A Unified and Efficient Approach for Free-form Deformable Registration

Ali Khamene, Fred Azar
Imaging and Visualization Dept.
Siemens Corporate Research
755 College Rd East
Princeton, NJ 08540 USA
<firstname>.<lastname>@siemens.com

Loren Schwarz, Darko Zikic, Nassir Navab
Chair for Computer Aided Medical Procedures
Technical University of Munich
Boltzmannstr. 3
85748 Garching bei Munich Germany
<lastname>@in.tum.de

Eike Rietzel
Siemens Medical Solutions
Particle Therapy
Hofmannstr. 26
91050 Erlangen Germany
eike.rietzel@siemens.com

Abstract

We propose a novel numerical approach for solving the free-form deformable registration problem. The central idea is to utilize the well understood techniques from variational deformable registration problems. We demonstrate that it is possible to formulate the free-form deformable registration problem as the optimization of an energy functional as in the dense deformation case. This energy functional possesses image distance and regularization terms, which are both functions of the free-form deformation control points. We then setup a semi-backward (implicit) partial differential equation that optimizes the established energy functional. In addition to being mathematically justified, this approach provides both accuracy and speed. Our evaluation on synthetic, real, two dimensional, and three dimensional data demonstrates accuracy and computational effectiveness.

1. Introduction and Background

Establishing dense correspondence or mapping between two images is a fundamental problem in both computer vision and medical imaging analysis. Numerous methods have been proposed that are primarily based on pioneering works in [2, 6]. Motion analysis through estimating the optical flow has been the primary application for such techniques in the computer vision realm, whereas image registration has been the driving application from the medical image analysis side. The underlying approach has always been similar: minimization of a two-term energy functional balancing off an image distance metric with some constraint on the underlying correspondence map. Calculus of variations enables the derivation of an iterative solution scheme for the problem of minimizing the energy functional. The condition of minimal energy is reached in a state when two conceptual forces acting against each other are in an equilibrium. One of these forces pulls points in one of the images towards a position that decreases the distance between the two images. The second force, determined by the regularizer, can be thought of as a constraint on the pixel movements which counteracts the effects of the former force. Image distance metrics have evolved from linear optic flow constraints to non-linear sum of squared differences on image intensities or gradients, and finally to information theoretic distances/similarities over the years [6, 1, 11]. For motion analysis applications, smoothness constraints are added primarily to resolve the ill-posedness of the problem. This is despite the fact that discontinuities in the motion field, due to occlusions must be allowed [8, 9]. On the other hand in medical image registration techniques, deformation fields should be free of discontinuities such as folding or tearing. These so-called diffeomorphic properties can be enforced by incorporating regularization terms inspired from the material properties of the deformable objects [2, 3]. Almost all such approaches boil down to solving a non-linear partial differential equation in an iterative setting [1, 3].

The computational complexity is mainly due to the underlying large dimension of the sought-after deformation or flow. In [5], authors have proposed a more computationally

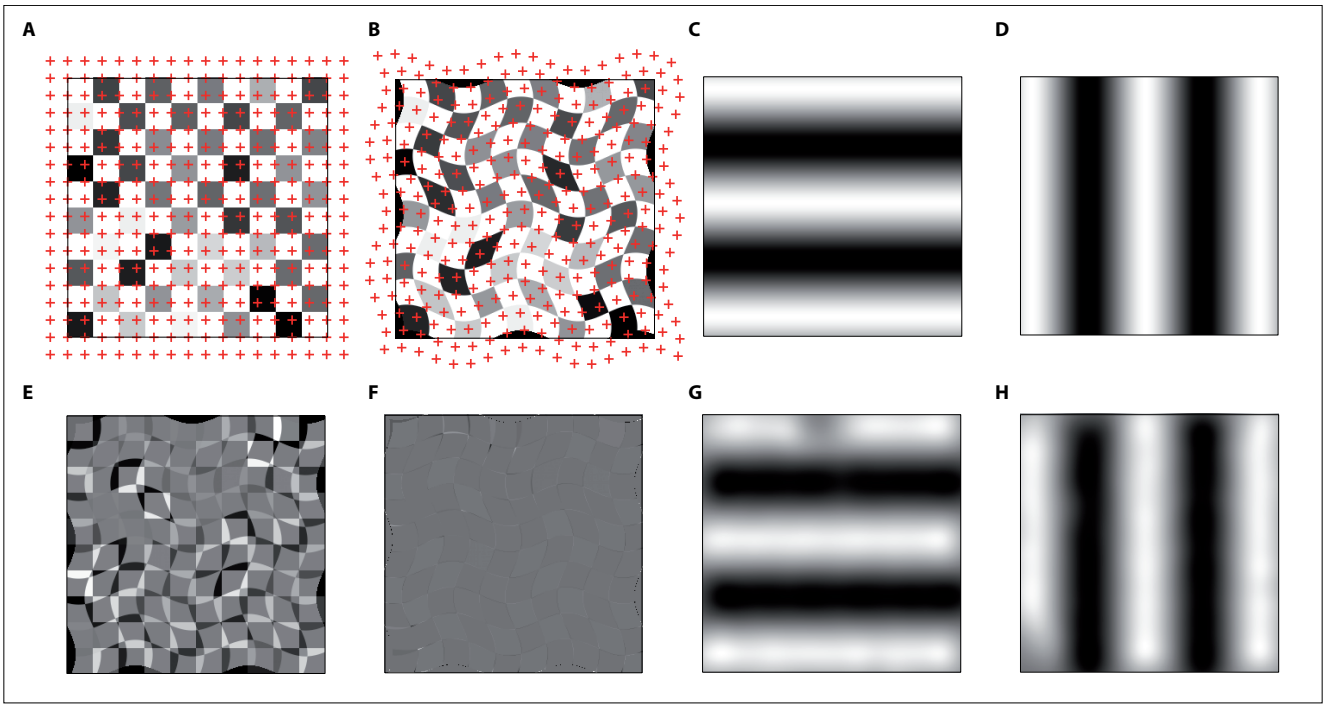

Figure 1. Evaluation of registration quality for deformable registration based on free-form deformations. Starting from a fixed control point grid (A), a known control point configuration (B) is used to generate a ground-truth displacement field. Its x and y components are shown in (C) and (D). A synthetic checker board image is deformed and treated as an artificial reference for registration. The template is given by the undeformed image. (E) and (F) are the difference images before and after registration, (G) and (H) are the components of the reconstructed displacement field after registration. A comparison with the ground truth displacement field helps to asses registration quality. Detailed results are given in section 3.1.

efficient method based on the multi-grid approach. On the other hand, parameterization of the deformation field has the potential to reduce the computational complexity and increase the practicality of the technique. Free Form Deformation (FFD) based approaches are popular mainly due to the decreased number of the unknown parameters, and to the inherent smoothness caused by the interpolating kernel (e.g., B-spline functions) [11, 12]. The general idea of free-form deformations is to warp an image by manipulating a regular grid of control points distributed across the image at an arbitrary mesh resolution. Smoothness is introduced since the displacement of individual pixels between the control points is interpolated from the displacement of the surrounding control points. The behavior of this interpolation step is determined by the degree of the B-splines. Displacement fields based on cubic B-splines are continuous up to the second derivative yielding smoothness, which only depends on the chosen control point grid resolution. Linear B-splines, which are in fact simply a linear interpolation between control points, can also be employed although their use is rarely documented in the literature. In both [11, 12], the solution is formulated as the result of the gradient descent or Levenberg-Marquardt optimization with a forward Euler update formula, where the spatial dependency of control points is ignored. Furthermore, in both approaches computations are still performed on a per-pixel basis. In this paper, we make a complete analogy between the variational solution strategy for dense deformation fields and the FFD setting with control point grids. All concepts from the variational solution strategy, such as the two antagonistic forces are transferred from the level of individual image pixels to the control points. Since the number of control points is typically significantly smaller than the number of image pixels, there is a potential performance advantage. In particular, we show that regularization performed on the control points is completely equivalent to the traditional way of applying a regularizer to the dense deformation field. We provide a numerically efficient solution for the diffusion FFD registration problem.

2. General Problem Formulation

Given the fixed and moving images I_f and I_m, both defined over a domain $\Omega \subset \mathbb{R}^n$ (n is the image dimension), the energy functional of the minimization problem for recovering the dense correspondence map $u : \mathbb{R}^n \mapsto \mathbb{R}^n$ can be formulated as $\mathcal{E}(\mathbf{u}) = \mathcal{M}(\mathbf{u}) + \lambda\mathcal{R}(\mathbf{u})$ The first term denotes the image metric and the second term is the reg-

ularization. In the case, where we use the sum of squared differences (SSD) and the diffusion regularization as the image metric and regularization, the energy can be computed as:

$$\int_{\Omega}(I_f(\mathbf{x})-I_m(\mathbf{x}+\mathbf{u}))^2dx+\lambda\int_{\Omega}trace(\nabla\mathbf{u}^{\top}\nabla\mathbf{u})dx \tag{1}$$

where $\top$ denotes transpose operation. The Euler Lagrange equation for dimension $d\in\{x,y,z\}$ associated with energy at 1 is then given by:

$$(I_m(\mathbf{x}+\mathbf{u})-I_f(\mathbf{x}))\nabla_d I_m(\mathbf{x}+\mathbf{u})+\lambda\Delta u_d=0 \tag{2}$$

The nonlinear parabolic partial differential equation (PDE) in 2 can be solved using a simple forward Euler update formula:

$$u_d^{t+1}=u_d^t+\delta t(f_d(\mathbf{u}^t)+\lambda\Delta u_d^t) \tag{3}$$

where δt is the time step, t is the time index, and $f_d(\mathbf{u})$ denotes with the first term in equation 2. To ensure stability the δt should be chosen small, which in turn may have an adverse effect on the convergence speed. In the discrete case where images are of size $M\times N$, we can represent the Laplacian operator Δ as a square circulant matrix operator $\mathbf{R}$ with size $MN\times MN$. Therefore, a semi-backward Euler update formula can be written as follows:

$$\bar{\mathbf{u}}_d^{t+1}=(\mathbf{I}+\lambda\delta t\mathbf{R})^{-1}(\bar{\mathbf{u}}_d^t+\delta t\bar{\mathbf{f}}_d(\mathbf{u}^t)); \tag{4}$$

where $\mathbf{I}$ is the identity matrix with same size as $\mathbf{R}$. $\bar{\mathbf{u}}_d$ and $\bar{\mathbf{f}}_d(\mathbf{u})$ are discrete version of u_d and $f_d(\mathbf{u})$ re-arranged into a vector. The update equation 4 presents the advantages of greater stability and faster convergence [4]. The downside is that it involves inverting a large matrix. Depending on the matrix form, accelerated inversion approaches such as the Fourier or multi-grid method may be used [5]. It should also be noted that, in order to cope with large displacements, equation 4 is solved in a scale-space setting. Here the equation must be solved over a number of resolutions starting with zero displacements as the initial value at the lowest resolution. Using this approach, the initial values at each resolution are computed based on the prior resolution estimates. This is an attempt to ensure covergence to a better local optima over a larger range of parameters of the energy functional .

2.1. Free-form Deformations

In the case of free form deformations, the underlying deformation field is parameterized based on a set of control points placed at the nodes of a discrete grid over the image domain. Without loss of generality, we consider the three dimensional case, where a set of $N=n_xn_yn_z$ control points $\phi=[\varphi_x,\varphi_y,\varphi_z]^{\top}$ with a spacing of s_x,s_y,s_z is introduced. The deformation field $\mathbf{u}(\phi,x)$ uses a set of B-spline basis functions to interpolate displacements on individual points based on that of the control points. Let p_x,p_y,p_z be the indices of the control point cell that contains a particular image point $\mathbf{x}=(x,y,z)^{\top}$, e.g. $p_x=\lfloor x/s_x\rfloor$. The cubic B-spline transformation is then:

$$\mathbf{u}(\phi,\mathbf{x})=\sum_{a=0}^{3}\sum_{b=0}^{3}\sum_{c=0}^{3}B_3^a(\alpha)B_3^b(\beta)B_3^c(\gamma)\,\phi(i,j,k), \tag{5}$$

where $(i,j,k)=(p_x+a-1,p_y+b-1,p_z+c-1)$ are the control point indices. The parameters α,β,γ are the fractional remainders of a pixel between its surrounding control points, for instance $\alpha=x/s_x-p_x$. The basis functions $B_3^k(\alpha)$ are the cubic B-splines [11]. If linear B-splines $B_1^k(\alpha)$ are used instead, the transformation simplifies to a linear interpolation between the eight control points surrounding an image point. Two types of spatial neighborhoods can be considered in free-form deformations for discretized images. The neighborhood Φ of a particular image point $\mathbf{x}$ contains all control points used for the computation of the deformation at that point. This neighborhood comprises a block of $2\times2\times2$ control points for linear B-splines and $4\times4\times4$ control points for cubic B-splines. The neighborhood Ψ consists of all image points affected by a particular control point ϕ_i. For instance with linear B-splines, the control point neighborhood contains $2s_x\times2s_y\times2s_z$ image points, whereas for cubic B-splines it contains $4s_x\times4s_y\times4s_z$ image points.

2.2. Free-form Registration

As described in section 2, the registration objective is posed as the problem of finding the optimal deformation that maps the moving (template) to the fixed (reference) image. Since in the context of free-form deformations the deformation is parameterized by the control points, this aim can be equivalently formulated as that of finding the optimal control point configuration.

2.2.1 Computation of External Image Force

Once we substitute the $\mathbf{u}$ in equation 1 with the one from equation 5, the dissimilarity term $\mathcal{M}(\phi)$ can be computed as follows:

$$\mathcal{M}(\phi)=\int_{\Omega}\left(I_f(\mathbf{x})-I_m(\mathbf{x}+\mathbf{u}(\phi,\mathbf{x}))\right)^2d\mathbf{x}. \tag{6}$$

The force field is computed as the partial derivative of the dissimilarity term with respect to the d component of the

control point (i, j, k)

$$f_d^{\phi}(\phi(i,j,k)) = \frac{\partial \mathcal{M}(\phi)}{\partial \varphi_d(i,j,k)} = \int_{\Psi} f_d(\mathbf{x}) \cdot \frac{\partial}{\partial \varphi_d(i,j,k)} u_d(\phi, \mathbf{x}) d\mathbf{x}. \quad (7)$$

where $f_d(\mathbf{x})$ is:

$$(I_m(\mathbf{x} + \mathbf{u}(\phi, \mathbf{x})) - I_f(\mathbf{x})) \cdot \frac{\partial}{\partial d} I_m(\mathbf{x} + \mathbf{u}(\phi, \mathbf{x})). \quad (8)$$

In equation (7), the partial derivative of the transformation function with respect to the d component of the control point ϕ_i, is non-zero for image points within the neighborhood of Ψ as defined in the previous section. Furthermore, its value is independent of the control point position. This allows the force term to be written as a convolution of the force at the image level with a kernel $\mathcal{H}(\mathbf{x})$. The kernel turns out to be a smoothing one with a compact support, built upon the basis functions. In the case of cubic B-splines the kernel has the following form:

$$\mathcal{H}(\mathbf{x}) = \begin{cases} B_3^a(\alpha) B_3^b(\beta) B_3^c(\gamma) & 0 \leq x_d \leq 4 s_d \\ 0 & \text{elsewhere} \end{cases} \quad (9)$$

where $a = |i + 1 - \lfloor x/s_x \rfloor|$, $b = |j + 1 - \lfloor y/s_y \rfloor|$, $c = |k + 1 - \lfloor z/s_z \rfloor|$ and and α, β, γ are defined in the previous sections. Therefore, the external image force acting on the control point (i, j, k) can be simply computed by means of a convolution and a downsampling step as follows:

$$f_d^{\phi}(\phi(i,j,k)) = (f_d * \mathcal{H})(\mathbf{x}) \mid_{\mathbf{x}=\phi(i,j,k)} \quad (10)$$

where $*$ denotes the convolution operation.

2.2.2 Computation of Regularization Force

Similar to the image dissimilarity term, the regularization term $\mathcal{R}(\phi)$ may be computed by substituting the $\mathbf{u}$ in equation 1 with the one from equation 5:

$$\mathcal{R}(\varphi) = \int_{\Omega} \sum_d \left(|\frac{\partial}{\partial_x} u_d(\mathbf{x}, \phi)|^2 + |\frac{\partial}{\partial_y} u_d(\mathbf{x}, \phi)|^2 + |\frac{\partial}{\partial_z} u_d(\mathbf{x}, \phi)|^2 \right) d\mathbf{x}. \quad (11)$$

Assuming a set of cubic B-spline basis functions, using the Cauchy-Schwarz integral inequality, we can derive the upper bound of the regularization. For instance, the first term corresponding to the x dimension (i.e., $d = x$) from equation 11 can be written as follows:

$$\sum_{i,j,k} \int_U \left| \sum_{a,b,c} \frac{dB_3^a(x)}{dx} B_3^b(y) B_3^c(z) \varphi_x(i^*, j^*, k^*) \right|^2 d\mathbf{x} \leq \beta \sum_{i,j,k} \left| \sum_{a,b,c} \varphi_x(i^*, j^*, k^*) \int_U \frac{dB_3^a(x)}{dx} B_3^b(y) B_3^c(z) d\mathbf{x} \right|^2 = \gamma \sum_{i,j,k} |\mathcal{D}_x \circ \varphi_x(i,j,k)|^2 \quad (12)$$

where U denotes a unit cube and $i^* = i + a - 1$ (similarly for j^* and k^*). Furthermore, β and γ are real positive numbers. In the case of cubic B-splines, the operator $\mathcal{D}_x$ resembles a discrete gradient operator in the x directio. The elements of the operator are as follows:

$$\mathcal{D}_x(i,j,k) = (B_3^i(1) - B_3^i(0)) \int_0^1 B_3^j(y) dy \int_0^1 B_3^k(z) dz \quad \text{for} \quad 0 \leq i, j, k \leq 3 \quad (13)$$

By performing the same operations as in equation 12 for $d = y, z$ and then compounding them together, we arrive at the upper bound for the regularization:

$$\mathcal{R}(\varphi) = \sum_{i,j,k} \sum_d ||\nabla \mathcal{D} \varphi_d(i,j,k))||^2. \quad (14)$$

where $\nabla \mathcal{D} = (\mathcal{D}_x, \mathcal{D}_y, \mathcal{D}_z)$. Equation 14 demonstrates a direct link between the regularization on the dense deformation field and the control point movements.

2.2.3 Numerical Solution

The Euler-Lagrange equation associated with the energy functional of $\mathcal{E}(\phi) = \mathcal{M}(\phi) + \lambda \mathcal{R}(\phi)$ can be formulated as follows:

$$f_d^{\phi}(\phi) + \lambda \Delta_{\mathcal{D}} \varphi_d = 0 \quad (15)$$

where d is the vector component and $\Delta_{\mathcal{D}} = \nabla \mathcal{D} \cdot \nabla \mathcal{D}^{\top}$ denotes a discrete laplacian like operator built based on the interpolating kernel.

Similar to the equation 3, we can re-arrange f_d^{ϕ} and φ_x into a $N \times 1$ vector as $\bar{\mathbf{f}}_d^{\phi}$ and $\bar{\varphi}_x$, respectively. Furthermore, the $\Delta_{\mathcal{D}}$ can be represented by a square circulant matrix of $N \times N$ dimension, called $\mathbf{D}$. Therefore, a semi-backward Euler update formula can be written as follows:

$$\bar{\varphi}_d^{t+1} = (\mathbf{I} + \lambda \delta t \mathbf{D})^{-1} (\bar{\varphi}_d^t + \delta t \bar{\mathbf{f}}_d^{\phi}(\phi)); \quad (16)$$

where $\mathbf{I}$ is the identity matrix same size as $\mathbf{D}$. The update equation 16 has the advantage of having the smaller dimension comparing to the one in equation 4. Similar to

the dense case, a multi resolution scheme must be used in order to converge to a better solution.

3. Evaluation

A number of experiments have been conducted in order to evaluate the proposed registration algorithm. Measurements on synthetic data sets are used to demonstrate the effectiveness and performance of the algorithm. Medical data sets combined with ground-truth data illustrate the applicability of the registration method to typical problems in the clinical setting.

3.1. Synthetic Data

The experiments on synthetic data are performed in the two-dimensional case. A synthetic image is deformed using a known control point configuration. The resulting deformed image is treated as the reference for a series of measurements while the template image is given by the original, undeformed image. Registration accuracy can now be assessed by comparing the displacement field obtained by the registration algorithm with the ground truth displacement computed from the known control point configuration. The synthetic image that is used for the experiments, shown in Figure 1, is a checker board of size 300×300 pixels with random intensities. A control point configuration that is sinusoidal in the x and y directions and that is based on a fixed spacing is used to generate the ground truth displacement field and the artificial reference image. Registration is performed for various initial control point spacings ranging from 5 to 50 pixels at increments of 5. All measurements are repeated for linear and cubic B-splines using otherwise identical parameters, allowing to make several observations.

3.1.1 Similarity after Registration

Registration based on both linear and cubic B-splines can yield comparably low SSD dissimilarity values after registration for moderate control point spacings. As can be seen in the left graph of Figure 2, control point spacings between 5 and 20 pixels give similar results for both types of B-splines. For larger control point spacings, cubic B-splines show better tolerance and only start to produce significantly worse registration results at spacings around 35 pixels. This aspect can be explained by the smoothness properties of displacement fields generated using cubic B-splines. However, for practical applications linear B-splines are attractive nonetheless, since typically relatively small control point spacings are used in order to capture small image details.

3.1.2 Processing Time

The main advantage of linear B-splines is their lower computational complexity that can significantly increase registration speed. Using the efficient B-spline coefficient precomputation technique proposed in [10], the difference in registration time between linear and cubic B-splines is reduced to an additive constant independent of the control point spacing. The right hand side of Figure 2 shows that in the 2D case this constant is around 3 seconds. Total registration durations for linear B-splines and the data described above are between 4 and 7 seconds. Shorter times are achieved on coarser control point grids. It can also be deduced from the figure that control point spacings below 15 pixels yield a disproportionate increase in registration time. In 3D, a typical registration run using linear B-splines and images of size $256 \times 256 \times 142$ voxels takes 70 to 150 seconds, depending on various parameter settings. Registration with identical settings and cubic B-splines takes roughly 140 seconds longer. All given durations refer to experiments performed on a 2.4 GHz Intel Core 2 Duo system equipped with 2 GB of memory.

3.1.3 Deformation Field Reconstruction

Several measures have been proposed in the literature to compare displacement fields. Two such measures that originate from the context of optical flow reconstruction [7] are the angular error and the magnitude of vector difference. Both being simple and intuitive, these two measures allow to quantify how well the ground truth displacement field is reconstructed using the registration algorithm. While the angular error measures the directional difference for corresponding vectors in the correct and the reconstructed displacement field, the second measure quantifies the difference in vector length. Evaluated for all positions in the dense deformation fields obtained for the aforementioned registration series, several statistics can be collected, as illustrated in Figure 3. Best results in the sense of both measures are achieved using cubic B-splines and a control point spacing between 15 and 35 pixels. The magnitude of vector difference is for these parameters less than one pixel on average and the mean angular error is around 4 degrees. Given the fact that the synthetic images used for experiments contain regions of constant intensities, offering less ”grip” for the registration algorithm, these values are competitive.

3.2. Medical Data

In order to evaluate the performance of the registration algorithm in the 3D case and its applicability to real patient data, a set of validation measurements is performed as follows. A data set of 11 computer tomography (CT) thorax scans of one patient is used where each scan is acquired at a different breathing stage. Pairs of volumes from the data

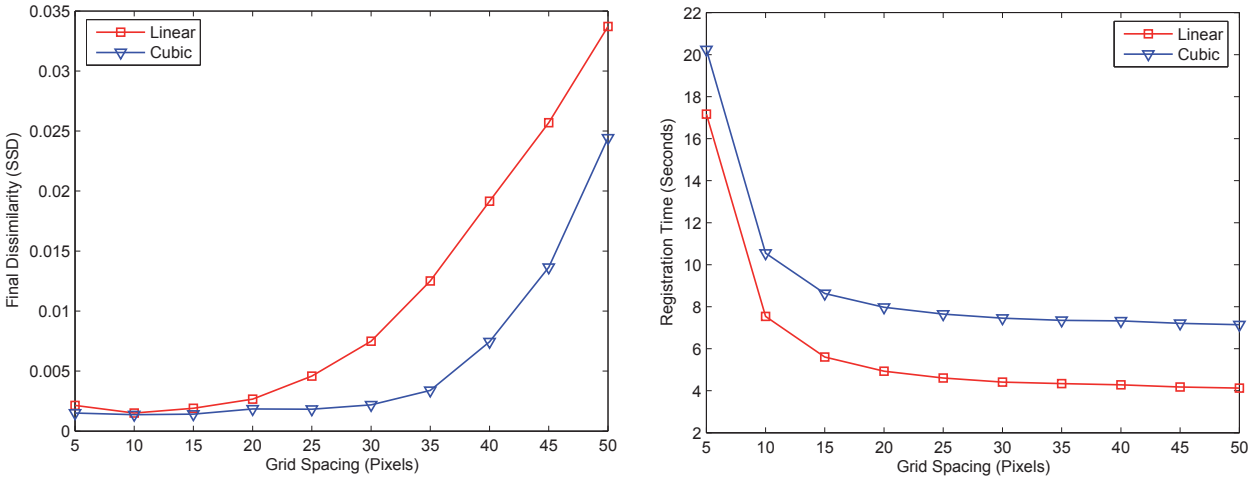

Figure 2. Evaluation of registration performance for the 2D ground-truth experiment series. SSD dissimilarity after registration (left) and computation time (right) for linear and cubic B-splines and for different control point grid resolutions are shown. Images are of size 300×300 pixels. The same, sufficient number of iterations (200) is performed in all measurements for comparability. Linear B-splines offer better computational efficiency and yield final dissimilarity values that are comparable to those achieved using cubic B-splines for moderate control point spacings up to 20 pixels.

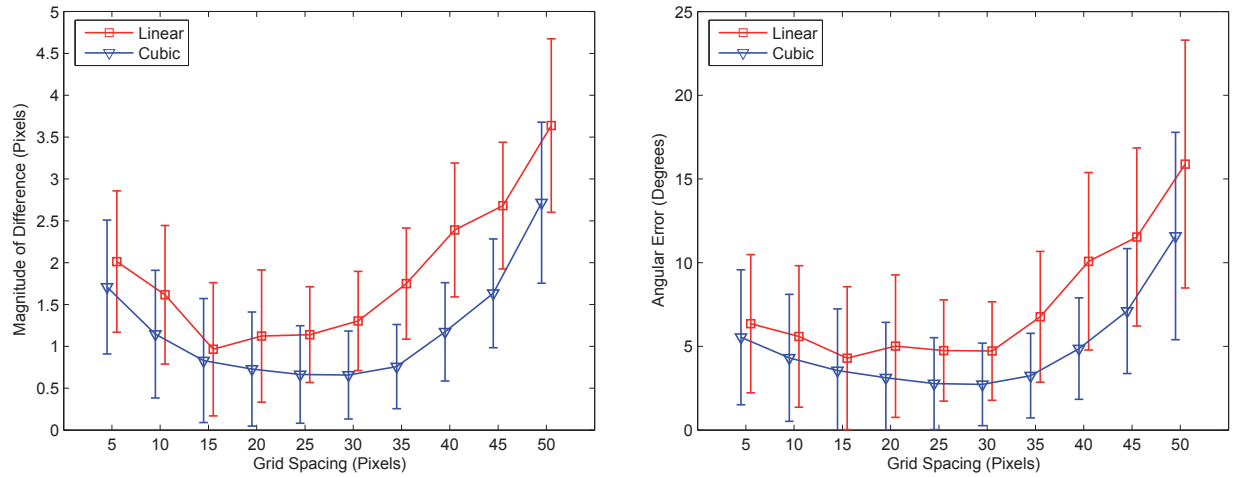

Figure 3. Assessment of 2D registration precision based on a comparison of reconstructed dense displacement fields with ground truth displacements. The magnitude of vector difference (left) and angular error (right) metrics are evaluated for all corresponding positions in the displacement fields. Cubic B-splines can generally account for marginally better results than linear B-splines. Best values are achieved for control point spacings between 15 and 35 pixels with average difference vector magnitudes of less than one pixel and average angular errors around 4 degrees. The data points in the graphs represent the respective mean values and the standard deviations for all experiments are indicated as vertical bars.

set are registered, resulting in a total of 55 registrations. A qualitative evaluation can be given based on the difference images before and after registration, a few sample slices of which are shown in Figure 4. Since a more quantitative assessment of registration success is desirable, ground truth data is generated. A tumor on the left side of the patient's chest is manually segmented in all CT scans, resulting in 11 binary segmentations. The dense deformation fields from the registrations that link the images from different breathing stages are then applied to the appropriate segmentation masks. For instance, the deformation field that maps image 1 to image 6 is applied to the segmentation corresponding to image 1. The resulting deformed segmentation is compared to the ground truth, the manual segmentation for image 6. In an ideal case these two segmentations overlap perfectly. To measure the degree of overlap that is achieved, sensi-

tivity and specificity are evaluated on a per-voxel basis and statistics over all 55 correspondences in the registration series are collected. Sensitivity is in this context defined for a pair of segmentation masks and gives the fraction of voxels in the ground-truth tumor region that are correctly matched in the reconstructed segmentation. A perfect overlap results in a sensitivity value of 1. Specificity gives the fraction of correctly identified non-tumor voxels. Since on average the tumor regions cover only 0.04% of the whole volume, specificity is practically 1 in all experiments. The sensitivity mean is at 0.879 with a standard deviation of 0.045 and a median at 0.894. Minimal and maximal sensitivity values among the data set are 0.769 and 0.937, respectively.

4. Summary and Conclusion

In this paper, we described a numerically efficient technique for the free-form deformable registration problem. We demonstrated the link between the formulations of the deformable registration technique in the variational and the free-form settings. We specifically showed that both the image dissimilarity and regularization in dense case translate directly to the free-form case. That led to the translation of the approaches already investigated for the variational dense deformation case. The proposed approach has the advantage of being numerically efficient for two reasons. The first is that the core iterative scheme is based on semi-backward Euler, and second the update formulation is done at the control point level, with a fraction of the number compared to that of image points. The approach is tested on set of synthetic and real data sets and the results are presented showing the effectiveness of the approach.

References

[1] L. Alvarez, J. Weickert, and J. Sanchez. Reliable estimation of dense optical flow fields with large displacements. *International Journal of Computer Vision*, 12(1):234–778, 2002.

[2] R. Baicsy and C. Broit. Matching of deformed images. pages 351–353, 1982.

[3] G. E. Christensen, R. D. Rabbit, and M. I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(10):1435–1447, 1996.

[4] C. Gonzales, A. Ostermann, C. Palencia, and M. Thalhammer. Backward euler discretization of fully nonlinear parabolic problems. *Mathematics of Computation*, 71(237):125–145, 2001.

[5] S. Henn and K. Witsch. Iterative multigrid regularization techniques for image matching. *SIAM J. Sci. Comput.*, 23(4):1077–1093, 2001.

[6] B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1):185–203, 1981.

[7] B. McCane, K. Novins, D. Crannitch, and B. Galvin. On benchmarking optical flow. *Computer Vision and Image Understanding*, 84:126–143, 2001.

[8] H.-H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.

[9] P. Nesi. Variational approach to optical flow estimation managing discontinuities. *Image and Vision Computing*, 11(7):419–439, 1993.

[10] T. Rohlfing and C. R. Maurer. Nonrigid image registration in shared-memory multiprocessor environments with application to brains, breasts and bees. *IEEE Transactions on Information Technology in Biomedicine*, 7/1:16–25, 2003.

[11] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes. Non-rigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.

[12] R. Szeliski and J. Coughlan. Hierarchical spline-based image registration. page 194201, 1994.

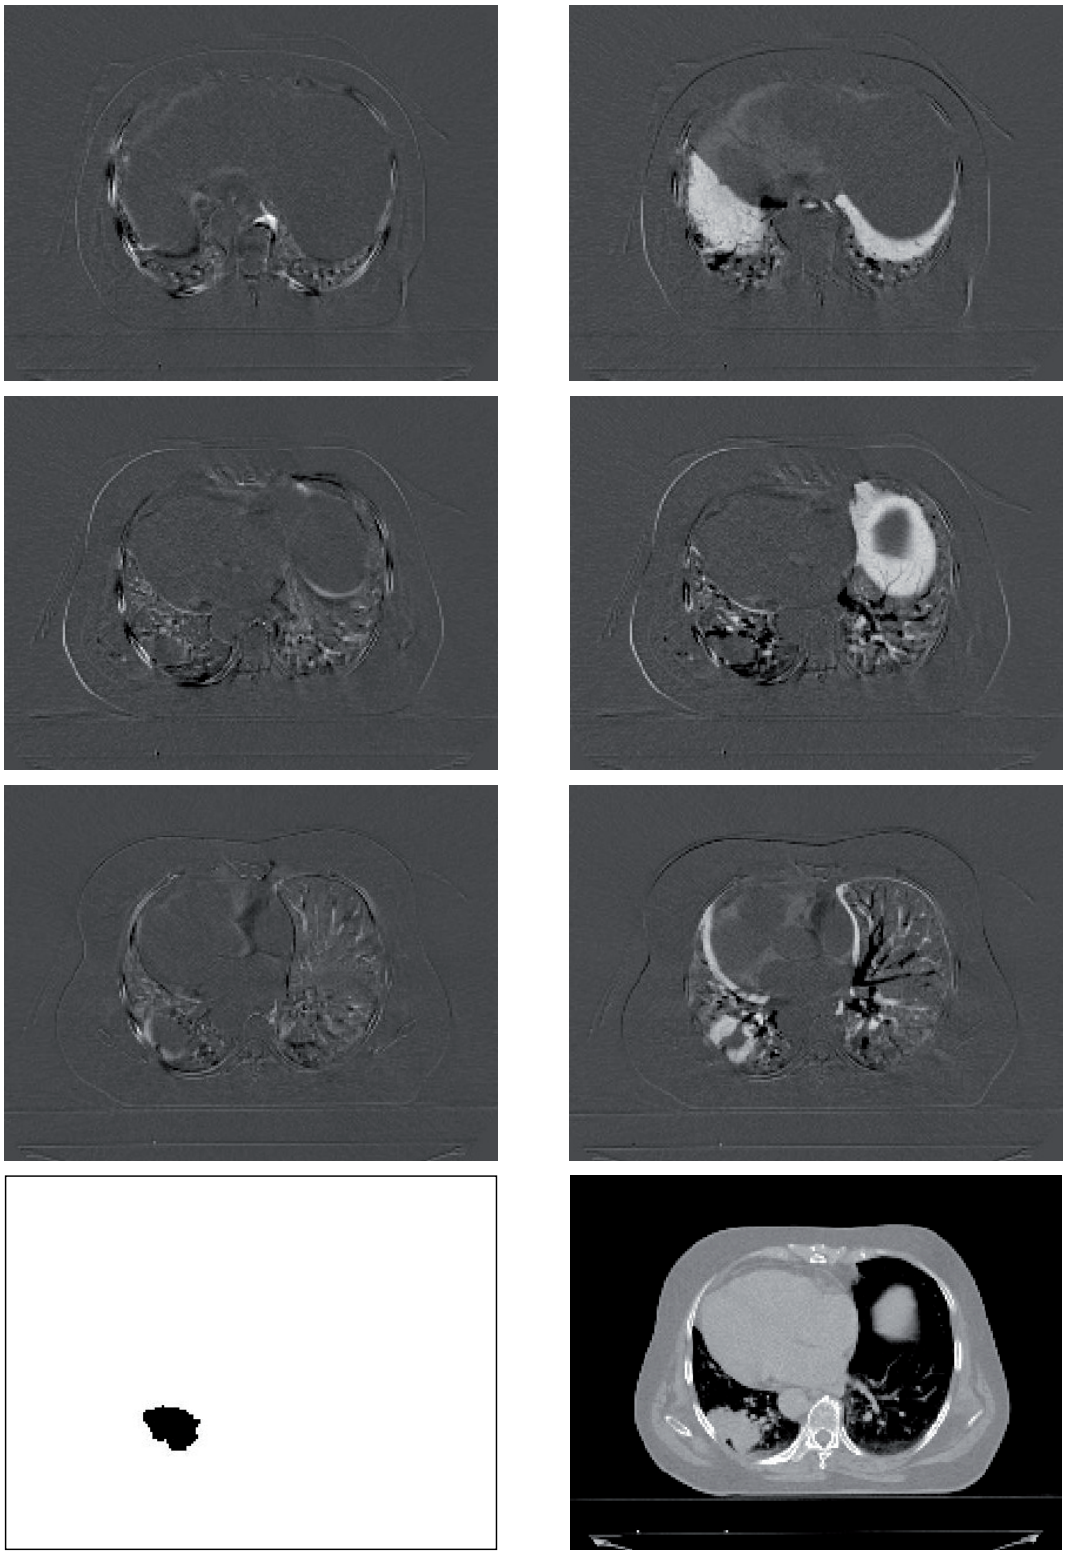

Figure 4. Sample slice from CT thorax data set and corresponding tumor segmentation (last row). The data set consists of 11 3D images of the same patient acquired at different breathing stages. Difference slice images after (left) and before (right) registering two of the volumes across breathing stages are shown in the remaining rows. The white areas in the right images are due to vertical movement of organs (in the direction of z in the images) caused by breathing. These movements are completely compensated by registration.