Large Deformation Diffeomorphic Registration Using Fine and Coarse Strategies

Laurent Risser[1,2], François-Xavier Vialard[1], Maria Murgasova[2], Darryl Holm[1], and Daniel Rueckert[2]

[1] Institute for Mathematical Science, Imperial College London, 53 Prince's Gate, SW7 2PG, London, UK*

[2] Visual Information Processing, Imperial College London, Huxley Building, Department of Computing, SW7 2BZ, London, UK

Abstract. In this paper we present two fine and coarse approaches for the efficient registration of 3D medical images using the framework of Large Deformation Diffeomorphic Metric Mapping (LDDMM). This formalism has several important advantages since it allows large, smooth and invertible deformations and has interesting statistical properties. We first highlight the influence of the smoothing kernel in the LDDMM framework. We then show why approaches taking into account several scales simultaneously should be used for the registration of complex shapes, such as those treated in medical imaging. We then present our fine and coarse approaches and apply them to the registration of binary images as well as the longitudinal estimation of the early brain growth in preterm MR images.

1 Introduction

Non-rigid image registration has various applications such as motion tracking, shape comparison, atlas creation or image segmentation. Recent years have seen the development of new non-rigid registration techniques allowing large diffeomorphic deformations. Diffeomorphic deformations are by definition smooth and invertible, properties that are highly desirable in image registration and that most of the classical registration techniques have. Importantly, since the pioneering work of [12], an increasing number of registration techniques transform the images using the concept of deformation flow characterized by a velocity vector field. This makes possible large deformations while preserving diffeomorphic properties. Note that the velocity fields can be either steady or time-dependent. Such approaches have led to new problems: First, how to find the optimal deformation flow between two shapes and secondly how to regularize spatially the deformations. Similarly, the issues of the computational complexity and memory requirements are also important, especially in the context of 3D medical image registration.

* We thank the Imperial College Strategic Initiative Fund for partial support. The work of DDH was also partially supported by the Royal Society of London Wolfson Research Merit Award.

B. Fischer, B. Dawant, and C. Lorenz (Eds.): WBIR 2010, LNCS 6204, pp. 186–197, 2010.

In the work of [12], the registration was formulated in a Bayesian setting. Statistical prior distributions of the deformations were modelled using stochastic PDEs controlling the displacement field according to driving forces. Later, the framework of Large Deformation Diffeomorphic Metric Mapping (LDDMM) [21, 23, 4, 27, 18], in which the optimal velocity fields are time-dependent and geodesic, was developed. Finding geodesic flows between two registered shapes is fundamental in the framework of LDDMM since it ensures that the optimal flow is the shortest path between the shapes according to a metric regularizing the deformation. We will discuss later the influence of this metric. The LDDMM has therefore convenient properties for the statistical comparison of images and shapes as well as the creation of atlases. A practical implementation of the LDDMM for image registration, considered as the reference, has been proposed in [4]. It solves the registration between two images in an Euler-Lagrange framework using a gradient descent to minimize the registration energy as a function of the velocity vector field of the deformation flow. Algorithms based on [4] have been applied to inter-subject local shape comparison or atlas creation. In [6, 16], the approach has been used to measure shape variations between segmented hearts, in order to highlight the structural remodelling of dyssynchronous failing heart. In [9, 8] the LDDMM has been extended to vector- and tensor-valued images. Finally a symmetric extension of [4] has been proposed in [5]. However, despite its interesting statistical properties, the LDDMM approach is particularly time and memory consuming. Similarly, although it were designed to allow very large deformations, its practical use remains limited to relatively small deformations in the literature.

Interestingly, alternatives to the LDDMMs, faster, requiring less memory or adapted to multi-modal images have been proposed: A symmetric interpretation of [4] using cross correlation to measure the image similarity between the source and target images was proposed in [3]. This interpretation was used in [11] to measure the cortical grey matter thickness in segmented brain images. Another interpretation [19] allowing multimodal registration for atlas creation estimates the transformations in a Bayesian framework. Correspondences between the underlying tissue classes are found by using Kullback-Leibler divergence on the space of posteriors probabilities. More recently, [4] has been formulated as an optimal control problem in [15] leading to an improvement of the convergence speed and robustness. Note finally that interesting approaches making use of the Navier-Stokes equation of fluid dynamics, have been proposed [10]. Such approaches allow large diffeomorphic deformations but are not designed to provide geodesic transformations.

Another class of large diffeomorphic registration techniques, using stationary velocity fields, emerged with [1]. Such parameterizations, have been applied as an evolution of the LDDMM framework [2, 17] and an extension of the demons algorithm [25]. Stationary parameterizations are efficient in terms of memory required and computational time while providing registrations similar to those obtained using time-dependent velocity fields in most cases. However, the optimal flows found using these techniques are usually not geodesic at convergence

except for very simple deformations. Statistical comparisons may therefore be less valid using these techniques than what can be provided by the LDDMM framework and the estimation of mean shapes out of an atlas is not as straightforward as using geodesic deformations. Moreover, the range of possible deformations is theoretically limited compared to time-varying velocity fields.

In an attempt to extend the utility of the LDDMM framework in the context of 3D medical image registration, we discuss in this paper one of its fundamental aspects: the choice of the smoothing kernel. Indeed, the kernel is directly related to the metric of the deformations and therefore controls its spatial regularization. In practice, small kernels favour deformations that match local details and large kernels favour deformations that match global structures. Since 3D medical images often contain complex shapes, such as the cortical surface of the brain, small kernels may provide unsatisfactory deformations. Similarly, commonly used coarse-to-fine strategies, are inappropriate to register complex shapes in the LDDMM framework. In addition, statistics on the deformations depend on the kernel used at the finest scale. Moreover, due to the use of gradient descent during the search for the optimal path, first at large scales and then at small scales, the algorithm can converge to local minima that would be unreachable by using only small kernels. Deformation statistics are biased in that case. In this paper we present two different multiscale extensions of [4]: The first one consists of the use of kernels that are the sum of Gaussian kernels of different scales while the second one consists in using a series of such kernels. Note, that the second extension can be related to the work of J. Glaunes where time-dependent kernels were used in the context of surface registration [13]. Similarly, the idea underlying these extensions is close to the coarse to fine strategy developed in [14]. Though these natural ideas might not be new in the literature they deserve a detailed comparison as provided in this paper.

In section 2, we present the method of [4] and discuss the influence of the kernel. Two multiscale extensions are then presented in section 3 and tested on synthetic and real images in section 4. Finally, the methodology and tests are discussed in section 5.

2 3D Image Registration Using LDDMM

2.1 Registration Technique

We give here an overview of the LDDMM approach and the classical algorithm to find optimal registrations described in [4]. This framework enables the registration of a source image I_S on a target image I_T defined on a spatial domain Ω through a time dependent diffeomorphic transformation of Ω, ϕ_t, $t \in [0,1]$. Such a transformation is generated by a time dependent velocity field v as follows:

$$\partial_t \phi_t = v_t(\phi_t), \quad t \in [0,1], \tag{1}$$

where $\phi_0 = Id$. The velocity field v_t deforms the image coordinates at time t and ϕ_t is the induced deformation. For notational convenience we introduce

$\phi_{t,s} \doteq \phi_s \circ \phi_t^{-1}$. The LDDMM framework assumes the velocity field at each time to be smooth enough so that the flow ϕ_t is well defined, as described in [26]. The registration problem then consists in finding the velocity field v_t that minimizes the energy $E(v)$, defined by:

$$E(v) = \int_0^1 \frac{1}{2}||v_t||_V^2 \mathrm{d}t + \frac{1}{2}||I_S \circ \phi_1^{-1} - I_T||_{L^2}^2. \tag{2}$$

The similarity measure here is the sum of the squared differences between the intensities in the target and the deformed image. The time dependent velocity field v is assumed to lie in $L^2([0,1],V)$, where V is a Hilbert space of vector fields, and the norm on V can be any norm which satisfies for any $u \in V$: $\|u\|_{1,\infty} \leq m\|u\|_V$, for a positive constant m (we denote by $\|\,.\,\|_{1,\infty}$ the sup norm on the vector field and its first derivative). Importantly, underlying the space of velocity fields V, there exists a smooth matrix-valued kernel $k(.,.)$ which describes the velocity fields that can be used for the registration. Conveniently, the space V can be defined from this kernel as the completion of the linear space spanned by $v(x) = \sum_{i=1}^n k(x,y_i)p_i$, where $y_i \in \Omega$ and $p_i \in \mathbb{R}^d$ (d the spatial dimension of Ω) with respect to the norm defined as $\|v\|_V^2 = \sum_{i,j=1}^n p_j^T k(y_j,y_i)p_i$. Therefore any vector field in V is efficiently approximated by using a finite sum of elementary vector fields $k(.,y)p$. Interestingly, there exists a wide family of available kernels. Among various properties, the space of admissible kernels must be stable under addition and under multiplication by a positive constant. For critical computational issues, we consider here Gaussian kernels that are translation invariant ($k(x,y) = K(x-y)$) and separable:

$$K(x) = (2\pi)^{-d/2}|\Sigma|^{-1/2} \exp\left(-\frac{1}{2}x^T \Sigma^{-1} x\right), \tag{3}$$

where Σ is the covariance matrix. We restrict our study to isotropic covariances, *i.e.* such that $\Sigma = \sigma Id_{\mathbb{R}^d}$. The key parameter σ then controls the spatial correlation of the deformations. We discuss this point in Subsection 2.2.

The minimization algorithm is described hereafter. We denote $J_t^S = I_S \circ \phi_{t,0}$, $J_t^T = I_T \circ \phi_{1,t}$ and $|D\phi_{t,1}|$ the Jacobian of $\phi_{t,1}$ at time t. We consider an homogeneous discretization of the time $t_i = i\Delta t$, $i \in \{1, \cdots, I\}$. The minimization of the variational problem of Eq. 2 is performed by using a steepest gradient descent approach. Practical resolution then involves the iterative use of the gradient of the functional E, denoted by $\nabla_v E_t$, in the space $L^2([0,1],V)$ at time t:

$$K \star \nabla_v E_t = v_t - K \star \left(|D\phi_{t,1}^v| \nabla J_t^S (J_t^S - J_t^T)\right), \tag{4}$$

where $\star$ denotes the convolution operator. The velocity field is then updated by computing:

$$v^{k+1} = v^k - \epsilon K \star \nabla_{v_{t_j}^k} E\,, \tag{5}$$

where ϵ is the step size during the gradient descent step. The optimal time dependent diffeomorphism should be a geodesic path in the group of diffeomorphisms for which the associated velocity field satisfies the Euler-Lagrange equation:

$$\widehat{v}_t - K \star \left(|D\phi^{\hat{v}}_{t,1}| \nabla J^S_t (J^S_t - J^T_t) \right) = 0, \quad \forall t \in [0,1]\,. \tag{6}$$

In such a case, the optimal path has shooting properties [20] from the velocity field at time $t = 0$ that can be used to statistically compare shapes [24]. We emphasize that these statistics depend critically on the choice of the metric on V as shown in the next subsection.

2.2 Influence of the Kernels

As discussed briefly in the Introduction, the kernel controls the spatial behaviour of the deformations. For isotropic Gaussian kernels of standard deviation σ, the parameter σ is the characteristic length defining the scale at which the registration is performed. To give a practical interpretation of the influence of σ, let us focus on the right hand side of Eq. 4, which provides the energy gradient that is used to update the velocity field at each iteration (Eq. 5). Its second term $|D\phi^{v}_{t,1}| \nabla J^S_t (J^S_t - J^T_t)$ pushes locally the source image I_S onto the target image I_T since the direction is defined by the gradient ∇J^S_t. This term is then smoothed by using the filter K so the deformations are more or less correlated according to the choice of σ. Figures 2 and 3 illustrate the influence of the kernel on the 2D binary images I^1_S, I^1_T, I^2_S, I^2_T of Fig. 1. In Fig. 2, we register the

Fig. 1. 2D binary images to illustrate the influence of the kernel. **From left to right:** Source image I^1_S, target image I^1_T, source image I^2_S, target image I^2_T

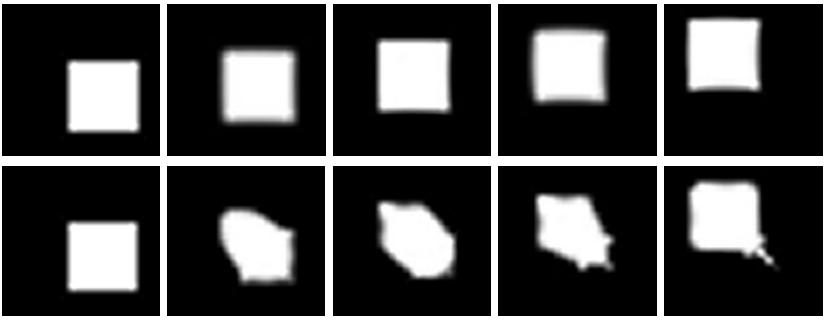

Fig. 2. Deformation of I^1_S to I^1_T for (from left to right) $t = 0$, $t = 0.25$, $t = 0.5$, $t = 0.75$, $t = 1.0$. **Row 1:** I^1_S to I^1_T using a large kernel ($\sigma = 8$ pixels). **Row 2:** I^1_S to I^1_T using a small kernel ($\sigma = 2$ pixels).

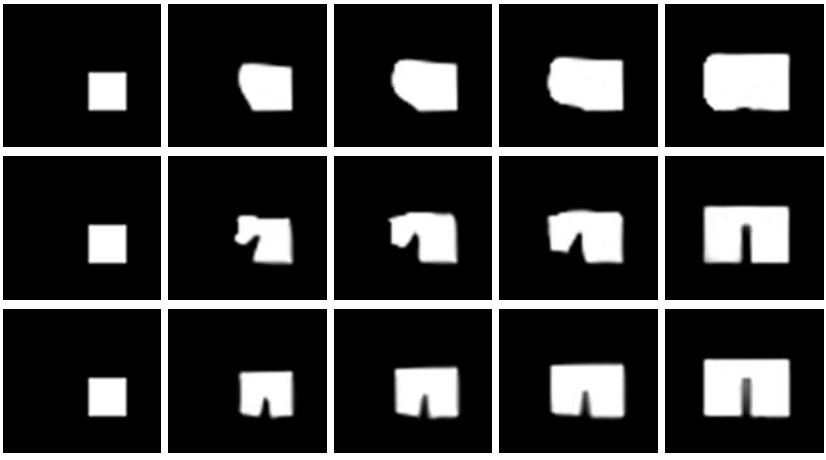

Fig. 3. Deformation of I_S^2 to I_T^2 for (from left to right) $t = 0$, $t = 0.25$, $t = 0.5$, $t = 0.75$, $t = 1.0$. **Row 1:** I_S^2 to I_T^2 using a large kernel ($\sigma = 5$ pixels). **Row 2:** I_S^2 to I_T^2 using a small kernel ($\sigma = 1.5$ pixels). **Row 3:** I_S^2 to I_T^2 registration with a small kernel ($\sigma = 1.5$ pixels) and initialized by the output of Row 1.

square I_S^1 on its translated I_T^1. In order to obtain a translation-like deformation using LDDMMs, the standard deviation σ of the gaussian kernel must be of a size at least similar to the size of the square ($\sigma = 8$ pixels here). Using a smaller kernel ($\sigma = 2$ pixel), the square is not considered as a whole object but as a set of details. Therefore some parts of I_S^1 are compressed and other are expanded during the registration. In Fig. 3, the square I_S^2 is now registered on a rectangle containing a slot I_T^2. As expected, at a large scale ($\sigma = 5$ pixels) the square is registered on the rectangle with a small perturbation while at a small scale ($\sigma = 1.5$ pixels) the registration fully takes into account the slot. Note that, in the latter case, the registration makes use of invertible large deformations. We show the corresponding deformation grid at $t = 1$ in Fig. 5. Even though the transformation appears non-invertible, a closer inspection shows that the deformation is actually invertible but with very high values of the Jacobian. Such behaviour is not desirable. We have registered I_S^2 on I_T^2 using the same small kernel as in row 1 by taking as initial guess the flow of row 2, computed with a large kernel, instead of a null flow. Theoretically, the initial guess should have no influence on the final result, so the estimated flows should be the same in rows 2 and 3. However, even though the final deformations look similar, the estimated flows are completely different, the one obtained with the coarse to fine technique obviously not enjoying shooting properties at $t = 0$. This deformation is therefore not the global optimum. The convergence to a local minimum is due here to the use of a gradient descent. Hence, in the context of LDDMM, the size of the kernels has a strong influence on the scale of the registration. Moreover, the use of inappropriate kernels or coarse to fine strategies can lead to unrealistic deformations. To tackle the presented issues in the context of medical image

registration, where the registered shapes may contain several scales of interest, we propose two fine and coarse extensions of [4] in the next section.

3 Multi-kernel LDDMM

3.1 Sum of Gaussian Kernels

An immediate extension of [4] to simultaneously perform the registration at several scales is to define the kernel K as the sum of several kernels having different scales. Here, instead of using the kernel of [4] or a simple Gaussian function, we build K as the sum of N Gaussian kernels as follows:

$$K(x) = \sum_{n=1}^{N} a_n K_{\sigma_n}(x) = \sum_{n=1}^{N} a_n (2\pi)^{-3/2} |\Sigma_n|^{-1/2} \exp\left(-\frac{1}{2} x^T \Sigma_n^{-1} x\right), \quad (7)$$

where Σ_n and a_n are respectively the covariance matrix and the weight of the n^{th} Gaussian function. An important property of reproducing kernel Hilbert spaces is:

$$|w|^2_{K_1+K_2} = \inf\left\{|u|^2_{K_1} + |v|^2_{K_2} \,\middle|\, w = u + v \text{ with } u \in H_{K_1} \text{ and } v \in H_{K_2}\right\}. \quad (8)$$

The optimization is then performed simultaneously at the fine and coarse scales. Note that an extension of this idea to the space of diffeomorphisms can be argued according to [7]. Equation (7) then allows one to construct a wide range of kernels with several scales of interest while preserving all the promising statistical properties of the LDDMM. Note that the choice of the weights $(a_n)_{n\in[1,N]}$ is a key issue here since it controls the influence of the structures at different scales for the deformation. For instance, one would want an equivalent influence of the large and small structures in the registration. This point is discussed in Subsection 3.3.

3.2 Chain of Gaussian Kernels

Our second extension of [4] consists of using a chain of N deformations between I_S and I_T, each having its own kernel. The idea of a time dependent kernel already appeared in [13] for surface registration where the width of the kernel was chosen C^1 and decreasing in time. We still register I_S on I_T through a time dependent diffeomorphic transformation ϕ_t, where $t \in [0, 1]$, related to the velocity vector field v_t by equation (1). Here, however, we minimize the following energy as a function of v_t, $t \in [0, 1]$:

$$\underset{v=v_t, t\in[0,1]}{\operatorname{argmin}} \frac{1}{2}\|I_S \circ \phi_1^{-1} - I_T\|^2_{L^2} + \sum_{n=1}^{N} \int_{(n-1)/N}^{n/N} \frac{1}{2}\|v_t\|^2_{V_{\sigma_n}} \mathrm{d}t \quad (9)$$

where $||.||_{V_{\sigma_n}}$, $n \in \{1, \cdots, N\}$ represents the norm related to Gaussian kernels $K_{\sigma_n}(x)$ of width σ_n and weighted by a_n. As in a coarse-to-fine approach,

$(\sigma_n)_{n\in[1,N]}$ decreases as n increases, which implies that the sequence of Hilbert spaces is increasing for the inclusion. Therefore the group of diffeomorphisms is the same as the one generated by V_{σ_N}. However, the cost of the transformation in (9) does not give a Riemannian metric on this group contrary to the classical framework and it does not even give in general a distance due to this non-symmetric cost. Despite these drawbacks, the shooting property of the initial setting is still conserved. Implementation of this scheme differs slightly from that of [4], as described below. Table 1 presents the resolution algorithm where, for readability, we recall that the symbol $\forall$ means "for all". Note that, in this algorithm, the gradient descent parameter ϵ must be small enough to ensure the invertibility of the deformation.

Table 1. Gradient descent for multi-kernel LDDMM

(1) Initialize the velocity field $v(x,t_i)=0$, $\forall x$, $\forall i$.
(2) While not convergence
(3) Estimate $\phi(x,t_i)$ and $\phi^{-1}(x,t_i)$ by forward/backward integration of v, $\forall x$, $\forall i$.
(4) For $i=1:I$
(5) Estimate $J^S_{t_i}(x)$ and $J^T_{t_i}(x)$ by using ϕ and ϕ^{-1}, $\forall x$.
(6) Compute $M(x)=|D\phi(x,t_i)|\nabla J^S_{t_i}(x)(J^S_{t_i}(x)-J^T_{t_i}(x))$, $\forall x$.
(7) Smoothing: $M_s(x,t_i)=K_{\sigma_n}(M(.))$, $\forall x$ and where $\frac{n-1}{N}\leq i\Delta t<\frac{n}{N}$.
(8) Update: $v(x,t_i)=v(x,t_i)-\epsilon(v(x,t_i)-M_s(x,t_i))$, $\forall x$, $\forall i$.

3.3 Weight of the Kernels

The fine and coarse registration techniques presented in Subsections 3.1 and 3.2 depend on a set of parameters a_n, $n\in[1,N]$ each of them controlling the weight of the deformations at scale n. The deformations are strongly related to the velocity field updates iteration after iteration. These updates do not only depend on the values of a_n but also on the kernels as well as the registered images. To set the weights, we then introduce the apparent weights a'_n, $n\in[1,N]$ such that $a_n=a'_n/g(K_{\sigma_n},I_S,I_T)$ where g represents the typical amplitude of the velocity field updates for a given smoothing kernel and two registered images. We then set g as the norm of the maximum gradient computed (Eq. 4) at the first iteration of the algorithm, when registering I_S on I_T using K_{σ_n}:

$$g(K_{\sigma_n},I_S,I_T)=K_{\sigma_n}\star(\nabla I_S(I_S-I_T))\,. \tag{10}$$

In the context of fine and coarse registration, all a'_n should then be similar in order to have equivalent deformations at each scale considered. If the value of a'_n is significantly higher than the other apparent weights, the registration will be almost the same as the registration at the scale n only. This technique has the advantage of being simple to use and was shown to be efficient both on 2D synthetic images and 3D CT and MR images. Importantly, for images of the same type (*e.g.* MR brain images with the same acquisition protocol) and the same kernels K_{σ_n}, the values of g were observed to be stable. The method can then be used for an atlas creation by systematically using the same kernels with same weights.

4 Results

4.1 Evaluation on Synthetic Images

We evaluate here the behaviour of the two techniques presented in Section 3 in the registration of I_S^2 on I_T^2 (cf. Fig. 4). Note that the interpretation of the deformations of I_S^1 on I_T^1 is similar. We use the characteristic scales to $\sigma_1 = 5$ and $\sigma_2 = 1.5$ pixels, the values at which the registration were performed using simple Gaussian kernels in Subsection 2.2, and equal apparent weights at both scales. One can observe that the flow obtained using the sum of kernels (Fig. 4, 1st row) looks similar to the one obtained using small kernels (Fig. 3, 2nd row) for an almost equivalent overlap between the deformed and target images. The flow obtained using the chain of kernels (Fig. 4, 2nd row) matches the shape at a large scale between $t = 0$ and $t = 0.5$ and then matches the details between $t = 0.5$ and $t = 1$ and the final matching is good even if slightly less accurate than using a small kernel or the sum of kernels. The comparison of grid deformations at $t = 1$ in Fig. 5 is particularly interesting. It shows that the deformations obtained using the multi-kernel approaches are smoother and visually more natural than using small kernels. This key property and the good final matching are due to the simultaneous consideration of two scales that are pertinent to the registered shape.

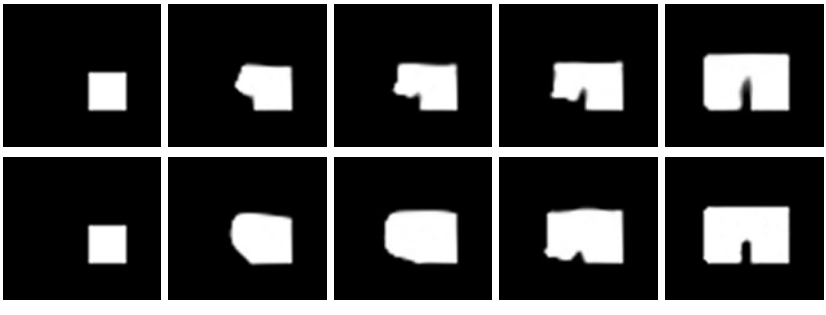

Fig. 4. Deformation of I_S^2 to I_T^2 for (from left to right) $t = 0$, $t = 0.25$, $t = 0.5$, $t = 0.75$, $t = 1.0$ by using sum (Row 1) and chain (Row 2) of kernels, each having $\sigma_1 = 5$ and $\sigma_2 = 1.5$ pixels

4.2 Evaluation on Brain MR Images

We apply now our techniques to the longitudinal estimation of the early brain development out of MR brain images. Here we limit our study to the comparison of two brains of 29.86 and 33.86 weeks of gestational age for the same preterm infant. The images have a spatial resolution of 0.85mm and bias field correction has been performed using N3 [22]. Two characteristic lengths are considered here, a large one ($\sigma_1 = 5$mm) and a small one ($\sigma_2 = 1.5$mm) and equal apparent weights are used for the fine and coarse deformations. As shown in Fig. 6, using a

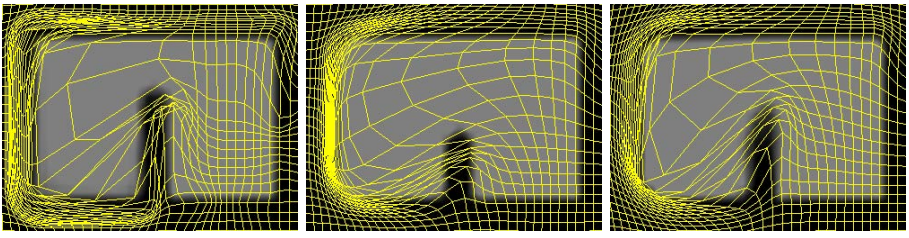

Fig. 5. Deformation grids at time $t = 1$ of the registration of I_S^2 to I_T^2 using a small kernel (left) the chain (center) and the sum (right) of kernels. Each square of the grid represents a deformed pixel.

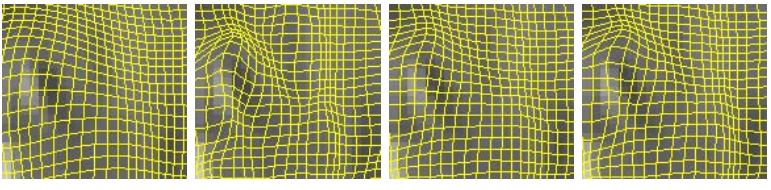

Fig. 6. Effect of the kernel on the deformation of a MR brain image. Grid size is one voxel. **From left to right:** large kernel ($\sigma_1 = 5$mm), small kernel $\sigma_2 = 1.5$mm), chain of the two kernels, sum of the two kernels.

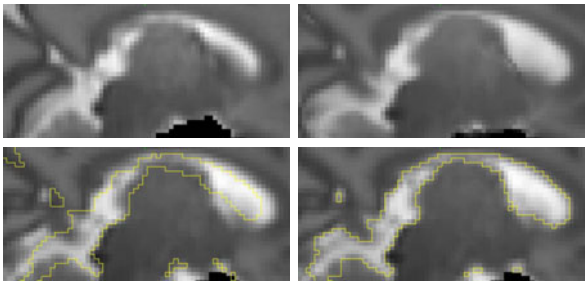

Fig. 7. Registration using the sum of kernels ($\sigma_1 = 5$mm, $\sigma_2 = 1.5$mm). **Top-left:** Source image. **Top-right:** Deformed source image at $t = 1$. **Bottom-left:** Target image and isovalues of the source image showing the surface of a lateral ventricle. **Bottom-right:** Target image and isovalues of the deformed source image at $t = 1$ showing the surface of a lateral ventricle.

small kernel leads to unnatural looking deformations while the deformations look more plausible when using a large kernel. The matching is however much better using small kernels instead of large ones. The multi-kernel approaches provide visually more natural deformations that also match the details. We can observe that they deform the grid at fine and coarse scales simultaneously. In Fig. 7, we can first observe that the deformations obtained using the sum of kernels look

natural and are accurate. In comparison, when using the chain of kernels with the same number of iterations the matching seems slightly less accurate both at large and small scales. It appears therefore that the sum of kernels produces more accurate deformations and higher statistical power.

5 Conclusion

We have presented examples of the use of fine and coarse approaches in the context of LDDMM for the registration of 3D medical images. The approaches we developed make use of either a time dependent kernel, or a constant kernel defined by the sum of several Gaussians. Our tests have shown that these methods estimate natural deformations on complex images with diffeomorphic properties. In particular, using the sum of Gaussian kernels leads to natural-looking, accurate registrations that have a strong statistical power, even on complex images. Future work will pursue the development of extensions with other similarity measures and use of a multi-resolution approach. More experiments and applications will also be carried out on MR cerebral images, as well as on CT cardiac images.

References

1. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-Euclidean framework for statistics on diffeomorphisms. In: MICCAI 2006, Part I. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006)
2. Ashburner, J.: A fast diffeomorphic image registration algorithm. NeuroImage 38, 95–113 (2007)
3. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. Medical Image Analysis 12, 26–41 (2008)
4. Beg, F.M., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision 61(2), 139–157 (2005)
5. Beg, M., Khan, A.: Symmetric data attachment terms for large deformation image registration. IEEE Transactions on Medical Imaging 26(9), 1179–1189 (2007)
6. Beg, M.F., Helm, P.A., McVeigh, E., Miller, M.I., Winslow, R.L.: Computational cardiac anatomy using MRI. Magnetic Resonance in Medecine 52(5), 1167–1174 (2004)
7. Bruveris, M., Gay-Balmaz, F., Holm, D., Ratiu, T.: The momentum map representation of images. J. Nonlin. Sci. (February 20, 2010) (submitted), JNLS-D-10-00024
8. Cao, Y., Miller, M.I., Mori, S., Winslow, R.L., Younes, L.: Diffeomorphic matching of diffusion tensor images. In: CVPRW 2006: Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop, p. 67 (2006)
9. Cao, Y., Miller, M.I., Winslow, R.L., Younes, L.: Large deformation diffeomorphic metric mapping of vector fields. IEEE Trans. Med. Imaging 24(9), 1216–1230 (2005)
10. Crum, W., Tanner, C., Hawkes, D.: Anisotropic multi-scale fluid registration: evaluation in magnetic resonance breast imaging. Physics in Medicine and Biology 50(21), 5153–5174 (2005)

11. Das, S.R., Avants, B.B., Grossman, M., Gee, J.C.: Registration based cortical thickness measurement. NeuroImage 45, 867–879 (2009)
12. Dupuis, P., Grenander, U., Miller, M.I.: Variational problems on flows of diffeomorphisms for image matching. Q. Appl. Math. LVI(3), 587–600 (1998)
13. Glaunes, J.: Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l'anatomie numérique. Ph.D. thesis, Université Paris 13 (2005)
14. Haber, E., Modersitzki, J.: Cofir: coarse and fine image registration. In: SIAM Real-Time PDE-Constrained Optimization, pp. 37–49 (2007)
15. Hart, G., Zach, C., Niethammer, M.: An optimal control approach for deformable registration. In: Computer Vision and Pattern Recognition Workshop, pp. 9–16 (2009)
16. Helm, P., Younes, L., Beg, M., Ennis, D., Leclercq, C., Faris, O., McVeigh, E., Kass, D., Miller, M., Winslow, R.: Evidence of structural remodeling in the dyssynchronous failing heart. Circulation Research 98, 125–132 (2006)
17. Hernandez, M., Bossa, M.N., Olmos, S.: Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. Int. J. Comput. Vision 85(3), 291–306 (2009)
18. Younes, L., Arrate, F., Miller, M.I.: Evolutions equations in computational anatomy. Neuroimage (November 2008)
19. Lorenzen, P., Prastawa, M., Davis, B., Gerig, G., Bullitt, E., Joshi, S.: Multi-modal image set registration and atlas formation. Med. Image. Anal. 10(3), 440–451 (2006)
20. Miller, M., Trouvé, A., Younes, L.: Geodesic shooting for computational anatomy. J. Math. Imaging Vis. 24(2), 209–228 (2006)
21. Miller, M., Younes, L.: Group actions, homeomorphisms, and matching: A general framework. International Journal of Computer Vision 41(1-2), 61–84 (2001)
22. Sled, J., Zijdenbos, A., Evans, A.: A nonparametric method for automatic correction of intensity nonuniformity in MRI data. IEEE Transactions on Medical Imaging 17(1), 87–97 (1998)
23. Trouvé, A., Younes, L.: Metamorphoses through lie group action. Foundations of Computational Mathematics 5(2), 173–198 (2005)
24. Vaillant, M., Miller, M., Trouvé, A., Younes, L.: Statistics on diffeomorphisms via tangent space representations. Neuroimage 23(S1), S161–S169 (2004)
25. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. NeuroImage 45(1), S61–S72 (2009)
26. Younes, L.: Shapes and Diffeomorphisms. Springer, Heidelberg (2008)
27. Younes, L., Qiu, A., Winslow, R.L., Miller, M.I.: Transport of relational structures in groups of diffeomorphisms. Journal of Mathematical Imaging and Vision 32, 41–56 (2008)