Diffeomorphic demons: Efficient non-parametric image registration.

Tom Vercauteren, Xavier Pennec, Aymeric Perchant, Nicholas Ayache

HAL Id: inserm-00349600

https://inserm.hal.science/inserm-00349600

Submitted on 2 Jan 2009

Diffeomorphic Demons: Efficient Non-parametric Image Registration

Tom Vercauteren[a,*], Xavier Pennec[b], Aymeric Perchant[a], Nicholas Ayache[b]

[a]*Mauna Kea Technologies,9 rue d'Enghien, 75010 Paris, France*
[b]*INRIA Sophia Antipolis, Asclepios Research Project, 2004 route des Lucioles - BP 93, 06902 Sophia Antipolis Cedex, France*

Abstract

We propose an efficient non-parametric diffeomorphic image registration algorithm based on Thirion's demons algorithm. In the first part of this paper, we show that Thirion's demons algorithm can be seen as an optimization procedure on the entire space of displacement fields. We provide strong theoretical roots to the different variants of Thirion's demons algorithm. This analysis predicts a theoretical advantage for the symmetric forces variant of the demons algorithm. We show on controlled experiments that this advantage is confirmed in practice and yields a faster convergence. In the second part of this paper, we adapt the optimization procedure underlying the demons algorithm to a space of diffeomorphic transformations. In contrast to many diffeomorphic registration algorithms, our solution is computationally efficient since in practice it only replaces an addition of displacement fields by a few compositions. Our experiments show that in addition to being diffeomorphic, our algorithm provides results that are similar to the ones from the demons algorithm but with transformations that are much smoother and closer to the gold standard, available in controlled experiments, in terms of Jacobians.

1. Introduction

Although non-rigid image registration has been a very active area of research for some time, it is still widely accepted that more work is needed. Typical research objectives aim at lowering the computational requirements, increasing the robustness, and better constraining the registration. In this paper, we build on the efficient tools presented in (Benhimane and Malis, 2004; Vercauteren et al., 2007a,c) and show that Thirion's demons algorithm (Thirion,

---

*Corresponding author.
Email addresses: `tom.vercauteren@maunakeatech.com` (Tom Vercauteren), `xavier.pennec@sophia.inria.fr` (Xavier Pennec), `aymeric.perchant@maunakeatech.com` (Aymeric Perchant), `nicholas.ayache@sophia.inria.fr` (Nicholas Ayache)

1998) can be extended to provide a very efficient non-parametric diffeomorphic image registration algorithm. To make this work available to a wide audience, we proposed in (Vercauteren et al., 2007b) an open-source implementation based on the Insight Toolkit (ITK) (Ibáñez et al., 2005). It can be downloaded at `http://hdl.handle.net/1926/510`.

We have shown in (Vercauteren et al., 2007a) that, for the problem of linear image registration, the efficient second-order minimization (ESM) algorithm (Benhimane and Malis, 2004, 2007; Malis, 2004) provides a very interesting alternative to classical schemes. Looking at non-rigid image registration, one of the most efficient methods is the demons algorithm proposed by Thirion (Thirion, 1998). Several variants of the algorithm have been proposed depending on how the *forces* are computed. In (Wang et al., 2005; Rogelj and Kovačič, 2006) an *ad hoc* symmetrization of the demons forces similar to the one proposed by Thirion was shown to improve the results of the original demons algorithm. In (Pennec et al., 1999; Cachier et al., 2003; Modersitzki, 2004), it has been shown that the demons algorithm had connections with gradient descent and variational schemes. However, to the best of our knowledge, the different variants of the demons have not been given a strong unified theoretical justification.

Our first contribution in this paper, Section 2, is to show that we can provide a strong theoretical justification to the demons algorithm and that the different variants are related to the use of different optimizers. One of the main results of this theoretical analysis we initially presented in (Vercauteren et al., 2007a) is to show that the symmetric forces variant can be related to the ESM scheme of (Benhimane and Malis, 2004; Malis, 2004). This study thus explains, from a theoretical point of view, why the symmetric forces demons algorithm appears to be more efficient in practice. *Another contribution* of this paper in Section 2.5 is to provide evidence that, in practice, using symmetric forces indeed leads to a higher convergence rate.

One of the main limitations of the demons algorithm is that it does not provide diffeomorphic transformations contrarily to the algorithms developed in (Ashburner, 2007; Avants et al., 2008; Beg et al., 2005; Hernandez et al., 2007; Marsland and Twining, 2004; Trouvé, 1998). Diffeomorphic transformations are a requirement in the growing field of computational anatomy. Invertibility is also needed for the registration of diffusion tensor images (DTI) since deforming a DTI typically uses a decomposition of the Jacobian of the spatial transformation to reorient the tensors (Yeo et al., 2008c). Such a decomposition can only be performed if the Jacobian is non-singular. Diffeomorphisms are powerful for other problems as well, as they preserve the topology of the objects and prevent from introducing folding which is often physically impossible. Finally, diffeomorphisms are considered to be a good working framework when no additional information about the spatial transformation is available. *The main contribution* of this paper, in Section 3, is to show that the alternate optimization scheme of the demons algorithm can be used in combination with the Lie group structure on diffeomorphic transformations of (Arsigny et al., 2006). In contrast to many diffeomorphic registration schemes, our diffeomorphic demons algorithm, which we initially proposed in (Vercauteren et al., 2007c), is computationally efficient.

Typical 3D MR images can indeed be registered in less than five minutes on a desktop computer. In practice, our algorithm only replaces, within the classical demons algorithm, an addition of displacement fields by a few compositions of non-parametric transformations. Our approach is evaluated in Section 4 in both a simulated and a realistic registration setup. We show that in addition to being diffeomorphic, our algorithm provides results that are similar to the ones from the demons but with transformations that are much smoother and closer to the true ones (on controlled experiments) in terms of Jacobians.

In Section 5 some of the theoretical and practical properties of our diffeomorphic demons algorithm are discussed and compared to that of some alternative approaches. Finally Section 6 concludes the paper and shows that our approach has already been extended to handle different types of images (Peyrat et al., 2008; Yeo et al., 2008c,b) and to make it symmetric with respect to the order of the input images (Vercauteren et al., 2008).

2. An Insight into the Demons Algorithm

In (Thirion, 1998), the author proposed to consider non-parametric non-rigid registration as a diffusion process. He introduced *demons* that push according to local characteristics of the images in a similar way Maxwell did for solving the Gibbs paradox. The forces are inspired from the optical flow equations (Barron et al., 1994) and the method alternates between computation of the forces and regularization by a simple Gaussian smoothing. This results in a computationally efficient algorithm compared to other non-rigid registration procedures such as those based on linear elasticity (Christensen et al., 1997). Several teams (Bro-Nielsen and Gramkov, 1996; Cachier et al., 2003; Modersitzki, 2004; Pennec et al., 1999) have worked towards providing a theoretical framework for the demons in order to understand and potentially modify the underlying assumptions.

The goal of this section is not to propose a novel non-rigid image registration algorithm but rather to build on previous work and improve the understanding of the demons algorithm. We first expand the alternate optimization framework of (Pennec et al., 1999) and show that the different variants of the algorithm can all be cast into a unique image registration framework. The symmetric forces variant was first proposed (but not analyzed) by Thirion as one possible expression of the demons forces. In (Rogelj and Kovačič, 2006; Wang et al., 2005) a similar *ad hoc* symmetrization of the forces proved to boost the results of the demons algorithm. One of the main results of our theoretical analysis is to show that the symmetric forces demons can be cast to the ESM framework of (Benhimane and Malis, 2004; Malis, 2004). We therefore have both empirical and theoretical evidences that this variant should be the most efficient one. Our second goal is to verify this evidence in a practical case study.

2.1. A Deeper Understanding of the Alternate Optimization of the Demons

Given a *fixed image* $F(.)$ and a *moving image* $M(.)$, non-parametric image registration is treated as an optimization problem that aims at finding the dis-

placement of each pixel so as to get a reasonable alignment of the images. The transformation $s(.)$, $p \mapsto s(p)$, models the spatial mapping of points from the fixed image space to the moving image space. In many cases, non-parametric spatial transformations will be described by a displacement field $\boldsymbol{s}$ which is simply added to an identity transformation to get the non-parametric transformation s:

$$s : p \mapsto p + \boldsymbol{s}(p)$$

The similarity criterion $\mathrm{Sim}(.,.)$ measures the resemblance of two images. In this paper we will only consider the mean squared error which forms the basis of intensity-based registration:

$$\begin{aligned} \mathrm{Sim}\left(F, M \circ s\right) &= \frac{1}{2} \left\| F - M \circ s \right\|^2 \\ &= \frac{1}{2 \left| \Omega_P \right|} \sum_{p \in \Omega_P} \left| F(p) - M(s(p)) \right|^2, \end{aligned} \tag{1}$$

where Ω_P is the region of overlap between F and $M \circ s$. Our approach could also be extended to other intensity relationships by borrowing ideas from (Avants et al., 2008; Cachier et al., 2003; Hermosillo et al., 2002).

A simple optimization of (1) over the space of non-parametric transformations leads to an ill-posed problem with unstable and non-smooth solutions. To avoid this and possibly add some *a priori* knowledge, a regularization term $\mathrm{Reg}(s)$ is often introduced to get the global energy

$$E(s) = \frac{1}{\sigma_i^2} \mathrm{Sim}\left(F, M \circ s\right) + \frac{1}{\sigma_T^2} \mathrm{Reg}\left(s\right),$$

where σ_i accounts for the noise on the image intensity, and σ_T controls the amount of regularization we need.

This energy indeed lead to a mathematically sound framework but the mixing of the similarity and the regularization terms leads in general to computationally intensive optimization steps. On the other hand the demons algorithm of (Thirion, 1998) provides a very efficient registration scheme. It has however often been considered as somewhat *ad hoc*.

In order to cast the demons algorithm into a minimization of a well posed criterion, it was proposed in (Cachier et al., 2003) to introduce a hidden variable in the registration process: correspondences. The idea is to consider the regularization criterion as a prior on the smoothness of the transformation s. Instead of requiring that point correspondences between image pixels, a non-parametric spatial transformation c, be exact realizations of the spatial transformation s, one allows some error at each image point. Considering a Gaussian noise on displacements, we end up with the global energy:

$$E(c, s) = \left\| \frac{1}{\sigma_i} (F - M \circ c) \right\|^2 + \frac{1}{\sigma_x^2} \mathrm{dist}\left(s, c\right)^2 + \frac{1}{\sigma_T^2} \mathrm{Reg}\left(s\right) \tag{2}$$

where σ_x accounts for a spatial uncertainty on the correspondences. We classically have $\mathrm{dist}\left(s, c\right) = \left\| c - s \right\|$ and $\mathrm{Reg}\left(s\right) = \left\| \nabla s \right\|^2$ but the regularization can,

for example, also be modified to address fluid-like constraints (Cachier et al., 2003).

The interest of this auxiliary variable is that an alternate optimization over c and s decouples the complex minimization into simple and very efficient steps. The first step solves for the correspondences by optimizing $\|\frac{1}{\sigma_i}(F - M \circ c)\|^2 + \frac{1}{\sigma_x^2} \text{dist}(s,c)^2$, with respect to c and with s being given, by making a step from $c = s$. The second step solves for the regularization by optimizing $\frac{1}{\sigma_x^2} \text{dist}(s,c)^2 + \frac{1}{\sigma_T^2} \text{Reg}(s)$, with respect to s and with c being given.

This second minimization has a closed-form solution using a single convolution when the regularization is quadratic and uniform. Given the harmonic regularization criteria $\|\nabla s\|^2$, for example, it can be shown that the optimal regularized deformation field is the convolution of the correspondence field by a Gaussian kernel. More elaborate regularization terms can lead to advanced vectorial filters (Cachier and Ayache, 2004). In this work, we focus on the first step of this alternate minimization and refer the reader to (Bro-Nielsen and Gramkov, 1996; Cachier et al., 2003; Modersitzki, 2004) for a detailed coverage of the regularization questions.

2.2. Compositive and Additive Demons

Let us consider the complete space of non-parametric spatial transformations. As with any type of spatial transformations, mapping a point through a first transformation and then through a second one is equivalent to mapping the point through the composition of the two spatial transformations. As such, the most natural and geometrically meaningful operation we can endow the space of non-parametric spatial transformations with, is the composition. To optimize for the correspondences, it is thus relevant to look for a small deformation that will be composed with the current estimate. Let this small deformation be described by a dense displacement field $\boldsymbol{u}$. In his seminal paper (Thirion, 1998), the author proposed to make, at each iteration what we call a compositive adjustment: $s \circ (\text{Id} + \boldsymbol{u})$. The compositive demons algorithm can then be described by the following iterations:

Algorithm 1 (Compositive Demons Iterations)
- Given the current transformation s, compute a correspondence update field $\boldsymbol{u}$ by minimizing $E_s^{\text{corr}}(\boldsymbol{u}) = \|F - M \circ s \circ (\text{Id} + \boldsymbol{u})\|^2 + \frac{\sigma_i^2}{\sigma_x^2} \|\boldsymbol{u}\|^2$ with respect to $\boldsymbol{u}$

- For a fluid-like regularization let $\boldsymbol{u} \leftarrow K_{\text{fluid}} \star \boldsymbol{u}$. The convolution kernel will typically be a Gaussian kernel.

- Let $c \leftarrow s \circ (\text{Id} + \boldsymbol{u})$

- For a diffusion-like regularization let $\boldsymbol{s} \leftarrow K_{\text{diff}} \star \boldsymbol{c}$ (else let $\boldsymbol{s} \leftarrow \boldsymbol{c}$). The convolution kernel will also typically be a Gaussian kernel. ∎

A different approach relies on the fact that the space of dense displacement fields forms a simple vector space with respect to the addition. Several teams

5

have thus been using an additive adjustment rule $s + u$ instead of the compositive update rule (Cachier et al., 2003; Ibáñez et al., 2005; Modersitzki, 2004; Pennec et al., 1999). This type of additive update is closer to the classical update rules used in Newton methods on vector spaces. It should however be noted that it disregards the fact that we work on spatial transformations. While it is natural to compose spatial transformations, their addition has no geometric meaning. From a more practical point of view though, it can be argued that the composition $s \circ (\text{Id} + u)$ requires to warp the dense displacement field s with u and to add the result with u. The addition rule is therefore less computationally expensive. We show however in our experiments that using this additive structure (which is not consistent with the spatial transformation space) leads to slower convergence with solutions that are further away from the gold standard available on controlled experiments. Let us however give an overview of how these additive adjustments have been used within the demons algorithm. From an initial non-parametric transformation s, the following iterations are performed until convergence:

Algorithm 2 (Additive Demons Iterations)

- Given the current transformation s, compute a correspondence update field u by minimizing $E_s^{\text{corr}}(u) = \|F - M \circ (s + u)\|^2 + \frac{\sigma_i^2}{\sigma_x^2} \|u\|^2$ with respect to u

- For a fluid-like regularization let $u \leftarrow K_{\text{fluid}} \star u$. The convolution kernel will typically be a Gaussian kernel.

- Let $c \leftarrow s + u$

- For a diffusion-like regularization let $s \leftarrow K_{\text{diff}} \star c$ (else let $s \leftarrow c$). The convolution kernel will also typically be a Gaussian kernel. ∎

In the current work, the demons iterations in Algorithm 1 and Algorithm 2 are simply run for a specified number of iterations. Future work should aim at evaluating more advanced stopping criteria.

2.3. Demons Forces

We see that, in Algorithm 1 and Algorithm 2, the minimization of $E_s^{\text{corr}}(u)$ is very close to a simple mean squared error image registration problem. The goal is to find an optimal update field u. Since we deal with a least-square problem, the methods we use in this paper rely on a linearization of the first inner term in $E_s^{\text{corr}}(u)$ and are based on Gauss-Newton-like approaches.

Let us consider the intensity difference at a given point, $\varphi_p(s) = F(p) - M \circ s(p)$. Let $\varphi_p^s(u) = F(p) - M \circ s \circ (\text{Id} + u)(p)$ in the compositive case and $\varphi_p^s(u) = F(p) - M \circ (s + u)(p)$ in the additive case. Let us assume that the following linearization is available:

$$\varphi_p^s(u) \approx \varphi_p^s(0) + J^p.u(p) = F(p) - M \circ s(p) + J^p.u(p). \tag{3}$$

Such a linearization can be used to rewrite the correspondence energy used in the demons algorithm:

$$E_s^{\text{corr}}(\boldsymbol{u}) \approx \frac{1}{2\,|\Omega_P|} \sum_{p \in \Omega_P} \left\| \left[\begin{smallmatrix} F(p)-M\circ s(p) \\ 0 \end{smallmatrix}\right] + \left[\begin{smallmatrix} J^p \\ \frac{\sigma_i(p)}{\sigma_x}I \end{smallmatrix}\right] . \boldsymbol{u}(p) \right\|^2,$$

where we recall that Ω_P is the overlap between F and $M \circ s$.

As opposed to global transformation cases (e.g. 2D rigid body transformations) we see that here, the approximations given for each pixel are independent from each other. This greatly simplifies the minimization of E_s^{corr} by splitting it into very simple systems for each pixel. We indeed only need to solve, at each pixel p, the following normal equations:

$$\left[\begin{smallmatrix} J^{pT} & \frac{\sigma_i(p)}{\sigma_x}I \end{smallmatrix}\right] . \left[\begin{smallmatrix} J^p \\ \frac{\sigma_i(p)}{\sigma_x}I \end{smallmatrix}\right] . \boldsymbol{u}(p) = - \left[\begin{smallmatrix} J^{pT} & \frac{\sigma_i(p)}{\sigma_x}I \end{smallmatrix}\right] . \left[\begin{smallmatrix} F(p)-M\circ s(p) \\ 0 \end{smallmatrix}\right]$$

which simplifies into

$$\left(J^{pT}.J^p + \frac{\sigma_i^2(p)}{\sigma_x^2}I \right).\boldsymbol{u}(p) = -(F(p) - M \circ s(p)).J^{pT}.$$

From the Sherman-Morrison formula, a.k.a. matrix inversion lemma, we finally get (Vercauteren, 2008):

$$\boldsymbol{u}(p) = - \frac{F(p) - M \circ s(p)}{\|J^p\|^2 + \frac{\sigma_i^2(p)}{\sigma_x^2}} J^{pT} \tag{4}$$

If we use the local estimation of the image noise $\sigma_i(p) = |F(p) - M \circ c(p)|$, we end up exactly with the expression of the demons forces proposed by Thirion. Note that this choice of image noise allows us to control the maximum step length by choosing σ_x (Cachier et al., 2003):

$$\|\boldsymbol{u}(p)\| \leq \frac{\sigma_x}{2}. \tag{5}$$

2.4. Linearization of the Intensity Difference

In the above explanation of the demons forces we have assumed that a linearization of φ_p^s was available. In (Vercauteren, 2008) we have demonstrated that the different demons forces could all be cast into a unified framework. We established a correspondence between different types of optimizers and different variants of the demons forces.

Obviously the linearization depends on the type of update rule, compositive or additive, that we use. The simplest linearization is given by a first-order Taylor expansion: $\varphi_p^s(\boldsymbol{u}) = \varphi_p^s(0) + J_s^{\varphi_p}.\boldsymbol{u} + O(\|\boldsymbol{u}\|^2)$, where $[J_s^{\varphi_p}]_i = \frac{\partial}{\partial u_i}\varphi_p^s(\boldsymbol{u})\big|_{\boldsymbol{u}=0}$.

We see that for the compositive update rule:

$$\frac{\partial \varphi_p^s(\boldsymbol{u})}{\partial \boldsymbol{u}(q)^T}\bigg|_{\boldsymbol{u}=0} = -\frac{\partial M \circ s\big((\mathrm{Id} + \boldsymbol{u})(p)\big)}{\partial \boldsymbol{u}(q)^T}\bigg|_{\boldsymbol{u}=0}$$
$$= -\frac{\partial M \circ s(p + \boldsymbol{u}(p))}{\partial \boldsymbol{u}(q)^T}\bigg|_{\boldsymbol{u}=0}$$
$$= -\delta_{p,q}\frac{\partial M \circ s(\rho)}{\partial \rho^T}\bigg|_{\rho=p}$$
$$= -\delta_{p,q}\nabla_p^T(M \circ s),$$

where $\delta_{p,q}$ is the Kronecker delta. This gives us a simple first order expansion of φ_p^s:

$$\varphi_p^s(\boldsymbol{u}) = \varphi_p^s(0) - \nabla_p^T(M \circ s).\boldsymbol{u}(p) + O(\|\boldsymbol{u}\|^2). \qquad (6)$$

By plugging $J^p = -\nabla_p(M \circ s)$ into (4), we thus get a Gauss-Newton step for the compositive update rule.

In (Vercauteren, 2008), we provided a full mathematical derivation of several other variants of the demons forces. In practice, changing the expression of J^p in the demons forces can be explained by different optimization strategies. To simplify the reading of this manuscript, we only provide the outcome of the derivations and refer the interested reader to (Vercauteren, 2008) for the technical part. One of the most interesting conclusions of these derivations was to show that the symmetric forces could be linked to the efficient second-order minimization (ESM) framework.

In (Benhimane and Malis, 2004; Malis, 2004), the authors have shown that, in the specific context of image registration, a Hessian-free second order approximation of the mean squared error similarity criterion could be obtained thanks to some nice properties of the image registration problem. More specifically, the ESM scheme uses the fact that when the images are aligned with the optimal spatial transformation, the fixed image and the warped image as well as their gradient should be very close to each other. The main idea is that we can use this information to improve the search direction of the Newton methods. Classical optimizers work by building a polynomial approximation of the cost function. The Newton-Raphson uses the value of the cost function, its first and second derivatives around 0 to build a second-order polynomial approximation which is typically unstable and expensive to compute. To alleviate the computational burden and improve numerical stability, the Gauss-Newton and Levenberg-Marquardt methods discard the second derivative and can thus only build a *first-order* polynomial approximation. In contrast, the ESM uses the value of the cost function, its first derivative around 0 as well as its first derivative around the optimum to build a numerically stable second-order polynomial without the need of a computationally expensive second derivative information

Extending the work of (Benhimane and Malis, 2004; Malis, 2004), we have shown in (Vercauteren, 2008) that by plugging $J^p = -\frac{1}{2}(\nabla_p^T(M \circ s) + \nabla_p^T F)$ into (4), we obtained an ESM-like step for the compositive update rule. This
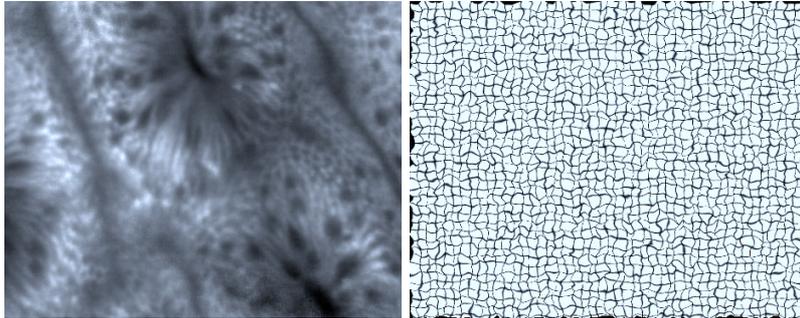
Figure 1: Original image (FCM) of a normal human colonic mucosa image (image courtesy of PD. Dr. A. Meining, Klinikum rechts der Isar, Munich) and one example random warp used in our controlled experimental setup. Size: 465×375 pixels. FOV: 229×184 µm.

implies that we get a second-order approximation in the demons forces instead of a first-order when using a Gauss-Newton step.

In Section 3.2, we will see that the demons algorithm can be extended to provide diffeomorphic spatial transformations. We will also see that the same optimization framework can be used to derive several variants of the forces for the diffeomorphic demons algorithm. A summary of the correspondence between demons forces, update rules (compositive, additive or diffeomorphic) and optimizers will be shown in Table 2.

2.5. Experiments: Practical Advantage of the Symmetric Forces

We have just highlighted the theoretical advantage of the symmetric forces variant of the demons algorithm in the compositive adjustment case. On the other hand, previous studies that have used (but not theoretically justified) a very similar symmetrization of the demons forces have reported the practical advantage of this variant (Rogelj and Kovačič, 2006; Wang et al., 2005) even with the additive adjustment rule. In this section, we provide some more evidences by comparing the different variants of the demons algorithm on both a synthetic and a realistic case study.

Choosing a right combination of parameters is always a difficult and application dependent problem. In the particular case of the demons algorithms, the most important parameter to be tuned is the sigma, σ_{diff}, of the Gaussian *diffusion-like* regularization. Standard values for the maximum step length (betwen 0.25 and 2 pixels) and the Gaussian *diffusion-like* regularization (between no *diffusion-like* regularization at all and $\sigma_{\text{diff}} = 1$) work well in most of the cases. In this work we follow the example of (Thirion, 1998; Wang et al., 2005) and use $\sigma_{\text{diff}} = 1$. A more elaborate way of tuning this parameter would be to evaluate the registration result for a variety of values and choose the best trade-off. Such an approach is for example clearly presented in (Yeo et al., 2008a). The same set of parameters is used for all the experiments: a maximum step
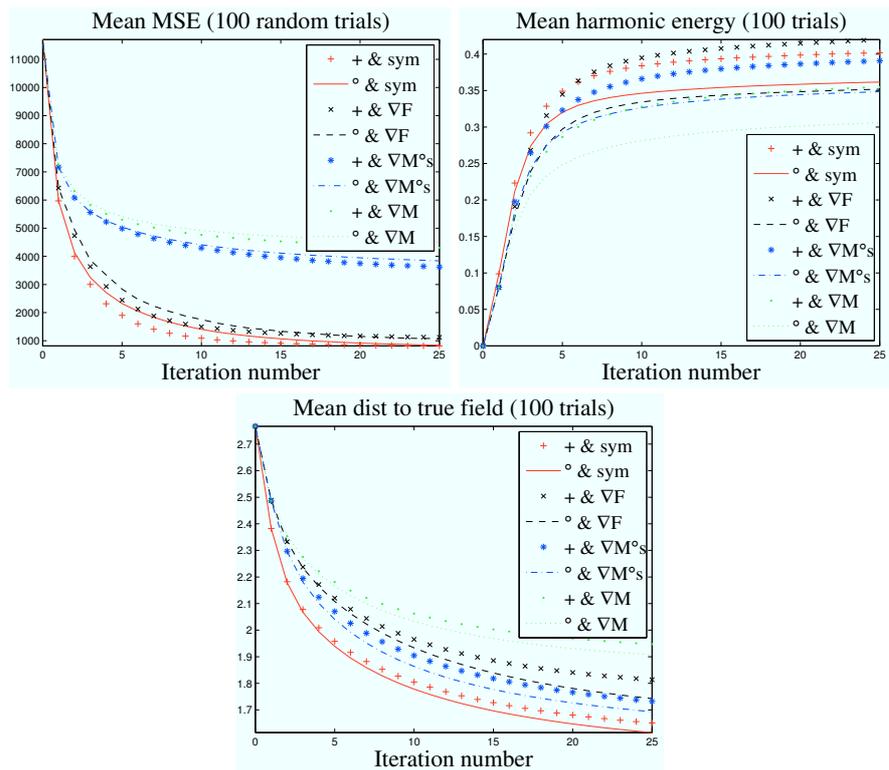
9

Figure 2: Registration on 100 random experiments such as the one in Fig. 1. Note the faster convergence of the symmetric forces demons in terms of images intensities agreement (MSE), smoothness of the non-rigid spatial transformation (harmonic energy) and more importantly in terms of distance to the actual spatial transformation (measured in pixels).

Table 1: Comparison (Dice similarity coefficient * 100) of the discrete segmentations obtained from the registration with the additive demons (most common update type) of the synthetic T1-weighted MR images shown in Fig. 7. For each tissue class, the best segmentation is obtained with the symmetric forces variant.

	CSF	GM	WM	Fat	Muscle	Skin	Skull	Vessels	Fat2	Dura	Marrow
No reg.	41.73	63.06	61.51	19.30	20.14	66.65	42.75	14.26	6.06	14.74	28.19
Add Fix.	63.41	78.99	79.23	47.74	36.40	78.57	64.91	27.21	14.75	23.13	45.05
Add Sym.	69.75	83.78	84.58	52.81	41.41	82.94	71.28	35.21	17.25	29.79	51.85

length of 2 pixels, a Gaussian *fluid-like* regularization with $\sigma_{\text{fluid}} = 1$ and a Gaussian *diffusion-like* regularization with $\sigma_{\text{diff}} = 1$. The emphasis is on the comparison of the various schemes and not on the final performance. Therefore, no multi-resolution scheme was used.

The first experiments provide a completely controlled setup. We use a fibered confocal microscopy image as our original image. For each random experiment, we generate a smooth displacement field with a Markov random field (MRF) sampler and warp the original image. We add some random noise to the intensities of both the original and the warped image. We then run the different demons algorithms starting with an identity spatial transformation. Two conclusions can be drawn from Fig. 2. First, we see that, for a given type of demons forces, the compositive and additive demons converge almost at the same pace in terms of MSE. However, if we look at the harmonic energy and at the distance to the actual field, we see that, for all demons forces, the compositive demons converges faster than its additive counterpart. The second interesting fact we see here, for both the compositive and the additive demons, is the fact that the symmetric forces variant converges faster in terms of visual appearance, i.e. MSE. It also behaves well in terms of smoothness of the displacement field. Finally, it is the fastest to converge in terms of distance to the actual field.

Our second setup is a more realistic case study were a gold standard is still available. We use synthetic T1 MR images from two different anatomies available from BrainWeb (Aubert-Broche et al., 2006). These datasets are distributed along with a segmentation of eleven different tissue classes. The goal of this experiment is to show that on this more realistic case we can also draw the conclusion that symmetric forces are an advantageous option in practice. To ease the reading of the results, we simply compare the most widely used version of the demons, additive update rule and fixed image force, with its symmetric forces counterpart. We see on Fig. 3, that, on this dataset also, the symmetric forces variant converges faster in terms of MSE on the images intensities and smoothness of the displacement field. In Table 1, we compare the agreement between the segmentation of the fixed image and the segmentation of the moving image warped by the non-parametric spatial transformation found by the registration of the T1-weighted MR images. Segmentation agreement is commonly measured with the Dice similarity coefficient (DSC). Given two segmentations A and B, the DSC is calculated as $\frac{2|A \cap B|}{|A|+|B|}$. It can be viewed as a measure of
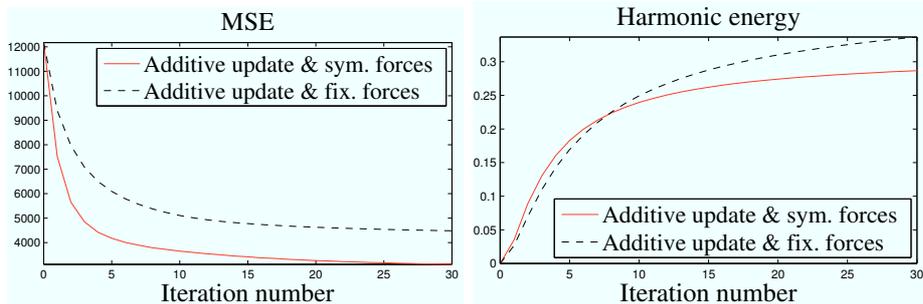
Figure 3: Comparison of the fixed image forces demons algorithm with the symmetric forces demons algorithm on the BrainWeb images (shown in Fig. 7). To ease the reading, only the most widely used update type, i.e. the additive one, is shown. Note the faster convergence of the symmetric forces demons in terms of images intensities agreement (MSE) and smoothness of the non-rigid transformation (harmonic energy).

overlap between the two segmentations. If A and B completely overlap, the DSC is equal to one. If there is no overlap, then the DSC is equal to zero. In Table 1, we show the Dice similarity coefficient of the eleven tissue classes before any registration and after registration with the two variants of the demons we consider. We see that, for each tissue class, the best segmentation is obtained with the symmetric forces variant.

3. Introducing Diffeomorphisms into the Demons

One of the main limitations of both the additive and compositive demons algorithm is that it does not ensure the invertibility of the output transformations contrarily to diffeomorphic image registration algorithms. It may not be a requirement to get such diffeomorphic transformations. However, this framework may be relevant and powerful for many image registration problems. It indeed preserves the topology of the objects in the image and prevents from introducing folding which is often physically impossible. Diffeomorphisms are also considered to be a good working framework when no additional information about the spatial transformation is available. With the development of computational anatomy and in the absence of a justified physical model of inter-subject variability, statistics on diffeomorphisms also become an important topic (Arsigny et al., 2006; Lepore et al., 2008; Vaillant et al., 2004; Xue et al., 2006). Diffeomorphic registration algorithms are at the core of this research field since they often provide the *input data*. Invertible spatial transformations are also required for the registration of diffusion tensor images (DTI) as shown in (Yeo et al., 2008c).

Diffeomorphic image registration usually relies on the computationally heavy solution of some partial differential equations (Avants et al., 2008; Beg et al.,

12

2005; Christensen et al., 1996; Joshi and Miller, 2000; Leow et al., 2007; Marsland and Twining, 2004; Miller et al., 1998; Trouvé, 1998; Vaillant et al., 2004) or uses very small optimization steps such as in the approach of (Chefd'hotel et al., 2002). In (Rueckert et al., 2006), the authors proposed a parametric approach by composing a set of constrained B-spline transformations. Since the composition and inversion of B-spline transformations cannot be expressed on a B-spline basis, the advantage of using a parametric approach is not clear in this case. A complex optimization scheme on constrained B-splines has also been proposed in (Noblet et al., 2005) for this problem.

In this section, we propose a non-parametric diffeomorphic image registration algorithm based on the demons algorithm. We have just shown that the original demons algorithm could be seen as an optimization procedure on the entire space of displacement fields. The main idea of our algorithm is to adapt this optimization procedure to a space of diffeomorphic transformations. We show that a Lie group structure on diffeomorphic transformations that has recently been proposed in (Arsigny et al., 2006) can be used in combination with optimization tools on Lie groups to derive our diffeomorphic image registration algorithm.

3.1. A Lie Group Structure on Diffeomorphisms

Like most spatial transformation spaces, diffeomorphisms do not form a vector space with respect to classical operations, e.g. adding two diffeomorphisms may lead to non-invertibility. However, diffeomorphisms can be smoothly composed and inverted. It is therefore possible to provide a Lie group structure on the space of diffeomorphisms[1]. The most straightforward way to adapt the demons algorithm to make it diffeomorphic, is to optimize the global energy (2) over a space of diffeomorphisms instead of the complete space of non-parametric spatial transformations. We thus need to perform an optimization procedure on a Lie group such as in (Benhimane and Malis, 2004; Mahony and Manton, 2002).

Newton methods for Lie groups: Optimization problems on Lie groups can often be related to constrained optimization by embedding the Lie group in an Euclidean space. The classical way of dealing with the geometric structure of the Lie group is to use Lagrange multipliers or when the constraints are simple to have an *ad hoc* procedure to preserve the constraints (e.g. renormalize a quaternion to have a unit quaternion). In this work we use an alternative strategy known as geometric optimization which uses local canonical coordinates (Mahony and Manton, 2002). This strategy intrinsically takes care of the geometric structure of the group and allows the use of unconstrained optimization routines.

Let us first recall that a Lie group $\mathcal{G}$ is a smooth manifold together with a smooth composition map usually denoted as multiplication ($x \mapsto s \circ x$ for x and

---

[1]In theory, we should clearly state that we do not have an actual Lie group but only a pseudo-group. These technical details should be addressed by future work.
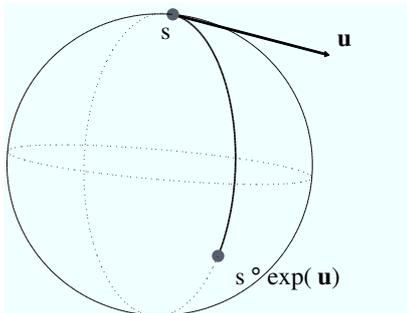
Figure 4: Schematic view of the intrinsic update step used in Newton methods for Lie groups. An unconstrained update $\boldsymbol{u}$ is computed on the Lie Algebra and is projected back onto the Lie group through the exponential map.

s in $\mathcal{G}$), and a smooth inverse map ($x \mapsto x^{-1}$ for x in $\mathcal{G}$), that satisfy the group axioms: closure, associativity, existence of a neutral element (denoted hereafter as Id) and existence of an inverse. We refer the reader to the standard textbooks for a detailed treatment of Lie groups, see e.g. (Helgason, 2001). To any Lie group can be associated a Lie algebra $\mathfrak{g}$, whose underlying vector space is the tangent space of $\mathcal{G}$ at the neutral element Id. This Lie algebra captures the local structure of $\mathcal{G}$. The Lie group and the Lie algebra are related through the group exponential which is a smooth mapping from a neighborhood of 0 in $\mathfrak{g}$ to a neighborhood of Id in $\mathcal{G}$[2]. Canonical coordinates provide local coordinate charts so that for any $x \in \mathcal{G}$ in some neighborhood of s, there exists a vector $\boldsymbol{u}$ such that $x = s \circ \exp(\boldsymbol{u})$. These coordinates can be used to get the Taylor expansion of a smooth function φ on the Lie group $\mathcal{G}$: $\varphi\left(s \circ \exp(\boldsymbol{u})\right) \approx \varphi(s) + J_s^\varphi . \boldsymbol{u}$ where $[J_s^\varphi]_i = \frac{\partial}{\partial u_i} \varphi(s \circ \exp(\boldsymbol{u}))\big|_{\boldsymbol{u}=0}$

In (Vercauteren et al., 2007a), we presented a number of Newton methods for optimization problems on Lie groups. The main idea of these methods is to find, from the current transformation s, an update step $\boldsymbol{u}$ on the Lie algebra based on an expansion of the cost function and to use an intrinsic update rule on the Lie group through the exponential map (cf. Fig. 4):

$$s \leftarrow s \circ \exp(\boldsymbol{u}).$$

Newton methods for diffeomorphic registration: Newton methods for Lie groups seem in theory to be a good candidate for diffeomorphic image registration. In practice however, they can only be used if a fast and tractable numerical

---

[2]In the finite dimensional case, the exponential mapping is one-to-one and invertible. We should however mention that, on the entire infinite dimensional space of diffeomorphisms, this is not the case anymore, even around zero. This may be problematic for some algorithms, for example if the logarithm is needed. In our case, we do not rely on such a mathematical property since we only require the exponential to be well defined.

scheme for the computation of the exponential is available. We would indeed have to use it at each iteration. Such an efficient scheme clearly relies on a good parameterization of the Lie group and the Lie algebra.

In the context of image registration, it has been proposed, in (Miller et al., 1998) for landmarks and (Beg et al., 2005) for images, to parameterize the space of diffeomorphic transformations using *time-varying* speed vector fields. This has the advantage of fully using the group structure. However, the computation of a displacement field requires the numerical integration of a time-varying ODE (Trouvé, 1998). In (Arsigny et al., 2006) the authors proposed a practical approximation of such a Lie group structure on diffeomorphisms by using *stationary* speed vector fields only. This has the significant advantage of yielding very fast computations of exponentials. It becomes indeed possible to use the scaling and squaring method and compute the exponential with just a few compositions of spatial transformations (Arsigny et al., 2006).

By generalizing to vector fields the equivalence that exists in the finite dimensional case between one-parameter subgroups and the exponential map, the exponential $\exp(\boldsymbol{u})$ of a smooth vector field $\boldsymbol{u}$ is defined in (Arsigny et al., 2006) as the flow at time one of the stationary ODE,

$$\frac{\partial p(t)}{\partial t} = \boldsymbol{u}(p(t)).$$

From the properties of one-parameters subgroups $(t \mapsto \exp(t\boldsymbol{u}))$, we know that for any scalars t and t' we have $\exp\big((t + t')\boldsymbol{u}\big) = \exp(t\boldsymbol{u}) \circ \exp(t'\boldsymbol{u})$. We thus see that, for any integer K,

$$\exp(\boldsymbol{u}) = \exp(K^{-1}\boldsymbol{u})^K,$$

where the power operation relates to the composition of spatial transformations. Using this property in a *recursive* manner, this yields the following efficient algorithm for the computation of vector fields exponentials:

Algorithm 3 (Fast Vector Field Exponentials)

- Choose N such that $2^{-N}\boldsymbol{u}$ is close enough to 0, e.g. $\max_p \big\|2^{-N}\boldsymbol{u}(p)\big\| \leq 0.5$

- Perform an explicit first order integration: $\boldsymbol{v}(p) \leftarrow 2^{-N}\boldsymbol{u}(p)$ for all pixels. As previously we use $v = \mathrm{Id} + \boldsymbol{v}$

- Do N (not 2^N!) *recursive* squarings of v: $v \leftarrow v \circ v$ ∎

It should be noted that since we deal with an infinite dimensional group, this framework poses some theoretical problems that are yet to be completely solved. For example, we do not strictly deal with a Lie group but only with a pseudo-group. We do however have sufficient evidence of its well-foundedness and its effectiveness to consider this framework as one of the most efficient ways of dealing with diffeomorphisms.

3.2. Diffeomorphic Demons Algorithm

Let us now derive our non-parametric diffeomorphic image registration algorithm. As in Section 2.1, we focus on the first step of the minimization rather

than on the regularization step. Let us assume that we have a first order expansion of the intensity difference of the form:

$$F(p) - M \circ s \circ \exp(\boldsymbol{u})(p) \approx F(p) - M \circ s(p) + J^p.\boldsymbol{u}(p).$$

To get a computationally tractable expression of the correspondence energy and optimize (2), we choose the following distance between two diffeomorphisms: $\mathrm{dist}\,(s, c) = \left\| \mathrm{Id} - s^{-1} \circ c \right\|$. We then get

$$\mathrm{dist}\,(s, s \circ \exp(\boldsymbol{u})) = \left\| \mathrm{Id} - \exp(\boldsymbol{u}) \right\| \approx \left\| \boldsymbol{u} \right\|.$$

By using these two expansions, we see that we get the following expression for the approximation of the correspondence energy:

$$E_s^{\mathrm{corr}}(\boldsymbol{u}) = \left\| F - M \circ s \circ \exp(\boldsymbol{u}) \right\|^2 + \frac{\sigma_i^2}{\sigma_x^2} \mathrm{dist}\,(s, s \circ \exp(\boldsymbol{u}))^2$$

$$\approx \frac{1}{2\,|\Omega_P|} \sum_{p \in \Omega_P} \left\| \left[\begin{smallmatrix} F(p) - M \circ s(p) \\ 0 \end{smallmatrix} \right] + \left[\begin{smallmatrix} J^p \\ \frac{\sigma_i(p)}{\sigma_x} I \end{smallmatrix} \right].\boldsymbol{u}(p) \right\|^2.$$

We thus have the same expression (4) as in the classical demons for the demons forces. The difference is in how we compute the Jacobian J^p and how we consider $\boldsymbol{u}$. In the classical demons $\boldsymbol{u}$ is a dense displacement field whereas in the diffeomorphic demons, $\boldsymbol{u}$ is considered as a speed vector field.

From these derivations, we thus obtain our non-parametric diffeomorphic image registration algorithm:

Algorithm 4 (Diffeomorphic Demons Iterations)

- Given the current transformation s, compute a correspondence update field $\boldsymbol{u}$ by minimizing $E_s^{\mathrm{corr}}(\boldsymbol{u})$ with respect to $\boldsymbol{u}$

- For a fluid-like regularization let $\boldsymbol{u} \leftarrow K_{\mathrm{fluid}} \star \boldsymbol{u}$. The convolution kernel will typically be a Gaussian kernel.

- Let $c \leftarrow s \circ \exp(\boldsymbol{u})$

- For a diffusion-like regularization let $\boldsymbol{s} \leftarrow K_{\mathrm{diff}} \star \boldsymbol{c}$ (else let $\boldsymbol{s} \leftarrow \boldsymbol{c}$). The convolution kernel will also typically be a Gaussian kernel.　■

As in Section 2.4, we now need a linearization of the intensity difference to put in the expression of the demons forces (4). In (Vercauteren, 2008) we have shown that the unified framework we used to explain the different demons forces in the compositive and additive case, could also be extended to relate different variants of the demons forces used in the diffeomorphic demons with different types of optimizers. The complete set of variants of the demons is summarized in Table 2. We also showed in (Vercauteren, 2008) that the symmetric forces can be explained with the ESM framework.

Even though from a practical point of view these forces are rather obvious, their mathematical derivation is quite technically involved. To ease the reading of this manuscript, we omit it from the current discussion and refer the interested reader to (Vercauteren, 2008).

Table 2: Demons variants according to the different adjustment rules and the demons forces. In this table, *Thirion* refers to the variants proposed in (Thirion, 1998), *ITK* refers to the implementation of the Insight Toolkit (Ibáñez et al., 2005), *Diffeo. Dem.* refers to our proposed diffeomorphic demons (Vercauteren et al., 2007c) and *Pennec 99* refers to the variant proposed in (Pennec et al., 1999). The schemes that can readily be explained by a Gauss-Newton scheme are appended with *(GN)* while those that fit in the ESM framework are denoted by *(ESM)*. The variants that are implemented in ITK can only be seen as approximate gradient schemes as they use an additive adjustment rule with forces derived for the compositive demons.

Adjustment / Used J^p	Compositive $c \leftarrow s \circ (\mathrm{Id} + \boldsymbol{u})$	Additive $\boldsymbol{c} \leftarrow \boldsymbol{s} + \boldsymbol{u}$	Diffeomorphic $c \leftarrow s \circ \exp(\boldsymbol{u})$
$-\nabla_p^T (M \circ s)$	Thirion Mov. (GN)		Diffeo. Mov. (GN)
$-\nabla_p^T F$	Thirion Fix.	ITK Fix.	Diffeo. Fix.
$-\left(\nabla_p^T F + \nabla_p^T (M \circ s)\right)/2$	Thirion Sym. (ESM-like)	ITK Sym.	Diffeo. Sym. (ESM)
$-\nabla_{s(p)}^T M$		Pennec 99 (GN)	

4. Experiments: Diffeomorphic Registration Can Be Fast

As a first test to evaluate the usefulness of the diffeomorphic demons with respect to the additive demons, we have used the classical "Circle to C" registration problem. We used a multi-resolution strategy and used the same set of parameters for both the additive and diffeomorphic demons: Thirion's rule ($J^p = \nabla_p^T F$) with a maximum step length of 2 pixels was used in the demons force (4), a Gaussian *fluid-like* regularization with $\sigma_{\text{fluid}} = 1$ and a Gaussian *diffusion-like* regularization with $\sigma_{\text{diff}} = 1$. As is clearly shown in Fig. 5, the additive demons gets trapped and is not able to correctly register the images. Moreover, we can see that the final spatial transformation is spurious and shows foldings. Contrastingly, on this classical "Circle to C" example, the diffeomorphic demons algorithm is able to register the images and does it with a smooth invertible transformation.

To further measure the performance of the diffeomorphic demons algorithm with respect to the additive and compositive demons algorithm, two sets of results are presented. We used the same set of parameters for all the experiments: Thirion's rule with a maximum step length of 2 pixels, $\sigma_{\text{fluid}} = 1$ and $\sigma_{\text{diff}} = 1$. Since the emphasis is on the comparison of the various schemes and not on the final performance, no multi-resolution scheme was used.

The first experiments are very similar to the ones presented in Section 2.5 and Fig. 1. In this section, we use this setup to compare the different variants of the adjustment rule instead of the different demons forces. These experiments provide a completely controlled setup. We use a fibered confocal microscopy image as our original image. For each experiment, we generate a random diffeomorphic spatial transformation (by passing a Markov random field through the exponential) and warp the original image. We add some noise both to the original and the warped image. We then run the registration algorithms starting with an identity spatial transformation. We see on Fig. 6 that, in terms of MSE, the performance of the additive demons, the compositive demons and
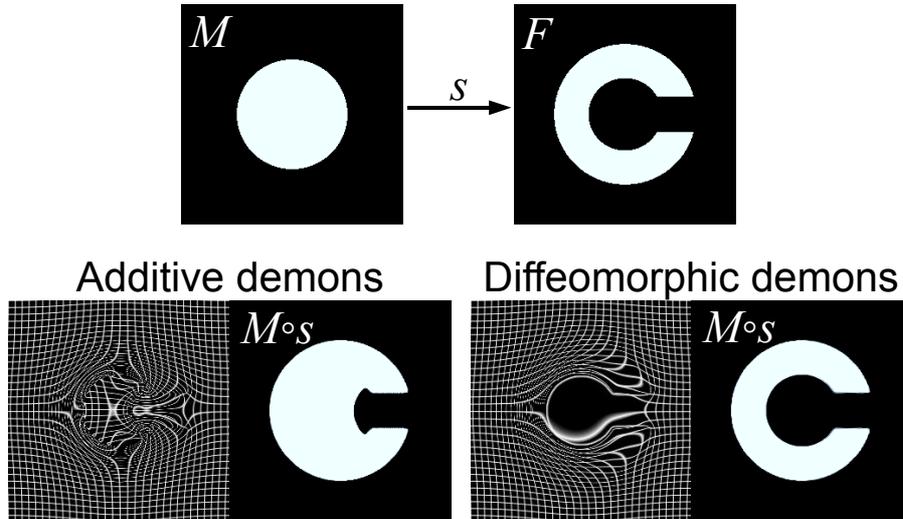
Figure 5: Classical Circle to C registration example. With the same set of parameters the additive demons fails to converge and shows foldings in the registration results whereas the diffeomorphic demons converges with a smooth invertible transformation.

the diffeomorphic demons algorithm are similar. However, the distance to the true field, the harmonic energy and the minimum and maximum values of the determinant of the Jacobian of the transformation show that both the compositive demons and the diffeomorphic demons clearly outperform the additive scheme by providing much smoother spatial transformations. We also see that our diffeomorphic algorithm provides better results in terms of distance to the true Jacobian of the transformation. Note that this is accomplished with a reasonable 50% increase of computation time per iteration with respect to the computationally efficient additive demons algorithm. We could also have presented for all the adjustment rules, a comparison of the different demons forces. We chose however not to show these results as the conclusions are very similar to the ones found in Section 2.5. Namely, the symmetric forces variant outperforms the other demons forces.

Our second setup is the same as the second setup of Section 2.5. It is a more realistic case study were a gold standard is still available. We use synthetic T1 MR images from two different anatomies available from BrainWeb (Aubert-Broche et al., 2006). These datasets are distributed along with a segmentation of eleven different tissue classes. The goal of this experiment is to show that, on this more realistic case, we can also draw the conclusion that diffeomorphic demons are an advantageous option in practice. To ease the reading of the results, we simply compare the most widely used version of the demons, additive update rule and fixed image force with its diffeomorphic counterpart.
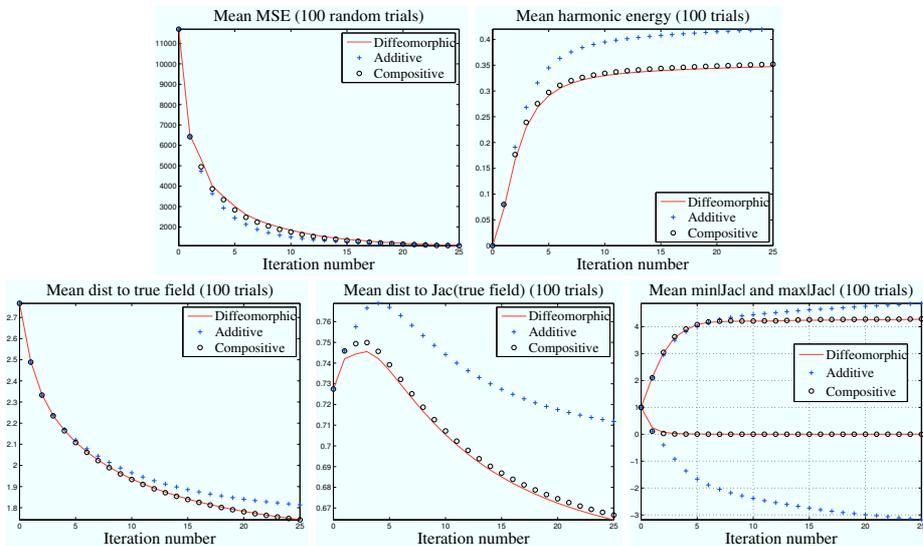
Figure 6: Registration on 100 random experiments such as the one presented in Fig. 1. We compare the three types of update rule with Thirion's demons forces based on the gradient of the fixed image. Note that for similar performance in terms of MSE and distance to the true field (measured in pixels), the compositive and diffeomorphic demons algorithm provides much smoother results than the additive demons algorithm. We also see that we provide diffeomorphic transformations whereas $\min(|\mathrm{Jac}(s)|)$ goes way below zero with the additive demons algorithm. Furthermore, in terms of distance to the true Jacobian of the transformation, the diffeomorphic demons provide a large gain with respect to the additive demons and a small gain with respect to the compositive demons.

We see on Fig. 8 and Table 3 that, on this dataset also, the additive demons algorithm and our algorithm provide very similar results in terms of visual appearance, MSE and segmentation accuracy. However, we see that our algorithm does it with much better spatial transformations. We indeed get smoother deformations that are diffeomorphic.

Thanks to our open-source implementation of the diffeomorphic demons algorithm (Vercauteren et al., 2007b), our registration scheme has also been tested by an independent group, cf. (Urschler et al., 2007). They proposed a non-rigid registration evaluation framework and benchmarked six different open-source non-rigid registration algorithms. According to their findings, "the overall best results on the evaluation experiments are given by the diffeomorphic demons algorithm." Given the small number of tested algorithms and the current evaluation framework used in (Urschler et al., 2007), these results do not, by any means, provide a definite ranking. They can however be considered as a further evidence to support the value of our approach. A more advanced tool for benchmarking non-rigid registration algorithms might also be available in the future (Christensen et al., 2006).
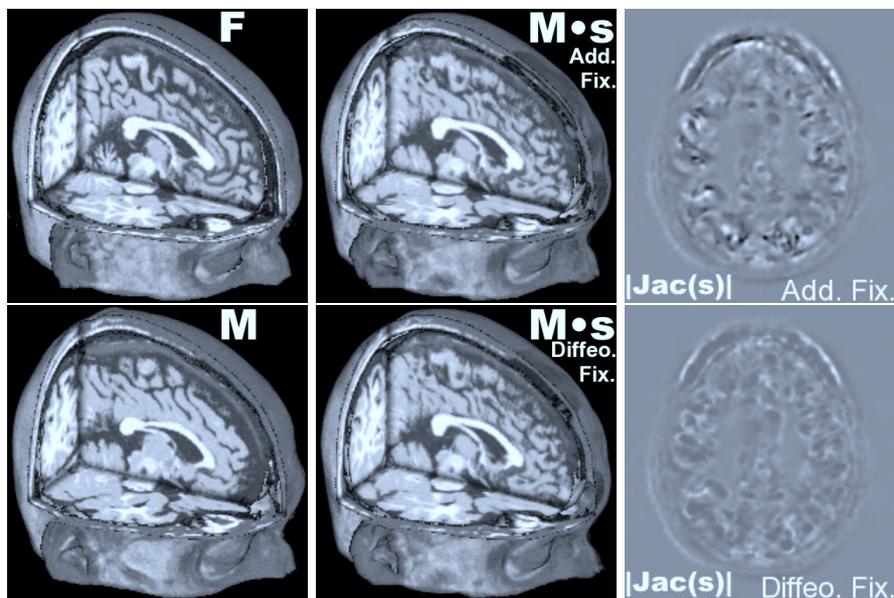
Figure 7: Registration of two synthetic T1 MR images of distinct anatomies. The focus is on the evaluation of our proposed update rule (diffeomorphic) with the most popular one (additive). For visually similar results, our algorithm provides smoother diffeomorphic transformations.

Table 3: Comparison (Dice similarity coefficient * 100) of the discrete segmentations obtained from the registration of the synthetic T1-weighted MR images in Fig. 7.

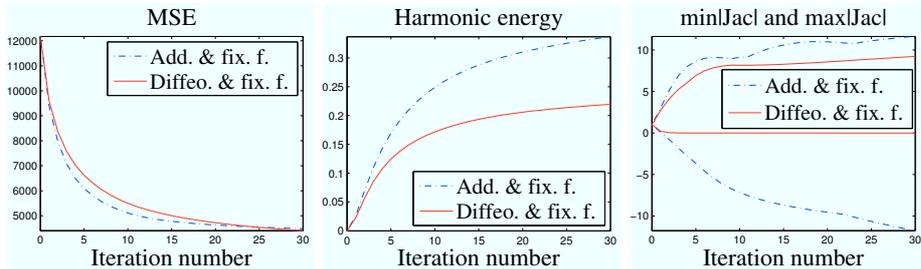|	CSF	GM	WM	Fat	Muscle	Skin	Skull	Vessels	Fat2	Dura	Marrow
No reg.	41.73	63.06	61.51	19.30	20.14	66.65	42.75	14.26	6.06	14.74	28.19
Add Fix	63.41	78.99	79.23	47.74	36.40	78.57	64.91	27.21	14.75	23.13	45.05
Diffeo Fix	64.37	78.94	78.43	47.22	36.11	79.39	65.02	27.25	14.70	24.56	43.92

Figure 8: Comparison of the additive demons algorithm with the diffeomorphic demons algorithm on the BrainWeb images shown in Fig. 7. To ease the reading, only the most widely used demons force (fixed image force) is used and the compositive update rule is omitted. For similar performance in terms of MSE, our algorithm provides much smoother transformations than the additive demons algorithm. Most importantly we see that we provide diffeomorphic transformations whereas $\min(|\mathrm{Jac}(s)|)$ goes way below zero with the additive demons.

Even though no multi-resolution scheme was used in the quantitative experiments we showed in this work, real-life studies usually benefit from such a strategy as it increases robustness and lowers the computational cost. With the setup given hereafter, the registration of two $256 \times 256 \times 181$ BrainWeb images runs in 2 minutes and 30 seconds on a 2×2.8 GHz quad-core Intel Xeon Apple Mac pro computer: 3 multi-resolution levels with 20 diffeomorphic demons iterations at the coarsest level, and 10 diffeomorphic demons iterations for the two others; symmetric forces with a maximum step length of 2 pixels; $\sigma_{\mathrm{fluid}} = 1$ and $\sigma_{\mathrm{diff}} = 1$. We believe that such computation times makes our algorithm well-fit for real-life problems.

5. Discussion

In this section, we discuss some of the theoretical properties of our diffeomorphic demons. Recent works proposed diffeomorphic image registration algorithms that are close to our proposition in the sense that they use the exponential of a stationary vector field to get a diffeomorphic spatial transformation (Ashburner, 2007; Hernandez et al., 2007). To provide a more objective view, the present discussion is organized around how the work in (Ashburner, 2007; Hernandez et al., 2007) and our diffeomorphic demons address some of the theoretical issues differently. A complementary discussion can also be found in (Hernandez et al., 2008).

Discretization Questions: We have provided a principled approach to the problem of non-parametric diffeomorphic registration. However our discussion is based on spatial transformations and vector fields defined on a continuous domain. In practice of course, we need to discretize the problem. In this work, this is simply done by using the sampling grid of the images to represent the spatial transformations and the vector fields. Trilinear interpolation is then commonly

used to interpolate between the sampling point. It should however be noted that this discretization is, in theory, not necessarily consistent with the continuous diffeomorphic framework. It might even be argued that diffeomorphisms are not well-posed in the context of discrete grid (Bazin et al., 2007). A somewhat more rigorous approach might be to integrate the interpolation, within the modeling of the problem such as in (Ashburner, 2007) and some finite element based methods. However, since our approach is based on the composition of transformations and that the composition of finite element based transformations cannot in general be represented by a similar finite element based transformations it is not clear what the advantage would be. Furthermore, this would induce a larger computational cost.

In our results, we have shown that, in general, the diffeomorphic demons provided spatial transformations whose Jacobians remain positive. It might however not be always the case. The way we compute the Jacobian is by finite difference on the sampling grid. This computational framework is however, in theory, not necessarily consistent with the continuous diffeomorphic setup. It might indeed well be the case that a true continuous diffeomorphism has negative Jacobians when evaluated by finite difference on a sampling grid. This would simply mean that, in some sense, our sampling grid does not meet the Nyquist criterion for the spatial transformation we consider. In (Ashburner, 2007), the author proposed to use a computation of the Jacobian of the transformation that is more consistent with the composition of spatial transformations. If we know the Jacobian of two transformations, the analytical Jacobian of the composition at a given point can be computed through matrix multiplication. This has the advantage of better estimating the true Jacobian. However, since we still represent the composed transformation on the sampling grid, it can in this case be argued that the Jacobian is not consistent with the representation of the transformation. This would also result in a more computationally intensive scheme as we would need to compute and store the Jacobian at each iteration. In our setting, we also use a smoothing of the transformation. Because of this smoothing, it is in the end difficult to get an analytical form of the Jacobian of the spatial transformation.

Since statistical analysis and its biological interpretation often relies on the assumption that the Jacobian determinant is positive, potential negative values might appear like a true possible problem. A real solution to this issue would rely on a more powerfull representation of spatial transformations that would be consistent with the composition and the computation of the Jacobian. In practice though, negative Jacobians appear very seldom and if they do, they are only so slightly negative that dismissing them should not lead to any problems for biological interpretation.

Compositive Demons as an Approximate Diffeomorphic Demons: In the results we showed in Section 4, we saw that the compositive demons algorithm provides results that are very close to the results of the diffeomorphic demons, with a slight advantage to the diffeomorphic demons.

This conclusion should in fact not be so surprising. A different way of looking at the compositive demons is to realize that the compositive and diffeomorphic

adjustment rules are equivalent up to the first-order:

$$s \circ \exp(\boldsymbol{u}) = s \circ (\mathrm{Id} + \boldsymbol{u}) + O(\|\boldsymbol{u}\|^2).$$

And since we only use small adjustments, the first-order approximation can provide good results. On a similar point, it is shown in (Arsigny et al., 2006; Ashburner, 2007) that, for a given speed vector field, there is an optimal number of steps to compute the exponential. If too many steps are used, we introduce too much numerical errors. Because, in our case, we only use small speed vector fields, it might be better to use very small number of scaling and squaring steps. As Id $+\boldsymbol{u}$ can also be seen as an exponential computed with zero scaling and squaring steps, it might explain that the compositive demons performs well.

In contrast, if we try to get such a retrospective view of the additive demons, we see that a crude approximation is necessary. We indeed have at the first order,

$$s \circ \exp(\boldsymbol{u})(p) \approx s(p) + \mathrm{Jac}(s)(p).\boldsymbol{u}(p),$$

which basically implies that the additive demons can be seen as an approximation of the diffeomorphic demons if we disregard the influence of the Jacobian of the current transformation. This Jacobian can actually be quite far from the identity.

All in all, we do believe that even if the results of the diffeomorphic demons are only slightly better when compared to the compositive demons, the algorithm we developed leads to a mathematically sound framework for dealing with diffeomorphisms. Future work should aim at finding how to automatically compute the number of steps in the exponential. It might well be the case that, since we only use small velocity fields, most of the time the optimum number should be zero.

Different Ways of Using the Exponential: We have seen that the parameterization of diffeomorphic transformations through a stationary speed vector field presented in (Arsigny et al., 2006), provides a very efficient framework for dealing with diffeomorphisms. It is thus not surprising to see that several groups have started using it. In the diffeomorphic demons, we chose to use this parameterization to encode an adjustment that needs to be made to the current transformation. In (Ashburner, 2007; Hernandez et al., 2007), the authors propose to use a complete transformation parameterized by a speed vector field.

The approaches of (Ashburner, 2007) and (Hernandez et al., 2007) have several advantages as well as several drawbacks. First of all, it is argued in (Arsigny et al., 2006) that the exponential is a smooth one-to-one mapping between an open neighborhood of the zero speed vector field and an open neighborhood of the identity transformation. It is still not clear which elements of the group of diffeomorphic transformations can be reached with the exponential, i.e., what is the size of the target open neighborhood. In our approach, as we iteratively compose the current transformation with a small exponential update, we don't suffer from this problem. In the finite-dimensional case, for example, even if all the elements of the Lie group cannot be directly reached

through the exponential, they can still all be reached through a composition of two exponentials (Wüstner, 2003). When using a single parameterization in the Lie algebra, it might therefore be useful to investigate the image of the exponential map.

In the diffeomorphic demons, we chose to use a regularization through a Gaussian smoothing on the spatial transformation. This may, in theory, not always be consistent with the diffeomorphic framework. It does however provide a very efficient regularization in terms of computation time. It is also familiar to many people in the field of biomedical image analysis. The most important benefit we get from it is, however, that it allows us to decouple the minimization problem into an alternate scheme composed of very easy steps. In (Ashburner, 2007), the author chose to introduce a regularization term on the speed vector field that is used to parameterize the complete diffeomorphisms. While this is consistent with its use of the exponential and allows him to propose an inverse consistent scheme, it is somewhat difficult to interpret the meaning of this regularization and, most importantly, it introduces coupling in the minimization problem. Similarly, since the exponential is not used around zero, we need to compute the derivative of the exponential far from the identity. This introduces additional tight coupling. It is therefore necessary to solve a very large system of equations with a non-trivial sparsity pattern. For typical MR images, in our case, one optimization step requires a time of the order of a second, whereas in (Ashburner, 2007) the author reports a computation time of around a minute.

In between these approaches we recently proposed in (Vercauteren et al., 2008) a symmetric diffeomorphic registration algorithm that also parameterizes the entire spatial transformation with a stationary speed vector field. An efficient demons-based scheme could however still be applied thanks to the Baker-Campbell-Hausdorff (BCH) approximation proposed in (Bossa et al., 2007).

It should also be mentioned that the exponential of (Arsigny et al., 2006) is not the only option available to work with diffeomorphisms when time-varying velocity fields are considered to be too demanding from a computational and memory requirement point of view. Another approach, could for example, rely on constraining the large deformation setting by using transformations that satisfy the initial momentum conservation (Marsland and McLachlan, 2007; Younes, 2007). Such representations may also help reducing the influence of the numerical issues involved with the exponential but may at the same time have a computational price.

6. Conclusions

We showed that the demons algorithm could be seen as an optimization procedure on the entire space of displacement fields. By using the efficient second-order minimization (ESM) framework of (Benhimane and Malis, 2004; Malis, 2004), the matrix inversion lemma and a local estimation of the image noise, we improved the current understanding of the demons algorithm. A pertinent comparison between the different variants of the demons was provided.

Our analysis predicted a theoretical advantage for the symmetric forces variant of the demons algorithm which we confirmed on the practical side.

The final goal of understanding an algorithm is to improve it. One of the main limitations of the demons algorithm is that it does not provide diffeomorphic transformations. By combining a recently developed Lie group framework on diffeomorphisms and an optimization procedure for Lie groups, we showed that the framework in which we cast the demons algorithm could be adapted to provide non-parametric diffeomorphic transformations. Our experiments have shown that our algorithm provides, with respect to the additive demons algorithm, very similar results in terms of visual appearance. This is however achieved with diffeomorphic transformations that are smoother and closer to the true transformations in terms of Jacobians.

Thanks to the open-source implementation of our diffeomorphic demons we proposed in (Vercauteren et al., 2007b), our algorithm has been successfully tested by several independent groups. In (Urschler et al., 2007), the authors reported that our algorithm outperformed several other non-rigid registration schemes. Our algorithm has also been integrated into MedINRIA, the free medical image navigation and research tool of Asclepios research group, INRIA Sophia Antipolis (Toussaint et al., 2007).

In this paper, we have focused on a rather specific registration problem since we only used the simplest image similarity criterion, the mean squared error. Nevertheless, the framework we proposed showed to be versatile enough to be extended to many types of images such as DTI (Yeo et al., 2008c), Cortical surfaces (Yeo et al., 2008b) or 4D time series of cardiac images (Peyrat et al., 2008). By borrowing some ideas from (Avants et al., 2008; Cachier et al., 2003; Hermosillo et al., 2002), other similarity criteria could also be used. Finally, a symmetric extension of the diffeomorphic demons was recently proposed in (Vercauteren et al., 2008).

References

Arsigny, V., Commowick, O., Pennec, X., Ayache, N., 2006. A Log-Euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (Eds.), Proc. MICCAI'06. Vol. 4190 of LNCS. Springer-Verlag, pp. 924–931.

Ashburner, J., Oct. 2007. A fast diffeomorphic image registration algorithm. Neuroimage 38 (1), 95–113.

Aubert-Broche, B., Griffin, M., Pike, G. B., Evans, A. C., Collins, D. L., Nov. 2006. Twenty new digital brain phantoms for creation of validation image data bases. IEEE Trans. Med. Imag. 25 (11), 1410–1416.

Avants, B. B., Epstein, C. L., Grossman, M., Gee, J. C., Feb. 2008. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. Med. Image Anal. 12 (1), 26–41.

Barron, J. L., Fleet, D. J., Beauchemin, S. S., Feb. 1994. Performance of optical flow techniques. Int. J. Comput. Vis. 12 (1), 43–77.

Bazin, P.-L., Ellingsen, L. M., Pham, D. L., Jul. 2007. Digital homeomorphisms in deformable registration. In: Karssemeijer, N., Lelieveldt, B. P. F. (Eds.), Proc. IPMI'07. Vol. 4584 of LNCS. Springer-Verlag, Kerkrade, The Netherlands, pp. 211–222.

Beg, M. F., Miller, M. I., Trouvé, A., Younes, L., Feb. 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. Int. J. Comput. Vis. 61 (2), 139–157.

Benhimane, S., Malis, E., Sep. 2004. Real-time image-based tracking of planes using efficient second-order minimization. In: Proc. IROS'04. Sendai, Japan, pp. 943–948.

Benhimane, S., Malis, E., Jul. 2007. Homography-based 2d visual tracking and servoing. Int. J. Rob. Res. 26 (7), 661–676.

Bossa, M., Hernandez, M., Olmos, S., Oct. 2007. Contributions to 3D diffeomorphic atlas estimation: Application to brain images. In: Ayache, N., Ourselin, S., Maeder, A. J. (Eds.), Proc. MICCAI'07. Vol. 4792 of LNCS. Springer-Verlag, Brisbane, Australia, pp. 667–674.

Bro-Nielsen, M., Gramkov, C., Sep. 1996. Fast fluid registration of medical images. In: Proc. VBC'96. pp. 267–276.

Cachier, P., Ayache, N., May 2004. Isotropic energies, filters and splines for vectorial regularization. J. Math. Imag. Vision 20 (3), 251–265.

Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N., Feb. 2003. Iconic feature based nonrigid registration: The PASHA algorithm. Comput. Vis. Image Underst. 89 (2–3), 272–298.

Chefd'hotel, C., Hermosillo, G., Faugeras, O., 2002. Flows of diffeomorphisms for multimodal image registration. In: Proc. ISBI'02. pp. 753–756.

Christensen, G. E., Geng, X., Kuhl, J. G., Bruss, J., Grabowski, T. J., Pirwani, I. A., Vannier, M. W., Allen, J. S., Damasio, H., Jul. 2006. Introduction to the non-rigid image registration evaluation project (NIREP). In: Pluim, J. P., Likar, B., Gerritsen, F. A. (Eds.), Proc. WBIR'06. Vol. 4057/2006 of LNCS. Springer, Utrecht, The Netherlands, pp. 128–135.

Christensen, G. E., Joshi, S. C., Miller, M. I., Dec. 1997. Volumetric transformation of brain anatomy. IEEE Trans. Med. Imag. 16 (6), 864–877.

Christensen, G. E., Rabitt, R. D., Miller, M. I., Oct. 1996. Deformable templates using large deformation kinematics. IEEE Trans. Image Process. 5 (10), 1435–1447.

Helgason, S., 2001. Differential Geometry, Lie Groups, and Symmetric Spaces. American Mathematical Society.

Hermosillo, G., Chefd'Hotel, C., Faugeras, O., Dec. 2002. Variational methods for multimodal image matching. Int. J. Comput. Vis. 50 (3), 329–343.

Hernandez, M., Bossa, M. N., Olmos, S., Oct. 2007. Registration of anatomical images using geodesic paths of diffeomorphisms parameterized with stationary vector fields. In: Proc. MMBIA Workshop of ICCV'07. Rio de Janeiro, Brazil, pp. 1–8.

Hernandez, M., Olmos, S., Pennec, X., Sep. 2008. Comparing algorithms for diffeomorphic registration: Stationary LDDMM and diffeomorphic demons. In: Proc. MFCA Workshop of MICCAI'08. New York, NY.

Ibáñez, L., Schroeder, W., Ng, L., Cates, J., 2005. The ITK Software Guide. Kitware, Inc., 2nd Edition.

Joshi, S. C., Miller, M. I., Aug. 2000. Landmark matching via large deformation diffeomorphisms. IEEE Trans. Image Process. 9 (8), 1357–1370.

Leow, A. D., Yanovsky, I., Chiang, M.-C., Lee, A. D., Klunder, A. D., Lu, A., Becker, J. T., Davis, S. W., Toga, A. W., Thompson, P. M., 2007. Statistical properties of jacobian maps and the realization of unbiased large-deformation nonlinear image registration. IEEE Trans. Med. Imag. 26 (6), 822–832.

Lepore, N., Brun, C., Chou, Y.-Y., Chiang, M.-C., Dutton, R. A., Hayashi, K. M., Luders, E., Lopez, O. L., Aizenstein, H. J., Toga, A. W., Becker, J. T., Thompson, P. M., Jan. 2008. Generalized tensor-based morphometry of HIV/AIDS using multivariate statistics on deformation tensors. IEEE Trans. Med. Imag. 27 (1), 129–141.

Mahony, R., Manton, J. H., Aug. 2002. The geometry of the Newton method on non-compact Lie-groups. J. Global Optim. 23 (3), 309–327.

Malis, E., Apr. 2004. Improving vision-based control using efficient second-order minimization techniques. In: Proc. ICRA'04.

Marsland, S., McLachlan, R. I., Jul. 2007. A Hamiltonian particle method for diffeomorphic image registration. In: Karssemeijer, N., Lelieveldt, B. P. F. (Eds.), Proc. IPMI'07. Vol. 4584 of LNCS. Springer-Verlag, Kerkrade, The Netherlands, pp. 396–407.

Marsland, S., Twining, C., 2004. Constructing diffeomorphic representations for the groupwise analysis of non-rigid registrations of medical images. IEEE Trans. Med. Imag. 23 (8), 1006–1020.

Miller, M. I., Joshi, S. C., Christensen, G. E., 1998. Large deformation fluid diffeomorphisms for landmark and image matching. In: Toga, A. (Ed.), Brain Warping. Elsevier, pp. 115–131.

27

Modersitzki, J., 2004. Numerical Methods for Image Registration. Oxford University Press.

Noblet, V., Heinrich, C., Heitz, F., Armspach, J.-P., May 2005. 3-D deformable image registration: A topology preservation scheme based on hierarchical deformation models and interval analysis optimization. IEEE Trans. Image Process. 14 (5), 553–566.

Pennec, X., Cachier, P., Ayache, N., 1999. Understanding the *demon's algorithm*: 3D non-rigid registration by gradient descent. In: Proc. MICCAI'99. Vol. 1679 of LNCS. Springer-Verlag, pp. 597–605.

Peyrat, J.-M., Sermesant, M., Pennec, X., Delingette, H., Xu, C., Ayache, N., Sep. 2008. Registration of 4D time-series of cardiac images with multichannel diffeomorphic demons. In: Proc. MICCAI'08. LNCS. Springer-Verlag, New York, USA.

Rogelj, P., Kovačič, S., Jun. 2006. Symmetric image registration. Med. Image Anal. 10 (3), 484–493.

Rueckert, D., Aljabar, P., Heckemann, R. A., Hajnal, J. V., Hammers, A., 2006. Diffeomorphic registration using B-splines. In: Larsen, R., Nielsen, M., Sporring, J. (Eds.), Proc. MICCAI'06. Vol. 4191 of LNCS. Springer-Verlag, pp. 702–709.

Thirion, J.-P., 1998. Image matching as a diffusion process: An analogy with Maxwell's demons. Med. Image Anal. 2 (3), 243–260.

Toussaint, N., Souplet, J.-C., Fillard, P., Nov. 2007. MedINRIA: Medical image navigation and research tool by INRIA. In: Proc. IMIAV Workshop of MICCAI'07. Brisbane, Australia.

Trouvé, A., Jul. 1998. Diffeomorphisms groups and pattern matching in image analysis. Int. J. Comput. Vis. 28 (3), 213–221.

Urschler, M., Kluckner, S., Bischof, H., Nov. 2007. A framework for comparison and evaluation of nonlinear intra-subject image registration algorithms. Insight Journal – ISC/NA-MIC Workshop on Open Science at MICCAIAvailable online with code source at `http://hdl.handle.net/1926/561`.

Vaillant, M., Miller, M. I., Younes, L., Trouvé, A., 2004. Statistics on diffeomorphisms via tangent space representations. Neuroimage 23 (S1), 161–169.

Vercauteren, T., Jan. 2008. Image registration and mosaicing for dynamic in vivo fibered confocal microscopy. PhD thesis, École Nationale Supérieure des Mines de Paris.

Vercauteren, T., Pennec, X., Malis, E., Perchant, A., Ayache, N., Jul. 2007a. Insight into efficient image registration techniques and the demons algorithm. In: Karssemeijer, N., Lelieveldt, B. P. F. (Eds.), Proc. IPMI'07. Vol. 4584 of LNCS. Springer-Verlag, Kerkrade, The Netherlands, pp. 495–506.

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., Nov. 2007b. Diffeomorphic demons using ITK's finite difference solver hierarchy. Insight Journal – ISC/NA-MIC Workshop on Open Science at MICCAIAvailable online with source code at `http://hdl.handle.net/1926/510`.

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., Oct. 2007c. Nonparametric diffeomorphic image registration with the demons algorithm. In: Ayache, N., Ourselin, S., Maeder, A. J. (Eds.), Proc. MICCAI'07. Vol. 4792 of LNCS. Springer-Verlag, Brisbane, Australia, pp. 319–326.

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., Sep. 2008. Symmetric log-domain diffeomorphic registration: A demons-based approach. In: Proc. MICCAI'08. LNCS. Springer-Verlag, New York, USA.

Wang, H., Dong, L., O'Daniel, J., Mohan, R., Garden, A. S., Ang, K. K., Kuban, D. A., Bonnen, M., Chang, J. Y., Cheung, R., 2005. Validation of an accelerated 'demons' algorithm for deformable image registration in radiation therapy. Phys. Med. Biol. 50 (12), 2887–2905.

Wüstner, M., 2003. A connected Lie group equals the square of the exponential image. J. Lie Theory 13 (1), 307–309.

Xue, Z., Shen, D., Karacali, B., Stern, J., Rottenberg, D., Davatzikos, C., Nov. 2006. Simulating deformations of MR brain images for validation of atlas-based segmentation and registration algorithms. Neuroimage 33 (3), 855–866.

Yeo, B. T. T., Sabuncu, M., Desikan, R., Fischl, B., Golland, P., Oct. 2008a. Effects of registration regularization and atlas sharpness on segmentation accuracy. Med. Image Anal. 12 (5), 603–615.

Yeo, B. T. T., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., Golland, P., Sep. 2008b. Spherical demons: Fast surface registration. In: Proc. MICCAI'08. LNCS. Springer-Verlag, New York, USA.

Yeo, B. T. T., Vercauteren, T., Fillard, P., Pennec, X., Golland, P., Ayache, N., Clatz, O., May 2008c. DTI registration with exact finite-strain differential. In: Proc. ISBI'08. IEEE, Paris, France.

Younes, L., Feb. 2007. Jacobi fields in groups of diffeomorphisms and applications. Quart. Appl. Math. 65, 113–134.