Stability of Multidimensional Persistent Homology

September 16, 2016

Reference

Interactive Visualization of 2-D Persistence Modules, [pdf]

The rank invariant stability via interleavings [pdf]

Seminar: Matthew Wright, Visualizing 2-Dimensional Persistent Homology [video]

Notation

Let **Simp** denote the category of simplicial complexes and simplicial maps.

Define a partial order on $\mathbb{R}^n$ by taking $(a_1, \cdots, a_n) \leq (b_1, \cdots, b_n)$ if and only if $a_i \leq b_i$ for all i.

Finite Multidimensional Filtrations

Definition. Let n-D filtration to be a functor $\mathcal{F} : \mathbb{R}^n \to \mathbf{Simp}$ such that for $a \leq b$, $\mathcal{F}(a) \subset \mathcal{F}(b)$.

Definition. n-D filtration stabilizes if there exists $a_0 \in \mathbb{R}^n$ such that $\mathcal{F}(a) = \mathcal{F}(a_0)$ whenever $a \geq a_0$. We write $\mathcal{F}_{\max} = \mathcal{F}(a_0)$.

We say that a simplex s in $\mathcal{F}_{\max}$ appears a finite number of times if there is a finite set $A \subset \mathbb{R}^n$ such that for each $a \in A$, $s \in \mathcal{F}(a)$, and for each $b \in \mathbb{R}^n$ with $s \in \mathcal{F}(b)$, there is some $a \in A$ with $a \leq b$.

Definition. We say an n-D filtration $\mathcal{F}$ is finite if

1. $\mathcal{F}$ stabilizes,
2. $\mathcal{F}_{\max}$ is finite,
3. each simplex $s \in \mathcal{F}_{\max}$ appears a finite number of times

Multidimensional Persistence Modules

Definition. For each n-D filtration $\mathcal{F}$ and for each $i \in \mathbb{N}_0$, we let associated n-D persistence module be $H_i\mathcal{F} : \mathbb{R}^n \to \mathbf{Vect}$ by $H_i\mathcal{F}(a) = H_i(\mathcal{F}(a))$.

If $\mathcal{F}$ is a finite n-D filtration, then $H_i\mathcal{F}$ is pointwise finite dimensional, i.e. $\dim(H_i\mathcal{F}(a)) < \infty$ for all $a \in \mathbb{R}^n$.

Barcodes of Multi-D Persistence Modules?

The set of isomorphism classes of indecomposable multi-D persistence modules is, in contrast to the 1-D case, extremely complicated.

Three Simple Invariants of a Multidimensional Persistence Module

1. The dimension function of $H_i\mathcal{F}$, i.e. the function which maps $a \in \mathbb{R}^2$ to $\dim H_i\mathcal{F}(a)$.

2. The fibered barcode of M, i.e. the collection of barcodes of 1-D affine slices of $H_i\mathcal{F}$.

3. The (multigraded) Betti numbers of $H_i\mathcal{F}$.

The Fibered Barcode

Let $\bar{\mathcal{L}}$ denote the collection of all lines in $\mathbb{R}^2$ with non-negative (possibly infinite) slope.

Definition. For $L \in \bar{\mathcal{L}}$ and For a finite n-D filtration $\mathcal{F}$, we define $H_i\mathcal{F}^L : L \to$ **Vect** by $H_i\mathcal{F}^L(a) = H_i\mathcal{F}(a)$.

Define an interval I in L to be a non-empty subset of L such that $a < b < c \in L$ and $a, c \in I$, we also have $b \in I$. As L is isomorphic to $\mathbb{R}$, the structure theorem for pointwise finite dimensional 1-D persistent modules yields a definition of the barcode $\mathcal{B}(H_i\mathcal{F}^L)$ as a collection of intervals in L.

Definition. We define $\mathcal{B}(H_i\mathcal{F})$, the fibered barcode of $H_i\mathcal{F}$, to be the map which sends each line $L \in \bar{\mathcal{L}}$ to the barcode $\mathcal{B}(H_i\mathcal{F}^L)$.

Stability of the Fibered Barcode

Let d_B be the bottleneck distance on the barcode.

Definition. We say two n-D filtration $\mathcal{F}$ and $\mathcal{G}$ are ϵ-interleaved if we let $\vec{\epsilon} = (\epsilon, \cdots, \epsilon) \in \mathbb{R}^n$, $\mathcal{F}(a) \subset \mathcal{G}(a+\vec{\epsilon}) \subset \mathcal{F}(a+2\vec{\epsilon})$ and $\mathcal{G}(a) \subset \mathcal{F}(a+\vec{\epsilon}) \subset \mathcal{G}(a+2\vec{\epsilon})$. We define interleaving distance d_I as $d_I(\mathcal{F}, \mathcal{G}) = \inf\{\epsilon \in [0, +\infty) : \mathcal{F} \text{ and } \mathcal{G} \text{ are } \epsilon\text{-interleaved}\}$.

We consider that a line L in $\mathbb{R}^n$ is parametrized by $u = s\vec{m} + b$ with $m^* = \min_i m_i > 0$.

Lemma. *If $\mathcal{F}$ and $\mathcal{G}$ are ϵ-interleaved, then $\mathcal{F}^L$ and $\mathcal{G}^L$ are $\frac{\epsilon}{m'}$-interleaved.*

Definition. We define $d_{match}(H_i\mathcal{F}, H_i\mathcal{G}) = \sup\limits_{L: u = s\vec{m}+b} m^* \cdot d_B(\mathcal{B}(H_i\mathcal{F}^L), \mathcal{B}(H_i\mathcal{G}^L))$.

Theorem. *For any pointwise finite dimensional $H_i\mathcal{F}$ and $H_i\mathcal{G}$,*

$$d_{match}(H_i\mathcal{F}, H_i\mathcal{G}) \leq d_I(\mathcal{F}, \mathcal{G}).$$