Geometric dissimilarities for shape matching

Joan Alexis Glaunès

Summer school "Geometry and data"
Part 2

Matching functional for discrete problems

- For a matching functional the optimal trajectories must follow geodesics. So the optimal vector fields v_t depend only on the initial momentum vectors $p(0)$. So we rewrite the functional as

$$J(p(0)) = \gamma \langle K_V(q(0), q(0))p(0) \, , \, p(0) \rangle + A(q(1))$$

  where $p(t)$ and $q(t)$ are constrained to follow the geodesic equations.
- The geodesic shooting algorithm consists in optimizing this function J. It is done via gradient descent or more advanced optimization techniques (e.g. LBFGS).

Data attachment terms : landmarks

- **Landmark matching** : assume correspondences between discretization points $(x_i)_{1\leq i\leq n}$ and $(y_i)_{1\leq i\leq n}$ of the two shapes are known (each x_i must get close to y_i).

$$\Rightarrow \text{ use } A(q(1)) = \|q(1) - y\|^2 = \sum_{i=1}^{n} \|q_i(1) - y_i\|^2.$$

Curves as measure or currents

- Let C be a curve in $\Omega \subset \mathrm{R}^d$, parametrized by $\gamma_C : [0,1] \to \Omega$. The **uniform measure** associated to C is the following linear form, defined by its action on test **functions** $f : \Omega \to \mathrm{R}$:

$$\mu_C(f) = \int_0^1 f(\gamma_C(s)) \|\gamma_C'(s)\| ds.$$

- The **current** associated to C is the following linear form, defined by its action on test **1-forms** $\omega : \Omega \to (\mathrm{R}^d)^*$:

$$\vec{\mu}_C(\omega) = \int_0^1 \langle \omega(\gamma_C(s)) \mid \gamma_C'(s) \rangle \, ds.$$

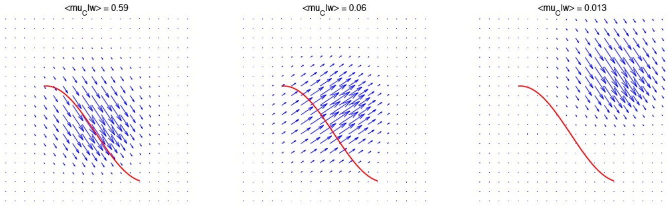

Submanifolds as currents

Let S be a regular (rectifiable), oriented and bounded m-submanifold in Ω

- The **uniform measure** μ_S is defined for every test function f as :

$$\mu_S(f) = \int_S f(x) d\mathcal{H}^m(x).$$

- The **current** $\vec{\mu}_S$ is defined for every test m-form ω by:

$$\vec{\mu}_S(\omega) = \int_S \langle \omega(x) \mid \vec{\tau}_S(x) \rangle \ d\mathcal{H}^m(x)$$

where $\tau_S\vec{(}x)$ is a unit m-vector associated to the tangent space at x of S.

Hilbert norms on measures and currents

- **Dual norms.** We consider scalar measures μ (resp. currents $\vec{\mu}$) as elements of a Hilbert space H^* (resp. W^*) which is dual to a space H of regular functions (resp. a space W of regular m-forms) on R^d.
- These dual norms give our measure of dissimilarity between curves:

$$\mathcal{A}(\phi(S), T) = \|\mu_{\phi(S)} - \mu_T\|^2_{H^*},$$

or

$$\mathcal{A}(\phi(S), T) = \|\vec{\mu}_{\phi(S)} - \vec{\mu}_T\|^2_{W^*},$$

Data attachment term for discrete measures

- We assume that source and target measures are combinations of Dirac functionals :

$$\mu_x = \frac{1}{n}\sum_{i=1}^{n}\delta_{x_i}, \quad \mu_y = \frac{1}{m}\sum_{j=1}^{m}\delta_{y_j},$$

- This allows to match two sets of points $(x_i)_{1\leq i\leq n}$, $(y_j)_{1\leq j\leq m}$, without any knowledge of correspondences between points (as opposed to the landmarks case), and with possibly different number of points.
- We evaluate the distance between these measures using the dual RKHS norm :

$$A(q(1)) = \|\mu_{q(1)} - \mu_y\|_{H^*}^2.$$

Expanding this squared norm and using the reproducing formula, we get

$$A(q(1)) = \|\mu_{q(1)}\|_{H^*}^2 - 2\langle \mu_{q(1)}, \mu_y\rangle_{H^*} + \|\mu_y\|_{H^*}^2$$

$$= \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n}K_H(q_i(1), q_j(1)) - 2\frac{1}{nm}\sum_{i=1}^{n}\sum_{j=1}^{m}K_H(q_i(1), y_j) + \frac{1}{m^2}\sum_{i=1}^{m}\sum_{j=1}^{m}K_H(y_i, y_j)$$

Data attachment terms for discrete currents

- Meshes are approximated in the space of currents as sums of vectorial Diracs:

$$\vec{\mu}_S = \sum_{i=1}^{n} \delta_{c_i^S}^{\eta_i^S}, \quad \vec{\mu}_T = \sum_{j=1}^{m} \delta_{c_j^T}^{\eta_j^T}.$$

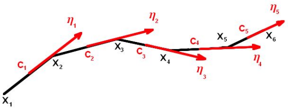

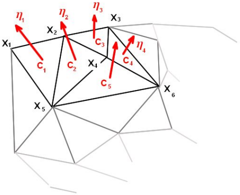

$$A(q(1)) = \|\vec{\mu}_S - \vec{\mu}_T\|_{W^*}^2 = \sum_{i=1}^{n}\sum_{i'=1}^{n} \left\langle \eta_i^S \, , K_W(c_i^S, c_{i'}^S)\eta_{i'}^S \right\rangle$$

$$-2\sum_{i=1}^{n}\sum_{j=1}^{m} \left\langle \eta_i^S \, , K_W(c_i^S, c_j^T)\eta_j^T \right\rangle + \sum_{j=1}^{m}\sum_{j'=1}^{m} \left\langle \eta_j^T \, , K_W(c_j^T, c_{j'}^T)\eta_{j'}^T \right\rangle$$

Diffeomorphic matchings via currents

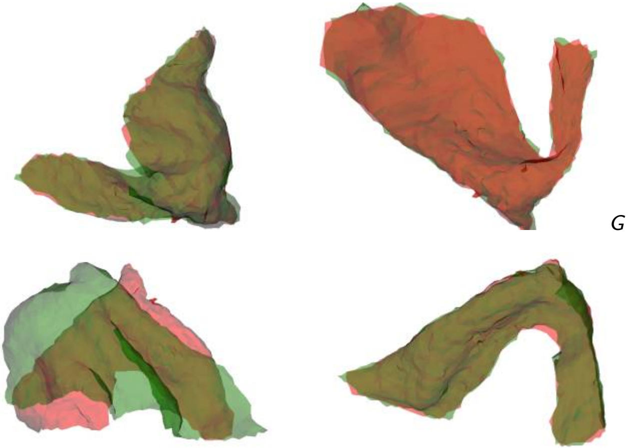

Diffeomorphic matchings via currents

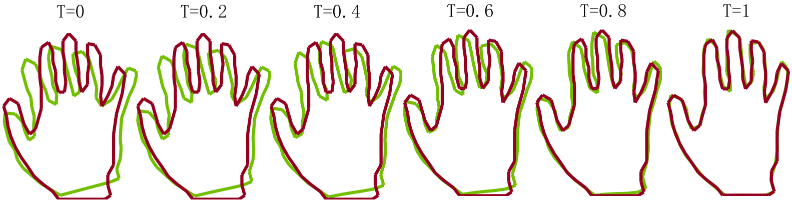

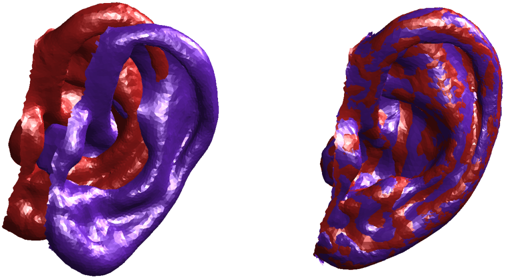

Outline

Outline

Measures and currents : properties

- Both models can handle changes in topology between shapes (e.g. one can compare and match a closed curve to an open one)
- The currents model is a priori more complete since it encodes both location and tangential information of the curves. One may think about it as a first-order model, while the measure model is zero-order.
- As a counterpart currents require to define an orientation on each curve, and on each subpart of the curve when one has to deal with disconnected or branching curves.
- Due to this orientation sensitivity, specific parts like spikes in curves are filtered out in the currents model. Depending on the application this can be seen as a good or bad property.

Some practical issues with currents norms

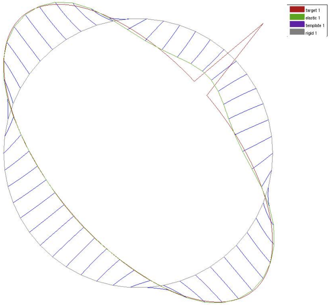

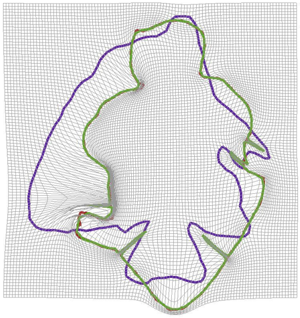

The varifolds model

- The **varifolds** model can be seen as an extension of the model of currents.
- We model shapes as linear forms over scalar functions on $\mathrm{R}^d \times \mathrm{S}^{d-1}$, then define a dual Hilbert norm on the varifolds space.
- Associated varifold :

$$\mu_S(\omega) := \int_S \omega(x, \tau_S(x)) d\mathcal{H}^m(x)$$

- dual norm :

$$\|\mu_S\|^2_{W^*} = \int_S \int_S k_e(x,y) k_t(\tau_S(x), \tau_S(y) d\mathcal{H}^m(y) d\mathcal{H}^m(x)$$

- dissimilarity :

Normal cycles : Tube formula and curvature measures

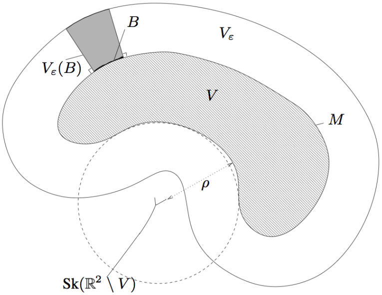

- For a set $V \in \mathrm{R}^d$ such that $M = \partial V$ is smooth, the volume of the ε-offset V_ε is a polynomial in ε which coefficients give integrals of curvatures of $M = \partial V$ when ∂V is smooth.
  - ex: in R^3,

$$Vol(V_\varepsilon) = Vol(V) + Area(M)\varepsilon + H(M)\frac{\varepsilon^2}{2} + G(M)\frac{\varepsilon^3}{3},$$

where $H(M)$ and $G(M)$ are the integrals of mean and Gauss curvatures.

Normal cycles : Tube formula and curvature measures

- This formula can be localized so that we get integrals of curvatures restricted to any Borel subset.
- If V is only assumed to be of positive reach, $Vol(V_\varepsilon)$ (and its localized version) is still a polynomial in ε; hence its coefficients define curvature measures in this general setting.

Ref for this part : P. Roussillon, J. Glaunès : Kernel Metrics on Normal Cycles and Application to Curve Matching. SIAM Journal on Imaging Sciences. 2016

Definitions

- ε-**offset** around a compact set $C \subset \mathrm{R}^d$: $C_\varepsilon = \{x \in \mathrm{R}^d, d(x, C) \leq \varepsilon\}$.

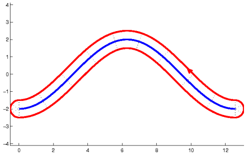

- **Normal cone** at $x \in C$:

$$\hat{\mathcal{N}}(C, x) = \{u \in \mathrm{R}^d, \exists \varepsilon > 0, \forall y \in C \cap B(x, \varepsilon), \langle x - y\ , u\rangle \leq 0\}.$$

- **Unit normal vectors** at $x \in C$: $\mathcal{N}(C, x) = \hat{\mathcal{N}}(C, x) \cap S^{d-1}$.

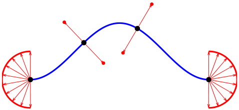

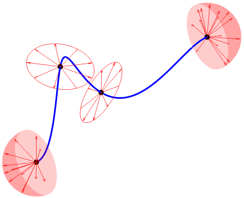

Definitions

- **Unit Normal bundle** associated to a set:

$$\mathcal{N}(C) = \{(x, \xi) \in C \times S^{d-1}, \xi \in \mathcal{N}(C, x)\}.$$

- Formally, we can see $\mathcal{N}(C)$ as the "derivative" of C_ε at $\varepsilon = 0$.
- $\mathcal{N}(C)$ is a closed sub-manifold of dimension $d-1$ in $\mathrm{R}^d \times S^{d-1}$.
- The **normal cycle** associated to C is the current $\vec{\mu}_{\mathcal{N}(C)}$ associated to $\mathcal{N}(C)$ (which is canonically oriented).

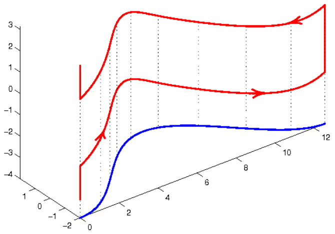

The addition formula

- For any subsets C_1, C_2, whenever it has sense,

$$\vec{\mu}_{\mathcal{N}(C_1 \cup C_2)} = \vec{\mu}_{\mathcal{N}(C_1)} + \vec{\mu}_{\mathcal{N}(C_2)} - \vec{\mu}_{\mathcal{N}(C_1 \cap C_2)}.$$

- This allows to extend the definition of normal cycles to any finite union of smooth curves (in fact to any finite union of sets of "positive reach")
- We can even define the normal cycle of a curve deprived of its end-points by simply substracting the normal cycles associated to them - which correspond to circles.

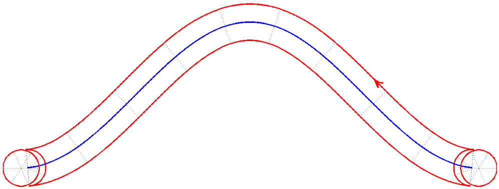

Properties

- The normal cycle is a second-order model; it encodes curvature information of the set. By computing specific integrals of the normal cycle over a small area, one gets the exact integrated values of the curvature of C on this area.
- The normal cycle does not depend on any choice of orientation on the curve, and there is no need to specify any,
- Since "spikes" are parts of high curvature; they get highly weighted in the model.
- Normal cycles are in fact a model for subsets of R^d and not for

  submanifolds of a specific dimension. Hence one can think about comparing a curve with a surface, or to model "hybrid" objects.

Designing Hilbert norms for normal cycles

- Since $\vec{\mu}_{\mathcal{N}(C)}$ is a current in the product space $\mathrm{R}^d \times S^{d-1}$, we need to define a kernel in $\mathrm{R}^d \times S^{d-1}$. This can be done by considering a product of two kernels:

$$k(\mathrm{x},\mathrm{y}) = k((x,u),(y,v)) = k_p(x,y)k_n(u,v),$$

where $k_p(x,y)$ is a reproducing kernel in R^d (e.g. $k_p(x,y) = \frac{1}{1+\|x-y\|^2/\sigma^2}$), and $k_n(u,v)$ is a reproducing kernel in S^{d-1} (e.g. the kernel given by a Sobolev metric on S^{d-1})

- Let $T(\mathrm{x}) = \tau_1(\mathrm{x}) \wedge \cdots \wedge \tau_{d-1}(\mathrm{x})$, where $(\tau_i(\mathrm{x}))_{1\leq i\leq d-1}$ is an orthonormal basis of the tangent space $T_{\mathrm{x}}\mathcal{N}(C)$ for any $\mathrm{x} \in \mathcal{N}(C)$. Then we have

$$\|\vec{\mu}_{\mathcal{N}(C)}\|^2_{W^*} = \int_{\mathcal{N}(C)}\int_{\mathcal{N}(C)} k(\mathrm{x},\mathrm{y}) \left\langle T(\mathrm{x})\,,\, T(\mathrm{y})\right\rangle \, d\sigma_{\mathcal{N}(C)}(\mathrm{x})\, d\sigma_{\mathcal{N}(C)}(\mathrm{y}),$$

where $d\sigma_{\mathcal{N}(C)}(\mathrm{x})$ is the volume element on the submanifold $\mathcal{N}(C)(\mathrm{x})$

Implementation for piecewise linear curves

- Let C be a piecewise linear curve, which we look at as a collection of segments which may be connected at their end-points.
- We can further decompose C as the disjoint union of open segments S_i and points P_j. The additive property for normal cycles then writes

$$\vec{\mu}_{\mathcal{N}(C)} = \sum_i \vec{\mu}_{\mathcal{N}(S_i)} + \sum_j \vec{\mu}_{\mathcal{N}(P_j)}.$$

- We decompose further again into space and angular components by writing each $\vec{\mu}_{\mathcal{N}(S_i)}$ as a sum of three terms. The tangent spaces of these space and angular components are orthogonal.

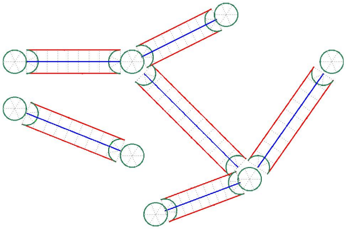

Implementation for piecewise linear curves

- Hence the whole squared dual norm of $\vec{\mu}_{\mathcal{N}(C)}$ can be computed as a sum of two parts, one involving only scalar products between "space" elements (located on edges) and the other involving only scalar products between "angular" elements (located on vertices).
- The "space" part of the metric is very similar to the usual metric on currents, except that it is an orientation-free representation of curves. In fact it corresponds exactly to the varifolds model (Charon et al, 2013). To compute the scalar product between two such elements we use the same approximation by vector-valued Dirac located at the center of each edge.
- For the angular part we need to compute double integrals of k_n over half-spheres in S^{d-1}; which can be computed either analytically (for $d = 2$) or via pre-computing look-up tables.

Experiments : currents metric

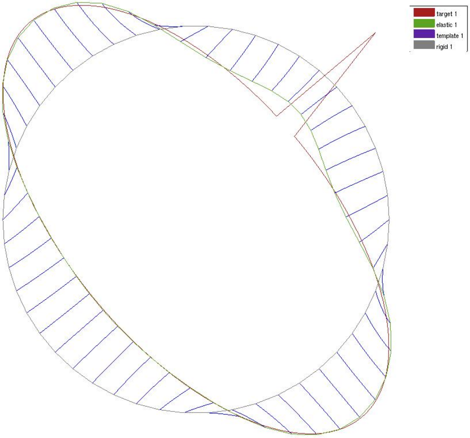

Experiment : normal cycle metric

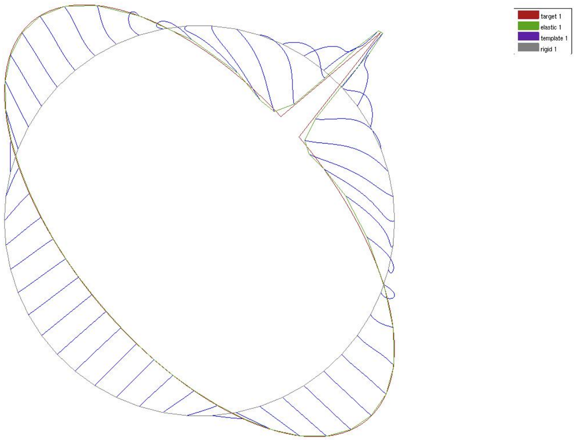

Experiment : currents metric

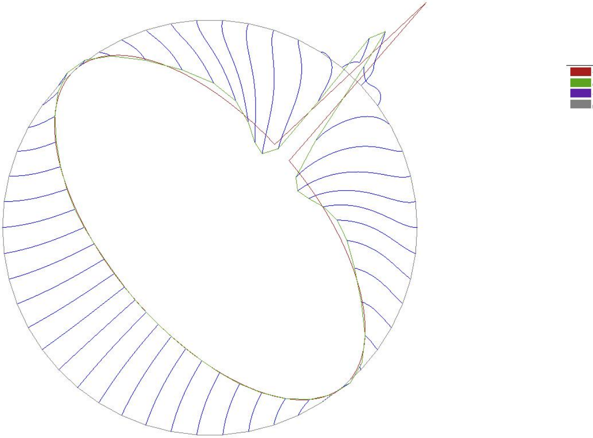

Experiment : normal cycle metric

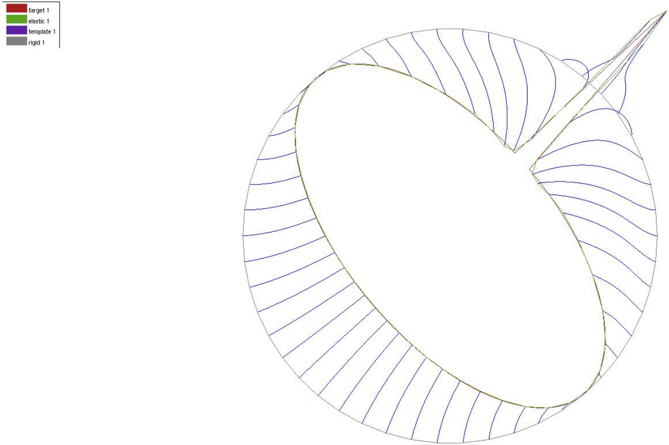

Experiments

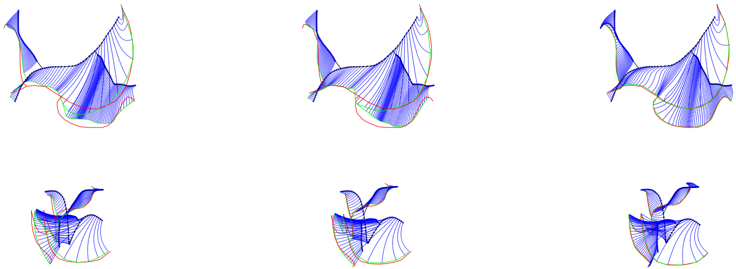

Experiment : normal cycles metric

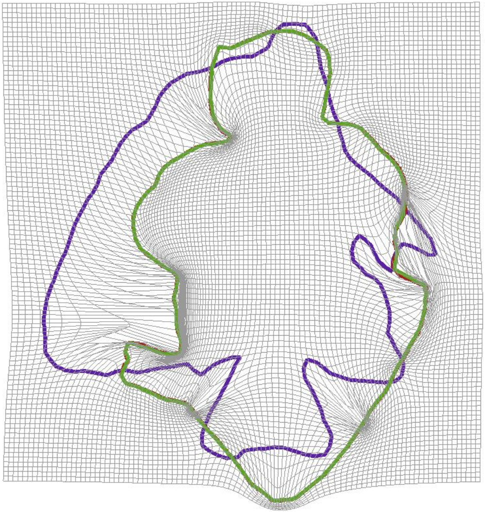

Other experiments

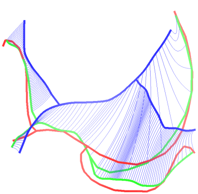

courants

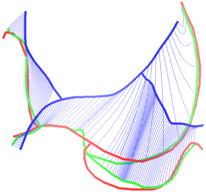

varifolds

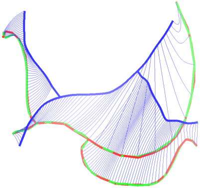

cycles normaux

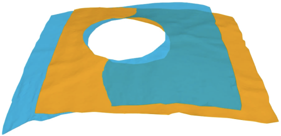

varifolds

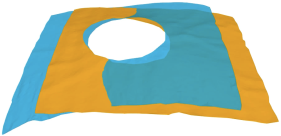

cycles normaux

The varifolds model

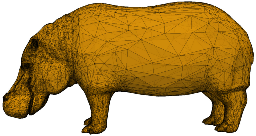

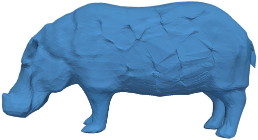

varifolds

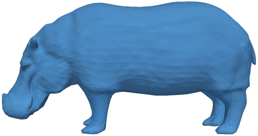

cycles normaux