

Discrete mechanics and optimal control for image registration

Robert I. McLachlan¹ Stephen Marsland²

(Received 14 August 2006; revised 17 March 2007)

Abstract

Diffeomorphic image registration, where images are aligned using diffeomorphic warps, is a popular subject for research in medical image analysis. We introduce a novel algorithm for computing diffeomorphic warps that solves the Euler equations on the diffeomorphism group explicitly, rather than using an optimiser. The result is an algorithm that is many times faster than those considered previously.

Contents

1 Introduction	C2
2 Problem formulation	C3

See <http://anziamj.austms.org.au/ojs/index.php/ANZIAMJ/article/view/82> for this article, © Austral. Mathematical Soc. 2007. Published April 23, 2007. ISSN 1446-8735

<i>1</i>	<i>Introduction</i>	C2
3	A particle method for image registration	C5
3.1	Position of point particles	C6
3.2	Initialisation of point momenta	C7
3.3	Optimisation method	C7
3.4	Choice of integrator	C7
3.5	Test particles and interpolation	C8
3.6	Choice of metric	C9
3.7	Adding more points	C9
4	Experiments	C10
5	Conclusions and open questions	C12
	References	C13

1 Introduction

Image registration has received much research over the past few years, not least because of its many applications in medicine. For example, it is useful for removing motion artefacts caused by patient movement [13], aligning to an atlas [4], assisting in disease diagnosis [10], and for measuring anatomical variability between subjects [14].

For applications in disease diagnosis and measuring anatomical variability, some form of measurement on the space of images is essential, to allow statistical analysis of the image warps. This generally requires using diffeomorphic image registration, whereby the choice of image warps that can be used to solve the registration problem are constrained to be diffeomorphisms, that is, smooth functions that have smooth inverses. One approach to this problem, known as Computational Anatomy, is to introduce group actions as deformable templates that are warped via the actions of a group onto other images [5]. This work has been fundamental to a large amount of research

on aligning images through landmark matching, where corresponding points are defined on a set of images, and diffeomorphic warps used to align them. The corresponding group for image analysis is the full diffeomorphism group, not the volume preserving subgroup that is used in fluid mechanics.

Under a right-invariant Riemannian metric it can be shown that the geodesics of the motion of a set of landmarks can be computed as an optimisation problem [11, 10]. In this article we introduce a novel formulation for the problem that is based on implicitly solving the partial differential equations that govern the motion. These partial differential equations are the Euler equations for the full diffeomorphism group, given by equations (1) and (2); for derivations, following Arnold [1], see [11, 7]. We introduce a particle method that enables us to solve for the diffeomorphism directly. This results in an algorithm that is orders of magnitude faster than previous ones (taking several minutes rather than an hour or two). We demonstrate the algorithm using the standard forward-Euler and Runge–Kutta algorithms, and discuss the benefits of using a symplectic integrator.

2 Problem formulation

We begin by defining the problem of diffeomorphic image registration:

Assume that there is a diffeomorphism ϕ that takes a free image F to a reference image R , that is, $R = F \circ \phi$. The aim of diffeomorphic image registration is to discover ϕ .

The diffeomorphism ϕ is defined on some domain $\Omega \subset \mathbb{R}^2$ or \mathbb{R}^3 , and the images are typically grayscale, so that $R, F : \mathbb{R}^2 \rightarrow \mathbb{R}$ or $R, F : \mathbb{R}^3 \rightarrow \mathbb{R}$. The method used to find the desired ϕ is generally optimisation of some norm $\|R - F \circ \phi\|$. Typical choices include the L_2 norm (sum-of-squares error) and mutual information (which is related to the KL-divergence) [15].

In this article we describe a novel method of constructing the diffeomorphisms. The standard approach in the literature is to use an energy minimisation, which produces the diffeomorphism as a geodesic. There are effectively three different approaches: (a) considering the problem as one of inexact matching on the boundary values (the start and end points of the landmarks) [8, 2]; (b) considering the problem as one of inexact matching using the initial values of position and momentum [3]; (c) considering the problem as one of exact matching [10]. The first approach optimises the warp over the end points of the spline, which are not guaranteed to be reached precisely; the second optimises the warp over the initial momenta of the landmarks; whereas the third approach can consider either method equally.

For the case of the full diffeomorphism group, $\mathfrak{G} = \text{Diff}(\mathbb{R}^n)$, which we consider here, the Euler equations are [6, 12, for details]:

$$\dot{m} + u \cdot \nabla m + \nabla u^T \cdot m + m(\text{div}u) = 0, \quad (1)$$

$$m = \mathcal{A}u, \quad (2)$$

where \dot{m} denotes differentiation with respect to time, $u(x, t)$ ($u, x \in \mathbb{R}^n$, $t \in \mathbb{R}$) is a velocity field, $m(x, t)$ its associated momentum, and \mathcal{A} is an elliptic operator (for example, $\mathcal{A} = (1 - \nabla^2)^k$) called the inertia operator. The inverse of the inertia operator \mathcal{A} is given by convolution with the Green's function \mathbf{G} of \mathcal{A} , that is, $u = \mathbf{G} * m$, where $*$ denotes convolution and $\mathcal{A}\mathbf{G}(x, x') = \delta(x - x')$ for $x, x' \in \mathbb{R}^n$. We shall only consider rotationally invariant and diagonal \mathcal{A} ; in this case $\mathbf{G}(x, x') = G(\|x - x'\|)$ for a scalar function G .

A striking feature of Euler equations on diffeomorphism groups is that they admit (formally, at least) exact solutions in which the momentum is concentrated at a finite set of points. For fluid equations these are point vortices, which are widely studied both in their own right and as a means of approximating the evolution of smooth or other vorticities. In analogy with the point vortices of fluid dynamics, we call the particles used for the image registration *point particles*.

3 A particle method for image registration

We are considering the deformation of an image F , with the deformation defined by a set of points i (some subset of the pixels of the image) with position and momentum $(q^i(t), p^i(t))$, where $p_i = \dot{q}_i$ as they move from their initial state (q_0^i, p_0^i) to their endpoint at $t = 1$. Starting from the Euler equations on the diffeomorphism group (1) and (2) we first compute the Hamiltonian, which is the kinetic energy, and then discretise it by introducing the particle ansatz $m(x, t) = \sum_{j=1}^N p_j(t) \delta(x - q_j(t))$, where $\delta(\cdot)$ is Kronecker delta function. The evolution of the particles is then determined by the Hamiltonian

$$H = \frac{1}{2} \sum_{i,j} p_i \cdot p_j G(q_i - q_j), \quad (3)$$

where $G(\cdot)$ is the Green's function corresponding to the chosen metric on $\text{Diff}(\Omega)$. The most common choice in image registration, and the one that we use in this article is the H^∞ metric, which corresponds to using a Gaussian Green's function $G(r) = (1/\epsilon^2) \exp(-r^2/\epsilon^2)$, where ϵ is the length-scale in the metric. Other choices include the thin-plate spline and clamped-plate spline [10, for a review].

Solutions to (1) of this particle form obey Hamilton's equations for (3), in which the components of q_i and p_i are canonically conjugate variables. Here q_1, \dots, q_N represent the positions of the N particles that define the deformation, and p_1, \dots, p_N their momenta. The equations of motion of the point particles are

$$\dot{q}_i = \sum_{j=1}^N G(\|q_i - q_j\|) p_j, \quad (4)$$

$$\dot{p}_i = - \sum_j (p_i \cdot p_j) G'(\|q_i - q_j\|) \frac{q_i - q_j}{\|q_i - q_j\|}. \quad (5)$$

Computing the diffeomorphism defined by q_i and p_i , $i = 1, \dots, N$, is then simply a case of integrating the motion forward in time using (4) and (5),

and then interpolating the motion of the rest of the image in some way. Here, we use *test particles*, which have zero momentum and so follow the flow—see Section 3.5. This leads us to the complete Algorithm 1 for image registration. The references on each line of Algorithm 1 give the section where it is discussed in more detail. Section 4 gives some examples of 2D registrations using the algorithm.

Algorithm 1 Our image registration algorithm

- Choose point particle positions q on image F (§3.1)
 - Initialise the particle momenta p randomly (§3.2)
 - Optimise $\|R - F \circ \phi\|$ over p (§3.3):
 - For current p , integrate point particles forward in time (§3.4)
 - Integrate positions of the test particles (§3.5)
 - Interpolate between the test particles (§3.5)
 - Compute $\|R - F \circ \phi\|$ for chosen distance measure (§3.6)
 - Add more point particles and reoptimise $\|R - F \circ \phi\|$ (§3.7)
-

3.1 Position of point particles

There are several possible choices for placing the point particles [9]. For registration of brains we initially place some points around the skull of head images, and after optimising them, place more points using the discrepancy image method (which is based on where there are still errors in the objective function). For the hand images we use a uniform grid and the discrepancy image method.

3.2 Initialisation of point momenta

In the current implementation, the momenta of each point particle are initialised with a uniformly random direction, and with a small (10% of average warp size) uniform random magnitude for the warp. The sensitivity of the algorithm to the initialisation does not seem to be large, but one option that does appear to improve the results, although at a moderate computational cost, is to perform a coarse search over this relatively small number of parameters (two for each of the point particles, of which there may be 10–20 on the initial pass).

3.3 Optimisation method

The choice of a suitable optimiser is obviously crucial. In the current implementation we use the sum-of-squares distance measure, which leads fairly naturally to a least-squares non-linear optimiser. We currently use the `lsqnonlin` function in Matlab 7.1, which is a subspace trust region method based on the interior-reflective Newton method. Experimentation has found that 100 iterations is usually sufficient for the algorithm to converge reasonably, although further work will investigate this more thoroughly.

3.4 Choice of integrator

The principal component of our method is the computation of the current geodesic, based on q and the current p . This is calculated by numerically integrating the particle dynamics forward in time. The Hamiltonian ordinary differential equations are discretised in time, and then integrated forward. We choose a time step for the integration, and the method of numerical integration. The standard choices would generally be Euler's method, or a second order improvement, such as second order Runge–Kutta.

The factors that affect the computation of the diffeomorphism are the number of point particles and test particles (see §3.5), the number of time steps, and the order of the integrator (how errors accumulate during the integration). One improvement that can be considered is that the equations of motion (4, 5) are Hamiltonian and their flow is therefore symplectic. In long time simulations of Hamiltonian systems (such as celestial and molecular mechanics) it has been found extremely advantageous to use symplectic integrators, which preserve the symplectic structure. This leads to good energy behaviour and a lack of dissipation. In fact, some implementations of image registration by diffeomorphisms used symplectic integrators, because calculating geodesics by minimising a discrete path length does give a symplectic integrator [10]. However, the diffeomorphism itself, calculated from the motion of the test particles, has never been done symplectically. Using a symplectic integrator on this problem is work in progress, particularly as it requires an implicit symplectic integrator. The computational costs look promising, requiring only a constant time overhead, although the benefits of the symplectic condition itself are less clear.

3.5 Test particles and interpolation

We can induce the value of the actual diffeomorphism $\phi(x)$ by the current geodesic by placing particles with zero momentum (so $q(0) = x$, $p(0) = 0$) at the centre point of each pixel in the image, and computing its trajectory under the induced velocity field (that is, solving the ODE $\dot{q} = f(q, t)$). These are known as test particles.

Assuming that the deformation is not too large, we make some computational savings by placing a test particle every k pixels, and interpolating ϕ between the test particles. This saves a factor of k^2 computations, but changes the computed diffeomorphism from the exact one that relates to the flow (indeed, it may actually stop the warp being diffeomorphic, although this does not seem to be a problem in general). Using $k = 4$ and bi-linear

interpolation has negligible effect on the accuracy in real registrations.

3.6 Choice of metric

Inherent in the choice of Green's function $G(r)$ is a choice of the metric under which the particle dynamics occur. There is complete freedom of choice over this metric. By far the most common choice to date for image registration has been to use a Gaussian metric, that is, Green's function $G(r) = (1/\epsilon^2) \exp(-r^2/\epsilon^2)$, where ϵ is the length scale in the metric. The role of this length scale is important. Clearly, if it is set too small (say smaller than the pixel spacing) then the kernels will not overlap, and the movement of each particle will be entirely independent of the rest of the image. This will require the number of point particles to tend to infinity to represent an arbitrary diffeomorphism. One way around this problem is to use a function such as the clamped-plate spline [10], which uses a bounded domain with strict boundary conditions to avoid the problem of defining a length-scale; effectively the length-scale is the size of the image. The metric used for the clamped-plate spline is of the form ∇^{2k} . In the results reported in section 4, a value of $\epsilon = 1$ was used, which is half the width of the image (which is scaled into $[-1, 1]^2$).

3.7 Adding more points

In our implementation we position new point particles for further levels of optimisation using the discrepancy image. This uses the objective function (here the sum-of-squares error) to find regions where the two images do not match, and then placing new point particles there, initialising their momenta to zero.



FIGURE 1: The registration of the two hands. The reference image, together with the positions of the particles and their momenta are shown on the left, the final result is shown in the middle, and the effect of the warp on a grid is shown on the right.

4 Experiments

We now present two different image registrations. The first is of a pair of hand images, while the second are two 2D T1-weighted MR scans of the human brain. The hand registration was initialised using 9 point particles, positioned in a 3×3 grid on the image. The optimiser ran for 40 iterations before converging, and then an additional 7 points were added to the image using the method described by Marsland and Twining [9], which adds points where the error is currently large (the positions of the particles can be seen in Figure 1, together with the positions of the points and the initial momenta on the reference image, the final output, and the effect of the warp on a regular grid). Another 37 iterations were then performed by the optimiser, with the final result being that shown. Figure 2 provides a different way to interpret the results, showing a chequer-board overlay of the two images before and after the registration, as well as the change. See that even after this relatively small amount of computation, the registration is very good. Computing this registration took 251 seconds on a 1.8 GHz G5 Apple Macintosh, compared



FIGURE 2: Chequer-board plots showing the difference between the initial images of the hands (left), the final images (centre), and the change between the initial and final versions of the free image (right).

to well over an hour for the optimisation version, with results being slightly better (that is, the objective function was lower) for the particle method.

Figure 3 show a sample registration of two brains. A set of 10 points were positioned evenly around the skull, and the result optimised for 20 iterations. Following this, an additional set of 11 knot points, with 50 iterations of optimisation then being performed. This registration took under 7 minutes on the same computer, and it can be see that the final result is not bad. There is still work to be done on the interior (and further optimisations do indeed correct this), but the skull and major structures have all been brought into alignment.

These results are generally much faster than using an optimisation method for finding the diffeomorphism—the method described by Marsland and Twinning [10] took just under two hours to perform the brain registration described above. One of the main reasons is that those optimisation methods take many more time steps to find the diffeomorphism, usually 20 time steps are used to guarantee a diffeomorphism. With our current method, for relatively small deformations, only one time step is needed.




FIGURE 3: Chequer-board plots showing the different between the initial images of the T1-weighted brains (left), the final images (centre), and the change between the initial and final versions of the free image (right). The registration has lined up the skulls and the major structures within the brain, but there is still more fine-scale work to be done.

5 Conclusions and open questions

We have presented a method of performing diffeomorphic image registration using the methods of discrete mechanics and optimal control. The implementation described performs high quality registrations in reasonably short computational time—orders of magnitude less than using energy minimisation methods. It is still a matter for debate whether diffeomorphic methods are suitable for general image registration. However, for applications where it is *variation* that is of interest, for example in disease diagnosis or measurement of anatomical variability, the access to a right invariant Riemannian metric on the diffeomorphism group makes diffeomorphic registration methods essential.

There are a great many unanswered questions and areas for future research. We are particularly interested in the dynamical behaviour of the Euler equations on the diffeomorphism group, and how it relates to point vortices in fluid dynamics, which act on the volume preserving subgroups.

However, with regard to using the method for image registration, there are also several areas for further work. Firstly, we are currently investigating the use of the midpoint rule symplectic integrator, as discussed above, and a second question that we highlighted earlier in the article is that of a suitable choice of metric. In this article we used the Gaussian metric, which is equivalent to H^k , $k \rightarrow \infty$. These H^k metrics have the form $(1 - \epsilon^2 \nabla^2)^k$, which has an inherent length-scale ϵ . The role of both the metric and the length-scale need further investigation—there is no guarantee that the Gaussian and any chosen length scale are the correct choices for any particular problem, and some methods of comparing the results of using different metrics on a set of different problems is needed.

Acknowledgments We are grateful to Jan Modersitzki for the use of the hand images and to the Royal Society of New Zealand Marsden fund and NZIMA for their financial support.

References

- [1] V. I. Arnold. Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. *Annales de L'Institut Fourier (Grenoble)*, 16(1):319–361, 1966. **C3**
- [2] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157, February 2005. doi:10.1023/B:VISI.0000043755.93987.aa **C4**
- [3] V. Camion and L. Younes. Geodesic interpolating splines. In M. Figueiredo, J. Zerubia, and A. K. Jain, editors, *EMMCVPR'01*,

- volume 2134 of *Lecture Notes in Computer Science*, pages 513–527. Springer-Verlag, 2001. C4
- [4] J. Gee, M. Reivich, and R. Bajcsy. Elastically deforming 3D atlas to match anatomical brain images. *Journal of Computer Assisted Tomography*, 17(2):225–236, 1993.
[doi:10.1097/00004728-199303000-00010](https://doi.org/10.1097/00004728-199303000-00010) C2
- [5] U. Grenander and M. Miller. Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics*, 56(4):617 – 694, 1998. C2
- [6] D. D. Holm, J. E. Marsden, and T. S. Ratiu. The Euler–Poincaré equations and semidirect products, with applications to continuum theories. *Advances in Mathematics*, 137(1):1–81, 1998.
[doi:10.1006/aima.1998.1721](https://doi.org/10.1006/aima.1998.1721) C4
- [7] D. D. Holm, J. T. Ratnanather, A. Trouvé, and L. Younes. Soliton dynamics in computational anatomy. *NeuroImage*, 23(Supplement 1):S170–S178, 2004.
[doi:10.1016/j.neuroimage.2004.07.017](https://doi.org/10.1016/j.neuroimage.2004.07.017) C3
- [8] S. Joshi and M. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing*, 9(8):1357–1370, 2000. [doi:10.1109/83.855431](https://doi.org/10.1109/83.855431) C4
- [9] S. Marsland and C. Twining. Constructing data-driven optimal representations for iterative pairwise non-rigid registration. In *Second International Workshop on Biomedical Image Registration*, number 2717 in *Lecture Notes in Computer Science*, pages 50–60. Springer, 2003. C6, C10
- [10] S. Marsland and C. Twining. Constructing diffeomorphic representations for the groupwise analysis of non-rigid registrations of medical images. *IEEE Transactions on Medical Imaging*,

- 23(8):1006–1020, 2004. doi:10.1109/TMI.2004.831228 C2, C3, C4, C5, C8, C9, C11
- [11] M. Miller, A. Trouvé, and L. Younes. On metrics and the Euler–Lagrange equations of computational anatomy. *Annual Reviews in Biomedical Eng.*, 4:375 – 405, 2002. C3
- [12] D. Mumford. *Pattern Theory and Vision*, chapter 3, pages 7–13. Institute Henri Poincaré, Paris, 1998. C4
- [13] D. Rueckert, L. Sonoda, C. Hayes, D. Hill, M. Leach, and D. Hawkes. Non-rigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999. doi:10.1109/42.796284 C2
- [14] A. W. Toga. *Brain Warping*. Academic Press, 1999. C2
- [15] P. Viola and W. M. Wells III. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24(2):137–154, 1997. doi:10.1023/A:1007958904918 C3

Author addresses

1. **Robert I. McLachlan**, Institute of Fundamental Sciences, Massey University, Palmerston North, NEW ZEALAND
<mailto:r.mclachlan@massey.ac.nz>
2. **Stephen Marsland**, Institute of Information Sciences & Technology, Massey University, Palmerston North, NEW ZEALAND
<mailto:s.r.marsland@massey.ac.nz>