

Left-Invariant Metrics for Diffeomorphic Image Registration with Spatially-Varying Regularisation

Tanya Schmah¹, Laurent Risser², and François-Xavier Vialard³

¹ Rotman Research Institute, Baycrest, Toronto, Canada

² CNRS, Institut de Mathématiques de Toulouse (UMR 5219), France*

³ CEREMADE (UMR 7534), Université Paris Dauphine, France

Abstract. We present a new framework for diffeomorphic image registration which supports natural interpretations of spatially-varying metrics. This framework is based on *left-invariant diffeomorphic metrics* (LIDM) and is closely related to the now standard *large deformation diffeomorphic metric mapping* (LDDMM). We discuss the relationship between LIDM and LDDMM and introduce a computationally convenient class of spatially-varying metrics appropriate for both frameworks. Finally, we demonstrate the effectiveness of our method on a 2D toy example and on the 40 3D brain images of the LPBA40 dataset.

1 Introduction

Medical image registration often consists in estimating the transformation ϕ which “best” maps images I and J . In diffeomorphic registration frameworks, ϕ is constrained to be a diffeomorphism, and in particular invertible. Successful diffeomorphic approaches include the Large Deformation Diffeomorphic Metric Matching framework (LDDMM) [13,3], and the closely-related Symmetric Normalisation (SyN) algorithm [2], as well as LogDemons [14]. In the LDDMM framework, we seek a path of diffeomorphisms $\phi(t)$, such that $\phi(0)$ is the identity and $\phi(1)$ is the final transformation of the image. The *spatial (Eulerian) velocity* of a path $\phi(t)$ is the time-varying vector field v defined by

$$\partial_t \phi(t) = v(t) \circ \phi(t), \quad (1)$$

where the symbol $\circ$ denotes composition. Given a Hilbert space V of smooth vector fields with norm $\|\cdot\|_V$, the matching problem is to find a time-varying vector field v that minimises the functional

$$\mathcal{J}(v) = \frac{1}{2} \int_0^1 \|v(t)\|_V^2 dt + \frac{\lambda}{2} \|I \circ \phi(1)^{-1} - J\|_{L^2}^2, \quad (2)$$

under the constraint (1). The minimisation problem (2) is well-posed provided that the norm on V is sufficiently strong in terms of smoothness. The first term

* The authors thank the AO1 grant from Université Paul Sabatier (Toulouse, France) and the ANR DEMOS project for funding.

of (2) is a “regularisation term” (or “energy term”) that serves both to guarantee a well-posed problem and to force ϕ to stay “small”. In practice, V is defined by its reproducing kernel K which is often chosen to be Gaussian.

Spatially-varying or non-isotropic regularisation is of great interest in medical applications, where it can represent variable deformability of tissue. For example, [7] models sliding conditions between the lungs and the ribs using piecewise-diffeomorphic transformations, *i.e.* transformations which are diffeomorphic in different regions only, and not in the whole image domain. Another recent example is the direction-dependent regularisation in [8], which computes displacement fields directly (and so large deformations may be non-invertible). Both of these papers use a fixed regularisation scheme based on prior anatomical knowledge. Neither paper uses fully diffeomorphic transformations.

Spatially-varying (or non-isotropic) regularisation within a diffeomorphic framework is clearly of interest, however it has not appeared in the literature until now. In LDDMM and SyN, we believe that this is because their standard interpretation in terms of moving source images does not support a natural interpretation of a spatially-varying regularisation kernel. Indeed in LDDMM, consider a deformation path $\phi(t)$ and a point X in the source image. As X moves along the path $\phi(t)(X)$, the contribution of its spatial velocity (defined by (1)) to the functional $\mathcal{J}(v)$ in (2) depends on the value of the regularisation kernel K at the (moving) point $\phi(t)(X)$. Conversely the value of K at a single point x affects the regularisation along a whole curve of points in the source image. Thus there is no sense in which $K(x)$ at a single value of x can be said to describe deformability at a single point of the source image.

With this motivation, we propose a new diffeomorphic registration framework that *does* support natural interpretations of spatially-varying metrics: Left-Invariant Diffeomorphic Matching (LIDM). This framework is analogous to LDDMM but based on a *left*-invariant metric, *i.e.* based on a norm in the body (Lagrangian) coordinates of the source image. This means that instead of the norm being applied to the spatial (Eulerian) velocity defined by (1), it is applied to the *convective velocity* defined by

$$\partial_t \phi(t) = d\phi(t) \cdot v(t), \quad (3)$$

where $d\phi(t)$ is the spatial derivative of $\phi(t)$ and the symbol $\cdot$ denotes the multiplication of a matrix and a vector. The matching problem in LIDM is to minimise the same functional as in LDDMM (2) but under the “new” constraint (3). In this framework, the $v(t)$ in (2) is a convective velocity, which is expressed in body coordinates at all times t . At any given point X in the source image, the contribution of $v(t)(X)$ to the energy term is controlled by $K(X)$ for all t . Hence $K(X)$ regularises the deformation at X , and in this sense describes *a priori* deformability of the source image at X . The parameters of K can be learnt from data, a point to which we return in the Discussion.

In Section 2, we present the gradient calculation of the matching functional in the LIDM model. We then develop in Section 3 the correspondence between LIDM and LDDMM. Finally, the performance of the LIDM model is tested on synthetic and real data in Section 5.

2 A Gradient Descent Algorithm for LIDM

In this section, we apply a standard adjoint calculation in order to compute the gradient of the LIDM functional (2), which is the same as in LDDMM, subject to Eq. (3), and not Eq. (1). Our method is very similar to that in [11]. The first step is to write the constraint (3) in the following form:

$$\partial_t \phi_i(t) = \langle \nabla \phi_i, v(t) \rangle, \quad (4)$$

for $i = 1, \dots, d$, where d is the dimension of the ambient space, ϕ_i is the i^{th} coordinate of ϕ , and $\langle \cdot, \cdot \rangle$ the standard scalar product on $\mathbb{R}^d$. As usual, we introduce time dependent Lagrange multipliers denoted by $P_i(t) \in L^2(\mathbb{R}^d, \mathbb{R})$ in order to compute the adjoint equations, the augmented functional is then

$$\tilde{\mathcal{J}}(v, P) = \mathcal{J}(v) + \sum_{i=1}^d \int_0^1 \langle P_i, \partial_t \phi_i(t) - \langle \nabla \phi_i, v(t) \rangle_{\mathbb{R}^d} \rangle_{L^2(\mathbb{R}^d, \mathbb{R})} dt. \quad (5)$$

Note that the scalar product $\langle \cdot, \cdot \rangle_{L^2(\mathbb{R}^d, \mathbb{R})}$ is the usual pairing in $L^2(\mathbb{R}^d, \mathbb{R})$. The gradient of (5) can be computed by taking free variations of the augmented functional. Variations w.r.t. v lead to

$$\nabla \tilde{\mathcal{J}}(v)(t) = v(t) - K \langle \nabla \phi(t), P(t) \rangle_{\mathbb{R}^d}, \quad (6)$$

where K denotes the isomorphism from V^* (the dual of V) to V and $\langle \nabla \phi(t), P(t) \rangle_{\mathbb{R}^d} = \sum_{i=1}^d P_i \nabla \phi_i$. In addition, $P(t) = [P_i]_{i=1, \dots, d}$ is equivalent to a vector field and solves $\dot{P}_i(t) = \nabla \cdot (P_i(t)v(t))$ where $\nabla \cdot$ stands for the divergence of a vector field. The previous equation is given by taking variations of functional (5) w.r.t. ϕ . Note that variations w.r.t. P_i lead to the reconstruction equation (3). Last, the boundary condition at time 1 for the case of the square of the L^2 norm that appears in (2) is:

$$P(1) = d\phi^{-T}(1)[|D\phi(1)|\Delta \circ \phi(1)\nabla I], \quad (7)$$

where $\Delta = \lambda(I \circ \phi^{-1}(1) - J)$ and $|D\phi(1)|$ denotes the Jacobian determinant of $\phi(1)$. The notation $d\phi^{-T}(1)$ denotes the inverse transpose of $d\phi(1)$. Combining equations (6) and (7), we obtain

$$\nabla \tilde{\mathcal{J}}(v)(t) = v(t) - K m(t), \quad (8)$$

where $m(t) = \nabla \hat{I}(t)|D\phi(t)| \left(\hat{I}(t) - J \circ \phi(t) \right)$ and $\hat{I}(t) := I \circ (\phi(1)^{-1} \circ \phi(t))$.

Perhaps not unexpectedly, those equations are very close to the gradient of the standard LDDMM functional. In the next section, we detail the relations between the two models.

3 Relation to Standard LDDMM

The previous calculation guides us towards the strong relation between LIDM and LDDMM models. By comparing the LDDMM gradient (see [6,12] for instance) and the LIDM gradient, the reader may infer the following proposition.

Proposition 1. *The optimal LIDM path $\phi(t)$ is given by $\phi(t) = \psi_1 \circ \psi_{1-t}^{-1}$ for $\psi(t)$ the LDDMM optimal path. In particular, the final diffeomorphic mappings are the same in the two models, $\phi(1) = \psi(1)$.*

Proof (Outline). The minimization of the functional (2) can be written as:

$$\mathcal{J}(v) = \frac{1}{2} \int_0^1 \|v(t)\|_V^2 dt + \frac{\lambda}{2} \|I \circ \psi(1) - J\|_{L^2}^2, \quad (9)$$

with $\partial_t \psi(t) = -v(t) \circ \psi(t)$. This is close to the standard LDDMM formulation but the two differences are (1) the use of ψ^{-1} instead of ψ and (2) the minus sign in the previous flow equation. Moreover, the first term of the functional is the square of the right-invariant distance d on the group of diffeomorphisms so that $d(Id, \psi) = d(\psi^{-1}, Id)$. Indeed, if $v_R(t)$ is a geodesic vector field for LDDMM between Id and ψ , then $-v_R(1-t)$ is a geodesic between Id and ψ^{-1} . We therefore have shown that to an optimal path $v(t)$ of functional (2) corresponds an optimal path $v_R(t)$ of the corresponding LDDMM functional such that $v(t) = v_R(1-t)$.

In summary, the final diffeomorphic mapping is the same in both approaches but the diffeomorphism paths do differ. There are two optimal paths from Id to ϕ_1 : one left- and one right- geodesic. Fig. 1 illustrates the different optimal paths given by the two models in an exact matching problem of points (landmarks).




Fig. 1. Deformations from the left-most source image to the right-most target image. Green and blue curves show the optimal LDDMM and LIDM paths, respectively.

4 Spatially-Varying Metrics

The LIDM model opens up the opportunity to use spatially-varying and non-isotropic metrics, as explained in the Introduction. If some *a priori* information on the deformation intensity is known in body coordinates, then this can be modelled in the regularisation kernel. We now give a simple example of this.

As an idealised situation of interest, one can consider a partition of the template image domain $\Omega \subset \mathbb{R}^d$ into n parts $(\Omega_i)_{i=1}^n$. To each part, one can associate a Gaussian kernel k_{σ_i} with a given smoothing parameter σ_i that incorporates some knowledge on the template. We also introduce a smooth partition of unity $\chi_i : \Omega \mapsto [0, 1]$ satisfying $\sum_{i=1}^n \chi_i(x) = 1$ for $x \in \Omega$ and χ_i vanishing outside a compact containing Ω_i . We can build the new kernel K :

$$K(x, y) = \sum_{i=1}^n \chi_i(x) k_{\sigma_i}(x, y) \chi_i(y) \text{ where } x, y \in \Omega.$$

In addition, this kernel has a variational interpretation (omitted here to save space) that justifies its form. This is the type of kernels we use for our experimental results on synthetic and real data in Section 5. Such a definition was motivated (1) by its simplicity to introduce different smoothing on the partition and (2) by computational considerations. Indeed, the computational cost of the kernel K is n times the cost of a single Gaussian kernel so that having two or three regions of interest is feasible in practical situations.

5 Results

5.1 Results on 2D Phantom Images

In this subsection, we register the phantom images shown in Fig. 2. We registered these images using: LIDM, Symmetric Normalization (SyN)¹ [2] and multi-kernel LDDMM² [6]. Note that SyN is closely-related to LDDMM, and its implementation in ANTS is considered to be state-of-the-art for neuroimaging [5]. Specifically, we used the following techniques: **(LIDM)** Considering the partition of unity shown in Fig. 2, we used Gaussian smoothing kernels with standard deviations 33 and 4 pixels in the white and black regions, respectively. These values were chosen because they are half of the width of the dark grey circle and the small black structures protruding from its inner surface, respectively. **(SyN _{σ_f, σ_d})** We used the command "ANTS 2 -m MSQ[target,source,1,0] -i 100x100x1000 -r Gauss[σ_f, σ_d] -t SyN[0.4]". Gaussian kernels with various standard deviations to perform fluid-like (σ_f) and diffusion-like (σ_d) regularization were tested: After registering the images using a large kernel ($\sigma_f = 32$, $\sigma_d = 1$: SyN_{init}), we composed the deformation with those obtained using finer kernels ($\sigma_f = \{1, 2, 4, 8, 16\}$ and $\sigma_d = 0$, or $\sigma_f = 1$ and $\sigma_d = \{\frac{1}{8}, \frac{1}{4}, \frac{1}{2}, 1, 2, 4\}$). **(MK-LDDMM)** We used a kernel constructed as the sum of 7 Gaussian kernels having standard deviations linearly sampled between 1 and 50 and weights automatically tuned (option -M_Gauss_easier of uTilzReg-LDDMM).

To evaluate the registration quality, we measured the sum of the square differences (SSD) between the registered images in the whole image domain Ω and in the ROI shown in Fig. 2, as well as the maximum determinant of the Jacobians (DetJ). SSDs were normalised by the corresponding SSD between the original images. Results are presented in Fig. 2 and Table 1. For SyN, we present a selection of results giving the best balance between SSD and DetJ.

The lowest SSDs, 0.14 for Ω and 0.21 in the ROI, were obtained by LIDM. A competitive strategy is (SyN) with fluid regularization ($\sigma_f > 0$), where $\text{SSD}_\Omega \in [0.20, 0.23]$ and $\text{SSD}_{\text{ROI}} \in [0.36, 0.38]$, for $\max \text{DetJ} \in [3.59, 4.06]$. The matching is however higher using (LIDM) for a similar $\max \text{DetJ}$.

¹ <http://www.picsl.upenn.edu/ANTS/>

² <http://sourceforge.net/projects/utilzreg/>




Fig. 2. Results of image registration tests on a synthetic example

Table 1. Quantitative results obtained on the synthetic images shown in Fig. 2. Subscripts after SyN indicate the degree of fluid- and diffusion-like regularisation (σ_f, σ_d).

	LIDM	SyN _{init}	SyN _{2,0}	SyN _{8,0}	SyN _{1,0.25}	SyN _{1,2}	MK-LDDMM
SSD _Ω	0.14	0.46	0.21	0.23	0.19	0.28	0.29
SSD _{ROI}	0.21	0.73	0.36	0.37	0.37	0.63	0.68
max DetJ	3.70	2.06	4.06	3.59	3.67	2.35	2.6

5.2 Results on 3D Brain Images

We performed additional tests on the 40 subjects of the LONI Probabilistic Brain Atlas (LPBA40) [9], using the probabilistic tissue maps (white matter, grey matter and CSF). We first resampled all images to a resolution of 1 mm and aligned them to subject 09 using non-rigid registration with a very large smoothing kernel (SyN with the option `-r Gauss[50,1]`). We then constructed a partition of unity by dilating (structuring element of 2 mm) and smoothing (Gaussian kernel of 3 mm) the grey matter density map of subject 09. Finally, we registered subject 09 to all other ones using two strategies: (**LIDM**) We first used LIDM with the predefined partition of unity and Gaussian kernels having standard deviations of 7 mm around the grey matter and 33 mm elsewhere. The underlying idea here is that we allow more flexibility for the registration around the grey matter than in the rest of the image. (**SyN_{σ_f,σ_d}**) We also performed SyN registration with different regularisation parameters: `-r Gauss[σf,σd]`, where $(\sigma_f, \sigma_d) = \{(7, 2), (3, 3), (10, 4), (5, 1), (2, 2)\}$ as well as the parameters of Klein *et al* study [5]. We computed the SSD between each pair of registered images, the target overlaps between the segmented brain regions (see [5]), and the maximum DetJ of the deformations. SSDs were normalised by the SSD before registration.

As shown Table 2, SyN gives slightly better average results than LIDM here. However, normalised SSDs, overlaps and max DetJs are within the range of values produced by varying the SyN parameters. This shows the validity of LIDM for the analysis of 3D brain images and opens up interesting perspectives for further investigations with applications in medical imaging, as discussed in the next section.

Table 2. Average results obtained on the 40 3D brain images of the LPBA40 dataset

	LIDM	SyN _{7,2}	SyN _{3,3}	SyN _{10,4}	SyN _{5,1}	SyN _{2,2}	SyN _[5]
SSD	0.28	0.21	0.23	0.30	0.17	0.21	0.21
Overlap	0.723	0.717	0.716	0.713	0.715	0.708	0.728
max DetJ	5.02	4.55	3.20	2.81	6.38	4.85	5.36

6 Discussion

We have introduced a novel diffeomorphic matching framework, Left Invariant Diffeomorphic Matching (LIDM), in which spatially-varying and directionally-dependent regularisation kernels can encode local deformability properties of the source image. Through the relationship between LIDM and LDDMM described in Section 3, it also follows that spatially-varying and directionally-dependent kernels in LDDMM are interpretable in the same way, which has not been remarked upon before.

We have demonstrated the value of spatially-varying kernels in registration, in experiments with both synthetic and real data (brain MRI). In both experiments, we applied LIDM with a fixed spatially-varying kernel, chosen on the basis of observed feature scales (for the synthetic example) and our experience with other algorithms (for the real data). We compared LIDM with two state-of-the-art algorithms: SyN (implemented in ANTS) and MK-LDDMM. For SyN, we explored a wide range of regularisation parameters. In the synthetic example, LIDM produced superior matches, as judged by sum of squared differences (SSD), while for the real data, LIDM produced results similar to SyN but not as good as the best SyN result. It is significant that in both cases, LIDM gave a good match with the first (and only) parameters that we chose, which suggests an important advantage in ease of use.

The main motivation driving our work is to automatically learn spatially-varying and directionally-dependent regularisation parameters, as has been done by Simpson et al. [10] for global regularisation parameters. Our contribution justifies this project in the diffeomorphic context, in LIDM and LDDMM and also, by extension, in SyN. Various methods could be used to optimise the parameters for a population of targets, including Bayesian methods related to those in [1,4].

References

1. Allasonnière, S., Amit, Y., Trouné, A.: Towards a coherent statistical framework for dense deformable template estimation. *J. R. Statist. Soc. B* 69(1), 3–29 (2007)
2. 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)
3. Beg, M.F., Miller, M.I., Trouné, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vision* 61(2), 139–157 (2005)
4. Cotter, C.J., Cotter, S.L., Vialard, F.X.: Bayesian data assimilation in shape registration. *ArXiv e-prints* (December 2012)
5. Klein, A., Ghosh, S.S., Avants, B.B., Yeo, B.T.T., Fischl, B., Ardekani, B.A., Gee, J.C., Mann, J.J., Parsey, R.V.: Evaluation of volume-based and surface-based brain image registration methods. *NeuroImage* 51(1), 214–220 (2010)
6. Risser, L., Vialard, F.X., Wolz, R., Murgasova, M., Holm, D.D., Rueckert, D.: Simultaneous Multi-scale Registration Using Large Deformation Diffeomorphic Metric Mapping. *IEEE Transactions on Medical Imaging* 30(10), 1746–1759 (2011)
7. Risser, L., Vialard, F.X., Baluwala, H.Y., Schnabel, J.A.: Piecewise-diffeomorphic image registration: Application to the motion estimation between 3D CT lung images with sliding conditions. *Medical Image Analysis* 17, 182–193 (2012)
8. Schmidt-Richberg, A., Werner, R., Handels, H., Ehrhardt, J.: Estimation of slipping organ motion by registration with direction-dependent regularization. *Medical Image Analysis* 16, 150–159 (2012)
9. Shattuck, D.W., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K.L., Poldrack, R.A., Bilder, R.M., Toga, A.W.: Construction of a 3D probabilistic atlas of human cortical structures. *NeuroImage* 39, 1064–1080 (2008)
10. Simpson, I.J.A., Schnabel, J.A., Groves, A.R., Andersson, J.L.R., Woolrich, M.W.: Probabilistic inference of regularisation in non-rigid registration. *NeuroImage* 59(3), 2438–2451 (2012)
11. Singh, N.P., Hinkle, J., Joshi, S., Fletcher, P.T.: A vector momenta formulation of diffeomorphisms for improved geodesic regression and atlas construction. In: *IEEE Proceedings of ISBI 2013* (2013)
12. Sommer, S., Nielsen, M., Lauze, F., Pennec, X.: A multi-scale kernel bundle for LD-DMM: Towards sparse deformation description across space and scales. In: Székely, G., Hahn, H.K. (eds.) *IPMI 2011. LNCS*, vol. 6801, pp. 624–635. Springer, Heidelberg (2011)
13. Trouné, A.: Diffeomorphic groups and pattern matching in image analysis. *Int. J. Comput. Vision* 28, 213–221 (1998)
14. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric log-domain diffeomorphic registration: A demons-based approach. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I. LNCS*, vol. 5241, pp. 754–761. Springer, Heidelberg (2008)