Distance [?]

An *extended pseudometric* on a class X is a function $d : X \times X \to [0, \infty]$ with the following three properties:

1. $d(x, x) = 0$ for all $x \in X$,
2. $d(x, y) = d(y, x)$ for all $x, y \in X$,
3. $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$

Note that an extended pseudometric d is a metric if

1. d is finite
2. $d(x, y) = 0$ implies $x = y$.

distance = extended pseudometric.

Decorated Endpoints and Intervals

	t^-	t^+	∞
$-\infty$	$(-\infty, t)$	$(-\infty, t]$	$(-\infty, \infty)$
s^-	$[s, t)$	$[s, t]$	$[s, \infty)$
s^+	(s, t)	$(s, t]$	(s, ∞)

$$\langle s, t \rangle = (s^*, t^*)$$

2.1 Persistence Modules Over a Real Parameter [?]

Example: **Real line**

Objects: real numbers

Morphisms: $s \to t$ if $s \leq t$

Example: **Vec**

Objects: Vector spaces

Morphisms: Linear maps

Example: Persistence module $\mathbb{V}$ over $\mathbb{R}$ is a functor from the **Real line** into **Vec**

I.e., $\mathbb{V} = \{V_t \mid t \in \mathbb{R}\}$ with linear maps $\{v_t^s : V_s \to V_t \mid s \leq t\}$

2.3 Module Categories [?]

Defn: Given **T**-persistent modules $\mathbb{U}, \mathbb{V} : \mathbf{T} \to \mathbf{Vec}$, a **homomorphism** $\phi : \mathbb{U} \Rightarrow \mathbb{V}$ consists of:

- a collection of linear maps $\{\phi_t : U_t \to V_t \mid t \in \mathbf{T}\}$

such that:

- for any morphism $s \leq t$ in $\mathbb{U}$, this diagram commutes:

$$\begin{array}{ccc} U_s & \xrightarrow{u_t^s} & U_t \\ \downarrow{\scriptstyle \phi_s} & & \downarrow{\scriptstyle \phi_t} \\ V_s & \xrightarrow[v_t^s]{} & V_t \end{array}$$

$Hom(\mathbb{U}, \mathbb{V}) = \{homomorphisms\ \mathbb{U} \Rightarrow \mathbb{V}\}$

$End(\mathbb{V}) = \{homomorphisms\ \mathbb{V} \Rightarrow \mathbb{V}\}$

$\mathbf{Vect}^{\mathbf{R}}$

Objects: Persistence modules

i.e., functors from the **Real line** into **Vec**).

Morphisms: Natural transformations = homomorphisms.

2.4 Interval Modules [?]

Let **T** be a totally ordered set.

$J \subset \mathbf{T}$ is an **interval** if $r, t \in J$ and if $r < s < t$, then $s \in J$.

The **interval module** $\mathcal{I} = \mathbf{k}^J$ is the **T**-persistence module with vector spaces

$$I_t = \begin{cases} \mathbf{k} & \text{if } t \in J \\ 0 & \text{otherwise} \end{cases}$$

and linear maps

$$i_t^s = \begin{cases} 1 & \text{if } s, t \in J \\ 0 & \text{otherwise} \end{cases}$$

In informal language, $\mathbf{k}^J$ represents a 'feature' which 'persists' over exactly the interval J and nowhere else.

I.e, $\mathbf{k}^J$ represents a bar in the barcode.

Shift Functors

the *shift functor* $(\cdot)(\delta) : \mathbf{Vect}^{\mathbf{R}} \to \mathbf{Vect}^{\mathbf{R}}$

$$\mathbb{V} = (V_t, v_t^s) \quad \to \quad \mathbb{V}(\delta) = (V(\delta)_t = V_{t+\delta}, v_{t+\delta}^{s+\delta})$$

For a morphism $f \in \hom(\mathbf{Vect}^{\mathbf{R}})$ we define $f(\delta)$ by $f(\delta)_t = f_{t+\delta}$.

If a persistence module $\mathbb{V}$ indexed over $\mathbb{R}$ can be decomposed,

$$\mathbb{V} \cong \bigoplus_{\ell \in L} \mathbf{k}\langle p_\ell, q_\ell \rangle$$

$$\text{then } \mathbb{V}(\delta) \cong \bigoplus_{\ell \in L} \mathbf{k}(p_\ell^* - \delta, q_\ell^* - \delta)$$

That is the barcode $\mathcal{B}_{M(\delta)}$ is obtained from $\mathcal{B}_M$ by shifting all intervals to the left by δ, as in Fig. 1.

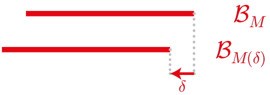

Figure: *Corresponding intervals in $\mathcal{B}_M$ and $\mathcal{B}_{M(\delta)}$.*

Transition Morphisms

For V a persistence module and $\delta \geq 0$, let the

δ-transition morphism $\varphi^{\delta}_{\mathbb{V}} : \mathbb{V} \Rightarrow \mathbb{V}(\delta)$ be the morphism whose restriction to V_t is the linear map $v^t_{t+\delta}$

That is $\varphi^{\delta}_{\mathbb{V}}$ consists of

- a collection of linear maps $\{v^t_{t+\delta} : V_t \to V_{t+\delta} \mid t \in \mathbb{R}\}$

such that:

- for any morphism $s \leq t$ in $\mathbb{U}$, this diagram commutes:

$$\begin{array}{ccc} V_s & \xrightarrow{v^s_t} & V_t \\ {\scriptstyle v^s_{s+\delta}}\downarrow & & \downarrow{\scriptstyle v^t_{t+\delta}} \\ V_{s+\delta} & \xrightarrow[v^s_t]{} & V_{t+\delta} \end{array}$$

Interleavings

We say that two persistence modules $\mathbb{U}$ and $\mathbb{V}$ are *δ-interleaved* if there exist morphisms $f : \mathbb{U} \to \mathbb{V}(\delta)$ and $g : \mathbb{V} \to \mathbb{U}(\delta)$ such that

$$g(\delta) \circ f = \varphi_{\mathbb{U}}^{2\delta},$$
$$f(\delta) \circ g = \varphi_{\mathbb{V}}^{2\delta}.$$

We refer to such f and g as *δ-interleaving* morphisms. The definition of δ-interleaving morphisms was introduced in [**?**].

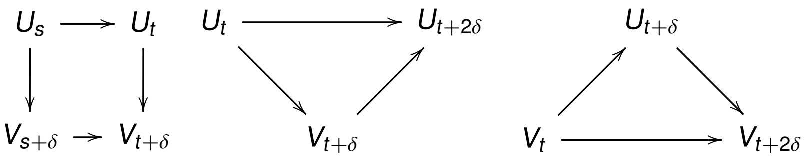

If L and M are δ-interleaved, and M and N are δ'-interleaved, then L and N are $(\delta + \delta')$-interleaved.

If $0 \leq \delta \leq \delta'$ and M and N are δ-interleaved, then M and N are also δ'-interleaved.

For T a topological space and functions $\gamma, \kappa : T \to \mathbb{R}$, let

$$d_\infty(\gamma, \kappa) = \sup_{y \in T} |\gamma(y) - \kappa(y)|.$$

Suppose $d_\infty(\gamma, \kappa) = \delta$, and let $\mathcal{S}^\gamma_t = \{x \in T \mid \gamma(x) \leq t\}$

Then for each $t \in \mathbb{R}$, we have inclusions

$$\mathcal{S}^\gamma_t \subseteq \mathcal{S}^\kappa_{t+\delta},$$
$$\mathcal{S}^\kappa_t \subseteq \mathcal{S}^\gamma_{t+\delta}.$$

Applying the i^{th} homology functor with coefficients in K to the collection of all such inclusion maps yields a δ-interleaving between $H_i(\mathcal{S}^\gamma)$ and $H_i(\mathcal{S}^\kappa)$.

The Interleaving Distance

Define the *interleaving distance*

$d_I : \mathrm{obj}(\mathbf{Vect}^{\mathbf{R}}) \times \mathrm{obj}(\mathbf{Vect}^{\mathbf{R}}) \to [0, \infty]$, by taking

$$d_I(M, N) = \inf \{\delta \in [0, \infty) \mid M \text{ and } N \text{ are } \delta\text{-interleaved}\}.$$

In addition, if M, M', and N are persistence modules with $M \cong M'$, then $d_I(M, N) = d_I(M', N)$.

Thus d_I descends to a distance on isomorphism classes of persistence modules.

Mch

Objects: Sets

Morphisms: matchings.

A *matching* from S to T (written as $\sigma : S \nrightarrow T$) is a bijection $\sigma : S' \to T'$, for some $S' \subseteq S$, $T' \subseteq T$;

For matchings $\sigma : S \nrightarrow T$ and $\tau : T \nrightarrow U$,

$$\tau \circ \sigma = \{(s,u) \mid (s,t) \in \sigma, (t,u) \in \tau \text{ for some } t \in T\}.$$

Note that any injective function is a matching.

$\operatorname{coim} \sigma = S'$ and $\operatorname{im} \sigma = T'$.

Matchings and the Bottleneck Distance [?]

For $\mathcal{D}$ a barcode and $\epsilon \geq 0$, let

$$\mathcal{D}^\epsilon = \{\langle b, d\rangle \in \mathcal{D} \mid b+\epsilon < d\} = \{I \in \mathcal{D} \mid [t, t+\epsilon] \subseteq I \text{ for some } t \in \mathbb{R}\}.$$

Note that $\mathcal{D}^0 = \mathcal{D}$. We define a δ-*matching* between barcodes $\mathcal{C}$ and $\mathcal{D}$ to be a matching $\sigma : \mathcal{C} \nrightarrow \mathcal{D}$ such that

1. $\mathcal{C}^{2\delta} \subseteq \operatorname{coim} \sigma$,
2. $\mathcal{D}^{2\delta} \subseteq \operatorname{im} \sigma$,
3. if $\sigma\langle b, d\rangle = \langle b', d'\rangle$, then

$$\langle b, d\rangle \subseteq \langle b' - \delta, d' + \delta\rangle,$$
$$\langle b', d'\rangle \subseteq \langle b - \delta, d + \delta\rangle.$$

Matchings and the Bottleneck Distance [?]

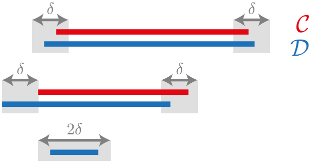

Figure: *A δ-matching between two barcodes $\mathcal{C}$ and $\mathcal{D}$. Endpoints of matched intervals are at most δ apart, and unmatched intervals are of length at most 2δ.*

Matchings and the Bottleneck Distance [?]

If $\sigma_1 : \mathcal{C} \nrightarrow \mathcal{D}$ is a δ_1-matching, and $\sigma_2 : \mathcal{D} \nrightarrow \mathcal{E}$ is a δ_2-matching, then $\sigma_2 \circ \sigma_1 : \mathcal{C} \to \mathcal{E}$ is a $(\delta_1 + \delta_2)$-matching.

We define the bottleneck distance d_B by

$$d_B(\mathcal{C}, \mathcal{D}) = \inf \{\delta \in [0, \infty) \mid \exists \text{ a } \delta\text{-matching between } \mathcal{C} \text{ and } \mathcal{D}\}.$$

The triangle inequality for d_B follows immediately from the composition above.

d_B is the most commonly considered distance on barcodes in the persistent homology literature. This is in part because d_B is especially well behaved from a theoretical standpoint.

The Isometry Theorem

A persistence module M is said to be *pointwise finite dimensional (p.f.d.)* if each of the vector spaces M_t is finite dimensional.

[Isometry Theorem for p.f.d. Persistence Modules] Two p.f.d. persistence modules M and N are δ-interleaved if and only if there exists a δ-matching between $\mathcal{B}_M$ and $\mathcal{B}_N$. In particular,

$$d_B(\mathcal{B}_M, \mathcal{B}_N) = d_I(M, N).$$

[**? ?**] For any topological space T, functions $\gamma, \kappa : T \to \mathbb{R}$, and $i \geq 0$ such that $H_i(\mathcal{S}^\gamma)$ and $H_i(\mathcal{S}^\kappa)$ are p.f.d.,

$$d_B(\mathcal{B}_{H_i(\mathcal{S}^\gamma)}, \mathcal{B}_{H_i(\mathcal{S}^\kappa)}) \leq d_\infty(\gamma, \kappa).$$