Shape Warping and Statistics

Guillaume CHARPIAT [1], Olivier FAUGERAS [2], Renaud KERIVEN [3]

[1] Projet Odyssée, École Normale Supérieure, France

[2] Projet Odyssée, INRIA Sophia Antipolis, France

[3] Projet Odyssée, École Nationale des Ponts et Chaussées, France

Reference

G. Charpiat, O. Faugeras, and R. Keriven. Shape Metrics, Warping and Statistics. In *Proc. International Conference on Image Processing*, Barcelona, September 2003.

Introduction

The recognition of a shape in an image can become easier if statistics on the previous occurrences of the shape in other images are available and usable.
The set of all plane curves is far more complicated than the usual data sets in $\mathbb{R}^n$. This induces an intricate problem, the one of shape warping, and makes the expression of statistics on curves a little more difficult. We present a general method to compute good approximations of them, assuming the different shapes to compare have already been aligned with respect to the Euclidean group (by classical techniques).

Basic Concepts

- We first need a practical workspace, big enough to contain usual curves, and restricted enough to have good properties, for example the manifold

$$\mathcal{S} = \left\{ f \in \mathcal{C}^2(\mathbb{S}^1 \to \mathbb{R}^2) \text{ without double point; } \frac{d}{d\sigma}\left\|\frac{d}{d\sigma} f\right\| = 0 \right\}.$$

- This workspace should be endowed with an explicit distance, for example the Hausdorff distance :

$$d_H(C_1, C_2) = \max\left(\sup_{x\in C_1} \inf_{y \in C_2} d(x,y);\ \sup_{y\in C_2} \inf_{x \in C_1} d(x,y) \right).$$

- The tangent space T_A at each curve A of $\mathcal{S}$ should be endowed with a scalar product.

$$T_A \equiv \{ f \in \mathcal{C}^1(\mathbb{S}^1 \to \mathbb{R})\}; \qquad \langle f_1 | f_2 \rangle_{T_A} = \int_{\mathbb{S}^1} f_1\ f_2.$$

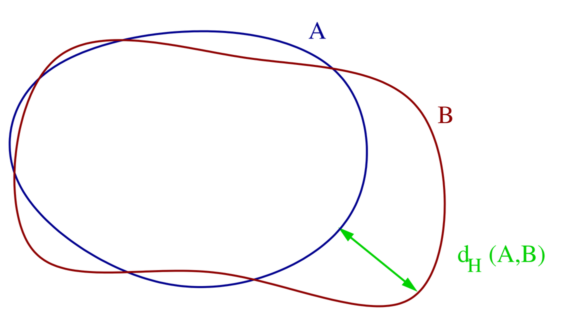

The Hausdorff distance between two curves.

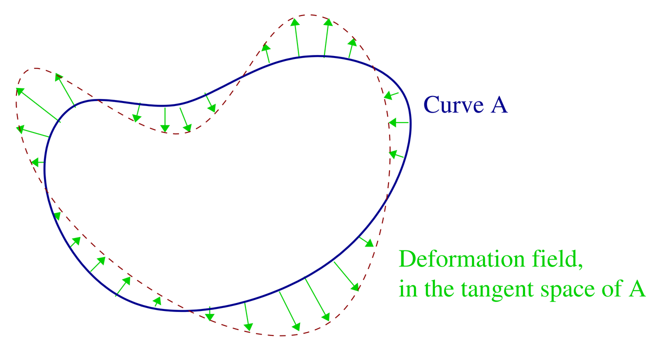

The tangent space of a curve is the set of its deformation fields.

Useful Tools for Warping

- A continuous warping between two shapes A and B of $\mathcal{S}$ is a geodesic path in $\mathcal{S}$, whose end points are A and B, that is the path in $\mathcal{S}$ with minimal length which links A to B. Thus, we would like to compute a gradient descent with respect to variable $X(t)$ on the distance $d_H(X(t), B)$ with initial condition $X(t) = A$, which yields the PDE $\partial_t X(t) = -\nabla_X d_H(X, B)$.

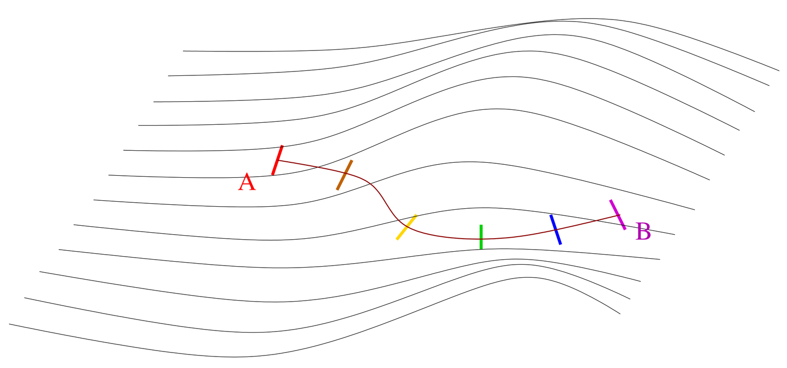

To minimize the distance, follow a geodesic path (each point of the manifold is a curve).

- Unfortunately, the gradient of such a distance is a sum of Dirac peaks, and such a gradient descent would immediately lead us out of the regular set of curves $\mathcal{S}$.
- We have to consider a computable, regular, derivable approximation of the distance. For the Hausdorff distance we consider :

$$E_H(A,B) = \left\langle \left\langle \langle d(\cdot,\cdot)\rangle_A^\phi \right\rangle_B^\psi, \left\langle \langle d(\cdot,\cdot)\rangle_B^\phi \right\rangle_A^\psi \right\rangle^\Psi.$$

where $\langle f \rangle_A^\phi = \phi^{-1}\left(\frac{1}{|A|}\int_A \phi \circ f\right)$ and $\langle a, b\rangle^\Psi = \Psi^{-1}\left(\frac{\Psi(a)}{2} + \frac{\Psi(b)}{2}\right)$. In order to avoid some kinds of local minima, we add to E_H a term depending on the length of the curves. The quality of this approximation can be as good as desired with a suitable choice of ϕ, ψ and Ψ.

Warping Computation

- We can now solve the PDE $\partial_t X(t) = -\nabla_X E_H(X, B)$ with initial condition $X = A$. Problems with local minima appear sometimes when the two curves to match are too different.

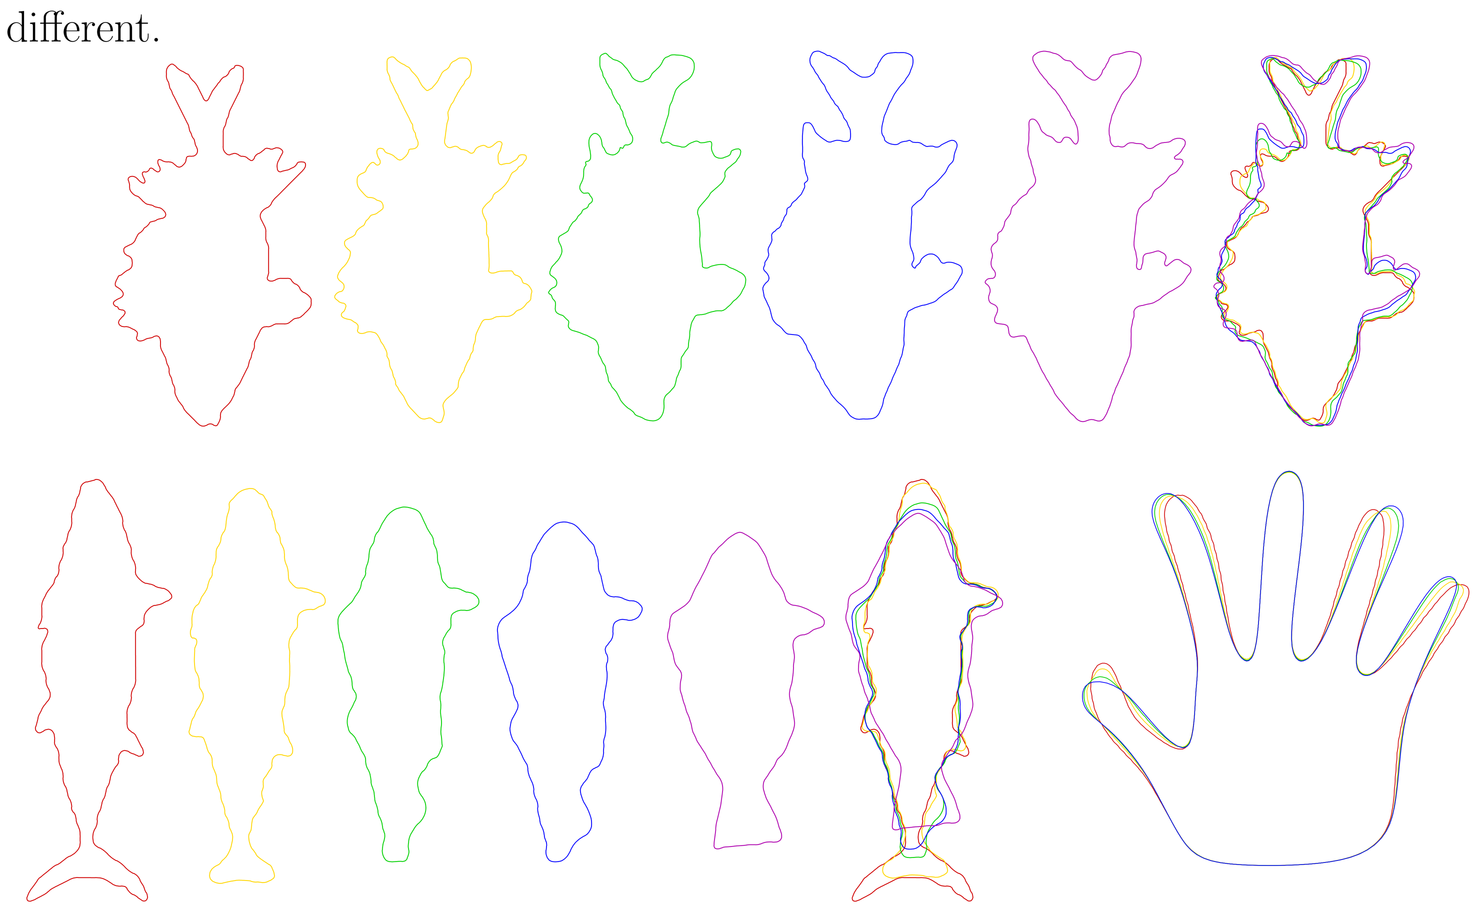

Three examples of continuous warping (from the red curve, onto the purple one).

- A mean M of n curves $(C_i)_{1\leqslant i \leqslant n}$ is then defined by being a curve of $\mathcal{S}$ which is as close as possible to all curves C_i, therefore minimizing $\sum_{i\leqslant n} E_H^2(M, C_i)$, which experimentaly appears to have a unique solution in common simple cases.

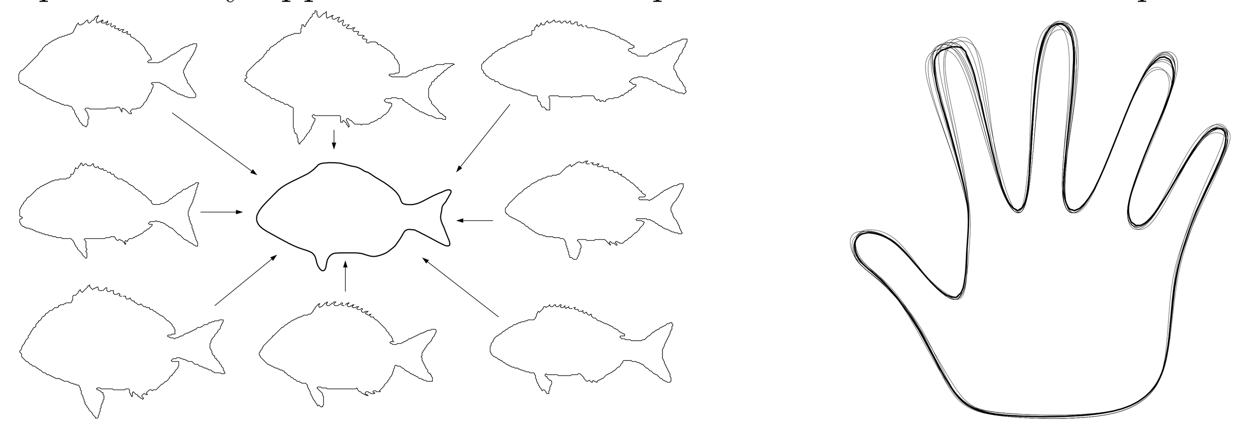

Two examples of the mean of several curves.

Warping and Statistics

- Let us note $\delta_i(A) = -\nabla_A E_H^2(A, C_i)$ the deformation field induced by C_i on the tangent T_A of the curve A, which would deform A so that A would draw nearer to C_i. Then the mean M is a special curve on which the mean deformation field $\frac{1}{n}\sum_i \delta_i(M)$ is null. The distribution of the fields $\delta_i(M)$ can be studied with classical techniques, by diagonalizing the covariance matrix Λ and extracting from it the principal modes β_k of variation of the deformation fields.

$$\Lambda_{i,j} = \langle \delta_i | \delta_j \rangle_{T_M}, \quad n \cdot n \text{ real symmetrical matrix.}$$

- Lastly, we can start from the mean M and track a mode β_k with an PDE in order to visualize it. The set (C_i) is not anymore described only by the mean curve, but also characterized by the principal modes of deformation of the set. A typical example of the set can be artificially constructed by following each mode β_k simultaneously with suitable coefficients.

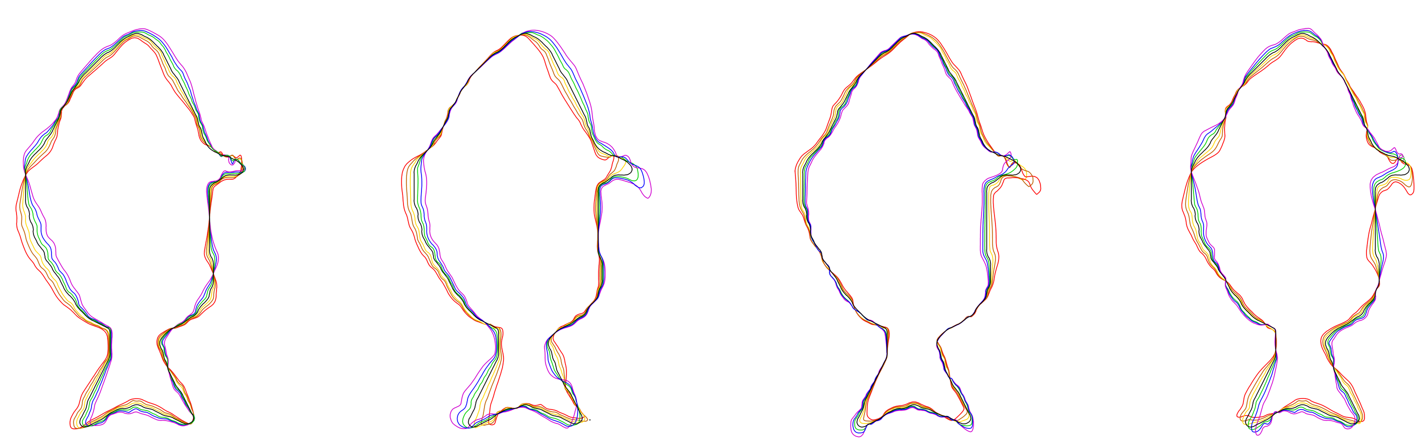

The first four modes of variation, for the same set of eight fishes.

Future Work

- The local minima are a major problem when the curves are quite different from one another and show large oscillations, like fingers in a hand. A possible solution might consist in adding other descriptors to the distance.
- Besides, the principal modes that have been observed are not always very satisfying, but it may be due only to the small number of samples.
- Finally, the notion of distance between any curve and a set of curves, which would be useful for classification, has still to be worked upon, but the track of considering histograms of the $\langle \beta_k | \delta_i \rangle$ seems to be promising.