

VI.1 Persistent Homology

A main purpose of persistent homology is the measurement of the scale or resolution of a topological feature. There are two ingredients, one geometric, assigning a function to a space, the other algebraic, turning the function into measurements. The measurements make sense only if the function does. In this section, we focus on the second step and simplify the scenario by substituting an ordering of the simplices for the function.

Filtrations. Let K be a simplicial complex. A *filtration* is a nested sequence of subcomplexes,

$$\emptyset = K_0 \subset K_1 \subset \dots \subset K_n = K.$$

We may think of the filtration as a description of how to construct K by adding chunks at a time. We have seen an example in Section III.3 where we constructed the Delaunay complex in a sequence of alpha complexes. More than in the sequence of complexes, we are interested in their topological evolution expressed by the corresponding sequence of homology groups. Since $K_{i-1} \subset K_i$, the inclusion map defined by $f(x) = x$ induces a homomorphism between the homology groups, $f_* : H_p(K_{i-1}) \rightarrow H_p(K_i)$. The nested sequence of complexes thus corresponds to sequences of homology groups connected by homomorphisms,

$$0 = H_p(K_0) \rightarrow H_p(K_1) \rightarrow \dots \rightarrow H_p(K_n) = H_p(K),$$

one for each dimension p . The filtration defines a partial ordering on the simplices with $\sigma \in K_i - K_{i-1}$ preceding $\tau \in K_j - K_{j-1}$ if $i < j$. We can extend this to a total ordering by deciding on the ordering of the simplices within each $K_i - K_{i-1}$. We do this such that each simplex is preceded by its faces. Equivalently, we may assume that $K_i - K_{i-1}$ consists of a single simplex, σ_i , for each i . In other words, the simplices of K are ordered as $\sigma_1, \sigma_2, \dots, \sigma_n$ such that $K_i = \{\sigma_1, \sigma_2, \dots, \sigma_i\}$ for each $0 \leq i \leq n$.

Incremental algorithm. We consider the problem of updating the Betti numbers while adding a single simplex to a complex, $K_i = K_{i-1} \cup \{\sigma_i\}$ with $\dim \sigma_i = p$. The addition of σ_i changes only two boundary matrices, the p -th and the $(p+1)$ -st. Since K_{i-1} is a complex it contains none of the cofaces of σ_i . The additional row in the $(p+1)$ -st boundary matrix is therefore zero, as in Figure VI.1. This implies that the ranks of Z_{p+1} and B_p remain unchanged. However, the additional column in the p -th boundary matrix is generally non-zero and we distinguish two cases.

1. The column is a linear combination of prior columns. We can use row operations to zero-out the new column. The rank of Z_p therefore increases by one and the rank of B_{p-1} stays the same. Hence, $\beta_p(K_i) = \beta_p(K_{i-1}) + 1$ and all other Betti numbers remain as before.
2. The additional column is not a linear combination of prior columns. We can use row and column operations to extend the diagonal of ones by one position. The rank of Z_p remains unchanged and the rank of B_{p-1} increases by one. Hence, $\beta_{p-1}(K_i) = \beta_{p-1}(K_{i-1}) - 1$ and all other Betti numbers remain as before.




Figure VI.1: Adding a p -simplex adds a column to the p -th and a row to the $(p+1)$ -st boundary matrices.

The computation of Betti numbers thus reduces to deciding whether a new p -simplex gives birth to a new p -cycle and thus increases β_p or it gives death to a $(p-1)$ -cycle (changes it to a $(p-1)$ -boundary) and thus decreases β_{p-1} . Calling the former simplices *positive* and the latter *negative*, we can express the idea of persistence as pairing positive with negative simplices and this way assessing homology classes in terms of their lifetime within a filtration.

Persistent homology groups. Recall that the filtration of complexes defines a sequence of homology groups connected by homomorphisms for each dimension. We simplify the notation by writing $H_p^i = H_p(K_i)$ and add the zero homology group at the end, giving

$$0 = H_p^0 \rightarrow H_p^1 \rightarrow \dots \rightarrow H_p^n \rightarrow H_p^{n+1} = 0.$$

The homomorphisms can be composed giving maps $f_p^{i,j} : H_p^i \rightarrow H_p^j$. The image of $f_p^{i,j}$ consists of all p -dimensional homology classes that are born at or before

K_i and die after K_j . The effect of adding the zero group at the end is that every class eventually dies.

DEFINITION. The *dimension p persistent homology groups* are the images of the homomorphisms induced by inclusion, $H_p^{i,j} = \text{im } f_p^{i,j}$, for $0 \leq i \leq j \leq n+1$. The corresponding *dimension p persistent Betti numbers* are the ranks of these groups, $\beta_p^{i,j} = \text{rank } H_p^{i,j}$.

Note that $H_p^{i,i} = H_p^i$. The persistent homology groups consist of the homology classes of K_i that are still alive at K_j or, more formally, $H_p^{i,j} = Z_p^i / (B_p^j \cap Z_p^i)$, where Z_p^i and B_p^j are the p -th cycle and boundary groups of K_i and K_j . Correspondingly, the persistent Betti numbers count the independent homology classes in K_i that are still alive and independent in K_j . Equivalently, they count the independent homology classes in K_j that are born at or before K_i . We have such a number for each dimension p and each index pair $i \leq j$. To visualize all these numbers we introduce multiplicities,

$$\mu_p^{i,j} = (\beta_p^{i,j-1} - \beta_p^{i,j}) - (\beta_p^{i-1,j-1} - \beta_p^{i-1,j}),$$

for all $i < j$. We have added the parentheses to suggest the following interpretation of this formula. The first difference counts the classes in K_{j-1} born at or before K_i that die entering K_j . The second difference counts the classes in K_{j-1} born at or before K_{i-1} that die entering K_j . It follows that $\mu_p^{i,j}$ counts the p -dimensional homology classes born at K_i that die entering K_j . Since we add only one simplex at every step, there is at most one class born at K_i . For trivial reasons, this implies that there is at most one class born at K_i that dies entering K_j . Hence $\mu_p^{i,j}$ is either zero or one for each choice of p, i, j . We draw the non-zero multiplicities as points in the plane, getting a collection for each dimension p .

DEFINITION. The *dimension p persistence diagram* of the filtration, denoted as Dgm_p , is the set of points $(i, j) \in \mathbb{R}^2$ with $\mu_p^{i,j} = 1$.

Since the multiplicities are defined only for $i < j$ all points lie above the diagonal. For technical reasons which will become clear later, we usually add the points on the diagonal to the diagram. The definition of $\mu_p^{i,j}$ may be viewed as an inclusion-exclusion formula for Betti numbers. Specifically, we associate $\beta_p^{k,l}$ with the point (k, l) and do inclusion-exclusion on the four vertices of a unit square, as illustrated in Figure VI.2. Adding up the multiplicities represented by points in an upper, left quadrant cancels all terms other than that at the corner of the quadrant. This implies that a persistent Betti number can be