Technical Report

TR-2005-009

Image registration with guaranteed displacement regularity

by

Eldad Haber, Jan Modersitzki

Mathematics and Computer Science

EMORY UNIVERSITY

Image registration with guaranteed displacement regularity

Eldad Haber* Jan Modersitzki†

July 20, 2005

Abstract

The goal of image registration is twofold. One goal is to enforce a certain similarity of two images by geometrically transforming one of the images. The second goal is to keep this transformation meaningful or regular. There is a large amount of approaches aiming for regularity. Most of those are based on certain regularization techniques, others use so-called regridding options.

Here, we present a mathematically sound formulation that explicitly controls the deformation in terms of the determinant of the Jacobian of the transformation. In contrast to similar work, we use pointwise inequality constraints, i.e., the volume is controlled voxel by voxel and and not by integral measures. This approach guaranties grid regularity and prevent folding.

As it turns out, the discretization of the volume constraint inequality is not straightforward. Therefore, we present a new type of discretization enabling the detection of twists in a pixel or a voxel. Such detection is crucial since a twists indicates that a transformation is physically meaningless.

To solve the large-scale inequality constrained optimization problem, we present a numerical approach based on an interior point method. We finally present some numerical examples that demonstrate the advantage of including inequality constraints explicitly.

*Dept of Mathematics and Computer Science, Emory University, Atlanta GA 30322 haber@mathcs.emory.edu.

†Institute of Mathematics, University of Lübeck modersitzki@math.uni-luebeck.de,
This work was supported by NSF grant CCF-0427094 and NIH grant NIH R01 HL068904.

1 Introduction

Registration is the determination of a geometrical transformation that aligns points in one view of an object with corresponding points in another view of the same or a similar object. There exist many instances particularly in a medical environment which demand for registration. Examples include the treatment verification of pre- and post-intervention images, the study of temporal series of cardiac images, and the monitoring of the time evolution of an agent injection subject to patient motion. Another important application is the combination of information from multiple images, acquired using different modalities, like for example computer tomography (CT) and magnetic resonance imaging (MRI), a technique also known as fusion. The problem of fusion and registration arises whenever images acquired from different subjects, at different times, or from different scanners need to be combined for analysis or visualization. In the last two decades, computerized non-rigid image registration has played an increasingly important role in medical imaging, see, e.g., [16, 9, 24, 17] and references therein.

A numerical treatment of the problem is typically based on two basic building blocks. The first one is a so-called distance measure $\mathcal{D}$, quantifying distance or similarity of two given images R and T and the second one is a so-called regularizer $\mathcal{S}$ which penalizes unwanted and/or unreasonable solutions. Since image registration is an ill-posed problem (see, e.g., [17]), regularization is inevitable. A common treatment of the registration problem is based on the following variational approach. Find a smooth transformation $\boldsymbol{u} = (u^1, \ldots, u^d)^\top : \mathbb{R}^d \to \mathbb{R}^d$ minimizing the joint energy

$$\mathcal{J}(\boldsymbol{u}) := \mathcal{D}(R, T(\boldsymbol{u})) + \alpha \mathcal{S}(\boldsymbol{u}). \tag{1}$$

Here, $\alpha > 0$ is a regularization parameter and compromises between similarity and regularity. The functional $\mathcal{D}$ measure the distance between the images and can be based e.g. on the Sum of Squares Difference (SSD), mutual information [22], or normal gradient fields [15, 12]. For ease of presentation, in this paper we focus on SSD,

$$\mathcal{D}(R, T(\boldsymbol{u})) = \tfrac{1}{2}\|T(\boldsymbol{u}) - R\|^2. \tag{2}$$

The regularization operator is design to yield a unique deformation field $\boldsymbol{u}$. For ease of presentation, we focus on the most common so-called elastic regularization (cf., e.g., [3, 5, 17]),

$$\mathcal{S}(\boldsymbol{u}) = \tfrac{\omega}{2}\textstyle\sum_j \| \nabla \boldsymbol{u}^j \|^2 + \frac{\lambda+\omega}{2} \| \nabla \cdot \boldsymbol{u} \|^2, \tag{3}$$

where λ, ω are the so-called Lamé-constants; cf., e.g., [17]. It is worthwhile noticing that our framework enables the usage of any distance measure combined with any regularizer as long as the Gâteaux-derivatives exist.

A registration result is typically evaluated by looking at the image distance and the regularity of the displacement; cf. Figure 5 and 6. If the displacement is more or less regular, the mapping might be considered as reasonable. If, on the other hand, the displacement is irregular, one may want to reject the results and start with a larger value of α. The worst case is that the displacement shows twists (or folding), indicating that the transformation is not even bijective. Note that the variational formulation (1) ensures a "smooth" solution but has no build-in mechanism to prevent a very irregular displacement.

The basic idea of the following new approach is to integrate this evaluation practice directly into the mathematical framework. In our new formulation, we explicitly demand for transformation regularity and ensure bijectivity. Mimicking the human evaluation, we constrain the wanted transformation in terms of minimal and maximal expansion. For any set $V \subset \mathbb{R}^d$ and $\boldsymbol{\varphi}(\boldsymbol{x}) = \boldsymbol{x} + \boldsymbol{u}(\boldsymbol{x})$ we compute the volume and the transformed volume

$$\mathrm{vol}(V, \boldsymbol{\varphi}) := \int_{\boldsymbol{\varphi}(V)} d\boldsymbol{y} = \int_V \det(I_d + \nabla \boldsymbol{u})\, d\boldsymbol{x} \tag{4}$$

and require that the ratio $\mathrm{vol}(V, \boldsymbol{\varphi})/\mathrm{vol}(V, \boldsymbol{x})$ is reasonable. For a smooth displacement $\boldsymbol{u}$ it is therefore equivalent to require

$$0 < \mathcal{C}(\boldsymbol{u}) := \det(I_d + \nabla \boldsymbol{u}) < \infty \quad \text{for all} \quad x \in \Omega. \tag{5}$$

However, for most practical considerations, the above bounds are not sufficient. Some registration algorithms monitor the size of the Jacobian $\mathcal{C}(\boldsymbol{u})$ and re-initialize or stop the registration if its value is small (cf., e.g., [17]). This monitoring is usually done implicitly and is not a part of the mathematical formulation. Therefore, registration algorithms require manual intervention and visual inspection of the distorted grids in order to achieve physically acceptable results. Another straightforward idea to prevent grid folding is to use high values of the regularization parameter α. However such values can cause an inferior distance between the images.

A seemingly simple approach to enforce a "reasonable" Jacobian is to add an additional regularization to the objective functional,

$$\mathcal{J}(\boldsymbol{u}) := \mathcal{D}(R, T(\boldsymbol{u})) + \alpha \mathcal{S}(\boldsymbol{u}) + \beta \|\mathcal{C}(\boldsymbol{u}) - 1\|^2.$$

Similar approaches were suggested in [4, 19, 6]. However, this approach has some disadvantages. For a finite, not too large β, we have only an overall penalty for "non-reasonable" transformations; locally, the transformation can still be irregular. The penalty approach with a moderate β's does not prevent vanishing Jacobians. On the other hand, it is well-known that increasing β generally leads to ill-conditioned optimization problems; see, e.g., [18]. Furthermore, for large β's one gets a volume preserving constraint which may be too tight for many applications: in practice, one is often interested in a softer inequality constraint which allows parts of the image to shrink in a prescribed band.

In this paper we examine inequality constraints applied to the transformed volume. Thus, an additional regularity of the transformation enters the mathematical formulation of the registration problem. As a consequence, our numerical approach explicitly controls the displacement and therefore no manual intervention is required in order to obtain application conform reliable displacements. Furthermore, this approach prevents folding of the grid even for very small regularization parameters α. Our approach is based on the inequality constraint minimization problem

$$\begin{aligned} \text{minimize} \quad & \mathcal{J}(\boldsymbol{u}) := \mathcal{D}(R, T(\boldsymbol{u})) + \alpha \mathcal{S}(\boldsymbol{u}) && \text{(6a)} \\ \text{subject to} \quad & \kappa_m(\boldsymbol{x}) \leq \mathcal{C}(\boldsymbol{u}) \leq \kappa_M(\boldsymbol{x}), && \text{(6b)} \end{aligned}$$

where we use non-negative *compressibility functions* $\kappa_m \leq \kappa_M$. For $\kappa_m \equiv \kappa_M \equiv 1$, we obtain equality constraints and thus volume preserving image registration with hard equality constraints; cf. e.g. [19, 23, 14].

The constraints are phrased as a determinant of the Jacobian. Hence, the constraints are polynomials of degree d in derivatives of $\boldsymbol{u}$. The solution of continuous optimization problems with differential inequality constraints is not trivial. Here we use the discretize then optimize approach. In the first stage we discretize the optimization problem (6a) and the constraints (6b). We then solve the finite dimensional (but relatively large) discrete optimization problem. Although the fundamental structure here may be considered discrete, it is highly useful to view it as an instance of a family of finer and finer discretizations of a continuous problem; see, e.g., [1, 11, 2]. In an abuse of notation, we use the same notation for discrete and continuous variables.

The rest of the paper is organized as follows. In Section 2 we discuss the underlying discretization. Particularly the discretization of the Jacobian is a delicate matter and care must be taken in order to obtain meaningful results.

In Section 3 we summarize and discuss an interior point method for the solution of the optimization problem. We elaborate on the solution of the linear system obtained in each iteration and describe a multilevel approach that enable to accelerate the computation. Finally, in Section 4 we present 2D and 3D examples that demonstrate the effectiveness and superiority of our approach.

2 Consistent Discretization

Choosing a stable discretization method for a system of partial differential equations (PDE's) with mixed derivatives is a delicate matter. In particular, the discretization of the constraint is not straightforward. We start by discussing the discretization of the displacement field followed by the discretization of the constraint which consists of the main difficulty. We then briefly describe the discretization of the objective function.

2.1 Discretizing the displacement

We assume that our discrete images have $m_1 \times \ldots \times m_d$ pixels, where $d = 2, 3$ is the image dimensionality. We also assume that each pixel/voxel is a box of lengthes $h_1, \ldots, h_d$. In our description we allow for half step indices. As usual in image processing, we identify pixels/voxels with cell centered grid points $x_{i_1+\frac{1}{2},\ldots,i_d+\frac{1}{2}}$. Given a box centered at $x_{i_1+\frac{1}{2},\ldots,i_d+\frac{1}{2}}$, the four/eight corners are numbered by full integer indices $i_k, i_k + 1$, $k = 1, \ldots, d$. The displacement $\boldsymbol{u} = (u^1, \ldots, u^d)$ is discretized in the nodal grid (corner of each box; see Figure 1). This discretization is different from the one proposed in [13] and it is related to the special structure of the constraints; cf. Section 2.2.

2.2 Discretizing the constraints

Similar to our previous work [14] we are motivated by a finite volume approach. We discuss the difficulties of finite volume discretization in 2D in some length. We then also discuss the (non-trivial) extension to 3D.

In the volume preserving approach [13], we demanded that the volume (or area for 2D) of every deformed box is preserved. We therefore discretized the volume of every deformed box v, by the cross-product of the diagonals

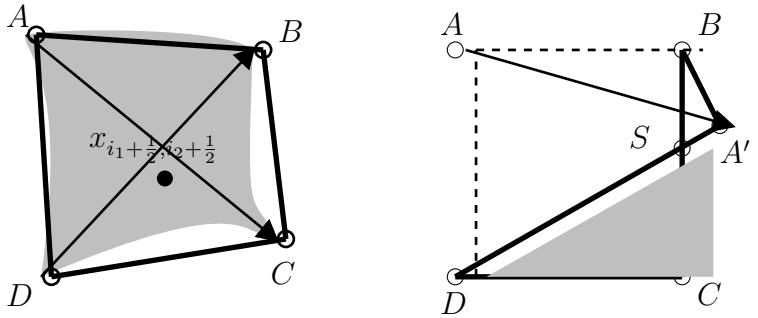

(a) nodal ∘ and cell centered • grid (b) twist

$$
\begin{aligned}
A &= [x^1_{i_1,i_2} + u^1_{i_1,i_2},\ x^2_{i_1,i_2} + u^2_{i_1,i_2}],\\
B &= [x^1_{i_1,i_2+1} + u^1_{i_1,i_2+1},\ x^2_{i_1,i_2+1} + u^2_{i_1,i_2+1}],\\
C &= [x^1_{i_1+1,i_2+1} + u^1_{i_1+1,i_2+1},\ x^2_{i_1+1,i_2+1} + u^2_{i_1+1,i_2+1}],\\
D &= [x^1_{i_1+1,i_2} + u^1_{i_1+1,i_2}\ , x^2_{i_1+1,i_2} + u^2_{i_1+1,i_2}].
\end{aligned}
$$

Figure 1: Cell center $x_{i_1+\frac{1}{2},i_2+\frac{1}{2}}$ and transformed cell for $d = 2$.

(cf. Figure 1(a))

$$\mathrm{vol}(v, \boldsymbol{\varphi}) = \int_{\boldsymbol{\varphi}(v)} d\boldsymbol{y} \approx \tfrac{1}{2}(C - A) \times (B - D).$$

Here, for ease of presentation, we drop the dependency on the location $i_1 + 1/2, i_2 + 1/2$. This approximation is motivated by the fact that for smooth and small deformations, the volume of the deformed box (gray area in Figure 1(a)) can be approximated by order h^2 by the volume of the box spanned by the deformed corners (surrounded by the bold lines in Figure 1(a)).

For volume preserving registration, this discretization is consistent. However, if we allow the transformation to shrink or enlarge volume in a certain amount, the transformation is allowed to be much more irregular and this discretization can cause difficulties. An intuitive example is illustrated in Figure 1(b). Here, the top left point A moved by more than the box diameter to A'. This movement results in a "*twist*" of the box. This twist can not be observed by measuring the volume of the box: The volume is given by the volume of the triangle (DCS) minus the volume of the triangle $(A'BS)$.

Particular, if A' is close to B, the volume may even be large although a twist has occurred. In fact, the situation is unstable since an arbitrarily small shift of A' can result in a twists. Therefore, box based volume measures do not detect twists.

It is obvious, that a twisted grid is physically incorrect (at least for all applications we are aware off). We therefore propose a different discrete measure that can detect twists of the grid. This new measure is crucial if we move from volume preservation equalities to volume constraining inequalities. As in [14], we assume that the transformation is smooth enough such that a deformed box can be approximated by the box spanned by the deformed corners.

Based on the previous considerations, we based our discretization on a triangulation. If a triangle (or tetrahedron in 3D) twists, its volume becomes negative. Therefore, to prevent twists and singular Jacobians, one has to consider a discretization based on triangulation. Since a triangle cannot twist without its volume to change sign, such a discretization is consistent even in cases of large deformations.

2.2.1 Triangulation in 2D

Every box is divided into two triangles $T^{1,2}$; cf. Figure 2(a). We compute the volume of both triangles separately,

$$\begin{aligned}
\frac{2}{h_1h_2}V^1_{i_1+\frac{1}{2},i_2+\frac{1}{2}} &= \tfrac{1}{h_1h_2}(C-B)\times(B-A)\\
&= 1+\tfrac{u^1_{i_1+1,i_2+1}-u^1_{i_1,i_2+1}}{h_1}+\tfrac{u^2_{i_1,i_2+1}-u^2_{i_1,i_2}}{h_2}\\
&\quad+\tfrac{u^1_{i_1+1,i_2+1}-u^1_{i_1,i_2+1}}{h_1}\tfrac{u^2_{i_1,i_2+1}-u^2_{i_1,i_2}}{h_2}-\tfrac{u^1_{i_1,i_2+1}-u^1_{i_1,i_2}}{h_2}\tfrac{u^2_{i_1+1,i_2+1}-u^2_{i_1,i_2+1}}{h_1}\\
\frac{2}{h_1h_2}V^2_{i_1+\frac{1}{2},i_2+\frac{1}{2}} &= \tfrac{1}{h_1h_2}(D-A)\times(C-D)\\
&= 1+\tfrac{u^1_{i_1+1,i_2}-u^1_{i_1,i_2+}}{h_1}+\tfrac{u^2_{i_i+1,i_2+1}-u^2_{i_1+1,i_2}}{h_2}\\
&\quad+\tfrac{u^1_{i_1+1,i_2}-u^1_{i_1,i_2}}{h_1}\tfrac{u^2_{i_1+1,i_2+1}-u^2_{i_1+1,i_2}}{h_2}-\tfrac{u^1_{i_1+1,i_2+1}-u^1_{i_1+1,i_2}}{h_2}\tfrac{u^2_{i_1+1,i_2}-u^2_{i_1,i_2}}{h_1}
\end{aligned}$$

and obtain $2m_1m_2$ inequality constraints for $d=2$.

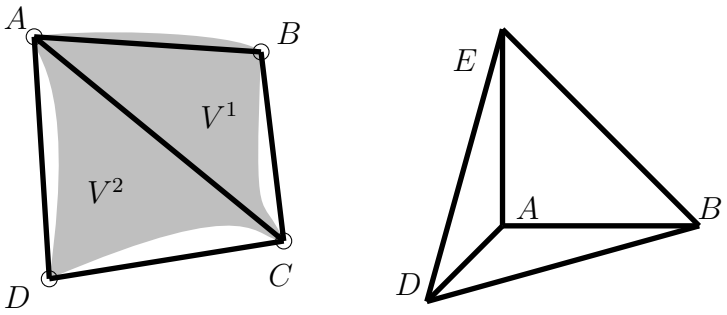

(a) 2D with two triangles (b) 3D with one tetrahedron

Figure 2: Triangulation of a pixel (a) or voxel (b).

2.2.2 Triangulation in 3D

In 3D, the computation becomes more involved. The displacement field $\boldsymbol{u}$ is discretized at the corner of each voxel. Once again, a discretization based on the deformed voxel cannot detect twists. Therefore, every voxel is divided into six tetrahedrons; cf. Figure 2(b). The volume of a tetrahedron with corners A, B, D, and E is given by

$$V_{i+\frac{1}{2}}^{A,B,D,E} \;=\; \tfrac{1}{6}\det(D-A,\ B-A,\ E-A).$$

Similar formula are used for the remaining five tetrahedrons. Hence, we end up with $6m_1m_2m_3$ inequality constraints for $d=3$.

Remark 1 *The number of constraints can be larger than the number of unknowns. Particularly in 3D, we have roughly $3m_1m_2m_3$ unknowns and $6m_1m_2m_3$ constraints. However, for $\kappa_m < 1 < \kappa_M$ the constraints are consistent, i.e., there exists a non-trivial transformation satisfying the constraints.*

2.3 Discretizing $\mathcal{S}$

Since many regularizers are phrased in terms of the more complex differential operators gradient ∇ and divergence $\nabla\cdot$, we introduce the notation ∇^h and $\nabla^h\cdot$ for the discrete analogs,

$$\nabla^h u^j = (\partial_1^h u^j, ..., \partial_d^h u^j)^\top, \tag{7}$$

where with $i=(i_1,\dots,i_d)$ and $e_k \in \mathbb{R}^d$ the kth unit vector,

$$(\partial_k^h u^j)_{i+\frac{1}{2}e_k} := \frac{1}{h_k}(u^j_{i+e_k} - u^j_i), \quad j,k=1,\dots,d.$$

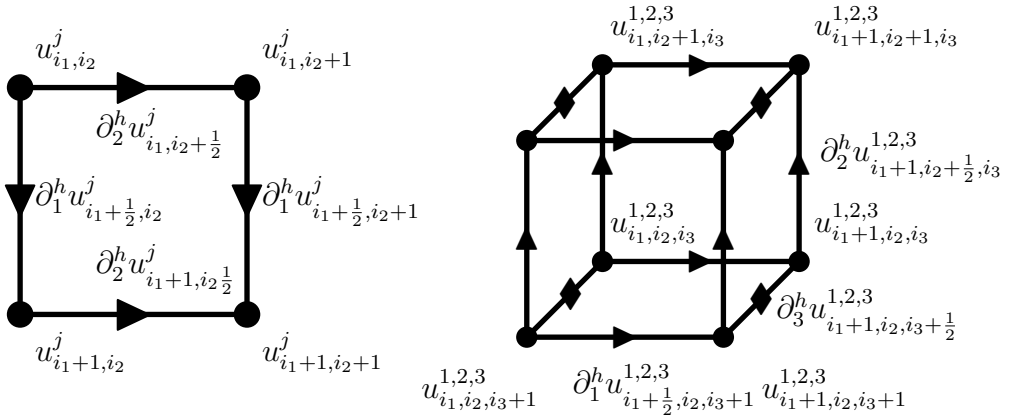

(a) 2D (b) 3D

Figure 3: Edge staggered grids for 2D (a) and 3D (b).

The unknowns u^j are discretized on the nodal grid, whereas the derivatives are discretized on face staggered grids; cf. Figure 3.

Note that the partial derivatives of u^j are located at different positions. Thus, for the divergence $\nabla^h\cdot$, we average to the cell centered grid. Particularly, for $d = 3$, we end up with

$$
\begin{aligned}
&(4\,\nabla^h\cdot\,\boldsymbol{u})_{i_1+\frac{1}{2},i_2+\frac{1}{2},i_3\frac{1}{2}} = \qquad (8)\\
&\quad (\partial_1^h u^1)_{i_1+\frac{1}{2},i_2,i_3} + (\partial_1^h u^1)_{i_1+\frac{1}{2},i_2,i_3+1} + (\partial_1^h u^1)_{i_1+\frac{1}{2},i_2+1,i_3} + (\partial_1^h u^1)_{i_1+\frac{1}{2},i_2+1,i_3+1}\\
+\;& (\partial_1^h u^1)_{i_1,i_2+\frac{1}{2},i_3} + (\partial_1^h u^1)_{i_1,i_2+\frac{1}{2},i_3+1} + (\partial_1^h u^1)_{i_1+1,i_2+\frac{1}{2},i_3} + (\partial_1^h u^1)_{i_1+1,i_2+\frac{1}{2},i_3+1}\\
+\;& (\partial_1^h u^1)_{i_1,i_2,i_3+\frac{1}{2}} + (\partial_1^h u^1)_{i_1,i_2+1,i_3+\frac{1}{2}} + (\partial_1^h u^1)_{i_1+1,i_2,i_3+\frac{1}{2}} + (\partial_1^h u^1)_{i_1+1,i_2+1,i_3+\frac{1}{2}}
\end{aligned}
$$

Based on this discrete analogs, the elastic potential (3) is discretized by

$$S^h(\boldsymbol{u}) = \|\boldsymbol{B}\boldsymbol{u}\|_2^2 := \tfrac{\lambda+\omega}{2}\|\,\nabla^h\cdot\,\boldsymbol{u}\|^2 + \tfrac{\omega}{2}\textstyle\sum_{j=1}^d \|\,\nabla^h\,\boldsymbol{u}^j\|^2. \qquad (9)$$

In our registration process we like to exploit fast optimization techniques. Therefore, we also have to consider the derivative

$$S^h_{\boldsymbol{u}}(\boldsymbol{u}) = (\lambda+\omega)(\,\nabla^h\cdot\,)^\top\,\nabla^h\cdot\,\boldsymbol{u} - \omega\Delta^h\boldsymbol{u} =: \boldsymbol{A}\boldsymbol{u}, \qquad (10)$$

where $\boldsymbol{A}$ is a discretization of the well-known Navier-Lamé operator and Δ^h is the usual seven points discrete vector Laplacian.

Remark 2 *With the nodal discretization of u^j, the divergence operator is not the discrete adjoint (transpose) of the gradient operator. The transpose of the discrete divergence $(\ \nabla^h\cdot\ )^\top$ is still an $\mathcal{O}(h^2)$ approximation to the gradient but the resulting discrete operator $\boldsymbol{A}^h$ is not h-elliptic. Therefore, we take advantage of the above consistent and h-elliptic discretization.*

2.4 Discretizing T and $\mathcal{D}$

Since we are heading for a continuously differentiable objective function, we need to have an appropriate image model. Particularly, d-linear image approximations can not be used. As discussed in [14], we use a smoothing B-spline approximation to the given discrete date. The smoothing parameter is used for noise reduction and automatically chosen according to the Generalized Cross Validation method (GCV) [10]. For data interpolation using B-splines see [21]. Since the grid is regular, the spline coefficients can be computed efficiently using a discrete cosine transform. The continuous smooth approximation is denoted by T^{spline}.

Given the nodal grid representation of $\boldsymbol{u}$ we use averaging operators P for the transfer to the cell centered positions,

$$(Pu^j)_{i_1+\frac{1}{2},\dots,i_d+\frac{1}{2}} := 2^{-d}\sum_{k_1,\dots,k_d=0,1} u^j_{i_1+k_1,\dots,i_d+k_d},$$

the discretization of T is given by

$$T(\boldsymbol{u}) := T^{\text{spline}}(x^1+Pu^1,\dots,x^d+Pu^d),$$

and the Jacobian $T_{\boldsymbol{u}}$ of T by

$$T_{\boldsymbol{u}} := \frac{\partial T}{\partial \boldsymbol{u}}(\boldsymbol{u}) = \Big(\operatorname{diag}(P^\top\partial_1 T),\dots\operatorname{diag}(P^\top\partial_d T)\Big), \tag{11}$$

where the partial derivatives $\partial_j T$ are evaluated at the spatial positions $(x^1+Pu^1,\dots,x^d+Pu^d)$. Note that using a spline approximation for T, $T_{\boldsymbol{u}}$ is a sparse matrix with only eight non-zero diagonals.

Our discretization of the SSD distance measure (2) is straightforward,

$$D(\boldsymbol{u}) := \tfrac{1}{2}\|T(\boldsymbol{u})-R\|_2^2 \quad\text{and thus}\quad D_{\boldsymbol{u}}(\boldsymbol{u}) = T_{\boldsymbol{u}}(\boldsymbol{u})^\top(T(\boldsymbol{u})-R).$$

3 Solution of the optimization problem

3.1 The Log Barrier Framework

To solve the discretized optimization problem (1), we use a variant of a log-barrier method [8, 18]. Rather than solving the constrained optimization problem, we replace it by a sequence of unconstrained optimization problems J^μ where μ is gradually decreased. Here,

$$J^\mu = D(R, T(\boldsymbol{u})) + \alpha S(\boldsymbol{u}) - \mu e^\top \Big(\log(\boldsymbol{C}(\boldsymbol{u}) - \kappa_m) + \log(\kappa_M - \boldsymbol{C}(\boldsymbol{u}) \Big), \quad (12)$$

where μ is the barrier parameter and $e = (1, \ldots, 1)^\top$. In classical optimization algorithms that barrier parameter is chosen large at first and slowly reduced to zero. The algorithm for solving the optimization problem is summarized in Algorithm 1.

Algorithm 1 Constrained Image Registration: $\boldsymbol{u} \leftarrow \mathrm{CIR}(\boldsymbol{u}, \alpha)$;

Choose $\mu > 0$, $\quad 0 < \gamma < 1$.
while true **do**
1 Correction: Approximately minimize J^μ for $\boldsymbol{u}^\mu$ starting at $\boldsymbol{u}$; cf. (15).
2 Set $\mu \leftarrow \gamma\mu$.
3 Prediction: Calculate $\boldsymbol{s}^\mu$ and update $\boldsymbol{u} \leftarrow \boldsymbol{u}^\mu + \boldsymbol{s}^\mu$; cf. Sec. 3.3.
4 Check for convergence.
end while

In the following we discuss a few non-trivial steps in this algorithm. Firstly, we discuss the minimization of J^μ for a particular choice of μ (Step 1). Secondly, we discuss the correction to the approximate solution after μ is decreased (Step 3).

3.2 Solving the optimization problem for a fixed μ

To solve the optimization problem for a fixed μ we use a variant of Gauss-Newton's method; see, e.g., [18]. The gradient of the objective function is

$$\nabla_{\boldsymbol{u}} J^\mu = T_{\boldsymbol{u}}(\boldsymbol{u})^\top (T(\boldsymbol{u}) - R) + \alpha \boldsymbol{A}\boldsymbol{u} - \mu \boldsymbol{C}_{\boldsymbol{u}}^\top \left(\frac{1}{\boldsymbol{C}(\boldsymbol{u}) - \kappa_m} - \frac{1}{\kappa_M - \boldsymbol{C}(\boldsymbol{u})} \right). \quad (13)$$

We approximate the Hessian of the objective function J^μ using the Gauss-Newton approximation

$$\boldsymbol{H}^\mu = T_{\boldsymbol{u}}^\top T_{\boldsymbol{u}} + \alpha \boldsymbol{A} + \mu \boldsymbol{C}_{\boldsymbol{u}}^\top \mathrm{diag}\left(\frac{1}{(\boldsymbol{C}(\boldsymbol{u}) - \kappa_m)^2} - \frac{1}{(\kappa_M - \boldsymbol{C}(\boldsymbol{u}))^2} \right) \boldsymbol{C}_{\boldsymbol{u}}. \quad (14)$$

The Gauss-Newton direction is then calculated by solving the linear system

$$\boldsymbol{H}^\mu \boldsymbol{s} = -\nabla_{\boldsymbol{u}} J^\mu. \quad (15)$$

For the solution of the system (15) we use the Conjugate Gradient method with a multigrid V-cycle preconditioner [20].

3.3 Prediction step

The optimization problem (12) is approximately solved for a particular choice of μ. To obtain an approximate solution of the underlying constrained optimization problem (6), μ has to be decreased to zero.

It is well known [18] that if one initializes every optimization problem with the numerical solution obtained for the previous μ and wants to ensure fast converges of the algorithm, then small changes of μ are required. In order to speedup the reduction process for μ, we use a predictor-corrector approach [18].

Since we relax μ by a fraction, it is natural to change to a logarithmic scale. Replacing μ by e^{-t}, for the solution $\boldsymbol{u}(t)$ for a fixed t we have $\nabla_{\boldsymbol{u}} J(\boldsymbol{u}(t), t) = 0$ and using the chain rule it follows that

$$e^{-t} \boldsymbol{C}_{\boldsymbol{u}}^\top \left(\frac{1}{\boldsymbol{C}(\boldsymbol{u}) - \kappa_m} - \frac{1}{\kappa_M - \boldsymbol{C}(\boldsymbol{u})} \right) + \boldsymbol{H}^\mu \partial_t \boldsymbol{u} = 0. \quad (16)$$

Discretizing equation (16) we obtain the following rule for the update $\boldsymbol{s} = \boldsymbol{u}(t_{k+1}) - \boldsymbol{u}(t_k)$,

$$e^{-t_k} \boldsymbol{C}_{\boldsymbol{u}}^\top \left(\frac{1}{\boldsymbol{C}(\boldsymbol{u}_k) - \kappa_m} - \frac{1}{\kappa_M - \boldsymbol{C}(\boldsymbol{u}_k)} \right) + \frac{1}{t_{k+1} - t_k} \boldsymbol{H}^\mu \boldsymbol{s} = 0. \quad (17)$$

This results in a linear system with the same Hessian matrix as in (15) but with a different right hand side. This fact can be exploited if inexact factorization is used as a preconditioner or a smoother.

4 Numerical Experiments

To demonstrate the effectiveness of our method we use SPECT data from the Emory Hospital in Atlanta. The data are 3D images of $128 \times 128 \times 36$ voxels of the heart of a patient through different phases of the cardiac cycle. The overall goal of the clinical study is to evaluate and track cardiac displacement from systolic to diastolic.

It's obvious, that the deformation of the heart is highly non-linear and thus presents a great challenge to any registration routine. Applying an unconstrained registration algorithm results in a dilemma: Either the image distance stays large or the grid folds. For this application, volume or mass preserving equality constraints as consider in [23, 14] are physiologically inappropriate. On the other hand, it is common to assume that the heart does not change its volume too much [7]. We therefore apply our volume constrained VCIR approach.

We preform two sets of experiments. In the first set of experiments we use a 2D slice which allows us to visualize some of the concepts discussed above in more detail. In the second set of our experiments we process the 3D data and demonstrate that the concepts carry over to 3D.

We experiment with different regularization parameters α. For both, the 2D and 3D example, we used $\kappa_m = \kappa_M^{-1} \equiv 0.3$. All computations are performed using MATLAB.

4.1 2D cardiac example

Figure 4 shows two SPECT images of a heart in systolic and diastolic phases. For the registration of the two phases, we attempt to use two different approaches. First, we run an unconstrained code that uses elastic registration [13]. Both algorithms are stopped as soon as $\|\boldsymbol{u}^{k+1} - \boldsymbol{u}^k\| < \text{voxelsize}/10$. In addition, the unconstrained algorithm is also stopped if the minimum value of the determinant of the Jacobian is less then 10^{-5}.

Using the unconstrained code and starting with $\alpha = 10^{-2}$ we obtained a minimal value of the Jacobian of 0.28 which is close to the lower bound $\kappa_m = 0.3$. The SSD reduction was 66%,

$$\text{red} := D(\boldsymbol{u}^{\text{final}})/D(\boldsymbol{u}^0 = 0).$$

In order to decrease the image difference further, we then rerun the unconstrained algorithm with a smaller regularization parameter $\alpha = 10^{-3}$ and

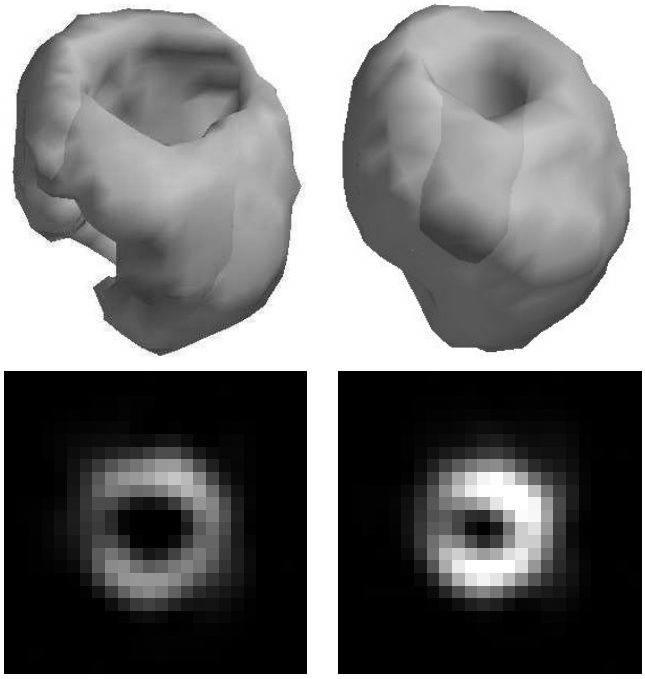

Figure 4: Heart in two phases: diastolic (reference, left) and systolic (template, right); top: 3D visualization, bottom: two 2D slices.

obtained a 47% reduction; the algorithm terminates because of the Jacobian becoming close to singular. However, as can be seen in Figure 5, the grid starts to degenerate, i.e. the minimum value of the determinant of the Jacobian is less then 10^{-5}.

We than run our VCIR algorithm. For $\alpha = 10^{-2}$, we obtain the same SSD reduction to 66% but here the minimal value of the determinant of the Jacobian is 0.32. Obviously, the results of the constrained and unconstrained approaches are nearly indistinguishable, both in terms of the image difference as well as in terms of the displacements.

Spectacular results are obtained for $\alpha = 10^{-3}$. Here, the SSD reduction of the constrained approach is 43%, which is even better than the reduction of the unconstrained approach with the same α and much better than the reduction for $\alpha = 10^{-2}$. Moreover, in contrast to the unconstrained approach, the minimum value of the determinant of the Jacobian was $0.31 \geq \kappa_m$. This is of course not a surprise since our algorithm is designed to fulfil the constraints. The regularity of the transformations is visualized in Figure 6.

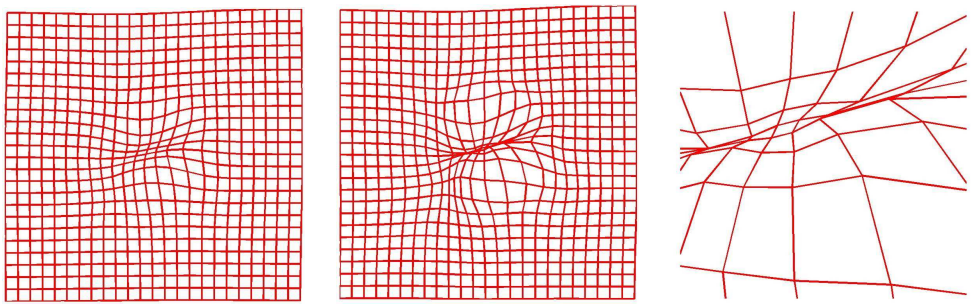

Figure 5: Grids obtained from the unconstrained algorithm; left: $\alpha = 10^{-2}$, red = 66%, $\min \boldsymbol{C} = 0.28$, middle: $\alpha = 10^{-3}$, red = 47%, $\min \boldsymbol{C} < 10^{-5}$, right: detail of the grid for 10^{-3}.

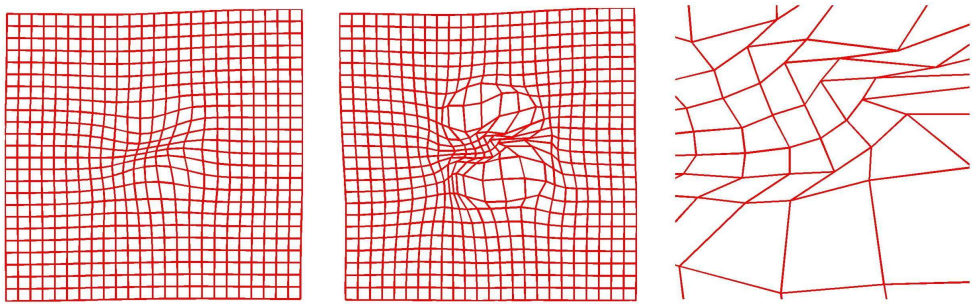

Figure 6: Grids obtained from VCIR; left: $\alpha = 10^{-2}$, red = 66%, $\min \boldsymbol{C} = 0.32$, middle: $\alpha = 10^{-3}$, red = 43%, $\min \boldsymbol{C} = 0.31$, right: detail of the grid for 10^{-3}.

4.2 3D cardiac example

In this example we test our algorithm on the 3D data; cf. Figure 4. The results for various regularization parameters, the unconstrained and the VCIR algorithms are summarized in Table 1. Note that for the constrained approach, we are able to reduced the regularization parameter down to $\alpha = 10^{-7}$.

Figure 7 displays a visualization of the volume change for the unconstrained and VCIR approaches. We observed that the unconstrained approach does change the volume significantly particularly in a small area next to the heart wall (see zoom of slice 10). In an integral measure this small spot is almost undetectable. However, for this particular medial application the volume change of the heart wall is crucial and therefore our VCIR approach

Table 1: Results for 3D registration various regularization parameters, unconstrained and VCIR approaches.

α	SSD red	$\min \boldsymbol{C}$
unconstrained		
10^{-4}	75%	$0.29 \cdot 10^{0}$
10^{-5}	68%	$0.14 \cdot 10^{-2}$
10^{-6}	52%	< 0
VCIR		
10^{-7}	47%	$0.30 \cdot 10^{0}$

is clearly preferable.

5 Conclusions

Regularity of the displacement field is a very important feature in image registration. In most applications, an extreme expansion or shrinkage and particularly folding of objects is non-physical. Therefore, registration algorithms should not produce such solutions. Rigridding strategies as introduced in [5] could be used. However, there is no physical or mathematical justification for this procedure. There also exists approaches aiming for volume preservation [19, 6, 23, 14]. But volume preservation is too restrictive for a wide range of applications. For example, for problems arising in cardiac imaging it is well-known that the volume of the heart does change within some bands and a volume preservation is not physically meaningful.

We present a novel registration approach, where the main idea is to add additional explicit volume inequality constraints. Thus, our mathematical model takes the constraints into and prevents large changes of the volume and folding. An analytic solution of the registration problem is not known and therefore numerical schemes have to be applied. A major concern of this paper is to point out that a proper treatment of volume inequality constraints is not straightforward. Here we suggest a nodal grid based discretization for the unknowns and evaluate the constraints on a triangulation.

Our numerical results indicate that there is a tremendous difference between the unconstrained and inequality constrained approach.

Acknowledgements: We are indebted to Prof. Tracy Faber from Emory

Hospital Atlanta and Petunia for providing exciting data.

References

[1] U. Ascher and E. Haber. Grid refinement and scaling for distributed parameter estimation problems. *Inverse Problems*, 17:571–590, 2001.

[2] U. Ascher, E. Haber, and H. Haung. On effective methods for implicit piecewise smooth surface recovery. Technical Report TR-2004-017-A, Dept of Math. & CS., Emory University, Atlanta GA 30322, Jun 2004.

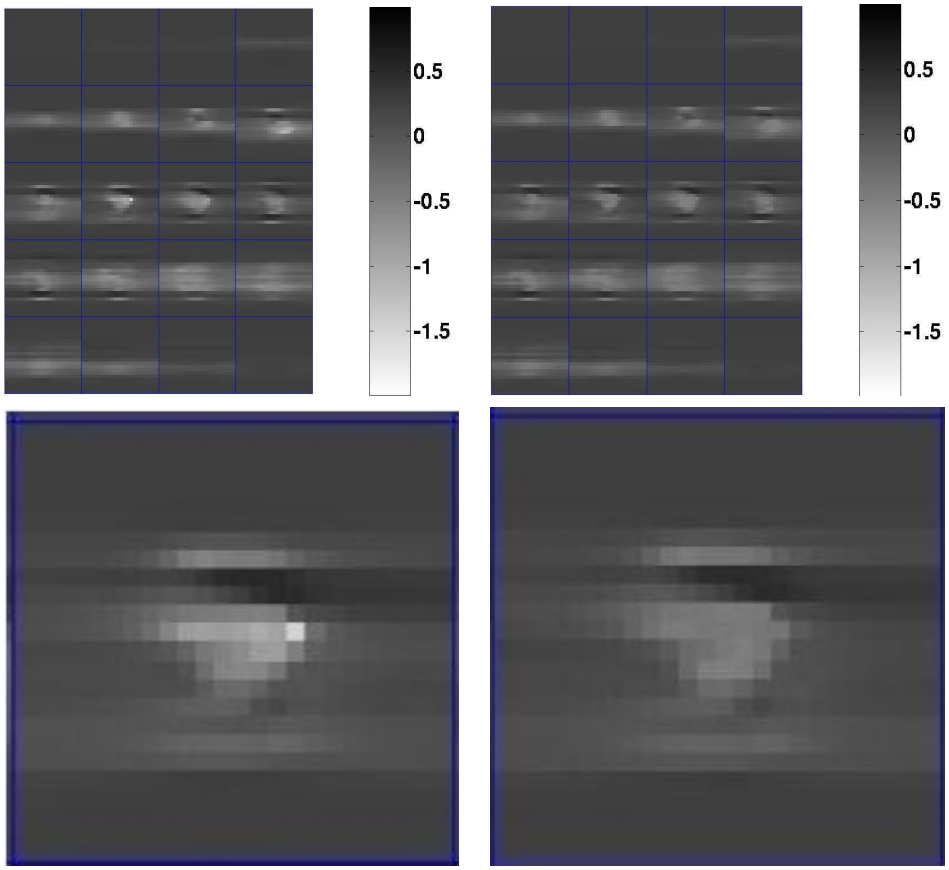

Figure 7: Logarithmic map of the determinant of the Jacobian, top: for all 20 slices; bottom slice 10, left: unconstraint with $\alpha = 10^{-5}$ and $\min \boldsymbol{C} = 0.014$, right: VCIR with $\alpha = 10^{-7}$ and $\min \boldsymbol{C} = 0.31$.

[3] C. Broit. *Optimal Registration of Deformed Images*. PhD thesis, Computer and Information Science, University of Pensylvania, 1981.

[4] G. E. Christensen. Consistent linear-elastic transformations for image matching. In *Information Processing in Medical Imaging*, volume 1613 of *LCNS*, pages 224–237. Springer-Verlag, 1999.

[5] G. E. Christensen. *Deformable Shape Models for Anatomy*. PhD thesis, Sever Institute of Technology, Washington University, 1994.

[6] M. Droske and M. Rumpf. A variational approach to non-rigid morphological registration. *SIAM Appl. Math.*, 64(2):668–687, 2004.

[7] T. Faber. *Private communications*. Emory Hospital, Atlanta, 2005.

[8] A. V. Fiacco and G. P. McCormick. Nonlinear Programming: Sequential Unconstrained Minimization Techniques (Classics in Applied Mathemaitcs) (Paperback) *SIAM Philadelphia*, 1991.

[9] J. M. Fitzpatrick, D. L. G. Hill, and C. R. Maurer Jr. Image registration. In M. Sonka and J. M. Fitzpatrick, editors, *Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis*, pages 447–513. SPIE, 2000.

[10] G. Golub, M. Heath, and G. Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21:215–223, 1979.

[11] E. Haber. A multilevel, level-set method for optimizing eigenvalues in shape design problems. *JCP*, 115:1–15, 2004.

[12] E. Haber and J. Modersitzki. Intensity gradient based registration and fusion of multi-modal images. Technical Report TR-2004-027-A, Department of Mathematics and Computer Science, Emory University, Atlanta GA 30322, Jun 2004. Submitted to IEEE TMI.

[13] E. Haber and J. Modersitzki. A multilevel method for image registration. Technical Report TR-2004-005-A, Department of Mathematics and Computer Science, Emory University, Atlanta GA 30322, May 2004. Accepted for publication in SIAM J. of Scientific Computing.

[14] E. Haber and J. Modersitzki. Numerical methods for volume preserving image registration. *Inverse Problems, Institute of Physics Publishing*, 20(5):1621–1638, 2004.

[15] J. B. A. Maintz, J. P. W. Pluim, and M. A. Viergever. Image registration by maximization of combined mutual information and gradient information. *IEEE TMI*, 19(8):809–814, 2000.

[16] J. B. A. Maintz and M. A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.

[17] J. Modersitzki. *Numerical Methods for Image Registration*. Oxford University Press, 2004.

[18] J. Nocedal and S. J. Wright. *Numerical optimization*. Springer, New York, 1999.

[19] T. Rohlfing, C. R. Maurer, Jr., D. A. Bluemke, and M. A. Jacobs. Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE Transactions on Medical Imaging*, 22(6):730–741, 2003.

[20] U. Trottenberg, C. Oosterlee, and A. Schüller. *Multigrid*. Academic Press, 2001.

[21] G. Wahba. *Spline Models for Observational Data*. SIAM, Philadelphia, 1990.

[22] W. M. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–51, 1996.

[23] L. Zhu, S. Haker, and A. Tannenbaum. Area preserving mappings for the visulization of medical structures. *MICCAI 2003*, pages 277–284, 2003.

[24] B. Zitová and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, 2003.