Establishing Local Correspondences towards Compact Representations of Anatomical Structures *

Xiaolei Huang[1], Nikos Paragios[2], and Dimitris Metaxas[1]

[1] Division of Computer and Information Sciences,
Rutgers, the State University of New Jersey, Piscataway, NJ, USA

[2] Real Time Vision & Modeling Department,
Siemens Corporate Research, Princeton, NJ, USA

Abstract. Computer-aided diagnosis is often based on comparing a structure of interest with prior models. Such a comparison requires automatic techniques in determining prior models from a set of examples and establishing local correspondences between the structure and the model. In this paper we propose a variational technique for solving the correspondence problem. The proposed method integrates a powerful representation for shapes (implicit functions), a state-of-the art criterion for global registration (mutual information) and an efficient technique to recover local correspondences (free form deformations) that guarantees one-to-one mapping. Local correspondences can then be used to build compact representations for a structure of interest according to a set of training examples. The registration and statistical modeling of Systolic Left Ventricle in Ultrasonic images demonstrate the potential of the proposed technique.

1 Introduction

Organ modeling is a critical component of medical image analysis. To this end, one would like to recover a compact representation that can capture the variation of the structure of interest across individuals. Building such representations requires establishing correspondences across the set of training examples.

Quite often, correspondences are manually determined by users, which is a non-efficient and time consuming process. Once correspondences have been established, the modeling can take place according to various statistical methods leading to compact representations that can be used for detection, segmentation, tracking, etc. of structures of interest.

Shape/image registration [7] is an evolving research activity in medical imaging [12]. One can define the registration problem as follows: recover a transformation between a source and a target shape that results in meaningful correspondences between their basic elements. Such a definition involves three aspects.

The first refers to the selection of an appropriate representation for the structures of interest. Cloud of points, parametric (concrete) structures (e.g. B-splines), compact representations (e.g. medial axis), etc. are often considered. Compact representations are appropriate when seeking global registration and quite in-efficient when seeking for local correspondences. Concrete parametric representation refer to a reasonable selection for local registration. Their limitations are related with the selection of the representation, the number of basic components and the ability to describe multi-component

* A color version of this paper can be found at http://paul.rutgers.edu/ xiaolei/Research.htm

structures. Space coordinates (cloud of points) are simple but rather inconsistent representations that do not provide sufficient support in estimating the local geometric characteristics.

The second aspect in registration refers to the selection of a transformation. The transformation can be either global or local. Global registration techniques aim to recover a linear motion model that is applicable to the entire structure of interest. Such techniques are a compromise between robustness, low complexity and acceptable registration performance. Medical image analysis is a domain where exact correspondences are required. Local registration techniques focus on recovering an element-based flow that creates correspondences in the pixel level.

Last, but not least, given a feature space and a selection of the form of transformation one should define an appropriate mathematical framework to recover the optimal registration parameters. Stochastic, variational, graph-based optimization functions are examples with known strengths and limitations.

In this paper we propose an hierarchical registration method for establishing local correspondences that can deal with the ill-poseness of the local registration problem. We represent the structures of interest (shapes) in a higher dimension using an implicit representation that is derived from the powerful space of distance transforms. Global registration for an arbitrary motion model is obtained using the mutual information criterion [2, 13]. Local deformations are considered to be a complementary (to the global motion model) deformation field. In order to preserve one-to-one correspondence, a cubic B-spline basis free form deformation model (FFD) [10] is used to perform local registration. The task of generating a compact representation using a principal component analysis technique [3] from a set of training examples is considered to validate and to demonstrate the potentials of the proposed registration algorithm.

2 Distance Transforms, Mutual Information & Global Registration

The definition of the feature space is a critical component of the registration process. The use of point clouds [1], deformable models [14], fourier descriptors [11] are some alternatives. Such representations are powerful enough to capture a certain number of local deformations. However, their extension to describe structures of higher dimension than curves and surfaces is in most of the cases not trivial. More advanced techniques are based on implicit geometric characteristics of a structure of interest, like the curvature, medial axes, normals , etc. or combination of them. The estimation of such implicit properties is a difficult task that often requires the parameterization of the structure.

Within the proposed framework, we consider an implicit representation for the source and the target structures [8]. Euclidean distance transforms are used to embed a structure of interest into a higher dimension. Let $\Phi_{\mathcal{S}} : \Omega \rightarrow R^{+}$ be a distance transform representation for a given shape $\mathcal{S}$:

$$\Phi_{\mathcal{S}}(x,y) = \begin{cases} 0, & (x,y) \in \mathcal{S} \\ D((x,y),\mathcal{S}), & (x,y) \in \Omega - \mathcal{S} \end{cases}$$

where $D((x,y),\mathcal{S})$ refers to the min distance between a pixel (x,y) in the embedding space and the shape $\mathcal{S}$ [3]. The use of implicit representations provides additional

[3] The signed distance function is a more powerful representation and can be used to describe closed structures

support to the registration process since one would like to align the original structures as well as their clones that are positioned coherently in the image/volume plane.

The selected implicit shape representation is inherently translation/rotation invariant, as demonstrated in [8]. When a shape undergoes scale variations, the intensity values of its associated distance map (distances to the shape) scales accordingly. Therefore registration of the distance maps of a shape in various scales are analogous to matching images in multiple modalities that refer to the same underlying anatomy. Mutual Information, can address such matching objective. Such criterion is based on the global characteristics of the structures of interest. In order to facilitate the notation let us denote: (**i**) the source representation $\Phi_{\mathcal{D}}$ as f, and (**ii**) the target representation $\Phi_{\mathcal{S}}$ as g.

In the most general case, registration is equivalent with recovering the parameters $\mathbf{\Theta} = (\theta_1, \theta_2, ..., \theta_N)$ of a parametric transformation A such that the mutual information between $f_\Omega = f(\Omega)$ and $g_\Omega^A = g\big(A(\mathbf{\Theta}; \Omega)\big)$ is maximized for a given sample domain Ω;

$$MI(X^{f_\Omega}, X^{g_\Omega^A}) = \mathcal{H}\left[X^{f_\Omega}\right] + \mathcal{H}\left[X^{g_\Omega^A}\right] - \mathcal{H}\left[X^{f_\Omega, g_\Omega^A}\right]$$

where $\mathcal{H}$ represents the differential entropy. Such quantity represents a measure of uncertainty, variability or complexity and consists of three components: (i) the entropy of the model $\mathcal{H}\left[X^{f_\Omega}\right]$, (ii) the entropy of the projection of the model given the transformation $\mathcal{H}\left[X^{g_\Omega^A}\right]$, and (iii) the joint entropy $\mathcal{H}\left[X^{f_\Omega, g_\Omega^A}\right]$ between the model and the projection that encourages transformations where f explains g. One can use the above criterion and an arbitrary transformation (rigid, affine, homographic, quadratic) model to perform global registration that is equivalent with minimizing:

$$E(A(\mathbf{\Theta})) = -MI(X^{f_\Omega}, X^{g_\Omega^A}) = -\iint_{\mathcal{R}^2} p^{f_\Omega, g_\Omega^A}(l_1, l_2) log \frac{p^{f_\Omega, g_\Omega^A}(l_1, l_2)}{p^{f_\Omega}(l_1) p^{g_\Omega^A}(l_2)} dl_1 dl_2$$

where (i) p^{f_Ω} corresponds to the probability density in f_Ω $\big([\Phi_{\mathcal{D}}(\Omega)] \big)$, (ii) $p^{g_\Omega^A}$ corresponds to density in g_Ω^A $\big([\Phi_{\mathcal{S}}(A(\mathbf{\Theta}; \Omega))] \big)$, and (iii) p^{f_Ω, g_Ω^A} is the joint density. Intuitively, the criterion aims to maximize the mutual information between the model intensity ($f(\Omega)$) and its projection ($g(A(\Omega))$), which is equivalent to maximize the Kullback-Leibler divergence between their joint distribution and the direct product of the two separate distributions.

Towards a continuous form of the criterion, a non-parametric Gaussian Kernel density model can be considered to approximate the joint density, leading to the following expression:

$$p^{f_\Omega, g_\Omega^A}(l_1, l_2) = \frac{1}{V(\Omega)} \iint_\Omega G(l_1 - f(\mathbf{x}), l_2 - g(A(\mathbf{\Theta}; \mathbf{x}))) d\mathbf{x}$$

where $[G(l_1 - f(\mathbf{x}), l_2 - g(A(\mathbf{\Theta}; \mathbf{x})))]$ represents a two dimensional zero-mean differentiable Gaussian kernel. A similar approach can be considered in defining $p^{f_\Omega}(l_1)$ and $p^{g_\Omega^A}(l_2)$ using a 1D Gaussian kernel. The calculus of variations with a gradient descent method can be used to minimize the cost function and recover the registration parameters θ_i [6]. Examples of such approach for rigid registration are given in [Fig. (1)]. Left Ventricle hand-drawn contours (40) from 2/4-chambers view have been considered and registered to the same target.

Medical imaging is an area where quite often global motion is not a valid answer when solving the registration [5]. Local deformations are a complementary component

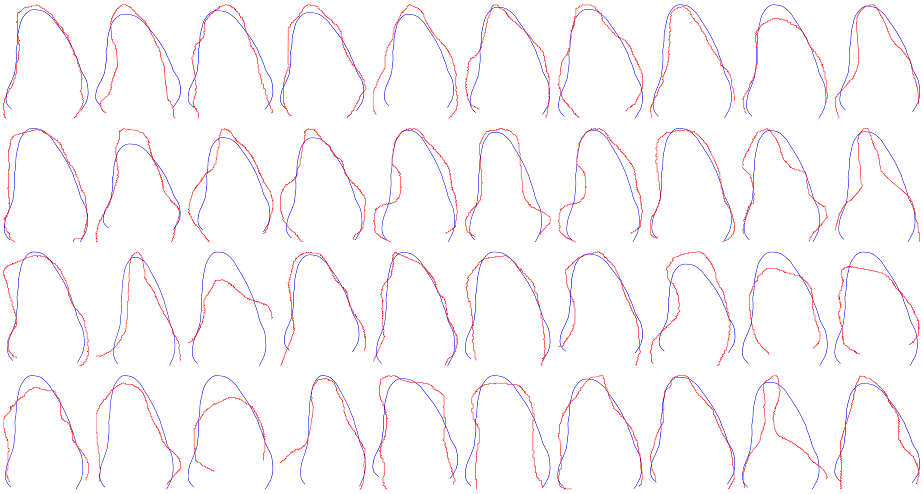

Fig. 1. Rigid Registration for User-Determined Ground Truth (Systole) from the Left Ventricle of Ultrasonic Images (multiple views). (blue) target mean shape, (red) registered source shape.

to the global registration model. Dense local motion (warping fields) estimation is an ill-posed problem since the number of variables to be recovered is larger than the number of available constraints. Smoothness as well as other form of constraints were employed to cope with this limitation.

3 Free-Form Deformations & Local Registration

In the proposed framework, a global motion model (τ) is recovered using the mutual information criterion. One can use such model to transform the source shape $\mathcal{D}$ to a new shape $\hat{\mathcal{D}} = \tau(\mathcal{D})$. Then, local registration is equivalent with recovering a pixel-wise deformation field that creates visual correspondences between the implicit representation [$\Phi_{\mathcal{S}}$] of the target shape $\mathcal{S}$ and the implicit representation [$\Phi_{\hat{\mathcal{D}}}$] of the transformed source shape $\hat{\mathcal{D}}$.

Such local deformation field $L(\mathbf{\Theta}; \mathbf{x})$ can be recovered either using standard optical flow constraints [8], or through the use of space warping techniques like the free form deformations method [9], which is a popular approach in graphics, animation and rendering [4]. Opposite to optical flow techniques, FFD techniques provide a dense registration paradigm that have flexibility (e.g, support various basis splines), implicitly enforce smoothness constraints, exhibit robustness to noise and are suitable for modelling large and small non-rigid deformations. Furthermore, the recovered deformation field is continuous, and under certain conditions, guarantees a one-to-one mapping.

The essence of FFD is to deform an object by manipulating a regular control lattice P overlaid on its volumetric embedding space. We consider an Incremental Cubic B-spline Free Form Deformation (FFD) to model the local transformation L. To this end, dense registration is achieved by evolving a control lattice P according to a deformation improvement [δP]. The inference problem is solved with respect to - the parameters of FFD - the control lattice coordinates.

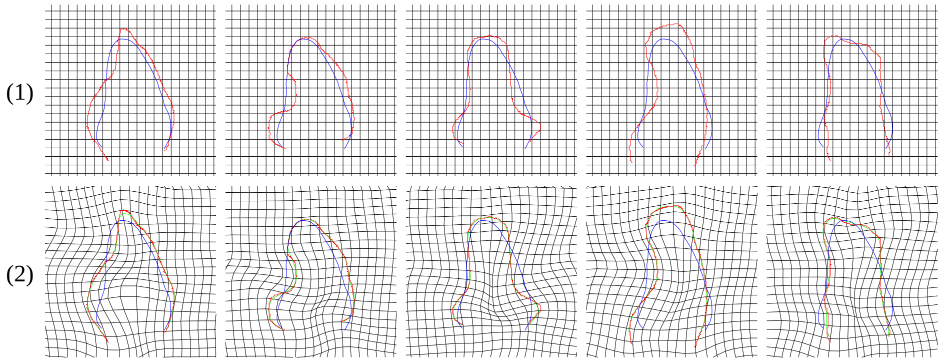

Fig. 2. Local Non-rigid registration using Incremental FFD. (1) initial undeformed grid overlaid on global rigid registration result (blue - mean reference shape), (2) deformed grid to map the reference shape to various training shapes. Each column corresponds to a different trial.

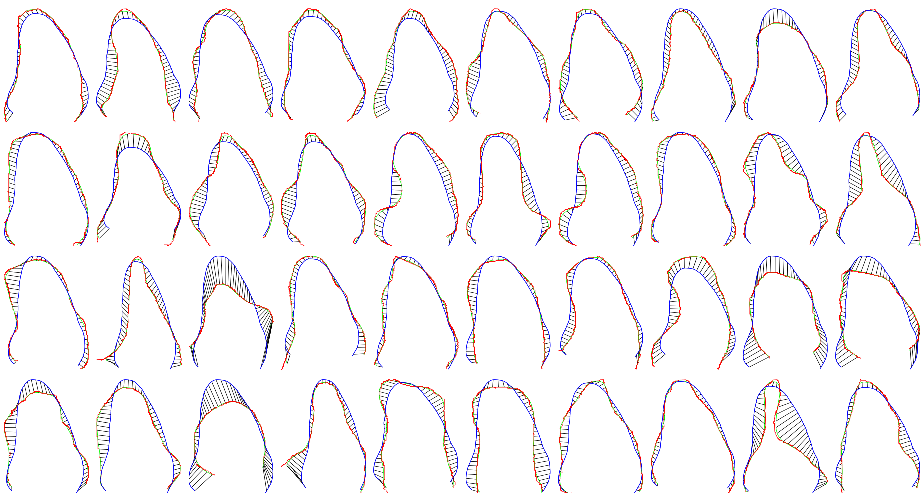

Fig. 3. Examples of establishing correspondences using Incremental FFD. (red) Global registration result, (blue) target mean shape, (dark lines) correspondences for a fixed set of points on the mean shape & (green) the projection of mean shape on each training shape.

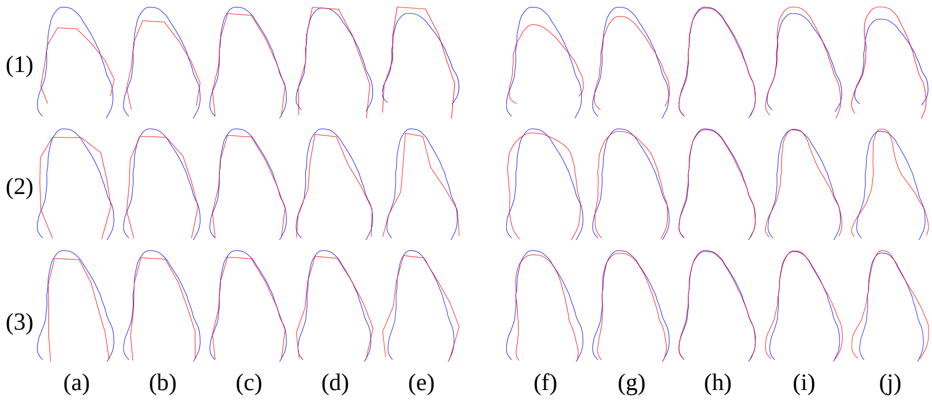

Fig. 4. Principle Component Analysis modelling for the systolic Left Ventricle shapes in Ultrasonic images using the established local correspondences. Changing modes of variation from $-2\sqrt{\lambda_i}$ to $2\sqrt{\lambda_i}$: (1) first mode, (2) second mode, (3) third mode; (a)-(e) 10-points based model, (f)-(j) 80-points based model.

Let us consider a regular lattice of control points

$$P_{m,n} = (P^x_{m,n}, P^y_{m,n});\ m = 1, ..., M, \quad n = 1, ..., N$$

overlaid to a region

$$\Gamma_c = \{\mathbf{x}\} = \{(x, y)|1 \leq x \leq X, 1 \leq y \leq Y\}$$

in the embedding space that encloses the source structure. Let us denote the initial configuration of the control lattice as P^0, and the deforming control lattice as $P = P^0 + \delta P$. Under these assumptions, the incremental FFD parameters are the deformations of the control points in both directions (x, y);

$$\mathbf{\Theta} = \{(\delta P^x_{m,n}, \delta P^y_{m,n})\};\ (m, n) \in [1, M] \times [1, N]$$

The deformed position of a pixel $\mathbf{x} = (x, y)$ given the deformation of the control lattice from P^0 to P, is defined in terms of a tensor product of Cubic B-spline:

$$L(\mathbf{\Theta}; \mathbf{x}) = \mathbf{x} + \delta L(\mathbf{\Theta}; \mathbf{x}) = \sum_{k=0}^{3} \sum_{l=0}^{3} B_k(u) B_l(v) (P^0_{i+k,j+l} + \delta P_{i+k,j+l})$$

where $i = \lfloor \frac{x}{X} \cdot (M-1) \rfloor + 1,\ j = \lfloor \frac{y}{Y} \cdot (N-1) \rfloor + 1$. The terms of the deformation component refer to (i) $\delta P_{i+l,j+l}$, $(k, l) \in [0, 3] \times [0, 3]$ are the deformations of pixel $\mathbf{x}$'s (sixteen) adjacent control points, (ii) $\delta L(\mathbf{x})$ is the incremental deformation at pixel $\mathbf{x}$, and (iii) $B_k(u)$ is the k^{th} basis function of a Cubic B-spline ($B_l(v)$ is similarly defined) [6].

Local registration now is equivalent with finding the lattice P configuration such that the overlaid structures coincide. Since the structures correspond to the distance transform of the target shape - [$\Phi_{\mathcal{S}}$], and the distance transform of its globally aligned source shape - [$\Phi_{\hat{\mathcal{D}}}$], the Sum of Squared Differences (SSD) can be considered as the data-driven term to recover the deformation field $L(\mathbf{\Theta}; \mathbf{x})$);

$$E_{data}(\mathbf{\Theta}) = \iint_{\Omega} \left(\Phi_{\hat{\mathcal{D}}}(\mathbf{x}) - \Phi_{\mathcal{S}}(L(\mathbf{\Theta}; \mathbf{x})) \right)^2 d\mathbf{x}$$

In order to further preserve the regularity of the recovered registration flow, one can consider an additional smoothness term on the local deformation field δL. We consider a computationally efficient smoothness term:

$$E_{smoothness}(\mathbf{\Theta}) = \iint_{\Omega} \left(\left|\left| \frac{\partial \delta L(\mathbf{\Theta}; \mathbf{x})}{\partial x} \right|\right|^2 + \left|\left| \frac{\partial \delta L(\mathbf{\Theta}; \mathbf{x})}{\partial y} \right|\right|^2 \right) \, d\mathbf{x}$$

Such smoothness term is based on an error norm with known limitations. One can replace this smoothness component with more elaborated norms. Within the proposed framework, an implicit smoothness constraint is also imposed by the B-Spline FFD, which guarantees C^1 continuity at control points and C^2 continuity everywhere else. Therefore there is no need for introducing complex and computationally expensive regularization components.

The data-driven term and the smoothness constraints term can now be integrated to recover the local deformation component of the registration and solving the correspondence problem: $E(\mathbf{\Theta}) = E_{data}(\mathbf{\Theta}) + \alpha E_{smoothness}(\mathbf{\Theta})$, where α is the constant balancing the contribution of the two terms. The calculus of variations and a gradient descent method can be used to optimize such objective function [6]. The performance of the proposed framework on the Systolic Left Ventricle dataset is demonstrated in [Fig. (2)] (FFD grid deformations) and [Fig. (2),(3)] (established local correspondences).

4 Building Compact Representations

Let us now assume the existence of n ground truth examples $\phi_{i=1...n}$ in a training set for a structure of interest [Fig. (1)]. Registering these examples to a common pose and establishing local correspondences is required prior of creating a compact statistical representation. Then Principle Component Analysis (PCA) can be applied to capture the statistics of the corresponding elements across the training examples. PCA refers to a linear transformation of variables that retains - for a given number n of operators - the largest amount of variation within the training data, according to: $\phi = \overline{\phi} + \Sigma_{j=1}^{m} b_j \; U_j$, where $\overline{\phi}$ is the mean shape, m is the number of retained modes of variation, U_j are these modes (eigenvectors), and b_j are linear weight factors within the allowable range defined by the eigenvalues.

The most critical part of such analysis process is the representation of the training examples using the same number of elements. Each element corresponds to the same location on the standard atlas of the structure of interest. One simplistic approach on establishing local correspondences is based on uniform (equal-distance) sampling. Parametric approximation of the training set using the same number of basic components is an alternative. Such methods require explicit parameterization of the shapes, finding at least one pair of correspondence between landmark points, or finding the correspondence between parameterization schemes. One can argue that finding straightforwardly extensible solutions in handling structures of high dimensionality and those with complex topology (e.g. multi-parts) in such methods is an unrealistic assumption. Furthermore, without the support of dense optimal local registration, the resulting correspondences from these methods are often non-intuitive and prone to noise.

The proposed global-to-local registration framework can cope with the above limitations. To this end, first all contours are registered to the same target using the global rigid transformation model. Such selection introduces bias on the modelling phase. To overcome this limitation, we apply an iterative schema where the mean shape approximated as in [8] can be used as target to perform registration. Once global registration is completed, local correspondences between the reference shape and the examples of the training set are established using the free-form deformation approach. Then according to the desired dimensionality of the model, one can sample the reference shape, use the one-to-one dense local deformation field to recover the corresponding positions within the training set, and to extract the compact representation for the set of training examples using PCA [Fig. (4)]. There is a compromise between the model complexity (number of elements) and the accuracy of the compact statistical representation.

5 Discussion

In this paper we have proposed a novel variational technique for establishing local correspondences between shapes. Registration has been approached in an hierarchical manner. First, a rigid motion model has been determined between the target and the source and then a dense registration field was recovered, supplementary to the global motion model. Shapes were considered in a higher dimension, the space of distance transforms. Such space when combined with mutual information results to a powerful global registration paradigm. The use of free-form deformations was considered to recover the local registration component in the space of distance transforms leading to one-to-one mapping between the source and the target. Principal component analysis was considered for modelling of the registered set of training examples. The proposed framework is completely automatic, efficient. It can be used for handling very large training datasets of anatomical structures (2D or 3D). The compact statistical representations (prior models) thus built can be used in a similar manner to the Active Shape Models (ASM) for model search in image segmentation and tracking applications.

Acknowledgements

This work was supported by NIH/NLM grant.

References

1. H. Chui and A. Rangarajan. A New Algorithm for Non-Rigid Point Matching. In *IEEE CVPR*, pages II: 44–51, 2000.
2. A. Collignon, F. Maes, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multi-modality image registration using information theory. In *IPMI*, pages 263–274, 1995.
3. T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61:38–59, 1995.
4. P. Faloutsos, M. van de Panne, and D. Terzopoulos. Dynamic Free-Form Deformations for Animation Synthesis. *IEEE Transactions on Visualization and Computer Graphics*, 3:201–214, 1997.
5. J. Feldmar and N. Ayache. Rigid, affine and locally affine registration of free-form surfaces. *IJCV*, 18:99–119, 1996.
6. X. Huang, N. Paragios, and D. Metaxas. Registration of Structures in Arbitrary Dimensions: Implicit Representations, Mutual Information & Free-Form Deformations. Technical Report DCS-TR-0520, Division of Computer & Information Science, Rutgers University, 2003.
7. J. Maintz and M. Viergever. A Survey for Medical Image Registration. *Medical Image Analysis*, 2:1–36, 1998.

8. N. Paragios, M. Rousson, and V. Ramesh. Matching Distance Functions: A Shape-to-Area Variational Approach for Global-to-Local Registration. In *ECCV*, pages II:775–790, 2002.
9. D. Rueckert, L.I. Sonda, C. Hayes, D. Hill, M. Leach, and D. Hawkes. Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE TMI*, 18:712–721, 1999.
10. T. Sederberg and S. Parry. Free-Form Deformation of Solid Geometric Models. In *ACM SIGGRAPH*, volume 4, pages 151–160, 1986.
11. L. staib and J. Duncan. Boundary Finding with Parametrically Deformable Models. In *PAMI*, 14(11), pp. 1061-1075, 1992
12. B. Vemuri, S. Huang, S. Sahni, C.M. Leonard, C. Mohr, R. Gilmore, and J. Fitzsimmons. An efficient motion estimator with application to medical image registration. *Medical Image Analysis*, 2:79–98, 1998.
13. P. Viola and W. Wells. Aligment by Maximization of Mutual Information. In *IEEE ICCV*, pages 16–23, 1995.
14. Y. Wang and L. Staib. Elastic Model Based Non-rigid Registration Incorporating Statistical Shape Information. In *MICCAI*, pages 1162–1173, 1998.