Let X be a compact topological space embedded in $\mathbf{R}^N$.

0.1 Reminder of last time

Recall a *filtration* of X is a sequence of embedded topoloigcal spaces

$$\emptyset = X_0 \subseteq X_1 \subseteq \cdots \subseteq X_n = X.$$

The inclusions induce maps maps on homology $H_k(X_i) \to H_k(X_i)$ for all i. The *persistence module* is the sequence on homology groups induced by the filtration. Persistence modules may be decomposed uniquely as the direct sum of sequences $k \to \cdots \to k$ of varying length, where every map is the identity map.

The indices where each summand starts and ends form the *persistence pair* of the homology class represented by the summand. A *persistence diagram* D is the persistence pairs of a filtration represented on $\mathbf{R}^2_+$.

Example: Recall the example of the torus from last time. A plane is chosen in $\mathbf{R}^3$ and a function $f : \mathbf{R}^3 \to \mathbf{R}$ measures the distance from x to the plane. A filtration is induced between the homological critical points.

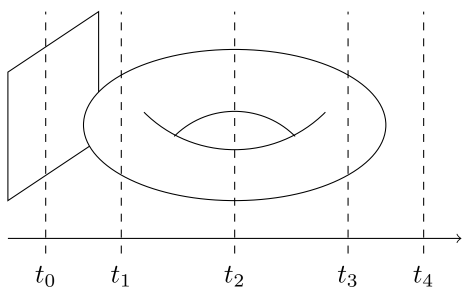

$$\begin{aligned} X_0 &= f^{-1}(-\infty, t_0] = \emptyset \\ X_1 &= f^{-1}(-\infty, t_1] \\ X_2 &= f^{-1}(-\infty, t_2] \\ X_3 &= f^{-1}(-\infty, t_3] \\ X_4 &= f^{-1}(-\infty, t_4] = X \end{aligned}$$

The sequence induced on homology groups is given below. Last time we talked about extended homology, but this time we do not mention it. Persistence pairs are given to the right of each summand.

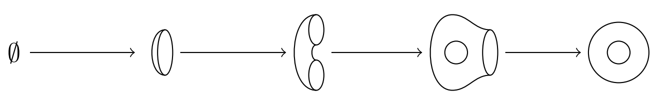

$$H_k(X_0) \longrightarrow H_k(X_1) \longrightarrow H_k(X_2) \longrightarrow H_k(X_3) \longrightarrow H_k(X_4)$$

0-class:		$\mathbf{Z}$	$\longrightarrow$		$\mathbf{Z}$	$(1, \infty)$ or (t_1, ∞)
1-class:			$\mathbf{Z}$	$\longrightarrow$	$\mathbf{Z}$	$(2, \infty)$ or (t_2, ∞)
1–class:				$\mathbf{Z} \longrightarrow$	$\mathbf{Z}$	$(3, \infty)$ or (t_3, ∞)
2-class:					$\mathbf{Z}$	$(4, \infty)$ or (t_4, ∞)

For classes that still exist at $f^{-1}(-\infty, a)$ for $a \gg 0$, they are given an end position of ∞.

0.2 Different height functions

What if we chose a different plane to measure from? How can the two filtrations be compared?

Definition: Let D, D' be two persistence diagrams with $|D| = |D'|$. The *bottleneck distance* between D and D' is

$$d(D, D') = \min_{\substack{\text{bijections} \\ \varphi: D \to D'}} \left\{ \max_{p \in D} \{ ||p - \varphi(p)||_\infty \} \right\}.$$

Example: Consider the torus of inner and outer radius 1 embedded in the natural way. Let $f, g : T^2 \to \mathbf{R}$ be height functions projecting to the planes $z = -2$ and $z = x - 4$, respectively. All critical points occur on the plane $y = 0$. Below, the slice at this plane is given (distances along planes from the first critical point are shown), as well as D_f, D_g on the same diagram (degrees of homology classes are shown).

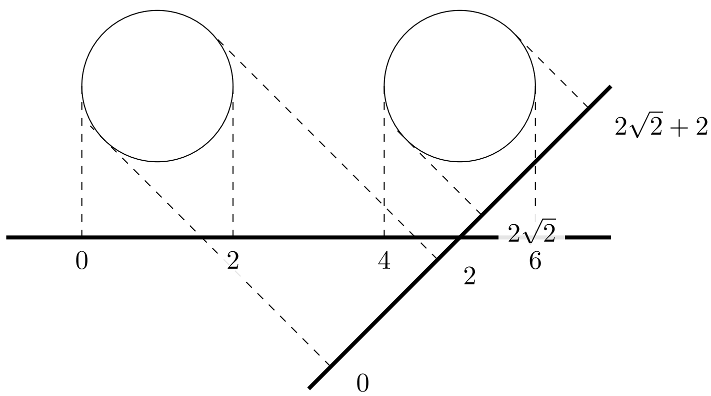

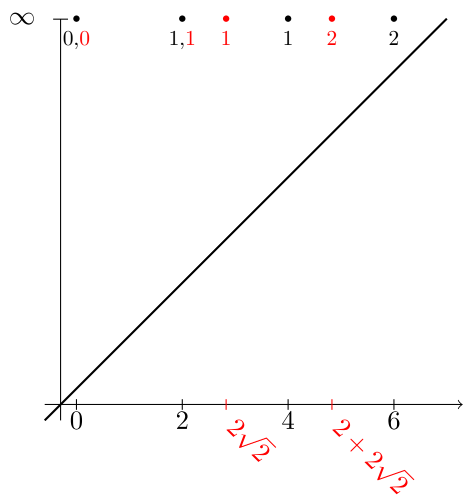

Note that

$$D_f = \{(0, \infty), (2, \infty), (4, \infty), (6, \infty)\},$$
$$D_g = \{(0, \infty), (2, \infty), (2\sqrt{2}, \infty), (2 + 2\sqrt{2}, \infty)\}.$$

Some thought indicates the identity matching $\varphi(p_{f,i}) = p_{g,i}$ will be the best matching. The bottleneck distance is $4 - 2\sqrt{2}$.

What if $|D| \neq |D'|$? In general, what if we have a different embedding of the same shape? For example:

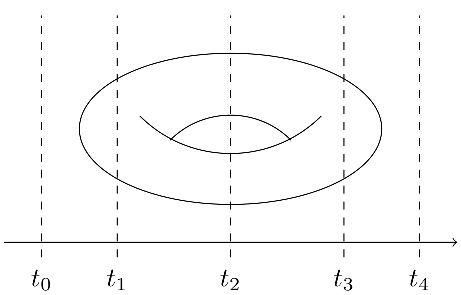

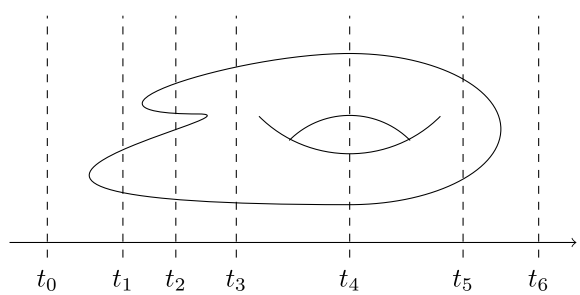

How to fix this? What if we have a filtration that does not come from a height function?

0.3 Different filtrations

Often we don't have a topological space, just some (finitely many) points sampled on this space. There is a natural filtration associated to this setup. Now let $X = \{x_1, \ldots, x_n\} \subset \mathbf{R}^n$.

Definition: Let $\epsilon > 0$. The *ϵ-nerve* of X is a simplicial complex N_ϵ with 0-skeleton X, and containing a k-cell $\sigma = \{x_{i_0}, \ldots, x_{i_k}\}$ iff $d(x_{i_j}, x_{i_\ell}) \leqslant \epsilon$ for all vertices x_{i_j}, x_{i_ℓ} of σ.

The above construction is the *Rips* complex. There is an alternate construction, called the *Čech* complex, where the complex N_ϵ contains a k-cell $\sigma = \{x_{i_0}, \ldots, x_{i_k}\}$ iff $\bigcap_{j=0}^k B(x_{i_j}, \epsilon) \neq \emptyset$. This is more difficult computationally, so we stick with the Rips complex.

Example: Suppose we are sampling from the unit circle in $\mathbf{R}^2$.

$N_{.1}$ $\qquad$ $N_{\sqrt{2}}$ $\qquad$ N_2

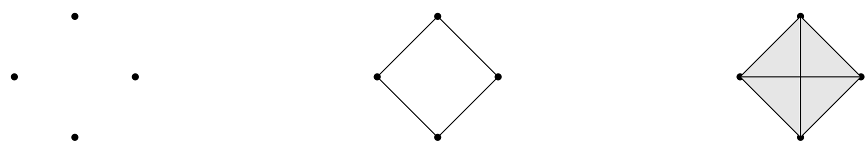

$H_0(N_{.2}) = \mathbf{Z}^4$ $\qquad$ $H_0(N_{\sqrt{2}}) = \mathbf{Z}$ $\qquad$ $H_0(N_{\sqrt{2}}) = \mathbf{Z}$

$H_1(N_{.2}) = 0$ $\qquad$ $H_1(N_{\sqrt{2}}) = \mathbf{Z}$ $\qquad$ $H_1(N_{\sqrt{2}}) = 0$

$H_2(N_{.2}) = 0$ $\qquad$ $H_2(N_{\sqrt{2}}) = 0$ $\qquad$ $H_2(N_{\sqrt{2}}) = 0$

It is possible to sample different points from the same shape.

$N_{.1}$ $\qquad$ $N_{\sqrt{2-\sqrt{2}}}$ $\qquad$ $N_{\sqrt{2}}$ $\qquad$ $N_{\sqrt{2+\sqrt{2}}}$ $\qquad$ N_2

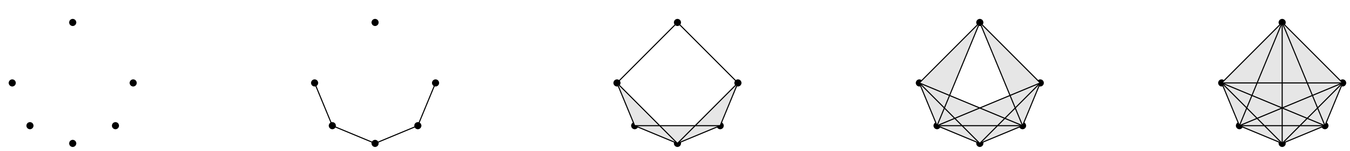

$H_0(N_{.2}) = \mathbf{Z}^6$ $\qquad$ $H_0(N_{\sqrt{2-\sqrt{2}}}) = \mathbf{Z}^2$ $\qquad$ $H_0(N_{\sqrt{2}}) = \mathbf{Z}$ $\qquad$ $H_0(N_{\sqrt{2+\sqrt{2}}}) = \mathbf{Z}$ $\qquad$ $H_0(N_2) = \mathbf{Z}$

$H_1(N_{.2}) = 0$ $\qquad$ $H_1(N_{\sqrt{2-\sqrt{2}}}) = 0$ $\qquad$ $H_1(N_{\sqrt{2}}) = \mathbf{Z}$ $\qquad$ $H_1(N_{\sqrt{2+\sqrt{2}}}) = \mathbf{Z}$ $\qquad$ $H_1(N_2) = 0$

$H_2(N_{.2}) = 0$ $\qquad$ $H_2(N_{\sqrt{2-\sqrt{2}}}) = 0$ $\qquad$ $H_2(N_{\sqrt{2}}) = 0$ $\qquad$ $H_2(N_{\sqrt{2+\sqrt{2}}}) = 0$ $\qquad$ $H_2(N_2) = 0$

Note we have 5 homology classes in the first case and 7 in the second case.

Example: A filtration could be much more general. Let $X = T^n = (S^1)^n$, the n-torus, with $X_0 = \emptyset$, and $X_i = (S^1)^i$ for $i > 0$. By the Kunneth formula, we get $H_k(X_i) = \mathbf{Z}^{\binom{i}{k}}$.

How can different filtrations be compared?

0.4 Filtered spaces and roofs

One solution to fix these problems is to define a better map between D and D'.

Definition: A *filitered space* is a pair (S, F_S), where S is a poset and $F_S : S \to \mathbf{R}$ is an order-preserving map ($A \preccurlyeq B$ in S implies $F_S(A) \leqslant F_S(B)$), called a *filtration.* For every $a \in \mathbf{R}$ set

$$L_a := \{A \in S \ : \ F_S(A) \leqslant a\},$$

which form a nested sequence

$$L \ : \ \emptyset \subseteq \cdots \subseteq L_\epsilon \subseteq L_{\epsilon'} \subseteq \cdots \subseteq$$

for some increasing sequence of indices. The set S will usually be the set of all simplices of X, or the set of all subsets of X. Hence the L_a have some topology, so for all $k \geqslant 0$ there is a sequence of vector spaces

$$H_k \ : \ 0 \to \cdots \to H_k(L_\epsilon) \to H_k(L_{\epsilon'}) \to \cdots .$$

Example: The previous filtration setups can be viewed as filtered spaces.

height function:	Set $S = \mathbf{P}(X)$ be the set of of all subsets of X. Set $F_S(A) = \max_{x\in A} f(x)$, so then $L_a \cong f^{-1}(-\infty, a]$.
ϵ-balls:	Set $S = \mathbf{P}(X)$ to be the set of all subsets of the finite point sample. Let $F_S(A) = \epsilon$ such that $\epsilon = \max_{x,y\in A}\{d(x,y)\}$. Then $\sigma \in L_a$ iff $\sigma \in N_a$.
filtration of a simplicial complex:	Set S to be the set of all simplices of X. Given a filtration $X_0 \subset X_1 \subseteq \cdots$, set $F_S(A) = i$ where i is the lowest index for which $A \subset X_i$.

To compare the persistence of different spaces, we describe a sort of "common denominator" for the two different posets that we have.

Definition: Let (S, F_S) and $(S', F_{S'})$ be two filtered spaces with S and S' finite. The *roof* of these spaces is a poset Z with order-preserving maps $\varphi_S : Z \to S$ and $\varphi_{S'} : Z \to S'$. The *pullback filtration* of φ_S is given by

$$\begin{array}{rccl} \varphi_S^* F_S \ : & Z & \to & \mathbf{R}, \\ & A & \mapsto & F_S(\varphi_S(A)), \end{array}$$

and similarly for $\varphi_{S'}$. The name of Z comes from the diagram

$$\begin{array}{ccccc} & & Z & & \\ & \varphi_S \swarrow & & \searrow \varphi_{S'} & \\ S & & & & S'. \end{array}$$

The *distance* between S and S' is

$$d_{\mathcal{F}}(S, S') := \inf_{\varphi_S, \varphi_{S'}} \left\{ \max_{A\in Z} \{|\varphi_S^* F_S(A) - \varphi_{S'}^* F_{S'}(A)|\} \right\}.$$