

COFIR: Coarse and Fine Image Registration

Jan Modersitzki* Eldad Haber*

May 28, 2004

Abstract

A new approach to image registration is proposed. The approach enables the proper segregation of the transformation space into a low dimensional so-called *coarse* space and a remaining non-linear so-called *fine* part. The segregation concept is combined with a multilevel approach, where the two parts are treated simultaneously.

Modeling and computation issues are discussed. Finally, we present a highly non-linear real life example.

1 Introduction

Image registration is one of today's challenging image processing problems. The problem can be formulated as follows: Given two images, find a "*reasonable*" transformation such that a transformed version of the so-called template image becomes "*similar*" to the so-called reference image. Image registration is applied whenever images resulting from different times, devices, and/or perspectives need to be compared or integrated; see, e.g. [26, 20] and references therein.

A registration procedure is typically based on two main building blocks. The first one is a distance measure. The distance measure quantifies the meaning of similarity or proximity of images. A distance measure can be based on image features (e.g., moments [1], landmarks [23, 20], or markers [19]) on image intensities (e.g. L_2 -norm or sum of squared differences [7],

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

correlation, mutual information [10, 25]), on image surfaces [3, 4], on level sets [11], or on combinations hereof [8, 13]. For an overview and comparisons, see also [22, 17, 20].

The second building block is a regularizing term. Since image registration is an ill-posed problem, regularization is inevitable and becomes a central topic [20]. Typical regularization techniques are a restriction to a low dimensional transformation space (e.g., the space of rigid or affine linear transformations) or an explicit regularization of the problem. The regularizer of the latter approach is typically based on a bilinear form

$$\mathcal{S}[\mathbf{u}] = \frac{1}{2} \int_{\Omega} \langle \mathcal{B}\mathbf{u}, \mathcal{B}\mathbf{u} \rangle dx, \quad (1)$$

where \mathcal{B} is a partial differential operator. This includes well-known regularizers like, for example, the elastic [6, 2], diffusion [24, 12], or curvature [14], and, in a qualified sense, also the fluid registration [9, 5]; see also [20] for an overview.

In many applications one has to deal with large linear but also important non-linear deformations of the displayed object. Typical examples arising from medical applications include the registration of a histological serial sectioning or the registration of two images from a joint taken at different positions; see also Fig. 1. Therefore one would prefer a registration scheme which handles both the linear (more general: the *coarse*) and the remaining non-linear (more general: the *fine*) part simultaneously. However, as it is well-known to practitioners, existing registration schemes can not handle this goal in an appropriate way, and in some cases such registration does not converge to an acceptable solution.

In the literature either of the following two ways is used to bypass this problem. The first one is to assume that the overall registration problem decouples into in a linear and non-linear part and to solve for the two parts separately and sequentially. However, though this approach seems to work quite well in practise, the decoupling assumption is artificial and it is an open question whether this assumption is true or not. The second remedy is to design a regularizer that does not penalize the affine linear part. In [14] it is suggested to used $\mathcal{B} = \Delta$, since a second order derivative operator does not penalize affine linear transformation. However, not all applications allow for an arbitrary regularization. Higher order derivative based operators result in smoother transformations which may not be physically correct. Finally,

from an implementation point of view, boundary conditions and the solution of the resulting system of equations become more troublesome.

Therefore we proposed a new way of dealing with these difficulties. We split the transformation space into two disjoint subspaces, a *coarse* space and a *fine* space. The coarse space is a low dimensional linear space and the fine space is the remaining orthogonal complement. Only the components of the transformation belonging to the fine space are regularized. We call this novel COarse and Fine Image Registration scheme COFIR. In this new approach, the coarse and the fine spaces are guaranteed to be disjoint and therefore the above mentioned decoupling assumption holds by construction.

Different applications may demand for different choices of the coarse space \mathcal{C} . For example, in studying tumor growth, one may want to consider only volume preserving transformations and thus choose \mathcal{C} to be the set of rigid transformations. In our mathematical formulation we therefore consider a general coarse space but concentrate on the space of affine linear transformations in our examples.

The paper is organized as follows. In Section 2 we introduce the continuous mathematical framework which is of particular interest for our multilevel approach and in Section 3 we discuss discretization issues. The optimization schemes and the implementation of COFIR are discussed in Section 4. Finally, in Section 5 we present results for a highly non-linear real life application; see also [18].

2 The mathematical setting

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. To be precise, we interpolate or approximate the given discrete data by a smooth spline functions R and T , respectively. Thus, for any spatial position $\mathbf{x} = (x_1, \dots, x_d)$, $T(\mathbf{x})$ gives the corresponding gray value. 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” deformation $\mathbf{u} = (u^1, \dots, u^d)^T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that the “distance” between the reference image R and the deformed template image $T(\mathbf{x} + \mathbf{u}(\mathbf{x}))$ becomes small. As explained in the introduction, we divide the space of feasible transformations into two disjoint subspaces, a coarse space \mathcal{C} and a fine space \mathcal{F} . The space \mathcal{C} is an m dimensional linear space spanned by some basis functions c_j , $j = 1, \dots, m$, such that a generic

element reads

$$\mathbf{u}^{\mathcal{C}} = \sum_{j=1}^m z_j \mathbf{c}_j(\mathbf{x}) =: \mathbf{C}(\mathbf{x})\mathbf{z}, \quad \mathbf{z} \in \mathbb{R}^m. \quad (2)$$

Example 1 *Probably the most important example is when \mathcal{C} is the space of affine linear transformations. For spatial dimension $d = 2$ we have*

$$\mathbf{u}^{\mathcal{C}} = \begin{pmatrix} z_1 + z_2 x_1 + z_3 x_2 \\ z_4 + z_5 x_1 + z_6 x_2 \end{pmatrix}$$

and for $d = 3$,

$$\mathbf{u}^{\mathcal{C}} = \begin{pmatrix} z_1 + z_2 x_1 + z_3 x_2 + z_4 x_3 \\ z_5 + z_6 x_1 + z_7 x_2 + z_8 x_3 \\ z_9 + z_{10} x_1 + z_{11} x_2 + z_{12} x_3 \end{pmatrix}.$$

To find $\mathbf{u} = \mathbf{u}^{\mathcal{C}} + \mathbf{u}^{\mathcal{F}}$, we formulate the following optimization problem

$$\text{minimize} \quad \mathcal{D}[R, T; \mathbf{u}^{\mathcal{C}} + \mathbf{u}^{\mathcal{F}}] + \alpha \mathcal{S}[\mathbf{u}^{\mathcal{F}}] \quad (3a)$$

$$\text{subject to} \quad \langle \mathbf{u}^{\mathcal{F}}, \mathbf{c}_j \rangle = 0 \quad j = 1, \dots, m, \quad (3b)$$

where \mathcal{D} is some distance measure, \mathcal{S} is some regularization term, and $\alpha > 0$ is a regularization parameter that compromises between similarity and regularity. In contrast to previous approaches the transformation has been divided into two parts and only the fine part is regularized. Moreover and most importantly, the fine part is constrained to be orthogonal to the coarse part and therefore unlike previous work, this splitting is disjoint.

Any differentiable choice of building blocks \mathcal{D} and \mathcal{S} is feasible. However, for ease of presentation, we restrict to the sum of squared differences (SSD) and the elastic regularizer,

$$\mathcal{D}[R, T; \mathbf{u}] = \frac{1}{2} \int_{\Omega} \left(T(\mathbf{x} + \mathbf{u}(\mathbf{x})) - R(\mathbf{x}) \right)^2 dx \quad (4)$$

$$\mathcal{B}[\mathbf{u}] = [\sqrt{\mu} \nabla u_1, \dots, \sqrt{\mu} \nabla u_d, \sqrt{\lambda + \mu} \nabla \cdot \mathbf{u}]^{\top}, \quad (5)$$

$$\mathcal{A} := \mathcal{B}^* \mathcal{B} = \mu \Delta + (\lambda + \mu) \nabla \nabla \cdot, \quad (6)$$

where λ and μ are the so-called Lamé constants and natural boundary conditions are imposed; see, e.g., [20].

Remark 1 *Let \mathcal{P} denote the projector onto the orthogonal complement of \mathcal{C} . Hence, (3) can be replaced by*

$$\text{minimize} \quad \mathcal{D}[R, T; \mathbf{u}] + \alpha \mathcal{S}[\mathcal{P}\mathbf{u}]. \quad (7)$$

3 Discretization

The image registration problem 3 is discretized using staggered grids for the components of \mathbf{u} . Since the images may be noisy and derivatives are needed for efficient optimization procedures, we use a smoothing B-spline approximation R^{spline} and T^{spline} for the images R and T , respectively; see also [16]. In order to reduce the notational overhead, we denote discrete quantities with the the same symbol as their continuous analogs. The precise meaning will be clear from the context. For ease of presentation, we also only discuss the two-dimensional case.

Let x_j denote the knots of the j th staggered grid, u_j the values of the j th component of \mathbf{u} on x_j , and \mathbf{u} the collecting of all these values. Our discretization of T is given by

$$T(\mathbf{u}) := T^{\text{spline}}(x_1 + P_1 u_1, x_2 + P_2 u_2),$$

where P_j denotes an averaging from the j th staggered to the cell centered grid and we use a B-spline interpolation or approximation scheme; see [16] for details. We denote the Jacobian of T by

$$T_{\mathbf{u}} := \frac{\partial T}{\partial \mathbf{u}}(\mathbf{u}) = \left(\text{diag}(P_1^\top \partial_1 T), \text{diag}(P_2^\top \partial_2 T) \right), \quad (8)$$

where the partial derivatives $\partial_j T$ have to be evaluated at $(x_1 + P_1 u_1, x_2 + P_2 u_2)$. For a discretization of the SSD distance measure (4), we use

$$\mathbf{D}(\mathbf{u}) := \frac{1}{2} \|T(\mathbf{u}) - R\|_2^2 \quad \text{and} \quad \mathbf{D}_{\mathbf{u}}(\mathbf{u}) = T_{\mathbf{u}}(\mathbf{u})^\top (T(\mathbf{u}) - R)$$

and for the regularizer (1) we take

$$\mathbf{S}(\mathbf{u}) = \frac{1}{2} \mathbf{u}^\top \mathbf{A} \mathbf{u} \quad \text{and} \quad \mathbf{S}_{\mathbf{u}}(\mathbf{u}) = \mathbf{A} \mathbf{u},$$

where $\mathbf{A} = \mathbf{B}^\top \mathbf{B}$ is based on a finite difference based discretization of \mathcal{B} . For details concerning the discretization of the elastic regularizer (5) see [16].

4 Optimization scheme

In order to have an efficient optimization scheme we use a multilevel strategy (or grid continuation) where the problem is discretized and solved on a sequence of grids starting from a coarse grid. On level ℓ , the discretized analog

of the image registration problem (3) for $\mathbf{u}_\ell = \mathbf{u}_\ell^{\mathcal{F}} + \mathbf{C}_\ell \mathbf{z}_\ell$ reads:

$$\text{minimize} \quad D_\ell(\mathbf{u}_\ell^{\mathcal{F}} + \mathbf{C}_\ell \mathbf{z}_\ell) + \alpha \mathbf{S}_\ell(\mathbf{u}_\ell^{\mathcal{F}}) \quad (9a)$$

$$\text{subject to} \quad \mathbf{C}_\ell^\top \mathbf{u}_\ell^{\mathcal{F}} = \mathbf{0}, \quad (9b)$$

The discrete version \mathbf{C}_ℓ on the ℓ th level is computed directly by evaluating the basis functions c_j on the underlying grids.

On each level, we use an iteration alternating between the minimization with respect to \mathbf{z} and with respect to $\mathbf{u}^{\mathcal{F}}$. In the next subsections we describe the three parts of our algorithm: the minimization with respect to \mathbf{z} , the minimization with respect to $\mathbf{u}^{\mathcal{F}}$ and the multilevel approach. The algorithm is summarized in Alg. 1. For the discussion of the optimization schemes we consider a fixed level and therefore drop the level index in the description.

Algorithm 1 Multilevel COarse and Fine Image Registration:

$[\mathbf{u}^{\mathcal{F}}, \mathbf{z}] \leftarrow \text{COFIR}(R, T, \mathbf{C})$.

for $\ell = 1, \dots, \max_\ell$ **do**

 transfer the images to the level ℓ ;

 solve the coarse problem for \mathbf{z} with $\mathbf{u}_\ell^{\mathcal{F}}$ fixed;

 solve the fine problem for $\mathbf{u}_\ell^{\mathcal{F}}$ with \mathbf{z} fixed;

 prolongate $\mathbf{u}_\ell^{\mathcal{F}}$ to the finer level;

 set $\ell \leftarrow \ell + 1$;

end for

4.1 Optimizing the coarse part

Assuming we have an approximation to the fine part $\mathbf{u}^{\mathcal{F}}$ we compute a solution to the problem

$$\text{given } \mathbf{u}^{\mathcal{F}}, \mathbf{C}, \quad \text{minimize} \quad D^{\mathcal{C}}(\mathbf{z}) := D(\mathbf{u}^{\mathcal{F}} + \mathbf{C}\mathbf{z}), \quad (10)$$

we apply a Gauß-Newton scheme with an Armijo line search; see, e.g., [21]. Note that the computational expensive parts of this algorithm are the computation of $T(\mathbf{v})$ and $T_{\mathbf{u}}(\mathbf{v})$. Particularly the linear system in the Gauss-Newton approach are just m -by- m , where m is the dimension of the linear space \mathcal{C} , i.e., $m = 6$ for the 2D linear space.

4.2 Optimizing the fine part

Assuming we have an approximation to the coarse part, i.e., \mathbf{z} is given, we compute a solution to the problem

$$\begin{aligned} & \text{minimize} && \mathbf{D}^{\mathcal{F}}(\mathbf{u}^{\mathcal{F}}) := \mathbf{D}(\mathbf{u}^{\mathcal{F}} + \mathbf{C}\mathbf{z}) + \alpha\mathbf{S}(\mathbf{u}^{\mathcal{F}}) \\ & \text{subject to} && \mathbf{C}^{\top}\mathbf{u}^{\mathcal{F}} = \mathbf{0}. \end{aligned} \quad (11)$$

To solve the problem we consider the Lagrangian

$$L(\mathbf{u}^{\mathcal{F}}, \mathbf{p}) := \mathbf{D}^{\mathcal{F}}(\mathbf{u}^{\mathcal{F}}) + \mathbf{p}^{\top}\mathbf{C}^{\top}\mathbf{u}^{\mathcal{F}},$$

where \mathbf{p} denotes a vector of Lagrange multipliers. The necessary conditions for a minimizer $\mathbf{u}^{\mathcal{F}}$ of (11) read

$$\begin{pmatrix} \mathbf{D}_{\mathbf{u}}^{\mathcal{F}}(\mathbf{u}^{\mathcal{F}} + \mathbf{C}\mathbf{z}) + \mathbf{C}\mathbf{p} \\ \mathbf{C}^{\top}\mathbf{u}^{\mathcal{F}} \end{pmatrix} = \mathbf{0} \quad (12)$$

and a numerical solution of (11) is computed using a Sequential Quadratic Programming (SQP) scheme (cf., e.g., [21]). Starting with a feasible initial guess $\mathbf{u}^{\mathcal{F}}$ (i.e. $\mathbf{C}^{\top}\mathbf{u}^{\mathcal{F}} = 0$), we solve

$$\begin{pmatrix} \mathbf{H} & \mathbf{C} \\ \mathbf{C}^{\top} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \delta_{\mathbf{u}} \\ \mathbf{p} \end{pmatrix} = - \begin{pmatrix} \mathbf{D}_{\mathbf{u}}^{\mathcal{F}}(\mathbf{u}^{\mathcal{F}} + \mathbf{C}\mathbf{z}) \\ \mathbf{0} \end{pmatrix} \quad (13)$$

and update $\mathbf{u} \leftarrow \mathbf{u} + \delta_{\mathbf{u}}$. Note that typical to SQP algorithms the Lagrange multipliers are computed directly. For the Hessian we use a Gauß-Newton approximation with respect to the data term

$$\mathbf{H} = T_{\mathbf{u}}^{\top}T_{\mathbf{u}} + \alpha\mathbf{A}.$$

For the computation of a numerical solution of the so-called KKT system (13) we use the multigrid based solution scheme as proposed in [15]. The optimization procedure is summarized in Alg. 2.

Note that in contrast to the linear part we now have to solve a $(n + m)$ -by- $(n + m)$ linear system where m is the dimension of the linear space and n is approximately the number of pixels/voxels.

Algorithm 2 Fine Image Registration: $[\mathbf{u}_t^{\mathcal{F}}] \leftarrow \text{FIR}(R, T, \mathbf{C}, \mathbf{u}^{\mathcal{F}}, \mathbf{z})$.

for $k = 0, \dots$ **do**

Set $\mathbf{u} \leftarrow \mathbf{u}^{\mathcal{F}} + \mathbf{C}\mathbf{z}$ and compute $T(\mathbf{u})$, $T_{\mathbf{u}}(\mathbf{u})$, $\mathbf{D}^{\mathcal{F}}(\mathbf{u}^{\mathcal{F}})$, $\mathbf{D}_{\mathbf{u}}^{\mathcal{F}}(\mathbf{u}^{\mathcal{F}})$.

Set $\mathbf{H} \leftarrow T_{\mathbf{u}}^{\top} T_{\mathbf{u}} + \alpha \mathbf{A}$ and solve

$$\begin{pmatrix} \mathbf{H} & \mathbf{C} \\ \mathbf{C}^{\top} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \delta_{\mathbf{u}} \\ \mathbf{p} \end{pmatrix} = - \begin{pmatrix} \mathbf{D}_{\mathbf{u}}^{\mathcal{F}} \\ \mathbf{0} \end{pmatrix}.$$

set $\gamma \leftarrow 1$

while line_search **do**

$\mathbf{u}_t^{\mathcal{F}} \leftarrow \mathbf{u}^{\mathcal{F}} + \gamma \delta_{\mathbf{u}}$

if $\mathbf{D}^{\mathcal{F}}(\mathbf{u}_t^{\mathcal{F}}) < \mathbf{D}^{\mathcal{F}}(\mathbf{u}^{\mathcal{F}})$ **then**

STOP line search

end if

$\gamma \leftarrow \gamma/2$

end while

if $\|\mathbf{u}_t^{\mathcal{F}} - \mathbf{u}^{\mathcal{F}}\| \leq \text{tol}$ **then**

return

end if

$\mathbf{u}^{\mathcal{F}} \leftarrow \mathbf{u}_t^{\mathcal{F}}$.

end for

4.3 The multilevel approach

For our multilevel approach we require two types grid transfer operators. Firstly, we need a grid transfer operator which coarsen the images, i.e. transfers an image from a fine level $\ell - 1$ to a coarser level ℓ . Secondly, we need two grid transfer operators which restrict and prolongate the transformation $\mathbf{u}^{\mathcal{F}}$.

For image transfer we use the down-sampled and convolved image

$$R_{\ell} = \text{transfer}(R_{\ell-1})$$

where the convolution stencil corresponds to a smoothing with $(1, 3, 3, 1)/8$ in each coordinate direction. For the prolongation of $\mathbf{u}^{\mathcal{F}}$ we use a linear interpolation scheme, as it is common for multigrid approaches. We denote the matrix representation of the prolongation by \mathbf{P}_{ℓ} and use \mathbf{P}_{ℓ}^{\top} as a restriction operator. We assemble \mathbf{B}_1 on the finest level and define $\mathbf{B}_{\ell} := \mathbf{B}_{\ell-1} \mathbf{P}_{\ell-1}$ for the coarser levels. Our numerical approach is summarized in Alg. 1.

For the minimization on the coarsest level \max_{ℓ} we initialize $\mathbf{u}_{\max_{\ell}}^{\mathcal{F}} = \mathbf{0}$ and $\mathbf{z}_{\max_{\ell}} = \mathbf{0}$. The numerical solutions of the optimization problems are

transferred to the finer grid and used as starting values. For the linear part we set $\mathbf{z}_\ell = \mathbf{z}_{\ell+1}$ and for the non-linear part we use the prolongation operator $\mathbf{P}_{\ell+1}$. However, since our prolongation scheme may not maintain the constraints, we apply an additional explicit orthogonalization step

$$\mathbf{u}_\ell = \text{proj}(\mathbf{P}_\ell \mathbf{u}_{\ell+1}, \mathbf{C}_\ell) \quad \text{where} \quad \text{proj}(\mathbf{u}, \mathbf{C}) = \mathbf{u} - \mathbf{C}(\mathbf{C}^\top \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{u}.$$

The optimization scheme on the ℓ th level is considered as converged if $\mathbf{u}_\ell^{\mathcal{F}}$ and/or \mathbf{z}_ℓ converges.

5 Numerical Example

Fig. 1 shows two two-dimensional sections of two three-dimensional magnetic resonance (MR) scans of a human knee; see also [18]. In this example, we choose \mathbf{C} to be the space of affine linear transformations, such that the coarse and the linear registrations schemes coincides.

We start with some general remarks. Tab. 1 summarizes our results for four different multilevel registration of the two 256×256 images, where five levels have been used such that the images on the coarsest level are 16×16 . For comparison reasons, we used the same stopping criteria for the building blocks. The coarse registration on an ℓ th level is considered as converged if $\|\mathbf{z}^{(k)} - \mathbf{z}^{(k-1)}\|_2 \leq 1\%$ or the number of iterations k exceeds 100. For the non-linear/fine part, we stopped if $\|\mathbf{u}^{(k)} - \mathbf{u}^{(k-1)}\|_2 \leq 0.01\%$ or the number of iterations k exceeds 5.

In the following we discuss three registration approaches, the linear (or coarse), the common sequential linear/non-linear, and the new COFIR registrations.

Plain coarse/linear The most common approach is to use a linear preregistration. The results of the linear registration are depicted in Fig. 1 (d,e,f). We observe that the overall result is fair in average but considerable deformations are still observable particularly at the upper part of the Tibia. This is in accordance to the chosen integral based distance measure; cf. (4). The computed deformation field is almost a pure rotation about 13 degrees.

Plain non-linear The results of a non-linear registration without an affine linear preregistration are depicted in Fig. 1 (g,h,i). Note that the results are meaningless (see, e.g., the deformation of the Tibia).

The new Cofir approach Fig. 1 (j,k, ℓ) shows the results of the new COFIR registration scheme. As is apparent from this figure, this registration gives reasonable results for the deformation as well as for the transformed image T^{COFIR} .

Remark 2 *The new COFIR approach combines two features, the clear segregation of the transformation space to coarse and fine parts and their treatment in a combined multilevel strategy.*

In all our examples, the new COFIR registration and the usual sequential linear/non-linear registration lead to visually undistinguishable results. Therefore, one may conjecture that the separation assumption is true.

6 Conclusions and further discussion

In this paper we present a novel and general framework for image registration based on a variational principle. As compared to other variational based registration schemes, the new approach has two new features. Firstly, a clean separation of the transformation space in a *coarse* and its orthogonal complement, the *fine* space, is performed. Only the fine part of the transformation is regularized. Secondly, rather than optimizing the parts separately using a multilevel scheme for each of the parts, we combine and solve for both parts simultaneously.

We present an implementation which is based on a multilevel sequential Gauss-Newton scheme that allows for the introduction of coarse and fine transformations on coarse grids. For the fine part, we used a constrained optimization procedure to ensure the proper segregation of spaces.

For convex optimization problems, the optimization procedure is of minor importance. For non-convex problems, like image registration, where one generally encounters many local minima, the procedure does play an important role. Therefore, the simultaneous treatment in the multilevel approach becomes essential.

Though the differences between the usual sequential linear/non-linear and the COFIR registration appear to be very small, we prefer the COFIR scheme for two reasons. The first reason is that the COFIR scheme is based on a sound mathematical formulation. The second reason is, that this clear approach comes almost for free: the computation times for a COFIR and a sequential linear/non-linear registration are about the same.

Acknowledgements: Jan Modersitzki was supported by the US National Institutes of Health under Grant NIHR01 HL 068904. We are indebted to Michele Benzi for suggestions and discussions.

References

- [1] N. M. Alpert, J. F. Bradshaw, D. Kennedy, and J. A. Correia, *The principal axes transformation – A method for image registration*, Journal of nuclear medicine **31** (1990), no. 10, 1717–1722.
- [2] Ruzena Bajcsy and Stane Kovačič, *Toward an individualized brain atlas elastic matching*, Tech. Report MS-CIS-86-71 Grasp Lap 76, Dept. of Computer and Information Science, Moore School, University of Philadelphia, 1986.
- [3] Barequet and Sharir, *Partial surface and volume matching in three dimensions*, IEEE TPAMI: IEEE Transactions on Pattern Analysis and Machine Intelligence **19** (1997).
- [4] Paul J. Besl and Neil D. McKay, *A method for registration of 3-d shapes*, IEEE Trans. Pattern Anal. Mach. Intell. **14** (1992), no. 2, 239–256.
- [5] Morten Bro-Nielsen, *Medical image registration and surgery simulation*, Ph.D. thesis, IMM, Technical University of Denmark, 1996.
- [6] Chaim Broit, *Optimal registration of deformed images*, Ph.D. thesis, Computer and Information Science, University of Pennsylvania, 1981.
- [7] Lisa Gottesfeld Brown, *A survey of image registration techniques*, ACM Computing Surveys **24** (1992), no. 4, 325–376.
- [8] Gary E. Christensen and H. J. Johnson, *Consistent image registration*, IEEE Transaction on Medical Imaging **20** (2001), no. 7, 568–582.
- [9] Gary Edward Christensen, *Deformable shape models for anatomy*, Ph.D. thesis, Sever Institute of Technology, Washington University, 1994.
- [10] A. Collignon, A. Vandermeulen, P. Suetens, and G. Marchal, *3d multi-modality medical image registration based on information theory*, Kluwer Academic Publishers: Computational Imaging and Vision **3** (1995), 263–274.

- [11] M. Droske, M. Rumpf, and C. Schaller, *Non-rigid morphological registration and its practical issues*, Proc. ICIP '03, IEEE International Conference on Image Processing, Sep. 2003, Barcelona, Spain, 2003.
- [12] Bernd Fischer and Jan Modersitzki, *Fast diffusion registration*, AMS Contemporary Mathematics, Inverse Problems, Image Analysis, and Medical Imaging **313** (2002), 117–129.
- [13] ———, *Combining landmark and intensity driven registrations*, PAMM **3** (2003), 32–35.
- [14] ———, *Curvature based image registration*, J. of Mathematical Imaging and Vision **18** (2003), no. 1, 81–85.
- [15] Eldad Haber and Jan Modersitzki, *A multilevel methods for image registration*, Technical Report TR-2004-014-A, Department of Mathematics and Computer Science, Emory University, Atlanta GA 30322, May 2004, Submitted to SIAM J. of Scientific Computing.
- [16] ———, *Numerical methods for volume preserving image registration*, Technical Report TR-2004-012-A, Department of Mathematics and Computer Science, Emory University, Atlanta GA 30322, Feb 2004, Submitted to Inverse Problems.
- [17] G. Hermosillo, *Variational methods for multimodal image matching*, Ph.D. thesis, Université de Nice, France, 2002.
- [18] Sven Kabus, Thomas Netsch, Bernd Fischer, and Jan Modersitzki, *B-spline registration of 3d images with Levenberg-Marquardt optimization*, Proceedings of the SPIE 2004, Medical Imaging, San Diego Februar 14-19, 2004, SPIE, 2004, Accepted for publication, pp. 1–10.
- [19] Calvin R. Maurer and J. Michael Fitzpatrick, *Interactive image-guided neurosurgery*, ch. A Review of Medical Image Registration, pp. 17–44, Park Ridge, IL, American Association of Neurological Surgeons, 1993.
- [20] Jan Modersitzki, *Numerical methods for image registration*, Oxford University Press, 2004.
- [21] J. Nocedal and S.J. Wright, *Numerical optimization*, Springer, New York, 1999.

- [22] A. Roche, *Recalage d'images médicales par inférence statistique*, Ph.D. thesis, Université de Nice, Sophia-Antipolis, France, 2001.
- [23] Karl Rohr, *Landmark-based image analysis*, Computational Imaging and Vision, Kluwer Academic Publishers, Dordrecht, 2001.
- [24] Jean-Philippe Thirion, *Non-rigid matching using demons*, 1996, IEEE 1996, pp. 245–251.
- [25] Paul A. Viola, *Alignment by maximization of mutual information*, Ph.D. thesis, Massachusetts Institute of Technology, june 1995, pp. 1–155.
- [26] Barbara Zitová and Jan Flusser, *Image registration methods: a survey*, Image and Vision Computing **21** (2003), no. 11, 977–1000.

Table 1: SSD reduction and convergence results for the knee example. The tables show the image sizes and the reductions of the image difference versus the discretization level. The reduction in the difference is obtained from either the coarse grid correction, the coarse/linear, and the fine/non-linear registration. The reduction is given in percent (%), where the initial difference on a particular level is considered as 100%. We also give the number of iterations spend for the coarse/linear and fine/non-linear parts.

| coarse/linear | | | | |
|----------------------|------------------|--------|-------|------|
| level | size | coarse | red. | iter |
| 1 | 16×16 | – | 46.4% | 21 |
| 2 | 32×32 | 53.1% | 52.0% | 10 |
| 3 | 64×64 | 58.4% | 58.0% | 29 |
| 4 | 128×128 | 63.4% | 62.8% | 62 |
| 5 | 256×256 | 66.7% | 66.6% | 100 |

| sequential linear/non-linear | | | | |
|-------------------------------------|------------------|--------|-------|------|
| level | size | coarse | red. | iter |
| 1 | 16×16 | 52.1 | 35.2% | 5 |
| 2 | 32×32 | 38.5% | 33.4% | 5 |
| 3 | 64×64 | 38.2% | 34.1% | 5 |
| 4 | 128×128 | 39.1% | 36.6% | 5 |
| 5 | 256×256 | 42.9% | 41.6% | 5 |

| COFIR | | | | | | |
|--------------|------------------|--------|--------|------|-------|------|
| level | size | coarse | coarse | | fine | |
| | | | red. | iter | red. | iter |
| 1 | 16×16 | – | 46.4% | 21 | 28.2% | 5 |
| 2 | 32×32 | 35.5% | 33.6% | 8 | 29.9% | 5 |
| 3 | 64×64 | 37.2% | 32.6% | 5 | 29.9% | 5 |
| 4 | 128×128 | 34.4% | 34.0% | 9 | 32.2% | 5 |
| 5 | 256×256 | 38.8% | 38.5% | 17 | 37.0% | 5 |

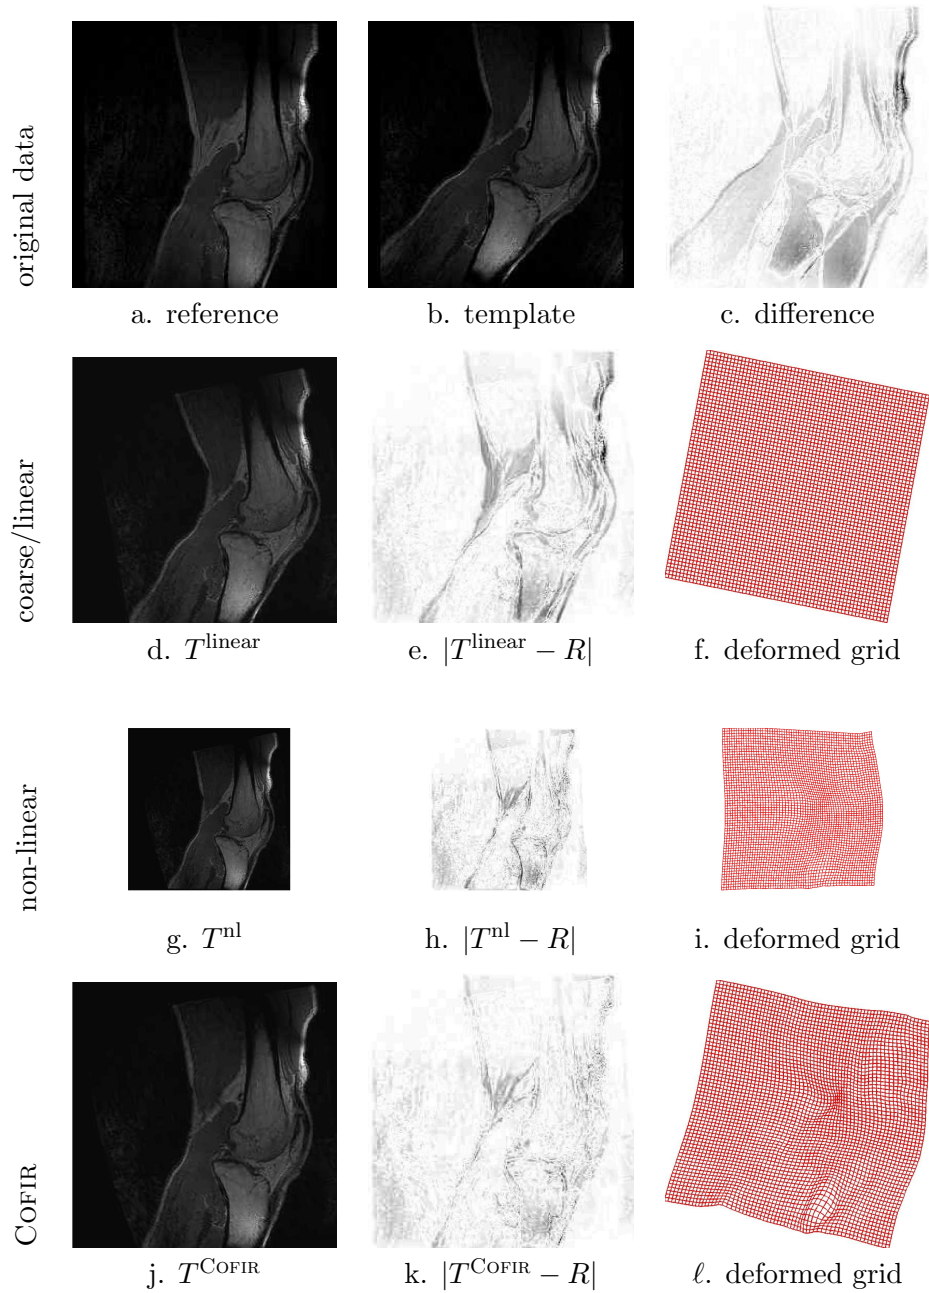


Figure 1: Two MR images of a human knee (a,b) and the difference image (c) as well as results for the linear or coarse (d,e,f), the non-linear (g,h,i), and the new COFIR (j,k,l) registrations: final transformed images (first column), difference image (second column), and deformed grid (third column).