Geodesic interpolating splines

Vincent Camion, Laurent Younes

CMLA, ENS de Cachan, CNRS UMR 0876
94235 Cachan CEDEX, France
email: younes@cmla.ens-cachan.fr
fax: 33 1 47 40 59 01
tel: 33 1 47 40 59 18

Abstract. We propose a simple and efficient method to interpolate landmark matching by a non-ambiguous mapping (a diffeomorphism). This method is based on spline interpolation, and on recent techniques developed for the estimation of flows of diffeomorphisms. Experimental results show interpolations of remarkable quality. Moreover, the method provides a Riemannian distance on sets of landmarks (with fixed cardinality), which can be defined intrinsically, without refering to diffeomorphisms. The numerical implementation is simple and efficient, based on an energy minimization by gradient descent. This opens important perspectives for shape analysis, applications in medical imaging, or computer graphics

1 Introduction

This paper proposes a new, efficient and consistent method for generating dense diffeomorphisms within an image from sparse information on the displacements of a finite number of points (landmarks). This is an important issue for image processing and computer graphics, and the problem has generated a large number of publications, starting with the seminal papers of Bookstein (see [3] and references therein). There are numerous applications: generating deformations from the position of control points is used, for example, to synthesize facial expressions, or to compute morphings; analyzing variations of shape has application in medical imaging or face recognition, matching is essential for the construction of anatomical atlases. Jointly with the purpose of interpolating from landmark-matching, comes the issue of measuring the discrepancy between two groups of matched landmarks. This is not an obvious problem, and it seems quite intuitive that the smoothness of the underlying, unobserved, global displacement comes as an essential part for the perceptive impression of discrepancy. A third, important, feature is the consistency of the interpolated displacement, in the sense that it should be one-to-one, ensuring that there cannot be two distinct parts of the original picture which are matched to the same zone in the target.

The method which is described here addresses the three problems simultaneously. It does provide a way to interpolate from landmark-matching, while providing a distance between configurations of landmarks which takes into account the smoothness of the underlying warping, generated as a diffeomorphism defined on the image grid. As a fourth, non-negligible property, comes the fact that this method is easy to implement, and numerically efficient.

To fix notations, let Ω be a bounded set in the plane. When $(x_1, \dots, x_N)$, $(y_1, \dots, y_N)$, two sets of N labeled landmarks in Ω, are given, we shall deal with the problem of finding a diffeomorphism g of Ω, with minimal size (in a sense to be defined), such that, for all i, $g(x_i) \simeq y_i$ (*inexact matching*).

The method which is developed in the sequel takes its roots from three main ideas:

- Interpolating splines, as pioneered by Bookstein in computer vision ([3]), and widely used to generate dense warpings from sparse information.
- Generation of diffeomorphisms as flows (solutions of an ODE), in a framework which guarantees smoothness and consistency, as in [10, 4]
- Computation of geodesic distances (minimal path length) on deformable data, as used in [11, 8].

The analysis results in a simple algorithm to compute diffeomorphisms from landmark data.

The paper is organized as follows. We start by reviewing the elements of spline theory which will be needed, and relate them to the (non-diffeomorphic) interpolation introduced by Bookstein. This forms the first ingredient for our method. In a second step, we give a presentation of the theory of groups of diffeomorphisms, generated as flows (solutions of ODEs) on a set Ω, and show how this framework can be used to generate geodesic distances on structures acted on by diffeomorphisms (ie. deformable patterns). The last step will be to use this framework on the very simple deformable structure which are sets of landmarks, to derive an efficient algorithm for simultaneously computing distances between sets of landmarks and generating a diffeomorphism to interpolate the pointwise matching. The paper ends with a presentation of some experimental facts and data.

2 Landmark matching and splines

2.1 Splines

Like for all landmark-matching methods, the numerical efficiency of the algorithm that we propose relies on spline interpolation theory. For completeness of the presentation, we spend some time in describing the foundations of this theory, exhibiting in particular its remarkable algebraic simplicity.

Formally speaking, spline fitting can be considered a particular case of what follows. let $\mathcal{H}$ be a Hilbert space, let $f_1, \dots, f_N \in \mathcal{H}$, and $c_1, \dots, c_N \in \mathbb{R}$ be given. Denote by $\langle . \, , . \rangle$ the inner product on $\mathcal{H}$. Consider the following problems:

1. Find $h \in \mathcal{H}$ such that $||h||$ is minimum subject to the constraints $\langle f_i \, , \, h \rangle = c_i$ for $i = 1, \dots, N$.

2. Fixing $\lambda > 0$, find h such that $||h||^2 + \lambda \sum_{i=1}^{N} (\langle f_i \, , \, h \rangle - c_i)^2$ is minimum.

The first problem corresponds to interpolation, or exact matching, the second one to smoothing, or approximate matching, and both are solved by elementary linear algebra. It is indeed clear that, in both cases, the constraints are not affected if h is replaced by $h + v$ where v is orthogonal to all the f_i, so that the solution must in fact be searched in the linear space spanned by $f_1, \ldots, f_N$: so, introduce the $N \times N$ matrix S with $S_{ij} = \langle f_i \, , \, f_j \rangle$, and express the unknown h as a linear combination $h = \sum_{i=1}^{N} \alpha_i f_i$.

Problem 1 now requires to minimize ${}^t\alpha S \alpha$ subject to the constraint $S\alpha = c$ (where α and c are vectors with components α_i and c_i respectively), and problem 2 to minimize ${}^t\alpha S \alpha + \lambda \, {}^t(S\alpha - c)(S\alpha - c)$

Assume that S is invertible, so that no linear constraint can be deduced one from another [1]. The solution of problem 1 is in fact uniquely specified by the constraint: it is $\alpha = S^{-1}.c$. For the second problem, routine computations shows that it is $\alpha = S_\lambda^{-1}.c$ with $S_\lambda = S + I/\lambda$ (the invertibility of S is here not required).

Let us see how this applies to splines. In this context, a set of points $(x_1, \ldots, x_N)$ in Ω is given, together with real numbers $c_1, \ldots, c_N$, and spline interpolation corresponds to finding a real-valued function h (defined on Ω), as smooth as possible, such that $h(x_i) = c_i$ or $h(x_i) \simeq c_i$. The smoothness of h is evaluated through a norm of the kind

$$||h||_L = \int_\Omega |Lh|^2 dx$$

where L is, say, a differential operator. This norm defines a Hilbert space of functions $\mathcal{H}_L$, with inner-product

$$\langle h \, , \, g \rangle_L = \int_\Omega LhLg dx \, .$$

The constraints $h(x_i) = c_i$ are linear in h, and the issue, to fit in the previous abstract setting, is whether there exists an element f_{x_i} in $\mathcal{H}_L$ such that, for all $h \in \mathcal{H}_L$ $h(x_i) = \langle f_{x_i} \, , \, h \rangle_L$. If this can be done[2], the solution of the interpolation problem is given by a linear combination of the f_{x_i}, the coefficients of which are simply obtained by applying the inverse of the matrix of inner-products of the f_{x_i} to the values of the constraints c_i. A similar conclusion can be drawn if we replace this exact interpolation problem by an inexact form, which consists in minimizing

$$||h||_L + \lambda \sum_{i=1}^{N} (h(x_i) - c_i)^2$$

[1] Problem 1 may have no solution when S is not invertible

[2] This is equivalent, by the Riesz representation theorem, to the continuity of the evaluation mapping $h \mapsto h(x)$ for the norm $||.||_L$

It must be noted that the inner-products $\langle f_{x_i}, f_{x_j} \rangle$ are, by construction, given by $f_{x_i}(x_j)$ (f is self-reproducing), so that their computation is immediate.

So everything depends on the existence of f_x. Theoretical arguments for proving this existence can be given (linked, for example, to Sobolev's inclusion theorems when L is a differential operator) but, for practical purposes, it is also necessary to have an analytical and computable expression for them. The functions f are obtained from the Green kernel of the operator $L^*.L$. [3] So, practical applicability of spline interpolation depends on whether the Green kernel of K is known or not, provided of course it exists at all.

Well-know cases in which explicit expressions for Green functions are available are when L is a variant of the Laplacian (with simple enough boundary conditions). However, this can be too specific in some applications, and noting that the method never requires explicit knowledge of L, it is possible to start directly with a function $(x, y) \mapsto f_x(y)$ provided that one knows that f is the Green kernel of some positive operator K (not necessarily a differential operator), which needs not be specified. Non constructive existence assumptions exist, for example one may require that f is symmetric ($f_x(y) = f_y(x)$), continuous, squared-integrable and induces a positive operator on $\mathcal{L}^2$, the last requirement being satisfied when $f_x(y) = F(x - y)$, F being the Fourier transform of a positive, even function.

In [2], [1], these requirements are further specialized to ensure that f a radial basis function: $f_x(y) = G(|x - y|)$, the simplest example being the Gaussian $G(t) = \exp\left(-t^2/\sigma^2\right)$. This additional assumption has the advantage to provide rotation invariant interpolation (which is also true when L is a linear combination of powers of the Laplacian).

2.2 Landmark matching and Bookstein's splines

For diffeomorphic matching, it must be taken into account, to apply the previous, that the unknown function is vector valued. In fact, if $x_1, \ldots, x_N$ and $y_1, \ldots, y_N$ are two matched sets of landmarks, one should find a diffeomorphism $h : \Omega \to \Omega$ such that $h(x_i) = y_i$ (equivalently, one searches a displacement $u = h - \mathrm{id}$ such that $u(x_i) = y_i - x_i$).

Bookstein (see [3]) proposes to apply spline interpolation to each component of u. This is the simplest approach, and we will keep to this setting in this

[3] Let L^* be the dual operator of L, which is such that, for all g and h with compact support in Ω,

$$\int_\Omega (Lh)g = \int_\Omega (L^* g)h$$

and let $K = L^*.L$. For all x, one has

$$h(x) = \langle f_x , h \rangle_L = \int_\Omega L f_x L h dy = \int_\Omega f_x K h dy$$

which is precisely the definition of the fact that the function $(x, y) \mapsto f_x(y)$ is the Green kernel K. (K is also affected by boundary conditions, but we omit the complications here).

paper. But it must be remarked that this is not a general point of view, the Green function being, in this context, naturally expressed as a *matrix* of kernels.

Returning to Bookstein splines, consider the inner-product

$$\langle h\,,\,g\rangle = \int_{\mathbb{R}^2} \Delta h \Delta g dx\,. \tag{1}$$

$\mathcal{H}$ being the space of functions with square integrable second derivatives (a Beppo-Levi space). This space is not, strictly speaking, a Hilbert space, because the inner product is degenerate (it vanishes for affine functions), but can considered as one provided one consideres that functions which only differ by an affine function are equal.

With this in mind, let $U(r) = r^2 \log r$: this function is such that, for any smooth function $h \in \mathcal{H}$,

$$h(x) = \int_{\mathbb{R}^2} U(|x-y|)\Delta^2 h(y) dy$$

where Δ^2 is the iterated Laplacian, and equality being up to the addition of an affine function (cf [7]). Letting $f_x(y) = U(|x-y|)$, this is $\langle f_x\,,\,h\rangle = h(x)$ (up to an affine function). So let c_i be one of the components of $y_i - x_i$, so that the constraints $h(x_i) = c_i$ for $i = 1$ to N write: there exist $a = (a^1, a^2) \in \mathbb{R}^2$, $b \in \mathbb{R}$ such that, for all $i : \langle h\,,\,f_{x_i}\rangle = c_i - {}^t a x_i - b$.

Let $S_{ij} = \left\langle f_{x_i}\,,\,f_{x_j}\right\rangle = U(|x_i - x_j|)$. The interpolation problem becomes: minimize ${}^t\alpha S \alpha$ with the constraint $S\alpha + Q\gamma = c$ where $\gamma = {}^t(a^1 a^2 b)$ is a 3×1 matrix and Q is a $N \times 3$ matrix, given by, letting $x_i = (x_i^1, x_i^2)$:

$$Q = \begin{pmatrix} x_1^1 & x_1^2 & 1 \\ \vdots & \vdots & \vdots \\ x_N^1 & x_N^2 & 1 \end{pmatrix}$$

solving this problem in (α, γ) yields $\hat{\gamma} = \left({}^tQS^{-1}Q\right)^{-1} {}^tQc$ and $\hat{\alpha} = S^{-1}(c - Q\hat{\gamma})$.

The smoothing problem requires minimizing

$${}^t\alpha S \alpha + \lambda^t(S\alpha + {}^t a x_i + b - c)(S\alpha + {}^t a x_i + b - c)$$

and its solution is formally similar to the previous one, simply replacing S by $S_\lambda = S + (1/\lambda) I$ in the formulas.

When this is applied to both components of $y_i - x_i$, one obtains a function u such that $u(x_i) = y_i - x_i$ for exact matching, which thus provides a smooth interpolation of the landmark correspondence. However, *there is no constraint in this approach, which ensures that* $h(x) = x + u(x)$ *is one-to-one*: folding is indeed possible, and examples will be given in the last section of this paper. We shall obtain a rigorous one-to-one matching using flows of diffeomorphisms, as introduced in the next section.

3 Diffeomorphic landmark matching

One way to introduce the groups of diffeomorphisms we shall be dealing with starts from standard methods for generating metrics and distances on sets acted on by groups. We address this in the next section.

3.1 Distances and group actions

General facts We start with a short algebraic section in which basic facts on how inducing a distance from a group action are obtained, introducing in particular a "least-action principle". First recall that a distance on a set $\mathcal{I}$ is a mapping $d : \mathcal{I}^2 \mapsto \mathbb{R}_+$ such that, for all $I, I', I'' \in \mathcal{I}$: D1) $d(I, I') = 0 \Leftrightarrow I = I'$, D2) $d(I, I') = d(I', I)$, D3) $d(I, I'') \leq d(I, I') + d(I', I'')$.

If D1) is not true and $d(I, I) = 0$ for all I, we use the term *pseudo-distance*.

A group G is acting on $\mathcal{I}$ if an operation $(g, I) \to g.I$ is defined on $G \times \mathcal{I}$ with values in $\mathcal{I}$ such that $\mathrm{id}_G.I = I$ and $g.(h.I) = (gh).I$ for all $I \in \mathcal{I}$ and all $g, h \in G$. If G is a group acting on $\mathcal{I}$, one says that a distance d on $\mathcal{I}$ is G-equivariant if and only if, for all $g \in G$, for all $I, I' \in \mathcal{I}$, $d(g.I, g.I') = d(I, I')$.

We shall be dealing with the following construction. Let G act on $\mathcal{I}$, and consider the product $\mathcal{O} = G \times \mathcal{I}$, so that G also acts on $\mathcal{O}$ (simply letting, for $k \in G$, $o = (h, I) \in \mathcal{O}$: $k.o = (kh, k.I)$).

For $o = (h, I) \in \mathcal{O}$, let $\pi(o) = h^{-1}.I$. Assume that $d_{\mathcal{O}}$ is a distance on $\mathcal{O}$, and let, for $I, I' \in \mathcal{I}$

$$d(I, I') = \inf\{d_{\mathcal{O}}(o, o'), o, o' \in \mathcal{O}, \pi(o) = I, \pi(o') = I'\} \tag{2}$$

Proposition 1. *If $d_{\mathcal{O}}$ is G-equivariant, then d in (2) is a pseudo-distance on $\mathcal{I}$*

We shall use these results for landmarks, letting diffeomorphisms act on them. This however asks the problem of building an invariant distance on $\mathcal{O}$: one simple approach for this is to use geodesics in this space, as presented in the next section.

Infinitesimal approach A standard way for building distances on sets like $\mathcal{O}$ is to compute shortests paths. Assume that we are able to give a meaning of the speed $V_{\mathbf{o}}(t) = \frac{d\mathbf{o}}{dt}$ of a path $\mathbf{o} : t \mapsto \mathbf{o}(t)$ on $\mathcal{O}$. Assume also that, for each $o \in \mathcal{O}$, we have a way to quantify the speeds of paths passing through o with the help of a norm $V \mapsto ||V||_o$ (the norm depends on o). Then, let the associated path energy be given by

$$E(\mathbf{o}) = \int_0^1 ||V_{\mathbf{o}}(t)||^2_{\mathbf{o}(t)}\, dt \tag{3}$$

and the *geodesic distance* on $\mathcal{O}$ be then defined by

$$d_{\mathcal{O}}(o, o') = \inf\{\sqrt{E(\mathbf{o})}, \mathbf{o}(0) = o, \mathbf{o}(1) = o'\}\,. \tag{4}$$

To build a G-equivariant distance (as required by proposition 1), it suffices to start with a family of norms $(||.||_o, o \in \mathcal{O})$ which shares this property, in the sense that, if $\mathbf{o}$ is a path on $\mathcal{O}$ and $h \in G$, then the translated path $h.\mathbf{o}$ and $\mathbf{o}$ both have the same speeds at the same times: this writes

$$||V_{h.\mathbf{o}}(t)||_{h.\mathbf{o}(t)} = ||V_{\mathbf{o}}(t)||_{\mathbf{o}(t)} \tag{5}$$

One can interpret this formula with the help of the differential of the action of G on $\mathcal{O}$, but the meaning and the consequences of (5) will be easily derived, and we will not need to introduce the usual machinery of differential geometry.

It is important to notice that this condition provides norms of velocities at translated objects $h.o$ as soon as these norms are known at o. Since $\mathcal{O} = G \times \mathcal{I}$, it thus suffices to define $||.||_o$ for $o \in \mathcal{O}$ of the kind $o = (\mathrm{id}_G, I)$.

3.2 Mixing deformations and object variations: landmark matching

We specialize the point of view of section 3.1, by letting G be a group of diffeomorphisms of Ω and $\mathcal{I}$ be the set of all collections of N landmarks on Ω. An element of $\mathcal{I}$ is thus a N-tuple $I = (p_1, \dots, p_N) \in \Omega^N$. We use on G the product $gh = h \circ g$ and define the action of G on $\mathcal{I}$ to be

$$g.I = (g^{-1}(p_1), \dots, g^{-1}(p_N))$$

which does provide a left-action: $(gh).I = g.(h.I)$.

Following the lines of section 3.1, we consider paths on $\mathcal{O}$. Such a path takes the form $\mathbf{o}(t) = (\mathbf{g}(t,.), \mathbf{p}_1(t), \dots, \mathbf{p}_N(t))$ where $\mathbf{g}(t,.)$ is a time dependent diffeomorphism and $\mathbf{p}_i(t)$ is a curve in Ω for $i = 1, \dots, N$ (landmark trajectory). The velocity at $\mathbf{o}(t)$ is

$$V_{\mathbf{o}}(t) = \left(V_{\mathbf{g}}(t), \frac{d\mathbf{p}_1}{dt}(t), \dots, \frac{d\mathbf{p}_N}{dt}(t)\right)$$

with $V_{\mathbf{g}}(t) = \frac{\partial \mathbf{g}}{\partial t}$.

Let h be a diffeomorphism; the path $h.\mathbf{o}$ is given by

$$h.\mathbf{o}(t) = (\mathbf{g}(t, h(.)), h^{-1}(\mathbf{p}_1(t)), \dots, h^{-1}(\mathbf{p}_N(t)))$$

and its speed is

$$V_{h.\mathbf{o}}(t) = \left(V_{\mathbf{g}} \circ h(t), D_{\mathbf{p}_1(t)}h^{-1}\frac{d\mathbf{p}_1}{dt}(t), \dots, D_{\mathbf{p}_N(t)}h^{-1}\frac{d\mathbf{p}_N}{dt}(t)\right)$$

where $D_p h^{-1}(t)$ is the differential of $h^{-1}(.)$ with respect to spatial coordinates, evaluated at a point $p \in \Omega$. Equation (5), requires that

$$||V_{h.\mathbf{o}}(t)||_{h.\mathbf{o}(t)} = ||V_{\mathbf{o}}(t)||_{\mathbf{o}(t)}$$

This is true in particular when h is the inverse of $\mathbf{g}(t,.)$, yielding

$$||V_{\mathbf{o}}(t)||_{\mathbf{o}(t)} = ||V_{\mathbf{g}^{-1}\mathbf{o}}(t)||_{\mathbf{g}^{-1}.\mathbf{o}(t)}$$

so that it is only necessary to define norms at elements $o = (id, p_1, \ldots, p_N)$. We have,

$$V_{\mathbf{g}^{-1}\mathbf{o}}(t) = \left(\mathbf{v_g}(t), D_{\mathbf{p}_1(t)}\mathbf{g}(t)\frac{d\mathbf{p}_1}{dt}(t), \ldots, D_{\mathbf{p}_N(t)}\mathbf{g}(t)\frac{d\mathbf{p}_N}{dt}(t) \right)$$

in which we have let $\mathbf{v_g}(t, y) = \mathbf{V_g}(t, \mathbf{g}^{-1}(t, y))$. Making the change of variables $\mathbf{q}_i(t) = \mathbf{g}(t, \mathbf{p}_i(t))$, this writes

$$\mathbf{V}_{\mathbf{g}^{-1}\mathbf{o}}(t) = \left(\mathbf{v_g}(t), \frac{d\mathbf{q}_1}{dt}(t) + \mathbf{v_g}(t, \mathbf{q}_1(t)), \ldots, \frac{d\mathbf{q}_N}{dt}(t) - \mathbf{v_g}(t, \mathbf{q}_N(t)) \right)$$

and the energy of the path $o(t)$ takes the form

$$E(\mathbf{o}) = \int_0^1 \left\| \mathbf{v_g}(t), \frac{dq_1}{dt}(t) + \mathbf{v_g}(t, \mathbf{q}_1(t)), \ldots, \frac{d\mathbf{q}_N}{dt}(t) - \mathbf{v_g}(t, \mathbf{q}_N(t)) \right\|^2 dt$$

with a certain norm, which may depend on the current position $(\mathbf{q}_1(t), \ldots, \mathbf{q}_N(t))$.

The essential trick, for building diffeomorphisms (in theory and in practice), is to replace the unknown time-dependent diffeomorphism $\mathbf{g}$ by its so-called Eulerian velocity $\mathbf{v_g}$ on which everything now depends. Noting that, by definition:

$$\frac{\partial \mathbf{g}}{\partial t} = \mathbf{v_g}(t, \mathbf{g}(t))$$

knowing $\mathbf{v}$ allows to compute $\mathbf{g}$ by integration of an ODE, providing in that way a flow of diffeomorphisms (under smoothness conditions on $\mathbf{v}$ which have been studied in detail in [10] and [4]).

So, we know think in terms of $\mathbf{v}$ rather of $\mathbf{g}$. To compute the distance between two elements of $\mathcal{O}$, it suffices to minimize the energy of the paths which link them. We now restrict to the particular case when

$$E(\mathbf{o}) = \int_0^1 \int_\Omega |L\mathbf{v_g}(t)|^2 dt dx + \sum_{i=1}^N \int_0^1 \left| \frac{d\mathbf{q}_i}{dt}(t) - \mathbf{v_g}(t, \mathbf{q}_i(t)) \right|^2 dt$$

and L being some operator acting on $v(t)$ for all t. We are interested by the distance between two sets of landmarks I and I', which is given, according to (2)

$$d(I, I') = \inf\{d_{\mathcal{O}}(o, o'), o, o' \in \mathcal{O}, \pi(o) = I, \pi(o') = I'\}$$

where $\pi(g, p_1, \ldots, p_N) = g^{-1}(p_1, \ldots, p_N) = (g(p_1), \ldots, g(p_N))$. It is not difficult to check that in fact, $d(I, I')$ is the infimum of

$$\int_0^1 \int_\Omega |L\mathbf{v}(t)|^2 dt dx + \sum_{i=1}^N \int_0^1 |\frac{d\mathbf{q}_i}{dt}(t) - \mathbf{v}(t, \mathbf{q}_i(t))|^2 dt \tag{6}$$

over all time dependent velocities $\mathbf{v}$ on Ω, and over all curves $\mathbf{q}_1(.), \ldots, \mathbf{q}_N(.)$ such that $I = (\mathbf{q}_1(0), \ldots, \mathbf{q}_N(0))$ and $I' = (\mathbf{q}_1(1), \ldots, \mathbf{q}_N(1))$. We thus obtain a new landmark-based matching formula, which only involves the velocity, and which, in the same time, provides a distance between sets of landmarks. For fixed trajectories $\mathbf{q}_i(.)$, $i = 1, \ldots, N$, the optimal $\mathbf{v}$ can be explicitely computed at each time t, in function of the Green kernel of L^*L, as developped in section 2.1: this is the basis of the numerical algorithm which is detailed in section 4.

The final form of the energy is somehow reminiscent from S. Joshi's landmark matching method, which is also based on flows of diffeomorphisms ([6]). The main difference comes from the fact that Joshi's method does not optimize landmark trajectories, but rather uses an end-point matching penalty which leads to an optimal control formulation. There are two consequences of this: the first one is that this does not provide a metric between landmark configurations, whereas the methods derived here provides this feature as an initial motivation. The second one is that the numerical problem in our case is much simpler, as will be seen in section 4.

3.3 A Riemannian metric on deformable landmark configurations

An interesting feature of the previous construction can be pointed out here: it is that minimizing (6) with boundary conditions only in $\mathbf{q}$ is equivalent to compute a *geodesic path* for a specific metric on the set of configurations of N landmarks. Indeed, for $I = (q_1, \ldots, q_N) \in \Omega^N$, and for $h = (h_1, \ldots, h_N) \in \left(\mathbb{R}^2\right)^N$, define

$$||h||_I^2 = \min_v \int_\Omega |Lv|\, dx + \lambda \sum_{i=1}^N |h_i - v(q_i)|^2$$

It is easy to prove that $||h||_I$ is a norm as a function of h, which therefore provides a Riemannian metric on Ω^N, and to check that the optimal trajectories $\mathbf{I}(.) = (\mathbf{q}_1(.), \ldots, \mathbf{q}_N(.))$ minimize the energy

$$E(\mathbf{I}) = \int_0^1 \left\| \frac{d\mathbf{I}}{dt} \right\|_{\mathbf{I}(t)}^2 dt$$

with fixed boundary conditions at time 0 and at time 1, yielding the fact that they are geodesics. Noting, moreover, that, for a given h, the minimizing v in the definition of $||h||_I$ can be explicitely computed, provided that the Green function of L^*L is known (cf section 2.1), we finally obtain an explicit view of Ω^N as a Riemanian manifold, in which diffeomorphisms are now only implicit.

4 Experiments

4.1 Implementation details

The simplest numerical scheme is not to directly work with the geodesic energy of the previous section. We give ourselves a Green function, denoted f, as in section 2.1, associated to some operator L we do not need to compute. For each t, the optimal $\mathbf{v}$ must have the form

$$\mathbf{v}(t,x) = \sum_{k=1}^{N} \alpha_k(t) f(q_k(t), x)$$

where for each k and t, $\alpha_k(t)$ is a 2D vector. The energy in (6) writes, as a function of α and q:

$$\begin{aligned} E(\alpha, q) = &\sum_{k,l=1}^{N} \int_0^1 \langle \alpha_k(t)\,,\, \alpha_l(t)\rangle f(q_k(t), q_l(t)) dt \\ &+ \lambda \sum_{k=1}^{n} \int_0^1 \left\| \frac{dq_k}{dt} - \sum_{l=1}^{N} \alpha_l(t) f(q_l(t), q_k(t)) \right\|^2 dt \end{aligned} \tag{7}$$

in which we use norms and inner products in $\mathbb{R}^2$. Now, for fixed q, the optimal α is explicit. Let $S(t)$ be the $N \times N$ matrix with coefficients $f(q_l(t), q_k(t))$, and let $S_\lambda(t) = S(t) + I/\lambda$. Letting $\alpha^i = (\alpha_1^i, \ldots, \alpha_N^i)$ and $q^i = (q_1^i, \ldots, q_N^i)$, $i = 1, 2$, we have, for all t

$$\alpha^i(t) = [S_\lambda(t)]^{-1} . q^i(t)$$

Minimizing in q with fixed α is not explicit, but the gradient of E in q can be computed. We assume a time discretization of order T, and set, for $t = 0, \ldots, T-1$: $D_t u = T(u(t+1) - u(t))$ for a time dependent function u. The discretized energy takes the form

$$\begin{aligned} E(\alpha, q) = &\sum_{k,l=1}^{N} \sum_{t=0}^{T} \langle \alpha_k(t)\,,\, \alpha_l(t)\rangle f(q_k(t), q_l(t)) \\ &+ \lambda \sum_{k=1}^{N} \sum_{t=0}^{T-1} \left\| D_t q_k - \sum_{l=1}^{N} \alpha_l(t) f(q_l(t), q_k(t)) \right\|^2 \end{aligned} \tag{8}$$

Introduce the notation $Z_k(t) = D_t q_k - \sum_{l=1}^{N} \alpha_l(t) f(q_l(t), q_k(t))$.

The partial derivative of E with respect to $q_k(t)$ is a 2D vector, given by, for $t = 1, \ldots T-1$:

$$\begin{aligned} \frac{\partial E}{\partial q_k(t)} = &\, 2 \sum_{l=1}^{N} \langle \alpha_k(t)\,,\, \alpha_l(t)\rangle \nabla_1 f(q_k(t), q_l(t)) - 2\lambda T D_{t-1} Z_k \\ &- 2\lambda \sum_{l=1}^{N} (\langle \alpha_k(t)\,,\, Z_l(t)\rangle + \langle \alpha_l(t)\,,\, Z_k(t)\rangle) \nabla_1 f(q_k(t), q_l(t)) \end{aligned} \tag{9}$$

where $\nabla_1 f \in \mathbb{R}^2$ denotes the gradient of f with respect to one of its variables (it does not matter which one, f is symmetric). Note that $q(0)$ and $q(T)$ are clamped to the boundary conditions.

The optimal landmark trajectories are computed by iterating the following steps until convergence (initializing with $\alpha = 0$ and linear trajectories)

1. Gradient step for q: for all $t = 2, \ldots T-1$, substract to $q_k(t)$ a quantity proportional to $\dfrac{\partial E}{\partial q_k(t)}$
2. Velocity updating: set $\alpha^i(t) = [S_\lambda(t)]^{-1} . q^i(t)$, $t = 1, \ldots, T-1$, $i = 1, 2$

Thin-plate generalization We now show how this can be modified to incorporate affine invariance as in the thin-plate case. If L is the Laplacian, we have seen in section 2.2 that, for each t, $v(t, x)$ is defined up to an affine component which may in turn be optimized. This yields an expression

$$\mathbf{v}(t, x) = a(t)x + b(t) + \sum_{k=1}^{N} \alpha_k(t) f(q_k(t), x)$$

for the unknown velocity field, $a(t)$ being a 2×2 matrix and $b(t) \in \mathbb{R}^2$. Let $\gamma^i(t)$ be a column of the 3×2 matrix ${}^t[a(t)b(t)]$, $i = 1, 2$. When trajectories are fixed, the optimal coefficients are given like in section 2.2 by

$$\gamma^i(t) = \left({}^tQ(t) S_\lambda^{-1}(t) Q(t)\right)^{-1} {}^tQ(t) \frac{dq^i}{dt}$$

and $\alpha^i(t) = S_\lambda^{-1} \left(\dfrac{dq^i}{dt} - Q(t)\gamma^i(t) \right)$ where S_λ is as above and

$$Q(t) = \begin{pmatrix} q_1^1(t) & q_1^2(t) & 1 \\ \vdots & \vdots & \vdots \\ q_N^1(t) & q_N^2(t) & 1 \end{pmatrix}$$

The modification of the gradient equation in $q(t)$ is almost straightforward and we only give the result in discretized form, letting

$$Z_k(t) = D_t q_k - a(t) q_k(t) - b(t) - \sum_{l=1}^{N} \alpha_l(t) f(q_l(t), q_k(t))$$

so that the minimized energy is

$$E(\alpha, q) \;=\; \sum_{k,l=1}^{N} \sum_{t=0}^{T} \langle \alpha_k(t) \,,\, \alpha_l(t) \rangle f(q_k(t), q_l(t)) \;+\; \lambda \sum_{k=1}^{N} \sum_{t=0}^{T-1} \|Z_k(t)\|^2 \,. \quad (10)$$

We have

$$\frac{\partial E}{\partial q_k(t)} = 2\sum_{l=1}^{N} \langle \alpha_k(t)\,,\, \alpha_l(t)\rangle \nabla_1 f(q_k(t), q_l(t)) - 2\lambda T D_{t-1} Z_k$$
$$- 2\lambda \sum_{l=1}^{N} (\langle \alpha_k(t)\,,\, Z_l(t)\rangle + \langle \alpha_l(t)\,,\, Z_k(t)\rangle) \nabla_1 f(q_k(t), q_l(t)) - 2\lambda\, {}^t a(t) Z_k(t)\,. \tag{11}$$

Note that we wrote everything for 2D matching, but that the formulas can obviously generalize to any number of dimensions.

4.2 Experiments

In the proposed experiments, random displacements are attributed to points evenly distributed on a grid. Interpolated deformation and landmark trajectories are computed. This is compared to classical spline interpolation (which just corresponds to $T = 2$ in the previous algorithm). The Green kernel, f is Gaussian: $f(x, y) = \exp\left(-|x-y|^2/2\sigma^2\right)$.

Three experiments are presented. The first one generates small displacements, the last one very large ones. Progressively, one sees the singularities generated by classical interpolating splines increase, foldings being created, while geodesic splines remain one-to-one, and of rather impressive smoothness. The estimated landmarks trajectories start to bend in the second experiments, and are clearly curved in the last one.

References

1. N. Arad, N. Dyn, D. Reisfeld, and Y. Yeshurun, *Image warping by radial basis functions: application to facial expressions*, CVGIP: Graphical Models and Image Processing, 56 (1994), pp. 161–172.
2. N. Arad and D. Reisfeld, *Image warping using few anchor points and radial functions*, Computer Graphics forum, 14 (1995), pp. 35–46.
3. L. Bookstein, F, *Morphometric tools for landmark data; geometry and biology*, Cambridge University press, 1991.
4. P. Dupuis, U. Grenander, and M. Miller, *Variational problems on flows of diffeomorphisms for image matching*, Quaterly of Applied Math., (1998).
5. S. Helgason, *Differential Geometry, Liie groups and Symmetric spaces*, Academic Press, 1978.
6. S. Joshi and M. Miller, *Landmark matching via large deformation diffeomorphisms*, IEEE transactions in image processing, 9 (2000), pp. 1357–1370.
7. J. Meinguet, *Multivariate interpolation at arbitrary points made simple*, J. Appliet Math. and Physics, 30 (1979), pp. 292–304.
8. I. Miller, M and L. Younes, *Group action, diffeomorphism and matching: a general framework*, in Proceeding of SCTV 99, 1999. http://www.cis.ohio-state.edu/ szhu/SCTV99.html.

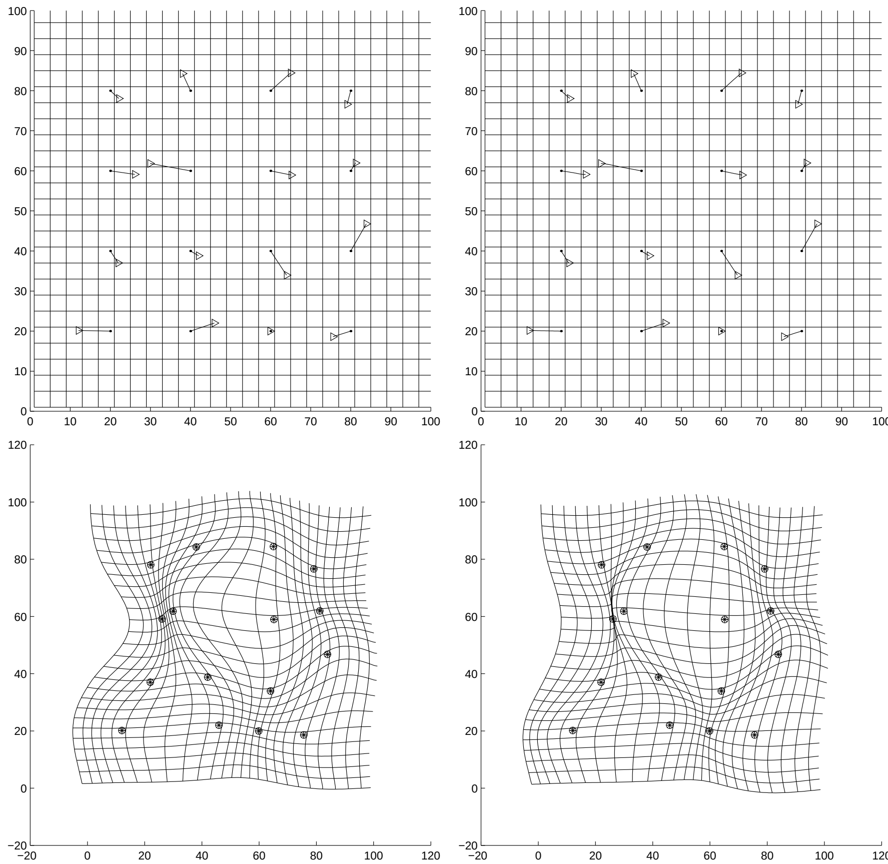

Fig. 1. Random point displacements: small deformations; upper-left: evenly spaced points and random displacements ; upper-right: optimal trajectories ; down-left: estimated diffeomorphism ; down-right: interpolated displacement field by classical splines.

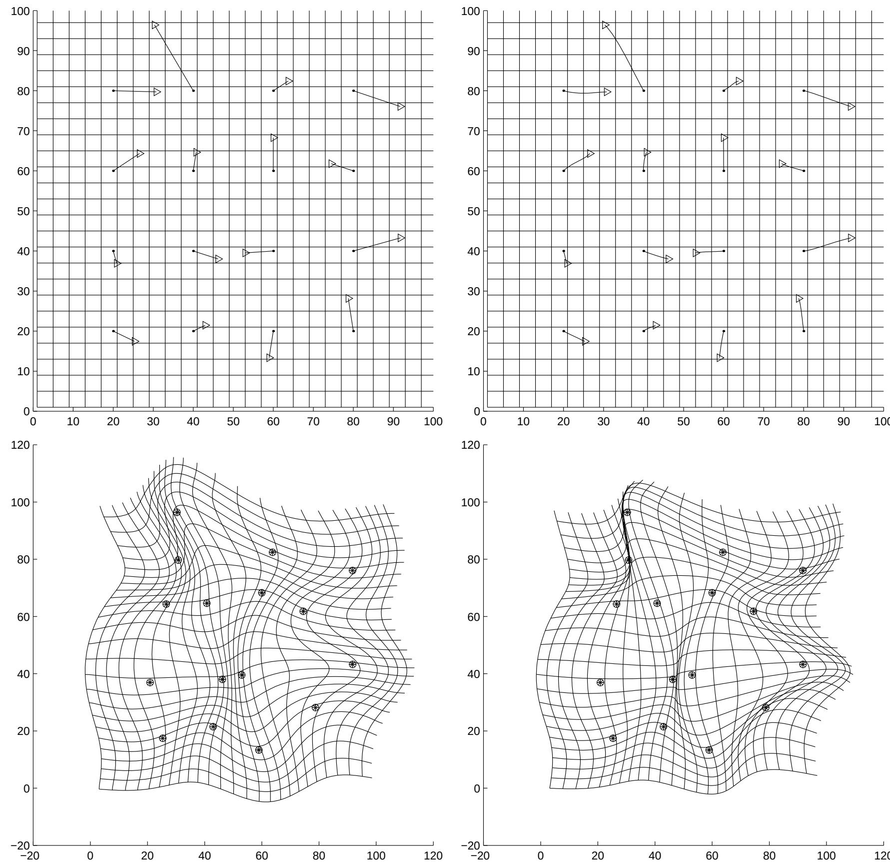

Fig. 2. Random point displacements: average deformations; upper-left: evenly spaced points and random displacements ; upper-right: optimal trajectories ; down-left: estimated diffeomorphism ; down-right: interpolated displacement field by classical splines

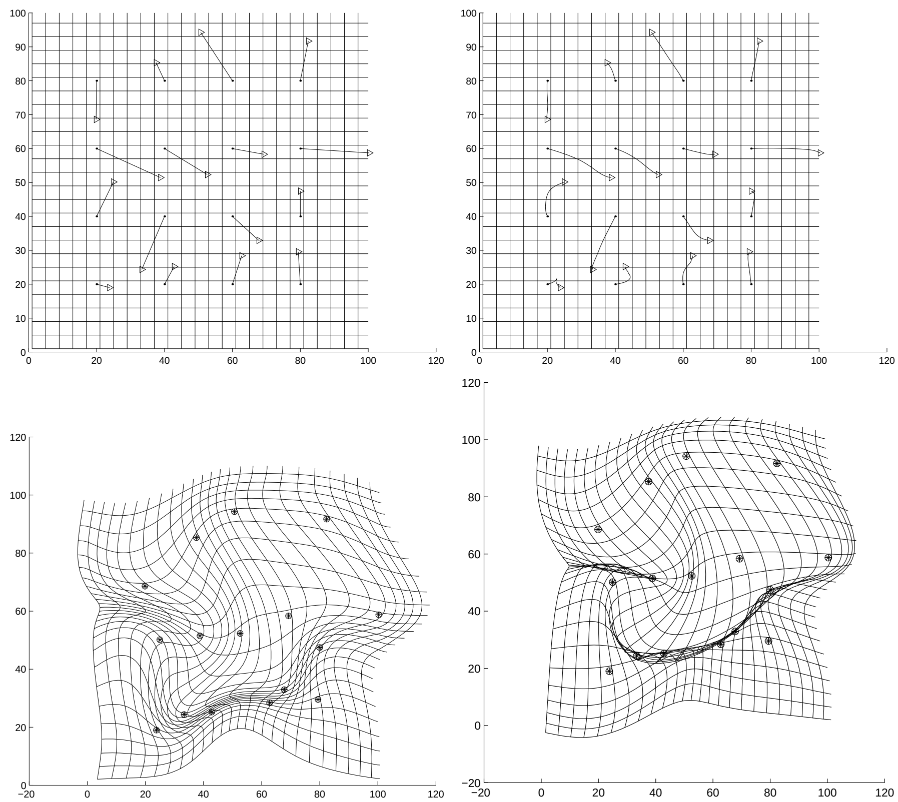

Fig. 3. Random point displacements: large deformations; upper-left: evenly spaced points and random displacements ; upper-right: optimal trajectories ; down-left: estimated diffeomorphism ; down-right: interpolated displacement field by classical splines

9. P. Olver, *Equivalence, Invarians and Symmetry*, Cambridge University Press, 1995.
10. A. Trouvé, *Infinite dimensional group action and pattern recognition*, Quaterly of Applied Math. (to appear), (2000).
11. L. Younes, *Computable elastic distances between shapes*, SIAM J. Appl. Math, 58 (1998), pp. 565–586.