Interleavings and Obstructions

Justin M. Curry

July 1, 2013

Abstract

Here we introduce the definition of interleaving of pre-sheaves on a metric space X. We then offer the easy observation that $H^0(X; -)$ provides an obstruction to interleaving of pre-sheaves. For constructible sheaves over $\mathbb{R}$ this shows that open and closed barcodes are infinite distance apart.

1 Interleavings for Pre-Sheaves

The following definition was communicated to the author by Amit Patel. We provide a slight re-wording of the definition along with some novel observations of the author.

Definition 1.1. Let (X, d) be a metric space. The **ϵ-thickening of open sets** is the map of posets

$$\epsilon : \mathbf{Open}(X) \to \mathbf{Open}(X)$$

given by

$$U \rightsquigarrow U^\epsilon := \{y \in X \,|\, \exists x \in U\,,\ d(x, y) \leqslant \epsilon\}.$$

Of course, any map of posets dualizes to a map of posets $\epsilon : \mathbf{Open}(X)^{op} \to \mathbf{Open}(X)^{op}$

Remark 1.2. For metric spaces like $\mathbb{R}^n$ with the Euclidean metric, $(U^\epsilon)^\epsilon = U^{2\epsilon}$. In general, the triangle inequality implies that $(U^\epsilon)^\epsilon \subseteq U^{2\epsilon}$. The reverse containment is also true if (X, d) is convex, for example. Convexity guarantees that intuitive results are true, but it is not strictly necessary for any of the foregoing arguments.

Definition 1.3. Let X be a topological space. A **pre-sheaf valued in** $\mathcal{D}$ is a functor

$$F : \mathbf{Open}(X)^{op} \to \mathcal{D}\,.$$

Dually, a pre-cosheaf is a covariant functor from $\mathbf{Open}(X)$. We call the map associated by F to each inclusion $V \subseteq U$, the **restriction map** $\rho^F_{V,U} : F(U) \to F(V)$, with the superscript sometimes suppressed for simplicity. A map of pre-sheaves is a natural transformation of functors.

Definition 1.4 (Thickened Pre-Sheaf)**.** Using the previous two definitions, we can define the **ϵ-thickening of a pre-sheaf** F via the formula

$$F^\epsilon := F \circ \epsilon \qquad \text{i.e.} \qquad F^\epsilon(U) := F(U^\epsilon).$$

Moreover, the thickening operation is functorial. If $\varphi : F \to G$ is a natural transformation, then we get for free a natural transformation between the thickened pre-sheaves $\varphi_\epsilon : F^\epsilon \to G^\epsilon$. Consequently, we can define the **ϵ-thickening functor** to be

$$\epsilon^* : \mathbf{Pre\text{-}Shv}(X) \to \mathbf{Pre\text{-}Shv}(X) \qquad F \rightsquigarrow F^\epsilon.$$

Because F is a pre-sheaf we get a canonical natural transformation

$$\eta^F_\epsilon : F^\epsilon \to F$$

coming from $\rho_{U,U^\epsilon} : F^\epsilon(U) = F(U^\epsilon) \to F(U)$. This follows by showing that for every pair $V \subseteq U$ the square

$$\begin{array}{ccc} F^\epsilon(U) & \xrightarrow{\rho_{U,U^\epsilon}} & F(U) \\ {\scriptstyle \rho_{V^\epsilon,U^\epsilon}}\downarrow & & \downarrow{\scriptstyle \rho_{V,U}} \\ F^\epsilon(V) & \xrightarrow[\rho_{V,V^\epsilon}]{} & F(V) \end{array}$$

commutes by virtue of F being a pre-sheaf:

$$\rho_{V,U} \circ \rho_{U,U^\epsilon} = \rho_{V,U^\epsilon} = \rho_{V,V^\epsilon} \circ \rho_{V^\epsilon,U^\epsilon}$$

Of course, the map

$$\eta^F_{2\epsilon} : F^{2\epsilon} \to F$$

always exists as well and for convex metric spaces it is equal to the composition of $\eta^F_\epsilon \circ \epsilon^* \eta^F_\epsilon$.

Definition 1.5 (Interleaving of Pre-Sheaves)**.** Let $F, G : \mathbf{Open}(X)^{op} \to \mathcal{D}$ be two pre-sheaves on a metric space X. We define an **ϵ-interleaving** of F and G to be the existence of two natural transformations

$$\varphi_\epsilon : F^\epsilon \to G \qquad \psi_\epsilon : G^\epsilon \to F$$

that satisfy the compatibility relations

$$\eta^F_{2\epsilon} = \psi_\epsilon \circ \varphi_{2\epsilon} \qquad \eta^G_{2\epsilon} = \varphi_\epsilon \circ \psi_{2\epsilon}$$

where for each open set U, the natural transformation $\varphi_{2\epsilon}$ is defined by

$$\varphi_{2\epsilon}(U) : F^{2\epsilon}(U) = F(U^{2\epsilon}) \to F((U^\epsilon)^\epsilon) = F^\epsilon(U^\epsilon) \to G(U^\epsilon) = G^\epsilon(U)$$

where the last map is defined using $\varphi_\epsilon(U)$. A similar definition holds for $\psi_{2\epsilon}$.

Remark 1.6. One should note that if F and G are ϵ-interleaved and G and H are ϵ'-interleaved, then F and H are $\epsilon + \epsilon'$-interleaved.

One should also note that if F and G are 0-interleaved, then they are isomorphic

Definition 1.7. We can define an extended pseudo metric on pre-sheaves by declaring

$$d(F, G) := \inf\{\epsilon \geqslant 0 \,|\, \exists \epsilon - \text{interleaving}\}.$$

If no interleaving exists, we define $d(F, G) = \infty$. This is what we mean by an extended metric.

Remark 1.8 (Not a Pseduo-Metric for Sheaves?). For sheaves, it is true that if a map induces isomorphisms on stalks $F_x := \varinjlim_{x \in U} F(U)$, then the map is an isomorphism of sheaves. This suggests that if for every $\epsilon > 0$ there is an ϵ-interleaving, then perhaps the sheaves are isomorphic.

2 An Application of $H^0(X; -)$ to Interleavings

Using Čech theory, one can associate to any cover $\mathcal{U}$ and any pre-sheaf F, a chain complex whose cohomology groups are of interest. When F is a sheaf, and the cover is fine enough, this complex computes cohomology groups naturally gotten by taking injective resolutions and computing limits (i.e. derived limits). With this in view, we will adopt the following definition.

Definition 2.1. Given a pre-sheaf F on a topological space X, define

$$H^0(X; F) := \varprojlim F \cong F(X).$$

Lemma 2.2 (H^0 Obstructs Interleavings)**.** *If* F *and* G *are pre-sheaves on a metric space* X *and* $F(X) \ncong G(X)$*, then there is no* ϵ*-interleaving, for any value of* ϵ*.*

Proof. Clearly $X^\epsilon = X$. Recalling the compatibility condition for interleavings

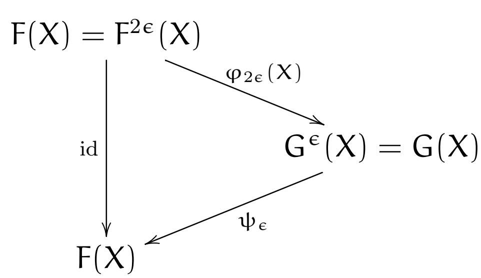

implies that $G(X) \cong F(X)$. Contraposition proves the result. □

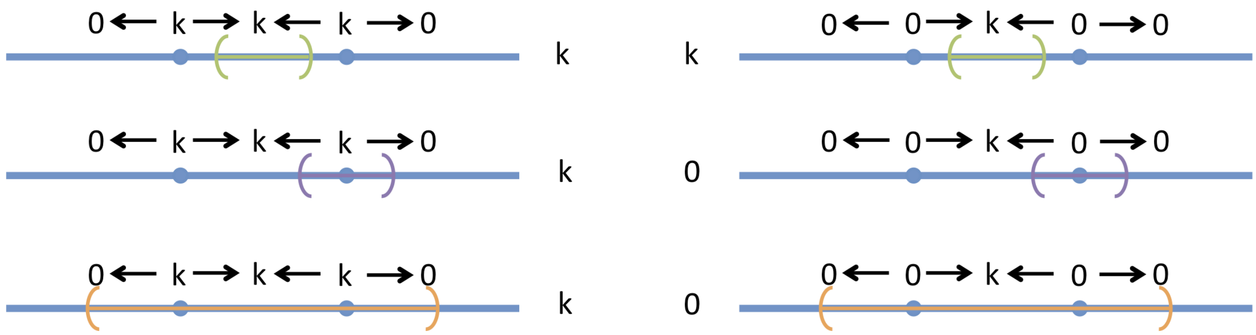

Figure 1: Cellular Description of $j_* k_Y$ and $j_! k_Y$ respectively

Example 2.3 (Sensitive to Topology)**.** Let $Y = (0, 1)$ be the open unit interval in $\mathbb{R}$. Denote the inclusion of Y into $\mathbb{R}$ by j. To the constant sheaf on Y, written k_Y, we can associate two sheaves: $j_* k_Y$ and $j_! k_Y$. To describe these sheaves, we can think constructibly. The inclusion of U defines a (non-proper) stratified map to $\mathbb{R}$. To this stratified map, we can associate two different cellular sheaves. This induces a cell structure on $\mathbb{R}$ and the associated sheaves can be viewed as a zig-zag of vector spaces where each vector space is assigned to an open cell. One can then define a sheaf that assigns values to an arbitrary open set by giving the open set the cell structure determined via restriction. One then takes the limit of the restricted zig-zag of vector spaces.[1] In figure 1 we have drawn the two cellular sheaves of interest and the value of the limit over each open set.

Since these two sheaves are constructible, they have a decomposition into indecomposables. In one dimension these indecomposables are viewed as barcodes that are sensitive to the underlying topology of the cell. Thus, viewed through the lens of barcodes, the support of the sheaf $j_* k_Y$ is the closed interval $[0, 1]$, while the support of $j_! k_Y$ is the open interval $(0, 1)$. The upshot is that

$$d(j_* k_Y, j_! k_Y) = \infty \qquad \text{or} \qquad d([-], (-)) = \infty.$$

Example 2.4 (Different than the bottleneck distance)**.** The sheaf supported on a closed interval is interleaved with a sheaf supported on a point. Neither are interleaved with the zero sheaf because the global sections are not isomorphic.

Example 2.5 (Open, Half-Open and the Zero Sheaf)**.** Sheaves supported on the half-open interval and the open interval do appear to be interleaved with the zero sheaf. This may seem distressing, but it agrees perfectly with the observation that cellular sheaf cohomology in degree zero counts closed barcodes and nothing else.

[1] One can easily dualize this discussion into cosheaves by turning the arrows around and taking colimits of the restricted diagram instead.

3 The Effect of Sheafification

In this section we investigate the impact of sheafification on interleaving of pre-sheaves. After all, every sheaf is a presheaf, so we can apply the notion of interleaving to both structures. As we will show by example, two pre-sheaves can be finite distance apart, but then be infinite distance apart (not ϵ-interleaved for any ϵ) after sheafification. Conversely, we can produce two pre-sheaves that are infinite distance apart, whose sheafifications are interleaved.

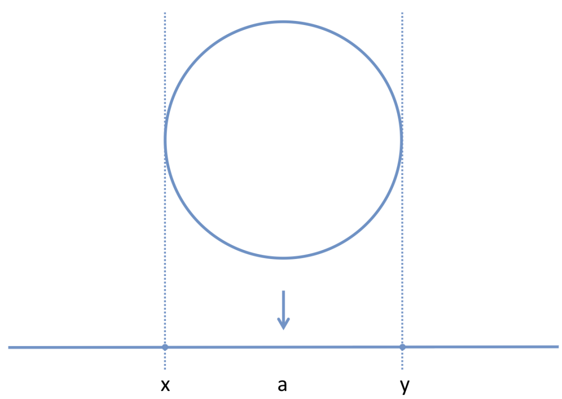

Figure 2: "High Noon" Projection f

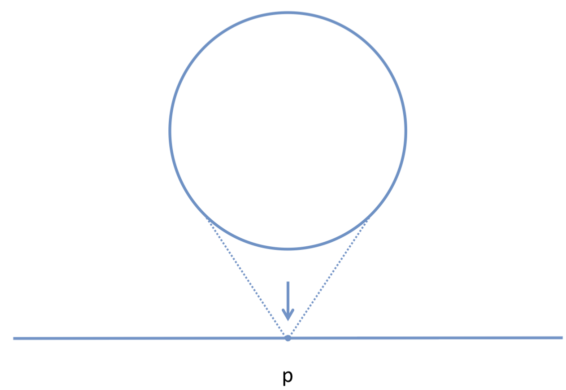

Figure 3: Constant Projection g

Example 3.1 ("High Noon" vs. Constant Projection)**.** Let us call the presheaf

$$F^1 : \quad U \rightsquigarrow H^1(f^{-1}(U); k)$$

associated to the "high Noon" projection in figure 2. Let us call the analogous presheaf associated to g in figure 3

$$G^1 : \quad U \rightsquigarrow H^1(g^{-1}(U); k).$$

One can easily check that these two presheaves are ϵ-interleaved for ϵ larger than the radius of the circle.

However, the sheafification of both of these produce radically different sheaves. Recall that the sheafification can be viewed as a sheaf of sections of the product over all the stalks mapping down to the base space.

$$\begin{array}{c} \prod_{x \in \mathbb{R}} F_x \\ \downarrow \pi \\ \mathbb{R} \end{array}$$

For figure 2, every stalk is the zero vector space. For figure 3, the stalk at p is non-zero and all other ones are zero. Consequently the sheafifications are

$$\tilde{F}^1 \cong 0 \qquad \text{and} \qquad \tilde{G}^1 \cong S_p$$

where S_p is the **skyscraper sheaf** at p, which makes the assignment

$$S_p(U) = k \qquad \text{if} \qquad p \in U$$

and zero otherwise. By the previous observations, these two sheaves are not interleaved.

One might conjecture in light of the above example that sheafification is a distance increasing operation. This is not the case.

Example 3.2 (Sheafification is not Distance Increasing)**.** Consider the presheaf F^1 from the example above. One can easily see that it is not interleaved with the zero sheaf. However, the sheafification of F^1 is the zero sheaf. So the sheafification took two presheaves that were infinite distance apart and returned isomorphic (distance zero) sheaves.