) = (0.25, 0.125, 0)$, as in [13, Section 4.1, Example 1]; we refer to this second dataset as ATTR2. In this case, regardless of the parameters chosen for local PCA, the first Stiefel–Whitney class of the approximate cocycle it gives is zero, as we see in Figure 3b.

7.3. Projective space of lines

Datasets of simple geometrical shapes, such as lines in \mathbb{R}^2 with different orientations, appear, for instance, as datasets of weights learned by autoencoders and other neural networks [29, Figure 2.A], [30, Figures 2 and 3] and as datasets of local patches of natural images [12]. It has been shown that topology can be used to understand these datasets by giving insight into the functions learned by the hidden layers of neural networks [10, 48], and by finding convenient parametrizations of spaces of image patches [52].

Here, we show how characteristic classes can be used to understand the topology of a set of lines in \mathbb{R}^2 . The dataset LINES consists of 160 grayscale images represented by a 10×10 real matrix. Each image represents a fuzzy line in \mathbb{R}^2 with a certain slope and offset. A sample of 56 elements of LINES is displayed in Figure 4a. We interpret this dataset as a point cloud $X \subseteq \mathbb{R}^{10 \times 10}$.

We proceed with the same pipeline as in Section 7.2. We compute the persistence diagram of $VR(X)$ with coefficients in $\mathbb{Z}/2$ and observe that there are two significant classes, one in H^1 and the other in H^2 (Figure 4b). By applying local PCA with a range of values k , we see that the dataset seems to have an intrinsic dimension of 2. This may lead one to suspect that the dataset is parametrized by the real projective plane. One way to confirm this suspicion is by computing persistent cohomology with coefficients in $\mathbb{Z}/3$ and seeing that the most persistent classes in H^1 and H^2 disappear.



(a) Grid displaying 56 images of the LINES dataset. (b) Persistence diagram of $\mathbb{Z}/2$ -cohomology of LINES with Stiefel–Whitney classes.

Figure 4. Samples of a dataset parametrized by the real projective plane and the persistence diagram of its Vietoris–Rips cohomology, with Stiefel–Whitney classes highlighted.

Another method to corroborate the hypothesis is the following. Local PCA gives a discrete approximate local trivialization Φ over $VR(X)$, and the 1-death of $\Omega = w(\Phi)$ is maximized using $k = 18$. We apply the algorithms sw_1 and sw_2 to Ω and see that the cohomology classes thus obtained coincide with the most persistent H^1 and H^2 classes of Figure 4b, respectively. This confirms the hypothesis that the dataset is parametrized by a projective plane, as we know that $H^1(\mathbb{R}P^2; \mathbb{Z}/2) \cong H^2(\mathbb{R}P^2; \mathbb{Z}/2) \cong \mathbb{Z}/2$ and that the Stiefel–Whitney classes of the tangent bundle of $\mathbb{R}P^2$ are nonzero.


7.4. Lack of global synchronization in cryo-EM

We give an example of how characteristic classes provide an obstruction to synchronization. In order to do this, we simulate an instance of the main problem of cryo-EM [18, 28]. In this problem, one is given a set of 2D projections of an unknown 3D shape and is asked to reconstruct the shape. One possible formalism is to assume that the shape is given by a density $S : \mathbb{R}^3 \rightarrow \mathbb{R}$. We think of the projection process as first acting on the molecule by an unknown element $v \in SO(3)$, which we call **projection angle**, giving $S \circ v : \mathbb{R}^3 \rightarrow \mathbb{R}$, and then integrating $S \circ v$ along the z -axis, yielding a map $S_v : \mathbb{R}^2 \rightarrow \mathbb{R}$.


One of the main difficulties is that one is only given a set of projected images and is not given the projection angle v corresponding to each image. Much attention has thus been given to the problem of recovering the projection matrices corresponding to each image, up to a global rigid automorphism of S^2 , that is, an element of $SO(3)$.

Let X be a set of 2D projections of an unknown 3D shape. A successful approach [23, 24, 69] to recovering the projection angles starts by computing rotations $g_{ij} \in SO(2)$ that best align $x_i, x_j \in X$. Formally, one chooses a distance d between maps $\mathbb{R}^2 \rightarrow \mathbb{R}$ and finds $g_{ij} \in SO(2)$ that minimizes the distance between $x_i : \mathbb{R}^2 \rightarrow \mathbb{R}$ and $g_{ij} \circ x_j : \mathbb{R}^2 \rightarrow \mathbb{R}$. One then fixes a threshold δ and constructs a two-dimensional simplicial complex K with 0-simplices given by the elements of X , 1-simplices given by pairs (ij) such that $d(x_i, g_{ij}x_j) < \delta$ and having a 2-simplex (ijk) anytime (ij) , (jk) and (ik) are 1-simplices of K . This construction can be interpreted as giving an approximate $SO(2)$ -cocycle over K . This cocycle approximates the tangent bundle of $S^2 \subseteq \mathbb{R}^3$ and can be used to recover the projection angles.

We are interested in simulating an instance of this problem and in computing the Euler class of the approximate cocycle. As we expect this class to be nonzero, this shows that characteristic classes of approximate vector bundles can be used for model validation and to detect nonsynchronizability of data. This also confirms the theoretical analysis of [58], which shows that global synchronization is not possible.



(a) Grid displaying 9 images of the PROJS dataset.



(b) Persistence diagram of $\mathbb{Z}/3$ -cohomology of PROJS, with mod 3 reduction of the Euler class.

Figure 5. A sample of a synchronization problem parametrized by a bundle over the sphere and the persistence diagram of the base of the bundle with the Stiefel–Whitney classes highlighted.

In order to construct an instance of the problem, we define $S : \mathbb{R}^3 \rightarrow \mathbb{R}$ to be the characteristic function of a union of four balls of different radii, centered at different points of \mathbb{R}^3 . We then compute 400 projections $\{S_{v_i} : \mathbb{R}^2 \rightarrow \mathbb{R}\}_{i=1, \dots, 400}$ for $\{v_i \in SO(3)\}$ a well-distributed random sample of $SO(3)$. We save these projections as 100×100 grayscale images, which constitute the dataset PROJS. A sample of 9 elements of PROJS is displayed in Figure 5a.

Since aligning images optimally is a difficult problem in itself, we simplify our computations by using our knowledge of the projection angles v_i and v_j to align the images x_i and x_j . More specifically, given $v \in SO(3)$, let $v^3 \in S^2$ denote its third column. We compose v_i with the rotation $r_{ij} \in SO(3)$ along a minimizing geodesic between v_i^3 and v_j^3 and define $\Omega_{ij} \in SO(2)$ such that

$$r_{ij}v_i = \begin{pmatrix} \Omega_{ij} & 0 \\ 0 & 1 \end{pmatrix}v_j.$$

We let X be a sample of 300 elements of PROJS and compute the Vietoris–Rips complex of X , where the dissimilarity function $d(i, j)$ is given by the distance between the image x_i and the rotated image $\Omega_{ij}x_j$.

Finally, we compute the persistence diagram of $VR(X)$ with coefficients in $\mathbb{Z}/3$, compute the Euler class of Ω at its 1-death, reduce it mod 3 and write it in the basis provided by the persistent cohomology computations. We see that the persistent class representing the reduction mod 3 of the Euler class, which appears circled in red in Figure 5b, is the most persistent class of $VR(X)$.

8. Future work

Robustness and applications.

The notions of vector bundle we have presented can tolerate a certain amount of noise but are not robust to outliers. What this means is that an ε -approximate cocycle on a simplicial complex can satisfy the δ -approximate cocycle condition for a very small δ on almost all 2-simplices, except on very few, and this can still make ε very large. This problem is similar to the robustness problem of persistent homology, where the addition of a few outliers to the dataset can completely alter the inferred persistent homology.

We are interested in studying extensions of our framework that lead to algorithms for the inference of topological information of approximate vector bundles even in the presence of outliers. The work in [59] is an example of such an algorithm, as it can be interpreted as a robust algorithm for inferring whether the first Stiefel–Whitney class of the tangent space of an embedded manifold is 0 or not.

Invariants of approximate cocycles.

A related problem is that of defining invariants of ε -approximate cocycles that do not require ε to be sufficiently small. A possible avenue is to consider the persistent cohomology of the thickened Grassmannians and use ε -approximate classifying maps to pull back the cohomology classes of these Grassmannians. The question of how long the universal Stiefel–Whitney classes persist in the thickened Grassmannians is related to the filling radius introduced by Gromov [22] and to the generalization considered by Lim, Mémoli and Okutan in [40]. A related question is whether better bounds for our results can be obtained by using the Vietoris–Rips complex of the Grassmannians, instead of the thickening.

Approximate sections.

In applications, one is interested in finding sections of approximate vector bundles. In order to do this in practice and to prove consistency theorems for this kind of computation, a notion of approximate section needs to be introduced and studied.

Approximate vector bundles with connection.

Discrete approximate cocycles over simplicial complexes, as in Definition 5.1, appear in the physics literature (e.g., [14]), where they have a different interpretation: The value $\Omega_{ij} \in O(d)$ of the (discrete) cocycle Ω on the 1-simplex (ij) represents parallel transport from the fiber of the vertex i to the fiber of the vertex j . It would be interesting to relate this interpretation to the results in this paper and to study to what extent a discrete approximate cocycle can be used to reconstruct a vector bundle together with a connection.

Algorithms for higher characteristic classes.

Finding algorithms for the effective computation of higher characteristic classes, at least in the exact case, has been the subject of much research in the past. For instance, [9] and [44] give cocycles representing Chern, Pontryagin, Stiefel–Whitney and Euler classes of vector bundles over compact manifolds. One should note that many of the formulas are not entirely algorithmic, as they require to determine, for example, whether certain singular cycles are 0 in homology or not.

Other structure groups.

Some synchronization problems are most naturally described using a structure group different from the orthogonal group $O(d)$. We hope our theory can be extended to bundles having other structure groups, such as $U(d)$, $PO(d)$, and $(\mathbb{R}^n, +)$.

A. Technical details***A.1. Reach and thickenings***

Let $\emptyset \neq X \subseteq \mathbb{R}^N$ be closed. For $\varepsilon \in [0, \infty]$, define the ε -**thickening** of X by $X^\varepsilon = X$ if $\varepsilon = 0$, and otherwise by

$$X^\varepsilon = \{y \in \mathbb{R}^N : \exists x \in X, \|x - y\|_2 < \varepsilon\}.$$

A **closest point** of $y \in \mathbb{R}^N$ in X is a point $x \in X$ that minimizes $\|x - y\|_2$. The **medial axis** of X , denoted $\text{med}(X)$, consists of all points $y \in \mathbb{R}^N$ that admit more than one closest point in X . The **reach** of X is the supremum of $\varepsilon \geq 0$ such that X^ε does not intersect $\text{med}(X)$. Define a function $\pi : \mathbb{R}^N \setminus \text{med}(X) \rightarrow X$ by mapping a point y to its closest point of X .

Lemma A.1. *If $\emptyset \neq X \subseteq \mathbb{R}^N$ is closed, then the map $\pi : \mathbb{R}^N \setminus \text{med}(X) \rightarrow X$ is continuous. In particular, if X has strictly positive reach and $\varepsilon \leq \text{reach}(X)$, then $\pi : X^\varepsilon \rightarrow X$ is well defined and continuous. It follows that the map $\pi : X^\varepsilon \rightarrow X$ is a homotopy equivalence, with inverse the inclusion $X \rightarrow X^\varepsilon$.*

Proof. We start by proving the first claim. For this, we show that π maps convergent sequences to convergent sequences. Let $y_n \rightarrow y$ be a convergent sequence in $\mathbb{R}^N \setminus \text{med}(X)$. By the triangle inequality, we have

$$\|\pi(y_n) - \pi(y)\| \leq \|\pi(y_n) - y_n\| + \|y_n - y\| + \|y - \pi(y)\| = d(y_n, X) + \|y_n - y\| + d(y, X) \rightarrow 2d(y, X).$$

Thus, for n sufficiently large, we have that y_n is inside the closed ball of radius $2d(y, X) + 1$ around $\pi(y)$, which is a compact set. Combining this with the fact that X is closed, we can assume, without loss of generality, that $\pi(y_n)$ converges to some point $x \in X$. We must show that $x = \pi(y)$, and to prove this it is enough to show that $\|y - x\|_2 = d(y, X)$, by uniqueness of minimizers. By definition, $\|y_n - \pi(y_n)\|_2 = d(y_n, X)$ and $d(\cdot, X)$ is a continuous function, so, taking a limit, we see that $\|y - x\|_2 = d(y, X)$.

For the second claim, it suffices to show that $X^{\text{reach}(X)}$ does not intersect $\text{med}(X)$, which is clear since any point of $X^{\text{reach}(X)}$ belongs to X^δ for some $\delta < \text{reach}(X)$.

Finally, to see that π is a homotopy equivalence, note that the inclusion followed by the projection is certainly the identity of X and the projection followed by the inclusion is homotopic to the identity of X^ε by means of a linear homotopy. □

Lemma A.2. *Let $X \subseteq \mathbb{R}^N$ be closed with strictly positive reach r . Let z be a convex combination of $x_1, \dots, x_n \in X$ such that $\|x_i - x_j\| < r$ for all $1 \leq i, j \leq n$. Then $\|z - x_j\| < r$ and $\|\pi(z) - x_j\| < 2r$ for all $1 \leq j \leq n$.*

Proof. Fix $1 \leq j \leq n$. The points x_1, \dots, x_n are all contained in the open ball of radius r around x_j . It follows that z is also contained in this ball by the convexity of open balls. In particular, $z \in X^r$ and $\|\pi(z) - z\| < r$. By the triangle inequality, we have $\|\pi(z) - x_j\| < 2r$, as required. □

A.2. Polar decomposition and orthogonal Procrustes problem

In this section, we collect a few standard facts about the orthogonal Procrustes problem; a standard reference is [31].

Let $n \geq m \geq 1 \in \mathbb{N}$. Let $N \in \mathbb{R}^{n \times m}$. A **polar decomposition** of N consists of matrices $U \in \mathbb{R}^{n \times m}$ and $P \in \mathbb{R}^{m \times m}$ such that $N = UP$, the matrix U has orthonormal columns and the matrix P is positive semidefinite. We will need the following well known fact about polar decompositions; for a reference, see [31, Theorem 7.3.1]. In the statement, for A a symmetric positive semidefinite matrix, $A^{1/2}$ denotes the unique positive semidefinite square root [31, Theorem 7.2.6], also called the **principal square root** of A .

Lemma A.3. *Any matrix admits a polar decomposition. The factor P is uniquely determined and satisfies $P = (N^t N)^{1/2}$. The factor U is uniquely determined if N has full rank.*

Let $V(m, n) \subseteq \mathbb{R}^{n \times m}$ denote the subset of matrices with orthonormal columns. Consider the map $Q : \mathbb{R}^{n \times m} \rightarrow V(m, n)$ that assigns to a matrix N a matrix U that is part of a polar decomposition $N = UP$. Note that, for general matrices, this map depends on a choice of polar decomposition.

Corollary A.4. *The map $Q : \mathbb{R}^{n \times m} \rightarrow V(m, n)$ is uniquely determined and continuous if we restrict its domain to matrices with full rank.*

Proof. By Lemma A.3, if N has full rank, we have $U = NP^{-1} = N((N^t N)^{1/2})^{-1}$. The map $A \mapsto A^{1/2}$ that takes a symmetric positive definite matrix to its principal square root is continuous [32, Chapter 6.2, Problem 26] and so is the map that takes an invertible matrix to its inverse. The result follows. □

For $M, N \in \mathbb{R}^{n \times d}$, the **orthogonal Procrustes problem** is the following optimization problem:

$$\min_{\Omega \in O(d)} \|M\Omega - N\|.$$

It is well known (see, e.g., [31, Section 7.4.5]) that the solutions to the above orthogonal Procrustes problem are precisely the orthogonal matrices U that are part of a polar decomposition $M^t N = UP$. In particular, if $M^t N$ has full rank, then the above problem admits exactly one solution. Moreover, this solution varies continuously in M^t and N , as long as we restrict ourselves to problems for which $M^t N$ has full rank.

A.3. Metrics on Grassmannians

In this section, we compare two metrics on the Grassmannians and we recall the computation of the reach of the Grassmannians due to Tinarrage [63].

Recall that $\text{Proj} : \mathcal{V}(d, n) \rightarrow \text{Gr}(d, n)$ is $O(d)$ -invariant, with $O(d)$ acting by matrix multiplication on the right. This means that we have $\text{Proj}(M) = \text{Proj}(M\Omega)$ for every $\Omega \in O(d)$. In fact, it is easy to see that, for $N, M \in \mathcal{V}(d, n)$, one has $\text{Proj}(N) = \text{Proj}(M)$ if and only if $N\Omega = M$ for some $\Omega \in O(d)$. This shows that, as a set, $\text{Gr}(d, n)$ is the quotient of $\mathcal{V}(d, n)$ by the action of $O(d)$ by matrix multiplication on the right. Since $O(d)$ acts by isometries, this induces a metric d_q on $\text{Gr}(d, n)$, given by

$$d_q(A, B) = \min_{\substack{N, M \in \mathcal{V}(d, n) \\ \text{Proj}(N)=A, \text{Proj}(M)=B}} \|N - M\|.$$

We now prove some useful comparisons between the metric and the Frobenius distance.

Lemma A.5. *Let $M \in \mathcal{V}(d, n)$, $A \in \mathbb{R}^{d \times m}$, and $B \in \mathbb{R}^{n \times m}$. Then $\|MA\| = \|A\|$ and $\|M^t B\| \leq \|B\|$.*

Proof. For the first claim, we compute

$$\|MA\|^2 = \text{tr}((MA)^t(MA)) = \text{tr}(A^t M^t MA) = \text{tr}(A^t A) = \|A\|^2.$$

For the second, note that $\|M^t B\|^2 = \text{tr}(MM^t BB^t)$. Since MM^t is an orthogonal projection to a subspace, there is an orthogonal change of basis such that MM^t is diagonal, with diagonal entries 1 repeated d times and 0 repeated $n - d$ times. The result follows by computing $\text{tr}(MM^t BB^t)$ in that basis. \square

Let d_{Fr} denote the metric on Grassmannians induced by the Frobenius norm.

Lemma A.6. *The map $\text{Proj} : \mathcal{V}(d, n) \rightarrow \text{Gr}(d, n)$ given by $\text{Proj}(M) = MM^t$ is $\sqrt{2}$ -Lipschitz. In particular, we have $d_{Fr} \leq \sqrt{2}d_q$.*

Simple examples with $d = 1$ and $n = 2$ show that this bound is tight.

Proof. The second claim is a consequence of the first one. For the first claim, we start by recalling that, for any matrix A , one has $\|A^t A\| = \|AA^t\|$. In particular, if A is square, $\|AA^t - \text{id}\|^2 = \|AA^t\|^2 + \|\text{id}\|^2 - 2 \text{tr}(AA^t) = \|A^t A\|^2 + \|\text{id}\|^2 - 2 \text{tr}(A^t A) = \|A^t A - \text{id}\|^2$.

Now, let $M, N \in \mathcal{V}(d, n)$. Since multiplication by an orthogonal matrix preserves the Frobenius norm, we may assume that M and N are matrices given by blocks:

$$M = \begin{bmatrix} A \\ B \end{bmatrix}, \quad N = \begin{bmatrix} \text{id} \\ 0 \end{bmatrix},$$

where the top block is $d \times d$ and the bottom one is $(n - d) \times d$. Note that, by assumption, we have $A^t A + B^t B = \text{id}$. We now bound as follows:

$$\begin{aligned} \|MM^t - \text{id}\|^2 &= \|AA^t - \text{id}\|^2 + \|AB^t\|^2 + \|BA^t\|^2 + \|BB^t\|^2 \\ &= \|A^tA - \text{id}\|^2 + 2\|BA^t\|^2 + \|BB^t\|^2 \\ &= 2\left(\|BA^t\|^2 + \|BB^t\|^2\right) = 2\|BM^t\|^2 = 2\|B\|^2 \\ &\leq 2\left(\|A - \text{id}\|^2 + \|B^t\|^2\right) = 2\|M - \text{id}\|^2, \end{aligned}$$

where we used the previous observations in the second and third equality and Lemma A.5 in the fifth equality. □

To prove a partial converse, we need the following lemma.

Lemma A.7. *If P is a symmetric and positive semidefinite $d \times d$ matrix, then $\|P^{1/2} - \text{id}\| \leq \|P - \text{id}\|$.*

Proof. Diagonalizing P in an orthonormal basis, we can assume that P is diagonal with nonnegative entries. Since then the square root of P corresponds to taking the square root of each diagonal entry, it is enough to prove that, for any nonnegative real number a one has $|\sqrt{a} - 1| \leq |a - 1|$, which is clear. □

The next lemma is not needed in what follows, but, since it is a direct consequence of the previous result, we give it here.

Lemma A.8. *Let $M \in \mathbb{R}^{n \times d}$ and $M = UP$ be a polar decomposition. Then $\|M - U\| \leq \|M^tM - \text{id}\|$.*

Proof. By definition, $P^2 = M^tM$. So the result follows from Lemma A.7. □

We are ready to prove the converse.

Lemma A.9. *Let $M, N \in V(d, n)$. Then there exists $\Omega \in O(d)$ such that $\|M\Omega - N\| \leq \|MM^t - NN^t\|$. In particular, $d_q \leq d_{Fr}$.*

Proof. The second claim is a consequence of the first one. For the first claim, let $\Omega \in O(d)$ minimize $\|M\Omega - N\|$. From Appendix A.2, we know that Ω is part of a polar decomposition $M^tN = \Omega(N^tMM^tN)^{1/2}$. Now, $\|M\Omega - N\| \leq \|\Omega - M^tN\| = \|\Omega - \Omega(N^tMM^tN)^{1/2}\| = \|(N^tMM^tN)^{1/2} - \text{id}\| \leq \|(N^tMM^tN)^{1/2} - \text{id}\|$, by Lemma A.7. □

We now recall the computation of the reach of the Grassmannians due to Tinarrage [63].

Construction A.10. Let $A \in \mathbb{R}^{n \times n}$. Define $A^s = (A + A^t)/2$ and let $A^s = \Omega D \Omega^t$ be a diagonalization of A^s by an ordered orthonormal basis $\Omega \in O(n)$, where the diagonal entries of D contain the eigenvalues of A^s in decreasing order. Let J_d be the diagonal $n \times n$ matrix with the first d diagonal entries equal to 1 and the rest equal to 0. Let $\pi(A) = \Omega J_d \Omega^t$.

Lemma A.11 [63, Lemma 2.1]. *If $d_{Fr}(A, \text{Gr}(d, n)) < \sqrt{2}/2$, then $\pi(A)$ is the unique minimizer of*

$$\min_{B \in \text{Gr}(d, n)} \|A - B\|.$$

Moreover, we have $\text{reach}(\text{Gr}(d, n)) = \sqrt{2}/2$.

In particular, if $A \in \text{Gr}(d, n)^\varepsilon$ for $\varepsilon \leq \sqrt{2}/2$, then $\pi(A)$ is independent of the choice of orthonormal basis Ω . We deduce the following.

Proposition A.12. *Let $\varepsilon \leq \sqrt{2}/2$. The inclusion $\text{Gr}(d, n) \subseteq \text{Gr}(d, n)^\varepsilon$ is a homotopy equivalence, with inverse given by the projection π . This is natural in both n and ε . Moreover, the projections assemble into a homotopy inverse of the inclusion $\text{Gr}(d) \subseteq \text{Gr}(d)^\varepsilon$.*

Proof. The first claim follows from Lemma A.1 and Lemma A.11. In the second claim, naturality means that the inclusion maps $i : \text{Gr}(d, n) \rightarrow \text{Gr}(d, n)^\varepsilon$ commute with the inclusions into $\text{Gr}(d, n + 1)$ and $\text{Gr}(d, n + 1)^\varepsilon$, which is clear, and that the projections do as well. The fact that projections commute with the inclusion i follows from Construction A.10. The third claim follows at once from naturality. □

A.4. Metrics on orthogonal groups

In this section, we will make use of basic Riemannian geometry; see, for instance, [39]. Our main goals are to calculate the reach (Appendix A.1) of the orthogonal groups, seen as a subspace of the metric space of square matrices, with metric given by the Frobenius distance; to compare the Frobenius distance to the geodesic distance; and to calculate the systole (Definition 6.4) of the orthogonal groups using the geodesic distance.

The geodesic distance we consider on the orthogonal groups is the one that comes from the bi-invariant metric given by the smooth inclusion $O(d) \subseteq \mathbb{R}^{d \times d}$, where the inner product on $\mathbb{R}^{d \times d}$ is the one associated to the Frobenius norm [39, Example 3.16(e)].

Lemma A.13. *Let $M, N \in O(d)$ such that $\det(M) = 1 = -\det(N)$. Then $\|M - N\| \geq 2$.*

Proof. Since $O(d)$ acts on itself by isometries, it suffices to show that $\|N - \text{id}\| \geq 2$ whenever $\det(N) = -1$. Since N is orthogonal, there is an orthogonal change of basis that takes it to be diagonal by blocks, where blocks are 1×1 and equal 1 or -1 , or 2×2 and rotation matrices. We can thus assume that N is diagonal by blocks of the form above. If any of the blocks is a 1×1 block of the form -1 we are done. Otherwise, there must be a block that consists of a rotation with negative determinant. In this case,

$$\|N - \text{id}\| \geq \sqrt{2(\sin(\theta))^2 + (1 - \cos(\theta))^2 + (1 + \cos(\theta))^2} = 2. \quad \square$$

The following lemma appears in [5] and is a key result that lets us compare the geodesic distance of the orthogonal groups to the Frobenius distance. Before giving the result, we recall that any metric space has an induced geodesic distance (also known as a path-length distance), where the distance between any two points is taken to be the infimum of the lengths of the continuous paths between them [5, Section 2], [11, Section 2.3.3]. If the metric space is not connected, this is an extended distance. Finally, if the metric space consists of a manifold smoothly embedded in Euclidean space, then this geodesic distance coincides with the geodesic distance associated to the Riemannian metric induced by the embedding into Euclidean space [11, Exercise 5.1.8].

Lemma A.14 [5, Lemma 3]. *Let $S \subseteq \mathbb{R}^N$ be a closed set such that $R = \text{reach}(S) > 0$. Let d_S denote the geodesic distance on S induced by the restriction of the Euclidean distance. If $x, y \in S$ are such that $\|x - y\|_2 < 2R$, then $d_S(x, y) \leq 2R \arcsin\left(\frac{\|x - y\|}{2R}\right)$.*

Lemma A.15. *Let $A \in \mathbb{R}^{d \times d}$ and $\Omega \in O(d)$ such that $\|A - \Omega\| < 1$, then A has full rank.*

Proof. It is enough to show that $\Omega^t A$ has full rank, so, since the Frobenius norm is $O(d)$ -invariant, it is sufficient to prove it for the case $\Omega = \text{id}$. Let $B = \text{id} - A$. Since $\|B\| < 1$, and the Frobenius norm is submultiplicative, the matrix $\sum_{n \geq 0} B^n$ is well defined. Finally, we have $(\text{id} - B)(\sum_{n \geq 0} B^n) = \text{id}$, so $\text{id} - B = A$ is invertible, as required. \square

Lemma A.16. *Consider $O(d) \subseteq \mathbb{R}^{d \times d}$ with the Frobenius distance. Then $\text{reach}(O(d)) = 1$.*

Proof. Let M_x denote the $d \times d$ diagonal matrix with 1 in all diagonal entries except the first one, which is x . Since a polar decomposition of M_0 is given by $M_0 = \text{id}M_0$, we have that the identity is a closest orthogonal matrix to M_0 , by Appendix A.2. Now, note that $\|M_0 - \text{id}\| = 1 = \|M_0 - M_{-1}\|$, so M_{-1} is a closest orthogonal matrix to M_0 too, and thus $\text{reach}(O(d)) \leq 1$.

To conclude, we must show that, if $\|M - \text{id}\| < 1$, then M has a unique closest orthogonal matrix. This is a consequence of Lemma A.15 and Lemma A.3. \square

Combined, Lemma A.14 and Lemma A.16 give us upper bounds for the geodesic distance of $O(d)$ in terms of the Frobenius distance. We will need the following specific bound.

Corollary A.17. *Let d_G be the geodesic distance on $O(d)$ induced by the embedding $O(d) \subseteq \mathbb{R}^{d \times d}$. For $M, N \in O(d)$, if $\|M - N\| < 1$, then $d_G(M, N) < \sqrt{2}\pi/4$.*

Proof. By inspection, $2 \arcsin(1/2) < \sqrt{2}\pi/4$. Note that a slightly tighter bound is possible, but we prefer this one for readability. \square

Lemma A.18. *Using the geodesic distance d_G , we have $\text{sys}(O(d)) = 2\sqrt{2}\pi$.*

Proof. Since $O(d)$ is a group acting on itself by isometries, to compute $\text{sys}(O(d))$, it suffices to consider loops on the identity matrix $\text{id} \in O(d)$. We observe that the constant speed, locally length-minimizing curves from the identity to itself are in bijection with the logarithms $L \in \mathfrak{so}(d)$ of the identity. More specifically, any such curve can be written as $t \mapsto \exp(tL)$ with $L \in \mathfrak{so}(d)$ [39, Proposition 5.19]. We also observe that the speed of such a curve is $\|L\|$, the Frobenius norm of L .

So let $L \in \mathfrak{so}(d)$ be a skew-symmetric matrix such that $\exp(L) = \text{id}$. Since L is skew-symmetric, there exists an orthogonal change of basis such that L is a block-diagonal matrix with either 1×1 blocks containing a 0 or 2×2 blocks with 0 in the diagonal and $\lambda, -\lambda \in \mathbb{R}$ in the antidiagonal. Since L is a logarithm of the identity, the number λ in any of these blocks must be an integer multiple of 2π . It follows that $\|L\| = \sqrt{2 \sum (2\pi n_i)^2} = 2\pi\sqrt{2} \sqrt{\sum n_i^2}$. Since we are considering nonnullhomotopic loops, at least one of the integers n_i must be nonzero, and thus the smallest value of $\|L\|$ is attained when one of the n_i is 1 and the rest are zero, and in such case we have $\|L\| = 2\sqrt{2}\pi$, concluding the proof. \square

A.5. Riemannian manifolds and covering spaces

In this section, we use the systole to give a lower bound for the distance between two distinct elements in the fiber of a covering map between Riemannian manifolds and we prove a metrically controlled lifting property for covering maps between Riemannian manifolds.

Lemma A.19. *Let $G \rightarrow H$ be a covering map between Riemannian manifolds. Fix $h \in H$, and let $F \subseteq G$ be the fiber of h . Then, the infimum of the distances between distinct elements of F is bounded below by $\text{sys}(H)$.*

Proof. Let $a \neq b \in F$ and consider a geodesic between them. The length of this path is equal to the length of the path mapped to H since $G \rightarrow H$ is a local isometry. The path mapped to H cannot be nullhomotopic since a and b are distinct points of the fiber, so its length is bounded below by the systole of H . \square

Lemma A.20. *Let $\zeta : G \rightarrow H$ be a covering map between Riemannian manifolds, with H compact.*

1. *Let $h, h' \in H$ be such that $d_H(h, h') < \text{sys}(H)/2$ and let $g \in G$ be a preimage of h . Then, there exists a unique preimage g' of h' such that $d_G(g, g') < \text{sys}(H)/2$ and g' has the property that $d_G(g, g') = d_H(h, h')$.*
2. *Let U be locally path connected and simply connected and let $v, w : U \rightarrow H$ continuous such that, for all $z \in U$, $d_H(v(z), w(z)) < \text{sys}(H)/2$. Then, there exist lifts $\bar{v}, \bar{w} : U \rightarrow G$ of v and w , respectively, such that, for all $z \in U$, $d_G(\bar{v}(z), \bar{w}(z)) = d_H(v(z), w(z))$.*

Proof. We start with the first claim. To prove existence, consider a shortest geodesic from h to h' , whose length must be strictly less than $\text{sys}(H)/2$. By lifting this path, we get a preimage g' of h' such that $d_G(g, g') \leq d_H(h, h')$. To prove uniqueness, suppose that $g'' \neq g'$ is a preimage of h' . By Lemma A.19, we have $d_G(g'', g') \geq \text{sys}(H)$, so $d_G(g, g'') \geq \text{sys}(H)/2$. Finally, since $G \rightarrow H$ is 1-Lipschitz, we have $d_H(h, h') \leq d_G(g, g')$, and thus $d_G(g, g') = d_H(h, h')$.

For the second claim, let $y \in U$ and use the first claim to choose lifts $g, g' \in G$ of $v(y)$ and $w(y)$, respectively, such that $d_G(g, g') = d_H(v(y), w(y))$. Since U is locally path connected and simply connected, there exist unique lifts \bar{v} and \bar{w} of v and w , respectively, such that $\bar{v}(y) = g$ and $\bar{w}(y) = g'$. Consider the subset $U' = \{z \in U : d_G(\bar{v}(z), \bar{w}(z)) = d_H(v(z), w(z))\} \subseteq U$. This subset is closed and nonempty, so it suffices to show that it is open, since U is connected. We conclude the proof by proving this fact.

Let $z \in U'$, and let $\varepsilon = \text{sys}(H)/2 - d_H(v(z), w(z)) > 0$. Let N be an open neighborhood of z such that $\bar{v}(N)$ is contained in the open ball of radius $\varepsilon/2$ around $\bar{v}(z)$, and $\bar{w}(N)$ is contained in the open ball of radius $\varepsilon/2$ around $\bar{w}(z)$. If $z' \in N$, then

$$\begin{aligned} d_G(\bar{v}(z'), \bar{w}(z')) &\leq d_G(\bar{v}(z'), \bar{v}(z)) + d_G(\bar{v}(z), \bar{w}(z)) + d_G(\bar{w}(z), \bar{w}(z')) \\ &= d_G(\bar{v}(z'), \bar{v}(z)) + d_H(v(z), w(z)) + d_G(\bar{w}(z), \bar{w}(z')) \\ &< d_H(v(z), w(z)) + 2\varepsilon/2 = \text{sys}(H)/2. \end{aligned}$$

From the first claim, it follows that $d_G(\bar{v}(z'), \bar{w}(z')) = d_H(v(z'), w(z'))$, so $N \subseteq U'$, and thus U' is open, concluding the proof. \square

Acknowledgements. We thank Dan Christensen, Peter Landweber, Fernando Martin and Raphaël Tinarrage for helpful conversations and comments, and Ximena Fernández for discussions and suggesting the example in Section 7.2. We also thank Peter Landweber for suggesting improvements to the proofs in Appendix A.1. Finally, we thank the anonymous reviewers for helpful feedback, which has improved this paper.

Conflict of Interest. The authors have no conflict of interest to declare.

Funding statement. This work was partially supported by the National Science Foundation through grants CCF-2006661 and CAREER award DMS-1943758.

References

- [1] H. D. I. Abarbanel, *Analysis of Observed Chaotic Data*. (Institute for Nonlinear Science. Springer-Verlag, New York, 1996), xiv+272. doi: [10.1007/978-1-4612-0763-4](https://doi.org/10.1007/978-1-4612-0763-4).
- [2] D. Attali, A. Lieutier and D. Salinas, ‘Vietoris–Rips complexes also provide topologically correct reconstructions of sampled shapes’, *Comput. Geom.* **46**(4) (2013), 448–465. doi: [10.1016/j.comgeo.2012.02.009](https://doi.org/10.1016/j.comgeo.2012.02.009).
- [3] A. P. Balachandran, G. Marmo, B.-S. Skagerstam and A. Stern, *Gauge Symmetries and Fibre Bundles*, Lecture Notes in Physics. Applications to Particle Dynamics, vol. 188 (Springer-Verlag, Berlin, 1983), 140.
- [4] U. Bauer, ‘Rips: Efficient computation of Vietoris–Rips persistence barcodes’, Preprint, 2021, [arXiv1908.02518v2](https://arxiv.org/abs/1908.02518v2).
- [5] J.-D. Boissonnat, A. Lieutier and M. Wintraecken, ‘The reach, metric distortion, geodesic convexity and the variation of tangent spaces’, *J. Appl. Comput. Topol.* **3**(1–2) (2019), 29–58. doi: [10.1007/s41468-019-00029-8](https://doi.org/10.1007/s41468-019-00029-8).
- [6] R. Bott, ‘On some recent interactions between mathematics and physics’, *Canad. Math. Bull.* **28**(2) (1985), 129–164. doi: [10.4153/CMB-1985-016-3](https://doi.org/10.4153/CMB-1985-016-3).
- [7] R. Bott and L. W. Tu, *Differential Forms in Algebraic Topology*, Graduate Texts in Mathematics, vol. 82 (Springer-Verlag, New York-Berlin, 1982), xiv+331.
- [8] T. Bröcker and T. t. Dieck, *Representations of Compact Lie Groups*, Graduate Texts in Mathematics, vol. 98 (Springer-Verlag, New York, 1995), x+313. isbn: 0-387-13678-9. Translated from the German manuscript, Corrected reprint of the 1985 translation.
- [9] J.-L. Brylinski and D. A. McLaughlin, ‘Čech cocycles for characteristic classes’, *Comm. Math. Phys.* **178**(1) (1996), 225–236. doi: [10.1007/BF02104916](https://doi.org/10.1007/BF02104916).
- [10] R. B. Gabrielsson and G. Carlsson, ‘Exposition and Interpretation of the Topology of Neural Networks’, in *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)* (IEEE, 2019), 1069–1076. doi: [110.1109/ICMLA.2019.00180](https://doi.org/10.1109/ICMLA.2019.00180).
- [11] D. Burago, Y. Burago and S. Ivanov, *A Course in Metric Geometry*, Graduate Studies in Mathematics, vol. 33 (American Mathematical Society, Providence, RI, 2001), xiv+415. doi: [10.1090/gsm/033](https://doi.org/10.1090/gsm/033).
- [12] G. Carlsson, T. Ishkhanov, V. de Silva and A. Zomorodian, ‘On the local behavior of spaces of natural images’, *Int. J. Comput. Vis.* **76**(1) (2008), 1–12. doi: [10.1007/s11263-007-0056-x](https://doi.org/10.1007/s11263-007-0056-x).
- [13] G. D. Charó, G. Artana and D. Sciamarella, ‘Topology of dynamical reconstructions from Lagrangian data’, *Physica D: Nonlinear Phenomena* **405** (2020), 132371. doi: [10.1016/j.physd.2020.132371](https://doi.org/10.1016/j.physd.2020.132371).
- [14] S. H. Christiansen and T. G. Halvorsen, ‘A simplicial gauge theory’, *J. Math. Phys.* **53**(3) (2012), 033501, 17. doi: [10.1063/1.3692167](https://doi.org/10.1063/1.3692167).
- [15] J. Curry, ‘The fiber of the persistence map for functions on the interval’, *J. Appl. Comput. Topol.* **2**(3–4) (2018), 301–321. doi: [10.1007/s41468-019-00024-z](https://doi.org/10.1007/s41468-019-00024-z).
- [16] X. Fernández, ‘Topology of fluid flows’, 2022, https://github.com/ximenafernandez/topology_fluids.
- [17] M. Feshbach, ‘The integral cohomology rings of the classifying spaces of $O(n)$ and $SO(n)$ ’, *Indiana Univ. Math. J.* **32**(4) (1983), 511–516. doi: [10.1512/iumj.1983.32.32036](https://doi.org/10.1512/iumj.1983.32.32036).

- [18] J. Frank, *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State: Visualization of Biological Molecules in Their Native State* (Oxford University Press, USA, 2006).
- [19] T. Gao, J. Brodzki and S. Mukherjee, 'The geometry of synchronization problems and learning group actions', *Discrete & Computational Geometry* **65** (2019), 150–211. doi: [10.1007/s00454-019-00100-2](https://doi.org/10.1007/s00454-019-00100-2).
- [20] R. Ghrist, 'Barcodes: the persistent topology of data', *Bull. Amer. Math. Soc. (N.S.)* **45**(1) (2008), 61–75. doi: [10.1090/S0273-0979-07-01191-3](https://doi.org/10.1090/S0273-0979-07-01191-3).
- [21] D. Govc, W. Marzantowicz and P. Pavešić, 'How many simplices are needed to triangulate a Grassmannian?' *Topological Methods in Nonlinear Analysis* **56**(2) (2020), 501–518. doi: [10.12775/TMNA.2020.027](https://doi.org/10.12775/TMNA.2020.027).
- [22] M. Gromov, *Metric Structures for Riemannian and Non-Riemannian Spaces, Modern Birkhäuser Classics* (Birkhäuser, Boston, MA, 2007), xx+585. English.
- [23] R. Hadani and A. Singer, 'Representation theoretic patterns in three dimensional cryo-electron microscopy I: The intrinsic reconstitution algorithm', *Ann. of Math. (2)* **174**(2) (2011), 1219–1241. doi: [10.4007/annals.2011.174.2.11](https://doi.org/10.4007/annals.2011.174.2.11).
- [24] R. Hadani and A. Singer, 'Representation theoretic patterns in three-dimensional cryo-electron microscopy II—The class averaging problem', *Found. Comput. Math.* **11**(5) (2011), 589–616. doi: [10.1007/s10208-011-9095-3](https://doi.org/10.1007/s10208-011-9095-3).
- [25] J. Hansen and R. Ghrist, 'Toward a spectral theory of cellular sheaves', *J. Appl. Comput. Topol.* **3**(4) (2019), 315–358. doi: [10.1007/s41468-019-00038-7](https://doi.org/10.1007/s41468-019-00038-7).
- [26] A. Hatcher, *Algebraic Topology* (Cambridge University Press, Cambridge, 2002), xii+544.
- [27] A. Hatcher, *Vector Bundles and K-Theory*, 2003, <https://pi.math.cornell.edu/~hatcher/VBKT/VBpage.html>.
- [28] M. van Heel, B. Gowen, R. Matadeen, E. V. Orlova, R. Finn, T. Pape, D. Cohen, H. Stark, R. Schmidt, M. Schatz and A. Patwardhan, 'Single-particle electron cryo-microscopy: towards atomic resolution', *Q Rev Biophys* **33**(4) (2000), 307–369.
- [29] G. Hinton and R. Salakhutdinov, 'Reducing the dimensionality of data with neural networks', *Science* **313**(5786) (2006), 504–507.
- [30] G. E. Hinton, A. Krizhevsky and S. D. Wang, 'Transforming auto-encoders', in *Proceedings of the 21th International Conference on Artificial Neural Networks – Volume Part I*. (Springer-Verlag, Berlin, Heidelberg, 2011), 44–51.
- [31] R. A. Horn and C. R. Johnson, *Matrix Analysis*, second edn., (Cambridge University Press, 2012).
- [32] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis* (Cambridge University Press, Cambridge, 1991), viii+607. doi: [10.1017/CBO9780511840371](https://doi.org/10.1017/CBO9780511840371).
- [33] D. Husemüller, M. Joachim, B. Jurčo and M. Schottenloher, *Basic Bundle Theory and K-Cohomology Invariants*, Lecture Notes in Physics, vol. 726 Springer, Berlin, 2008, xvi+340. With contributions by S. Echterhoff, S. Fredenhagen and B. Krötz.
- [34] D. Husemüller, *Fibre Bundles*, third edn., *Graduate Texts in Mathematics*, vol. 30 (Springer-Verlag, New York, 1994), xx+353. doi: [10.1007/978-1-4757-2261-1](https://doi.org/10.1007/978-1-4757-2261-1).
- [35] J. Jost, *Riemannian Geometry and Geometric Analysis*, second edn., Universitext (Springer-Verlag, Berlin, 1998), xiv+455. doi: [10.1007/978-3-662-22385-7](https://doi.org/10.1007/978-3-662-22385-7).
- [36] R. C. Kirby and L. R. Taylor, 'Pin structures on low-dimensional manifolds', in *Geometry of Low-Dimensional Manifolds*, 2 (Durham, 1989), London Math. Soc. Lecture Note Ser., vol. 151 (Cambridge Univ. Press, Cambridge, 1990), 177–242.
- [37] F. Knöppel and U. Pinkall, 'Complex line bundles over simplicial complexes and their applications', in *Advances in Discrete Differential Geometry*, edited by A. I. Bobenko (Berlin, Heidelberg: Springer Berlin Heidelberg, 2016), 197–239. doi: [10.1007/978-3-662-50447-5_6](https://doi.org/10.1007/978-3-662-50447-5_6).
- [38] H. B. Lawson Jr. and M.-L. Michelsohn, *Spin Geometry*, Princeton Mathematical Series, vol. 38 (Princeton University Press, Princeton, NJ, 1989), xii+427.
- [39] J. M. Lee, *Introduction to Riemannian Manifolds*, second edn., *Graduate Texts in Mathematics*, vol. 176 (Springer, Cham, 2018), xiii+437.
- [40] S. Lim, F. Memoli and O. B. Okutan, 'Vietoris–Rips persistent homology, injective metric spaces and the filling radius', Preprint, 2020, [arXiv:2001.07588](https://arxiv.org/abs/2001.07588).
- [41] A. V. Little, J. Lee, Y.-M. Jung and M. 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, Cardiff, UK, 2009, pp. 85–88. doi: [10.1109/SSP.2009.5278634](https://doi.org/10.1109/SSP.2009.5278634).
- [42] P. Lounesto, *Clifford Algebras and Spinors*, second edn., London Mathematical Society Lecture Note Series, vol. 286 (Cambridge University Press, Cambridge, 2001), x+338. doi: [10.1017/CBO9780511526022](https://doi.org/10.1017/CBO9780511526022).
- [43] G. Luke and A. S. Mishchenko, *Vector Bundles and Their Applications*, Mathematics and Its Applications, vol. 447 (Kluwer Academic Publishers, Dordrecht, 1998), viii+254. doi: [10.1007/978-1-4757-6923-4](https://doi.org/10.1007/978-1-4757-6923-4).
- [44] D. A. McLaughlin, 'Local formulae for Stiefel-Whitney classes', *Manuscripta Math.* **89**(1) (1996), 1–13. doi: [10.1007/BF02567501](https://doi.org/10.1007/BF02567501).
- [45] J. Milnor, 'Construction of universal bundles. II', *In Ann. of Math. (2)* **63** (1956), 430–436. doi: [10.2307/1970012](https://doi.org/10.2307/1970012).
- [46] J. Milnor, 'Curvatures of left invariant metrics on Lie groups', *Advances in Math.* **21**(3) (1976), 293–329. doi: [10.1016/S0001-8708\(76\)80002-3](https://doi.org/10.1016/S0001-8708(76)80002-3).
- [47] J. W. Milnor and J. D. Stasheff, *Characteristic Classes*, Annals of Mathematics Studies, no. 76 (Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1974), vii+331.
- [48] G. Naitzat, A. Zhitnikov and L.-H. Lim, 'Topology of deep neural networks', *J. Mach. Learn. Res.* **21** (2020), Paper No. 1, 40.
- [49] Partha Niyogi, S. Smale and S. Weinberger, 'Finding the homology of submanifolds with high confidence from random samples', *Discrete Comput. Geom.* **39**(1) (2008), 419–441.

- [50] S. Y. Oudot, *Persistence Theory: From Quiver Representations to Data Analysis*, Mathematical Surveys and Monographs, vol. 209 (American Mathematical Society, Providence, RI, 2015), viii+218. doi:[10.1090/surv/209](https://doi.org/10.1090/surv/209).
- [51] J. A. Perea, ‘Topological times series analysis’, *Notices Amer. Math. Soc.* **66**(5) (2019), 686–694.
- [52] J. A. Perea and G. Carlsson, ‘A Klein-bottle-based dictionary for texture representation’, *Int. J. Comput. Vis.* **107**(1) (2014), 75–97. doi:[10.1007/s11263-013-0676-2](https://doi.org/10.1007/s11263-013-0676-2).
- [53] M. A. Rieffel, ‘Vector bundles and Gromov–Hausdorff distance’, *Journal of K-Theory* **5**(1) (2010), 39–103. doi:[10.1017/is008008014jkt080](https://doi.org/10.1017/is008008014jkt080).
- [54] M. Robinson, ‘Assignments to sheaves of pseudometric spaces’, *Compositionality* **2** (2 2020). doi: [10.32408/compositionality-2-2](https://doi.org/10.32408/compositionality-2-2).
- [55] M. Robinson, ‘Sheaves are the canonical data structure for sensor integration’, *Information Fusion* **36** (2017), 208–224. <https://doi.org/10.1016/j.inffus.2016.12.002>.
- [56] L. Scoccola and J. Perea, ‘Approximate vector bundles, computational examples’. 2021, <https://github.com/LuisScoccola/approximate-vector-bundles-examples.git>.
- [57] S. C. Shadden, F. Lekien and J. E. Marsden, ‘Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows’, *Physica D: Nonlinear Phenomena* **212**(3) (2005), 271–304. <https://doi.org/10.1016/j.physd.2005.10.007>.
- [58] A. Singer, Z. Zhao, Y. Shkolnisky and R. Hadani, ‘Viewing angle classification of cryo-electron microscopy images using eigenvectors’, *SIAM J. Imaging Sci.* **4**(2) (2011), 723–759. doi: [10.1137/090778390](https://doi.org/10.1137/090778390).
- [59] A. Singer and H.-t. Wu, ‘Orientability and diffusion maps’, *Appl. Comput. Harmon. Anal.* **31**(1) (2011), 44–58. doi: [10.1016/j.acha.2010.10.001](https://doi.org/10.1016/j.acha.2010.10.001).
- [60] S. Smale, ‘Differentiable dynamical systems’, *Bull. Amer. Math. Soc.* **73** (1967), 747–817. doi: [10.1090/S0002-9904-1967-11798-1](https://doi.org/10.1090/S0002-9904-1967-11798-1).
- [61] N. Steenrod. ‘The Topology of Fibre Bundles’, Princeton Mathematical Series, vol. 14 (Princeton University Press, Princeton, NJ, 1951), viii+224.
- [62] F. Takens, ‘Detecting strange attractors in turbulence’, in *Dynamical Systems and Turbulence, Warwick 1980 (Coventry, 1979/1980)*, Lecture Notes in Math., vol. 898 (Springer, Berlin-New York, 1981), 366–381.
- [63] R. Tinarrage, ‘Computing persistent Stiefel–Whitney classes of line bundles’, *J. Appl. Comput. Topol.* **6**(1) (2022), 65–125. doi: [10.1007/s41468-021-00080-4](https://doi.org/10.1007/s41468-021-00080-4).
- [64] C. Tralie, T. Mease and J. Perea, *DREiMac*, 2017, <https://github.com/scikit-tda/DREiMac>.
- [65] C. Tralie, N. Saul and R. Bar-On, ‘Ripser.py: A lean persistent homology library for Python’, *The Journal of Open Source Software* **3**(29) (2018), 925. doi: [10.21105/joss.00925](https://doi.org/10.21105/joss.00925).
- [66] F. W. Warner. *Foundations of Differentiable Manifolds and Lie Groups*, Graduate Texts in Mathematics, vol. 94 (Springer-Verlag, New York-Berlin, 1983), ix+272. Corrected reprint of the 1971 edition.
- [67] R. F. Williams, ‘Expanding attractors’, *Inst. Hautes Études Sci. Publ. Math.* **43** (1974), 169–203.
- [68] B. Xu, C. J. Tralie, A. Antia, M. Lin and J. A. Perea, ‘Twisty takens: A geometric characterization of good observations on dense trajectories’, *J. Appl. Comput. Topol.* **3**(4) (2019), 285–313. doi: [10.1007/s41468-019-00036-9](https://doi.org/10.1007/s41468-019-00036-9).
- [69] K. Ye and L.-H. Lim, ‘Cohomology of cryo-electron microscopy’, *SIAM Journal on Applied Algebra and Geometry* **1**(1) (2017), 507–535.