FLIRT: A Flexible Image Registration Toolbox

Bernd Fischer and Jan Modersitzki

Institute of Mathematics
University of Lübeck, 23560 Lübeck, Germany
{fischer,modersitzki}@math.uni-luebeck.de
http://www.math.uni-luebeck.de

Abstract. Image registration is central to many challenges in medical imaging today and has a vast range of applications. The purpose of this note is to provide a toolbox for intensity based non-rigid registration problems. To do so, we review some of the most promising non-linear registration strategies currently used in medical imaging and show that all these techniques may be phrased in terms of a variational problem and allow for a unified treatment.
Depending on the application at hand, it is often desirable to constrain the wanted deformation. The idea is to incorporate higher level information about the expected deformation. We examine the most common constraints and show again that they may be conveniently phrased in a variational setting.
As a consequence, all of discussed modules allow for fast implementations and may be combined in any favorable order. We discuss individual methods for various applications, including the registration of histological serial sections of a human brain.

1 Introduction

In the last two decades, computerized non-rigid image registration has played an increasingly important role in medical imaging, see for example MAINTZ & VIERGEVER [13], FITZPATRICK ET AL. [10], and references therein. The problem of registration arises whenever images acquired from different subjects, at different times, or from different scanners need to be combined for analysis or visualization.

Due to the wide range of applications a variety of different registration techniques has been developed. Here, we focus on so-called *intensity-driven approaches*. These schemes aim to match intensity pattern between a deformed scan and the target based on a rigorous mathematical criteria. Depending on the application, different strategies may be employed. From a practical point of view, it is desirable to incorporate properties of the underlying problem into the registration scheme. Here, we provide a toolbox of registration routines which enables the user to choose in a consistent way building blocks for schemes which cover a wide range of applications. The idea is to phrase each individual block in terms of a variational formulation. This not only allows for a unified treatment but also for fast and reliable implementation. The various building blocks comprises three categories: smoother and internal forces, distances and external forces, and "hard" or "soft" constraints. The *internal forces*, are defined for the wanted *displacement field* itself and are designed to keep the displacement field smooth during deformation. In contrast, the *external forces* are computed from the image data and are defined to drive the displacement in order to arrive at the desired registration result. Whereas the internal forces implicitly constrain the displacement to obey a smoothness criteria, the additional *constraints* force the displacement to satisfy explicit criteria, like for example landmark or volume preserving imposed constraints.

The paper is organized as follows. In Section 2 we summarize the most popular choices for the above outlined building blocks. Furthermore, we set up a general and unified framework for automatic non-rigid registration. In Section 3 we show in more detail, how these building blocks can be translated into a variational setting. It is this formulation, which allows for a fast and reliable numerical treatment. In Section 4 we indicate on how to actually implement the registration schemes. An example in Section 5 highlights the importance of having more than one regularizer at hand.

2 A flexible image registration toolbox

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 object or a similar object. There exist many instances in a medical environment which demand for a registration, including the treatment verification of pre- and post-intervention images, study of temporal series of cardiac images, and the monitoring of the time evolution of an agent injection subject to patient motion. Another important area is the need for combining information from multiple images, acquired using different modalities, like for example computer tomography (CT) and magnetic resonance imaging (MRI).

To be successful, each individual application should be treated by a specific registration technique. It is the purpose of this note to provide a toolbox for non-linear registration schemes, which may be adapted to the special problem class under consideration. The main building blocks of this toolbox resemble typical user demands and may be assembled in a consistent and intuitive fashion.

Given two images, a reference R and a template T, the aim of image registration is to find a global and/or local transformation from T onto R in such a way that the transformed template matches the reference. Ideally there exists a coordinate transformation u such that the reference R equals the transformed template T_u. Given such a displacement u, the registration problem reduces to a simple interpolation task. However, in general it is not possible to come up with a perfect u, and the registration problem is to compute an application conformal transformation u, given the reference and template image.

It should be pointed out, that apart from the fact that a solution may not exist, it is not necessarily unique. For an example, see Modersitzki [14]. In other words, intensity based registration is inherently an ill-posed problem.

A displacements u which does produce a perfect or nearly perfect alignment of the given images is not necessarily a "good" displacement. For example, a computed displacement which interchanges the eyes of one patient when registered to a probabilistic atlas in order to produce a nearly perfect alignment, has obviously to be discarded. Also, folding and cracks introduced by the displacement are typically not wanted. Therefore it is desirable to have a possibility to incorporate features into the registration model, such that the computed displacement u does resemble the properties of the acquisition, like for example the elastic behavior of a human brain. To mimic the elastic properties of the objects under consideration is a striking example for internal forces. These forces constrain the displacement to physically meaningful movements.

In contrast, the external forces are designed to push the deformable template into the direction of the reference. These forces are based upon the intensities of the images. The idea is to design a similarity measure, which is ideally calculated from all voxel values. An intuitive measure is the sum of squares of intensity differences (SSD). This is a reasonable measure for some applications like the serial registration of histological sections. If the intensities of corresponding voxels are no longer identical, the SSD measure may perform poorly. However, if the intensities are still linearly related, a the correlation coefficient (CC) based measure is the measure of choice for monomodal situations. In contrast, the mutual information (MI) related measure is based on the cooccurrence of intensities in both images as reflected by their joint intensity histogram. It appears to be the most successful similarity measure for multimodal imaginary, like MR-PET.

Finally, one may want to guide the registration process by incorporating additional information which may be known beforehand. Among these are landmarks and fiducial markers. Sometimes it is also desirable to impose a local volume-preserving (incompressibility) constraint which may, for example, compensate for registration artifacts frequently observed by processing pre- and post-contrast images. Depending on the application and the reliability of the specific information, one may want to insist on a perfect fulfilment of these constraints or on a relaxed treatment. For examples, in practise, it is a tricky (and time consuming) problem to determine landmarks to subvoxel precision. Here, it does not make sense to compute a displacement which produces a perfect one to one match between the landmarks.

Summarizing, the general registration problem may be phrased as follows.

(IR) image registration problem:

$$\begin{aligned} \mathcal{J}[u] = \mathcal{D}[R, T; u] + \alpha \mathcal{S}[u] = \min, \\ \text{subject to} \quad \mathcal{C}_j[u] = 0, \quad j = 1, 2, \ldots, m. \end{aligned}$$

Here, $\mathcal{D}$ models the distance measure (external force, e.g., MI), $\mathcal{S}$ the smoother (internal force, e.g., elasticity), and $\mathcal{C}$ explicit constraints (e.g., landmarks). The regularization parameter α may be used to control the strength of the smoothness of the displacement versus the similarity of the images. In the following we will discuss these building blocks in more detail.

3 Toolbox building blocks

Our approach is valid for images of any spatial dimension d, i.e., there is no restriction to $d = 2, 3, 4$. The reference and template images are represented by the compactly supported mappings $R, T : \Omega \to \mathbb{R}$, where without loss of generality, $\Omega =]0, 1[^d$. Hence, $T(x)$ denotes the intensity of the template at the spatial position x, where for ease of discussion we set $R(x) = b_R$ and $T(x) = b_T$ for all $x \notin \Omega$. Here, b_R and b_T are appropriately chosen background intensities. The overall goal is to find a *displacement* u, such that ideally T_u is similar to R, where T_u is the deformed image, i.e., $T_u(x) = T(x - u(x))$. Note that $u = (u_1, \ldots, u_d)$ denotes a vector field.

The starting point of our numerical treatment is the minimization of problem (IR). In order to compute a minimizer we apply a steepest descent method, where we take advantage of the calculus of variations. To end up with an efficient and fast converging scheme, we require to have explicit expressions of the derivatives of building blocks $\mathcal{D}$, $\mathcal{S}$, and $\mathcal{C}$. In the following subsections we will exemplarily discuss the most popular building blocks as well as their derivatives.

Smoother and Internal Forces. The nature of the deformation depends strongly on the application under consideration. For example, a slice of a paraffin embedded histological tissue does deform elastically, whereas the deformation between the brains of two different individuals is most likely not elastically. Therefore, it is necessary to supply a model for the nature of the expected deformation.

We now present some of the most prominent smoothers $\mathcal{S}$ and discuss exemplarily the Gâteaux-derivatives for two of them. An important point is, that we are not restricted to a particular smoother $\mathcal{S}$. Any smoother can be incorporated into this toolbox, as long as it possesses a Gâteaux-derivative.

In an abstract setting, the Gâteaux-derivative looks like

$$d\mathcal{S}[u; v] := \lim_{h \to 0} \frac{1}{h}(\mathcal{S}[u + hv] - \mathcal{S}[u]) = \int_\Omega \langle \mathcal{A}[u], v \rangle_{\mathbb{R}^d} \, dx,$$

where $\mathcal{A}$ denotes the associated linear partial differential operator. Note that for a complete derivation one also has to consider appropriate boundary conditions. However, these details are omitted here for presentation purposes; see Modersitzki [14] for details.

Elastic registration This particular smoother measures the elastic potential of the deformation. In connection with image registration it has been introduced by Broit [3] and discussed by various image registration groups; see, e.g., Bajcsy & Kovačič [2] or Fischer & Modersitzki [6]. The partial differential operator is the well-known Navier-Lamé operator. For this smoother, two natural parameters, the so-called Lamé-constants can be used in order to capture features of the underlying elastic body. A striking example, where the underlying physics suggests to look for deformations satisfying elasticity constraints, is the three-dimensional reconstruction of the human brain from a histological sectioning. Details are given in Schmitt [17] and Modersitzki [14].

Fluid registration Due to the fact that an elastic body memorizes its non-deformed initial state (rubber band), elastic registration schemes are only able to compensate for small deformations. The situation changes for the viscous fluid model. Here the body adapts to its current state (honey) and consequently is much more flexible than an elastic body. The viscous fluid approach was introduced to image registration by CHRISTENSEN [4]. His derivation was based on a specific linearization of the NAVIER-STOKES equation. However, there is yet another derivation of the underlying partial differential equations, which does fit into “design rules of our toolbox. Roughly speaking, one obtains these equations by considering the elastic potential of the *velocity* of the displacement field. It should come as no surprise that the partial differential operator is again the NAVIER-LAMÉ operator, this time, however, applied to the velocity. The wanted deformation is related to the velocity via the material derivative and is straightforward to recover.

Since the viscous fluid approach is quite flexible, it is mainly used when the focus is more on similarity than on a “natural deformation process. For example, for the design of a probabilistic brain atlas, a biophysical model for the nature of the deformations is not available, but the fluid registration has been proven to be a valuable tool.

Diffusion registration For image registration problems FISCHER & MODERSITZKI [7] introduced the so-called diffusion regularization

$$\mathcal{S}^{\text{diff}}[u] := \frac{1}{2}\sum_{\ell=1}^{d}\int_{\Omega}\|\nabla u_\ell\|^2 \; dx, \tag{1}$$

which is well-known for *optical flow* applications; see HORN & SCHUNCK [12]. The associated GÂTEAUXderivative leads to the well-studied LAPLACE-operator, i.e., $\mathcal{A}^{\text{diff}}[u] = \Delta u = (\Delta u_1, \ldots, \Delta u_d)$, where $\Delta u_\ell = \partial_{x_1 x_1} u_\ell + \cdots + \partial_{x_d x_d} u_\ell$. The main reason for introducing this smoother was its exceptional computational complexity. FISCHER & MODERSITZKI [7] devised an $\mathcal{O}(N)$ (!) implementation of the registration scheme, where N denotes the number of image voxels. It is based on an additive operator splitting scheme (which parallelizes in a very natural way). Its outstanding computational speed makes the diffusion registration scheme to a very attractive option for high-resolution, high dimensional, and/or time critical applications. Examples include the registration of a time series of three-dimensional MRI’s or the online correction of the so-called brain shift during the surgery.

Curvature registration As a last example, we present the curvature smoother

$$\mathcal{S}^{\text{curv}}[u] := \frac{1}{2}\sum_{\ell=1}^{d}\int_{\Omega}(\Delta u_\ell)^2 \; dx, \tag{2}$$

introduced by FISCHER & MODERSITZKI [8]. The design principle behind this choice was the idea to make the non-linear registration phase more robust against a poor (affine linear) pre-registration. Since the smoother is based on second order derivatives, affine linear maps do not contribute to its costs, i.e.,

$$\mathcal{S}^{\text{curv}}[Cx+b]=0, \quad \text{for all} \quad C\in\mathbb{R}^{d\times d},\ b\in\mathbb{R}^d.$$

In contrast to other non-linear registration techniques, affine linear deformations are corrected naturally by the curvature approach. This advantage is illustrated by Figure 1, where the results of a fluid and a curvature based registration of two X-ray images are compared; see also Section 5. Again the GÂTEAUX derivative is explicitly known and leads to the so-called bi-harmonic operator $\mathcal{A}^{\text{curv}}[u]=\Delta^2 u$.

Distances and External Forces. Another important building block is the similarity criterion. As for the smoothing operators, we concentrate on those measures $\mathcal{D}$ which allow for differentiation. Moreover, we assume that there exists a function $f:\mathbb{R}^d\times\mathbb{R}^d\to\mathbb{R}^d$, the so-called *force field*, such that

$$\begin{aligned} d\mathcal{D}[R,T;u;v] &= \lim_{h\to 0}\frac{1}{h}(\mathcal{D}[R,T;u+hv]-\mathcal{D}[R,T;u]) \\ &= \int_\Omega \langle f(R,T,x,u(x)),v(x)\rangle_{\mathbb{R}^d}\ dx. \end{aligned}$$

Again, we are not restricted to a particular distance measure. Any measure can be incorporated into our toolbox, as long as it permits a GÂTEAUX-derivative.

The most common choices for distance measures in image registration are the *sum of squared differences*, *cross correlation*, *cross validation*, and *mutual information*. We give explicit formulae for only two of them; for more information see, e.g., MODERSITZKI [14] or ROCHE [15].

Sum of squared differences. The measure is based on a point-wise comparison of image intensities,

$$\mathcal{D}^{\text{SSD}}[R,T;u]:=\frac{1}{2}\int_\Omega \big(R(x)-T_u(x)\big)^2\ dx,$$

and the force-field is given by $f^{\text{SSD}}(R,T,x,y)=(T(x-y)-R(x))\cdot\nabla T(x-y)$. This measure is often used when images of the same modality have to be registered.

Mutual information. Another popular choice is mutual information. It basically measures the entropy of the joint density $\rho^{R,T}$, where $\rho^{R,T}(g_1,g_2)$ counts the number of voxels with intensity g_1 in R and g_2 in T. The precise formula is

$$\mathcal{D}^{\text{MI}}[R,T;u]:=-\int_{\mathbb{R}^2} p^{R,T_u}\log\frac{p^{R,T_u}}{p^R p^{T_u}}\ d(g_1,g_2),$$

where p^R and p^{T_u} denote the marginal densities. Typically, the density is replaced by a PARZEN-window estimator; see, e.g. VIOLA [18]. The associated force-field is given by

$$f^{\text{MI}}(R,T,x,y)=\int_\Omega [\Psi_\sigma * \partial_{g_2} L^{R,T_u}](R(x),T_u(x))\cdot\langle\nabla T_u(x),v(x)\rangle_{\mathbb{R}^d},$$

where $L^{R,T_u} := 1 + p^{R,T_u}(\log p^{R,T_u} - \log(p^R p^{T_u})$ and Ψ is the PARZEN-window function; see, e.g., HERMOSILLO [11] or D'AGOSTINO ET AL. [5]. This measure is useful when images of a different modality have to be registered.

Additional Constraints. Often it is desirable to guide the registration process by incorporating additional information which may be known beforehand, like for example fiducial marker. To incorporate such information, the idea is to add additional constraints to the minimization problem. For example, to restrict the deformation to volume preserving mappings, one has to add the quantity

$$\mathcal{C}[u] := \frac{1}{2} \int_{\Omega} (\det \nabla u)^2 \, dx$$

to the smoother; see also ROHLFING & MAURER [16]. Note that the JACOBIAN $\det \nabla u(x)$ has to vanish, if the deformation at x is incompressible.

In other applications, one may want to incorporate landmarks or fiducial markers. Let r^j be a landmark in the reference image and t^j be the corresponding landmark in the template image. The toolbox allows for either adding explicit constraints

$$\mathcal{C}_j[u] := u(t^j) - t^j + r^j, \quad j = 1, 2, \ldots, m,$$

which have to be precisely fulfilled $\mathcal{C}_j[u] = 0$ ("hard" constraints), or by adding an additional cost term

$$\mathcal{C}[u] := \sum_{j=1}^{m} \lambda_j \, \|\mathcal{C}_j[u]\|^2_{\mathbb{R}^d}$$

to the smoother ("soft" constraints, since we allow for deviations). For a more detailed discussion, we refer to FISCHER & MODERSITZKI [9].

4 Numerical Treatment

As already pointed out, our numerical approach is based on the EULER-LAGRANGE equations for the problem (IR)

$$\mathcal{A}[u](x) + f(R, T, x, u(x)) + \sum_{j=1}^{m} \lambda_j d\mathcal{C}_j[u](x) = 0 \quad \text{and} \quad \mathcal{C}_j[u] = 0, \; j = 1, \ldots, m,$$

where the λ_j's are LAGRANGE parameter. Roughly speaking, all associated GÂTEAUX-derivatives have to vanish. It remains to efficiently solve this system of non-linear partial differential equations. After invoking a time-stepping approach and after an appropriate space discretization, we finally end up with a system of linear equations. As it turns out, these linear systems have a very rich structure, which allows one to come up with very fast and robust solution schemes for all of the above mentioned building blocks. It is important to note that the system matrix does not depend on the force field and the constraints. Thus, changing the similarity measure or adding additional constraints does not change the favorable computational complexity. Moreover, fast and parallel solution schemes can be applied to even more reduce the computation time.

5 An example: X-rays of hands

We present a synthetic example in order to demonstrate the fact, that changing the smoother may dramatically affect the registration result. Here, we compare the fluid and curvature smoother, both accompanied with the SSD measure. In Figure 1, a reference (a) and a template image (b) are displayed (modified X-rays from human hands, images from AMIT [1]). Obviously, an affine linear pre-registration (rotation of about 45 degrees and re-scaling) would improve the similarity of the images considerably. However, in order to keep the issue of interest clear, we did not apply any pre-registration. For the fluid registration we end up with the deformed template displayed in Figure 1(e) and for the curvature registration we obtain the result shown in Figure 1(f) (Figure 1(c,d) show intermediate results of the time-stepping scheme). As it is apparent from this example, the fluid approach produced a miss-registration whereas the curvature approach produced a satisfactory result.

The main point is that the fluid registration (as well as the other approaches) does penalize affine linear deformations and may therefore privilege non-linear deformations, as is clearly visible in Figure 1(c). Due to its flexibility, the fluid method finally recovers the reference, where, however, the deformation field is unnatural. In contrast, the curvature approach does not penalize affine linear deformations as can be seen in Figure 1(d), which displays almost a rotated and re-scaled template image. Note that the deformation is not completely linear and note that this is an extreme example. For comparison reasons, we applied non-optimized methods on a single scale (though all building blocks can be applied in a multiscale resolution as well).

6 Conclusions

In this note we presented a general approach to image registration. Its flexibility enables one to integrate and to combine in a consistent way various different registration modules.We discussed the use of different smoothers, distance measures, and additional constraints. The numerical treatment is based on the solution of a partial differential equation related to the EULER-LAGRANGE equations. These equations are well studied and allow for fast, stable, and efficient schemes. Due to page limits, we reported on only one example, showing the effect of different smoothers. We will report in a forthcoming paper on an exhaustive comparison of the various building blocks.

References

1. Y. Amit, *A nonlinear variational problem for image matching*, SIAM J. Sci. Comput. **15**(1), 207–224, 1994.
2. R. Bajcsy and S. Kovačič, *Multiresolution elastic matching*, Computer Vision, Graphics and Image Processing **46**, 1–21, 1989.

3. C. Broit, *Optimal registration of deformed images*, Ph.D. thesis, Computer and Information Science, Uni Pensylvania, 1981.
4. G. E. Christensen, *Deformable shape models for anatomy*, Ph.D. thesis, Sever Institute of Technology, Washington University, 1994.
5. E. D'Agostino, J. Modersitzki, F. Maes, D. Vandermeulen, B. Fischer, and P. Suetens, *Free-form registration using mutual information and curvature regularization*, Preprint A-03-05, Institute of Mathematics, University of Lübeck, 2003.
6. B. Fischer and J. Modersitzki, *Fast inversion of matrices arising in image processing*, Num. Algo. **22**, 1–11, 1999.
7. B. Fischer and J. Modersitzki, *Fast diffusion registration*, AMS Contemporary Mathematics, Inverse Problems, Image Analysis, and Medical Imaging **313**, 117–129, 2002.
8. B. Fischer and J. Modersitzki, *Curvature based image registration*, JMIV **18**(1), 81–85, 2003.
9. B. Fischer and J. Modersitzki, *Combination of automatic non-rigid and landmark based registration: the best of both worlds*, Preprint A-03-01, Institute of Mathematics, University of Lübeck, 2003.
10. J. M. Fitzpatrick, D. L. G. Hill, and C. R. Maurer Jr., *Image registration*, Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis (M. Sonka and J. M. Fitzpatrick, eds.), SPIE, 447–513, 2000.
11. G. Hermosillo, *Variational methods for multimodal image matching*, Ph.D. thesis, Université de Nice, France, 2002.
12. B. K. P. Horn and B. G. Schunck, *Determining optical flow*, Artificial Intelligence **17**, 185–204, 1981.
13. J. B. Antoine Maintz and M. A. Viergever, *A survey of medical image registration*, Medical Image Analysis **2**(1), 1–36, 1998.
14. J. Modersitzki, *Numerical methods for image registration*, Oxford University Press, to appear 2003.
15. A. Roche, *Recalage d'images médicales par inférence statistique*, Ph.D. thesis, Université de Nice, Sophia-Antipolis, France, 2001.
16. T. Rohlfing and C. R. Maurer, Jr, *Volume-Preserving Non-Rigid Registration of MR Breast Images Using Free-Form Deformation with an Incompressibility Constraint*, IEEE TMI, to appear 2003.
17. O. Schmitt, *Die multimodale Architektonik des menschlichen Gehirns*, Habilitation (in German), Insitute of Anatomy, Medical University of Lübeck, Germany, 2001.
18. P. A. Viola, *Alignment by maximization of mutual information*, Ph.D. thesis, Massachusetts Institute of Technology, 1–155, 1995.

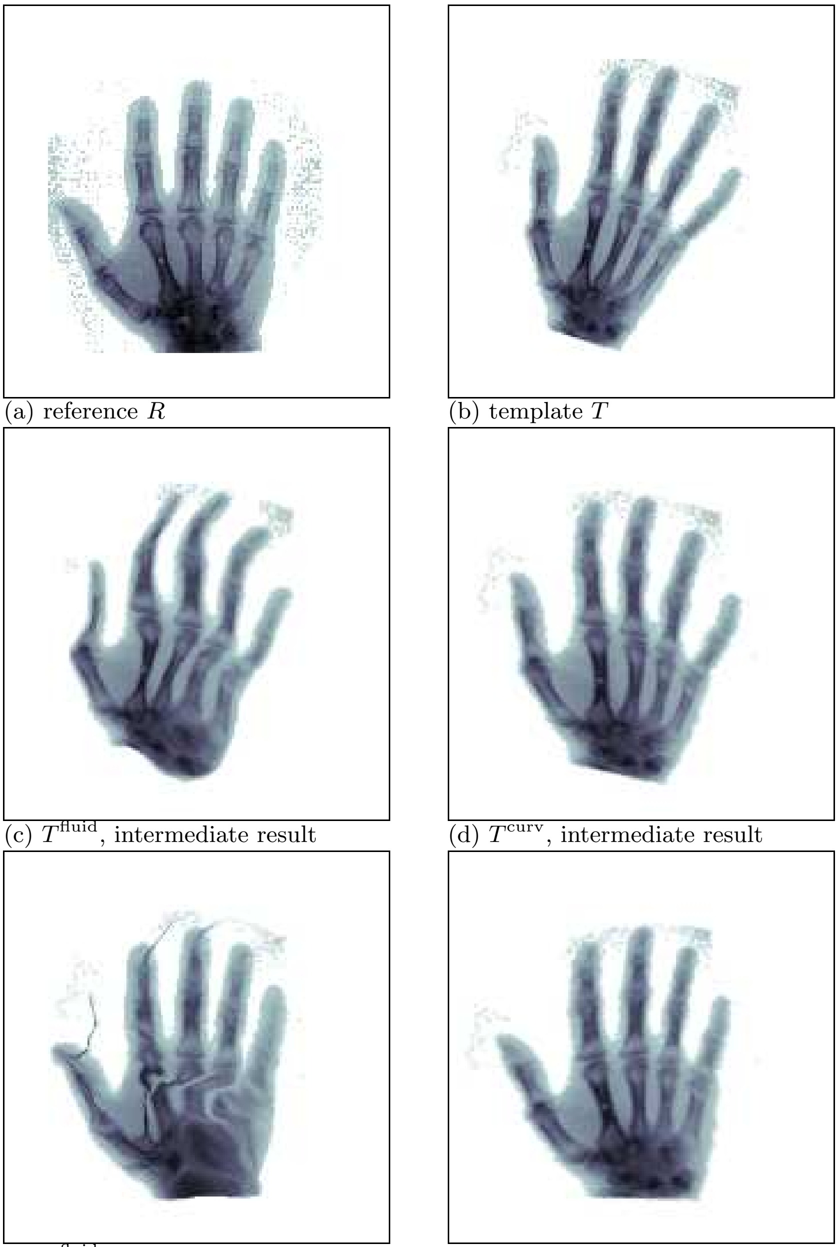

(a) reference R (b) template T

(c) T^{fluid}, intermediate result (d) T^{curv}, intermediate result

(e) T^{fluid}, final result (f) T^{curv}, final result

Fig. 1. Two modified X-ray images of human hands; see also Amit [1]. (a) reference image, (b) template image, (c) intermediate result of fluid registration, (d) intermediate result of curvature registration, (e) final result of fluid registration, and (f) final result of curvature registration.