

Assignments to sheaves of pseudometric spaces

Michael Robinson

Mathematics and Statistics, American University, Washington, DC, USA

An assignment to a sheaf is the choice of a local section from each open set in the sheaf's base space, without regard to how these local sections are related to one another. This article explains that the *consistency radius* — which quantifies the agreement between overlapping local sections in the assignment — is a continuous map. When thresholded, the consistency radius produces the *consistency filtration*, which is a filtration of open covers. This article shows that the consistency filtration is a functor that transforms the structure of the sheaf and assignment into a nested set of covers in a structure-preserving way. Furthermore, this article shows that consistency filtration is robust to perturbations, establishing its validity for arbitrarily thresholded, noisy data.

Contents

1	Introduction	1
2	Interpretation and applications	3
3	Historical context	4
4	Preliminaries	4
5	Sheaves paired with assignments	7
6	Coarsening filtrations	12
7	Consistency radius on subspaces	15
8	Consistency filtrations	18
9	Sheaves on partial orders	22

1 Introduction

An *assignment* of a sheaf of sets on a topological space selects a local section from each open set in the sheaf's base space without regard to how these local sections are related to one another. This article explores the structure of a *sheaf of pseudometric spaces on a topological space* by analyzing the structure of its assignments. More structure is present in a sheaf of pseudometric spaces than in a sheaf of sets, since each stalk is a set equipped with a *pseudometric*, and each restriction is required to be continuous with respect to the pseudometrics present on its domain and codomain.

A sheaf of pseudometric spaces is best thought of as a tool for measuring the self-consistency between the local sections in an assignment. When data are present in a *global section* on two overlapping open sets, their consistency is manifest by being equal upon restricting to the intersection of the open sets. For a sheaf of sets, an assignment is self-consistent when it is a global section, and that is all that can be said. But for a sheaf of pseudometric spaces, instead of requiring

Michael Robinson: michaelr@american.edu

equality, we can instead ask for the data to restrict to values within some distance of each other. The supremum¹ of these distances is called the *consistency radius* of the assignment. Since global sections evidently have zero consistency radius — the distance between data restricted to the same open set is always zero in a global section because they are equal — positive consistency radius is an obstruction to an assignment being a global section.

It may happen that some parts of an assignment are more consistent than others, which is to say that restricting the consistency radius to subspaces of the sheaf’s base space yields new information. The resulting *local consistency radius* can be used to identify open sets on which the assignment and sheaf are consistent to some specified level. Sweeping a threshold value over the local consistency radius identifies portions of the base space on which an assignment is more or less in agreement with the restrictions of the sheaf. Since increasing the threshold loosens the tolerance for data to restrict to the same value, the result is something akin to a filtration: any open set on which the assignment is consistent at a lower threshold will remain consistent at a higher threshold. Making this intuition precise, one obtains the *consistency filtration*.

Providing precise definitions of *consistency radius* and *consistency filtration* occupy some of our attention in this article, but their value is squarely rooted in the properties we prove about them. From the standpoint of applications, this article establishes that small perturbations in the data (the assignment) or the model (the sheaf) do not result in large changes in the consistency radius or in the consistency filtration. Furthermore, the consistency filtration preserves specific relationships between related sheaves and assignments — it is a covariant functor between appropriately constructed categories. Although the results in this article require a limited base topological space — sheaves on finite spaces are the focus, though a few results described in this article generalize to arbitrary topological spaces — we do not restrict to sheaves over a single topological space. *Sheaf morphisms* will be permitted to transform a sheaf on one space to a sheaf on another space.

This dual nature of the consistency filtration — it is both a continuous function between topological spaces *and* a covariant functor between categories — arises because we can view the class of all sheaf-assignment pairs through two lenses: topological and categorial. The stalks and restrictions of a given sheaf of pseudometric spaces are manifestly topological, but that same sheaf is also a special kind of functor. This dual nature is not present for a sheaf of sets, since most of its interesting properties arise from being a functor, nor is this dual nature present for a continuous map between pseudometric spaces, since most of its interesting properties arise from being topological.

The topological viewpoint of sheaf-assignment pairs addresses the question of how to threshold data based on self-consistency, using the sheaf of pseudometric spaces as a tool to measure the assignment. From this viewpoint, it is best to think of the consistency radius as a continuous function of the sheaf and of the assignment (jointly!). Since the consistency radius is a nonnegative real number that quantifies how compatible the sheaf and assignment are (smaller means more consistent), it has a straightforward interpretation. We can therefore turn the consistency radius on its head and ask, “for a given real number threshold, which open sets have local consistency radius smaller than that threshold?” The answer to this is the consistency filtration, which although it is not a filtration of subspaces in the usual sense, it is an order preserving function from the real numbers to the power set of the sheaf’s base space topology, in which the ordering is given by *coarsening*². There the topological story would seem to end, if not for the recent definition of a general *interleaving distance* [12]. With the interleaving distance in hand, this article shows that the consistency filtration is a continuous function as well. This establishes that the consistency filtration is robust to perturbations, with *persistent Čech cohomology* as a special case.

In contrast, the categorial viewpoint emphasizes transformations of sheaf-assignment pairs, first by generalizing the notion of a sheaf morphism to address the case where the base space of the domain and the base space of the codomain differ, and secondly by focusing attention on only those sheaf morphisms that preserve *both* the sheaf and assignment. These definitions put the category of assignments of a sheaf of pseudometric spaces on firm theoretical ground. With this specialized category in hand, consistency radius is not a functor, even though the morphisms transform it in a predictable way (Lemma 4). This suggests that there still is a functor at work,

¹or other suitable aggregate

²The opposite of the *refinement* of a cover!

which unsurprisingly turns out to be the consistency filtration when an appropriate category is constructed for its codomain.

In taking these two viewpoints, this article defines several apparently novel categories for sheaf-assignment pairs and filtrations of covers, and obtains three results about their dual topological/categorical natures:

Theorem 1 : Consistency radius is a continuous function,

Theorem 2 : Consistency filtration is a covariant functor, and

Theorem 3 : Consistency filtration is a continuous function.

2 Interpretation and applications

A *sheaf of sets on a topological space* is a mathematical structure for representing local data. While sheaves are general purpose and abstract, they do not require unrealistic assumptions about domain-specific data. As such, they provide a convenient language for describing how systems composed of interrelated parts can interact. In applications, this is valuable because sheaves can naturally represent systems composed of different types of subsystems. Since subsystems are usually modeled (or engineered) separately and largely in isolation from one another, composing the state of several subsystems into a global composite state is not trivial. If one requires that the composite system be represented as a sheaf, this imposes specific conditions on how subsystems may be assembled. If these conditions are met, then the idea of a global composite state — a *global section* of the sheaf — is well-defined and can be obtained algorithmically [31]. From a practical standpoint, this means that the process of asserting that the composite system be modeled as a sheaf simplifies and standardizes the bookkeeping tasks involved with modeling the composite system.

Over the years, these considerations have motivated various researchers to encode a number of systems as sheaves, such as

- Bayes nets and graphical models [30],
- Biochemical networks [28],
- Communication networks [9; 24],
- Control systems [18; 19],
- Differential equations [29] and their discretizations (also as *dual sheaves*) [30],
- Discrete and continuous dynamical systems (as *cosheaves*) [26],
- Flow networks [10; 21],
- Formal models of interacting software objects [11; 16; 17; 20; 34],
- Multi-dimensional weighted or labeled graphs [14; 15; 22],
- Quantum information systems [1; 2],
- Sensor networks [31], and
- Signal processing chains [25].

One may ask whether encoding any of these composite systems as a sheaves confers any particular benefit, or if sheaves are merely another universal encoding. The payoff is that sheaf theory provides *invariants*, which are general analytic tools that can be easily applied to perform certain data processing tasks, largely without imposing additional assumptions on the subsystem models. This article discusses two of these invariants, the *consistency radius* and *consistency filtration*, that quantify the compatibility between the system's representation (which may be hypothetical, incomplete, or inaccurate) and data collected about the state of subsystems (which may be noisy or subject to other kinds of systematic errors). Although this article does not directly address any of the specific systems mentioned above, the results obtained apply broadly to all of them.

3 Historical context

Recent years have seen a rapid growth in the field of *topological data analysis*, which is the study of the topological properties of datasets. The typical analysis involves transforming a *point cloud* — a collection of data points in a high-dimensional metric space — into a combinatorial model of a topological space, such as an abstract simplicial complex, whose properties are then studied. For instance, the *simplicial homology* of an abstract simplicial complex can be used to identify potentially interesting features. These combinatorial models typically depend on the choice of a threshold that is used to determine if two of the data points are “close enough” to be joined together. One of the primary insights of topological data analysis is *persistence*: one should not select a single threshold, but instead consider the ensemble of all thresholds. This idea led to the *persistent homology* of a point cloud [6; 7; 8; 35], an idea that has since become quite popular because of its broad applicability and straightforward usage.

Even though persistent homology has wide appeal, it is limited to relatively homogeneous data sets, essentially those that can be represented as point clouds. Data sets with more structure than that are difficult to study. Sheaves on a topological space provide a good framework for studying these kind of data sets, since they permit the structure of the data to vary over the base space. In most applications, it suffices to consider either *constructible sheaves* or sheaves on a *finite space* [4; 32]. Traditionally, sheaves on a given topological space are studied using *sheaf cohomology*, a global algebraic summary that respects the category of sheaves. The cohomology of a sheaf includes a representation of the space of global sections of the sheaf. Computing the cohomology is greatly simplified for constructible sheaves [33] and can be done efficiently using discrete Morse theory [5]. Given this happy situation, sheaf cohomology has been useful in applications to network science [9; 21; 29] and quantum information [1; 2]. Sheaf cohomology contains information beyond the global sections, namely obstructions to the existence of other global sections, and this too can be useful in certain cases [23].

Useful as it may be, sheaf cohomology cannot be defined if the stalks of a sheaf are merely sets, as it requires the stalks to have the structure of abelian groups [13]. Unfortunately, in many applications, such as those listed in the previous section, this algebraic structure is simply unavailable. However, for many systems, the stalks have geometric structure instead of algebraic structure. Specifically, the distance between two points within a stalk can be measured with a *pseudometric*. It is therefore convenient and appropriate to consider sheaves of pseudometric spaces instead of sheaves of abelian groups. In an earlier article [31], the author initiated the study of sheaves of pseudometric spaces, by showing how to encode a general model of a sensor system as a sheaf of pseudometric spaces, and how to encode the measurements made by such a system as an *assignment* to that sheaf. Under this framework, the *consistency radius* can be defined, which bounds the distance between the assignment and the nearest global section of the sheaf. This bound motivated an optimization algorithm to fuse potentially noisy or uncertain measurements made by the sensor system into a single globalized measurement, represented as a global section.

This article applies the idea of persistence — using an ensemble of thresholds rather than one threshold — to the consistency radius. The result is the definition of a new mathematical object called the *consistency filtration*. The consistency filtration fits within the landscape of topological data analysis tools, since persistent Čech cohomology for point clouds is the consistency filtration for a particular sheaf and assignment (Example 4), but it is considerably more general. Indeed, the two main results proven in this article about the consistency filtration — that it is both a covariant functor and a continuous function — rely on a generalization of the interleaving distance [12] developed for use in persistent homology.

4 Preliminaries

This article involves the study of *sheaves*, mathematical models of local consistency relationships between data. Our primary focus is a *sheaf of pseudometric spaces on a topological space*, which is a rather elaborate object. It is easiest to develop the definition of this object slowly, by unpacking each term in turn, as is done in this section.

Definition 1. Suppose (X, \mathcal{T}) is a topological space. A *presheaf* \mathcal{P} of sets on (X, \mathcal{T}) consists of the following specification:

1. For each open set $U \in \mathcal{T}$, a set $\mathcal{P}(U)$, called the *stalk* at U ,
2. For each pair of open sets $U \subseteq V$, there is a function $\mathcal{P}(U \subseteq V) : \mathcal{P}(V) \rightarrow \mathcal{P}(U)$, called a *restriction function* (or just a *restriction*), such that
3. $\mathcal{P}(U \subseteq U)$ is the identity function and
4. For each triple $U \subseteq V \subseteq W$ of open sets, $\mathcal{P}(U \subseteq W) = \mathcal{P}(U \subseteq V) \circ \mathcal{P}(V \subseteq W)$.

Intuitively, a presheaf is a data structure for holding a variety of data items, localized to each open set. The stalk $\mathcal{P}(U)$ at an open set U specifies the possible values of a datum localized at U . It is also the case that a presheaf \mathcal{P} can be succinctly defined as “a contravariant functor from the category of open sets of (X, \mathcal{T}) , with inclusion functions being morphisms, to the category **Set** of sets and functions.”

Recall that a *finite topological space* is a topological space (X, \mathcal{T}) in which \mathcal{T} is a finite set. Finite topological spaces are automatically *Alexandrov spaces*, which means that intersections of arbitrarily many open sets remains open. For Alexandrov spaces, the *star* of some subset $A \subseteq X$, given by

$$\text{star } A := \bigcap \{U \in \mathcal{T} : A \subseteq U\}$$

is the smallest open set that contains A .

Remark 1. Our usage of the word *stalk* differs from the traditional usage, in which $\mathcal{P}(U)$ is called *set of local sections*. The traditional definition of a *stalk* is a *direct limit* of sets of the form $\mathcal{P}(U)$ such that each open set U contains a given point $x \in X$. The distinction between these two usages becomes primarily linguistic for finite topological spaces: the traditional definition deems $\mathcal{P}(U)$ a stalk at $x \in X$ whenever an open set U is $\text{star } \{x\}$, making U minimal in the sense of inclusion. If the topology is not finite, there may not necessarily be such a minimal open U at each point, but the direct limit still exists. Our usage here merely removes the minimality requirement since it happens to be unimportant in this article.

The restrictions identify how to relate a datum on a large open set to one on a smaller open set. When the smaller open set already has a value, this may or may not agree with the value propagated by a restriction function from a larger open set. When these data agree is a special situation; the data are then called a *section* as the next definition explains.

Definition 2. A *global section* of a presheaf \mathcal{P} of sets on a topological space (X, \mathcal{T}) is an element s of the direct product³ $\prod_{U \in \mathcal{T}} \mathcal{P}(U)$ such that for all $U \subseteq V \in \mathcal{T}$ then $\mathcal{P}(U \subseteq V)(s(V)) = s(U)$. A *local section* is defined similarly, but refers to a subcollection \mathcal{U} of \mathcal{T} instead of \mathcal{T} .

The set of global sections of a presheaf \mathcal{P} may be quite different from the set $\mathcal{P}(X)$. It is for this reason that an additional *gluing axiom* can be included to remove this distinction.

Definition 3. Let \mathcal{S} be a presheaf of sets on a topological space (X, \mathcal{T}) . We call \mathcal{S} a *sheaf of sets* on (X, \mathcal{T}) if for every open set $U \in \mathcal{T}$ and every collection of open sets $\mathcal{U} \subseteq \mathcal{T}$ with $U = \bigcup \mathcal{U}$, then $\mathcal{S}(U)$ is isomorphic to the space of local sections over the set of elements \mathcal{U} .

This article is concerned with more than mere agreement or disagreement; more can be said if the distance between two data can be measured. The appropriate way to measure distance between two data is with a *pseudometric*.

Definition 4. A *pseudometric* d_X on a set X is a function $d : X \times X \rightarrow \mathbb{R}$ that satisfies the following axioms:

1. $d_X(x, y) \geq 0$,
2. $d_X(x, x) = 0$,

³Which is in general *not* the direct sum, since \mathcal{T} may be infinite!

- 3. $d_X(x, y) = d_X(y, x)$, and
- 4. $d_X(x, z) \leq d_X(x, y) + d_X(y, z)$

for all $x, y, z \in X$. The pair (X, d_X) is called a *pseudometric space*.

It is a standard fact that every pseudometric d_X on X induces a topology \mathcal{T} on X generated by the *open balls*

$$B_r(x) := \{y \in X : d_X(x, y) < r\}$$

for each $x \in X$ and each $r > 0$. We will usually abuse terminology slightly by regarding a pseudometric space as a topological space. We can deem a function $f : X \rightarrow Y$ to be a *continuous map* if d_X is a pseudometric on X and d_Y is a pseudometric on Y , and if f is a continuous function for the topologies induced by d_X and d_Y . We will write $f : (X, d_X) \rightarrow (Y, d_Y)$ to emphasize the situation where f is a continuous map.

Definition 5. The category **Pseud** is the category of *pseudometric spaces*, in which objects are pseudometric spaces and the morphisms are continuous maps.

In this article, it is often important to measure “how continuous” a continuous map is. If there is a $K > 0$ such that

$$d_Y(f(x), f(y)) \leq Kd_X(x, y)$$

for all $x, y \in X$, this K is called a *Lipschitz constant* for $f : (X, d_X) \rightarrow (Y, d_Y)$. It is easy to demonstrate that the existence of a finite Lipschitz constant for a function $f : X \rightarrow Y$ between two pseudometric spaces (X, d_X) and (Y, d_Y) implies that f is continuous. Such a function is said to be *Lipschitz continuous*.

With these tools in hand, namely sheaves of sets on a topological space and pseudometric spaces, we combine them into one concept that carries through the rest of the article.

Definition 6. A *sheaf \mathcal{S} of pseudometric spaces on a topological space (X, \mathcal{T})* consists of a sheaf of sets \mathcal{S} on (X, \mathcal{T}) and the choice of a pseudometric d_U on $\mathcal{S}(U)$ for each open set $U \in \mathcal{T}$, such that each restriction function is a continuous map. We will call each restriction a *restriction map* in this situation to emphasize that we are working with a sheaf of pseudometric spaces rather than a sheaf of sets.

Succinctly, a sheaf \mathcal{S} of pseudometric spaces on a topological space is a kind of contravariant functor from the category of open sets in (X, \mathcal{T}) to **Pseud**. With this interpretation in mind, we may simply think of each stalk $\mathcal{S}(U)$ as *being* a pseudometric space $(\mathcal{S}(U), d_U)$, even though this may be considered a slight abuse of notation.

It is immediate from the definitions that each space of (local or global) sections of a sheaf of pseudometric spaces is itself a topological space. In most of what follows, we will work with sheaves on a *finite topological space (X, \mathcal{T})* , one in which the topology \mathcal{T} is a finite set. Each space of sections for a sheaf of pseudometric spaces on a finite topological space is then a pseudometric space as well, by the usual constructions of a pseudometric on a product space. This idea leads to the next definition.

Definition 7. [31]. For a sheaf \mathcal{S} of sets on a topological space (X, \mathcal{T}) , an *assignment* is an element $a \in \prod_{U \in \mathcal{T}} \mathcal{S}(U)$. If \mathcal{S} is a sheaf of pseudometric spaces, then the set of assignments has the *assignment pseudometric* given by

$$D(a, b) := \sup_{U \in \mathcal{T}} d_U(a(U), b(U)).$$

For an assignment a to a sheaf \mathcal{S} , each value

$$d_U((\mathcal{S}(U \subseteq V)) a(V), a(U))$$

where $U \subseteq V \in \mathcal{T}$, is called a *critical threshold*. The *consistency radius* given by

$$c_{\mathcal{S}}(a) := \sup_{U \subseteq V \in \mathcal{T}} d_U((\mathcal{S}(U \subseteq V)) a(V), a(U)),$$

is the supremum of all critical thresholds.

The central relationship between global sections of \mathcal{S} and assignments is captured by the following bound.

Proposition 1. [31, Prop. 23] For an assignment a to a sheaf \mathcal{S} of pseudometric spaces on (X, \mathcal{T}) in which each restriction map of \mathcal{S} is Lipschitz with constant K , then

$$D(a, s) \geq \frac{c_{\mathcal{S}}(a)}{1 + K}$$

for every global section s of \mathcal{S} .

Remark 2. Assignments thresholded to a certain level of consistency are global sections of a (different) sheaf on a subdivision of the base space [27]. While this perspective has theoretical merit, it is useful to study the extent to which assignments are consistent to a given threshold. This perspective is introduced by Praggastis in [22], where it is shown that assignments supported on the vertices of an abstract simplicial complex yield a cover of that simplicial complex when they are thresholded to a certain level of consistency. We take a parallel approach in this article, but instead consider assignments supported on the entire space. The most prominent difference is in the definition of maximal consistent cover (Definition 17 and Lemma 5), which is analogous to the cover constructed in [22], but our cover is only over a subspace and is generally finer.

Remark 3. The consistency radius is a *radius*, rather than a diameter, because the assignment itself acts like a central point to which values propagated along the restrictions are compared.

We could instead define a quantity

$$d_{\mathcal{S}}(a) := \sup_{U \subseteq V_1 \in \mathcal{T}} \sup_{U \subseteq V_2 \in \mathcal{T}} d_U((\mathcal{S}(U \subseteq V_1))a(V_1), \mathcal{S}(U \subseteq V_2)a(V_2)),$$

which is a *diameter* in that

$$c_{\mathcal{S}}(a) \leq d_{\mathcal{S}}(a) \leq 2c_{\mathcal{S}}(a).$$

The left inequality arises simply by taking $U = V_2$ since $\mathcal{S}(U \subseteq U) = \text{id}_{\mathcal{S}(U)}$ by definition. The right inequality is a short calculation

$$\begin{aligned} d_{\mathcal{S}}(a) &= \sup_{U \subseteq V_1 \in \mathcal{T}} \sup_{U \subseteq V_2 \in \mathcal{T}} d_U((\mathcal{S}(U \subseteq V_1))a(V_1), \mathcal{S}(U \subseteq V_2)a(V_2)) \\ &\leq \sup_{U \subseteq V_1 \in \mathcal{T}} d_U((\mathcal{S}(U \subseteq V_1))a(V_1), a(U)) + \sup_{U \subseteq V_2 \in \mathcal{T}} d_U(a(U), \mathcal{S}(U \subseteq V_2)a(V_2)) \\ &\leq 2c_{\mathcal{S}}(a). \end{aligned}$$

5 Sheaves paired with assignments

This article discusses quite a few different categories, many of which are not extensively discussed in the literature. For convenience, all of the categories under discussion are listed below, with their descriptions (and mnemonic hints in boldface):

1. **Pseud**: the category of **pseudometric spaces**, Definition 5,
2. **Shv**: the category of all **sheaves**, Definition 8,
3. **ShvFP**: the category of all **sheaves** on **finite topological spaces** of **pseudometric spaces**, Definition 8,
4. **ShvA**: the category of all **sheaves** paired with **assignments**, Definition 9,
5. **ShvPA**(X, \mathcal{T}): the category of all **sheaves** of **pseudometric spaces** on a fixed topological space (X, \mathcal{T}) , each paired an **assignment**, Definition 9,
6. **ShvFPA**: the category of all **sheaves** on arbitrary **finite topological spaces** of **pseudometric spaces**, each paired an **assignment**, Definition 9,

7. **ShvFPA_L**: the subcategory of **ShvFPA** whose morphisms are sheaf morphisms along base space homeomorphisms whose component maps are Lipschitz, Theorem 2,
8. **PartCovers**: the category of **partial covers**, that consist of topological spaces with collections of open sets, Definition 10,
9. **CoarseFilt**: the category of **coarsening filtrations** of topological spaces, Definition 11, and
10. **PMod**: the category of **persistence modules**, Definition 12.

Definition 8. The category **Shv** is the category of *all sheaves*, which each object is a sheaf on some topological space. A morphism m in this category from a sheaf \mathcal{P} on (X, \mathcal{T}_X) to a sheaf \mathcal{Q} on (Y, \mathcal{T}_Y) consists of (1) a continuous map $f : (X, \mathcal{T}_X) \rightarrow (Y, \mathcal{T}_Y)$ and (2) a set of *component functions* m_U , one for each $U \in \mathcal{T}_Y$, such that

$$\begin{array}{ccc} \mathcal{P}(f^{-1}(U)) & \xrightarrow{m_U} & \mathcal{Q}(U) \\ \mathcal{P}(f^{-1}(U) \subseteq f^{-1}(V)) \uparrow & & \uparrow \mathcal{Q}(U \subseteq V) \\ \mathcal{P}(f^{-1}(V)) & \xrightarrow{m_V} & \mathcal{Q}(V) \end{array}$$

commutes for each pair of open sets $U \subseteq V$ in \mathcal{T}_Y .

Composition of morphisms in **Shv** is defined by composing both the continuous maps and the component functions. Explicitly, if $n : \mathcal{Q} \rightarrow \mathcal{R}$ is a morphism from the sheaf \mathcal{Q} defined above to a sheaf \mathcal{R} on (Z, \mathcal{T}_Z) along the continuous map $g : (Y, \mathcal{T}_Y) \rightarrow (Z, \mathcal{T}_Z)$ with component functions n_U for each $U \in \mathcal{T}_Z$, then the composition $n \circ m$ is given by the following commutative diagram

$$\begin{array}{ccccc} & & \xrightarrow{(nom)_U} & & \\ & \mathcal{P}((g \circ f)^{-1}(U)) & \xrightarrow{m_{g^{-1}(U)}} & \mathcal{Q}(g^{-1}(U)) & \xrightarrow{n_U} & \mathcal{R}(U) \\ & \uparrow \mathcal{P}((g \circ f)^{-1}(U) \subseteq (g \circ f)^{-1}(V)) & & \uparrow \mathcal{Q}(g^{-1}(U) \subseteq g^{-1}(V)) & & \uparrow \mathcal{R}(U \subseteq V) \\ & \mathcal{P}((g \circ f)^{-1}(V)) & \xrightarrow{m_{g^{-1}(V)}} & \mathcal{Q}(g^{-1}(V)) & \xrightarrow{n_V} & \mathcal{R}(V) \\ & & \xrightarrow{(nom)_V} & & & \end{array}$$

The subcategory **ShvFP** is the category of sheaves of pseudometric spaces on finite topological spaces. Because of Definition 6, a sheaf of pseudometric spaces always has continuous restriction maps. We will additionally assume that the component functions of every morphism of **ShvFP** are also continuous, and will emphasize this by calling them *component maps*. Under this convention, **ShvFP** is not a full subcategory of **Shv**.

Remark 4. The category **Shv** is a generalization of the more traditional **Shv**(X, \mathcal{T}), which is the category of sheaves on a fixed topological space (X, \mathcal{T}) in which each object is a sheaf on (X, \mathcal{T}) and all morphisms are along the identity map $X \rightarrow X$.

Definition 9. Each object of the category **ShvA** of *sheaf assignments* consists of a sheaf \mathcal{S} on (X, \mathcal{T}) and *assignment* $a \in \prod_{U \in \mathcal{T}} \mathcal{S}(U)$ to that sheaf. In this category, morphisms consist of sheaf morphisms that are compatible with the assignment, in the following way. Suppose that a is an assignment to a sheaf \mathcal{S} on a space X and b is an assignment to a sheaf \mathcal{R} on a space Y . A sheaf morphism $m : \mathcal{S} \rightarrow \mathcal{R}$ along a continuous map $f : X \rightarrow Y$ is morphism $(\mathcal{S}, a) \rightarrow (\mathcal{R}, b)$ in **ShvA** if for each open $U \subseteq Y$, the component function

$$m_U : \mathcal{S}(f^{-1}(U)) \rightarrow \mathcal{R}(U)$$

satisfies

$$m_U(a(f^{-1}(U))) = b(U).$$

We will also make use of the following subcategories:

- The subcategory **ShvFPA** of all sheaves on finite topological spaces of pseudometric spaces, each paired an assignment, and
- The subcategory **ShvPA**(X, \mathcal{T}) of all sheaves of pseudometric spaces on a fixed topological space (X, \mathcal{T}) , each paired an assignment.

Proposition 2. *There is a forgetful functor $\mathbf{ShvFPA} \rightarrow \mathbf{ShvFP}$, given by merely forgetting the assignment.*

Proof. Observe that without the assignments, the morphisms of **ShvFPA** are merely sheaf morphisms. \square

Proposition 3. *There is a functor $\mathbf{ShvFP} \rightarrow \mathbf{Pseud}$, given by replacing a sheaf by its space of assignments using the assignment pseudometric given in Definition 7.*

This proposition is not true for sheaves over infinite spaces. Perhaps surprisingly, requiring compactness of the base spaces is not sufficient to resolve the difficulty because assignments can lack all internal consistency.

Proof. Sheaf morphisms were defined to support this! Specifically, suppose that $m : \mathcal{S} \rightarrow \mathcal{R}$ is a morphism of sheaves of pseudometric spaces along the continuous map $f : (X, \mathcal{T}_X) \rightarrow (Y, \mathcal{T}_Y)$ between topological spaces, so that it is a morphism in **ShvFP**. Define a function $Fm : \prod_{U \in \mathcal{T}_X} \mathcal{S}(U) \rightarrow \prod_{V \in \mathcal{T}_Y} \mathcal{R}(V)$ between spaces of assignments for the two sheaves by its action on an arbitrary assignment a to \mathcal{S}

$$(Fm(a))(V) = m_V(a(f^{-1}(V))),$$

for each $V \in \mathcal{T}_Y$, which should be read as defining an assignment to \mathcal{R} .

Fm is continuous. Let $\epsilon > 0$ and an assignment $a \in \prod_{U \in \mathcal{T}_X} \mathcal{S}(U)$ be given. Observe that since each component map m_V of the sheaf morphism is continuous according to the convention established in Definition 8 for each open $V \in \mathcal{T}_Y$, there is a $\delta_V > 0$ such that for every $x \in \mathcal{S}(f^{-1}(V))$ with

$$d_{f^{-1}(V)}(x, a(f^{-1}(V))) < \delta_V,$$

it follows that

$$d_V(m_V(x), Fm(a)_V) < \epsilon.$$

Since objects in **ShvFP** consist of sheaves over finite spaces, the infimum


$$\delta := \inf_{V \in \mathcal{T}_Y} \delta_V$$

is strictly positive. (If \mathcal{T}_Y is infinite, this may be false.) Thus if b is an assignment to \mathcal{S} with $D(a, b) < \delta$, then $D(Fm(a), Fm(b)) < \epsilon$. As an aside, observe that $D(a, b) < \delta$ is stronger than necessary, since this constrains all stalks of \mathcal{S} not just those that are preimages of \mathcal{T}_Y -open sets.

This defines a functor. Let $n : \mathcal{R} \rightarrow \mathcal{P}$ be another morphism in **ShvFP**, this one along a continuous map $g : (Y, \mathcal{T}_Y) \rightarrow (Z, \mathcal{T}_Z)$. This implies that for an open set $V \in \mathcal{T}_Z$, $g^{-1}(V)$ and $(g \circ f)^{-1}(V)$ are also both open. We then compute

$$\begin{aligned} (F(n \circ m)(a))(V) &= (n \circ m)_V(a((g \circ f)^{-1}(V))) \\ &= (n \circ m)_V(a(f^{-1} \circ g^{-1}(V))) \\ &= (n \circ m)_V(a(f^{-1}(g^{-1}(V)))) \\ &= n_V(m_{g^{-1}(V)}(a(f^{-1}(g^{-1}(V))))) \\ &= n_V((Fm(a))(g^{-1}(V))) \\ &= (Fn(Fm(a)))(V). \end{aligned}$$

\square

Figure 1: The subspace (X, d) of \mathbb{R}^2 in Example 1

Proposition 4. *In \mathbf{ShvA} , any morphism whose domain (\mathcal{S}, s) is a sheaf \mathcal{S} and a global section s of \mathcal{S} will have a global section as its codomain.*

Proof. This follows immediately from the fact that images of global sections through sheaf morphisms are global sections. \square

Global sections are sometimes — but not always — isomorphic objects in \mathbf{ShvFPA} as the following example shows.

Example 1. Consider the metric space (X, d) that is a subspace of \mathbb{R}^2 with the usual metric d , with

$$X := \{(x, 0) : x \in \mathbb{R}\} \cup \{(0, 1)\}$$

as shown in Figure 1. The action of the group of homeomorphisms on (X, d) is not transitive because every homeomorphism of (X, d) must leave the point $(0, 1)$ fixed.

Now consider the sheaf \mathcal{S} on the single-point space $\{*\}$ (with the only possible topology for that space) with $\mathcal{S}(\{*\}) = (X, d)$. Both the space of global sections of \mathcal{S} and the space of assignments of \mathcal{S} are therefore precisely the points of (X, d) . Any isomorphism $\mathcal{S} \rightarrow \mathcal{R}$ in \mathbf{Shv} or \mathbf{ShvPA} will induce a homeomorphism on (X, d) . Thus (\mathcal{S}, a) and (\mathcal{R}, b) are isomorphic if and only if a and b are elements of the same connected component of (X, d) .

Proposition 5. *For any topological space (X, \mathcal{T}) , each isomorphism class of $\mathbf{ShvPA}(X, \mathcal{T})$ can be given a topology that is a subspace of the product of the assignment pseudometric and copies of the topology of uniform convergence (one copy for each restriction).*

Recall that the *topology of uniform convergence* for two pseudometric spaces $(X, d_X), (Y, d_Y)$ provides a pseudometric D on $C(X, Y)$ given by

$$D(f, g) := \sup_{x \in X} d_Y(f(x), g(x)).$$

Proof. Fix an isomorphism class of $\mathbf{ShvPA}(X, \mathcal{T})$ by specifying a representative sheaf \mathcal{S} of pseudometric spaces and a representative assignment a to \mathcal{S} . Every sheaf isomorphic to \mathcal{S} has stalks that are isomorphic to those of \mathcal{S} but possibly different restriction maps. Therefore, the space of assignments for every sheaf isomorphic to \mathcal{S} is the same. Given these two facts, the sheaves isomorphic to \mathcal{S} are parameterized by appropriate choices of restriction maps from the product

$$\prod_{U \subseteq V \in \mathcal{T}} C(\mathcal{S}(V), \mathcal{S}(U)),$$

in which each factor is given the topology of uniform convergence.

Thus the isomorphism class of (\mathcal{S}, a) in $\mathbf{ShvPA}(X, \mathcal{T})$ is merely a subset of the product

$$\prod_{U \subseteq V \in \mathcal{T}} C(\mathcal{S}(V), \mathcal{S}(U)) \times \prod_{U \in \mathcal{T}} \mathcal{S}(U),$$

in which the second factor in the product has the topology induced by the assignment pseudometric. \square

Corollary 1. For any topological space (X, \mathcal{T}) , $\mathbf{ShvPA}(X, \mathcal{T})$ can be made into topological space by simply taking the disjoint union of all the isomorphism classes.

Observe that each isomorphism class in $\mathbf{ShvPA}(X, \mathcal{T})$ is actually a closed subspace of the product

$$\prod_{U \subseteq V \in \mathcal{T}} C(\mathcal{S}(V), \mathcal{S}(U)) \times \prod_{U \in \mathcal{T}} \mathcal{S}(U),$$

because the commutativity and gluing axioms will prohibit many choices of restriction maps. This means that the topology of $\mathbf{ShvPA}(X, \mathcal{T})$ will likely be quite complicated.

Theorem 1. If (X, \mathcal{T}) is a finite topological space, then consistency radius is a continuous function $\mathbf{ShvPA}(X, \mathcal{T}) \rightarrow \mathbb{R}$.

This function is continuous in the assignment metric for assignments to a single sheaf [31]. But it is also continuous in the compact open topology for sheaves for a fixed assignment. Finiteness of \mathcal{T} is not essential, but makes for considerably less delicate argumentation.

Proof. Without loss of generality, we restrict attention to an isomorphism class of $\mathbf{ShvPA}(X, \mathcal{T})$. From the proof of Proposition 5, the elements of this isomorphism class are a subspace of

$$\prod_{U \subseteq V \in \mathcal{T}} C(\mathcal{S}(V), \mathcal{S}(U)) \times \prod_{U \in \mathcal{T}} \mathcal{S}(U).$$

Before we proceed with the main argument, suppose that a is an assignment to a sheaf \mathcal{S} , that \mathcal{R} is a sheaf isomorphic to \mathcal{S} , and that b is an assignment to \mathcal{R} . Recalling that

$$\| \|x\| - \|y\| \| \leq \|x - y\|$$

for any norm $\| \cdot \|$, we have that

$$\begin{aligned} |c_{\mathcal{R}}(b) - c_{\mathcal{S}}(a)| &= \left| \sup_{U_1 \subseteq V_1 \in \mathcal{T}} d_{U_1} ((\mathcal{R}(U_1 \subseteq V_1)) b(V_1), b(U_1)) \right. \\ &\quad \left. - \sup_{U_2 \subseteq V_2 \in \mathcal{T}} d_{U_2} ((\mathcal{S}(U_2 \subseteq V_2)) a(V_2), a(U_2)) \right| \\ &\leq \sup_{U \subseteq V \in \mathcal{T}} |d_U ((\mathcal{R}(U \subseteq V)) b(V), b(U)) - d_U ((\mathcal{S}(U \subseteq V)) a(V), a(U))| \\ &\leq \sup_{U \subseteq V \in \mathcal{T}} |d_U ((\mathcal{R}(U \subseteq V)) b(V), b(U)) - d_U ((\mathcal{S}(U \subseteq V)) a(V), a(U))| \\ &\quad + d_U (\mathcal{S}(U \subseteq V) b(V), b(U)) - d_U (\mathcal{S}(U \subseteq V) b(V), b(U)) \\ &\quad + d_U (\mathcal{S}(U \subseteq V) a(V), b(U)) - d_U (\mathcal{S}(U \subseteq V) a(V), b(U))| \\ &\leq \sup_{U \in \mathcal{T}} d_U (a(U), b(U)) \\ &\quad + \sup_{U \subseteq V \in \mathcal{T}} d_U (\mathcal{S}(U \subseteq V) b(V), \mathcal{S}(U \subseteq V) a(V)) \\ &\quad + \sup_{U \subseteq V \in \mathcal{T}} d_U (\mathcal{R}(U \subseteq V) b(V), \mathcal{S}(U \subseteq V) b(V)). \end{aligned}$$

Returning to the task at hand, suppose that a is an assignment to a sheaf \mathcal{S} and $\epsilon > 0$ is given. We show that there is an open set $Q \subseteq \mathbf{ShvPA}(X, \mathcal{T})$ containing (\mathcal{S}, a) whose consistency radii all lie in the interval

$$(c_{\mathcal{S}}(a) - \epsilon, c_{\mathcal{S}}(a) + \epsilon) \subseteq \mathbb{R}.$$

To this end, we construct Q as the intersection of three open subsets Q_1, Q_2, Q_3 of $\mathbf{ShvPA}(X, \mathcal{T})$. The set Q_i consists of the (\mathcal{R}, b) for which the i -th term in the last expression above is bounded by $\epsilon/3$, as follows:

- The first term merely requires that the distance from a to b in the assignment metric be not more than $\epsilon/3$; hence b lies in an open subset of the space of assignments

$$Q_1 := \{(\mathcal{R}, b) : D(a, b) < \epsilon/3\}.$$

- Since the restriction maps of \mathcal{S} are continuous, bounding the second term by $\epsilon/3$ requires that b lie within the intersection of finitely many open sets in the space of assignments

$$Q_2 := \left\{ (\mathcal{R}, b) : b(V) \in (\mathcal{S}(U \subseteq V))^{-1} (B_{\epsilon/3}(\mathcal{S}(U \subseteq V)a(V))) \text{ for all } U \subseteq V \in \mathcal{T} \right\},$$

where $B_{\epsilon/3}(x)$ is the open ball of radius $\epsilon/3$ centered at x . (If the topology \mathcal{T} is not finite, this can fail since there may be infinitely open subsets U of a given open set V . All of these need to be intersected in the space of assignments for the factor corresponding to the stalk $\mathcal{S}(V)$.)

- Bounding the last term above by $\epsilon/3$ requires that $\mathcal{R}(U \subseteq V)$ lies in an open set in the topology of uniform convergence

$$Q_3 := \{ (\mathcal{R}, b) : \mathcal{R}(U \subseteq V) \in B_{\epsilon/3}(\mathcal{S}(U \subseteq V)) \text{ for all } U \subseteq V \in \mathcal{T} \}.$$

Thus, if $(\mathcal{R}, b) \in Q = Q_1 \cap Q_2 \cap Q_3$, then $|c_{\mathcal{S}}(a) - c_{\mathcal{R}}(b)| < \epsilon$, which completes the argument. \square

Remark 5. Assuming the topology \mathcal{T} on a space (X, \mathcal{T}) is finite, consistency radius factors into two continuous maps,

$$\mathbf{ShvPA}(X, \mathcal{T}) \xrightarrow{F} \mathbb{R}^n \xrightarrow{N} \mathbb{R}$$

in which n is the number of open set inclusions $U \subseteq V$ in the topology \mathcal{T} , the second map is the infinity norm $N(v) = \|v\|_{\infty}$ and the first is given by

$$F(\mathcal{S}, a) := (d_U(\mathcal{S}(U \subseteq V)a(V), a(U)))_{U \subseteq V \in \mathcal{T}}.$$

Because of the equivalence of topologies of all finite-dimensional normed spaces, consistency radius is still continuous if we instead were to define the consistency radius using another norm as the second map N . Of particular interest is the 2-norm, resulting in a different definition for the consistency radius

$$c_{\mathcal{S}}(a) := \sqrt{\sum_{U \subseteq V \in \mathcal{T}} d_U((\mathcal{S}(U \subseteq V))a(V), a(U))^2}. \tag{1}$$

Proving bounds (such as appear in the proofs of Theorem 1, Lemma 4, and Theorem 3) with this definition is somewhat more cumbersome, but it preserves the smooth structure if the sheaves take values in the category of Riemannian manifolds.

6 Coarsening filtrations

In this section, we establish properties collections of nested collections of open sets in a topological space. The most prominent of these is the *consistency filtration* defined in Definition 18.

Lemma 1. (Standard) *If \mathcal{V} and \mathcal{U} are open covers of a topological space (X, \mathcal{T}) , when \mathcal{V} refines \mathcal{U} this induces a homomorphism $\check{H}^{\bullet}(X; \mathcal{U}) \rightarrow \check{H}^{\bullet}(X; \mathcal{V})$ on Čech cohomology.*

It is worth noting that the homomorphism — as usually defined (see for instance [13]) — uses a *refinement function* $\tau : \mathcal{V} \rightarrow \mathcal{U}$ for which $V \subseteq \tau(V)$ for all $V \in \mathcal{V}$. The homomorphism on Čech cohomology apparently depends on such a τ , but in fact the homomorphism is independent of the refinement function.

Definition 10. Each object of the category **PartCovers** of *partially covered topological spaces* is a triple $(X, \mathcal{T}, \mathcal{U})$, where (X, \mathcal{T}) is a topological space and $\mathcal{U} \subseteq \mathcal{T}$ is a collection⁴ of open sets. Each morphism $(X, \mathcal{T}_X, \mathcal{V}) \rightarrow (X, \mathcal{T}_Y, \mathcal{U})$ of **PartCovers** is given by a continuous map $f : (X, \mathcal{T}_X) \rightarrow (Y, \mathcal{T}_Y)$ such that each $V \in \mathcal{V}$ is a subset of $f^{-1}(U)$ for some $U \in \mathcal{U}$. We will sometimes abuse notation and write $f : (X, \mathcal{T}_X, \mathcal{V}) \rightarrow (X, \mathcal{T}_Y, \mathcal{U})$ in this case. Composition of morphisms in **PartCovers** is given by composition of the underlying continuous maps.

⁴ \mathcal{U} may not be a cover of X , but covers a subspace. This is important for Definition 17.

Lemma 2. *Suppose that \mathcal{V} is a collection of open sets in a topological space (X, \mathcal{T}) and \mathcal{U} is a collection of open sets in a topological space (Y, \mathcal{S}) . A continuous map $f : (X, \mathcal{T}) \rightarrow (Y, \mathcal{S})$ induces a homomorphism $\check{H}^\bullet(Y \cap \bigcup \mathcal{U}; \mathcal{U}) \rightarrow \check{H}^\bullet(X \cap \bigcup \mathcal{V}; \mathcal{V})$ if for all $V \in \mathcal{V}$, there is a $U \in \mathcal{U}$ such that $V \subseteq f^{-1}(U)$.*

Thus, Čech cohomology is a contravariant functor from **PartCovers** to the category of abelian groups.

Proof. Since \mathcal{U} contains open sets and f is continuous, then

$$f^{-1}(\mathcal{U}) = \{f^{-1}(U) : U \in \mathcal{U}\}$$

is a collection of open sets in X . Thus there is a chain map $\check{C}^k(Y \cap \bigcup \mathcal{U}; \mathcal{U}) \rightarrow \check{C}^k(X \cap \bigcup f^{-1}(\mathcal{U}); f^{-1}(\mathcal{U}))$ for every k . Observe that \mathcal{V} refines $f^{-1}(\mathcal{U})$ by assumption, so Lemma 1 means that the desired homomorphism is the composition

$$\check{H}^\bullet(Y \cap \bigcup \mathcal{U}; \mathcal{U}) \rightarrow \check{H}^\bullet(X \cap \bigcup f^{-1}(\mathcal{U}); f^{-1}(\mathcal{U})) \rightarrow \check{H}^\bullet(X \cap \bigcup \mathcal{V}; \mathcal{V}).$$

□

Definition 11. The category of *coarsening filtrations* **CoarseFilt** describes a collection of refinements of partial covers on fixed topological spaces. Each object is a triple $(X, \mathcal{T}, \mathbf{V})$ where (X, \mathcal{T}) is a topological space and \mathbf{V} is a function

$$\mathbf{V} : \mathbb{R} \rightarrow 2^{\mathcal{T}},$$

for which $\mathbf{V}(t)$ is a collection of open sets in (X, \mathcal{T}) for all t and $\mathbf{V}(t)$ refines $\mathbf{V}(t')$ whenever $t \leq t'$. If $(X, \mathcal{T}, \mathbf{V})$ and $(Y, \mathcal{S}, \mathbf{U})$ are objects of **CoarseFilt**, a morphism $(X, \mathcal{T}, \mathbf{V}) \rightarrow (Y, \mathcal{S}, \mathbf{U})$ in **CoarseFilt** consists of an order preserving function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ and a family of continuous functions $f_t : (X, \mathcal{T}) \rightarrow (Y, \mathcal{S})$ such that for all t , $\mathbf{V}(s)$ refines $f_t^{-1}(\mathbf{U}(t))$ for all $s \in \phi^{-1}(t)$.

Composition of morphisms in **CoarseFilt** is given by composing the order-preserving functions and the continuous maps. Explicitly, if $\phi_1, f_\bullet : (X, \mathcal{T}, \mathbf{V}) \rightarrow (Y, \mathcal{S}, \mathbf{U})$ and $\phi_2, g_\bullet : (Y, \mathcal{S}, \mathbf{U}) \rightarrow (Z, \mathcal{R}, \mathbf{W})$ are two morphisms in **CoarseFilt**, then their composition is given by $\phi = \phi_2 \circ \phi_1$ (still order preserving) and $h_t = g_t \circ f_{\sup \phi_2^{-1}(t)}$.

To see that composition in **CoarseFilt** is well-defined, observe that if $r \in \phi^{-1}(t) = (\phi_2 \circ \phi_1)^{-1}(t)$, then $\phi_1(r) \in \phi_2^{-1}(t)$. Therefore, $\mathbf{U}(\phi_1(r))$ refines $g_t^{-1}(\mathbf{W}(t))$. Since $\mathbf{V}(r)$ refines $f_{\phi_1(r)}^{-1}(\mathbf{U}(\phi_1(r)))$, these two refinements together mean that $\mathbf{V}(r)$ refines $f_{\phi_1(r)}^{-1}(g_t^{-1}(\mathbf{W}(t)))$, which itself refines $h_t^{-1}(\mathbf{W}(t))$ because $\phi_1(r) \leq \sup \phi_2^{-1}(t)$.

Proposition 6. *There is a fully faithful functor taking **CoarseFilt** into the category of sheaves of **PartCovers** on the Alexandrov space (\mathbb{R}, \leq) .*

Proof. Given an object $(X, \mathcal{T}, \mathbf{V})$ of **CoarseFilt**, define a sheaf \mathcal{S} by its stalks

$$\mathcal{S}([t, \infty)) := (X, \mathcal{T}, \mathbf{V}(t))$$

for each $t \in \mathbb{R}$ (the smallest open set containing t in the Alexandrov topology for (\mathbb{R}, \leq) is $[t, \infty)$) and restrictions

$$\mathcal{S}([t', \infty) \subseteq [t, \infty)) := \text{id}_X$$

for each $t \leq t'$. The restrictions are morphisms in **PartCovers** because $\mathbf{V}(t)$ refines $\mathbf{V}(t')$. Evidently this assignment is one-to-one on objects.

Each **CoarseFilt** morphism $(X, \mathcal{T}, \mathbf{V}) \rightarrow (Y, \mathcal{S}, \mathbf{U})$ consisting of an order preserving function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ and a family of continuous functions $f_t : (X, \mathcal{T}) \rightarrow (Y, \mathcal{S})$ is transformed into a sheaf morphism according to the following procedure. Suppose that $t \leq t'$. Since order preserving maps are continuous in the Alexandrov topology, this means that $\phi^{-1}([t, \infty))$ is an Alexandrov-open set of the form

$$\phi^{-1}([t, \infty)) = [\inf \phi^{-1}(t), \infty).$$

Additionally, if $t \leq t'$ this implies that $\inf \phi^{-1}(t) \leq \inf \phi^{-1}(t')$. Thus any morphism between any two such sheaves along an order preserving map $\phi : \mathbb{R} \rightarrow \mathbb{R}$ consists of a commutative diagram

$$\begin{array}{ccc} (X, \mathcal{T}_X, \mathbf{V}(\inf \phi^{-1}(t))) & \xrightarrow{m_t} & (Y, \mathcal{T}_Y, \mathbf{U}(t)) \\ \text{id}_X \downarrow & & \downarrow \text{id}_Y \\ (X, \mathcal{T}_X, \mathbf{V}(\inf \phi^{-1}(t'))) & \xrightarrow{m_{t'}} & (Y, \mathcal{T}_Y, \mathbf{U}(t')) \end{array}$$

where m_t and $m_{t'}$ are morphisms in **PartCovers**. Morphisms in **PartCovers** consist of continuous maps $f_t, f_{t'} : (X, \mathcal{T}_X) \rightarrow (Y, \mathcal{T}_Y)$ such that

- For every $V \in \mathbf{V}(\inf \phi^{-1}(t))$, there is a $U \in \mathbf{U}(t)$ such that $V \subseteq f_t^{-1}(U)$ and
- For every $V' \in \mathbf{V}(\inf \phi^{-1}(t'))$, there is a $U' \in \mathbf{U}(t')$ such that $V' \subseteq f_{t'}^{-1}(U')$.

This is precisely — not more nor less than — what is defined by a morphism in **CoarseFilt**, so the functor we are constructing must be full and faithful.

Composition of morphisms in **CoarseFilt** is preserved by their transformation to sheaf morphisms as follows: given a pair of **CoarseFilt** morphisms

$$(X, \mathcal{T}_X, \mathbf{V}) \xrightarrow{\phi_1, f_\bullet} (Y, \mathcal{T}_Y, \mathbf{U}) \xrightarrow{\phi_2, g_\bullet} (Z, \mathcal{T}_Z, \mathbf{W}),$$

we have just constructed the following pair of morphisms in **PartCovers**

$$(X, \mathcal{T}_X, \mathbf{V}(\inf(\phi_1^{-1} \circ \phi_2^{-1})(t))) \xrightarrow{m_{1, \inf \phi_2^{-1}(t)}} (Y, \mathcal{T}_Y, \mathbf{U}(\inf \phi_2^{-1}(t))) \xrightarrow{m_{2, t}} (Z, \mathcal{T}_Z, \mathbf{W}(t))$$

for each t . On the other hand, the composition of the **CoarseFilt** morphisms is given by the order preserving function $\phi = \phi_2 \circ \phi_1$ and the family of continuous functions $h_t = g_t \circ f_{\sup \phi_2^{-1}(t)}$. We need to show that our construction produces the same composition in **PartCovers** as the composition from **CoarseFilt**.

If $r \in (\phi_1^{-1} \circ \phi_2^{-1})(t)$, the well-definedness of composition in **CoarseFilt** requires that $\mathbf{V}(r)$ refines $h_t^{-1}(\mathbf{W}(t))$. Since $\inf(\phi_1^{-1} \circ \phi_2^{-1})(t) \leq r$, it follows that $\mathbf{V}(\inf(\phi_1^{-1} \circ \phi_2^{-1})(t))$ refines $h_t^{-1}(\mathbf{W}(t))$ as well. More plainly, we have shown that for every $V \in \mathbf{V}(\inf(\phi_1^{-1} \circ \phi_2^{-1})(t))$, there is a $W \in \mathbf{W}(t)$ such that $V \subseteq h_t^{-1}(W)$, namely that this composition in **CoarseFilt** yields a composition in **PartCovers** for each t . □

Note that our definition of the morphisms in **CoarseFilt** is consistent with the definition in [12, Def. 2.12] of a ϕ -shifted persistence module morphism.

Definition 12. (equivalent to [12, Def. 2.12]) The category **PMod** of *shifted persistence modules* has persistence modules as its objects and persistence module morphisms shifted by order preserving functions as its morphisms. Each object of **PMod** is a functor from the poset category (\mathbb{R}, \leq) to the category of vector spaces. Explicitly, each object is a choice of a vector space $\mathbf{E}(t)$ for each real number t and a choice of linear map $f_{\mathbf{E}, t \leq t'} : \mathbf{E}(t) \rightarrow \mathbf{E}(t')$ for each $t \leq t'$ such that

$$f_{\mathbf{E}, t \leq t''} = f_{\mathbf{E}, t' \leq t''} \circ f_{\mathbf{E}, t \leq t'}$$

and

$$f_{\mathbf{E}, t \leq t} = \text{id}_{\mathbf{E}(t)}.$$

Each morphism $\mathbf{E} \rightarrow \mathbf{F}$ of **PMod** is an order preserving map $\phi : \mathbb{R} \rightarrow \mathbb{R}$ and a choice of a set of maps $g_t : \mathbf{E}(\inf \phi^{-1}(t)) \rightarrow \mathbf{F}(t)$ such that the following diagram commutes

$$\begin{array}{ccc} \mathbf{E}(\inf \phi^{-1}(t)) & \xrightarrow{f_{\mathbf{E}, \inf \phi^{-1}(t) \leq \inf \phi^{-1}(t')}} & \mathbf{E}(\inf \phi^{-1}(t')) \\ g_t \downarrow & & \downarrow g_{t'} \\ \mathbf{F}(t) & \xrightarrow{f_{\mathbf{F}, t \leq t'}} & \mathbf{F}(t') \end{array}$$

for all $t \leq t'$.

Given this algebraic structure, a compatible geometric structure is given by *interleaving*. For our purposes, we use a definition that is consistent with the definition of δ -matchings in [3] and (τ, σ) -interleavings in [12, Def. 2.13].

Definition 13. A pair of morphisms $f_t : (X, \mathcal{T}, \mathbf{V}) \rightarrow (Y, \mathcal{S}, \mathbf{U})$, $g_t : (Y, \mathcal{S}, \mathbf{U}) \rightarrow (X, \mathcal{T}, \mathbf{V})$ along ϕ, ψ (respectively) in **CoarseFilt** is an ϵ -interleaving if each of the following hold for all t :

1. $|\phi(t) - t| < \epsilon$,
2. $|\psi(t) - t| < \epsilon$,
3. $\mathbf{V}(\inf(\psi \circ \phi)^{-1}(t))$ refines $\mathbf{V}(t)$, and
4. $\mathbf{U}(\inf(\phi \circ \psi)^{-1}(t))$ refines $\mathbf{U}(t)$.

An ϵ -interleaving between two morphisms $f_t : \mathbf{E} \rightarrow \mathbf{F}$, $g_t : \mathbf{E} \rightarrow \mathbf{F}$ along ϕ, ψ (respectively) in **PMod** may be defined in exactly the same way.

Definition 14. (compare the definition in [3]) For each pair of objects $(X, \mathcal{T}, \mathbf{V})$, $(Y, \mathcal{S}, \mathbf{U})$ in **CoarseFilt**, the function

$$D((X, \mathcal{T}, \mathbf{V}), (Y, \mathcal{S}, \mathbf{U})) := \inf\{\epsilon : \text{there is an } \epsilon\text{-interleaving } (X, \mathcal{T}, \mathbf{V}) \rightarrow (Y, \mathcal{S}, \mathbf{U})\}$$

is a pseudometric, called the *interleaving distance*.

The same definition works for objects in **PMod**, *mutatis mutandis*.

Proposition 7. *Persistent Čech cohomology is a contravariant functor $PH : \mathbf{CoarseFilt} \rightarrow \mathbf{PMod}$.*

Proof. If $t < t'$, a morphism $(X, \mathcal{T}, \mathbf{V}) \rightarrow (Y, \mathcal{S}, \mathbf{U})$ in **CoarseFilt** along an order preserving ϕ and a family of continuous functions $f_t : (X, \mathcal{T}) \rightarrow (Y, \mathcal{S})$ induces the following commutative diagram on Čech cohomologies

$$\begin{array}{ccc} \check{H}^\bullet(X \cap \bigcup \mathbf{V}(\inf \phi^{-1}(t)); \mathbf{V}(\inf \phi^{-1}(t))) & \longleftarrow & \check{H}^\bullet(X \cap \bigcup \mathbf{V}(\inf \phi^{-1}(t')); \mathbf{V}(\inf \phi^{-1}(t'))) \\ \uparrow & & \uparrow \\ \check{H}^\bullet(Y \cap \bigcup \mathbf{U}(t); \mathbf{U}(t)) & \longleftarrow & \check{H}^\bullet(Y \cap \bigcup \mathbf{U}(t'); \mathbf{U}(t')) \end{array}$$

in which the horizontal homomorphisms arise from Lemma 1 and the vertical homomorphisms arise from Lemma 2. Each row can be interpreted as a persistence module and the vertical maps are morphisms in **PMod**, which completes the argument. \square

An immediate consequence of [12, Prop. 4.8] is that if two persistence modules are ϵ -interleaved (in the sense of Definition 13), then the bottleneck distance between them is bounded above by ϵ . Taken with the proof of Proposition 7, this implies the next statement.

Corollary 2. *An ϵ -interleaving between objects in **CoarseFilt**, induces an ϵ -interleaving between their persistent Čech cohomologies, so the bottleneck distance between their persistence diagrams is bounded above by ϵ .*

7 Consistency radius on subspaces

The consistency radius of an assignment is a global property. A large consistency radius may mean that there is widespread inconsistency among the values of an assignment, or the inconsistency may be localized to a small part of the base space. To discriminate between these two situations, it is useful to restrict the consistency radius to a subset of the base space.

Definition 15. If a is an assignment to a sheaf \mathcal{S} of pseudometric spaces on a topological space (X, \mathcal{T}) , and $U \in \mathcal{T}$ is open, then the *local consistency radius on U* is

$$c_{\mathcal{S}}(a, U) := c_{i_U^* \mathcal{S}}(i_U^* a),$$

where $i : U \rightarrow X$ is the inclusion.

Intuitively, $c_{\mathcal{S}}(a, U)$ is only sensitive to the part of the assignment that touches U .

Proposition 8. *If a is an assignment to a sheaf \mathcal{S} of pseudometric spaces on a topological space (X, \mathcal{T}) , and $U \in \mathcal{T}$ is open, then*

$$c_{\mathcal{S}}(a, U) = \sup_{V_1 \subseteq V_2 \subseteq U} d_{V_1} ((\mathcal{S}(V_1 \subseteq V_2))a(V_2), a(V_1)).$$

Proof. The left side is by definition

$$c_{i_U^* \mathcal{S}}(i_U^* a) = \sup_{V_1 \subseteq V_2 \in \mathcal{T}} d_{V_1} ((i_U^* \mathcal{S}(V_1 \subseteq V_2))(i_U^* a(V_2)), (i_U^* a(V_1))),$$

where $i : U \rightarrow X$ is the inclusion map. Once we notice that $i^{-1}(V) = V$ for each open $V \subseteq U$, the result follows from three facts about pullbacks:

1. $i_U^* \mathcal{S}(V) = \mathcal{S}(V)$ for any open $V \subseteq U$,
2. $i_U^* \mathcal{S}(V_1 \subseteq V_2) = \mathcal{S}(V_1 \subseteq V_2)$ for any $V_1 \subseteq V_2 \subseteq U$, and
3. for any $V \subseteq U$,

$$\begin{aligned} (i_U^*)_V &: i_U^* \mathcal{S}(i^{-1}(V)) \rightarrow \mathcal{S}(V) \\ &: i_U^* \mathcal{S}(V) \rightarrow \mathcal{S}(V) \\ &: \mathcal{S}(V) \rightarrow \mathcal{S}(V) \end{aligned}$$

is the identity map. □

The local consistency radius on an open set is indeed local, in the sense that replacing the open set with a larger one cannot decrease the local consistency radius.

Lemma 3. *If a is an assignment to a sheaf \mathcal{S} of pseudometric spaces and $U \subseteq V$ are open subsets of the base space, then*

$$c_{\mathcal{S}}(a, U) \leq c_{\mathcal{S}}(a, V).$$

Proof. This follows from the fact that the supremum in the expression for $c_{\mathcal{S}}(a, V)$ is over a strictly larger set than the supremum in the expression for $c_{\mathcal{S}}(a, U)$. □

Corollary 3. *If a is an assignment to a sheaf \mathcal{S} of pseudometric spaces and U, V are open subsets of the base space, then*

$$\max\{c_{\mathcal{S}}(a, U), c_{\mathcal{S}}(a, V)\} \leq c_{\mathcal{S}}(a, U \cup V).$$

Caution: Corollary 3 does not ensure equality since $a(U \cup V)$ has no particular relationship to $a(U)$ or to $a(V)$.

Corollary 4. *For an assignment a to a sheaf \mathcal{S} of pseudometric spaces, the mean (or max, min, or any monotonic function of) consistency radius of a collection of open sets is monotonic with refinement. That is, if \mathcal{U}, \mathcal{V} are collections of open sets in the base space of \mathcal{S} and \mathcal{U} is a refinement of \mathcal{V} , then the mean consistency radius for \mathcal{U} will be less than the mean consistency radius for \mathcal{V} .*

Lemma 3 and its Corollaries indicate that the notion of consistency radius can be extended to an assignment that is only partially specified.

Definition 16. If $\mathcal{U} \subseteq \mathcal{T}$ is a collection of open sets for a topological space (X, \mathcal{T}) and \mathcal{S} is a sheaf of pseudometric spaces on (X, \mathcal{T}) , then an *assignment supported on \mathcal{U}* is an element of $\prod_{U \in \mathcal{U}} \mathcal{S}(U)$. The consistency radius of an assignment a supported on \mathcal{U} is written $c_{\mathcal{S}}(a, \mathcal{U})$, and is the infimum of all consistency radii of assignments b that restrict to a , namely

$$c_{\mathcal{S}}(a, \mathcal{U}) := \inf \left\{ c_{\mathcal{S}}(b) : b \in \prod_{V \in \mathcal{T}} \mathcal{S}(V) \text{ such that } b(U) = a(U) \text{ whenever } U \in \mathcal{U} \right\}.$$

We say that each such assignment b *extends* a .

While this definition ensures that

$$c_{\mathcal{S}}(a) = c_{\mathcal{S}}(a, X) = c_{\mathcal{S}}(a, \mathcal{T}),$$

the reader is warned that if an assignment a is supported on a collection $\mathcal{U} = \{U\}$ consisting of a single open set, then $c_{\mathcal{S}}(a)$ and $c_{\mathcal{S}}(a, U)$ are both meaningless if the base space topology \mathcal{T} is not trivial. Nevertheless, the following Proposition is true.

Proposition 9. *If a is an assignment to a sheaf \mathcal{S} on a topological space (X, \mathcal{T}) and $U \in \mathcal{T}$, then*

$$c_{\mathcal{S}}(a, \mathcal{T} \cap U) \geq c_{\mathcal{S}}(a, U).$$

Proof. Suppose b is another assignment such that $b(V) = a(V)$ if $V \in \mathcal{T} \cap U$. Thus

$$\begin{aligned} c_{\mathcal{S}}(b) &= \sup_{V_1 \subseteq V_2 \in \mathcal{T}} d_{V_1}(\mathcal{S}(V_1 \subseteq V_2)b(V_2), b(V_1)) \\ &\geq \sup_{V_1 \subseteq V_2 \in \mathcal{T} \cap U} d_{V_1}(\mathcal{S}(V_1 \subseteq V_2)b(V_2), b(V_1)) \\ &\geq \sup_{V_1 \subseteq V_2 \in \mathcal{T} \cap U} d_{V_1}(\mathcal{S}(V_1 \subseteq V_2)a(V_2), a(V_1)) \\ &\geq \sup_{V_1 \subseteq V_2 \subseteq U} d_{V_1}(\mathcal{S}(V_1 \subseteq V_2)a(V_2), a(V_1)) \\ &\geq c_{\mathcal{S}}(a, U). \end{aligned}$$

Thus $c_{\mathcal{S}}(a, U)$ is a lower bound for the consistency radius of any extension b of a . Since $c_{\mathcal{S}}(a, \mathcal{T} \cap U)$ is the greatest lower bound of these, the result follows. \square


A nonzero consistency radius for an assignment that is supported on \mathcal{U} rather than the entire space is thus still the obstruction to extending that assignment to a global section. It is for this reason that we need only consider assignments supported on the entire topology. If we are given an assignment not supported on a given open set, a value can be supplied that minimizes the overall consistency radius. Consistency is not assured over the entire base space by this process, so only a subspace will typically be covered by sets whose consistency radius is small. The extent to which consistency is obtained is formalized by the following definition.

Definition 17. Let a be an assignment to a sheaf \mathcal{S} of pseudometric spaces on (X, \mathcal{T}) . A collection of open sets \mathcal{U} is an ϵ -consistent collection for a if for every $U \in \mathcal{U}$, the consistency radius on U is less than ϵ

$$c_{\mathcal{S}}(a, U) < \epsilon.$$

Generally, ϵ -consistent collections only cover part of the base space.

Example 2. Consider the topological space (X, \mathcal{T}) for the finite set $X := \{A, B, C\}$ in which $\mathcal{T} := \{\emptyset, \{A\}, \{A, B\}, \{A, C\}, \{A, B, C\}\}$. We can define a sheaf \mathcal{S} on (X, \mathcal{T}) according to



where the diagram on the left shows the open sets in \mathcal{T} and the diagram on the right shows the restrictions of \mathcal{S} . (Each restriction of \mathcal{S} is a homomorphism $\mathbb{R} \rightarrow \mathbb{R}$, which is given by multiplication by the listed factor.) Observe that any sheafification of the sheaf on the partial order for the base $\{\{A\}, \{A, B\}, \{A, C\}\}$ will be of this form, with the $r \neq 0$ being left as a free parameter.

Consider the assignment a supported on the two sets $\mathcal{U} := \{\{A, B\}, \{A, C\}\}$ on the middle row above, given by

$$a(\{A, B\}) := 0, \quad a(\{A, C\}) := 1.$$

If we wish to compute $c_{\mathcal{S}}(a, \mathcal{U})$, we need to extend to assignments over all open sets. Namely, if

$$a(\{A\}) := y, \quad a(\{A, B, C\}) := x,$$

we must solve the optimization problem

$$c_{\mathcal{S}}(a, \mathcal{U}) = \min_{x, y} \max\{|y|, |1 - y|, |rx - y|, |2rx|, |1 - rx|\}.$$

By inspection (or a short application of the simplex algorithm), we find that $y = 1/2$ and $rx = 1/3$. Thus the critical thresholds are independent of r (and hence really only depend on the sheaf on the partial order rather than the topological space) and are given by

$$c_{\mathcal{S}}(a, \mathcal{U}) = \min_{x, y \in \mathbb{R}} \max\{1/2, 1/2, 1/6, 2/3, 2/3\} = 2/3.$$

The consistency radius on each open set is shown in the left frame of Figure 2.

Although there is an $2/3$ -consistent collection that contains X , it happens that X is covered by a $1/2$ -consistent collection, namely \mathcal{U} itself. If $\epsilon < 1/2$, the only consistent open set is $\{A\}$, which clearly does not cover X .

The next Lemma provides a useful bound on the consistency radius of the image of an assignment through a morphism of **ShvPA**.

Lemma 4. *Suppose $m : (\mathcal{S}, a) \rightarrow (\mathcal{R}, b)$ is a morphism in **ShvPA** along $f : X \rightarrow Y$, for which all component maps are Lipschitz continuous with constant K . Then for every open $U \subseteq Y$,*

$$c_{\mathcal{R}}(b, U) \leq K c_{\mathcal{S}}(a, f^{-1}(U)).$$

Proof. This is merely calculation:

$$\begin{aligned} c_{\mathcal{R}}(b, U) &= \sup_{V_1 \subseteq V_2 \subseteq U} d_{V_1}((\mathcal{R}(V_1 \subseteq V_2))b(V_2), b(V_1)) \\ &= \sup_{V_1 \subseteq V_2 \subseteq U} d_{V_1}((\mathcal{R}(V_1 \subseteq V_2) \circ m_{V_2})a(f^{-1}(V_2)), m_{V_1}(a(f^{-1}(V_1)))) \\ &= \sup_{V_1 \subseteq V_2 \subseteq U} d_{V_1}((m_{V_1} \circ \mathcal{S}(f^{-1}(V_1) \subseteq f^{-1}(V_2)))a(f^{-1}(V_2)), m_{V_1}(a(f^{-1}(V_1)))) \\ &\leq K \sup_{f^{-1}(V_1) \subseteq f^{-1}(V_2) \subseteq f^{-1}(U)} d_{f^{-1}(V_1)}(\mathcal{S}(f^{-1}(V_1) \subseteq f^{-1}(V_2))a(f^{-1}(V_2)), a(f^{-1}(V_1))) \\ &\leq K c_{\mathcal{S}}(a, f^{-1}(U)). \end{aligned}$$

□

Lemma 4 is a generalization of Proposition 4, that morphisms in **ShvPA** preserve global sections, because a global section of \mathcal{S} has consistency radius zero.

8 Consistency filtrations

Given an assignment to a sheaf of pseudometric spaces, there can be many possible ϵ -consistent collections. Of these, there is a distinguished ϵ -consistent collection that consists of open sets that are as large as possible.

Lemma 5. *For every $\epsilon > 0$ and every assignment a to a sheaf \mathcal{S} of pseudometric spaces on a finite space, there is a unique ϵ -consistent collection $\mathcal{M}_{\mathcal{S}, a}(\epsilon)$ such that every other ϵ -consistent collection is a refinement of $\mathcal{M}_{\mathcal{S}, a}(\epsilon)$.*

This Lemma is not the same as the Theorem in [22], where an assignment is made only to vertices on an abstract simplicial complex.




Figure 2: The consistency filtration for the assignment in Example 2; see Example 3 for interpretation.

Proof. Lemma 3 implies that any refinement of an ϵ -consistent collection is itself ϵ -consistent. Because of the finiteness of the base space, if for some open U

$$c_{\mathcal{S}}(a, U) < \epsilon,$$

then there is a unique maximal open set U' containing U for which

$$c_{\mathcal{S}}(a, U') < \epsilon.$$

This argument applied to each open set completes the argument. \square

Definition 18. The *consistency filtration* $\mathbf{CF} : \mathbf{ShvFPA} \rightarrow \mathbf{CoarseFilt}$ associates the coarsening filtration

$$(\mathbf{CF}(\mathcal{S}, a))(\epsilon) := \mathcal{M}_{\mathcal{S}, a}(\epsilon)$$

to each sheaf assignment (\mathcal{S}, a) on a finite space, in which $\epsilon > 0$ selects the threshold for the ϵ -consistent collection $\mathcal{M}_{\mathcal{S}, a}(\epsilon)$.

In what follows, we will show that the consistency filtration is *both* a functor (Theorem 2) *and* a continuous function (Theorem 3), depending on the interpretation of **ShvFPA** and **CoarseFilt**. Specifically, although **ShvFPA** was defined as a category (Definition 9), it is also a topological space according to Corollary 1 (by taking the disjoint union over all topological spaces). Similarly, although **CoarseFilt** was defined as a category (Definition 11), the interleaving distance is a pseudometric for its set of objects (Corollary 2), making it into a pseudometric space.

Example 3. The consistency filtration for the assignment in Example 2 (after extending to an assignment supported on \mathcal{T}) is given in Figure 2. The left frame of Figure 2 shows the ranges of ϵ for which each open set is ϵ -consistent. The frame on the right of Figure 2 shows the consistency filtration, namely the coarsest ϵ -consistent collections. Observe that this is an element of **CoarseFilt** because $\{\{A\}\}$ refines $\{\{A, B\}, \{A, C\}\}$ which refines $\{\{A, B, C\}\}$.

This particular example has rather uninteresting persistent Čech cohomology, since $\check{H}^0 \cong \mathbb{R}$ and $\check{H}^1 = 0$ for all thresholds.

Theorem 2. Consider \mathbf{ShvFPA}_L , the subcategory of **ShvFPA** consisting of sheaf morphisms along homeomorphisms whose component maps are Lipschitz. Consistency filtration is a covariant functor $\mathbf{CF} : \mathbf{ShvFPA}_L \rightarrow \mathbf{CoarseFilt}$.

Proof. First, we show that $\mathbf{CF}(\mathcal{S}, a)$ is a coarsening filtration. To this end, we need only observe that $\mathcal{M}_{\mathcal{S}, a}(t)$ is a refinement of $\mathcal{M}_{\mathcal{S}, a}(t')$ if $t < t'$ by Lemma 3.

Suppose $m : (\mathcal{S}, a) \rightarrow (\mathcal{R}, b)$ is a morphism in **ShvFPA** along the continuous map $f : (X, \mathcal{T}_X) \rightarrow (Y, \mathcal{T}_Y)$. Since we assumed all component maps are Lipschitz and the topologies are finite, we may assume that the maximum Lipschitz constant of any component map is K . Define

$$\phi(t) := \begin{cases} t & \text{if } K < 1, \\ Kt & \text{otherwise.} \end{cases}$$

which is evidently order-preserving.

We need to argue that $\mathcal{M}_{\mathcal{S}, a}(\inf \phi^{-1}(t))$ refines $f^{-1}(\mathcal{M}_{\mathcal{R}, b}(t))$. Let $V \in \mathcal{M}_{\mathcal{S}, a}(\inf \phi^{-1}(t))$, which means that $c_{\mathcal{S}}(a, V) < \inf \phi^{-1}(t) = t/K$. (If $K < 1$, the bound is evidently t . For brevity

we assume $K \geq 1$ in what follows.) But, due to the maximality of $\mathcal{M}_{\mathcal{S},a}(\inf \phi^{-1}(t))$, any open set containing V will have consistency radius larger than t/K . Since $\text{star } f(V)$ is open and (Y, \mathcal{T}_Y) is a finite topological space,

$$f^{-1}(\text{star } f(V)) = f^{-1}(\cap \{U \in \mathcal{T}_Y : f(V) \subseteq U\})$$

is an open set containing V , which means that

$$c_{\mathcal{S}}(a, f^{-1}(\text{star } f(V))) \geq t/K \geq c_{\mathcal{S}}(a, V).$$

Because any \mathcal{T}_Y -open set containing $f(V)$ will contain $\text{star } f(V)$, this inequality means that $f(V)$ being contained in some $U \in \mathcal{M}_{\mathcal{R},b}(t)$ implies that $V \subseteq f^{-1}(U)$.

Using the assumption that f is a homeomorphism and that the Lipschitz constants of the component maps are less than or equal to 1, Lemma 4 implies

$$c_{\mathcal{R}}(b, f(V)) \leq K c_{\mathcal{S}}(a, V) \leq t.$$

By the maximality of $\mathcal{M}_{\mathcal{R},b}(t)$, this means that there is an open $U \in \mathcal{M}_{\mathcal{R},b}(t)$ that contains $f(V)$.

Finally, we show that **CF** preserves composition of morphisms. Start with two morphisms in **ShvFPA** $_L$, $m : (\mathcal{S}, a) \rightarrow (\mathcal{R}, b)$ and $n : (\mathcal{R}, b) \rightarrow (\mathcal{Q}, c)$ along $f : (X, \mathcal{T}_X) \rightarrow (Y, \mathcal{T}_Y)$ and $g : (Y, \mathcal{T}_Y) \rightarrow (Z, \mathcal{T}_Z)$, respectively. Using the above construction of **CF**(m), an **CoarseFilt** morphism for m , suppose that ϕ_m is the order preserving function constructed for m so that $\mathcal{M}_{\mathcal{S},a}(\inf \phi_m^{-1}(t))$ refines $f^{-1}(\mathcal{M}_{\mathcal{R},b}(t))$. Similarly, suppose that ϕ_n is the order preserving function constructed for **CF**(n) so that $\mathcal{M}_{\mathcal{R},b}(\inf \phi_n^{-1}(t))$ refines $g^{-1}(\mathcal{M}_{\mathcal{Q},c}(t))$. Notice that the Lipschitz constants of the component maps of m and n impact our construction of ϕ_m and ϕ_n , respectively. We need to use these data to show that $\mathcal{M}_{\mathcal{S},a}(\inf \phi_m^{-1}(\phi_n^{-1}(t)))$ refines $(g \circ f)^{-1}(\mathcal{M}_{\mathcal{Q},c}(t))$. Since f is continuous,

$$f^{-1}(\mathcal{M}_{\mathcal{R},b}(\inf \phi_n^{-1}(t))) \text{ refines } f^{-1}(g^{-1}(\mathcal{M}_{\mathcal{Q},c}(t))) = (g \circ f)^{-1}(\mathcal{M}_{\mathcal{Q},c}(t)).$$

We know that

$$\mathcal{M}_{\mathcal{S},a}(\inf \phi_m^{-1}(\inf \phi_n^{-1}(t))) \text{ refines } f^{-1}(\mathcal{M}_{\mathcal{R},b}(\inf \phi_n^{-1}(t))).$$

Since ϕ_m and ϕ_n are order preserving

$$\inf \phi_m^{-1}(\inf \phi_n^{-1}(t)) = \inf \phi_m^{-1}(\phi_n^{-1}(t)),$$

which means that

$$\mathcal{M}_{\mathcal{S},a}(\inf \phi_m^{-1}(\phi_n^{-1}(t))) \text{ refines } (g \circ f)^{-1}(\mathcal{M}_{\mathcal{Q},c}(t))$$

as desired. \square

Remark 6. **CF** is not a faithful functor. Global sections of a sheaf \mathcal{S} may not be isomorphic objects in **ShvFPA**, as Example 1 shows, but they all have exactly the same consistency filtration.

Recalling that we can interpret **ShvPA**(X, \mathcal{T}) and **CoarseFilt** not as categories but as topological spaces, the consistency filtration can be interpreted as a function.

Theorem 3. *If (X, \mathcal{T}) is a finite topological space, then **CF** is a continuous function **ShvPA**(X, \mathcal{T}) \rightarrow **CoarseFilt** under the interleaving distance.*

Using Corollary 2, this means that the transformation of isomorphism classes of **ShvPA**(X, \mathcal{T}) to persistence diagrams of Čech cohomology is continuous.

Proof. Since continuity is a local property, we work within each isomorphism class of **ShvPA**(X, \mathcal{T}) separately.

Let $\epsilon > 0$ and $(\mathcal{S}, a) \in \mathbf{ShvPA}(X, \mathcal{T})$ be given. We aim to show that there is an open neighborhood $Q \subseteq \mathbf{ShvPA}(X, \mathcal{T})$ containing (\mathcal{S}, a) such that the consistency filtration of any $(\mathcal{R}, b) \in Q$ is ϵ -interleaved with the consistency filtration of (\mathcal{S}, a) .

Let $U \in \mathcal{T}$ be an arbitrary open set. Since $c_{\mathcal{S}}(a, U) = c_{i_U^* \mathcal{S}}(i_U^* a)$, this means that the local consistency radius on U is a continuous function **ShvPA**($U, \mathcal{T} \cap U$) $\rightarrow \mathbb{R}$, where $\mathcal{T} \cap U$ is the

subspace topology of U . This means that there is an open $Q_U \subseteq \mathbf{ShvPA}(U, \mathcal{T} \cap U)$ containing (\mathcal{S}, a) such that for every $(\mathcal{R}, b) \in Q_U$, it follows that $|c_{\mathcal{R}}(b, U) - c_{\mathcal{S}}(a, U)| < \epsilon$. This inequality still holds upon extending Q_U to an open neighborhood Q'_U in $\mathbf{ShvPA}(X, \mathcal{T})$, since $c_{\mathcal{R}}(b, U)$ is independent of the sheaf and assignment outside U . The set Q'_U consists of those elements of

$$Q_U \times \left(\prod_{V' \subseteq V \in \mathcal{T}, V \cap U = \emptyset} C(\mathcal{S}(V), \mathcal{S}(V')) \times \prod_{V \in \mathcal{T}, V \cap U = \emptyset} \mathcal{S}(V) \right)$$

that actually correspond to sheaves in $\mathbf{ShvPA}(X, \mathcal{T})$. We are largely free to select spaces of local sections and restrictions outside of U subject to commutativity of the diagram of restrictions and the gluing axiom. Notice that we need only choose these for open sets V that are disjoint from U , since the gluing axiom mandates the rest. Since \mathcal{T} is finite,

$$Q := \bigcap_{U \in \mathcal{T}} Q'_U$$

contains (\mathcal{S}, a) and is still open in $\mathbf{ShvPA}(X, \mathcal{T})$. Thus for any $(\mathcal{R}, b) \in Q$, the consistency radius of every open set measured with (\mathcal{R}, b) differs from that measured by (\mathcal{S}, a) by no more than ϵ .

We now show that $\mathbf{CF}(\mathcal{S}, a)$ and $\mathbf{CF}(\mathcal{R}, b)$ are ϵ -interleaved for any $(\mathcal{R}, b) \in Q$. Specifically, we construct morphisms $f : \mathbf{CF}(\mathcal{S}, a) \rightarrow \mathbf{CF}(\mathcal{R}, b)$ along the monotonic function ϕ and $g : \mathbf{CF}(\mathcal{R}, b) \rightarrow \mathbf{CF}(\mathcal{S}, a)$ along ψ in $\mathbf{CoarseFilt}$ that are each others' inverses and $|\phi(t) - t| \leq \epsilon$ and $|\psi(t) - t| \leq \epsilon$. This is easily done; let

$$\phi(t) := \psi(t) = t + \epsilon,$$

and

$$f_t := g_t := \text{id}_X$$

for all $t \in \mathbb{R}$. Without loss of generality, it remains to show that these are actually morphisms in $\mathbf{CoarseFilt}$, namely that

$$(\mathbf{CF}(\mathcal{S}, a))(\inf \phi^{-1}(t)) = \mathcal{M}_{\mathcal{S}, a}(t - \epsilon)$$

refines

$$(\mathbf{CF}(\mathcal{R}, b))(t) = f_t^{-1}(\mathcal{M}_{\mathcal{R}, b}(t)) = \mathcal{M}_{\mathcal{R}, b}(t)$$

for all $t \in \mathbb{R}$. Suppose that $U \in \mathcal{M}_{\mathcal{S}, a}(t - \epsilon)$, so that

$$c_{\mathcal{S}}(a, U) < t - \epsilon.$$

This means that because $(\mathcal{R}, b) \in Q$,

$$\begin{aligned} c_{\mathcal{R}}(b, U) &= c_{\mathcal{R}}(b, U) + c_{\mathcal{S}}(a, U) - c_{\mathcal{S}}(a, U) \\ &\leq c_{\mathcal{S}}(a, U) + |c_{\mathcal{R}}(b, U) - c_{\mathcal{S}}(a, U)| \\ &< c_{\mathcal{S}}(a, U) + \epsilon \\ &< t - \epsilon + \epsilon = t \end{aligned}$$

which implies that U is an element (or a subset of some element) of $\mathcal{M}_{\mathcal{R}, b}(t)$. □

Example 4. Consider a subset of N points $X := \{x_1, \dots, x_N\} \subseteq \mathbb{R}^M$ in M -dimensional Euclidean space. We can realize X as a partial assignment x to the constant sheaf \mathcal{K} of \mathbb{R}^M on a topological space built from a single simplex. To this end, let $Y := 2^X$ be the power set of X , which is also an abstract simplicial complex. Given the partial order by inclusion \subseteq on Y , form the Alexandrov topology \mathcal{T} on Y . Specifically, since each set U that is a star over $\{x_1, \dots, x_k\}$ is open, let

$$\mathcal{K}(\text{star } \{x_1, \dots, x_k\}) := \mathbb{R}^M$$

and let every restriction map be the identity map. Let the assignment x be supported on vertices (only), with

$$x(\text{star } \{x_i\}) := x_i.$$

Selecting the extension a of x that minimizes consistency radius,

$$a = \operatorname{argmin} \left\{ c_{\mathcal{K}}(y) : y \in \prod_{U \in \mathcal{T}} \mathbb{R}^M \text{ with } y(\operatorname{star} \{x_i\}) = x_i \text{ for each } x_i \in X \right\}$$

yields an assignment whose value on any \mathcal{T} -open set V is the circumcenter of that set of points in X contained in V . This is because the extension a minimizes

$$\|a(\operatorname{star} \{x_1, \dots, x_k\}) - x_i\|$$

for all $i = 1, \dots, k$. Given this observation, a set $U := \operatorname{star} \{x_1, \dots, x_k\}$ is ϵ -consistent if the intersection

$$\bigcap_{i=1}^k B_{\epsilon}(x_i)$$

of radius ϵ balls centered at each point is nonempty. A maximal ϵ -consistent collection $\mathcal{M}_{\mathcal{K},a}(\epsilon)$ has cohomology isomorphic to the simplicial cohomology of the the radius ϵ Čech complex $C_{\epsilon}(X)$ of the point cloud, namely

$$\check{H}^k(X \cap \bigcup \mathcal{M}_{\mathcal{K},a}(\epsilon); \mathcal{M}_{\mathcal{K},a}(\epsilon)) \cong H^k(C_{\epsilon}(X)).$$

Thus the Čech cohomology of the consistency filtration is isomorphic to the persistent Čech cohomology of the point cloud.

9 Sheaves on partial orders

Much of the subtlety of the examples in this article arises from the fact that the value of an assignment on two open sets has no relationship to the value on their union. While this is not a problem in principle, it presents a practical issue since sheaves on partial orders are generally most common in applications. Sheaves on finite partial orders are convenient to specify since they merely record the stalks and restrictions on stars over points⁵ and have no requirement aside from commutativity of the resulting diagram. In particular, the gluing axiom is satisfied implicitly since the spaces of sections over unions of stars are not explicitly specified. Practically, this means that a typical assignment for a sheaf on a partial order will be supported on the stars only. As Example 2 shows, the consistency radius for this kind of assignment may be strictly larger than the supremum of restrictions *between stars*. To mitigate this difficulty, we define the *star consistency radius* for an open set.

Definition 19. For an assignment a supported on the stars of a sheaf \mathcal{S} of pseudometric spaces on an Alexandrov space X , the *star consistency radius* on an open subset $U \subseteq X$ is given by

$$c_{\mathcal{S}}^*(a, U) := \max \left\{ \sup_{y \in U} \sup_{x \in \operatorname{star} y} d_{\operatorname{star} x}((\mathcal{S}(\operatorname{star} x \subseteq \operatorname{star} y)) a(\operatorname{star} y), a(\operatorname{star} x)), \right. \\ \left. \sup_{y \in U} \sup_{x \in U} \sup_{z \in \operatorname{star} x \cap \operatorname{star} y} \frac{1}{2} d_{\operatorname{star} z}((\mathcal{S}(\operatorname{star} z \subseteq \operatorname{star} y)) a(\operatorname{star} y), (\mathcal{S}(\operatorname{star} z \subseteq \operatorname{star} x)) a(\operatorname{star} x)) \right\}$$

Lemma 6. *If a is an assignment supported on the stars of a sheaf \mathcal{S} of pseudometric spaces on an Alexandrov space X , then*

$$c_{\mathcal{S}}^*(a, U) \leq c_{\mathcal{S}}(a, U).$$

Proof. Evidently the first quantity taken in the maximum (for the definition of star consistency radius) is less than the consistency radius since the consistency radius takes the supremum over a larger set. The result follows from the calculation made in Remark 3 since the second quantity is a lower bound on the consistency diameter. \square

⁵Which is to say that their *stalks* have the traditional meaning of a *stalk of a sheaf* — a direct limit.

Without the second quantity in the maximum for the definition of star consistency of Lemma 6, the Lemma is clearly still true, but the bound is less tight.

Corollary 5. *If \mathcal{U} is an ϵ -consistent collection, then it is an ϵ -star consistent collection as well.*

Restricting our attention to sheaves on Alexandrov spaces and assignments supported on the set of all stars (and making appropriate changes to the arguments) does not impact the validity of essentially all of the results (Theorem 1, Lemma 3, Corollary 3, Corollary 4, Proposition 9, Lemma 4, Theorem 2, and Theorem 3) obtained earlier in this article, though the proofs are somewhat more tedious.

Acknowledgements

This article is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) and SPAWAR Systems Center Pacific (SSC Pacific) under Contract No. N66001-15-C-4040 and the Office of Naval Research (ONR) under Contract Nos. N00014-15-1-2090 and N00014-18-1-2541. Any opinions, findings and conclusions or recommendations expressed in this article are those of the author and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA), SPAWAR Systems Center Pacific (SSC Pacific), or the Office of Naval Research (ONR).

References

- [1] Samson Abramsky and Adam Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13(11):113036, 2011. URL <https://doi.org/10.1088/1367-2630/13/11/113036>.
- [2] Samson Abramsky, Rui Soares Barbosa, Kohei Kishida, Raymond Lal, and Shane Mansfield. Contextuality, cohomology and paradox, [arxiv:1502.03097](https://arxiv.org/abs/1502.03097). 2015. URL <https://dx.doi.org/10.4230/LIPIcs.CSL.2015.211>.
- [3] Ulrich Bauer and Michael Lesnick. Induced matchings of barcodes and the algebraic stability of persistence. In *Proceedings of the Thirtieth Annual Symposium Computational Geometry*, page 355, 2014. URL <https://doi.org/10.1145/2582112.2582168>.
- [4] J. Curry. *Sheaves, Cosheaves and Applications*, [arXiv:1303.3255](https://arxiv.org/abs/1303.3255). PhD thesis, University of Pennsylvania, 2013. URL <https://arxiv.org/abs/1303.3255>.
- [5] Justin Curry, Robert Ghrist, and Vidit Nanda. Discrete Morse theory for computing cellular sheaf cohomology. *Foundations of Computational Mathematics*, 16(4):875–897, 2016. URL <https://doi.org/10.1007/s10208-015-9266-8>.
- [6] V. de Silva and G. Carlsson. Topological estimation using witness complexes. In M. Alexa and S. Rusinkiewicz, editors, *Eurographics Symposium on Point-based Graphics*, 2004.
- [7] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002. URL <https://doi.org/10.1109/SFCS.2000.892133>.
- [8] R. Ghrist. Barcodes: the persistent topology of data. *Bulletin-American Mathematical Society*, 45(1):61, 2008. URL <https://doi.org/10.1090/S0273-0979-07-01191-3>.
- [9] Robert Ghrist and Yasuaki Hiraoka. Network codings and sheaf cohomology. In *NOLTA*, 2011.
- [10] Robert Ghrist and Sanjeevi Krishnan. A topological max-flow-min-cut theorem. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 815–818. IEEE, 2013. URL <https://doi.org/10.1109/GlobalSIP.2013.6737016>.
- [11] Joseph A Goguen. Sheaf semantics for concurrent interacting objects. *Mathematical Structures in Computer Science*, 2(02):159–191, 1992. URL <https://doi.org/10.1017/S0960129500001420>.
- [12] Shaun Harker, Miroslav Kramár, Rachel Levanger, and Konstantin Mischaikow. A comparison framework for interleaved persistence modules. *Journal of Applied and Computational Topology*, 3(1-2):85–118, 2019. URL <https://doi.org/10.1007/s41468-019-00026-x>.
- [13] John H. Hubbard. *Teichmüller Theory, volume 1*. Matrix Editions, 2006.

- [14] Cliff Joslyn. On beyond graphs: Multidimensional graph structures for the data sciences. In *GraphFest*, 2016.
- [15] Cliff Joslyn, Kathleen Nowak, Brenda Praggastis, Emilie Purvine, and Michael Robinson. Computations for local and pseudo sections in real-world sheaves. In *Joint Mathematics Meetings Special Session on Sheaves in Topological Data Analysis*, Atlanta, GA, January 2017.
- [16] J. Lilius. Sheaf semantics for Petri nets. Technical report, Helsinki University of Technology, Digital Systems Laboratory, 1993.
- [17] Grant Malcolm. Sheaves, objects, and distributed systems. *Electronic Notes in Theoretical Computer Science*, 225:3–19, 2009. URL <https://doi.org/10.1016/j.entcs.2008.12.063>.
- [18] Seyed Mansourbeigi. Sheaf theory approach to distributed applications: Analysing heterogeneous data in air traffic monitoring. *International Journal of Data Science and Analysis*, 3(5):34–39, 2017. URL <http://dx.doi.org/10.11648/j.ijdsa.20170305.11>.
- [19] Seyed MH Mansourbeigi. Sheaf theory as a mathematical foundation for distributed applications involving heterogeneous data sets. In *2018 32nd International Conference on Advanced Information Networking and Applications Workshops (WAINA)*, pages 28–33. IEEE, 2018. URL <https://doi.org/10.11648/j.ijdsa.20170305.11>.
- [20] Saigopal Nelaturi, Johan de Kleer, and Vadim Shapiro. Combinatorial models for heterogeneous system composition and analysis. In *SoSE*, pages 1–6, 2016. URL <https://doi.org/10.1109/SYSOSE.2016.7542952>.
- [21] Miradain Atontsa Nguemo and Calvin Tcheka. Sheaf cohomology on network codings: maxflow-mincut theorem. *Applied General Topology*, 18(2):219–230, 2017. URL <http://dx.doi.org/10.4995/agt.2017.3371>.
- [22] Brenda Praggastis. Maximal sections of sheaves of data over an abstract simplicial complex, [arxiv:1612.00397](https://arxiv.org/abs/1612.00397), 2016. URL <https://arxiv.org/abs/1612.00397>.
- [23] Michael Robinson. Asynchronous logic circuits and sheaf obstructions. *Electronic Notes in Theoretical Computer Science*, pages 159–177, 2012. URL <https://doi.org/10.1016/j.entcs.2012.05.010>.
- [24] Michael Robinson. Analyzing wireless communication network vulnerability with homological invariants. In *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2014. URL <https://doi.org/10.1109/GlobalSIP.2014.7032250>.
- [25] Michael Robinson. *Topological Signal Processing*. Springer, January 2014. URL <http://doi.org/10.1007/978-3-642-36104-3>.
- [26] Michael Robinson. Tracking before detection using partially ordered sets and optimization. Technical Report AU-CAS-MathStats Technical Report No. 2015-5, American University, 2015. URL <http://auislandora.wrlc.org/islandora/object/techreports:24>.
- [27] Michael Robinson. Pseudosections of sheaves with consistency structures. Technical Report AU-CAS-MathStats Technical Report No. 2015-2, American University, 2015. URL <http://auislandora.wrlc.org/islandora/object/techreports:19>.
- [28] Michael Robinson. Finding cross-species orthologs with local topology. In *ACM-Biocomputing and Bioinformatics Workshop on Topological Data Analysis in Biomedicine*, Seattle, WA, October 2016.
- [29] Michael Robinson. Imaging geometric graphs using internal measurements. *J. Differential Equations*, 260:872–896, 2016. URL <http://dx.doi.org/10.1016/j.jde.2015.09.014>.
- [30] Michael Robinson. Sheaf and duality methods for analyzing multi-model systems. In I. Pesenson, Q.T. Le Gia, A. Mayeli, H. Mhaskar, and D.-X. Zhou, editors, *Novel Methods in Harmonic Analysis (in press)*. Birkhäuser, 2017. URL https://doi.org/10.1007/978-3-319-55556-0_8.
- [31] Michael Robinson. Sheaves are the canonical datastructure for information integration. *Information Fusion*, 36:208–224, 2017. URL <http://dx.doi.org/10.1016/j.inffus.2016.12.002>.
- [32] J. Schürmann. *Topology of singular spaces and constructible sheaves*, volume 63. Birkhauser, 2003.
- [33] Allen Shepard. *A cellular description of the derived category of a stratified space*. PhD thesis, Brown University, 1985.

- [34] Wlodek Zadrozny and Luciana Garbayo. A sheaf model of contradictions and disagreements. preliminary report and discussion, [arxiv:1801.09036](https://arxiv.org/abs/1801.09036). 2018. URL <https://arxiv.org/abs/1801.09036>.
- [35] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete and Computational Geometry*, 33(2):249–274, 2005. URL <https://doi.org/10.1007/s00454-004-1146-y>.