A Multi-scale Kernel Bundle for LDDMM: Towards Sparse Deformation Description across Space and Scales

Stefan Sommer[1], Mads Nielsen[1,2], François Lauze[1,2], and Xavier Pennec[3]

[1] Dept. of Computer Science, Univ. of Copenhagen, Denmark
sommer@diku.dk
[2] Synarc Imaging Technologies, Rødovre, Denmark
[3] Asclepios Project-Team, INRIA Sophia-Antipolis, France

Abstract. The Large Deformation Diffeomorphic Metric Mapping framework constitutes a widely used and mathematically well-founded setup for registration in medical imaging. At its heart lies the notion of the regularization kernel, and the choice of kernel greatly affects the results of registrations. This paper presents an extension of the LDDMM framework allowing multiple kernels at multiple scales to be incorporated in each registration while preserving many of the mathematical properties of standard LDDMM. On a dataset of landmarks from lung CT images, we show by example the influence of the kernel size in standard LDDMM, and we demonstrate how our framework, LDDKBM, automatically incorporates the advantages of each scale to reach the same accuracy as the standard method optimally tuned with respect to scale. The framework, which is not limited to landmark data, thus removes the need for classical scale selection. Moreover, by decoupling the momentum across scales, it promises to provide better interpolation properties, to allow sparse descriptions of the total deformation, to remove the trade-off between match quality and regularity, and to allow for momentum based statistics using scale information.

Keywords: diffeomorphic registration, computational anatomy, LDDKBM, LDDMM, scale, sparsity, kernels, momentum, landmarks.

1 Introduction

Among the many methods for non-rigid registration in medical imaging, the Large Deformation Diffeomorphic Metric Mapping framework (LDDMM) has the benefit of both providing good registrations and drawing strong theoretical links with Lie group theory and evolution equations in physical modeling [5,17]. The mathematical foundation ensures that optimization procedures find diffeomorphic optima and allows statistics to be performed on the results of registrations [9,15].

Diffeomorphisms in the LDDMM framework are regularized by a norm, which not only greatly influences the computed registration but also affects subsequent

G. Székely and H.K. Hahn (Eds.): IPMI 2011, LNCS 6801, pp. 624–635, 2011.

statistics. The norm is often connected to a kernel, and since deformation frequently occurs at different scales, the choice of appropriate kernel scale will in such cases involve compromises between regularity and registration quality. In addition, in order to reduce computational complexity and provide meaningful statistics, it would be desirable to introduce sparsity in the framework which further complicates the choice of kernels.

In this paper, we propose a generalization of LDDMM allowing multiple kernels at multiple scales to be incorporated into the registration. We extend the theory by introducing a multi-scale kernel bundle, and the resulting Large Deformation Kernel Bundle Mapping (LDDKBM) framework keeps the strong mathematical foundation while giving significant benefits. Unlike previous methods, our construction explicitly decouples the momentum across scales allowing the algorithm to adapt to the different scales in the input data. This makes classical scale selection unnecessary and removes the common trade-off between match quality and regularity of the diffeomorphism. The ability to describe the deformation at precisely the right scales promises to allow for sparse representations and statistics incorporating scale information.

1.1 Deformation at Multiple Scales; An Example

Figure 1 shows a simple example of landmark matching. In order to register the points, movement is needed at both large and small scales, and the standard LDDMM algorithm faces a compromise between regularity of the diffeomorphism and match quality. The registration in the fourth image is performed by the method developed in this paper incorporating kernels at different scales, and the result shows how the trade-off between regularity and quality has been removed. Furthermore, no cross validation for choosing the kernel size is needed and contrary to previous methods, the computed momentum is represented only at the appropriate scales.

1.2 Related Work

Besides LDDMM, many methods for non-rigid registration are in use today including elastic methods [11], parametrizations using static velocity fields [1,8] and the demons algorithm [13,16]. The deformable template model pioneered by Grenander in [7] and the flow approach by Christensen et al. [4] was paramount in the development of LDDMM together with the theoretical contributions of Dupuis et al. and Trouvé [6,14]. Algorithms for computing optimal diffeomorphisms have been developed in [2], and [15] uses the momentum representation for statistics and develops a momentum based algorithm for the landmark matching problem. The review paper [18] and the book [17] provide excellent overviews of the theory and applications of LDDMM in medical imaging.

The role of the kernel and deformation at different scales have been addressed by Risser et al. in [12], where the authors propose a multi-kernel LDDMM approach which constructs new kernel shapes by adding Gaussian kernels. Our method differs from this approach by extending LDDMM to allow decoupling of the energy and momentum at each scale, and it therefore enables the algorithm

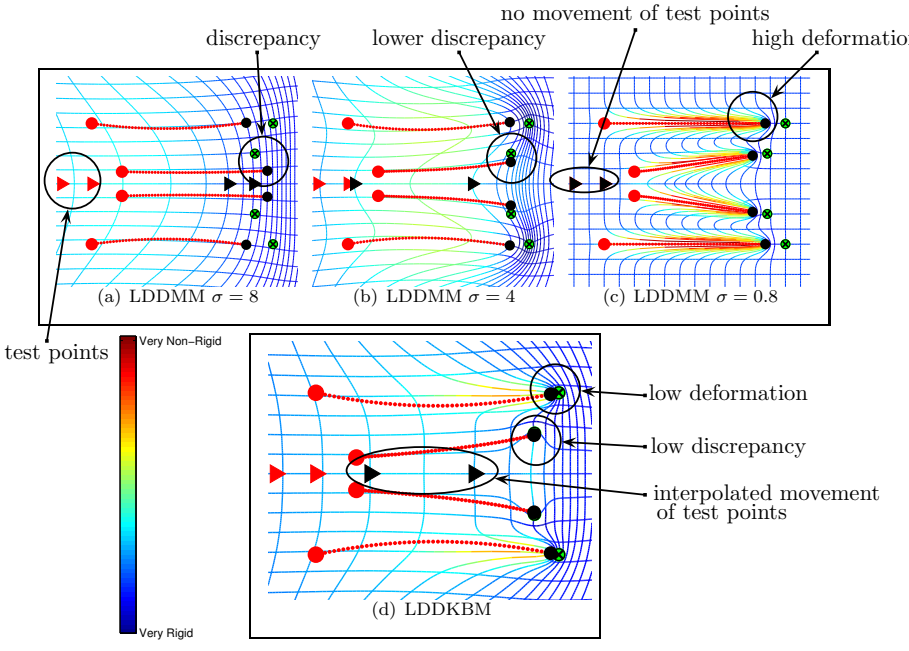

Fig. 1. Matching four landmarks (red) to four landmarks (green, crossed) with results (black) for Gaussian kernels of difference sizes (scale σ in grid units). Top row: standard LDDMM, bottom row: the proposed LDDKBM method. An initially square grid is shown deformed by each diffeomorphism along with two test points (red triangles) transported by the diffeomorphism to the black triangles. The grids are colored with the trace of Cauchy-Green strain tensor (log-scale). The compromise between regularity and match quality is visible for the standard LDDMM method. The poor match for the smallest scale is due to the lack of "carpooling" effect [10]. Notice that for standard LDDMM with the smallest scales, the test points are not moved at all or split up. The proposed method obtain both a regular diffeomorphism and a good match, and the test points are transported in an arguably natural way.

to select the appropriate deformation at each scale individually. The differences between the methods are detailed later in the paper.

1.3 Content and Outline

In the next section, we give an overview of the LDDMM framework and its use in non-rigid registration in medical imaging. In particular, the focus will be on the role of the kernel. We then progress to developing our key contribution, the multi-scale kernel bundle framework LDDKBM, and highlight the important properties of the construction. We will show how the decoupling of scales differentiates the construction from previous methods and promises sparse representations. The evaluation of the method on a dataset of landmarks obtained from CT

lung images will show the automatic incorporation of the appropriate scale and indicate why the method allows sparsity. The paper thus contributes by

(1) highlighting the role of the kernel in LDDMM and showing how the registration results depend on the choice of kernel,
(2) developing the multi-scale kernel bundle generalization LDDKBM to automatically incorporate multiple scales in the registration,
(3) and showing how the developed method removes the need for the classical scale selection while promising improved momentum based statistics and sparse representations.

We note that while the evaluation is performed on landmark data, the LDDKBM framework is very general, and it will apply to different types of data as well. We are currently working on an implementation for image registration, and we consider this paper a first step which eventually will lead to the method being useful for clinical applications.

2 LDDMM and Kernels

The Large Deformation Diffeomorphic Metric Mapping framework provides a well-defined metric structure on spaces of diffeomorphisms and gives convenient ways to parametrize such spaces. We here give a brief overview of the LDDMM construction before going into details about the use of kernels in the framework. Landmarks will be used for the examples, though both the LDDMM framework and the kernel bundle we will develop also apply to images, curves, surfaces, and tensors.

Registration of geometric objects is often performed by defining an action of diffeomorphisms on the objects before searching for diffeomorphisms matching the objects through the action. For example, in order to register landmarks $x_1, \dots, x_N$ and $y_1, \dots, y_N$ in $\mathbb{R}^d$, $d = 2, 3$, we search for a diffeomorphism $\varphi : \mathbb{R}^d \to \mathbb{R}^d$ such that $\varphi(x_i) = y_i$. Frequently, a perfect match is not possible or even not desirable because noisy data may force the diffeomorphism to be highly irregular. Instead, we search for φ minimizing

$$E(\varphi) = E_1(\varphi) + \lambda U(\varphi) \tag{1}$$

where $E_1(\varphi)$ is a regularization measure, $U(\varphi)$ a measure of the quality of the match, and $\lambda > 0$ a weight. A simple and often used choice for U are the L^2 error which takes the form $U(\varphi) = \sum_{i=1}^{N} \|\varphi(x_i) - y_i\|^2$ for landmarks. In the LDDMM framework, the regularization measure $E_1(\varphi)$ is defined as the minimum energy of paths of diffeomorphisms transporting the identity Id_Ω to φ, i.e.

$$E_1(\varphi) = \min_{v_t \in V, \varphi^v_{01} = \varphi} \int_0^1 \|v_s\|_V^2 \, ds \tag{2}$$

with $\|\cdot\|_V$ being a right invariant Riemannian metric on the tangent space V of derivatives of such paths, and φ^v_{0t} denoting the path starting at Id_Ω with

derivative $\partial_t \varphi_{0t}^v = v_t \circ \varphi_{0t}^v$. This norm is chosen to penalize highly varying paths and, therefore, a low value of $E_1(\varphi)$ implies that the path to reach φ, and hence φ itself, is regular. In addition, the norm and tangent space V is chosen to ensure several important properties, which include ensuring a minimizer for E exists and that integration in V is well-defined. The latter property enables us to represent many diffeomorphisms as endpoints of paths φ_{0t}^v.

2.1 The Tangent Space

The norm on the tangent space V is most often chosen to ensure V is a reproducing kernel Hilbert space [17]. If the domain Ω is a subset of $\mathbb{R}^d$ so that V is a subset of the maps $\Omega \to \mathbb{R}^d$, one can with an appropriate norm $\|\cdot\|_V$ on V show the existence of a kernel $K : \Omega \times \Omega \to \mathbb{R}^{d\times d}$ so that, for any constant vector $a \in \mathbb{R}^d$, the vector field $K(\cdot, x)a \in V$, and $\langle K(\cdot, x)a, K(\cdot, y)b\rangle_V = a^T K(x, y)b$ for all points $x, y \in \Omega$ and all vectors $a, b \in \mathbb{R}^d$. This latter property is denoted the reproducing property and it provides a way to evaluate the inner product on the span of the kernel elements. Tightly connected to the norm and kernels is the momentum, which is named so because of its connection to momentum in classical mechanics. The momentum operator L on V is defined by

$$\langle Lv, w\rangle_{L^2(\Omega)} = \int_\Omega \big(Lv(x)\big)^T w(x)dx = \langle v, w\rangle_V$$

and hence it connects the inner product on V with the inner product in $L^2(\Omega)$. As we will see in the landmark case, the value Lv might be singular and in fact $L\big(K(\cdot, y)a\big)(x)$ is the Dirac measure $\delta_y(x)a$. These relations lead to convenient equations for minimizers of the energy (1). In particular, the EPDiff equations for the evolution of the momentum a_t for optimal paths assert that if φ_t is a path minimizing $E_1(\varphi)$ with $\varphi_1 = \varphi$ minimizing $E(\varphi)$ and v_t is the derivative of φ_t then v_t satisfies the system

$$\begin{aligned} v_t &= \int_\Omega K(\cdot, x)a_t(x)dx \ , \\ \frac{d}{dt}a_t &= -Da_t v_t - a_t \nabla \cdot v_t - (Dv_t)^T a_t \ . \end{aligned}$$

The first equation connects the momentum a_t with the velocity v_t. The EPDiff equations reduces to particular forms for several objects. For landmarks $x_1, \ldots, x_N$, the momentum will be concentrated at $\varphi_t(x_i)$ as Dirac measures $a_{i,t}\delta_{\varphi_t(x_i)}$ leading to the finite dimensional system of ODE's

$$\begin{aligned} v_t &= \sum_{l=1}^N K(\cdot, \varphi_t(x_l))a_{l,t} \ , \quad \frac{d}{dt}\varphi_t(x_i) = v_t(\varphi_t(x_i)) \ , \\ \frac{d}{dt}a_{i,t} &= -\sum_{l=1}^N D_1 K(\varphi_t(x_i), \varphi_t(x_l))a_{i,t}^T a_{l,t} \ . \end{aligned} \tag{3}$$

2.2 Kernels

There is some freedom in the choice of kernel or, equivalently, the operator L but, in lack of general models for e.g. deformations of organs to be registered [11], it is hard to give satisfactory modelling arguments specifying the kernel shapes. Rotational and translational invariance is commonly assumed [17] and the Gaussian kernel $K(x,y) = \exp(\frac{|x-y|^2}{\sigma^2})\mathrm{Id}_d$ is a convenient and often used choice. The scaling factor σ is not limited to Gaussian kernels and a scaling factor needs to be determined for both Gaussian kernels and many other proposed kernel shapes.

Larger scales lead in general to higher regularization and smoother diffeomorphisms, whereas smaller kernels penalize higher frequencies less and often gives better matches. This phenomenon is in particular apparent for objects with sparse information, such as the point matching problem (3) illustrated in Figure 1 and images with e.g. areas of constant intensity.

3 LDDKBM: A Multi-scale Kernel Bundle

To incorporate multiple kernels in the framework, we extend LDDMM by a multi-scale kernel bundle. The construction, which we denote LDDKBM (Large Deformation Diffeomorphism Kernel Bundle Mapping) is very general and allows both multiple kernel shapes and scales though varying the scale of the kernel constitutes the most obvious use. We start by outlaying the framework and later discuss how it deviates from other proposals incorporating scale space ideas.

As described in the previous sections, a time varying vector field v_t in the tangent space V generates a path φ_{0t}^v of diffeomorphisms. This provides a view of the space of diffeomorphisms G_V obtainable in this way as a manifold having V as its tangent space at all points, confer Figure 2(a). The norm $\|\cdot\|_V$ works as a metric making G_V a Riemannian manifold. In this view, the regularization factor E_1 is the energy of the path φ_t, and this energy is tightly connected to the kernel arising from the norm on V.

In order to use more kernels, we extend the tangent space V to the family $W = \{V_r\}_{r\in I_W}$ of Hilbert spaces V_r parametrized by scalars r in a set I_W so that, for each r, V_r is a subspace of the tangent space V but equipped with a different metric. The key point here is to allow the norms $\|\cdot\|_{V_r}$ on V_r and connected kernels K_r to vary with r. In this way, the tangent bundle $G_V \times V$ is extended to a vector bundle $G_V \times W$ which we denote the multi-scale kernel bundle. The space W inherits the vector space structure from each V_r, and it is a pre-Hilbert space with norm $\|w\|_W^2 = \int_{I_W} \|w_r\|_{V_r}^2 dr$. Using the norm, we define the energy E_1^W of paths in W by $E_1^W(w_t) = \int_0^1 \|w_s\|_W^2 ds$.

With appropriate conditions on V_r and the norms $\|\cdot\|_{V_r}$, W is a Hilbert space and, by possibly restricting to a subspace, we can pass from W to V using the map $\Psi : W \to V$ defined by integration $\Psi(w) = \int_{I_W} w_r dr$ for elements $w = \{w_r\}_r \in W$. Since we as usual have a connection between paths in V and

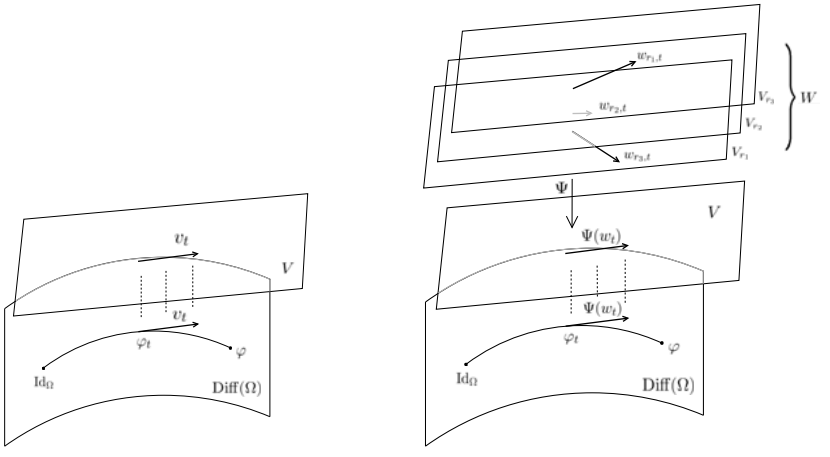

(a) In standard LDDMM, a path on the manifold $\mathrm{Diff}(\Omega)$ is generated by integrating the time varying vector field v_t in the tangent space V.

(b) A path w_t in the vector space W, here constructed from V_{r_1}, V_{r_2}, and V_{r_3}, sum through the map Ψ to a vector field $\Psi(w_t)$, which in turn generates a path on $\mathrm{Diff}(\Omega)$.

Fig. 2. The manifold view of standard LDDMM and LDDKBM

paths on the manifold G_V, we get using Ψ a similar relation between paths $w_t = \{w_{r,t}\}_r$ in W and paths in G_V by

$$w_t \mapsto \varphi_{0t}^{\Psi(w)} \ , \tag{4}$$

i.e. $\varphi_{0t}^{\Psi(w)}$ is the path starting at Id_Ω with derivative $\partial_t \varphi_{0t}^{\Psi(w)} = \Psi(w_t) \circ \varphi_{0t}^{\Psi(w)}$. The path energy E_1^W gives a regularization measure on diffeomorphisms

$$E_1^W(\varphi) = \min_{w_t \in W, \varphi_{01}^{\Psi(w)} = \varphi} \int_0^1 \|w_s\|_W^2 \, ds \tag{5}$$

which, together with a quality of match measure $U(\varphi)$, allows a reformulation of the registration problem as the search for a diffeomorphism minimizing

$$E^W(\varphi) = E_1^W(\varphi) + \lambda U(\varphi) \ . \tag{6}$$

The above formulation should be compared with the standard LDDMM formulation (1) using the regularization (2). It is immediately clear that the proposed method is an extension of standard LDDMM, since the original regularization is the special case with only one scale and hence $W = V$.

We note that the parameter space I_W can be a compact interval or finite set of scalars in which case the integral reduces to just a sum. Often, it will be an interval specifying a scale range, and a practical implementation will discretize the interval into a finite set of scalars.

3.1 Evolution Equations

Many of the properties of standard LDDMM are retained in the LDDKBM construction. Importantly, the evolution equations for solutions of the registration problem take on a form very similar to the EPDiff equations; if $\Psi(w_t)$ is the derivative of the path of diffeomorphisms φ_t minimizing (5) with $\varphi = \varphi_1$ minimizing (6) then

$$\begin{aligned} w_{r,t} &= \int_\Omega K_r(\cdot, x) a_{r,t}(x) dx \ , \\ \frac{d}{dt} a_{r,t} &= \int_{I_W} -D a_{r,t} w_{s,t} - a_{r,t} \nabla \cdot w_{s,t} - (D w_{s,t})^T a_{r,t} ds \ . \end{aligned} \tag{7}$$

with $a_{r,t}$ being the momentum for the part w_r of w. In essence, we just integrate the standard EPDiff equations over the parameter space I_W to obtain the evolution of the momentum at each scale. The system implies that the momentum conservation property holds with the kernel bundle construction.

An important difference from the standard framework relates to the energy along optimal paths. The relation to geodesics in standard LDDMM suggests that the norm $\|v_t\|_V$ is constant in t when v_t is optimal for $E_1(\varphi)$. This is in fact the case for standard LDDMM. For LDDKBM, momentum is conserved along optimal paths of $E_1^W(\varphi)$ though $\|w_t\|_W$ is not constant. This occurs because the new energy is not directly related to a metric in the Riemannian sense.

3.2 Scale Decoupling and Relation to Other Approaches

The energy E_1^W measures the contribution of each scale independently, and the total energy is the sum of the energy at each scale. Because of this, the momentum components can vary over scale, and the method may use any combination of small and large scale features at each spatial location to obtain the best match with least energy at each scale.

In contrast to this, the simultaneous coarse and fine method developed by Risser et al. in [12] builds a kernel by summing Gaussians of different scale. This effectively changes only the shape of the kernel and does not allow different momentum at different scales. Therefore, the method works simultaneously in the sense that no decoupling between scales is present, whereas the method proposed here exactly aims at decoupling the scales. The decoupling will be clear in the experiments and we believe it is desirable for several reasons. We expect better matches because movement at one scale does not imply undesired movement at all other scales. This will also allow sparsity to enter the representation: since movement is represented only at the right scale, the momentum components at all unrelated scales can be zero. Though we currently do not directly enforce sparsity, we will see this effect visually in the experiments. Last, doing statistics on the momentum as proposed in [15] will result in information at each scale because of the decoupling. With the simultaneous coarse and fine method, no scale-specific momentum is available.

3.3 Implementation

As for standard LDDMM, the choice of algorithm for optimizing (6) depends very much on the objects to be matched. For the experiments in this paper, we will match landmarks to which the shooting method of [15] applies. The algorithm does a second order optimization while integrating along critical paths specified by the system (3). We implement a much similar algorithm for the finite dimensional landmark version of the system (7). Because the energy along optimal paths is not constant, the energy must be integrated over the point and velocity trajectories and this rules out the most simple 2nd-order schemes. Therefore, at this point, we use a gradient descent algorithm but we are in the process of developing both 2nd-order and fast 1st-order algorithms.

4 Experiments

We perform experiments on the publicly available DIR [3] dataset of lung CT images and manually annotated landmarks. The dataset consists of five cases of CT images for different stages of the inhale and exhale phases and annotated landmarks for the maximum inhale and exhale phases, respectively. The images and landmarks are available on grids with voxel size varying slightly between the cases but close to $1 \times 1 \times 2.5\,\text{mm}$. Further details can be found in the reference.

For each case, the 300 publicly available landmarks, $x_1^I, \ldots, x_{300}^I$ for the maximum inhale phases and $x_1^E, \ldots, x_{300}^E$ for the maximum exhale phase, correspond pairwise. We will drive the registration using these landmarks and not incorporate image information at this point. We measure the fiducial registration error (FRE) of each computed registration. Since a small FRE indicates a good match on the registered landmarks only, we need an additional way of comparing the quality of the diffeomorphism. We do this by computing the match on random subsets of the landmarks and compute the target registration error (TRE) on the landmarks not included in the match. Choosing kernel scale based on this evaluation is akin to performing cross-validation tuning in machine learning.

4.1 Setup and Results

We choose random subsets of 75 landmarks to drive the registration. For each such choice of subset S and each of the five patient cases, we compute the FRE $(\sum_{j \in S} \|\varphi(x_j^I) - x_j^E\|^2)^{1/2}$ and the TRE $(\sum_{j \notin S} \|\varphi(x_j^I) - x_j^E\|^2)^{1/2}$. We find the relative size of these measures against the respective values before the match, and average over the patients and different choices of subsets. This setup is performed for standard LDDMM with Gaussian kernels with scale ranging between 10 and 170 voxels and with LDDKBM with five scales in the same range. For both methods, we let $\lambda = 8$ for the energy weight in (1) and (6), and, for the kernel bundle, we let each scale contribute with equal weight. The theoretical influence of these weights remains to be explored, though, from a practical point of view, we have tried different values without significant effects on the results.

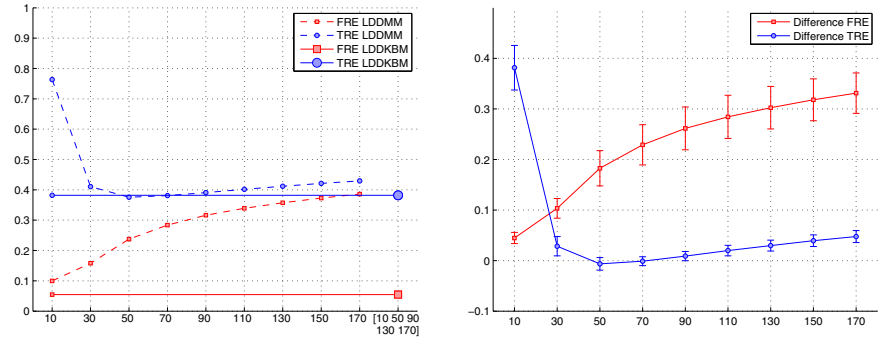

(a) LDDMM and LDDKBM relative errors (b) Error differences

Fig. 3. (a) Average relative FRE and TRE for different kernel scales (standard LDDMM) and LDDKBM (horizontal line and rightmost). Values of 0 indicate perfect matches and 1 indicates no reduction in error compared to before the match. Labels on the horizontal axis are kernel scale in voxels (9 different scales for standard LDDMM and the interval [10 170] discretized in five scales for LDDKBM). (b) FRE and TRE for standard LDDMM (scale in voxels on horizontal axis) subtracted the respective errors for LDDKBM with interval the [10 170] discretized in five scales. Positive values indicate superior performance of the proposed method. Error bars show standard deviation of the results.

As the results in Figure 3(a) shows, the FRE increases with scale for standard LDDMM. The LDDKBM method produces an FRE of around half the FRE for the smallest scale for standard LDDMM. Similarly, we see that the TRE decreases with increasing scale up to a scale of 50 units after which it starts increasing. This indicates that a kernel scale of 50 units will be appropriate for standard LDDMM. As displayed in Figure 3(b), the LDDKBM method attains an error well within one standard deviation of the best standard LDDMM result. Performing t-tests on the 5% significance level of the hypotheses that the methods have equal mean results in a failure to reject the hypotheses for the best standard LDDMM results, and thus we can only conclude a difference between the methods for the scales where LDDKBM clearly outperform standard LDDMM. The LDDKBM method therefore *automatically* adapts to the right scales. Furthermore, the results indicate that the same quality of match can be reached with less data since we potentially could use the entire dataset to drive the registration. Manual scale selection will allow only a part of the data as input for the registration as the rest is needed to select the kernel scale.

To show the decoupling between scales, Figure 4 displays the L_2-norm of w_t at $t = 0$ for three scales of LDDKBM and v_t at $t = 0$ with the best scale for standard LDDMM. The decoupling is apparent, and the expected sparsity is visible in particular for the smallest scales were movement is very localised.

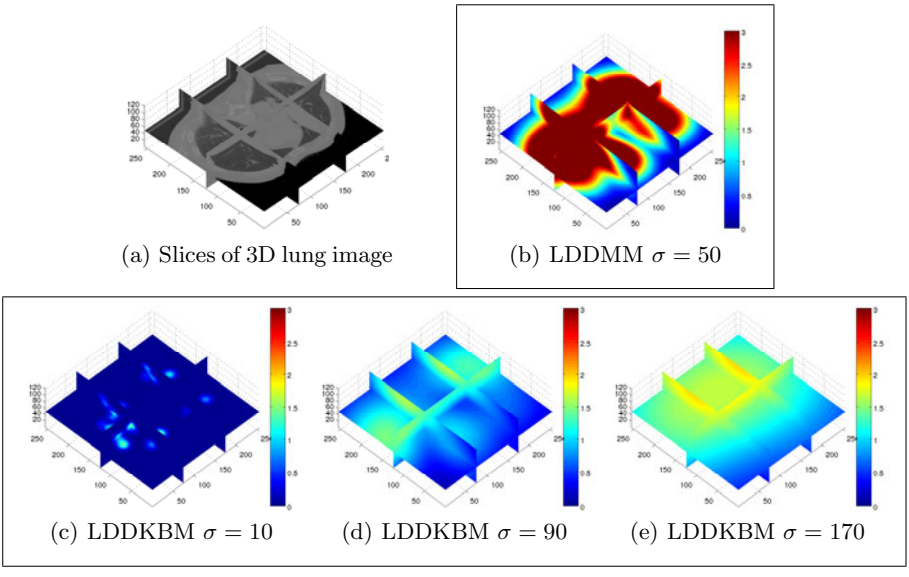

(a) Slices of 3D lung image (b) LDDMM $\sigma = 50$

(c) LDDKBM $\sigma = 10$ (d) LDDKBM $\sigma = 90$ (e) LDDKBM $\sigma = 170$

Fig. 4. Slices of 3D lung image and initial vector fields generating registration. Upper row, left to right: 3D image of lung to be registered, L^2-norm of vector field v_t generating φ at $t = 0$ for scale 50 (LDDMM). Lower row, left to right: L^2-norm of w_t generating φ at $t = 0$ for scales 10, 90, 170 (LDDKBM). Even for the best scale, standard LDDMM is not localized with high energy to account for the warp across scales, while LDDKBM use the appropriate scales localized and sparsely for the smallest scales.

5 Conclusion and Outlook

The multi-scale kernel bundle framework LDDKBM extends the LDDMM framework by incorporating multiple kernels at multiple scales in the registration. The method decouples scales and thereby allows the registration to act at the right scales only. As shown in the experiments, the method removes the need for classical scale selection by automatically, without user intervention or cross-validation, adapting to the right scales, and it thereby provides as good results as standard LDDMM optimally tuned with respect to scale. The decoupling of scales and the ability of the method to adapt at different scales are visualized and LDDKBM promises to allow for sparse description of deformation and momentum based statistics incorporating scale information.

The method opens up many directions for future work. Extending existing LDDMM algorithms to LDDKBM with non-landmark data is clearly important, and using the smoothness properties of different scales to speed up such algorithms will be beneficial. We expect to make use of the improved representation to compute sparse systems which will further improve computation efficiency and statistical results. Enforcing sparsity through sparse priors will be an important step for this.

References

1. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006)
2. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. IJCV 61(2), 139–157 (2005)
3. Castillo, R., Castillo, E., Guerra, R., Johnson, V.E., McPhail, T., Garg, A.K., Guerrero, T.: A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. Physics in Medicine and Biology 54(7), 1849–1870 (2009)
4. Christensen, G., Rabbitt, R., Miller, M.: Deformable templates using large deformation kinematics. IEEE Transactions on Image Processing 5(10) (2002)
5. Cotter, C.J., Holm, D.D.: Singular solutions, momentum maps and computational anatomy. nlin/0605020 (May 2006)
6. Dupuis, P., Grenander, U., Miller, M.I.: Variational problems on flows of diffeomorphisms for image matching (1998)
7. Grenander, U.: General Pattern Theory: A Mathematical Study of Regular Structures. Oxford University Press, USA (1994)
8. Hernandez, M., Bossa, M., Olmos, S.: Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. International Journal of Computer Vision 85(3), 291–306 (2009)
9. Joshi, S., Davis, B., Jomier, B.M., Guido Gerig, B.: Unbiased diffeomorphic atlas construction for computational anatomy. NeuroImage 23, 151–160 (2004)
10. Micheli, M.: The differential geometry of landmark shape manifolds: metrics, geodesics, and curvature. Ph.D. thesis, Brown University, Providence, USA (2008)
11. Pennec, X., Stefanescu, R., Arsigny, V., Fillard, P., Ayache, N.: Riemannian elasticity: A statistical regularization framework for non-linear registration. In: Duncan, J.S., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3750, pp. 943–950. Springer, Heidelberg (2005)
12. Risser, L., Vialard, F.-X., Wolz, R., Holm, D.D., Rueckert, D.: Simultaneous fine and coarse diffeomorphic registration: Application to atrophy measurement in alzheimer's disease. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6362, pp. 610–617. Springer, Heidelberg (2010), PMID: 20879366
13. Thirion, J.: Image matching as a diffusion process: an analogy with maxwell's demons. Medical Image Analysis 2(3), 243–260 (1998)
14. Trouvé, A.: An infinite dimensional group approach for physics based models in patterns recognition (1995)
15. Vaillant, M., Miller, M., Younes, L., Trouvé, A.: Statistics on diffeomorphisms via tangent space representations. NeuroImage 23(suppl. 1), S161–S169 (2004)
16. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: efficient non-parametric image registration. NeuroImage 45(suppl. 1), 61–72 (2009)
17. Younes, L.: Shapes and Diffeomorphisms. Springer, Heidelberg (2010)
18. Younes, L., Arrate, F., Miller, M.I.: Evolutions equations in computational anatomy. NeuroImage 45(1, suppl. 1), S40–S50 (2009)