

Simplicial Leray Cosheaf

Let $f : X \rightarrow Y$ be a continuous map.

Let \mathcal{U} be an open cover of $f(X) \subseteq Y$

Let $N_{\mathcal{U}}$ be nerve of corresponding to the cover \mathcal{U} .

For each integer $i \geq 0$ we have the **Leray simplicial cosheaf** over the nerve $N_{\mathcal{U}}$ via the assignment



$$\widehat{F}_i : \sigma \rightsquigarrow H_i(f^{-1}(U_{\sigma})).$$

If $\sigma \leq \tau$, the $U_{\tau} \subset U_{\sigma}$. Thus, we have induced inclusion:

$$i_* : H_i(f^{-1}(U_{\tau})) \rightarrow H_i(f^{-1}(U_{\sigma})).$$

Height Function on the Circle

Example: $f : S^1 \rightarrow \mathbb{R}$.



$$f(S^1) \subset V \cup W.$$

$$k \leftarrow k^2 \rightarrow k$$

Let v be the vertex representing V . Thus $U_v = V$.

Let w be the vertex representing W . Thus $U_w = W$.

Let e be the vertex representing $V \cap W$. Thus $U_e = V \cap W$.

$$H_0(f^{-1}(U_v)) = H_0(f^{-1}(V)) = k = H_0(f^{-1}(U_w)) = H_0(f^{-1}(W))$$

$$H_0(f^{-1}(U_e)) = H_0(f^{-1}(V \cap W)) = k^2$$

Height function on the Torus




Figure : Barcodes for Leray cosheaves coming from the height function on the torus.

$$\widehat{F}_1 : 0 \longleftarrow k_a \longrightarrow k_y^2 \longleftarrow k_b^2 \longrightarrow k_z^2 \longleftarrow k_c \longrightarrow 0$$

$$\widehat{F}_1: 0 \longleftarrow k_a \longrightarrow k_y^2 \longleftarrow k_b^2 \longrightarrow k_z^2 \longleftarrow k_c \longrightarrow 0$$

$$\text{Homology: } 0 \rightarrow \widehat{F}(a) \oplus \widehat{F}(b) \oplus \widehat{F}(c) \xrightarrow{\partial_1} 0 \oplus \widehat{F}(y) \oplus \widehat{F}(z) \oplus 0 \rightarrow 0$$

$$\text{Homology: } 0 \rightarrow k_a \oplus k_b^2 \oplus k_c \xrightarrow{\partial_1} 0 \oplus k_y^2 \oplus k_z^2 \oplus 0 \rightarrow 0$$

$$\partial_1 = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & -1 & -1 \end{bmatrix}$$

$$H_1(N_{\mathcal{U}}; \widehat{F}_1) = \left\langle \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \right\rangle \quad H_0(N_{\mathcal{U}}; \widehat{F}_1) \cong k$$

However, if we change our bases as follows

$$\begin{bmatrix} y'_1 = y_1 \\ y'_2 = y_1 + y_2 \end{bmatrix} \quad \begin{bmatrix} b'_1 = b_1 \\ b'_2 = b_1 + b_2 \end{bmatrix} \quad \begin{bmatrix} z'_1 = z_1 \\ z'_2 = z_1 + z_2 \end{bmatrix}$$

then our cosheaf \widehat{F}_1 can then be written as the direct sum of two interval modules:

$$0 \longleftarrow 0 \longrightarrow k_{y'_1} \longleftarrow k_{b'_1} \longrightarrow k_{z'_1} \longleftarrow 0 \longrightarrow 0$$

$$0 \longleftarrow k_a \longrightarrow k_{y'_2} \longleftarrow k_{b'_2} \longrightarrow k_{z'_2} \longleftarrow k_c \longrightarrow 0$$

This observation is, at the moment, a mere curiosity. However when wedded with the following classical theorem it provides a powerful result in homology:

Theorem (5.12)

Let $f : X \rightarrow Y$ be continuous. Assume a cover \mathcal{U} of the image $f(X) \subset Y$ whose nerve $N_{\mathcal{U}}$ is at most one-dimensional, i.e. the nerve has at most 1-simplices. For each $i \geq 0$, we have

$$H_i(X) \cong H_0(N_{\mathcal{U}}; \hat{F}_i) \oplus H_1(N_{\mathcal{U}}; \hat{F}_{i-1}).$$

The proof of this result is outside of the scope of this paper, but can be found in many references [?, ?, ?].

Suppose the simplicial complex K is linear: a graph where every vertex has degree at most two and contains no cycles.

For example, the simplicial leraf cosheaf where $f : X \rightarrow \mathbb{R}$ and \mathcal{U} is a “nice” cover of $f(X)$, i.e. at most double intersections.

$$K = N_{\mathcal{U}}$$

We can phrase these computations in terms of the barcode decomposition of a simplicial cosheaf over a linear complex:

$H_0(K; \widehat{F})$ counts **closed bars** and

$H_1(K; \widehat{F})$ counts **open bars**.

Recalling that the latter interval module is an open bar, we can read off the homology of the torus T by summing the vector spaces that lie in the same anti-diagonal slice, as described in Theorem 5.12.

$$\begin{array}{ccc}
 H_0(N_{\mathcal{U}}; \widehat{F}_1) = k & & H_1(N_{\mathcal{U}}; \widehat{F}_1) = k \\
 & \searrow \text{dotted line} & \\
 H_0(N_{\mathcal{U}}; \widehat{F}_0) = k & & H_1(N_{\mathcal{U}}; \widehat{F}_0) = k \\
 \\
 H_0(T) = k & & H_1(T) = k^2 & & H_2(T) = k
 \end{array}$$

One can also use Theorem 5.12 to obtain a non-obvious theorem in 1-D persistence: that level set persistence determines sub-level set persistence.

By making use of the above interpretation of barcodes and cosheaf homology, we illustrate how one can take the Leray cosheaves presented as a barcode and sweep from left to right to obtain the associated sub-level set persistence module (and its barcode in certain situations).

Stated formally, we have the following theorem.

Theorem

Suppose X is compact and $f : X \rightarrow Y \subset \mathbb{R}$ is continuous. Given a cover \mathcal{U} of the image with linear nerve and associated simplicial Leray cosheaves \widehat{F}_i , one can recover the sub-level set persistence module of f for any choice of $t_0 < \dots < t_n$ and integer $i \geq 0$ as follows:

- 1 For each t_j take the intersection of elements in \mathcal{U} with the interval $(-\infty, t_j]$ to form the restricted cosheaves $\widehat{F}_i|_{(-\infty, t_j]}$ and $\widehat{F}_{i-1}|_{(-\infty, t_j]}$.
- 2 The persistence module in degree i is then determined pointwise at t_j by

$$H_i(f^{-1}(-\infty, t_j]) \cong H_0(N_{\mathcal{U} \cap (-\infty, t_j]}; \widehat{F}_i) \oplus H_1(N_{\mathcal{U} \cap (-\infty, t_j]}; \widehat{F}_{i-1}).$$

One must first observe that Theorem 5.12 holds over the restriction.

$$\begin{array}{ccc}
 f^{-1}(-\infty, t_i] & \longrightarrow & X \\
 \downarrow & & \downarrow f \\
 (-\infty, t_i] & \longrightarrow & Y
 \end{array}$$

This proves that the i^{th} homology of the sub-level set can be computed via cosheaf homology. Now we must show that one can recover functoriality from the cosheaf perspective. If $\sigma \in N_{\mathcal{U}}$ is a simplex in the nerve and if $t < t'$, then there is a map

$$U_{\sigma} \cap (-\infty, t] \hookrightarrow U_{\sigma} \cap (-\infty, t'].$$

This implies that there is a map $\widehat{F}_i(U_{\sigma} \cap (-\infty, t]) \rightarrow \widehat{F}_i(U_{\sigma} \cap (-\infty, t'])$ and thus a map from chains valued in $\widehat{F}_i|_{(-\infty, t]}$ to chains valued in $\widehat{F}_i|_{(-\infty, t']}$. By functoriality of spectral sequences (maps of filtrations induce maps between spectral sequences) we get the desired map on homology. □

Sheaves as the Correct Foundation for Level Set Persistence

Level set persistence: take a cover \mathcal{U} of the image of $f : X \rightarrow Y$ and studying simplicial Leray cosheaves over the nerve $N_{\mathcal{U}}$.

Suppose we use a different cover \mathcal{U}' of the image.

Is there any way of comparing the Leray simplicial cosheaves over two different nerves?

Of course one could always refine the two covers \mathcal{U} and \mathcal{U}' to a common cover, but it would be convenient for proving theorems to work with all open sets at once.

This leads to the general notion of a cosheaf, which is the dual notion of a sheaf.

Definition (Pre-Cosheaves)

Any functor

$$\widehat{F} : \mathbf{Open}(X) \rightarrow \mathbf{D}$$

is called a **pre-cosheaf** valued in \mathbf{D} .

If $V \subset U$, then we usually write the **extension maps** of a cosheaf as $r_{U,V}^{\widehat{F}} : \widehat{F}(V) \rightarrow \widehat{F}(U)$. Often we omit the superscript F or \widehat{F} .

Definition (Pre-Sheaves)

Any functor

$$F : \mathbf{Open}(X)^{op} \rightarrow \mathbf{D}$$

is called a **pre-sheaf** valued in \mathbf{D} .

If $V \subset U$, then we usually write the **restriction maps** of a sheaf as $\rho_{V,U}^F : F(U) \rightarrow F(V)$. Often we omit the superscript F or \widehat{F} .

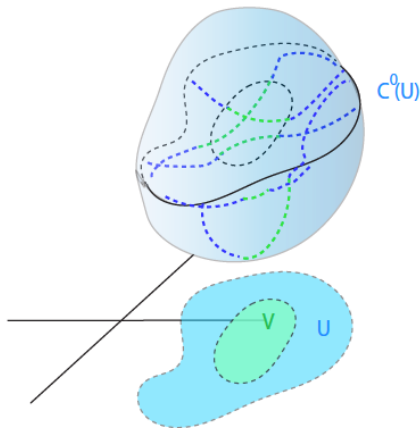


Figure 9.1: The elevated blue balloon is schematic representation of a presheaf of real valued functions over the open set $U \subset \mathbf{R}^2$. Each “function” is represented as blue and green dotted lines, where the green dash is the restriction of the function on V .

From: **A Gentle Introduction to Homology, Cohomology, and Sheaf Cohomology**
Jean Gallier and Jocelyn Quaintance, Book in Progress (2016)
<http://www.cis.upenn.edu/~jean/gbooks/sheaf-coho.html>

Definition (Leray Pre-Cosheaf)

Given a continuous map $f : X \rightarrow Y$ and an integer $i \geq 0$, one has the **Leray pre-cosheaf**:

$$\widehat{P}_i : U \subset Y \rightsquigarrow H_i(f^{-1}(U))$$

Dually, one has the **Leray pre-sheaf**:

$$P^i : U \subset Y \rightsquigarrow H^i(f^{-1}(U))$$




Figure : Visualizing the Leray pre-cosheaf H_0 for the height function on the circle.

Example (Height Function on the Circle)

Let $f : S^1 \rightarrow \mathbb{R}$ be the function that projects

$C = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$ onto the x -axis. For each open set U in \mathbb{R} , \widehat{P}_i assigns the i^{th} homology group $H_i(f^{-1}(U))$ to U .

Let us restrict our functor to the category of bounded open intervals $\mathbf{Int}(\mathbb{R})$, since they generate all of $\mathbf{Open}(\mathbb{R})$. Note that $\mathbf{Int}(\mathbb{R})$ can be visualized as the upper half-plane

$\mathbb{H}_+ = \{(m, r) \mid m \in \mathbb{R}, r > 0\}$ by letting each point (m, r) represent the midpoint and radius of an interval $I \subset \mathbb{R}$:

$$m(I) = \frac{x + y}{2} \quad r(I) = \frac{y - x}{2}$$

The partial order $I \leq J \Leftrightarrow I \subseteq J$ is then equivalent to the partial order on \mathbb{H}_+ where $(m, r) \leq (m', r')$ if and only if $|m' - m| \leq r' - r$. Thus, for maps to the real line, the Leray pre-cosheaf \widehat{P}_i assigns to each point in the upper-half plane the vector space $H_i(f^{-1}(I))$, and to each pair of inclusions $I \leq J$ the map $H_i(f^{-1}(I)) \rightarrow H_i(f^{-1}(J))$. For $i = 0$ and the height function on the circle, this assignment is depicted in Figure 3.


Obtaining Fibers via Stalks

One apparent disadvantage that Leray pre-cosheaves have is the restriction to open sets U prohibits directly recording the homology of the fiber $f^{-1}(y)$. However, there is a categorical construction that can be used in some cases to derive $H_i(f^{-1}(y))$ from the homology groups $H_i(f^{-1}(U))$. Moreover, this construction will work even better when we dualize to cohomology, which motivates the use of Leray pre-sheaves.

Definition (Limit)

The **limit** of a functor $F : \mathbf{I} \rightarrow \mathbf{C}$ is an object $\varprojlim F \in \mathbf{C}$ along with a collection of morphisms $\psi_x : \varprojlim F \rightarrow F(x)$ that commute with arrows in the diagram of F , i.e. if $g : x \rightarrow y$ is a morphism in \mathbf{I} , then $\psi_y = F(g) \circ \psi_x$ in \mathbf{C} .

We require that the limit is universal in the following sense: if there is another object L' and morphisms ψ'_x that also commute with arrows in F , then there is a unique morphism $u : L' \rightarrow \varprojlim F$ that commutes with everything in sight, i.e. $\psi'_x = \psi_x \circ u$ for all objects x in \mathbf{I} .



Example

Let \mathbf{I} be the category of open sets U that contain a point $y \in Y$ with morphisms corresponding to inclusions, which we call $\mathbf{Open}(Y)_y$. The limit of the restricted functor


$\widehat{P}_i : \mathbf{Open}(Y)_y \rightarrow \mathbf{Vect}$ is called the **costalk** of \widehat{P}_i at y .

Unfortunately, for a general continuous map it is unknown how the costalk at y is related to the homology of the fiber $f^{-1}(y)$.

The technical reason for this is that limits and homology do not commute [?, Prop. 2.5.19]. This is one traditional reason why many mathematicians prefer pre-sheaves over pre-cosheaves.

Definition (Colimit)

The **colimit** of a functor $F : \mathbf{I} \rightarrow \mathbf{C}$ is defined in a dual manner.



Example (Stalk)

Given a pre-sheaf $F : \mathbf{Open}(Y)^{op} \rightarrow \mathbf{Vect}$ and a point $y \in Y$ the **stalk** at y is defined to be the colimit of F over open sets containing y :

$$F_y := \varinjlim_{U \ni y} F(U)$$

In contrast to the Leray pre-cosheaves, the Leray pre-sheaves are traditionally considered better behaved by the following theorem.

Theorem (Thm. 6.2 [?])

Suppose $f : X \rightarrow Y$ is a proper map between locally compact spaces. For any point $y \in Y$ we have

$$P_y^i \cong H^i(f^{-1}(y)).$$

Proof.

The bulk of the proof appears in Theorem 6.2 of [?, pp. 176-7] where it is proved for the sheafification of P^i , which we will describe shortly. One can then observe that sheafification preserves stalks to get the desired result. □

Local to Global Properties of the (Co)Sheaf Axiom

If a topological space is equipped with a cover $\mathcal{U} = \{U_i\}_{i \in I}$ and a pre-cosheaf \widehat{F} , then we can define a simplicial cosheaf over $N_{\mathcal{U}}$ by restricting the assignment of \widehat{F} to only those open sets (and their intersections) appearing in \mathcal{U} :

$$\widehat{F} : \sigma \rightsquigarrow \widehat{F}(U_\sigma)$$

One can then compute simplicial cosheaf homology of \widehat{F} on this cover, which is also called the

Cech homology of \widehat{F} :

$$H_0(N_{\mathcal{U}}; \widehat{F}) \quad H_1(N_{\mathcal{U}}; \widehat{F}) \quad H_2(N_{\mathcal{U}}; \widehat{F}) \quad \dots$$

The first term $H_0(N_{\mathcal{U}}; \widehat{F})$ is used to define the cosheaf axiom, and its mirror term $H^0(N_{\mathcal{U}}; F)$ is used to define the sheaf axiom.

Definition

A pre-cosheaf \widehat{F} of vector spaces is a **cosheaf** if for every open set U and every cover \mathcal{U} of U

$$\widehat{F}(U) \cong H_0(N_{\mathcal{U}}; \widehat{F}).$$

Dually, a pre-sheaf F of vector spaces is a **sheaf** if for every open set U and every cover \mathcal{U} of U

$$F(U) \cong H^0(N_{\mathcal{U}}; F).$$

Remark (Local to Global)

It is often said that sheaves mediate the passage from local to global. This means that the value of $F(U)$ (the global datum) is completely determined by the values of $\{F(U_i)\}$ (the local data) where $\mathcal{U} = \{U_i\}$ is a cover of U . This perspective has powerful implications for parallel processing; in essence, the (co)sheaf axiom is a distributed algorithm.

The first observation one can make about the cosheaf axiom is that if $U = U_1 \cup U_2$ where $U_1 \cap U_2 = \emptyset$ and \widehat{F} is a cosheaf, then $\widehat{F}(U) \cong \widehat{F}(U_1) \oplus \widehat{F}(U_2)$. Many pre-cosheaves satisfy this property without being cosheaves themselves. For example, each of the Leray pre-cosheaves \widehat{P}_i satisfy this property without being cosheaves themselves.




Figure : The two Leray pre-cosheaves \widehat{P}_0 and \widehat{P}_1 for the height function on the circle. The figure on the right is an example of a pre-cosheaf that is not a cosheaf.

Example (\widehat{P}_1 is not a cosheaf)

In Figure 4 we consider side-by-side the two non-zero Leray pre-cosheaves associated to the height function on the circle $f : S^1 \rightarrow \mathbb{R}$. The pre-cosheaf \widehat{P}_1 fails to be a cosheaf because if one takes any cover $\mathcal{U} = \{U_i\}$ of $f(S^1)$ by open sets where no single open set contains the entire image, then the pre-cosheaf \widehat{P}_1 restricts to a collection of zero vector spaces and zero maps over the nerve $N_{\mathcal{U}}$. One immediately has that

$$\widehat{P}_1(\cup U_i) = k \neq H_0(N_{\mathcal{U}}; \widehat{P}_1) = 0,$$

which is required in order for \widehat{P}_1 to be a cosheaf. On the other hand, \widehat{P}_0 is always a cosheaf.

Example (\widehat{P}_0 is a cosheaf)

Suppose $f : X \rightarrow Y$ is a continuous map. The Leray pre-cosheaf $\widehat{P}_0 : U \rightsquigarrow H_0(f^{-1}(U))$ is a cosheaf. To see why, let $W = U \cup V$. By continuity of the map f and the Mayer-Vietoris long-exact sequence in homology, we have the exact sequence (meaning the kernel of one map is the image of the previous) of vector spaces

$$H_0(f^{-1}(U \cap V)) \rightarrow H_0(f^{-1}(U)) \oplus H_0(f^{-1}(V)) \rightarrow H_0(f^{-1}(W)) \rightarrow 0.$$

The first two terms are exactly the terms one writes down for computing

Cech homology of \widehat{P}_0 over the cover $\{U, V\}$, i.e.

$$\widehat{P}_0(U \cap V) \rightarrow \widehat{P}_0(U) \oplus \widehat{P}_0(V).$$

The cokernel of this map is precisely the

Cech homology of \widehat{P}_0 over $\{U, V\}$. The final two terms in the last row of the Mayer-Vietoris long exact sequence says precisely that $\widehat{P}_0(W)$ is isomorphic to this cokernel, i.e.

Sheafification and the Leray Sheaf

Both sheaves and cosheaves have the local-to-global properties described above and so either one should be preferred over their “pre”-cousins. Fortunately, there is a well understood procedure for turning any pre-sheaf into a sheaf called **sheafification**. It is a cruel asymmetry that there is not a similarly nice procedure for turning any pre-cosheaf into a cosheaf [?, Sec. 2.5.4].

Definition (Sheafification)

Let $F : \mathbf{Open}(X)^{op} \rightarrow \mathbf{Vect}$ be a pre-sheaf. The **sheafification** \tilde{F} of F assigns to every open set U the set of functions $s : U \rightarrow \sqcup_{x \in U} F_x$ that locally extend, i.e. for every $x \in U$ and $s(x) \in F_x$ there exists a $V \ni x$ with $V \subset U$ and a $t \in F(V)$ such that the image of $t \in F(V)$ in F_y agrees with $s(y)$ for all $y \in V$.

Definition (Leray Sheaves)

Suppose $f : X \rightarrow Y$ is a continuous map, then the i^{th} **Leray sheaf** F^i is the sheafification of the Leray pre-sheaf P^i associated to f .

The assertion of this paper is that the Leray sheaves are the proper object of study for understanding the level set persistence of a proper continuous map $f : X \rightarrow Y$. Unfortunately, the Leray sheaves are uncomputable in practice and are primarily good for proving theoretical results. In principle the cosheafification of the Leray pre-cosheaves \widehat{P}_i would be preferred, but there is no known cosheafification procedure.

Level Set Persistence for Definable Maps

In this section we restrict ourselves to a suitably tame class of maps and spaces so that most of the technical discrepancies between pre-sheaves and pre-cosheaves disappear. This class of maps and spaces is defined in terms of finitely many logical operations and includes most applications of interest, most notably point cloud persistence. Finally, we present the culmination of this paper: a collection of functors that can be reliably called the i^{th} level set persistence of a tame map.

An **o-minimal structure** on \mathbb{R} is a sequence of sets $\mathcal{O} = \{\mathcal{O}_n\}_{n \geq 0}$ satisfying

- 1 \mathcal{O}_n is a boolean algebra of subsets of \mathbb{R}^n , i.e. it is a collection of subsets of \mathbb{R}^n closed under unions and complements, with $\emptyset \in \mathcal{O}_n$;
- 2 If $A \in \mathcal{O}_n$, then $A \times \mathbb{R}$ and $\mathbb{R} \times A$ are both in \mathcal{O}_{n+1} ;
- 3 The sets $\{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_i = x_j\}$ for varying $i \leq j$ are in \mathcal{O}_n ;
- 4 If $A \in \mathcal{O}_{n+1}$ then $\pi(A) \in \mathcal{O}_n$ where $\pi : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$ is projection onto the first n factors;
- 5 For each $x \in \mathbb{R}$ we require $\{x\} \in \mathcal{O}_1$ and $\{(x, y) \in \mathbb{R}^2 \mid x < y\} \in \mathcal{O}_2$;
- 6 The only sets in \mathcal{O}_1 are the finite unions of open intervals and points.

When working with a fixed o-minimal structure, we say a set is **definable** if it belongs to some \mathcal{O}_n . A map is definable if its graph, viewed as a subset of the product, is definable.

The prototypical o-minimal structure is the class of semi-algebraic sets, which has become increasingly relevant in applied mathematics.

Definition

A **semi-algebraic** subset of \mathbb{R}^n is a subset of the form

$$X = \bigcup_{i=1}^p \bigcap_{j=1}^q X_{ij}$$

where the sets X_{ij} are of the form $\{f_{ij}(x) = 0\}$ or $\{f_{ij} > 0\}$ with f_{ij} a polynomial in n variables.

Proposition (Semi-algebraic Sets are Definable)

The collection of semi-algebraic subsets in \mathbb{R}^n for all $n \geq 0$ defines an o-minimal structure on \mathbb{R} .

Proof.

The only semi-algebraic subsets of \mathbb{R} are finite unions of points and open intervals. From the definition, one sees that the class of semi-algebraic sets is closed under finite unions and complements. The **Tarski-Seidenberg** theorem states that the projection onto the first m factors $\mathbb{R}^{m+n} \rightarrow \mathbb{R}^m$ sends semi-algebraic subsets to semi-algebraic subsets [?]. We can deduce from this theorem all of the conditions of o-minimality. \square

Semi-algebraic maps are defined to be those maps $f : \mathbb{R}^k \rightarrow \mathbb{R}^n$ whose graphs are semi-algebraic subsets of the product. The next example shows that the collection of augmented point clouds can be regarded as the fibers of a semi-algebraic map.

Example (Point-Cloud Data)

Suppose Z is a finite set of points in \mathbb{R}^n . For each $z \in Z$, consider the square of the distance function

$$d_z(x_1, \dots, x_n) = \sum_{i=1}^n (x_i - z_i)^2.$$

By the previously stated facts we know that the sets

$$B_z := \{x \in \mathbb{R}^{n+1} \mid d_z(x_1, \dots, x_n) \leq x_{n+1}^2\}$$

are semi-algebraic along with their unions and intersections. Denote by X the union of the B_z . The Tarski-Seidenberg theorem implies that the map

One of the nice features of a point cloud is that the topology of the union $X_r = \cup_{x_i \in Z} B(x_i, r)$ only changes for finitely many values of r . This behavior is common among all definable sets and maps.

Definition

A definable map $f : E \rightarrow B$ between definable sets is said to be **definably trivial** if there is a definable set F and a definable homeomorphism $h : E \rightarrow B \times F$ such that the diagram

$$\begin{array}{ccc} E & \xrightarrow{h} & B \times F \\ & \searrow f & \swarrow \pi \\ & B & \end{array}$$

commutes, i.e. $\pi \circ h = f$.

Remark

A definably trivial map is simple because the topology of the fiber $f^{-1}(b) \cong F$ does not change. In particular, there is a neighborhood U of b for which $H_i(f^{-1}(U)) \cong H_i(f^{-1}(b))$, so that the costalk of the Leray pre-cosheaf agrees with the homology of the fiber. In short, there is no advantage to studying the Leray pre-sheaves over the Leray pre-cosheaves for definably trivial maps.

Theorem (Trivialization Theorem [?])

Let $f : E \rightarrow B$ be a definable continuous map between definable sets E and B . Then B can be partitioned into definable sets B_1, \dots, B_k so that the restrictions

$$f|_{f^{-1}(B_i)} : f^{-1}(B_i) \rightarrow B_i$$

are definably trivial.




Figure : A point cloud consisting of three points in the plane on the edges of an equilateral triangle can be regarded as a definable map. Example 0.25 explains how the persistence modules are constructed.

In Example 0.21, we showed that the family of augmented spaces associated to a point cloud is a definable map. This example is crucial because it shows that point cloud persistence is a special case of level set persistence. By Theorem 0.24, there is a decomposition of \mathbb{R} into definable sets over which the map f is definably trivial. With some work, one can show that this decomposition is into half-open intervals $\{[s_i, s_{i+1})\}$. Let t_i denote a point strictly between s_i and s_{i+1} . Letting $X_r = f^{-1}(r)$, one can show that there is a sequence of fibers and maps

$$\cdots \leftarrow X_{t_i} \rightarrow X_{s_{i+1}} \leftarrow X_{t_{i+1}} \rightarrow X_{s_{i+2}} \leftarrow \cdots$$

where every map $X_{s_i} \leftarrow X_{t_i}$ is a homeomorphism and thus an isomorphism on homology. The fact there is such an isomorphism follows from Remark 0.23. The fact that there is a map $X_{t_i} \rightarrow X_{s_{i+1}}$ follows from the existence of a neighborhood U containing X_{t_i} and $X_{s_{i+1}}$ that deformation retracts onto $X_{s_{i+1}}$ [?, Prop. 11.1.26]. Taking homology in each degree produces the persistence modules depicted in Figure 5.

Stratified Spaces and Constructible Cosheaves

In this section we present the Leray (co)sheaves associated to a definable map in an entirely different way. This characterization is based on a folk-theorem of MacPherson [?] and is phrased in the language of Whitney stratified spaces, which includes definable sets as a special case [?], and constructible cosheaves, which we define below.

Definition (Whitney Stratified Spaces)

A **Whitney stratified space** is a space X that is a closed subset of a smooth manifold M along with a decomposition into pieces $\{X_\sigma\}_{\sigma \in P_X}$ such that

- each piece X_σ is a locally closed smooth submanifold of M , and
- whenever X_σ is in the closure of X_τ the pair satisfies **condition (b)**. This condition says if $\{y_i\}$ is a sequence in X_τ and $\{x_i\}$ is a sequence in X_σ converging to $p \in X_\sigma$ and the tangent spaces $T_{y_i}X_\tau$ converges to some plane T at p , and the secant lines ℓ_i connecting x_i and y_i converge to some line ℓ at p , then $\ell \subseteq T$. See Figure 6.

Remark

We have omitted **condition (a)** because it is implied by condition (b) [?, Prop. 2.4]. Condition (a) states that if we only consider a sequence y_i in X_τ converging to p such that the tangent planes $T_{y_i}X_\tau$ converge to some plane T , then the tangent plane to p in X_σ must be contained inside T .




Figure : Diagram for Whitney Condition (b)

The Whitney conditions are important because so many types of spaces admit Whitney stratifications, the most important being semi-algebraic and sub-analytic spaces. Remarkably, these conditions about limits of tangent spaces and secant lines imply strong structural properties of the space, such as being triangulable [?].

Definition (Entrance Path Category)

Suppose X is a stratified space. The **entrance path category** of X $\text{Entr}(X)$ has points of X for objects and equivalence classes of entrance paths for morphisms. An entrance path is a continuous map $\gamma : I = [0, 1] \rightarrow X$ with the property that the ambient dimension of the stratum containing $\gamma(t)$ is non-increasing with t . Two entrance paths γ and η connecting x to x' are equivalent if there is a map $h : [0, 1]^2 \rightarrow X$ such that for every $s \in [0, 1]$ the map $h(s, t)$ is an entrance path, $\gamma(t) = h(0, t)$ and $\eta(t) = h(1, t)$; see Figure 7. The **definable entrance path category** is similar with the added stipulation that X is definable and that all the paths and relations are definable in the sense of Definition 0.18.

Example

If X is the geometric realization of a simplicial complex, then it can be stratified by its open simplices. One can prove that $\mathbf{Entr}(X)$ is equivalent to a poset with the relation that there is a unique entrance path from τ to σ if and only if $\sigma \leq \tau$. We express this succinctly as

$$\mathbf{Entr}(X) \simeq (X, \leq)^{op}$$

The folk-theorem of MacPherson is that suitably behaved cosheaves defined on stratified spaces are equivalent to functors from the entrance path category. This equivalence would take us beyond the scope of this paper (see [?] for a more thorough treatment), so we will simply define these well-behaved cosheaves as functors from the entrance path category.




Figure : Two entrance paths in the plane related through a family of entrance paths.

Example

By Example 0.29, we see that a simplicial cosheaf on K is the same as a constructible cosheaf on the geometric realization of K , regarded as a stratified space.

The correspondence between constructible cosheaves and actual cosheaves is encapsulated in the following theorem.

Theorem (Correspondence with Cosheaves)

Given a constructible cosheaf \widehat{F} on a stratified space X one can associate an actual cosheaf, which we also call \widehat{F} , by observing that each open set U receives an induced stratification from X , and hence has an entrance path category, and letting

$$\widehat{F}(U) := \varinjlim_{\text{Entr}(U)} \widehat{F}|_U$$

Proof.

This is theorem 11.2.15 of [?]. It requires proving a Van Kampen theorem for the entrance path category, which is beyond the scope of this paper. □

Example

In Figure 8 we have two constructible cosheaves over the real line. For each constructible cosheaf we have picked the three open intervals and the corresponding colimit of the cosheaf over the entrance path category restricted to that open set.




Figure : Two constructible cosheaves and the associated colimits of the restriction to various intervals.

Now we can state a definable analog of the Leray sheaves that could be programmed on a computer.

Theorem (Constructible Cosheaves from Definable Maps [?])

If we are given a proper definable map $f : E \rightarrow B$ that comes from the restriction of a C^1 map between manifolds, then for each i the assignment

$$b \in B \rightsquigarrow H_i(f^{-1}(b))$$

defines a definable cosheaf.

Sketch of the Proof

This is a non-trivial theorem, which is proved in detail as Theorem 11.2.17 of [?]. The first observation to make is that the fiber $f^{-1}(b)$ over a point $b \in B$ has an open neighborhood U that retracts onto the fiber. This is because $f^{-1}(b)$ can be presented as a closed union of finitely many strata [?, p. 60] and the closed union of finitely many strata has a regular neighborhood that retracts onto it [?, Prop. 11.1.26].

Intuitively, if a path $\gamma : I \rightarrow B$ starts in a stratum B_τ that contains $b = \gamma(1)$ in its closure, then one can assign the homology of the zig-zag of inclusions

$$f^{-1}(\gamma(0)) \hookrightarrow U \hookleftarrow f^{-1}(\gamma(1))$$

to any morphism γ in $\mathbf{Entr}(B)$. However, we prefer a more inductive procedure by considering the pullback

$I \times_B E = \{(t, e) \mid \gamma(t) = f(e)\}$ as a definable set [?, Lem. 11.1.15] and the projection $\pi_1 : I \times_B E \rightarrow [0, 1]$ as a definable map.

To prove invariance under homotopy through entrance paths, one then considers a definable homotopy $h : I^2 \rightarrow B$ and pulls back to a definable map to the square I^2 . One then proves invariance for this restricted map.

