An introduction to persistent homology with applications to shape analysis

Marc Ethier

Department of Mathematical Sciences
Université de Saint-Boniface
Winnipeg, Manitoba

November 24, 2017

Shapes

Here is a hand shape :

Shapes

Here is a hand shape :

We would like to identify it as an instance of shapes such as

, and not such as

.

Shapes

Here is a hand shape :

We would like to identify it as an instance of shapes such as

, and not such as

.

How do we do this?

Shape signatures

To a shape we wish to attribute a *signature*, also called *shape descriptor*, that is a mathematical object encapsulating the shape's properties and allowing its comparison to other shapes.

Shape signatures

To a shape we wish to attribute a *signature*, also called *shape descriptor*, that is a mathematical object encapsulating the shape's properties and allowing its comparison to other shapes.
This signature shall have to be computable algorithmically.

Shape signatures

To a shape we wish to attribute a *signature*, also called *shape descriptor*, that is a mathematical object encapsulating the shape's properties and allowing its comparison to other shapes.
This signature shall have to be computable algorithmically.
For this we will use *persistent homology*.

Filtration

Definition

A *filtration of topological spaces* is a collection X_u of topological spaces where $u \in I$ with I an ordered set, with the property that $X_u \subseteq X_v$ and that the topology of X_u is induced by that of X_v if $u \leq v$.

Filtration

Definition

A *filtration of topological spaces* is a collection X_u of topological spaces where $u \in I$ with I an ordered set, with the property that $X_u \subseteq X_v$ and that the topology of X_u is induced by that of X_v if $u \leq v$.

The inclusion $j^{(u,v)} : X_u \hookrightarrow X_v$ induces for every q a linear map $H_q(j^{(u,v)})$ between the homology spaces $H_q(X_u)$ and $H_q(X_v)$.

Filtration

Definition

A *filtration of topological spaces* is a collection X_u of topological spaces where $u \in I$ with I an ordered set, with the property that $X_u \subseteq X_v$ and that the topology of X_u is induced by that of X_v if $u \leq v$.

The inclusion $j^{(u,v)} : X_u \hookrightarrow X_v$ induces for every q a linear map $H_q(j^{(u,v)})$ between the homology spaces $H_q(X_u)$ and $H_q(X_v)$. $H_q(j^{(u,v)})$ contains the homology q-cycles of X_u that still *persist* in X_v. If $v = u + t$, we may say that these cycles have a duration of persistence equal to or greater than t.

Persistence computation

$t = 0$

Persistence computation

$$t = 1$$

Persistence computation

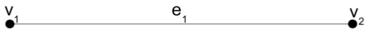

$$t = 2$$

Persistence computation

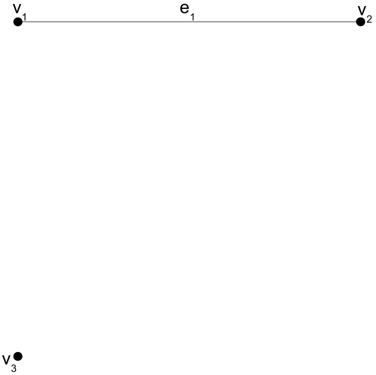

$t = 3$

Persistence computation

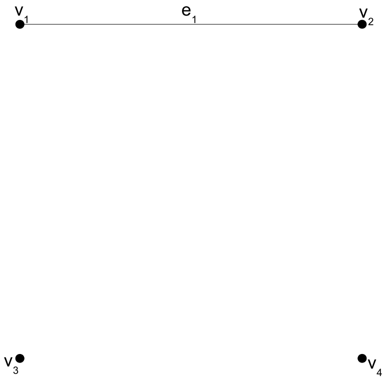

$t = 4$

Persistence computation

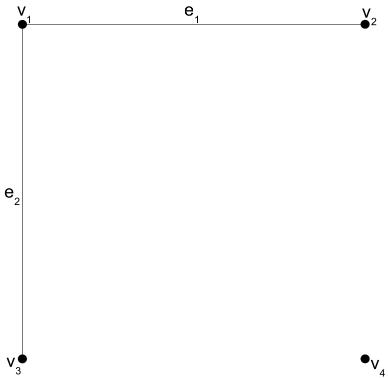

$t = 5$

Persistence computation

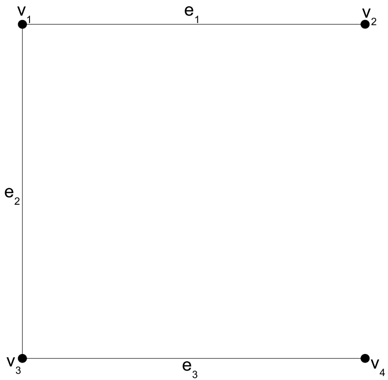

$t = 6$

Persistence computation

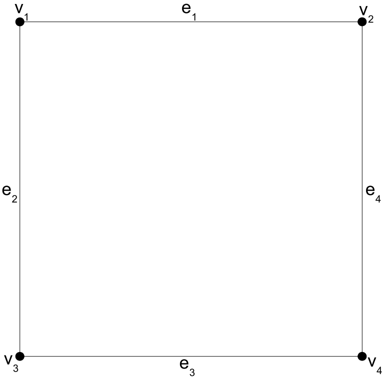

$t = 7$

Persistence computation

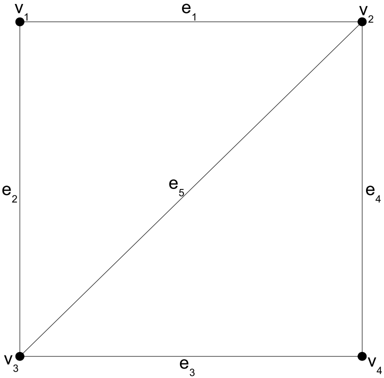

$t = 8$

Persistence computation

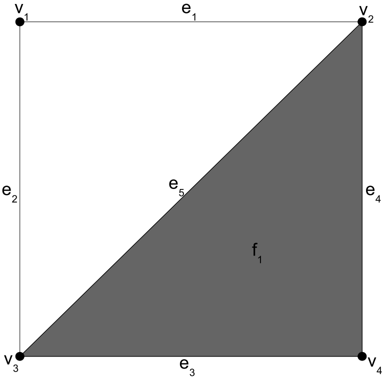

$t = 9$

Persistence computation

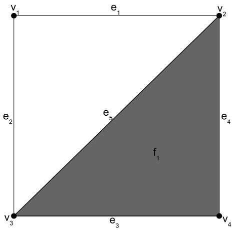

	v_1	v_2	v_3	v_4	e_1	e_2	e_3	e_4	e_5	f_1
v_1					1	1				
v_2					1			1	1	
v_3						1	1		1	
v_4							1	1		
e_1										
e_2										
e_3										1
e_4										1
e_5										1
f_1										

Persistence computation

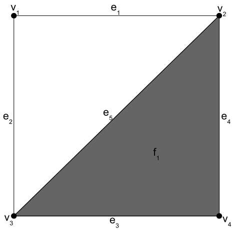

	v_1	v_2	v_3	v_4	e_1	e_2	e_3	e_4	e_5	f_1
v_1					1	1				
v_2					1			1	1	
v_3						1	1	1	1	
v_4							1			
e_1										
e_2										
e_3										1
e_4										1
e_5										1
f_1										

Persistence computation

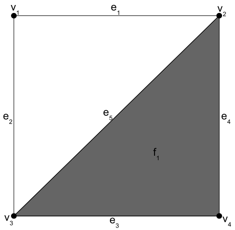

	v_1	v_2	v_3	v_4	e_1	e_2	e_3	e_4	e_5	f_1
v_1					1	1		1	1	
v_2					1			1	1	
v_3						1	1			
v_4							1			
e_1										
e_2										
e_3										1
e_4										1
e_5										1
f_1										

Persistence computation

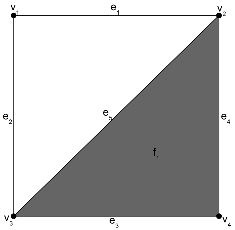

	v_1	v_2	v_3	v_4	e_1	e_2	e_3	e_4	e_5	f_1
v_1					1	1				
v_2					1					
v_3						1	1			
v_4							1			
e_1										
e_2										
e_3										1
e_4										1
e_5										1
f_1										

Persistence computation

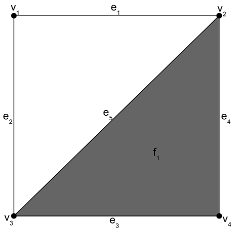

	v_1	v_2	v_3	v_4	e_1	e_2	e_3	e_4	e_5	f_1
v_1					1	1				
v_2					1					
v_3						1	1			
v_4							1			
e_1										
e_2										
e_3										1
e_4										1
e_5										1
f_1										

Simplex pairs : (v_2, e_1), (v_3, e_2), (v_4, e_3), (e_5, f_1), (v_1, ∞), (e_4, ∞)

Persistence computation

The preceding algorithm for computing persistence is taken from D. Cohen-Steiner, H. Edelsbrunner and D. Morozov. Vines and vineyards by updating persistence in linear time. In *Proc. 22nd Ann. Sympos. Comput. Geom.*, 119–126, 2006.

Sublevel filtration

Definition

Let $f : X \to \mathbb{R}$ be a continuous function. For $u \in \mathbb{R}$, we define

$$X_u = \{x \in X \mid f(x) \leq u\}.$$

f is called a *filtering* or *measuring function*, and X_u, $u \in \mathbb{R}$ is a filtration, called the *sublevel filtration* for f.

Definition

The *order q persistent Betti number function* for (X, f) is the function defined as

$$\beta_{f,q}(u, v) = \dim \operatorname{im} H_q(j^{(u,v)}).$$

Persistence diagrams

The persistent Betti number function β_f can be represented by a *persistence diagram* $\mathrm{Dgm}(\beta_f)$, that is, a *multiset* of *proper cornerpoints* and *cornerpoints at infinity*. The cornerpoints represent birth and death times for a homological cycle.

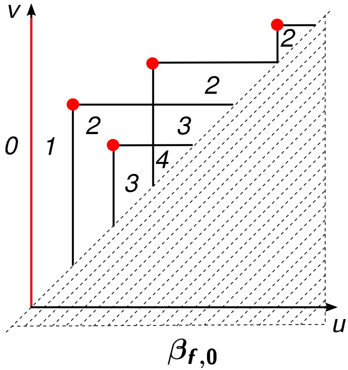

Persistence diagrams

The persistent Betti number function β_f can be represented by a *persistence diagram* $\mathrm{Dgm}(\beta_f)$, that is, a *multiset* of *proper cornerpoints* and *cornerpoints at infinity*. The cornerpoints represent birth and death times for a homological cycle.

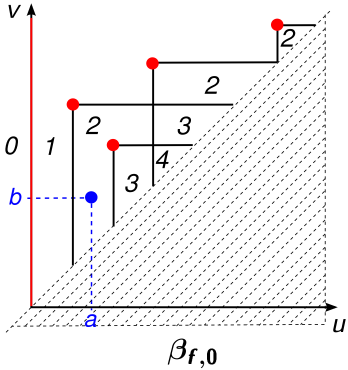

Persistence diagrams

The persistent Betti number function β_f can be represented by a *persistence diagram* $\mathrm{Dgm}(\beta_f)$, that is, a *multiset* of *proper cornerpoints* and *cornerpoints at infinity*. The cornerpoints represent birth and death times for a homological cycle.

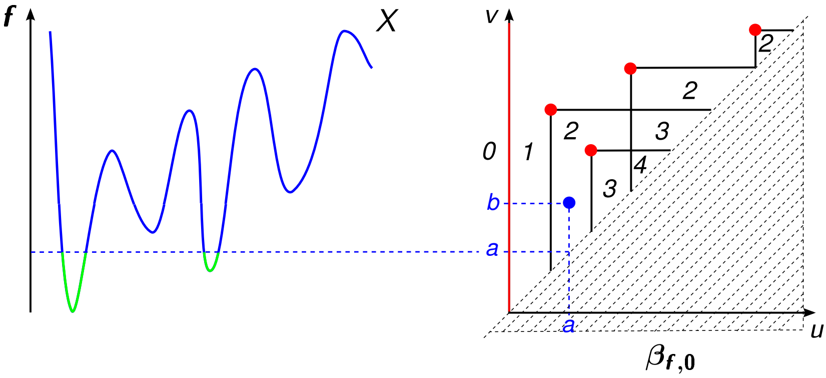

Persistence diagrams

The persistent Betti number function β_f can be represented by a *persistence diagram* $\mathrm{Dgm}(\beta_f)$, that is, a *multiset* of *proper cornerpoints* and *cornerpoints at infinity*. The cornerpoints represent birth and death times for a homological cycle.

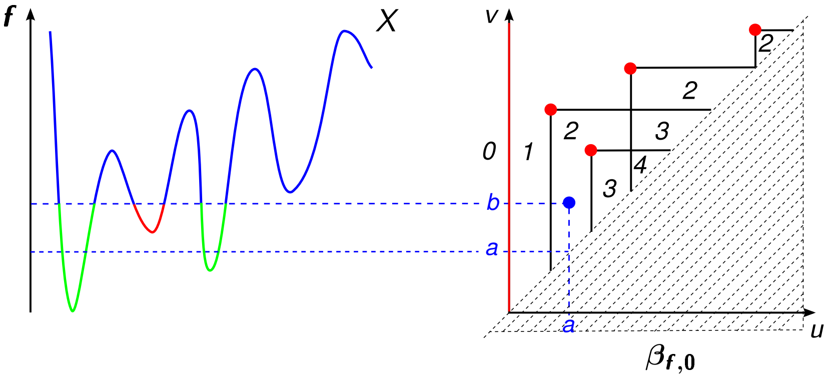

Persistence diagrams

The persistent Betti number function β_f can be represented by a *persistence diagram* $\mathrm{Dgm}(\beta_f)$, that is, a *multiset* of *proper cornerpoints* and *cornerpoints at infinity*. The cornerpoints represent birth and death times for a homological cycle.

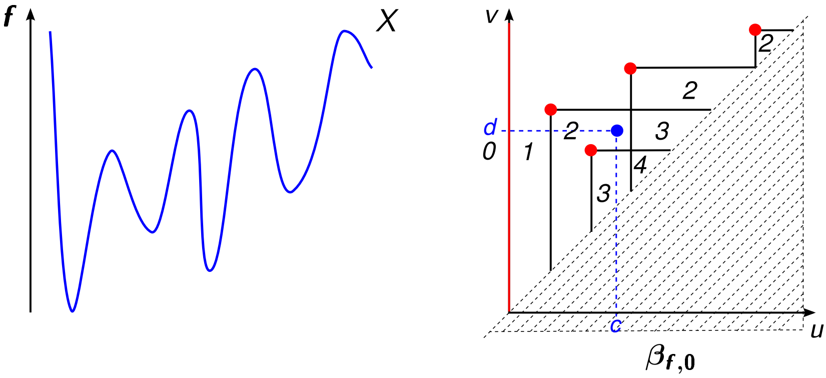

Persistence diagrams

The persistent Betti number function β_f can be represented by a *persistence diagram* $\mathrm{Dgm}(\beta_f)$, that is, a *multiset* of *proper cornerpoints* and *cornerpoints at infinity*. The cornerpoints represent birth and death times for a homological cycle.

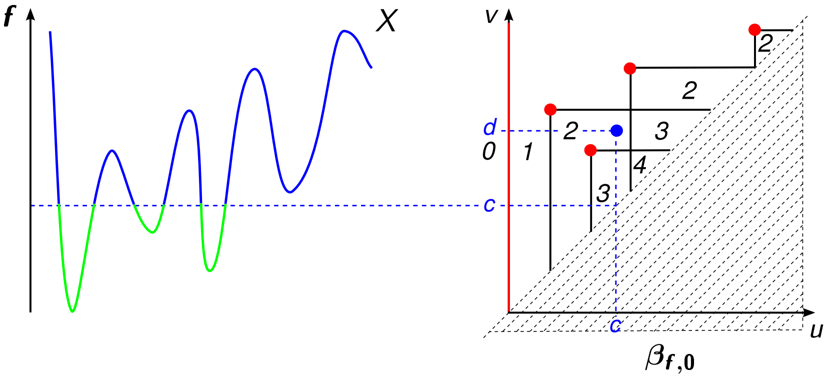

Persistence diagrams

The persistent Betti number function β_f can be represented by a *persistence diagram* $\mathrm{Dgm}(\beta_f)$, that is, a *multiset* of *proper cornerpoints* and *cornerpoints at infinity*. The cornerpoints represent birth and death times for a homological cycle.

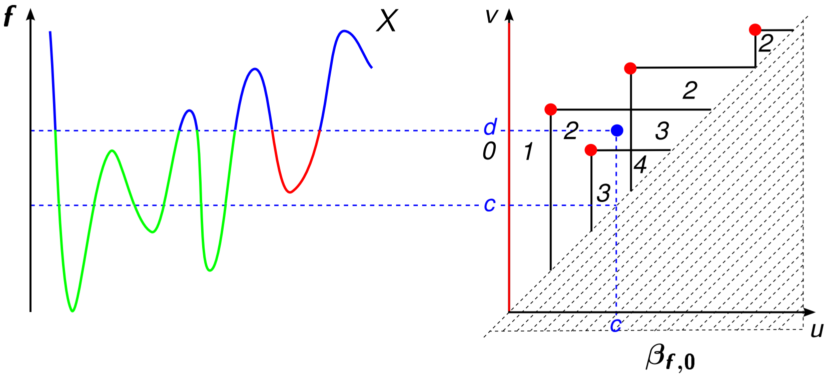

Bottleneck distance

Definition

Let β_f and β_g be two persistent Betti number functions for $f : X \to \mathbb{R}$ and $g : Y \to \mathbb{R}$. The *bottleneck distance* or *matching distance* between β_f and β_g is

$$d(\beta_f, \beta_g) = \min_{\sigma} \max_{p \in \mathrm{Dgm}(\beta_f)} \delta(p, \sigma(p))$$

where σ is taken along the set of bijections between $\mathrm{Dgm}(\beta_f)$ and $\mathrm{Dgm}(\beta_g)$ and where

$$\delta((u,v),(u',v')) = \min\left\{\max\{|u-u'|, |v-v'|\}, \max\left\{\frac{v-u}{2}, \frac{v'-u'}{2}\right\}\right\}.$$

Bottleneck distance

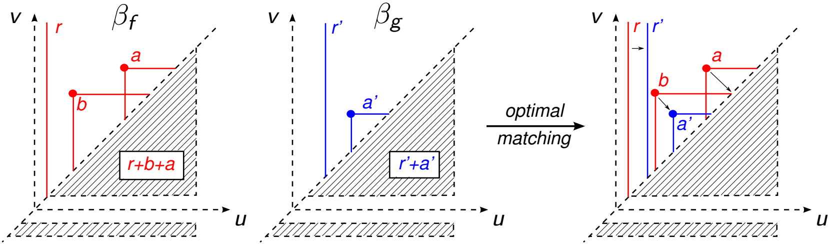

Call Δ the diagonal $\{(u,v) \in \mathbb{R} \times \mathbb{R} \mid u = v\}$, Δ^+ the set $\{(u,v) \in \mathbb{R} \times \mathbb{R} \mid u < v\}$, and Δ^* the set Δ^+ to which we add the points at infinity of the form (u,∞). Persistence diagrams' cornerpoints therefore exist in the space $\Delta^* \cup \{\Delta\}$.

Stability

Theorem

Let $f, g : X \to \mathbb{R}$ be two continuous filtering functions, and β_f and β_g associated persistent Betti number functions. Then

$$d(\beta_f, \beta_g) \leq \|f - g\|.$$

This result proves the stability of the matching distance with respect to perturbations in the measuring function.

Multifiltration

We would like to use the discriminatory power of more than one filtering function at a time.

Multifiltration

We would like to use the discriminatory power of more than one filtering function at a time.
Solution : multifiltration.

Definition

A *multifiltration* is an indexed collection of topological spaces $X_{\vec{u}}$ where $\vec{u} \in I^k$, with the property that $X_{\vec{u}} \subseteq X_{\vec{v}}$ and that the topology of $X_{\vec{u}}$ is induced by that of $X_{\vec{v}}$ if $\vec{u} \preceq \vec{v}$, that is, if $u_i \leq v_i$ for $i = 1, \ldots, k$.

Sublevel multifiltration

Definition

We shall consider a continuous function $\vec{f} = (f_1, \ldots, f_k) : X \to \mathbb{R}^k$. For $\vec{u} \in \mathbb{R}^k$, we shall define

$$X_{\vec{u}} = \{x \in X \mid \vec{f}(x) \preceq \vec{u}\}.$$

$X_{\vec{u}}$, $\vec{u} \in \mathbb{R}^k$ is called the *sublevel multifiltration* with respect to the filtering function $\vec{f}$.

Multidimensional persistence diagrams

There exists no shape signature as compact as persistence diagrams in the case of multidimensional persistent homology. However, we can reduce its computation to ordinary persistent homology computation for a parametrized family of $\mathbb{R}$-valued functions.

Reduction

Definition

We will call *admissible* a pair $(\vec{l}, \vec{b}) \in \mathbb{R}^k \times \mathbb{R}^k$ where $l_i > 0$ for $i = 1, \ldots, k$, $\sum l_i = 1$, $\sum b_i = 0$, and we will denote the set of admissible pairs Adm_k.

Reduction

Definition

We will call *admissible* a pair $(\vec{l}, \vec{b}) \in \mathbb{R}^k \times \mathbb{R}^k$ where $l_i > 0$ for $i = 1, \ldots, k$, $\sum l_i = 1$, $\sum b_i = 0$, and we will denote the set of admissible pairs Adm_k.

$(\vec{l}, \vec{b}) \in \mathrm{Adm}_k$ corresponds to a line in $\mathbb{R}^k$ whose direction vector has positive components (in $\mathbb{R}^2$, this means a line of positive slope).

Reduction

Definition

We will call *admissible* a pair $(\vec{l}, \vec{b}) \in \mathbb{R}^k \times \mathbb{R}^k$ where $l_i > 0$ for $i = 1, \ldots, k$, $\sum l_i = 1$, $\sum b_i = 0$, and we will denote the set of admissible pairs Adm_k.

$(\vec{l}, \vec{b}) \in \mathrm{Adm}_k$ corresponds to a line in $\mathbb{R}^k$ whose direction vector has positive components (in $\mathbb{R}^2$, this means a line of positive slope). We notice that for all $(\vec{u}, \vec{v}) \in \mathbb{R}^k \times \mathbb{R}^k$ with $\vec{u} \preceq \vec{v}$, there is a unique admissible pair $(\vec{l}, \vec{b})$ such that for well-chosen $s, t \in \mathbb{R}$, $\vec{u} = s\vec{l} + \vec{b}$ and $\vec{v} = t\vec{l} + \vec{b}$.

Reduction

Definition

Let $\vec{f}: X \to \mathbb{R}^k$ be a measuring function. For every pair $(\vec{l}, \vec{b}) \in \mathrm{Adm}_k$, we define

$$f_{(\vec{l},\vec{b})}(x) = \min_{i=1,\dots,k} l_i \max_{i=1,\dots,k} \frac{f_i(x) - b_i}{l_i}.$$

Reduction

Definition

Let $\vec{f} : X \to \mathbb{R}^k$ be a measuring function. For every pair $(\vec{l}, \vec{b}) \in \mathrm{Adm}_k$, we define

$$f_{(\vec{l},\vec{b})}(x) = \min_{i=1,\dots,k} l_i \max_{i=1,\dots,k} \frac{f_i(x) - b_i}{l_i}.$$

Theorem

For every $(\vec{u}, \vec{v}) \in \mathbb{R}^k \times \mathbb{R}^k$ with $\vec{u} \prec \vec{v}$, $\vec{u} = s\vec{l} + \vec{b}$ and $\vec{v} = t\vec{l} + \vec{b}$, we have

$$\beta_{\vec{f}}(\vec{u}, \vec{v}) = \beta_{f_{(\vec{l},\vec{b})}}(s, t).$$

Reduction

Definition

Let $\vec{f} : X \to \mathbb{R}^k$ be a measuring function. For every pair $(\vec{l}, \vec{b}) \in \mathrm{Adm}_k$, we define

$$f_{(\vec{l},\vec{b})}(x) = \min_{i=1,\dots,k} l_i \max_{i=1,\dots,k} \frac{f_i(x) - b_i}{l_i}.$$

Theorem

For every $(\vec{u}, \vec{v}) \in \mathbb{R}^k \times \mathbb{R}^k$ with $\vec{u} \prec \vec{v}$, $\vec{u} = s\vec{l} + \vec{b}$ and $\vec{v} = t\vec{l} + \vec{b}$, we have

$$\beta_{\vec{f}}(\vec{u}, \vec{v}) = \beta_{f_{(\vec{l},\vec{b})}}(s, t).$$

The computation of $\beta_{\vec{f}}$ may therefore be reduced to the computation of $\beta_{f_{(\vec{l},\vec{b})}}$ for each $(\vec{l}, \vec{b}) \in \mathrm{Adm}_k$.

Multidimensional matching distance

Definition

Let $\beta_{\vec{f}}$ et $\beta_{\vec{g}}$ be two k-dimensional persistent Betti numbers for $\vec{f} : X \to \mathbb{R}^k$ et $\vec{g} : Y \to \mathbb{R}^k$. The *$k$-dimensional matching distance* between $\beta_{\vec{f}}$ et $\beta_{\vec{g}}$ is defined as

$$D(\beta_{\vec{f}}, \beta_{\vec{g}}) = \max_{(\vec{l},\vec{b})\in \mathrm{Adm}_k} d\left(\beta_{f_{(\vec{l},\vec{b})}}, \beta_{g_{(\vec{l},\vec{b})}}\right).$$

Multidimensional matching distance

Definition

Let $\beta_{\vec{f}}$ et $\beta_{\vec{g}}$ be two k-dimensional persistent Betti numbers for $\vec{f}: X \to \mathbb{R}^k$ et $\vec{g}: Y \to \mathbb{R}^k$. The *$k$-dimensional matching distance* between $\beta_{\vec{f}}$ et $\beta_{\vec{g}}$ is defined as

$$D(\beta_{\vec{f}}, \beta_{\vec{g}}) = \max_{(\vec{l},\vec{b})\in \mathrm{Adm}_k} d\left(\beta_{f_{(\vec{l},\vec{b})}}, \beta_{g_{(\vec{l},\vec{b})}}\right).$$

This matching distance is also stable: for $\vec{f}, \vec{g}: X \to \mathbb{R}^k$,

$$D(\beta_{\vec{f}}, \beta_{\vec{g}}) \leq \|\vec{f} - \vec{g}\|$$

where $\|\vec{h}\| = \max_{x\in X} \max_{i=1,\dots,k} |h_i(x)|$.

Advantages

Multidimensional persistence allows the extraction of more information from a shape than ordinary persistence.

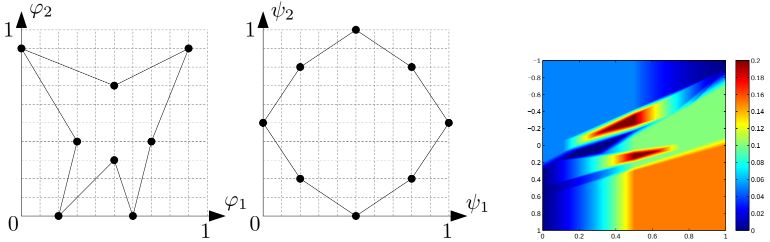

It also allows one to distinguish and compare noisy images, objects sampled by point clouds, and fuzzy sets.

Image retrieval tolerant to domain perturbation

Let X be a topological space, $K, K' \subset X$, and $\varphi : K \to \mathbb{R}^n$, $\varphi' : K' \to \mathbb{R}^n$ continuous filtering functions. If they represent point clouds or images subjected to noise, K and K' may differ in topology, making their comparison by means of persistent homology more problematic.

Image retrieval tolerant to domain perturbation

Let X be a topological space, $K, K' \subset X$, and $\varphi : K \to \mathbb{R}^n$, $\varphi' : K' \to \mathbb{R}^n$ continuous filtering functions. If they represent point clouds or images subjected to noise, K and K' may differ in topology, making their comparison by means of persistent homology more problematic.

However, extending φ, φ' so that they take all X as their domain, and substituting the sets K, K' with appropriate functions $f_K, f_{K'} : X \to \mathbb{R}$ so that perturbations of the sets become perturbations of these functions, we can then use persistence to compare the functions $\Phi = (f_K, \varphi) : X \to \mathbb{R}^{n+1}$ and $\Phi' = (f_{K'}, \varphi') : X \to \mathbb{R}^{n+1}$.

Examples of perturbed domains

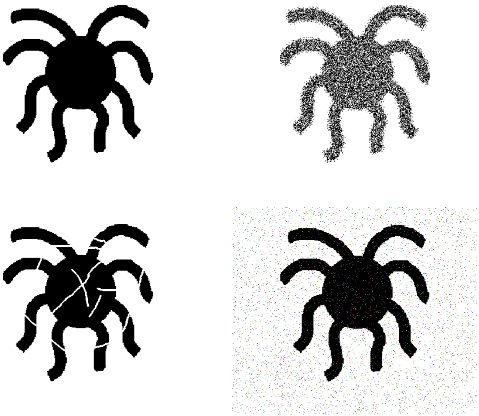

Four binary images of an octopus. Last three correspond to the first one subjected to different kinds of noise.

Choice of set distance function

The choice of f_K depends on what deformation is expected. For small perturbations, sets are comparable using the Hausdorff distance, and we take as f_K the distance from K (in any norm). In presence of outlying points, sets can be compared using the symmetric difference pseudometric, in which case f_K is taken as χ_K convolved with a ball.

Choice of set distance function

The choice of f_K depends on what deformation is expected. For small perturbations, sets are comparable using the Hausdorff distance, and we take as f_K the distance from K (in any norm). In presence of outlying points, sets can be compared using the symmetric difference pseudometric, in which case f_K is taken as χ_K convolved with a ball.
See *Persistent Betti numbers for a noise tolerant shape-based approach to image retrieval*, P. Frosini and C. Landi (2012) for details.

Bidimensional persistence

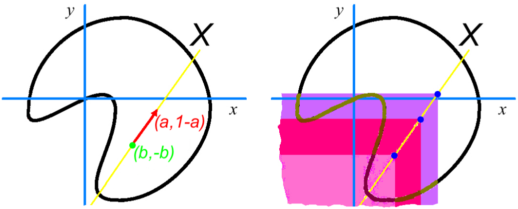

In the case of a bifiltration given by $\varphi = (\varphi_1, \varphi_2) : X \to \mathbb{R}^2$, we consider the filtration along a line $r_{(a,b)}$ defined by a unit vector (in $\|\cdot\|_\infty$) $(a, 1-a)$ and a point $(b, -b)$.

Bidimensional persistence

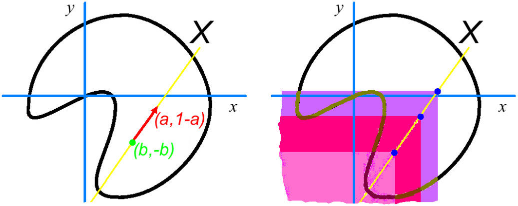

In the case of a bifiltration given by $\varphi = (\varphi_1, \varphi_2) : X \to \mathbb{R}^2$, we consider the filtration along a line $r_{(a,b)}$ defined by a unit vector (in $\|\cdot\|_\infty$) $(a, 1-a)$ and a point $(b, -b)$.
Each such choice of line defines a persistence diagram $\mathrm{Dgm}(\varphi_{a,b})$.

Bidimensional persistence

Formally, we consider the filtering function $\varphi_{a,b} : X \to \mathbb{R}$ where

$$\varphi_{a,b}(x) = \min\{a, 1-a\} \max\left\{\frac{\varphi_1(x) - b}{a}, \frac{\varphi_2(x) + b}{1-a}\right\}.$$

The normalization by $\min\{a, 1-a\}$ is important to ensure stability of the multidimensional matching distance.

Multidimensional matching distance

The (classical) *bidimensional matching distance* between $\varphi : X \to \mathbb{R}^n$, $\psi : Y \to \mathbb{R}^n$ is defined as

$$D(\varphi, \psi) = \sup_{(a,b)} d\left(\mathrm{Dgm}(\varphi_{a,b}), \mathrm{Dgm}(\psi_{a,b})\right)$$

where d is the one-dimensional bottleneck distance.

Multidimensional matching distance

The (classical) *bidimensional matching distance* between $\varphi : X \to \mathbb{R}^n$, $\psi : Y \to \mathbb{R}^n$ is defined as

$$D(\varphi, \psi) = \sup_{(a,b)} d\left(\mathrm{Dgm}(\varphi_{a,b}), \mathrm{Dgm}(\psi_{a,b})\right)$$

where d is the one-dimensional bottleneck distance.
This distance is stable :

$$D(\varphi, \psi) \leq \|\varphi - \psi\|_\infty.$$

Interesting results

This framework has been established for several years. In the course of our experiments with the computation of the bidimensional matching distance, we noticed that the supremum over (a, b) always seemed to be reached at a point where $a = 1/2$, that is, a line $r_{a,b}$ of slope 1.

Results

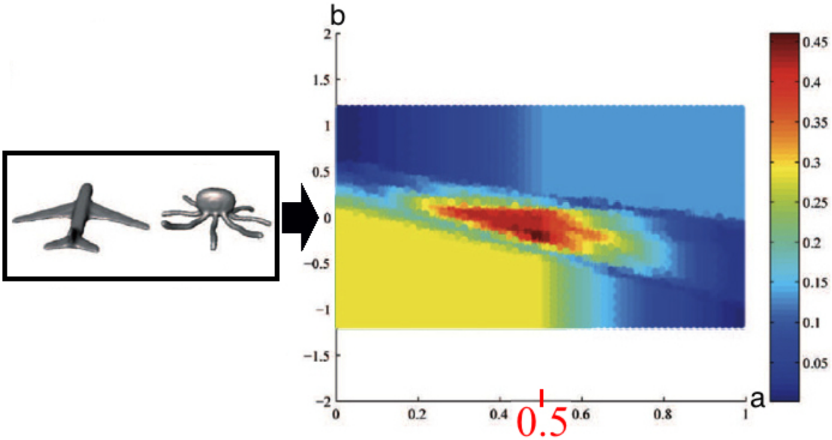
b
2
1.5
1
0.5
0
−0.5
−1
−1.5
−2
0
0.2
0.4
0.5
0.6
0.8
1
a
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05

Results

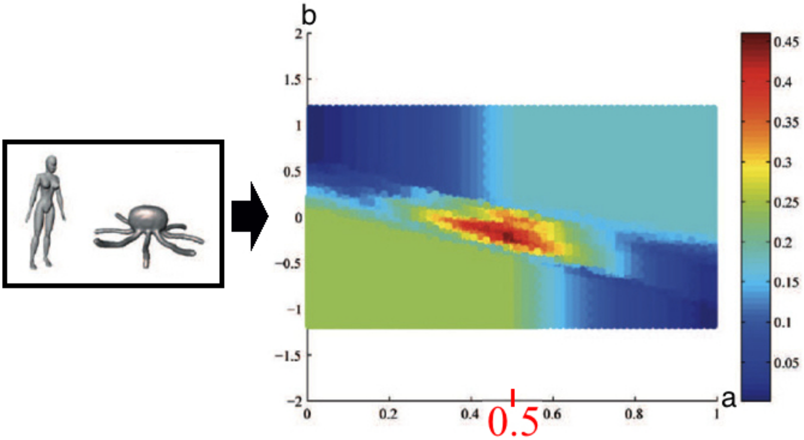
b
2
1.5
1
0.5
0
−0.5
−1
−1.5
−2
0
0.2
0.4
0.5
0.6
0.8
1
a
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05

Results

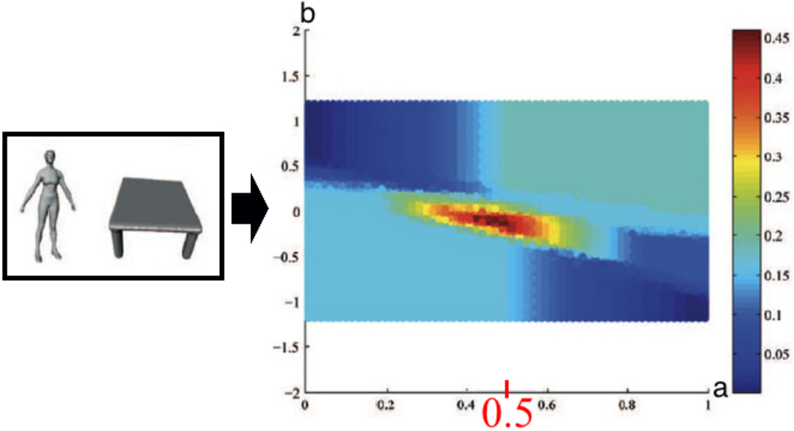
b
2
1.5
1
0.5
0
−0.5
−1
−1.5
−2
0
0.2
0.4
0.5
0.6
0.8
1
a
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05

Results

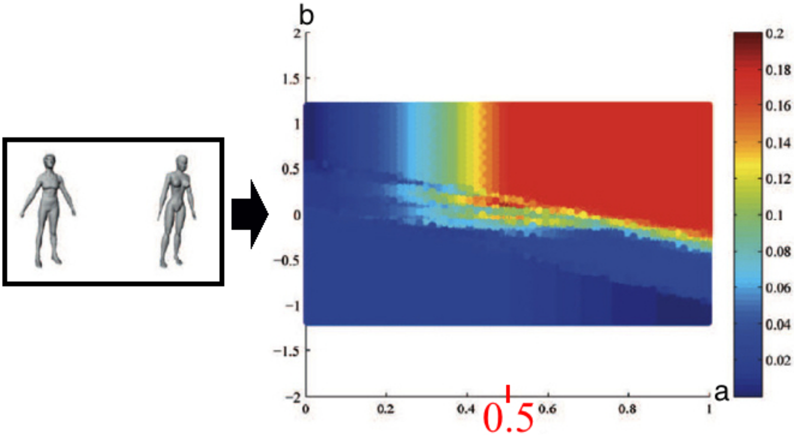
b
2
1.5
1
0.5
0
−0.5
−1
−1.5
−2
0
0.2
0.4
0.5
0.6
0.8
1
a
0.2
0.18
0.16
0.14
0.12
0.1
0.08
0.06
0.04
0.02

Conjecture

It seems to be the case that we should be able to compute $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$ only for $a = 1/2$.

Conjecture

It seems to be the case that we should be able to compute $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$ only for $a = 1/2$.
This is our conjecture:

$$\begin{aligned} D(\varphi,\psi) &= \sup_{(a,b)} d\left(\mathrm{Dgm}(\varphi_{a,b}), \mathrm{Dgm}(\psi_{a,b})\right) \\ &= \sup_{b} d\left(\mathrm{Dgm}(\varphi_{1/2,b}), \mathrm{Dgm}(\psi_{1/2,b})\right) \end{aligned}$$

Conjecture

It seems to be the case that we should be able to compute $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$ only for $a = 1/2$.
This is our conjecture:

$$D(\varphi,\psi) = \sup_{(a,b)} d\left(\mathrm{Dgm}(\varphi_{a,b}), \mathrm{Dgm}(\psi_{a,b})\right)$$
$$= \sup_{b} d\left(\mathrm{Dgm}(\varphi_{1/2,b}), \mathrm{Dgm}(\psi_{1/2,b})\right)$$

However, this conjecture showed itself hard to prove.

Coherent matching

In occurred to us that the way the classical bidimensional matching distance is defined might not be the most natural.

Coherent matching

In occurred to us that the way the classical bidimensional matching distance is defined might not be the most natural.
Consider the two parametrized sets $A_t = \{0, 1\}$ and $B_t = \{t, 1 - t\}$, $t \in [0, 1]$. The classical matching distance between A_t and B_t equals $1/2$, reached at $t = 1/2$.

Coherent matching

In occurred to us that the way the classical bidimensional matching distance is defined might not be the most natural.
Consider the two parametrized sets $A_t = \{0, 1\}$ and $B_t = \{t, 1-t\}$, $t \in [0, 1]$. The classical matching distance between A_t and B_t equals $1/2$, reached at $t = 1/2$.
However, there are two *coherent matchings* between A_t and B_t, and both actually have a cost of 1.

Coherent matching distance

To define a coherent matching distance between $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$, we thought of fixing $(\bar{a}, \bar{b})$ and a matching between $\mathrm{Dgm}(\varphi_{\bar{a},\bar{b}})$ and $\mathrm{Dgm}(\psi_{\bar{a},\bar{b}})$. We would then use the stability of normalized persistence diagrams to follow the points over a path c going from $(\bar{a}, \bar{b})$ to any (a, b).

First problem : tracking points

The first problem is that if at parameter values (a, b), $\mathrm{Dgm}(\varphi_{a,b})$ has points of multiplicity greater than 1, we cannot follow them further since they lose their identification.

First problem : tracking points

The first problem is that if at parameter values (a, b), $\mathrm{Dgm}(\varphi_{a,b})$ has points of multiplicity greater than 1, we cannot follow them further since they lose their identification.
A solution is to call such parameter values *singular* and exclude them from computation, that is, limit ourselves to the *regular* parameter values:

$$\mathrm{Reg}(\varphi) = \{(a, b) | \mathrm{Dgm}(\varphi_{a,b}) \text{ does not contain multiple points}\}.$$

Monodromy

However, even while moving only along paths $c : [0,1] \to \mathrm{Reg}(\varphi) \cap \mathrm{Reg}(\psi)$, coherent matchings between $\mathrm{Dgm}(\varphi_{c(t)})$ and $\mathrm{Dgm}(\psi_{c(t)})$ do not depend only on $c(0)$ and $c(1)$, but on the homotopy class of c relative to $c(0)$ and $c(1)$.

Monodromy

However, even while moving only along paths $c : [0, 1] \to \mathrm{Reg}(\varphi) \cap \mathrm{Reg}(\psi)$, coherent matchings between $\mathrm{Dgm}(\varphi_{c(t)})$ and $\mathrm{Dgm}(\psi_{c(t)})$ do not depend only on $c(0)$ and $c(1)$, but on the homotopy class of c relative to $c(0)$ and $c(1)$.
This is what we refer to as the *monodromy phenomenon* in bidimensional persistent homology.

Example in the visualizer

Example (Nontrivial monodromy)

Consider the function $\varphi : \mathbb{R}^2 \to \mathbb{R}^2$ *defined on the plane in the following way:* $\varphi_1(x, y) = x$, *and*

$$\varphi_2(x, y) = \begin{cases} -x & \text{if } y = 0 \\ -x + 1 & \text{if } y = 1 \\ -2x & \text{if } y = 2 \\ -2x + \frac{5}{4} & \text{if } y = 3 \end{cases},$$

$\varphi_2(x, y)$ *then being extended linearly for every* x *on the segment joining* $(x, 0)$ *with* $(x, 1)$, $(x, 1)$ *with* $(x, 2)$, *and* $(x, 2)$ *to* $(x, 3)$. *On the half-lines* $\{(x, y) \in \mathbb{R}^2 | y < 0\}$ *and* $\{(x, y) \in \mathbb{R}^2 | y > 3\}$, φ_2 *is then being taken with constant slope* -1 *in the variable* y.

Example in the visualizer

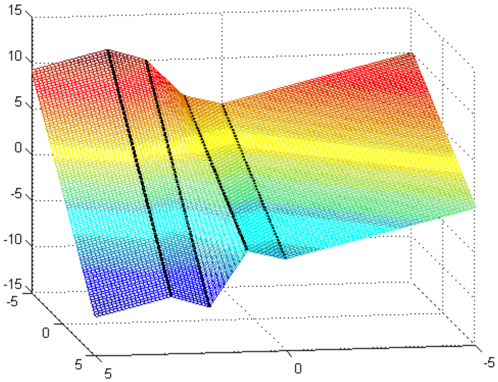

Figure: Function φ_2 of previous example. Depth is x, width is y.

Monodromy

Let $p : \widetilde{X} \to X$ a covering map onto the topological space X, and let $x \in X$. In algebraic topology, we refer to as *monodromy* the phenomenon by which, for a loop $\gamma : I \to X$ where $\gamma(0) = \gamma(1) = x$, and for $\tilde{x} \in p^{-1}(x)$ an element of the fibre of x, the associated continuous path $\tilde{\gamma}$ such that $\tilde{\gamma}(0) = \tilde{x}$ and $p \circ \tilde{\gamma} = \gamma$ might not be such that $\tilde{\gamma}(1) = \tilde{x}$.

Monodromy

Let $p : \widetilde{X} \to X$ a covering map onto the topological space X, and let $x \in X$. In algebraic topology, we refer to as *monodromy* the phenomenon by which, for a loop $\gamma : I \to X$ where $\gamma(0) = \gamma(1) = x$, and for $\tilde{x} \in p^{-1}(x)$ an element of the fibre of x, the associated continuous path $\tilde{\gamma}$ such that $\tilde{\gamma}(0) = \tilde{x}$ and $p \circ \tilde{\gamma} = \gamma$ might not be such that $\tilde{\gamma}(1) = \tilde{x}$.

In other words, as we turn around a singularity, it may be necessary to define applications on the cover $\widetilde{X}$ of X in order to guarantee their continuity.

Genericity condition

To alleviate the difficulty of monodromy, we will first assume that the sets of singular pairs for φ and ψ are discrete. We can then also assume that no parameter values are such that $\mathrm{Dgm}(\varphi_{a,b})$ or $\mathrm{Dgm}(\psi_{a,b})$ have points of multiplicity strictly greater than 2.

Genericity condition

To alleviate the difficulty of monodromy, we will first assume that the sets of singular pairs for φ and ψ are discrete. We can then also assume that no parameter values are such that $\mathrm{Dgm}(\varphi_{a,b})$ or $\mathrm{Dgm}(\psi_{a,b})$ have points of multiplicity strictly greater than 2. This ensures that loops in $\mathrm{Reg}(\varphi) \cap \mathrm{Reg}(\psi)$ can be factored into loops that induce only transpositions on $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$.

Simplifying assumptions

In our research, we have assumed that the topological space X is a manifold M homeomorphic to the m-sphere S^m, with $m \geq 2$. This ensures that there is a single point at infinity in $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$ for homology degree 0 and m, and none for other homology degrees.

Simplifying assumptions

In our research, we have assumed that the topological space X is a manifold M homeomorphic to the m-sphere S^m, with $m \geq 2$. This ensures that there is a single point at infinity in $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$ for homology degree 0 and m, and none for other homology degrees.

We also assume, for technical reasons, that a real value $\epsilon > 0$ exists such that if two proper points P_1, P_2 of $\mathrm{Dgm}(\varphi(\bar{a}, \bar{b}))$ have Euclidean distance less than ϵ from the diagonal $\Delta := \{(u, v) \in \mathbb{R}^2 | u = v\}$, then the Euclidean distance between P_1 and P_2 is not smaller than ϵ, for all regular (a, b), and that the same is also true for $\mathrm{Dgm}(\psi(\bar{a}, \bar{b}))$.

Transporting a matching

Given a homotopy $f_\tau = H(\tau, \cdot)$ between f and $g : M \to \mathbb{R}$, which in our case will be along the path c, we must specify what it means to transport a point $\bar{P}$ of $\mathrm{Dgm}(f)$ along this homotopy.

Transporting a matching

Given a homotopy $f_\tau = H(\tau, \cdot)$ between f and $g : M \to \mathbb{R}$, which in our case will be along the path c, we must specify what it means to transport a point $\bar{P}$ of $\mathrm{Dgm}(f)$ along this homotopy.
If $P(\tau)$ is a path, we call it *admissible* for H if

1. $P(\tau)$ belongs to $\mathrm{Dgm}(f_\tau)$ for every $\tau \in [0, \bar{\tau}]$;
2. $P(\tau)$ meets Δ at a finite number of points;
3. $P(\tau)$ "does not stop" at any point of Δ if it can "move on" in the set $\Delta^+ := \{(u, v) \in \mathbb{R}^2 | u < v\}$.

Transporting a matching

Proposition

Let $H(\tau, \cdot)$ be a homotopy between two continuous functions $f, g : M \to \mathbb{R}$.
For every point $\bar{P}$ that belongs to $Dgm(f)$ and has multiplicity 1, an $\epsilon > 0$ and a unique path $P : [0, \epsilon] \to \Delta^+ \cup \Delta$ exist, such that $P(0) = \bar{P}$ and the path $P(\tau)$ is admissible for the restriction of $H(\tau, \cdot)$ to the set $[0, \epsilon]$.

This means that we can follow the points in $\mathrm{Dgm}(f)$ along H.

Transporting a matching along a path

In our case, what we need is

Proposition

Let $c : [0,1] \to \mathit{Reg}(\varphi)$ be a continuous path with $c(0) = (a, b)$. For every proper point $\bar{P} \in \mathit{Dgm}(\varphi_{a,b})$, a unique path $P : [0,1] \to \Delta^+ \cup \Delta$ admissible for c exists, such that $P(0) = \bar{P}$.

We say that c transports $\bar{P}$ to $P(1)$ with respect to φ.

Transporting a matching along a path

Let $\sigma_{a,b}$ be a matching between $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$, with $(a,b) \in \mathrm{Reg}(\varphi) \cap \mathrm{Reg}(\psi)$. We can naturally associate to $\sigma_{a,b}$ a matching $\sigma_{c(1)} : \mathrm{Dgm}(\varphi_{c(1)}) \to \mathrm{Dgm}(\psi_{c(1)})$. We set $\sigma_{c(1)}(P') = Q'$ if and only if c transports $\bar{P}$ to P' with respect to φ and $\bar{Q}$ to Q' with respect to ψ. We also say that c *transports* $\sigma_{a,b}$ *to* $\sigma_{c(1)}$ *along* c *with respect to the pair* (φ, ψ).

Transporting a matching along a path

Let $\sigma_{a,b}$ be a matching between $\mathrm{Dgm}(\varphi_{a,b})$ and $\mathrm{Dgm}(\psi_{a,b})$, with $(a,b) \in \mathrm{Reg}(\varphi) \cap \mathrm{Reg}(\psi)$. We can naturally associate to $\sigma_{a,b}$ a matching $\sigma_{c(1)} : \mathrm{Dgm}(\varphi_{c(1)}) \to \mathrm{Dgm}(\psi_{c(1)})$. We set $\sigma_{c(1)}(P') = Q'$ if and only if c transports $\bar{P}$ to P' with respect to φ and $\bar{Q}$ to Q' with respect to ψ. We also say that c *transports* $\sigma_{a,b}$ *to* $\sigma_{c(1)}$ *along* c *with respect to the pair* (φ, ψ).
We are now ready to introduce the coherent 2D matching distance.

Coherent 2D matching distance

Definition

Fix $(a,b) \in \mathit{Reg}(\varphi) \cap \mathit{Reg}(\psi)$*. Let* Γ *be the set of all continuous paths* $c : [0,1] \to \mathit{Reg}(\varphi) \cap \mathit{Reg}(\psi)$ *with* $c(0) = (a,b)$*. Let* S *be the set of all matchings* $\sigma : \mathit{Dgm}(\varphi_{c(0)}) \to \mathit{Dgm}(\psi_{c(0)})$*. For every* $\sigma \in S$ *and every* $c \in \Gamma$*, let* $T_c^{(\varphi,\psi)}(\sigma)$ *be the matching obtained by transporting* σ *along* c *with respect to the pair* (φ,ψ)*. We define the coherent 2D matching distance* $CD_{match}(\varphi,\psi)$ *as*

$$CD_{match}(\varphi,\psi) = \max\left\{\min_{\sigma\in S}\sup_{c\in\Gamma} cost\left(T_c^{(\varphi,\psi)}(\sigma)\right), \gamma_\infty\right\},$$

where γ_∞ *is the maximum varying* (a,b) *of the distance between the point at infinity of* $\mathit{Dgm}(\varphi_{a,b})$ *and the point at infinity of* $\mathit{Dgm}(\psi_{a,b})$ *for degrees* 0 *and* m*, and* 0 *for the other degrees.*

Properties of CD_{match}

Proposition

The definition of $CD_{match}(\varphi, \psi)$ does not depend on the choice of the point $(a, b) \in Reg(\varphi) \cap Reg(\psi)$.

Proposition

$CD_{match}(\varphi, \psi)$ is a pseudo-distance.

Theorem

$D_{match}(\varphi, \psi) \leq CD_{match}(\varphi, \psi) \leq \|\varphi - \psi\|_\infty$.