

Numerical methods for volume preserving image registration

Eldad Haber and Jan Modersitzki

Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA

E-mail: haber@mathcs.emory.edu and modersit@mathcs.emory.edu

Received 18 February 2004, in final form 9 June 2004

Published 6 August 2004

Online at stacks.iop.org/IP/20/1621

doi:10.1088/0266-5611/20/5/018

Abstract

Image registration is one of today's challenging image processing problems, particularly in medical imaging. Since the problem is ill posed, one may like to add additional information about distortions. This applies, for example, to the registration of time series of contrast-enhanced images, where variations of substructures are not related to patient motion but to contrast uptake. Here, one may only be interested in registrations which do not alter the volume of any substructure. In this paper, we discuss image registration techniques with a focus on volume preserving constraints. These constraints can reduce the non-uniqueness of the registration problem significantly. Our implementation is based on a constrained optimization formulation. Upon discretization, we obtain a large, discrete, highly nonlinear optimization problem and the necessary conditions for the solution form a discretized nonlinear partial differential equation. To solve the problem we use a variant of the sequential quadratic programming method. Finally, we present results on synthetic as well as on real-life data.

1. Introduction

Image registration is one of the fundamental tasks in today's image processing and in particular in medical imaging; see, e.g., [15] and references therein. Given are a reference image R and a template image T , which are taken at different times, from different perspectives, and/or from different devices. Loosely speaking, the objective of image registration is to find a 'reasonable' deformation such that the 'distance' between R and a deformed version of T becomes small.

An application of particular clinical interest is the registration of pairs of images acquired before and after contrast administration; see, e.g., [32, 33] and references therein. A typical example is depicted in figure 1. In this application, magnetic resonance images of a female breast are taken at different times (images from Bruce Daniel, Lucas Center for Magnetic




Figure 1. MRIs of a female breast, left: during the wash-in phase, middle: during the wash-out phase, and right: difference image.




Figure 2. Synthetic images simulating contrast uptake, left: wash-in, right: wash-out.

Resonance Spectroscopy and Imaging, Stanford University). The left image shows an MRI section taken during the so-called wash-in phase of a marker and the middle image shows the analogous section during the so-called wash-out phase. A comparison of these two images indicates a suspicious region in the upper part of the images. This region can be detected easily since the images have been registered linearly: tissue located at a certain position in the wash-in image is almost related to the tissue at the same position in the wash-out phase. Generally, however, a quantitative analysis is a delicate matter since observable differences are not only related to contrast uptake but also due to motion of the patient, such as, for example, breathing or heart beat.

As pointed out by Rohlving *et al* [32], there is a substantial difficulty with the registration of pre- and post-contrast images. Bright regions seem to enlarge during the so-called wash-in phase. This enhancement is due to contrast uptake but not to movement of the patient. Figure 2 illustrates an ideal situation. Without external information, it is impossible to answer whether the white area has been enlarged or the grey area turned to white.

The idea is to restrict the set of feasible transformations *a priori* in a reasonable way. For this situation, we constrain the transformations to be volume preserving (VP). In contrast to [32], we use a variational setting. Therefore, we are not restricted to parametric and in particular B-spline based deformations. Also, we do not introduce an additional penalty term but use a constrained approach.

The VP approach is also connected to the so-called mass preserving (MP) registration which is related to the Monge–Kantorovich transport problem; see [21, 39]. From a mathematical point of view, the MP approach contains the VP approach (setting the density to 1). However, in the MP approach the images enter directly into the constraints, whereas in the VP they do not. In contrast to the MP approach, the VP approach aims to correct geometrical distortions but does not alter grey values. Another major difference between the above work and ours is the numerical treatment. For the MP approach, a Helmholtz decomposition is exploited, and, after eliminating the constraints, the necessary conditions for the optimization problem are solved using a steepest descent method. Here, we discretize the transformation and the constraints directly and use sequential quadratic programming (SQP) to solve the discrete optimality conditions.

In this paper, we present a flexible constrained image registration approach. It has three main ingredients: a distance measure, a regularizer and the constraints.

Our mathematical framework is general enough to handle a variety of distance measures, including the most popular ones, like those based on the sum of squared differences (SSD), mutual information (MI) or correlation, as long as a Gâteaux derivative exists; see, e.g., [2, 22, 24, 31]. For presentation purposes, we explicitly discuss the approach only for the SSD measure.

Image registration is an ill-posed problem (cf, e.g., [23, 28]). Adding additional constraints does not make the problem well posed. This can be seen by considering the following example: a registration of an image of a disc to a copy of this image. Any rotation about the centre of the disc gives just another copy of the image and is therefore a solution of the registration problem. Since a rotation (or, more generally, a rigid transformation) preserves volume, we also have multiple solutions to the volume preserving registration. For this reason, also the constrained image registration problem has to be regularized. The framework presented here is based on a general regularizer. Any regularizing term with a Gâteaux derivative can be used. This includes well-known choices, such as, for example, the elastic [6, 8, 11, 18], the diffusion [12, 25], and the biharmonic [3, 14] regularizer. For the ease of presentation, we emphasize the most popular of these, the elastic registration.

Also our approach to the constraints is very general. However, since the implementation of the volume preserving constraints is not straightforward, we restrict ourselves to these constraints.

Finally, we suggest the use of a staggered grid discretization. This is a well-known and often used technique in many fields, such as fluid dynamics or electromagnetics. However, we are not aware of any image registration algorithm based on this discretization.

The paper is organized as follows. In section 2, we present the continuous mathematical set-up of the constrained registration problem and derive the continuous Euler–Lagrange equations. Although we do not use these equations in our numerical implementation, they give insight into the problem’s structure and serve as a reference for the discrete analogue. In section 3, we discuss the discretization of the problem. For readers from image processing, who are not familiar with staggered grids, we give a brief and formal introduction. Staggered grid discretization is well known to be stable when dealing with tightly coupled systems of partial differential equations.

Our constraints are discretized using a finite volume discretization of each displaced pixel/voxel. We show that the resulting formula mimics the continuous one. In section 4, we discuss a numerical scheme for the optimization of the discretized image registration problem. In section 5, we present numerical results and finally, we summarize the paper and discuss future work in section 6.

2. Mathematical set-up

With $d \in \mathbb{N}$ we denote the spatial dimension of the given images $R, T : \mathbb{R}^d \rightarrow \mathbb{R}$ which are assumed to be sufficiently smooth. Thus, $T(\mathbf{x})$ gives a grey value at a spatial position \mathbf{x} . We assume that the supports of the images are contained in a bounded domain $\Omega := (0, L)^d$, i.e. $R(\mathbf{x}) = T(\mathbf{x}) = 0$ for $\mathbf{x} \notin \Omega$.

Our goal is to find a ‘reasonable’ transformation such that the ‘distance’ between the reference image R and the deformed template image $T(\mathbf{x} + \mathbf{u}(\mathbf{x}))$ is small. As usual in image registration, we work the displacement $\mathbf{u} = (u^1, \dots, u^d)^T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ rather than the transformation $\varphi(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x})$. It is well known that this problem is ill posed and therefore needs to be regularized. In general, it is common to use a Tikhonov-style regularization. A mathematical formulation of the regularized and constrained problem reads

$$\mathcal{D}[R, T; \mathbf{u}] + \alpha \mathcal{S}[\mathbf{u}] = \min \quad (1a)$$

$$\text{subject to } \mathcal{C}[\mathbf{u}](\mathbf{x}) = 0 \quad \text{for all } \mathbf{x} \in \Omega, \quad (1b)$$

where \mathcal{D} is some distance measure, \mathcal{S} is some regularization term, and \mathcal{C} are some constraints. Here, $\alpha > 0$ is a regularization parameter that compromises between similarity and regularity.

The three building blocks are discussed below. The constraints can be used to supply additional information about the registration problem. For example, in some applications it is of importance that particular points, such as, e.g., anatomical landmarks, are in a precise one-to-one correspondence. With the setting $\mathcal{C}[\mathbf{u}] = \ell^T + \mathbf{u}(\ell^T) - \ell^R$, this fits into our general framework. Using these constraints, one can guarantee the correspondence of the landmarks ℓ^R and ℓ^T ; see, e.g., [7, 13]. In this paper, we focus on applications which require a volume preserving transformation. However, one may set $\mathcal{C}[\mathbf{u}] \equiv 0$ to recover the unconstrained problem.

For the following analysis and numerics any choices and combinations of the building blocks \mathcal{D} , \mathcal{S} and \mathcal{C} are feasible as long as they do have a Gâteaux derivative. For the purpose of this presentation, we restrict the discussion to the following choices.

Distance measures. The distance measure used in this paper is the so-called sum of squared differences (or L_2 norm)

$$\mathcal{D}[R, T; \mathbf{u}] = \frac{1}{2} \int_{\mathbb{R}^d} (T(\mathbf{x} + \mathbf{u}(\mathbf{x})) - R(\mathbf{x}))^2 d\mathbf{x} \quad (2)$$

with Gâteaux derivative (cf, e.g., [28])

$$d_{\mathbf{u}, \mathbf{v}} \mathcal{D}[R, T; \mathbf{u}] = \int_{\mathbb{R}^d} \langle \mathbf{f}(\mathbf{x}, \mathbf{u}(\mathbf{x})), \mathbf{v}(\mathbf{x}) \rangle_{\mathbb{R}^d} d\mathbf{x}, \quad (3)$$

$$\mathbf{f}(\mathbf{x}, \mathbf{u}(\mathbf{x})) := (T(\mathbf{x} + \mathbf{u}(\mathbf{x})) - R(\mathbf{x})) \cdot \nabla T(\mathbf{x} + \mathbf{u}(\mathbf{x})), \quad \mathbf{x} \in \mathbb{R}^d. \quad (4)$$

Other distance measures, for example, those based on mutual information [9, 37], might be used as well.

Remark 1. Note that due to the chain rule any differentiable distance functional based on R and $T(\mathbf{x} + \mathbf{u}(\mathbf{x}))$ has a factor $\nabla T(\mathbf{x} + \mathbf{u}(\mathbf{x}))$. Thus, since we assume R and T to be zero for $\mathbf{x} \notin \Omega$, \mathbf{f} is also zero for $\mathbf{x} \notin \Omega$. Therefore, the integration in equations (2) and (3) reduces to an integration over the domain Ω , only.

Regularizer. In this work, we consider the well-known elastic regularizer [6, 8, 11]. However, our formulation is flexible enough that it can be used with other regularizers, such as, e.g., the

fluid [5, 8], diffusion [12], curvature regularizers [14], or any combinations of these. Each of the above regularizers is of the form

$$\mathcal{S}[u] = \int_{\Omega} \langle \mathcal{B}[u], \mathcal{B}[u] \rangle_{\mathbb{R}^d} dx, \tag{5}$$

for some differential operator \mathcal{B} , and therefore its Gâteaux derivative is

$$d_{u,v}\mathcal{S}[u] = \int_{\Omega} \langle \mathcal{A}[u](x), v(x) \rangle_{\mathbb{R}^d} dx, \quad \mathcal{A} = \mathcal{B}^*\mathcal{B}, \tag{6}$$

where, for the ease of presentation, we assume natural boundary conditions on u . For the elastic regularizer with Lamé constants λ and μ , we have

$$\mathcal{B}[u] = \begin{pmatrix} \sqrt{\mu} & 0 \\ 0 & \sqrt{2\mu + \lambda} \end{pmatrix} \begin{pmatrix} \nabla \times u \\ \nabla \cdot u \end{pmatrix}, \quad \mathcal{A}[u] = -\mu \Delta u - (\mu + \lambda) \nabla \nabla \cdot u \tag{7}$$

with $\nabla \cdot$ the divergence, $\nabla \times$ the curl and Δ the Laplacian.

Constraints. We require the transformation $\varphi(x) := x + u(x)$ to be volume preserving (see also [32]). The transformation φ is volume preserving if and only if for any domain V ,

$$\int_{\varphi(V)} dx = \int_V dx, \quad \text{where} \quad \varphi(V) = \{\varphi(x) : x \in V\}. \tag{8}$$

This constraint implies that not only is the overall volume conserved but also, and most importantly, that of each arbitrarily small subdomain. For a smooth transformation φ , we use the transformation rule to derive the pointwise constraints

$$1 = \det(\nabla \varphi) = \det(I_d + \nabla u) \tag{9}$$

or

$$\mathcal{C}[u] := \det(I_d + \nabla u) - 1 = 0. \tag{10}$$

Here, $I_d \in \mathbb{R}^{d,d}$ denotes the d -by- d identity matrix and $\delta_{j,k}$ its (j, k) th entry.

Remark 2. Volume preserving maps can also have a negative volume. Thus, it is possible to have $\mathcal{C}[u] = \pm 1$. For the application here we do not consider such mappings and consider mapping with a uniform sign alone. Note that if a mapping changes its sign, there is a discontinuity in the determinant that implies a ‘twist’ in $T(u)$.

The Gâteaux derivative of the volume preserving constraints is given by the following lemma.

Lemma 1. *Let v be a suitable perturbation of u . The Gâteaux derivative of \mathcal{C} (cf equation (10)) is given by*

$$d_{u,v}\mathcal{C}[u](x) = \det(I_d + \nabla u(x)) \langle (I_d + \nabla u(x))^{-T}, \nabla v(x) \rangle_{\mathbb{R}^{d,d}}. \tag{11}$$

Proof. A computation gives

$$\begin{aligned} d_{u,v} \det(I_d + \nabla u) &= \lim_{\varepsilon \rightarrow 0} (\det(I_d + \nabla u + \varepsilon \nabla v) - \det(I_d + \nabla u)) \\ &= \det(I_d + \nabla u) \lim_{\varepsilon \rightarrow 0} (\det(I_d + \varepsilon \nabla v (I_d + \nabla u)^{-1}) - 1) \\ &= \det(I_d + \nabla u) \cdot \text{trace}[\nabla v (I_d + \nabla u)^{-1}] \\ &= \langle \det(I_d + \nabla u) (I_d + \nabla u)^{-T}, \nabla v \rangle_{\mathbb{R}^{d,d}}. \end{aligned}$$

□

Remark 3. Applying Cramer’s rule, it can be verified that the Gâteaux derivative of \mathcal{C} is a polynomial in $\partial_j u^k$, $j, k = 1, \dots, d$.

Example 1. For $d = 2$ we have

$$\begin{aligned} \mathcal{C}[\mathbf{u}] &= \nabla \cdot \mathbf{u} + \partial_1 u_1 \partial_2 u_2 - \partial_2 u_1 \partial_1 u_2, \\ \det(I_2 + \nabla \mathbf{u})(I_2 + \nabla \mathbf{u})^{-\top} &= \begin{pmatrix} 1 + \partial_2 u_2 & -\partial_1 u_2 \\ -\partial_2 u_1 & 1 + \partial_1 u_1 \end{pmatrix}. \end{aligned}$$

We now investigate necessary conditions for a solution of the image registration problem (1). For computational purposes, we have to discretize either the optimization problem (1) or the resulting necessary conditions. In the following section, we choose to discretize the optimization problem directly and therefore, the continuous conditions are not used directly in our numerical scheme. However, the discrete conditions have to mimic the continuous analogues and therefore we find it useful to study the latter.

Introducing the Lagrange multiplier $p : \mathbb{R}^d \rightarrow \mathbb{R}$, the Lagrangian of (1) is given by

$$\mathcal{L}[\mathbf{u}, p] = \mathcal{D}[R, T; \mathbf{u}] + \alpha \mathcal{S}[\mathbf{u}] + \int_{\Omega} \mathcal{C}[\mathbf{u}](\mathbf{x}) \cdot p(\mathbf{x}) \, d\mathbf{x} \quad (12)$$

and the continuous Euler–Lagrange equations for (1) read

$$\begin{aligned} 0 &= d_{\mathbf{u};v} \mathcal{L}[\mathbf{u}, p] \\ &= d_{\mathbf{u};v} \mathcal{D}[R, T; \mathbf{u}] + \alpha d_{\mathbf{u};v} \mathcal{S}[\mathbf{u}] + \int_{\Omega} d_{\mathbf{u};v} \mathcal{C}[\mathbf{u}](\mathbf{x}) \cdot p(\mathbf{x}) \, d\mathbf{x}, \end{aligned} \quad (13a)$$

$$0 = d_{p;q} \mathcal{L}[\mathbf{u}, p] = \int_{\Omega} \mathcal{C}[\mathbf{u}](\mathbf{x}) \cdot q(\mathbf{x}) \, d\mathbf{x}, \quad (13b)$$

for any appropriate perturbations v and q . Thus, for any $\mathbf{x} \in \Omega$, we have

$$0 = \mathbf{f}(\mathbf{x}, \mathbf{u}(\mathbf{x})) + \alpha \mathcal{A}[\mathbf{u}](\mathbf{x}) - \nabla \cdot [\det(I_d + \nabla \mathbf{u}(\mathbf{x}))(I_d + \nabla \mathbf{u}(\mathbf{x}))^{-\top} \cdot p(\mathbf{x})] \quad (14a)$$

$$0 = \det(I_d + \nabla \mathbf{u}(\mathbf{x})) - 1, \quad (14b)$$

where we imposed zero Dirichlet boundary conditions for the Lagrange multiplier p .

The system (14) is a highly coupled system of nonlinear partial differential equations (PDE). The quantity \mathbf{f} in (14a), which depends nonlinearly on \mathbf{u} and the images, can be viewed as a force field pushing the template towards similarity. The differential operator \mathcal{A} in (14a) is a linear, elliptic operator. The last term in (14a) is related to the derivative of the constraints which also show up in (14b). For a simpler case, when the \mathbf{f} and \mathcal{A} in (14a) are replaced with the identity operator, existence and uniqueness of a solution can be shown [10, p 324]. However, for the registration problem, it is not easy to show either existence or uniqueness of a solution of the PDE (14). For the purpose of this paper, we therefore assume existence of a solution and remark that proving its existence is a subject of further research.

3. Discretization

There are two approaches for the discretization of the PDE constrained optimization problem (1). In the first so-called *optimize–discretize* approach one forms the Lagrangian (12), then differentiates to obtain the continuous Euler–Lagrange equations (13), which are finally discretized and solved.

The second approach, that we use here, is the so-called *discretize–optimize* approach. Here, one directly discretizes the problem (1) and then solves a constrained optimization problem in a finite (but typically high) dimensional space. Note that we still use the fact that

the discrete Euler–Lagrange equations are a discretization of some differential operators. The advantage of this approach is that we are able to use standard optimization methods for the solution of the problem.

Choosing a stable discretization method for an optimization problem with a differential constraint is a delicate matter. It is well known that such a discretization should fulfil the LBB conditions [1, 4]. It is also well known that some seemingly good discretization methods do not fulfil these conditions (see, e.g., [19] for an elaborate discussion for the Stokes system). Further complications arise in our case where we have differential operators such as the divergence and the curl. We would like to choose a conservative compact discretization scheme. This could be achieved by either mixed finite elements or by staggered grids. Staggered grids are very common for the stable discretization of fluid flow problems (see, e.g., [16]) where first-order differential constraints are discretized and in electromagnetics (see, e.g., [20, 38]) where operators such as curl and divergence are discretized. In the context of fluid flows and electromagnetics, it is well known that compact discretizations are crucial in order to obtain a stable linear system of equations and avoid spurious modes [19, 26, 27]. It is therefore most natural to choose such a discretization for our problem as well. Further investigation is needed to show that the LBB conditions are fulfilled and this will be done in a forthcoming paper.

Though a staggered grid seems to be natural for the discretization of the registration problem on a regular grid, we are not aware of any registration scheme where this discretization is used. We therefore give a brief but formal description in section 3.1. For a more elaborate discussion, see, e.g., [20].

It is important to note that staggered grids only require short differences for the approximation of the derivatives $\partial_j u^k$ and therefore no spurious modes are introduced by discretization; see, e.g., [38].

3.1. Staggered grid discretization

We assume that our discrete images have $m_1 \times \dots \times m_d$ pixels and, for the ease of presentation, that each pixel is square with lengths h . We allow for half step indices. As usual in image processing, we identify pixels/voxels with cell-centred grid points with are therefore labelled with full-integer indices. The knots of the nodal, cell centred, the d face staggered and (for dimension $d = 3$) the d -edge-staggered grids are collected in d -dimensional arrays as follows (see also figure 3 for an illustration),

$$\begin{aligned} X^n &= h(i_1 - \frac{1}{2}, \dots, i_d - \frac{1}{2})_{i_j = \frac{1}{2}, \dots, m_j + \frac{1}{2}}, \\ X^c &= h(i_1 - \frac{1}{2}, \dots, i_d - \frac{1}{2})_{i_j = 1, \dots, m_j}, \\ X^{f,j} &= h(i_1 - \frac{1}{2}, \dots, i_d - \frac{1}{2})_{i_j = \frac{1}{2}, \dots, m_j + \frac{1}{2}} \\ &\qquad\qquad\qquad i_k = 1, \dots, m_k, k \neq j \\ X^{e,j} &= h(i_1 - \frac{1}{2}, \dots, i_d - \frac{1}{2})_{i_j = 1, \dots, m_j}, \\ &\qquad\qquad\qquad i_k = \frac{1}{2}, \dots, m_k + \frac{1}{2}, k \neq j \end{aligned}$$

where $j = 1, \dots, d$. The nodal grid is numbered with half-integers, the cell-centred grid with integers, the j th face (edge) staggered grid with integers (half-integers) except for the j th direction for which half-integers (integers) are used.

We denote the discrete analogue of the continuous vector field $\mathbf{u} = (u^1, \dots, u^d)$ by $U = (U^1, \dots, U^d)$ where U^k is a grid function approximated on the face-staggered grid $X^{f,k}$. Thus, each of the u^k is approximated at different locations which are staggered. We approximate the derivatives $\partial_j u^k$ by

$$\partial_j u^k \approx D_M^{j,k} [U^k].$$




Figure 3. Staggered grids, nodal \blacksquare , cell centered \bullet , face-staggered grids (\blacktriangledown in the x_1 -, \blacktriangleright in the x_2 - and \blackstar in the x_3 -direction), and edge-staggered grids (\blacktriangle in the x_1 -, \blacktriangleleft in the x_2 - and \otimes in the x_3 -direction). Top left: $d = 2$, four pixels, top right: $d = 2$ pixel (i_1, i_2) with grids, bottom left: $d = 3$, voxel (i_1, i_2, i_3) with face-staggered grids and positions of U^1, U^2, U^3 , bottom right: $d = 3$, edge-staggered grids.

Here, the $D_M^{k,k}$ approximate derivatives in the normal direction and the $D_M^{j,k}, j \neq k$, derivatives in tangential directions. Neglecting obvious indices, the operators are defined by

$$\begin{aligned}
 D_M^{k,k}[U^k]_{\dots, i_k, \dots} &:= \frac{1}{h} (U^k_{\dots, i_k + \frac{1}{2}, \dots} - U^k_{\dots, i_k - \frac{1}{2}, \dots}), \\
 D_M^{j,k}[U^k]_{\dots, i_j + \frac{1}{2}, \dots, i_k + \frac{1}{2}, \dots} &:= \frac{1}{h_j} \begin{cases} 0 & i_j = 0, m_j, \\ U^k_{\dots, i_j + 1, \dots, i_k + \frac{1}{2}, \dots} - U^k_{\dots, i_j, \dots, i_k + \frac{1}{2}, \dots} & 0 < i_j < m_j. \end{cases} \quad (15)
 \end{aligned}$$

Hence, $D_M^{k,k}[U^k]$ is located on the cell-centered grid but $D_M^{j,k}[U^k]$ is located on the nodal grid for $d = 2$ and on an edge-staggered grid for $d = 3$. Note that no boundary conditions are needed to calculate derivatives in the normal directions and that we have assumed Neumann boundary conditions in the tangential directions.

Remark 4. To calculate our constraints we will need to approximate sums of products of the form $(\partial_j u_k)(\partial_k u_j)$, see remark 3. For $d = 2$ the normal derivatives are naturally approximated on the nodal grid. However for $d = 3$ these derivatives are located on edge-staggered grids and therefore are not centred at the same places. For this reason we introduce the averaging operators from the edge-staggered grids to the nodal grid. Ignoring obvious indices, for $\ell = 1, 2, 3$, we have

$$P_{M,\ell}^{e \rightarrow n}[V]_{i_1 + \frac{1}{2}, i_2 + \frac{1}{2}, i_3 + \frac{1}{2}} := \begin{cases} 0 & i_\ell = 0, m_\ell \\ \frac{1}{2} (V_{\dots, i_\ell, \dots} + V_{\dots, i_\ell + 1, \dots}) & 0 < i_\ell < m_\ell. \end{cases}$$

Thus, for any choice of three numbers with $\{j, k, \ell\} = \{1, 2, 3\}$, the projections $P_{M,\ell}^{e \rightarrow n}[D^{j,k}[U^k]]$ are positioned on the nodal grid.

Finally, we introduce an averaging operator $P_M^{n \rightarrow c}$ from the nodal grid to the cell-centred grid,

$$P_M^{n \rightarrow c}[V]_{i_1, \dots, i_d} := 2^{-d} \sum_{j_1, \dots, j_d = \pm \frac{1}{2}} V_{i_1+j_1, \dots, i_d+j_d}.$$

For later convenience, we denote by \vec{X} the vector assembled from the lexicographically ordered entries of the array X . Let $\vec{U} := (\vec{U}^{1\top}, \dots, \vec{U}^{d\top})^\top$,

$$D^{k, \ell} \vec{U}^\ell := (D_M^{k, \ell}[U^\ell])^\top \quad \text{and} \quad P^{x \rightarrow y} \vec{V} := (P_M^{x \rightarrow y}[V])^\top.$$

Remark 5. Although the above introduction gives some insight into the actions of the operators $D^{j, k}$ and $P^{x \rightarrow y}$, in particular for coding reasons it might be advantageous to have a compact formal description. Let \otimes denote the Kronecker product of matrices and I_k an identity matrix of appropriate size, then

$$\begin{aligned} D^{j, k} &= I_1 \otimes \dots \otimes I_{j-1} \otimes D_V^{j, k} \otimes I_{j+1} \otimes \dots \otimes I_d, \\ P_\ell^{e \rightarrow n} &= I_1 \otimes \dots \otimes I_{\ell-1} \otimes P_{V, \ell}^{e \rightarrow n} \otimes I_{\ell+1} \otimes \dots \otimes I_d, \\ P_\ell^{f \rightarrow c} &= I_1 \otimes \dots \otimes I_{\ell-1} \otimes P_{V, \ell}^{f \rightarrow c} \otimes I_{\ell+1} \otimes \dots \otimes I_d, \\ P^{n \rightarrow c} &= P_{V, d}^{f \rightarrow c} \otimes \dots \otimes P_{V, 1}^{f \rightarrow c}, \end{aligned}$$

where the matrices $D_V^{k, k} \in \mathbb{R}^{m_k, m_k+1}$, $D_V^{j, k} \in \mathbb{R}^{m_j+1, m_j}$, $P_{V, \ell}^{e \rightarrow n} \in \mathbb{R}^{m_\ell+1, m_\ell}$ and $P_{V, \ell}^{f \rightarrow c} \in \mathbb{R}^{m_\ell, m_\ell+1}$ are given by

$$\begin{aligned} D_V^{k, k} &:= \frac{1}{h} \begin{pmatrix} -1 & 1 & & & \\ & \ddots & \ddots & & \\ & & & -1 & 1 \\ & & & & & \end{pmatrix}, & D_V^{j, k} &:= \frac{1}{h} \begin{pmatrix} 0 & \dots & \dots & 0 \\ -1 & 1 & & \\ & \ddots & \ddots & \\ & & & -1 & 1 \\ 0 & \dots & \dots & 0 \end{pmatrix}, \\ P_{V, \ell}^{e \rightarrow n} &:= \frac{1}{2} \begin{pmatrix} 0 & \dots & \dots & 0 \\ 1 & 1 & & \\ & \ddots & \ddots & \\ & & & 1 & 1 \\ 0 & \dots & \dots & 0 \end{pmatrix}, & P_{V, \ell}^{f \rightarrow c} &:= \frac{1}{2} \begin{pmatrix} 1 & 1 & & \\ & \ddots & \ddots & \\ & & & 1 & 1 \end{pmatrix}. \end{aligned}$$

3.2. Discretization of the building blocks \mathcal{D} , \mathcal{S} and \mathcal{C}

For the particular building blocks we derive discrete analogues. Let $\vec{X}^c = (\vec{X}_1^{c\top}, \dots, \vec{X}_d^{c\top})^\top$, $\vec{R} = R(\vec{X}^c)$ and

$$\vec{T}(\vec{U}) = T(\vec{X}_1^c + P_1^{f \rightarrow c} \vec{U}^1, \dots, \vec{X}_d^c + P_d^{f \rightarrow c} \vec{U}^d)^\top.$$

Note that $\vec{T}(\vec{U})$ is the discrete analogue of the image $T(x + u(x))$ as a function of u . Since T is assumed to be a smooth function, $T(x)$ can be evaluated for any x . In our implementation we use a B-spline interpolation scheme. The assumption on T to be differentiable is for the ease of presentation only. In the continuous formulation (3), derivatives appear only in a weak formulation. Thus, all we need is the existence of a distributional derivative of T .

We denote the Jacobian of \vec{T} by

$$J(\vec{U}) := \frac{\partial \vec{T}}{\partial \vec{U}}(\vec{U}) = (\text{diag}((P_1^{f \rightarrow c})^\top \partial_1 T), \dots, \text{diag}((P_d^{f \rightarrow c})^\top \partial_d T)), \quad (16)$$

where the partial derivatives $\partial_j T$ have to be evaluated at $(\vec{Y}_1, \dots, \vec{Y}_d)$, $\vec{Y}_j = \vec{X}_j^c + Q_j \vec{U}_j$.

Discretizing D. The discretization of $\mathcal{D}[R, T; \mathbf{u}]$ (2) is straightforward,

$$D(\vec{U}) := \frac{1}{2} \|\vec{T}(\vec{U}) - \vec{R}\|_2^2$$

and its derivative, which is also known and interpreted as a force field \vec{F} (see, e.g., [28]), is

$$\vec{F}(\vec{U}) := D_{\vec{U}}(\vec{U}) = J(\vec{U})^\top (\vec{T}(\vec{U}) - \vec{R}). \tag{17}$$

Discretizing S. Following (7), all we need are discretizations $\nabla^h \times$ and $\nabla^h \cdot$ of $\nabla \times$ and $\nabla \cdot$, respectively. Since the discretization for $\nabla \times$ and $\nabla \cdot$ is composed from first-order derivatives, for $d = 2, 3$, we obtain

$$\begin{pmatrix} \nabla_2^h \times \\ \nabla_2^h \cdot \end{pmatrix} := \begin{pmatrix} D^{2,1} & -D^{1,2} \\ D^{1,1} & D^{2,2} \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \nabla_3^h \times \\ \nabla_3^h \cdot \end{pmatrix} := \begin{pmatrix} 0 & -D^{3,1} & D^{2,3} \\ D^{3,1} & 0 & -D^{1,3} \\ -D^{2,1} & D^{1,2} & 0 \\ D^{1,1} & D^{2,2} & D^{3,3} \end{pmatrix}.$$

With $A = B^\top B$, $B = \text{diag}(\sqrt{\mu}, \sqrt{2\mu + \lambda})(\nabla^h \times, \nabla^h \cdot)^\top$, the discretization of $\mathcal{S}[\mathbf{u}]$ (5) is given by

$$S(\vec{U}) := \frac{1}{2} \vec{U}^\top A \vec{U}, \quad S_{\vec{U}}(\vec{U}) = A \vec{U}.$$

Here, we imposed zero Neumann boundary conditions in tangential directions.

Discretizing C. In our discretization of the volume preserving constraints, we exploit a simple finite difference formulation. Discretizing the constraint involves a polynomial in the derivatives of \mathbf{u} . To approximate these derivatives we use (15). The approximation to the derivatives in the normal direction naturally averages in the cell centre, however, the derivatives with respect to tangential directions are averaged on nodes in two dimensions and on cell edges in three dimensions. To calculate products of derivatives in the normal directions, we use averages of short differences. For example, in two dimensions we use

$$\begin{aligned} C(\vec{U}) &= \frac{1}{h} (D_M^{1,1} \vec{U}^1 + D_M^{2,2} \vec{U}^2) + \frac{1}{h^2} D_M^{1,1} \vec{U}^1 \odot D_M^{2,2} \vec{U}^2 \\ &\quad + \frac{1}{h^2} (P^{n \rightarrow c} D_M^{1,2} \vec{U}^1) \odot (P^{n \rightarrow c} D_M^{2,1} \vec{U}^2). \end{aligned} \tag{18}$$

The derivatives can be calculated straightforwardly,

$$\begin{aligned} C_{\vec{U}^1}(\vec{U}) &= \frac{1}{h} (1 + \text{diag}(D_M^{2,2} \vec{U}^2)) D_M^{1,1} + \frac{1}{h^2} P^{n \rightarrow c} D_M^{1,2} \text{diag}(P^{n \rightarrow c} D_M^{2,1} \vec{U}^2), \\ C_{\vec{U}^2}(\vec{U}) &= \frac{1}{h} (1 + \text{diag}(D_M^{1,1} \vec{U}^1)) D_M^{2,2} + \frac{1}{h^2} P^{n \rightarrow c} D_M^{2,1} \text{diag}(P^{n \rightarrow c} D_M^{1,2} \vec{U}^1). \end{aligned}$$

Since the complete formula for the three-dimensional case is lengthy but its derivation is along the same lines as the one for two dimensions, we omit the details. However, it is important to note that for $d = 2$, the resulting nonlinear equations are in general quadratic, but for $d = 3$ they are cubic.

Remark 6. In the first iteration, $\vec{U} = 0$, and the VP constraints imply that the (discrete) divergence of \vec{U} vanishes. Together with the regularization operator, the resulting system is similar to the Stokes system. As is well known, a staggered-grid discretization is crucial for the stability of the system; cf, e.g., [35].

4. Solving the discrete optimization problem

We are now ready to phrase the discrete analogue of the image registration problem (1),

$$J(\vec{U}) := D(\vec{U}) + \alpha S(\vec{U}) = \min \tag{19a}$$

$$\text{subject to } C(\vec{U}) = 0. \tag{19b}$$

In order to solve this problem numerically we use the framework of SQP; see [29] for a detailed discussion. Let \vec{P} be a cell-centred vector of Lagrange multipliers. The Lagrangian of the problem is

$$L(\vec{U}, \vec{P}) = D(\vec{U}) + \frac{\alpha}{2} \vec{U}^\top A \vec{U} + C(\vec{U})^\top \vec{P}.$$

Differentiating with respect to \vec{U} and \vec{P} , we obtain the following discrete version of the Euler–Lagrange equations (13)

$$0 = L_{\vec{U}}(\vec{U}, \vec{P}) = \vec{F}(\vec{U}) + \alpha A \vec{U} + C_{\vec{U}}(\vec{U})^\top \vec{P}, \tag{20a}$$

$$0 = L_{\vec{P}}(\vec{U}, \vec{P}) = C(\vec{U}). \tag{20b}$$

We can now solve the nonlinear system (20) numerically by using a Newton-type method. Approximating the (1, 1) block of the Hessian by

$$H := \alpha A + J^\top J, \tag{21}$$

where J is defined in (16), we obtain the following linear system of equations to be solved at each iteration:

$$\begin{pmatrix} H & C_{\vec{U}}^\top \\ C_{\vec{U}} & 0 \end{pmatrix} \begin{pmatrix} s_{\vec{U}} \\ s_{\vec{P}} \end{pmatrix} = \begin{pmatrix} L_{\vec{U}}(\vec{U}, \vec{P}) \\ L_{\vec{P}}(\vec{U}, \vec{P}) \end{pmatrix}. \tag{22}$$

The system (22) is a so-called Karush–Kuhn–Tucker (KKT) system; it is symmetric but indefinite. Solving KKT systems is a well-known challenge. Similar systems arise for example in fluid dynamics (cf, e.g., [19, 36]) and the solution for this case has been addressed by many authors; see, e.g., [34–36]. Nevertheless, the robust and effective solution of such systems is still an open research topic. Here we have used MINRES [30] with a block diagonal preconditioning as proposed in [34]. This preconditioner can be written as

$$M = \begin{pmatrix} H & 0 \\ 0 & \hat{S} \end{pmatrix},$$

where \hat{S} is an approximation to the Schur complement

$$S := C_{\vec{U}} H^{-1} C_{\vec{U}}^\top.$$

Here, we use the approximation suggested in [34],

$$\hat{S}^{-1} = (C_{\vec{U}} C_{\vec{U}}^\top)^{-1} C_{\vec{U}} H C_{\vec{U}}^\top (C_{\vec{U}} C_{\vec{U}}^\top)^{-1}. \tag{23}$$

Note that $(C_{\vec{U}} C_{\vec{U}}^\top)^{-1} C_{\vec{U}}$ is the pseudo-inverse of $C_{\vec{U}}$. The application of the preconditioner only involves a multiplication of H^{-1} and \hat{S}^{-1} with a vector. However, an efficient numerical scheme is not straightforward and will be addressed in a forthcoming paper. Here, we use a multigrid approach.

After the KKT system has been solved, we update \vec{U} by setting

$$\vec{U} \leftarrow \vec{U} + \gamma s_{\vec{U}}.$$

As is common in SQP algorithms [29], the parameter γ is chosen such that the L_1 merit function

$$\text{merit}_{\text{KKT}}(\vec{U}) := \mathbf{D}(\vec{U}) + \alpha \mathbf{S}(\vec{U}) + \theta \|\mathbf{C}(\vec{U})\|_1 \quad (24)$$

decreases, where $\theta := \|\vec{P}\|_\infty + \theta_{\min}$ with a fixed parameter θ_{\min} , $\|x\|_1 = \sum_j |x_j|$, and \vec{P} is the least squares multiplier computed by solving

$$(\mathbf{C}_{\vec{U}} \mathbf{C}_{\vec{U}}^\top) \vec{P} = -\mathbf{C}_{\vec{U}} (\mathbf{D}_{\vec{U}} + \alpha \mathbf{S}_{\vec{U}}). \quad (25)$$

In order to avoid being too far from feasibility we project the intermediate \vec{U} to the constraints. Here we take the advantage of a second merit function

$$\text{merit}_{\mathbf{C}}(\vec{U}) := \|\mathbf{C}(\vec{U})\|_2^2. \quad (26)$$

If $\text{merit}_{\mathbf{C}}(\vec{U}) \geq \text{tol}_{\mathbf{C}}$, we calculate a step \vec{U}_P such that ideally

$$\mathbf{C}(\vec{U} + \vec{U}_P) \approx \mathbf{C}(\vec{U}) + \mathbf{C}_{\vec{U}}(\vec{U}) \vec{U}_P = 0.$$

If $\mathbf{C}_{\vec{U}}$ has full rank, a solution is given by $\vec{U}_P = \mathbf{C}_{\vec{U}}(\vec{U})^\top \vec{W}$, where \vec{W} is the solution of the projection system

$$(\mathbf{C}_{\vec{U}} \mathbf{C}_{\vec{U}}^\top) \cdot \vec{W} = -\mathbf{C}(\vec{U}). \quad (27)$$

Note that the very same system has already been addressed in the Schur complement approximation (23) and (25).

The above step is repeated until convergence of the optimization process, which is measured by the change of \vec{U} . In order to find an appropriate regularization parameter, we solve (19) for a few values of α , where we start with a large α and slowly decrease it until our stopping criterion is fulfilled. For the results presented in section 5, since choosing α depends on the statistics of the difference between the images (which is unknown) we base the stopping criterion on a visual inspection of the images. Our numerical scheme is summarized in algorithm 1.

5. Numerical examples

The above scheme has been implemented using MATLAB 6.5. The following computations are performed on a DELL Inspiron 8600 notebook (1.4 GHz, 1GB RAM, 60G disk space) under Windows XP. The computation time (CPU time) for the blob example 5.1 (64×64 images) is about 28 s for the unconstrained registration (38 iterations needed to fulfil our stopping criterion $\|\mathbf{u}^{\text{old}} - \mathbf{u}\| \leq \text{tol}_U := 10^{-2}$, for the constrained problem we spend about 69 s (21 iterations needed).

The computation time for the MRI example 5.2 (128×128 images) is about 296 s for the unconstrained registration (68 iterations needed to fulfil our stopping criterion $\|\mathbf{u}^{\text{old}} - \mathbf{u}\| \leq \text{tol}_U := 10^{-1}$, for the constrained problem we spend about 846 s (23 iterations needed).

5.1. The blob

To illustrate the potential of the volume preserving registration we present a synthetic example; see figure 4. The reference image (top right) shows an elliptic global structure which contains a small almost circular object. The template (top left) shows a rotated version of the global ellipse, where the inner structure is translated and considerably enlarged. Note that this example mimics the situation for contrast enhanced images: slightly deformed global structures, where inner structures may change drastically due to contrast uptake.

Algorithm 1. Volume preserving image registration: $\vec{U} \leftarrow \text{VPIR}(R, T)$.

```

1: Set  $k \leftarrow 0, \vec{U} \leftarrow 0, \vec{P} \leftarrow 0, n \leftarrow \text{length}(\vec{U})$ .
2: for  $k = 0, \dots$  do
3:   Compute  $D, J, D_{\vec{U}}, S, S_{\vec{U}}, C$  and  $C_{\vec{U}}$ , set  $\vec{U}_{\text{old}} = \vec{U}$ .
4:   Set  $\begin{pmatrix} L_{\vec{U}} \\ L_{\vec{P}} \end{pmatrix} \leftarrow \begin{pmatrix} D_{\vec{U}} + \alpha S_{\vec{U}} + C_{\vec{U}}^T \vec{P} \\ C \end{pmatrix}$ .
      {Computation of the SQP step}
5:   Solve the KKT system (22) for  $s_{\vec{U}}$  and  $s_{\vec{P}}$ 
6:   Solve for the Lagrange-multiplier  $\vec{P}$ , cf (25); set  $\theta \leftarrow \|\vec{P}\|_{\infty}$ 
7:    $\vec{U} \leftarrow \text{LS}(\vec{U}, s_{\vec{U}}, \text{merit}_{\text{KKT}}, L_{\vec{U}})$ ; see algorithm 2.
8:   Update  $C$  and  $C_{\vec{U}}$ .
      {Computation of the projection onto the constraints}
9:   while  $\text{merit}_C > \text{tol}_C$ , do
10:    solve  $(C_{\vec{U}} C_{\vec{U}}^T) \vec{W} = -C$ , set  $\vec{U}_P \leftarrow C_{\vec{U}}^T \vec{W}$ .
11:     $\vec{U} \leftarrow \text{LS}(\vec{U}, \vec{U}_P, \text{merit}_C, C_{\vec{U}}^T C)$ ; see algorithm 2.
12:    if  $\|\vec{U}_{\text{old}} - \vec{U}\| \leq \text{tol}_U$  then
13:      return  $\vec{U}$ , done.
14:    end if
15:  end while
16: end for

```

Algorithm 2. Armijo's line search: $\vec{U} \leftarrow \text{LS}(\vec{U}, s_{\vec{U}}, \text{merit}, \vec{G})$.

```

Set  $j \leftarrow 0, \gamma \leftarrow 1$  and  $\eta \leftarrow 10^{-5}$ .
while true do
  Set  $\vec{U}_t \leftarrow \vec{U} + \gamma s_{\vec{U}}$ .
  if  $\text{merit}(\vec{U}_t) < \text{merit}(\vec{U}) + \eta \gamma \vec{G}^T s_{\vec{U}}$  then
    break
    {step reduces merit function}
  end if
  if  $j > \max_{\ell}$  then
    error
    {step not successful after  $\max_{\ell}$  line search steps}
  end if
  Set  $\gamma \leftarrow \gamma/2$  and  $j \leftarrow j + 1$ .
end while
Set  $\vec{U} \leftarrow \vec{U}_t$ .

```

As is apparent from figure 4, the unconstrained registration gives very good results if we are looking at the difference between the reference and deformed template images alone. However, as expected, the inner structure has been reduced so as to fit the one in the reference image. This results in a drastic change of volume, which can be observed from the visualization of a part of the grid in figure 4 (middle right) corresponding to a region of interest emphasized in the template image (top left). Thus, for contrast enhanced images, the registration gives meaningless results, though the difference is small.

Figure 4 also shows the results of the volume preserving registration (bottom left). As is apparent from this figure, the global deformation has been resolved, the inner ellipse has been moved to match the inner ellipse in the reference image. However, the volume of the

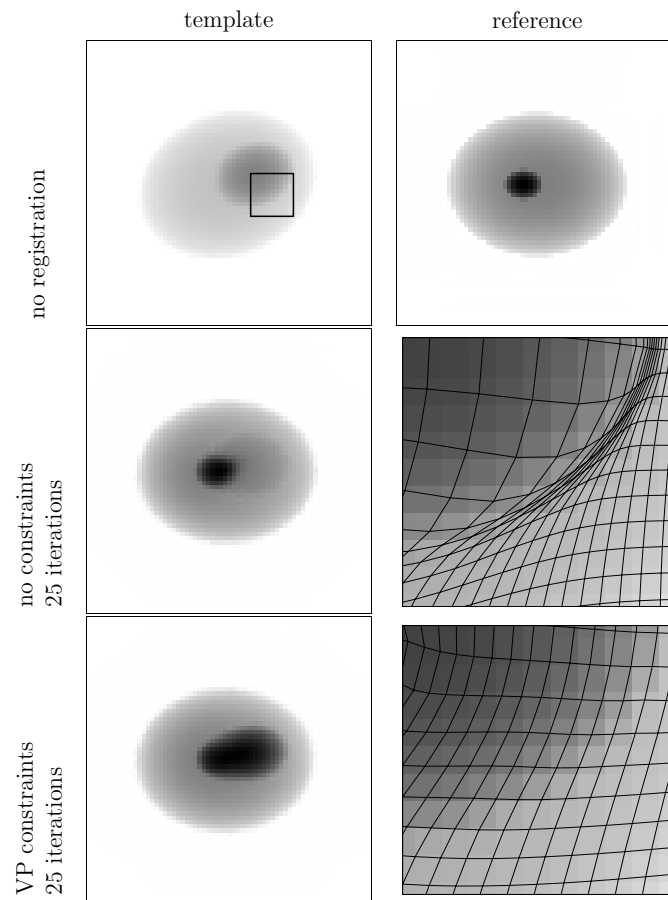


Figure 4. Synthetic example. Left column: deformed template, right column: reference and details of the deformed grid; top row: template and reference, no registration, middle row: deformed template and details with grid after unconstrained registration, bottom row: deformed template and details with grid after VP-constrained registration. For both schemes, we choose $\alpha = 10^3$ and stopped after 25 iterations.

inner ellipse has not been altered, which leads to a larger difference as in the unconstrained case but also to a more realistic registration (assuming the VP assumption holds); see also the deformed grid (bottom right).

In order to compare these results, we choose $\alpha = 10^3$ and stop after convergence, which occurs after at most 25 iterations for both registrations. Note that in order to make the results comparable, we decide to stop both schemes after 25 iterations ignoring the tolerance-based stopping rule. Also note that $h = 1$ in our implementation. The values for the difference D and the constraints C for the un- and VP-constrained registration are summarized in table 1.

5.2. MRI scans

In our second example, we discuss results obtained for the images shown in figure 1. Figure 5 shows the results after two (second row) and ten iterations (third row) of the unconstrained

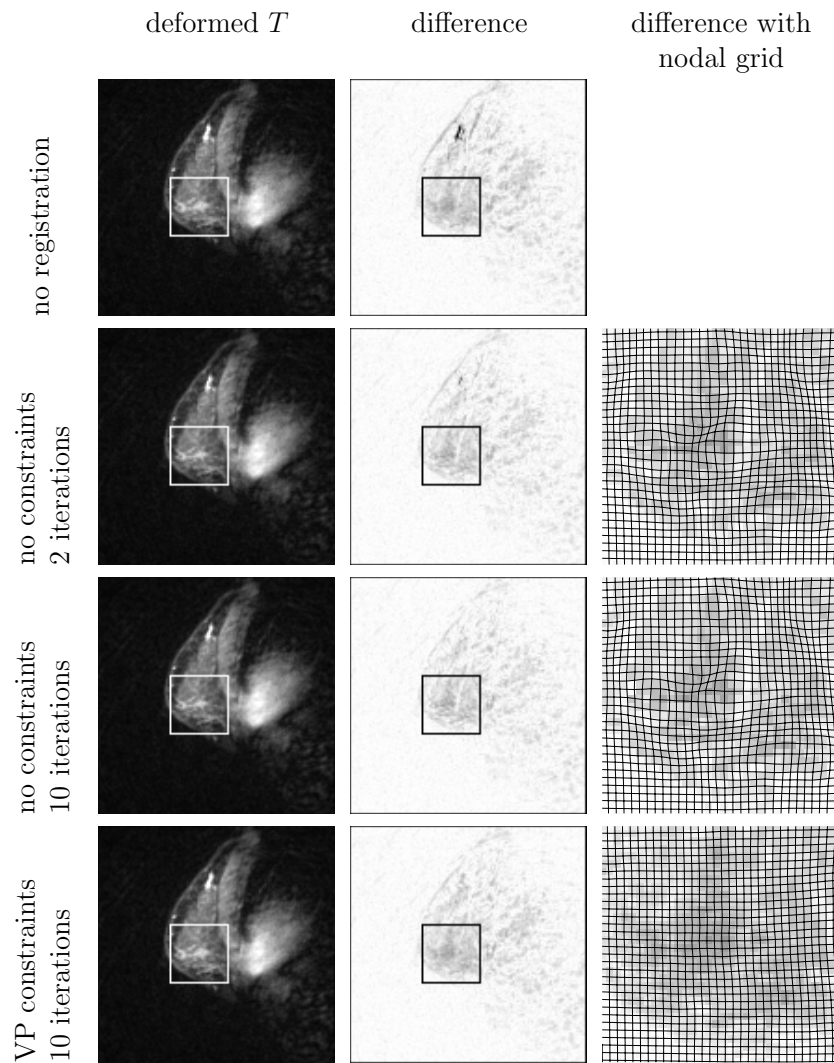


Figure 5. Registration results for the images of figure 1. Left column: deformed template images T_u , middle column: difference image $|R - T_u|$ with region of interest (ROI), right column: ROI with nodal grid, vertices connected by straight lines; row 1: no registration, row 2: no constraints, two iterations, row 3: no constraints, ten iterations, and row 4: volume preserving constraints, ten iterations.

registration as well as after ten iterations of the VP-constrained registration (fourth row). The numerical results are summarized in table 1. After ten iterations both schemes have converged.

Although the numbers (cf table 1) indicate a larger reduction of the difference by the unconstrained registrations, the ranking is not so clear if one looks at the difference images, cf figure 5. Here, the difference after ten steps un- and VP-constrained registration looks pretty much the same. After two steps of the unconstrained registration the bright spot in the top part of the image has not been resolved satisfactorily. The explanation is that small spots which are related to noise in the MRI images and hardly visible in the images are registered in the unconstrained registration. This leads to a large reduction though it is hardly visible. To

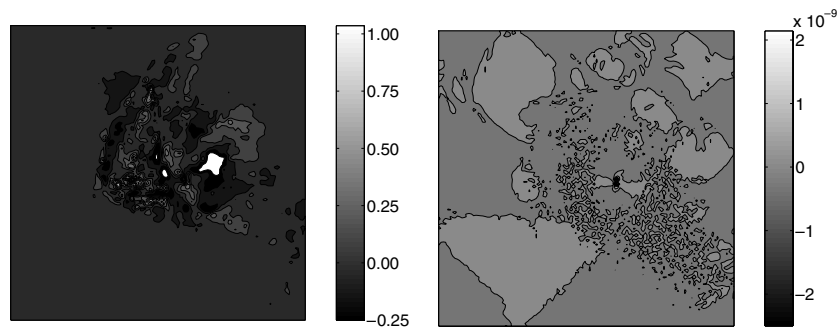


Figure 6. Volume preservation of the unconstrained (left) and constrained (right) registration results for the MRI example.

Table 1. Numerical results for the un- and VP-constrained registrations; k is the number of iterations performed.

	α	k	$D(\vec{U}^{(k)})/D(0)$	$\ C(\vec{U}^{(k)})\ _\infty$
Blob				
Unconstrained	10^3	25	0.21	0.87
VP constrained	10^3	25	0.73	$\leq 10^{-6}$
MRI				
Unconstrained	10^5	2	0.81	1.36
Unconstrained	10^5	10	0.78	1.36
VP constrained	10^5	10	0.87	$\leq 10^{-6}$

remove these small spots, the volume has to be changed locally. However, the registration of these small spots does not contribute to a meaningful solution for this problem.

In figure 6, we display the pointwise map of the change of volume. Using the unconstrained approach, we observe a considerable change of volume for the breast with a peak value of 1.36. Thus, part of the breast has been enlarged by a factor of 2.36. For the constrained approach, we observe that the volume change is below a user supplied threshold (here, $\text{tol}_C = 10^{-6}$) everywhere. In fact, since we used a quasi-Newton scheme for projection, the numbers are around 10^{-9} .

6. Summary

A general approach to image registration is based on a variational formulation, where the overall goal is to minimize a certain energy functional. However, it is well known that image registration is an ill-posed problem. Therefore, for a particular application, one aims to add as much information as possible.

One step in this direction is to add a designed regularizer, which penalizes unwanted solutions but which is also inevitable to make the problem well posed. In this paper, we consider a further step, which is to provide additional information by adding additional constraints. Constraints are sometimes treated by penalization [7, 32], which in a sense is similar to the regularized approach. However, in the penalized version it is not guaranteed that the constraints hold whereas in the constrained approach it is.

To our best knowledge, the first paper dealing with constrained image registration is [13]. However, there only a small number of so-called landmark constraints are considered. In this paper, we consider a more general approach to constrained registration, which allows for pointwise constraints.

We present an implementation of this general approach for volume preserving constraints. In particular, we present a stable and consistent staggered grid discretization of the continuous variational approach. We tested our algorithm on real medical data as well as on synthetic data.

From the numerical point of view, the constrained approach leads to a very challenging and highly nonlinear optimization problem where both the discretization and the numerical optimization are not obvious. We used state-of-the-art SQP optimization techniques that allow for an efficient solution of the registration problem. Our formulation opens up a few avenues of research such as effective solvers for the KKT systems and inexact SQP methods. These issues will be addressed in a subsequent paper.

From an imaging point of view, the addition of the VP constraint should be done judiciously. It is clear that such a constraint can lead to greater distance between the images thus one may want to apply the VP constraint only to a part of the image. Our formulation is flexible and can be modified to deal with this approach.

Acknowledgments

JM was supported by the US National Institutes of Health under grant NIHR01 HL 068904. We are indebted to Michele Benzi and the anonymous referees for suggestions and discussions.

References

- [1] Fortin M and Glowinski R 1983 *Augmented Lagrangian Methods: Applications in the Numerical Solution of Boundary-Value Problems (Studies in Mathematics and Its Applications)* (Amsterdam: North-Holland)
- [2] Droske M and Rumpf M 2004 A variational approach to non-rigid morphological registration *SIAM J. Appl. Math.* **64** 668–87
- [3] Bookstein F L 1989 Principal warps: thin-plate splines and the decomposition of deformations *IEEE Trans. Pattern Anal. Mach. Intell.* **11** 567–85
- [4] Brezzi F and Fortin M 1991 *Mixed and Hybrid Finite Element Methods* (Berlin: Springer)
- [5] Bro-Nielsen M 1996 Medical image registration and surgery simulation *PhD Thesis* IMM, Technical University of Denmark
- [6] Broit C 1981 Optimal registration of deformed images *PhD Thesis* Computer and Information Science, University of Pennsylvania, PA
- [7] Christensen G E and Johnson H J 2001 Consistent image registration *IEEE Trans. Med. Imaging* **20** 568–82
- [8] Christensen G E 1994 Deformable shape models for anatomy *PhD Thesis* Sever Institute of Technology, Washington University
- [9] Collignon A, Vandermeulen A, Suetens P and Marchal G 1995 *3D Multi-Modality Medical Image Registration Based on Information Theory (Computational Imaging and Vision vol 3)* pp 263–74
- [10] Courant R and Hilbert D 1962 *Methods of Mathematical Physics vol II* (New York: Wiley)
- [11] Fischer B and Modersitzki J 1999 Fast inversion of matrices arising in image processing *Numer. Algorithm* **22** 1–11
- [12] Fischer B and Modersitzki J 2002 Fast diffusion registration *AMS Contemp. Math. Inverse Probl. Image Anal. Med. Imaging* **313** 117–29
- [13] Fischer B and Modersitzki J 2003 Combination of automatic non-rigid and landmark based registration: the best of both worlds *Medical Imaging 2003: Image Processing* ed M Sonka and J M Fitzpatrick *Proc. SPIE* **5032** 1037–48
- [14] Fischer B and Modersitzki J 2003 Curvature based image registration *J. Math. Imaging Vis.* **18** 81–5
- [15] Fitzpatrick J M, Hill D L G and Maurer C R Jr 2000 Image registration *Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis* ed M Sonka and J M Fitzpatrick *SPIE* pp 447–513

- [16] Fletcher C A J 1988 *Computational Techniques for Fluid Dynamics* vol II (Berlin: Springer)
- [17] Fortin M and Glowinski R 1983 *Augmented Lagrangian Methods: Applications in the Numerical Solution of Boundary-Value Problems* (Amsterdam: North-Holland)
- [18] Gee J C, Haynor D R, Le Briquer L and Bajcsy R 1997 Advances in elastic matching theory and its implementation *CVR Med.* pp 63–72
- [19] Gresho P M and Sani R L 1987 On pressure boundary conditions for the incompressible Navier–Stokes equations *Int. J. Numer. Methods Fluids* **7** 1111–45
- [20] Haber E and Ascher U 2001 Fast finite volume simulation of 3D electromagnetic problems with highly discontinuous coefficients *SIAM J. Sci. Comput.* **22** 1943–61
- [21] Haker S and Tannenbaum A 2001 Optimal transport and image warping *MICCAI 2001* pp 120–7
- [22] Heldmann S, Mahnke O, Potts D, Modersitzki J and Fischer B 2004 Fast computation of mutual information in a variational image registration approach *Bildverarbeitung für die Medizin 2004* (Berlin: Springer) pp 1–5
- [23] S Henn and K Witsch 2001 Iterative multigrid regularization techniques for image matching *SIAM J. Sci. Comput.* **23** 1077–93
- [24] Hermosillo G 2002 Variational methods for multimodal image matching *PhD Thesis* Université de Nice, France
- [25] Horn B K P and Schunck B G 1981 Determining optical flow *Artif. Intell.* **17** 185–204
- [26] Hyman J M and Shashkov M 1999 The orthogonal decomposition theorems for mimetic finite difference methods *SIAM J. Numer. Anal.* **36** 788–909
- [27] Jin J 1993 *The Finite Element Method in Electromagnetics* (New York: Wiley)
- [28] Modersitzki J 2004 *Numerical Methods for Image Registration* (Oxford: Oxford University Press)
- [29] Nocedal J and Wright S J 1999 *Numerical Optimization* (New York: Springer)
- [30] Paige C C and Saunders M A 1975 Solution of sparse indefinite systems of linear equations *SIAM J. Numer. Anal.* **12** 617–29
- [31] Roche A 2001 Recalage d'images médicales par inférence statistique *PhD Thesis* Université de Nice, Sophia-Antipolis, France
- [32] Rohlfing T, Maurer C R Jr, Bluemke D A and Jacobs M A 2003 Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint *IEEE Trans. Med. Imaging* **22** 730–41
- [33] Schnabel J A, Tanner C, Castellano-Smith A D, Degenhard A, Leach M O, Hose D R, Hill D L G and Hawkes D J 2003 Validation of non-rigid image registration using finite element methods: application to breast MR images *IEEE Trans. Med. Imaging* **22** 865–73
- [34] Silvester D, Elman H, Kay D and Wathen A 2001 Efficient preconditioning of the linearized Navier–Stokes equations *J. Comput. Appl. Math.* **128** 261–79
- [35] Trottenberg U, Oosterlee C and Schüller A 2001 *Multigrid* (New York: Academic)
- [36] Turek S 1999 *Efficient Solvers for Incompressible Flow Problems* (New York: Macmillan)
- [37] Viola P A 1995 Alignment by maximization of mutual information *PhD Thesis* Massachusetts Institute of Technology pp 1–155
- [38] Yee K S 1966 Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media *IEEE Trans. Antennas Propag.* **14** 302–7
- [39] Zhu L, Haker S and Tannenbaum A 2003 Area preserving mappings for the visualization of medical structures *MICCAI 2003* pp 277–84