On the Representation of Shapes Using Implicit Functions

N. Paragios [1], M. Taron [2], X. Huang[3], M. Rousson[4] and D. Metaxas[5]

ABSTRACT In this chapter, we explore shape representation, registration and modeling through implicit functions. To this end, we propose novel techniques to global and local registration of shapes through the alignment of the corresponding distance transforms, that consists of defining objective functions that minimize metrics between the implicit representations of shapes.
Registration methods in the space of implicit functions like the Sum of Squares Differences (SSD) that can account for primitive transformations (similarity) and more advanced methods like mutual information that are able to handle more generic parametric transformations are considered. Toward addressing local correspondences we also propose an objective function on the space of implicit representations where the displacement field is represented with a free form deformation that can guarantee one-to-one mapping. In order to address outliers as well as introduce confidence in the registration process, we extend our registration paradigm to estimate uncertainties through the formulation of local registration as a statistical inference problem in the space of implicit functions. Validation of the method through various applications is proposed: (i) parametric shape modeling and segmentation through active shapes for medical image analysis, (ii) variable bandwidth non-parametric shape modeling for recognition and (iii) object extraction through a level set method. Promising results demonstrate the potentials of implicit shape representations.

1 Introduction

Shape modeling is a critical component in various applications of imaging and vision with registration[54] being its most challenging aspect. In the most gen-

[1]C.E.R.T.I.S. - Ecole Nationale des Ponts et Chaussees - 6-8 Avenue Blaise Pascal, France, 77455 Champs sur Marne, France - mailto:nikos.paragios@ieee.org

[2]C.E.R.T.I.S. - Ecole Nationale des Ponts et Chaussees - 6-8 Avenue Blaise Pascal, France, 77455 Champs sur Marne, France - mailto:maxime.taron@certis.enpc.fr

[3]Computer-Aided Diagnosis & Therapy - Siemens Medical Solutions - 51 Valley Stream Parkway-Malvern, PA 19355-1406 USA - mailto:xiaolei.huang@siemens.com

[4]Imaging and Visualization Department - Siemens Corporate Research - 755 College Road East, Princeton, NJ 08540, USA - mailto:mikael.rousson@siemens.com

[5]Departments of Computer Science & Biomedical Engineering Rutgers - The State University of New Jersey - 110 Frelinghuysen Road, Piscataway, NJ 08854, USA - mailto:dnm@cs.rutgers.edu

eral case, given a source and a target shape registration consists of recovering a transformation that creates some correspondence between the two shapes. Segmentation, recognition, indexing and retrieval, tracking and animations are some examples where registration is needed. Often, segmentation consists of deforming a prior shape model to the image while recognition consists of element-wise comparison between structures of interest that were aligned. Similar concept is applicable when addressing indexing and retrieval, while tracking can be formulated as a registration problem [58] of the target from one image to the next.

Global registration, refers to parametric transformations with a small number of degrees of freedom, while non-rigid local registration aims to establish dense correspondences between the two shapes and in principle can have an infinite number of parameters. The importance of shape registration/modeling in computational vision was a motivation for researchers and therefore one can find a vast prior art [1, 2, 11, 13, 15, 34, 43, 59]. Given the definition of the registration problem, one can classify existing methods according to three aspects: (i) representation for the structures of interest, (ii) nature of plausible transformations, and (iii) mathematical framework used to recover the optimal registration parameters.

1. **Shape representation** refers to the selection of an appropriate representation for the shapes. Clouds of points [1, 11], parametric curves & surfaces [15], Fourier descriptors [48], medial axes [44], and implicit distance functions [34] were considered.

2. **Transformation** can be either global or local. Global parametric models like rigid, similarity, affine and perspective among others, are applicable to the entire shape. On the other hand, local alignment is defined at the local shape element level and used to represent non-rigid deformations leading to dense correspondences between shapes; optical flow [7, 34], Thin Plate Spline (TPS) [1, 11], and space deformation techniques such as Free Form Deformations (FFD) [42, 45] are some examples.

3. **Registration criterion** is a mathematical framework used to recover the optimal registration parameters given the shape representation and the nature transformation. One can find in the literature two dominant techniques: (i) estimation of explicit geometric feature correspondences that are used in a second stage to determine the transformation parameters [1, 2, 11, 59], and (ii) recovery of the optimal transformation parameters through the minimization of an objective function [7, 12, 34, 55].

Point clouds is a quite popular and rather intuitive generic shape representation [1, 11] with certain strengths and numerous limitations. On one hand one can adopt certain freedom on the shape topologies either on 2D or in 3D while on the other hand the sampling rule used to determine the number of shape basic elements as well as their distribution can have a substantial effect on the registration process. In particular when addressing local registration one should be cautious toward introducing similar or quite dense representation both for the source and

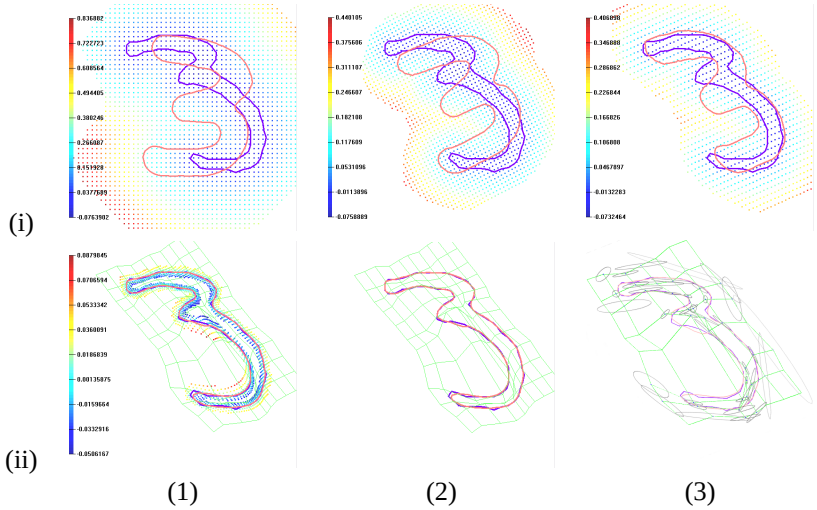

FIGURE 1. *Registration Pipeline; (i.1) implicit representations of the source and the target, (i.2) rigid registration, (i.3) affine registration, (ii.1) free form deformation of the grid, (ii.2) local correspondences between the source and the target, (ii.3) uncertainties estimates of the registration process.*

the target shape. In the opposite case correspondences could be meaningless in the presence of improper sampling and lead to erroneous registration results.

Parametric curve representations of shapes [15, 28] are more appropriate selections at least in the case of 2D & 3D that could provide meaningful correspondences since shape structure could be recovered through efficient interpolation techniques with the expense of being quite inefficient when handling complex topologies or shapes in higher dimensions. Fourier descriptors [48] and medial axis [44]while being promising shape representations to measure similarity between shapes they become quite inefficient for registration and in particular when seeking local correspondences between the basic shape elements that is in the general the case for parametric representations. Local correspondences require proper selection of the basis elements, the position and the number of control points that in order to be optimal has to be shape-driven and cannot be done a priori.

Rigid (translation and rotation), similarity (translation, rotation, and isotropic scaling) and affine (translation, rotation, isotropic or anisotropic scaling, as well as shearing) are the most frequent models to address global registration. The case of local registration is more complex since the number of constraints is inferior to the number of parameters to be recovered. Therefore additional constraints are often introduced like in the case of optical flow estimation through a regularization process. However, shape registration is a different problem that optical flow estimation since advanced regularization and smoothness constraints [7, 34] could fail to preserve the topology of the source shape while at the same time cannot guarantee an one-to-one mapping. Such a limitation can be addressed at some partial fashion through a thin plate spline (TPS) model [1, 11] with expense of re-

covering explicit correspondences between landmark points along the source and the target shape. Once such correspondences are recovered, one can estimate the local deformation field through a TPS interpolation on the landmarks. Selection of landmark points as well as establishing correspondences between these points are the most challenging steps within such an approach.

Numerous shape alignment methods were proposed to address global as well as local registration. The use of explicit feature correspondences toward estimating the transformation [1, 59, 2, 11] is the most primitive approach to recover registration parameters through robust optimization techniques. Such an approach is rather sequential and therefore it heavily depends on the feature extraction process. Furthermore, the registration problem can become under-constrained, especially in the case of non-rigid registration when many reliable correspondences are needed in order to solve for the local deformation parameters.

A different approach consists of addressing registration as a statistical estimation problem [25] through successive steps. Within each step the uncertainty in the estimates is being computed [50] and is used to guide further steps of the overall algorithm [36]. In [49] the covariance matrix is used within an Iterated Closest Point (ICP) algorithm to sample the correspondences so that registration is well-constrained in all directions in parameter space. Last, but not least in [47] local deformation and uncertainties are simultaneously recovered for the optical flow estimation problem through a Gaussian noise assumption on the observation. Prior art leads to the conclusion that shape modeling and registration are open research topics.

In this chapter we propose an alternative representation to the existing methods that introduces novel elements in each component of the registration process. To this end, first we assume shape representations to be implicit functions [34, 35] (Euclidean distance transforms). Such representations are invariant to translation and rotation, can account for scale variations and cope to some extent with noise and and local deformations [61]. Registration is addressed in a complete fashion through a global and a local component.

Objective functions that aim to account for global transformations in the space of implicit representations are introduced. Global registration models of increasing complexity are addressed like rigid, similarity [34] through an SSD approach or affine, homographic [23] through a mutual information criterion. Free form deformations and higher order polynomials [23] are used to encode local deformations in the space of implicit functions. Such a model is robust to the presence of outliers, can provide a one-to-one mapping between the source and the target shape while being able to preserve their topologies. Shapes refer to components of variable complexity and different degrees of freedom while at the same time are often corrupted by noise. Such information is to be accounted for, and therefore we propose a statistical inference approach that associates certain uncertainties on the local deformation field [52] leading to a complete local registration paradigm that is shown in [FIGURE (1)].

Toward validation of the proposed method, we consider parametric shape modeling for the segmentation of the left ventricle in ultrasound images, non-parametric

variable bandwidth shape modeling for shape recognition and implicit active contours for knowledge-driven object extraction within a level set approach

The remainder of the chapter is organized as follows: In section 2 we present the implicit shape representation, its properties and its use for global alignment. Local registration is introduced in section 3, along with the estimation of the deformation uncertainties. Validation of the method is part of section 4, while discussion and conclusions are presented in section 5.

2 Implicit Representation of Shapes & Global Registration

Distance transforms have been popular in image analysis for a while. One can refer to the famous Chamfer transform [3] often used for object extraction and to the use of implicit representations (often called level set methods [16, 17, 31]) for curve propagation.

Such representations are heavily considered in the domain of computational vision because of their intrinsic and parameter-free nature while being able to describe multi-component shapes and structures. Last, but not least a straightforward estimation is possible of various geometric properties of the shape (normal, curvature, skeleton) often needed for registration, curve propagation, etc. Let us consider a closed curve $\mathcal{S}$ that defines a bi-modal image partition of Ω; the region that is enclosed by $\mathcal{S}$, $[\mathcal{R}_{\mathcal{S}}]$, and the background $[\Omega - \mathcal{R}_{\mathcal{S}}]$ and an implicit level set representation of $\mathcal{S}$

$$\phi_{\mathcal{S}}(\mathbf{x}) = \begin{cases} 0, & \mathbf{x} \in \mathcal{S} \\ +\mathcal{D}(\mathbf{x}, \mathcal{S}) > 0, & \mathbf{x} \in \mathcal{R}_{\mathcal{S}} \\ -\mathcal{D}(\mathbf{x}, \mathcal{S}) < 0, & \mathbf{x} \in [\Omega - \mathcal{R}_{\mathcal{S}}] \end{cases}$$

that embeds $\mathcal{S}$ in a higher dimensional distance function $\phi : \Omega \rightarrow R^+$ that is assumed to be a Lipschitz function of the Euclidean distance from the shape $\mathcal{S}$,

$$\mathcal{D}(\mathbf{x}, \mathcal{S}) = \min_{\mathbf{y} \in \mathcal{S}} \{\|\mathbf{x} - \mathbf{y}\|_2\},$$

Such a representation can be constructed in various ways, simple two passes in the image [3] could provide an approximate form, while more advanced methods like the fast marching algorithm [46] or PDE-based techniques [51] can also be considered.

Such implicit shape representation provides a feature space in which objective functions that are optimized using a gradient descent method can be conveniently used. One can prove that the gradient of the embedding distance function is a unit vector in the normal direction of the shape and the representation satisfies a sufficient condition for the convergence of gradient descent methods, which requires continuous first derivatives. Furthermore, the use of the implicit representation

provides additional support to the registration process around the shape boundaries and facilitates the imposition of smoothness constraints, since one would like to align the original structures as well as their clones that are positioned coherently in the image/volume plane. Last, but not least, implicit shape representations are invariant to rigid transformations while the effect of isotropic scale changes can be accounted for.

Let us consider a global transformation $\mathcal{A}$. Suppose $\hat{\phi}$ is the level set obtained after transformation of ϕ by $\mathcal{A}$. The zero-crossing of $\hat{\phi}$ gives a shape $\hat{\mathcal{S}}$ which corresponds to the original shape $\mathcal{S}$ after being transformed by $\mathcal{A}$ that refer to a rigid transformation with a translation vector $\mathbf{T}$ and a rotation angle $\mathbf{R}$:

$$\mathcal{A}(\mathbf{x}) = \mathbf{R}\,\mathbf{x} + \mathbf{T}$$

One can prove that $\hat{\phi}$ is also the distance transform of $\hat{\mathcal{S}}$. Let us consider $\hat{\mathbf{x}}$ be the location of $\mathbf{x}$ after being displaced according to $\mathcal{A}$. Then, for all $\mathbf{x}$ in the image domain Ω, we have:

$$\begin{aligned}\mathcal{D}(\hat{\mathbf{x}}, \hat{\mathcal{S}}) &= \min_{\hat{\mathbf{y}} \in \hat{\mathcal{S}}} \{\|\hat{\mathbf{x}} - \hat{\mathbf{y}}\|_2\} \\ &= \min_{\mathbf{y} \in \mathcal{S}} \{\|\mathbf{R}\,\mathbf{x} + \mathbf{T} - (\mathbf{R}\,\mathbf{y} + \mathbf{T})\|_2\} = \min_{\mathbf{y} \in \mathcal{S}} \{\|\mathbf{R}\,(\mathbf{x} - \mathbf{y})\|_2\} = \mathcal{D}(\mathbf{x}, \mathcal{S})\end{aligned}$$

that is equivalent of saying that distance transforms are invariant to translation & rotation:

$$\hat{\phi}(\hat{\mathbf{x}}) = \phi(\mathbf{x}) = \mathcal{D}(\mathbf{x}, \mathcal{S})$$
$$\hat{\mathbf{x}} = \mathbf{R}\,\mathbf{x} + \mathbf{T} \Rightarrow \mathcal{D}(\hat{\mathbf{x}}, \hat{S}) = \mathcal{D}(\mathbf{x}, \mathcal{S}).$$

We can now also deduce the effect of adding a scale factor in the transformation: $\mathcal{A}(\mathbf{x}) = \mathbf{s}\,\mathbf{R}\,\mathbf{x} + \mathbf{T}$;

$$\mathcal{D}(\hat{\mathbf{x}}, \hat{\mathcal{S}}) = \min_{\hat{\mathbf{y}} \in \hat{\mathcal{S}}} \{\|\hat{\mathbf{x}} - \hat{\mathbf{y}}\|_2\} = \min_{\mathbf{y} \in \mathcal{S}} \{\|\mathbf{s}\,\mathbf{R}\,(\mathbf{x} - \mathbf{y})\|_2\} \; = \mathbf{s}\,\mathcal{D}(\mathbf{x}, \mathcal{S})$$

Since for the directly transformed level set image representation $\hat{\phi}$ we have $\hat{\phi}(\hat{\mathbf{x}}) = \mathcal{D}(\mathbf{x}, \mathcal{S}) = \frac{1}{\mathbf{s}}\mathcal{D}(\hat{\mathbf{x}}, \hat{\mathcal{S}})$, we can derive the distance transform of $\hat{\mathcal{S}}$ by simply multiplying the scale factor $\mathbf{s}$ to $\hat{\phi}$. One can now address a similarity invariant registration through the definition of an objective function in the space of implicit representations of shapes.

Global parametric registration consists of recovering a transformation $\mathcal{A}$ that creates pixel-wise intensity correspondences between the implicit representation of the source $\phi_{\mathcal{S}}$ and the target $\phi_{\mathcal{T}}$ shape. Similarity or affine nature of transformations have been mostly considered with either 4 or 6 degrees of freedom.

2.1 Similarity Registration of Shapes

In the case of similarity transformations [$\mathcal{A}(\mathbf{x}) = \mathbf{s}\,\mathbf{R}\,\mathbf{x} + \mathbf{T}$], given the explicit relation between the implicit representations of the source and the target [$\hat{\phi}(\hat{\mathbf{x}}) =$

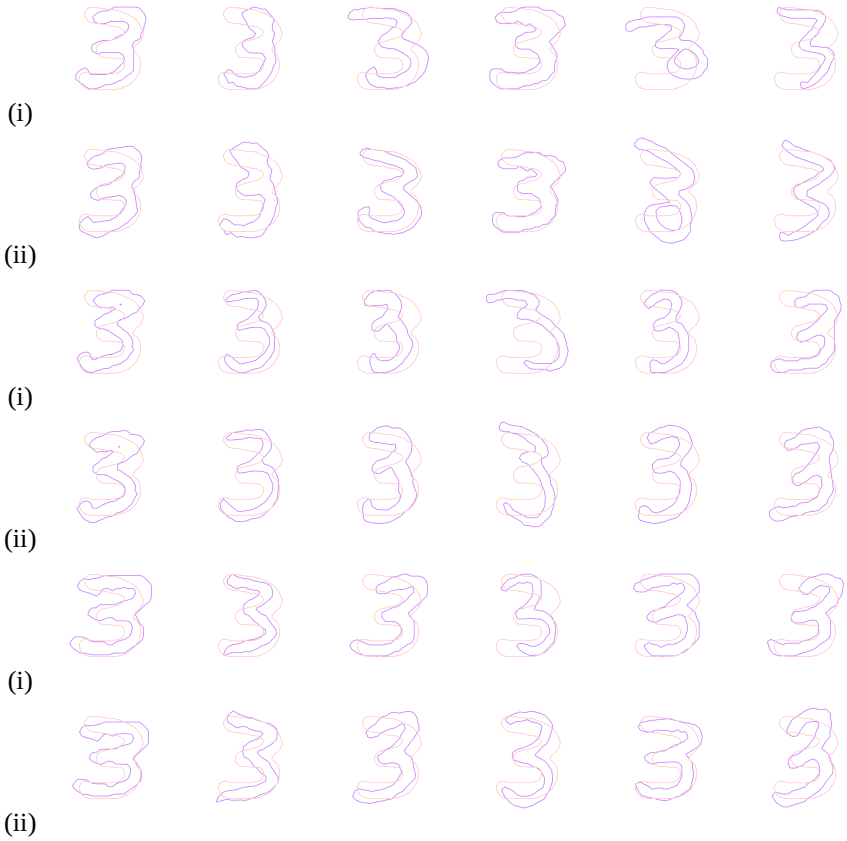

FIGURE 2. *Global registration using the similarity transformation model.* (i) initial pose (ii) similarity-invariant registration (the same target contour is used).

$\mathcal{D}(\mathbf{x}, \mathcal{S}) = \frac{1}{\mathbf{s}}\mathcal{D}(\hat{\mathbf{x}}, \hat{\mathcal{S}})$], a primitive criterion to recover registration parameters refer to the Sum-of-Squared-Differences (SSD);

$$E(\mathbf{s}, \mathbf{R}, \mathbf{T}) = \iint_{\Omega} \left(\mathbf{s}\,\phi_{\mathcal{S}}(\mathbf{x}) - \phi_{\mathcal{T}}(\mathcal{A}(\mathbf{x}))\right)^2 \, d\mathbf{x},$$

which measures the dissimilarity between the intensity values (i.e. distance values) of pixels in a sample image domain on the source representation and that of the projected pixels on the target representation according to the transformation $\mathcal{A}$ that is a computational expensive approach. One can address such a concern through the adoption of a narrow band in the embedding space as the sample domain. The use of a band indicator function $\chi_\alpha(\phi_{\mathcal{S}}(\mathbf{x}))$ can reduce registration domain

$$E(\mathbf{s}, \mathbf{R}, \mathbf{T}) = \iint_{\Omega} \chi_\alpha(\phi_{\mathcal{S}}(\mathbf{x})) \left(\mathbf{s}\,\phi_{\mathcal{S}}(\mathbf{x}) - \phi_{\mathcal{T}}(\mathcal{A}(\mathbf{x}))\right)^2 \, d\mathbf{x},$$

where $\chi_\alpha(\phi_{\mathcal{S}}(\mathbf{x}))$ is given by

$$\chi_\alpha(\phi) = \begin{cases} 0, & |\phi| > \alpha, \\ 1, & \text{otherwise} \end{cases}$$

Such a modified function accounts for pixels (isophotes) within a range of distance α from the source shape and their projections on the target are considered in the optimization process. One can consider the calculus of variations and a gradient-descent based method to recover the optimal registration parameters;

$$\frac{d}{dt}\mathbf{T} = 2 \iint\limits_{\Omega} \chi_\alpha(\phi_{\mathcal{S}}(\mathbf{x}))\, \psi(\mathbf{x})\, \nabla\phi_{\mathcal{T}}(\mathcal{A}(x))\, d\mathbf{x}$$

$$\frac{d}{dt}\mathbf{s} = -2 \iint\limits_{\Omega} \chi_\alpha(\phi_{\mathcal{S}}(\mathbf{x}))\, \psi(\mathbf{x})\, \left(\phi_{\mathcal{S}}(\mathbf{x}) - \nabla\phi_{\mathcal{T}}(\mathcal{A}(x)) \cdot \mathbf{R}\mathbf{x}\right) d\mathbf{x}$$

$$\frac{d}{dt}\theta_i = 2 \iint\limits_{\Omega} \chi_\alpha(\phi_{\mathcal{S}}(\mathbf{x}))\, \psi(\mathbf{x})\, \nabla\phi_{\mathcal{T}}(\mathcal{A}(x)) \cdot \nabla_{\theta_i}(\mathcal{A}(\mathbf{x}))\, d\mathbf{x}, \quad 1 \leq i \leq p$$

where $(\mathbf{s}\,\phi_{\mathcal{S}}(\mathbf{x}) - \phi_{\mathcal{T}}(\mathcal{A}(\mathbf{x})))$ is the residual error that has been replaced by $\psi(\mathbf{x})$ and p is the number of rotation angles [$\mathbf{R} = (\theta_i)$]. Examples of such a registration process are shown in [FIGURE (2)]. Based on the experimental results one can claim that shapes undergoing similarity transformations are properly registered.

On the other hand dealing with rather generic parametric transformations like affine is not straightforward. In principle the effect of such transformations cannot be predicted in the space of implicit representations and therefor the distance function of the transformed shape is not available and it has to be recomputed, that is a rather inefficient procedure within iterative processes like the one we have adopted. One can overcome such a limitation through the consideration of an alternative affine-invariant objective function that does not seek point-to-point correspondences (like the SSD case) between the implicit representations of the source and the target. Mutual Information [37] is a rather convenient registration paradigm that satisfies such constraints in particular when registering distance transforms [22, 21].

2.2 Affine-invariant Registration of Shapes

Scale variations can be considered as a global illumination changes in the space of distance transforms. Therefore, registration under scale variations is equivalent with matching different modalities that refer to the same structure of interest. Mutual Information is an information theoretic criterion that is an invariant technique according to a monotonic transformation of the two input random variables. 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 $\phi_{\mathcal{S}}$ as f, and **(ii)** the target representation $\phi_{\mathcal{T}}$ 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 $\mathcal{A}$ such that the mutual

information between $f_\Omega = f(\Omega)$ and $g_\Omega^{\mathcal{A}} = g\big(\mathcal{A}(\mathbf{\Theta};\Omega)\big)$ is maximized for a given sample domain Ω;

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

where $\mathcal{H}$ represents the differential entropy. Such a quantity represents a measure of uncertainty, variability or complexity and consists of three components: (i) the entropy of the model, (ii) the entropy of the projection of the model given the transformation, and (iii) the joint entropy 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) to perform global registration that is equivalent with minimizing:

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

where (i) p^{f_Ω} corresponds to the probability density in f_Ω $\big(\ [\phi_{\mathcal{S}}(\Omega)]\big)$, (ii) $p^{g_\Omega^{\mathcal{A}}}$ corresponds to density in $g_\Omega^{\mathcal{A}}$ $\big(\ [\phi_{\mathcal{T}}(\mathcal{A}(\mathbf{\Theta};\Omega))]\ \big)$, and (iii) $p^{f_\Omega, g_\Omega^{\mathcal{A}}}$ is the joint density. Such framework can account for various global motion models. Toward 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:

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

where $V(\Omega)$ represents the volume of Ω and $G(l_1 - f(\mathbf{x}), l_2 - g(\mathcal{A}(\mathbf{\Theta};\mathbf{x})))$ is a two dimensional zero-mean differentiable Gaussian kernel. A similar approach can be considered in defining $p^{f_\Omega}(l_1)$ and $p^{g_\Omega^{\mathcal{A}}}(l_2)$ using a 1D Gaussian kernel. The calculus of variations with a gradient descent method [23] can be used to minimize the cost function and recover the registration parameters θ_i:

$$\begin{aligned} \frac{\partial E}{\partial \theta_i} =& -\iint_{\mathcal{R}^2} \Big(1 + log \frac{p^{f_\Omega, g_\Omega^{\mathcal{A}}}(l_1, l_2)}{p^{f_\Omega}(l_1) p^{g_\Omega^{\mathcal{A}}}(l_2)}\Big) \\ & \Big[\frac{1}{V(\Omega)} \iint_\Omega -G_\beta(l_1 - \alpha, l_2 - \beta)\big(\nabla g(\mathcal{A}(\mathbf{\Theta};\mathbf{x})) \cdot \frac{\partial}{\partial \theta_i}\mathcal{A}(\mathbf{\Theta};\mathbf{x})\big)d\mathbf{x}\Big] dl_1 dl_2 \\ =& -\frac{1}{V(\Omega)} \iint_\Omega \Big[\iint_{\mathcal{R}^2} \Big(1 + log \frac{p^{f_\Omega, g_\Omega^{\mathcal{A}}}(l_1, l_2)}{p^{f_\Omega}(l_1) p^{g_\Omega^{\mathcal{A}}}(l_2)}\Big) \\ & \Big(-G_\beta(l_1 - \alpha, l_2 - \beta)\Big) dl_1 dl_2\Big] \big(\nabla g(\mathcal{A}(\mathbf{\Theta};\mathbf{x})) \cdot \frac{\partial}{\partial \theta_i}\mathcal{A}(\mathbf{\Theta};\mathbf{x})\big)d\mathbf{x} \end{aligned}$$

The resulting global registration protocol is the following: given a source and a target shape, implicit representations in the space of distance transforms are recovered. Then, the mutual information criterion is used to estimate the parameters

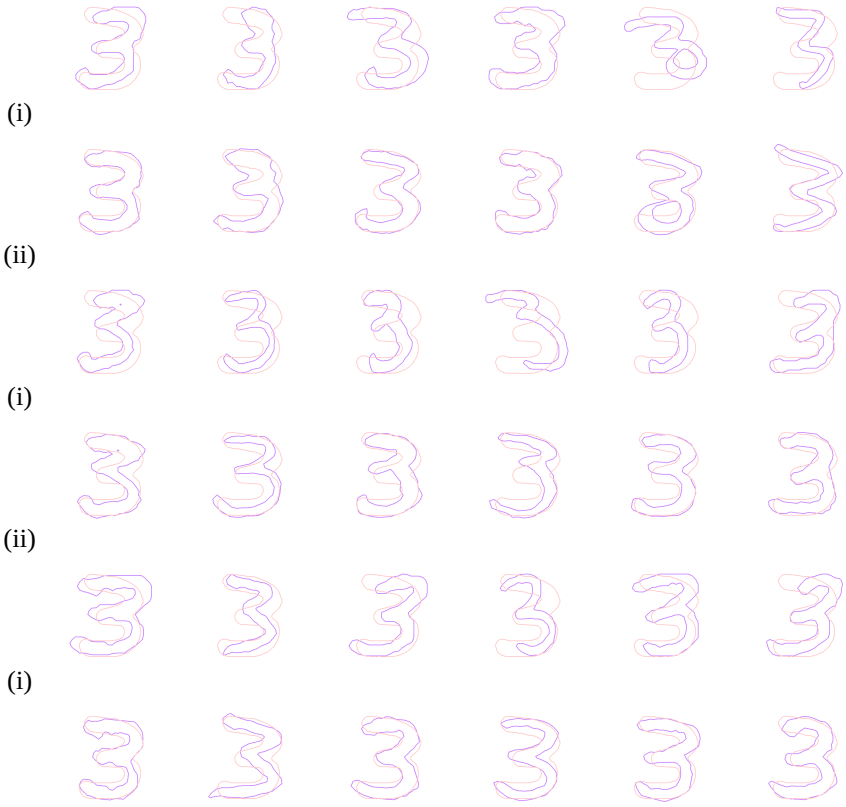

FIGURE 3. Global registration using the affine transformation model. *(i) initial pose (i) affine-invariant registration (the same target contour is used).*

of the optimal transformation between the source and the target implicit representations. Examples of such approach for rigid as well as affine registration are given in [FIGURE (3)]. However, in numerous application domains of computational vision, global transformations are not a proper answer when solving the registration problem like the case of medical imaging [19].

3 Local Registration

Local deformations are a complementary component 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.

In the proposed framework, a global motion model ($\mathcal{A}$) is recovered using dif-

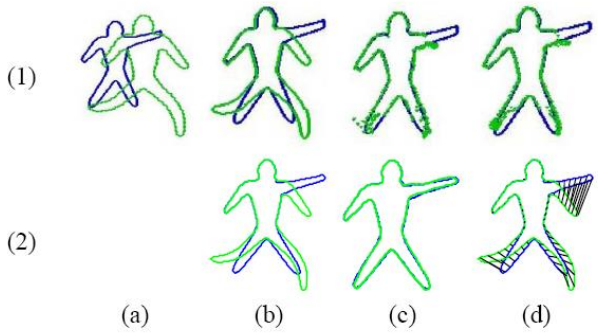

FIGURE 4. Comparison with the non-rigid shape registration algorithm presented in [22]. (1) Results from the algorithm in [34, 35]. (2) Results from our algorithm. (a) Initial poses of the source (in green) and target (in blue) shapes. (b) Alignment after global registration. (1.c) Local registration result; (1.d) local registration with regularization constraints. (2.c) Local registration result; (2.d) established correspondences.

ferent criteria. One can use such model to transform the source shape $\mathcal{S}$ to a new shape $\hat{\mathcal{S}} = \mathcal{A}(\mathcal{S})$ that is the projection of $\mathcal{S}$ to $\mathcal{T}$. Then, local registration is equivalent with recovering a pixel-wise deformation field that creates visual correspondences between the implicit representation $[\phi_{\mathcal{T}}]$ of the target shape $\mathcal{S}$ and the implicit representation $[\phi_{\hat{\mathcal{S}}}]$ of the transformed source shape $\hat{\mathcal{S}}$.

3.1 Free Form Deformations

Such a deformation field $\mathcal{L}(\mathbf{\Theta}; \mathbf{x})$ can be recovered either using standard optical flow constraints [34, 35] or through the use of warping techniques like the free form deformations method [41], which is a popular approach in graphics, animation and rendering [18]. Opposite to optical flow techniques (where smoothness is introduced in a form of an additional constraint), FFD techniques support smoothness constraints in an implicit fashion, exhibit robustness to noise and are suitable for modeling large and small non-rigid deformations. Furthermore, under certain conditions, it can support a dense registration paradigm that is continuous and guarantees a one-to-one mapping. Such a concept and a primitive comparison with the optical flow approach is presented in [FIGURE (4)].

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

Let us consider a regular lattice of control points

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

overlaid to a structure

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

in the embedding space that is the distance transform the source structure once the global registration parameters have been applied. Let us denote the initial configuration of the control lattice as $\mathbf{P}^0$, and the deforming control lattice as $\mathbf{P} = \mathbf{P}^0 + \delta\mathbf{P}$. Under these assumptions, the incremental FFD parameters are the deformations of the control points in both directions (x, y);

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

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

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

where

$$i = \lfloor \frac{x}{X} \cdot M \rfloor - 1,\ j = \lfloor \frac{y}{Y} \cdot N \rfloor - 1$$

and

$$u = \frac{x}{X} \cdot M - \lfloor \frac{x}{X} \cdot M \rfloor,\ v = \frac{y}{Y} \cdot N - \lfloor \frac{y}{Y} \cdot N \rfloor$$

The terms of the deformation component refer to

- $\delta\mathbf{P}_{i+l,j+l}$, $(k, l) \in [0, 3] \times [0, 3]$ consists of the deformations of pixel $\mathbf{x}$'s (sixteen) adjacent control points,
- $\delta\mathcal{L}(\mathbf{x})$ is the incremental deformation at pixel $\mathbf{x}$, and
- $B_k(u)$ is the k^{th} basis function of a Cubic B-spline ($B_l(v)$ is similarly defined):

$$B_0(u) = (1-u)^3/6, \qquad B_1(u) = (3u^3 - 6u^2 + 4)/6$$
$$B_2(u) = (-3u^3 + 3u^2 + 3u + 1)/6, \qquad B_3(u) = u^3/6$$

Local registration now is equivalent with finding the best lattice $\mathbf{P}$ configuration such that the overlaid structures coincide. Since structures correspond to distance transforms of globally aligned shapes, the Sum of Squared Differences (SSD) can be considered as the data-driven term to recover the deformation field $\mathcal{L}(\mathbf{\Theta}; \mathbf{x}))$;

$$E_{data}(\mathcal{L}(\Theta)) = \iint_{\Omega} \chi_\alpha(\phi_{\mathcal{S}}(\mathbf{x}))\big(\phi_{\mathcal{S}}(\mathbf{x}) - \phi_{\mathcal{T}}(\mathcal{L}(\mathbf{\Theta}; \mathbf{x}))\big)^2 d\mathbf{x}$$

The use of such technique to model the local deformation component of the registration introduces in an implicit form some smoothness constraint since displacement refers to a cubic spline interpolation. Therefore, it can deal with a limited level of deformation. In order to further preserve the regularity of the recovered registration flow, one can consider an additional smoothness term on the

FIGURE 5. Global registration using the FFD local transformation model. *(i) initial pose (after affine) (i) deformation of the grid (the same target contour is used).*

deformation field $\delta\mathcal{L}$. We consider a computationally efficient smoothness term:

$$E_{smoothness}(\mathcal{L}(\Theta)) = \iint_{\Omega} \left(|\mathcal{L}_x(\mathbf{\Theta};\mathbf{x})|^2 + |\mathcal{L}_y(\mathbf{\Theta};\mathbf{x})|^2 \right) d\mathbf{x}$$

Such smoothness term is based on a classic error norm that has certain known limitations. One can replace this smoothness component with more elaborated norms like the a regularization term motivated by the thin plate energy functional [57] :

$$E_{smoothness}(\mathcal{L}(\Theta)) = \iint_{\Omega} \left(|\mathcal{L}_{xx}(\mathbf{\Theta};\mathbf{x})|^2 + 2\,|\mathcal{L}_{xy}(\mathbf{\Theta};\mathbf{x})|^2 + |\mathcal{L}_{yy}(\mathbf{\Theta};\mathbf{x})|^2 \right) d\mathbf{x}$$

that can be further simplified in the case of the cubic B-spline and reduced to the quadratic form [52]. One can claim that second order regularization terms are more flexible than the ones of first order since they allow free affine transformations. Within the proposed framework, an implicit smoothness constraint is also

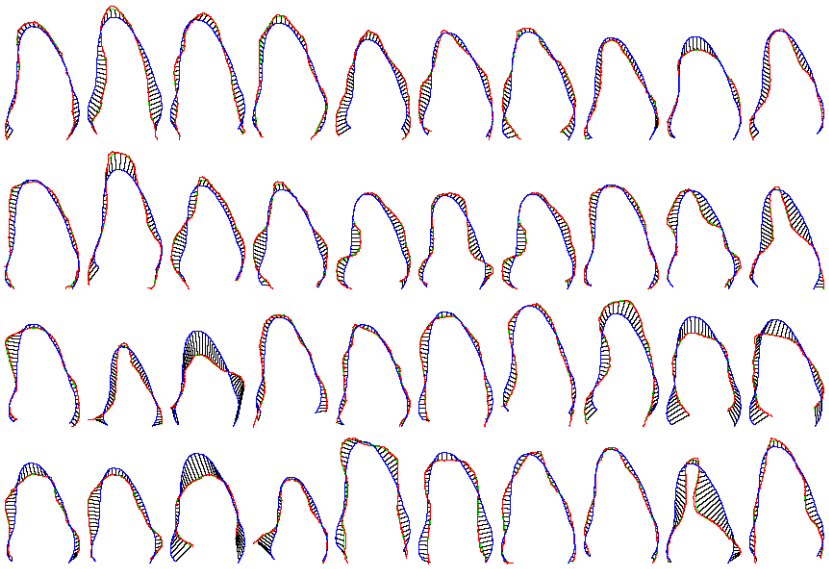

FIGURE 6. *Established correspondences using Incremental FFD. (red) source shapes after global transformations, (blue) target mean shape, (dark lines) correspondences for a fixed set of points on the mean shape. & (green) the projection of mean shape on each training shape.*

imposed by the Spline FFD. Therefore there is not need for introducing complex and computationally expensive regularization components.

The data-driven term and the smoothness constraint component can now be integrated to recover the local deformation component of the registration and solving the correspondence problem:

$$E(\mathbf{\Theta}) = E_{data}(\mathbf{\Theta}) + \beta E_{smoothness}(\mathbf{\Theta})$$

where β 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 [23] (the error-two norm is adopted to impose smoothness).

$$\frac{\partial}{\partial \theta_i} E(\mathbf{\Theta}) = -2 \iint_{\Omega} (\phi_{\mathcal{S}}(\mathbf{x}) - \phi_{\mathcal{T}}(\mathcal{L}(\mathbf{\Theta};\mathbf{x}))) \nabla \phi_{\mathcal{T}}(\mathcal{L}(\mathbf{\Theta};\mathbf{x})) \cdot \frac{\partial}{\partial \theta_i} \delta\mathcal{L}(\mathbf{\Theta};\mathbf{x})) d\mathbf{x} +$$
$$2\alpha \iint_{\Omega} \frac{\partial}{\partial x} \delta\mathcal{L}(\mathbf{\Theta};\mathbf{x}) \cdot \frac{\partial}{\partial \theta_i} \left(\frac{\partial}{\partial x} \delta\mathcal{L}(\mathbf{\Theta};\mathbf{x}) \right) + \frac{\partial}{\partial y} \delta\mathcal{L}(\mathbf{\Theta};\mathbf{x}) \cdot \frac{\partial}{\partial \theta_i} \left(\frac{\partial}{\partial y} \delta\mathcal{L}(\mathbf{\Theta};\mathbf{x}) \right) d\mathbf{x}$$

The flow consists of a data-driven update component and a diffusion term that constraints the parameters of the free form deformation to be locally smooth. The performance on recovering successful local deformations is demonstrated in character registration [FIGURE (5)] and in cardiac shape registration (systolic Left Ventricle dataset) [FIGURE (6)] (established local correspondences). Computational complexity is the most important limitation of such a method that can be

addressed through an efficient multi-level implementation, as shown in [FIGURE (5)].

To this end, multi-resolution control lattices are used according to a coarse-to-fine strategy. A coarser level control lattice is applied first to account for relatively global non-rigid deformations; then the space deformation resulted from the coarse level registration is used to initialize the configuration of a finer resolution control lattice; On this finer level, the local registration process continues to deal with highly local deformations and achieve better matching between the deformed source shape and the target. Generally speaking, the hierarchy of control lattices can have arbitrary number of levels, but typically $2 \sim 3$ levels are sufficient to handle both large and small non-rigid deformations. The layout of the control lattices in the hierarchy can be computed efficiently using a progressive B-spline subdivision algorithm [20]. At each level, we can solve for the incremental deformation of the control lattice while at the end, the overall dense deformation field is defined by these incremental deformations from all levels. In particular, the total deformation $\delta\mathcal{L}(\mathbf{x})$ for a pixel $\mathbf{x}$ in a hierarchy of r levels is:

$$\delta\mathcal{L}(\mathbf{x}) = \sum_{k=1}^{r} \delta\mathcal{L}^k(\Theta^k; \mathbf{x})$$

where $\delta\mathcal{L}^k(\Theta^k; \mathbf{x})$ refers to the deformation improvement at this pixel due to the incremental deformation Θ^k of the kth level control lattice.

The proposed registration framework and the derivations can be straightforwardly extended to 3D. For global registration, parameters of a 3D transformation model can be solved by optimizing the global registration criterion in the 3D sample domain; for local registration, free form deformations can be defined by the 3D tensor product of B-spline polynomials, and the SSD energy functional is defined in the 3D volumetric domain. More details on the 3D formulation for non-rigid FFD registration can be found in [24]. We can use the 3D registration framework to align, register, and stitch 3D face scans captured from range scanners. This problem plays an important role in face modeling, recognition, etc. We show one set of such registration result from our framework in [FIGURE (7)]. The global transformation model consists of translation, scaling, and quaternion-based rotation and the local incremental FFD model uses control lattices in the 3D space and a 3D tensor product of B-spline polynomials. Qualitatively the result after local non-rigid registration can be seen from two views: the front view, and the side view. Quantitatively, the FFD based local registration reduced the average registration error from 8.3 to 1.2, using three resolutions and 20 iterations for each resolution. The total time spent was 4.6 minutes.

However, one can claim that the local deformation field is not sufficient to characterize the registration between two shapes. Often data is corrupted by noise while at the same time outliers exist in the training set. Therefore recovering measurements that do allow the characterization of the quality of the registration process is an eminent condition of accurate shape modeling.

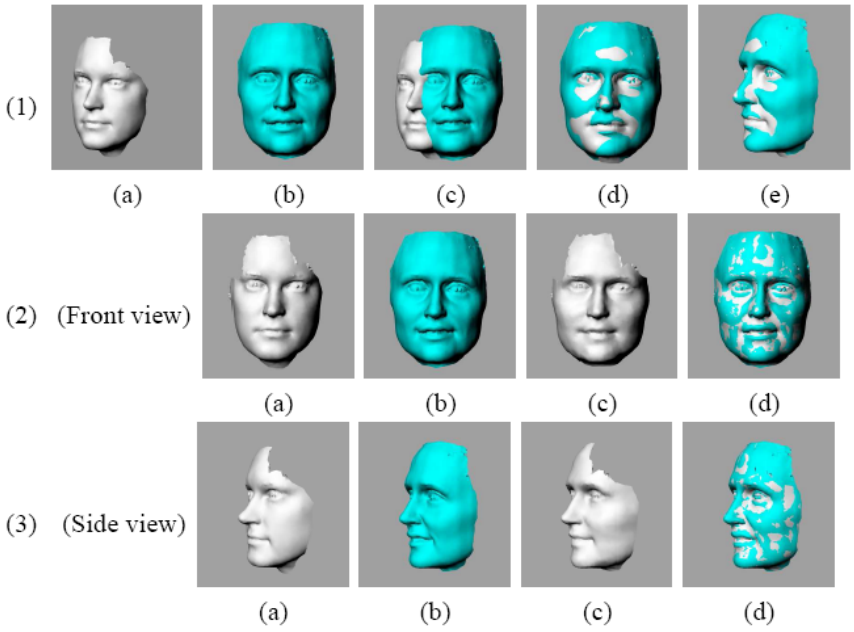

FIGURE 7. *Global-to-local Registration for Open 3D structures (both source and target shapes are from face range scan data). (1) Global Registration using the similarity transformation model: (a) Initial source shape; (b) Target shape; (c) Initial pose of the source shape relative to the target shape; (d & e) Globally transformed source shape shown overlaid on the target - front view (d) and side view (e). (2) Local registration using FFD. The results are shown from both a 3D front view in the second row and a side view in the third row. (Front view & Side view): (a) Source shape after rigid transformation; (b) Target shape; (c) Locally deformed source shape after FFD registration; (d) Locally deformed source shape shown overlaid on the target.*

4 Estimation of Registration Uncertainties

Several attempts to build statistical models on noisy set of data in order to infer the properties of a certain model were proposed in the former literature. In [25], various techniques to extract feature points in images along with uncertainties due to the inherent noise were reported while in [36] an iterative estimation method was proposed to handle uncertainties estimates of rigid motion on sets of matched points. Last, but not least in [49] an iterative technique to determine uncertainties within the Iterated Closest Point [9] registration algorithm was proposed. In a quite different context, [47] introduced uncertainties within the estimation of dense optical flow, that can be seen as a form of registration between images.

In the present case curves are considered using implicit representation, therefore uncertainty does not lie in the relative position of points but of an isosurface and therefore one can seek for equivalences with ”aperture problem” in optical

flow estimation. Inspired by the work in [4, 49] we aim to recover uncertainties on the vector Θ while being able to use only the zero iso-surface, defining the shape itself. To this end, we use a discrete formulation of:

$$E_{data}(\mathcal{L}(\Theta)) = \iint_{\Omega} \chi_{\alpha}(\phi_{\mathcal{S}}(\mathbf{x}))\big(\phi_{\mathcal{S}}(\mathbf{x}) - \phi_{\mathcal{T}}(\mathcal{L}(\boldsymbol{\Theta};\mathbf{x}))\big)^2 d\mathbf{x}$$

that can be re-written in the following fashion when $\alpha \to 0$:

$$E_{data}(\mathcal{L}(\Theta)) = \iint_{\Omega} \phi^2_{\mathcal{T}}(\mathcal{L}(\boldsymbol{\Theta};\mathbf{x}))d\mathbf{x} = \sum_{i=1}^{K} \rho(\phi_{\mathcal{T}}(\mathcal{L}(\Theta,\mathbf{x}_i)) = \sum_{i=1}^{K} \rho(\phi_{\mathcal{T}}(\mathbf{x}'_i))$$

with $\mathbf{x}' = \mathcal{L}(\Theta;\mathbf{x})$. Let us consider $\mathbf{q}_i$ to be the closest point on the target contour from $\mathbf{x}'_i$. Since $\phi_{\mathcal{T}}$ is assumed to be an Euclidean distance transform, it satisfies the condition $\|\nabla\phi_{\mathcal{T}}(\mathbf{x}'_i)\| = 1$. Therefore one can express the values of $\phi_{\mathcal{T}}(\mathbf{x}'_i)$:

$$\phi_{\mathcal{T}}(\mathbf{x}'_i) = (\mathbf{x}'_i - \mathbf{q}_i) \cdot \nabla\phi_{\mathcal{T}}(\mathbf{x}'_i)$$

Then, one has a first order approximation of $\phi_{\mathcal{T}}(\mathbf{x})$ in the neighborhood of $\mathbf{x}'_i$, in the form :

$$\begin{aligned}\phi_{\mathcal{T}}(\mathbf{x}'_i + \delta\mathbf{x}'_i) &= \phi_{\mathcal{T}}(\mathbf{x}'_i) + \delta\mathbf{x}'_i \cdot \nabla\phi_{\mathcal{T}}(\mathbf{x}'_i) \\ &= (\mathbf{x}'_i + \delta\mathbf{x}'_i - \mathbf{q}_i) \cdot \nabla\phi_{\mathcal{T}}(\mathbf{x}'_i)\end{aligned}$$

that reflects the condition that a point to curve distance is adopted rather than a point to point. Under the assumption that $E(\mathcal{L}(\Theta)) = \circ(1)$ we can neglect the second order term in the development of $\phi_{\mathcal{T}}$ and therefore write the following second order approximation of E_0 in quadratic form :

$$E(\mathcal{L}(\Theta)) = \sum \left[(\mathcal{L}(\Theta,\mathbf{x}_i) - \mathbf{q}_i) \cdot \nabla\phi_{\mathcal{T}}(\mathbf{x}'_i)\right]^2$$

Free form deformations is a linear transformation with respect to the parameters $\Theta = \delta\mathbf{P}_{i,j}$. Therefore one can rewrite this transformation over the image domain in a rather compact form:

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

where $\mathcal{X}(\mathbf{x})$ is a matrix of dimensionality $2 \times N$ with N being the size of Θ. One now can substitute this term in the objective function toward :

$$E(\mathcal{L}(\Theta)) = (\mathcal{X} \cdot \Theta - \mathbf{y})^T(\mathcal{X} \cdot \Theta - \mathbf{y})$$

with

$$\mathcal{X} = \begin{pmatrix} \eta_1^T\mathcal{X}(\mathbf{x}_1) \\ \vdots \\ \eta_K^T\mathcal{X}(\mathbf{x}_K) \end{pmatrix} \text{ and } \mathbf{y} = \begin{pmatrix} \eta_1^T(\mathbf{q}_1 - \mathbf{x}_1) \\ \vdots \\ \eta_K^T(\mathbf{q}_K - \mathbf{x}_K) \end{pmatrix}$$

and $[\eta_i = \nabla\phi_{\mathcal{T}}(\mathbf{x}'_i)]$ due to the distance transform nature of the implicit function. We assume that $\mathbf{y}$ is the only random variable. Such assumption is equivalent with

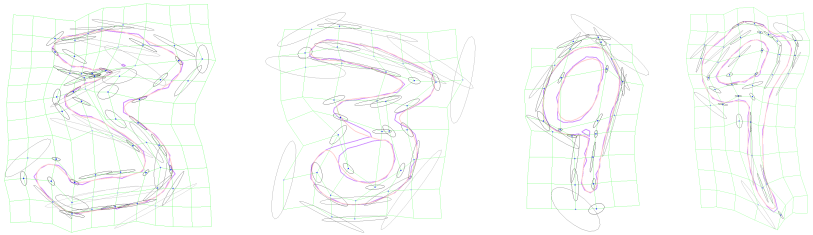

FIGURE 8. *Examples of registration with uncertainties estimation at the free form deformation Grid.*

saying that errors in the point positions are only quantified along the normal direction. This accounts for the fact that the point set is treated as samples extracted from a continuous manifold. One can take the derivative of the objective function in order to recover a linear relation between Θ and $\mathbf{y}$:

$$\mathcal{X}^T\mathcal{X}\Theta = \mathcal{X}^T\mathbf{y}$$

Last, assume that the components of $\mathbf{y}$ are independent and identically distributed. In that case, the covariance matrix of $\mathbf{y}$ has the form $\sigma^2\mathbf{I}$ of magnitude σ^2 with $\mathbf{I}$ being the identity. In the most general case one can claim that the matrix $\mathcal{X}^T\mathcal{X}$ is not invertible because due to the fact that the registration problem is underconstrained. Additional constraints are to be introduced toward the estimation of the covariance matrix of Θ through the use of an arbitrarily small positive parameter γ :

$$E(\mathcal{L}(\Theta)) = (\mathcal{X}\Theta - \mathbf{y})^T(\mathcal{X}\Theta - \mathbf{y}) + \gamma\,\Theta^T\Theta$$

Then the covariance matrix of the parameter estimate is :

$$\Sigma_\Theta = \sigma^2(\mathcal{X}^T\mathcal{X} + \alpha\mathbf{I})^{-1}$$

Some example of such estimates are shown in [FIGURE (8)]. Once registration between shapes has been addressed numerous computational vision tasks can be considered. Shape modeling with applications to and object extraction and recognition are the more frequent ones.

5 Applications

In this section, we present three applications of the proposed global-to-local registration framework to demonstrate its potential in the domains of grouping and recognition.

5.1 Statistical (Gaussian) modeling of anatomical structures & Segmentation of the Left Ventricle in Ultrasound Images

Organ modeling is a critical component of medical image analysis. To this end, one would like to learn a compact representation that can capture the variation

in an anatomical structure of interest across individuals. Building such representation requires establishing dense local correspondences across a set of training examples. The registration framework proposed in this book chapter can be use to solve the dense correspondence problem.

As an example, we show the statistical modeling of systolic left ventricle (LV) shapes from ultrasonic images, using 40 pairs of hand-drawn LV contours. We first apply global rigid registration to align all contours to a same target. Local registration based on free form deformations is then used to non-rigidly register all these contours to a common target. In order to establish dense one-to-one correspondences between all the aligned contours, we pick a set of sample points on the common target and compute their correspondences on each training contour based on the local registration result [FIGURE (6)] for established local correspondences). Similar procedure is applied for the endiastolic shape of left ventricle.

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 o_1, o_2 of operators - the largest amount of variation within the training data, according to:

$$\mathbf{d} = \mathbf{d}' + \sum_{k=1}^{o_1} \lambda_k^d \, (\mathbf{u}_k^d, \mathbf{v}_k^d), \quad \mathbf{s} = \mathbf{s}' + \sum_{k=1}^{o_2} \lambda_k^s \, (\mathbf{u}_k^s, \mathbf{v}_k^s)$$

where $\mathbf{d}'$ (resp. $\mathbf{s}'$) is the mean diastolic shape, o_1 (resp. o_2) is the number of retained modes of variation, $(\mathbf{u}_k^d, \mathbf{v}_k^d)$ (resp. $(\mathbf{u}_k^s, \mathbf{v}_k^s)$) are these modes (eigenvectors), and λ_j^d (resp. λ_j^s) are linear factors within the allowable range defined by the eigenvalues.

Once average models for the systolic and diastolic cases are considered, one can further assume that these models are registered, therefore there is an one-to-one correspondence between the points that define these shapes. To this end, their implicit representations $\phi_d,\ \phi_s$ are aligned using first a similarity transformation and then a free form deformation.

Let $\big(\overline{\mathbf{d}} = (\mathbf{x}_1^d, \mathbf{x}_2^d, ..., \mathbf{x}_m^d)\big)$ be the diastolic and $(\overline{\mathbf{s}} = (\mathbf{x}_1^s, \mathbf{x}_2^s, ..., \mathbf{x}_m^s))$ the systolic average model once global and local registration between them has been recovered. Then one can define a linear space of shapes as follows:

$$\begin{aligned} \overline{\mathbf{c}}(\mathbf{a}) &= \mathbf{a}\,\overline{\mathbf{s}} + (1-\mathbf{a})\,\overline{\mathbf{d}}, \quad 0 \le \mathbf{a} \le 1 \\ &= \big(\mathbf{a}\mathbf{x}_1^d + (1-\mathbf{a})\mathbf{x}_1^s, ..., \alpha\mathbf{x}_1^d + (1-\mathbf{a})\mathbf{x}_m^s)\big) \end{aligned}$$

and a linear space of deformations that can account for the systolic, the diastolic frame as well as the frames in between:

$$\mathbf{c}(\mathbf{a}, \lambda_k^d, \lambda_s^d) = \overline{\mathbf{c}}(\mathbf{a}) + \sum_{k=1}^{o_1} \lambda_k^d \, (\mathbf{u}_k^d, \mathbf{v}_k^d) + \sum_{k=1}^{o_2} \lambda_k^s \, (\mathbf{u}_k^s, \mathbf{v}_k^s)$$

The most critical issue to be addressed within this process is the global and local registration between the systolic and diastolic average shapes. The approach proposed in [23] that performs registration in the implicit space of distance functions

using a combination between mutual information criterion and a free-form deformation principle is used. The resulting composite model has limited complexity, can account for the systolic and the diastolic form of the endocardium as well as for the frames between the two extrema.

Rough Segmentation through Registration

The central idea behind active shape models is to recover (i) an approximate solution through a global registration $(\mathbf{a}, \mathcal{A})$ between the time varying average model and the image and (ii) the exact solution through a linear combination of the principal modes of variation $\mathbf{c}(\lambda_k^d, \lambda_s^d)$. To this end, given an initial position of the average model and a number of control points $\mathbf{c}_i$, the method seeks in a repetitive manner the most prominent correspondence of each control point in the image plane $\mathbf{p}_i$. Once such correspondences have been recovered, the registration parameters between the image and the model are updated so the recovered projection to the image approximates the desired image features;

$$E_{data}(\mathbf{a}, \mathcal{A}) = \sum_{i=0}^{m} \rho \left(|\mathcal{A}(\overline{\mathbf{c}}_i(\mathbf{a})) - \mathbf{p}_i| \right) \tag{1.1}$$

where $\mathcal{A}$ refers to translation, rotation and scale, $\mathbf{a}$ dictates the model space and ρ is a robust error metric.

However, recovering the correspondences $\mathbf{p}_i$ is a tedious task. Active shape models are based on generating an intensity profile for each control point along the normal to the model and seeking for a transition from tissue to the ventricle. We consider a probabilistic formulation of the problem. One would like to recover a density $p_{wall}(;)$ that can provide the probability of a given pixel ω being at the boundaries of the endocardium.

Let $p_{tissue}(;)$ being the probability of a given intensity being part of the endocardium walls and $p_{blood}(;)$ the density that describes the visual properties of the blood pool. Then correspondences between the model and the image are meaningful in places where there is a transition (tissue to blood pool) between the two classes. Given a local partition one can define a transition probability between these two classes.

Such a partition consists of two line segments $[\mathcal{L}(\mathcal{A}(\mathbf{x}_i)), \mathcal{R}(\mathcal{A}(\mathbf{x}_i))]$ along the normal $[\mathcal{A}(\mathcal{N}_i)]$. Their origin is the point of interest $\mathcal{A}(\mathbf{x}_i)$ and they have opposite directions. Therefore in statistical means one can write:

$$p_{wall}(\omega) = p\left([tissue|\phi \in \mathcal{L}(\omega)] \cap [blood|\phi \in \mathcal{R}(\omega)]\right)$$

Furthermore, independence between the two classes can be considered:

$$\begin{aligned} p_{wall}(\omega) &= p\left(tissue|\phi \in \mathcal{L}(\omega)\right)\ p\left(\text{blood}|\phi \in \mathcal{R}(\omega)\right) \\ &= \prod_{\phi \in \mathcal{L}(\omega)} p_{tissue}(I(\phi)) \prod_{\phi \in \mathcal{R}(\omega)} p_{blood}(I(\phi)) \end{aligned}$$

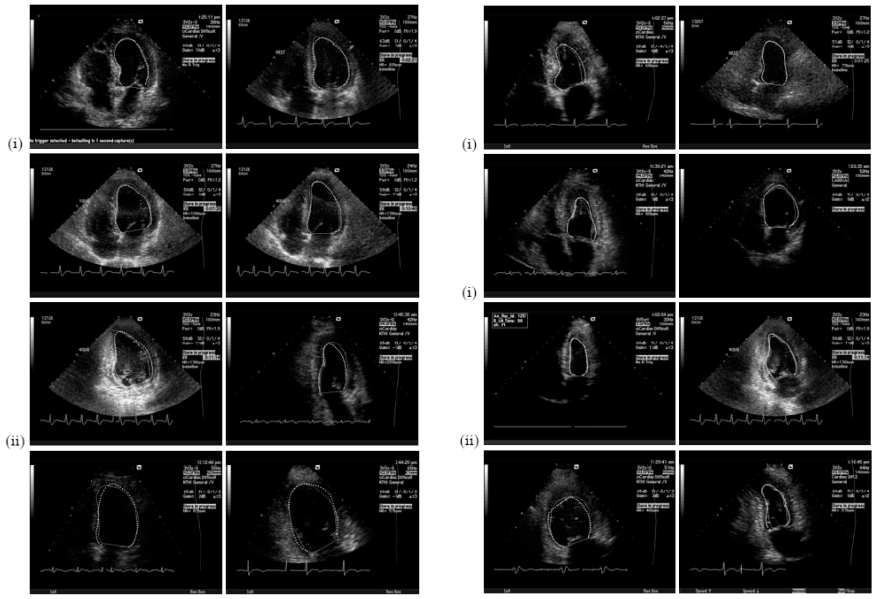

(a) (b)

FIGURE 9. *(a) Diastole Endocardium Segmentation, (b) Systole Endocardium Segmentation: (i) Apical 4 View , (ii) Apical 2 View.*

One can evaluate this probability under the condition that the blood pool and tissue densities are known. The use of -log function can be considered to overcome numerical constraints, leading to:

$$E(\omega) = \sum_{\phi \in \mathcal{L}(\omega)} \lambda_{blood}\, I(\phi) + \sum_{\phi \in \mathcal{R}(\omega)} \frac{(I(\phi) - \mu_{tissue})^2}{2\sigma^2_{tissue}}$$

after dropping out the constant terms. Thus, the best correspondence $(\mathbf{p}_i)$ is recovered by evaluating $E(\omega)$ for all ω (within a search window) that belong to the line segment that is normal to the latest solution $\mathcal{A}(\mathcal{N}_i)$ at a given control point $\mathbf{c}_i$.

Recovering the optimal transformation $\mathcal{A}$ can now be done in an incremental manner by solving a linear system

$$(m,n), \quad \frac{\partial}{\partial \alpha_{m,n}} E_{data}(\mathbf{a}, \mathcal{A}) = 0 \rightarrow$$
$$\sum_{i=0}^{m} \rho'(|\mathcal{A}(\overline{\mathbf{c}}_i(\mathbf{a})) - \mathbf{p}_i|) \frac{\partial}{\partial \alpha_{m,n}} (|\mathcal{A}(\overline{\mathbf{c}}_i(\mathbf{a})) - \mathbf{p}_i|) = 0$$

where $(\alpha_{m,n})$ are the parameters of the global transformation $\mathcal{A}$, four in the case of similarity that was considered. On the other hand, the estimation of the $\mathbf{a}$ is done through an exhaustive search. The search space $[0, 1]$ is quantized using a

uniform sampling rule and the cost function [eqn. (1)] is evaluated to determine the optimal measure of $(\mathbf{a})$. The process alternates between the estimation of the transformation ($\mathcal{A}$) and the blending parameter between the systolic and the endiastolic model until convergence to the endocardium boundaries.

Last, but not least the distribution parameters for the tissue and blood case are adaptively recovered using the latest position of the mean model. The inner region is used to determine the Laplacian parameter (λ_{blood}) while a narrow band in the outwards direction dictates the estimates of $(\mu_{tissue}, \sigma_{tissue})$.

Once, appropriate models and similarity transformations were recovered, the next step is precise extraction of the endocardium walls. Such a task is equivalent with finding a linear combination of the modes of variation that deforms globally the model toward the desired image features. The space of variations consists of the diastolic and the systolic models. We claim that the need of a blending parameter between systolic and diastolic modes of variation is not present. It can be easily shown that adding such a factor leads to a multiplication of the $(\lambda_0^d, ..., \lambda_0^s, ...)$ coefficients that are to be recovered and therefore such a multiplication factor can be omitted.

Under the assumption of existing correspondences $\mathbf{p}_i$ and the global transformation $(\mathbf{a}, \mathcal{A})$ these linear coefficients are recovered through:

$$E_{refine-data}(\lambda_0^d, ..., \lambda_0^s, ...) = \sum_{i=0}^{m} \rho \left(|\mathcal{A}\left(\overline{\mathbf{c}}_i(\mathbf{a})\right) + \sum_{k=1}^{o_1} \lambda_k^d \left(\mathbf{u}_k^d, \mathbf{v}_k^d\right) + \sum_{k=1}^{o_2} \lambda_k^s \left(\mathbf{u}_k^s, \mathbf{v}_k^s\right) - \mathbf{p}_i| \right)$$

The objective function is minimised using a robust incremental estimation technique. The calculus of Euler-Lagrange equations with respect to the unknown variables $(\lambda_0^d, ..., \lambda_0^s, ...)$ leads to a $[o_1 + o_2] \times [o_1 + o_2]$ linear system that has a closed form solution. Such step is repeated until convergence. Examples of such a segmentation process are shown in [FIGURE (9)]. While the form of the left ventricle can be well described using a Gaussian density, in the most general case shapes that refer to objects of particular interest are non-linear structures and therefore the assumption of simple parametric models likes Gaussians is rather unrealistic. Therefore within our approach we propose a non-parametric form of the *pdf*.

On top of that, our registration paradigm along with the deformation field it estimates confidence measures (uncertainties) that upon proper use can provide a better use of the information space. It is natural to assign less importance on the variations that appear in areas with low registration confidence. Such areas in principle are not related with the distribution of the training samples after registration. Kernels of variable bandwidth can be used to encode such a condition and provide a structured way for utilizing the variable uncertainties associated to the sample points

5.2 Variable Bandwidth Density Estimation & Shape Recognition

Let $\{\mathbf{x}_i\}_{i=1}^{M}$ denote a random sample with common density function f. The fixed bandwidth kernel density estimator consists of:

$$\begin{aligned}\hat{f}(\mathbf{x}) &= \frac{1}{M}\sum_{i=1}^{M} \mathrm{K}_{\mathbf{H}}\left(\mathbf{x}-\mathbf{x}_i\right) \\ &= \frac{1}{M}\sum_{i=1}^{M} \frac{1}{\|\mathbf{H}\|^{1/2}} \mathrm{K}\left(\mathbf{H}^{-1/2}(\mathbf{x}-\mathbf{x}_i)\right)\end{aligned}$$

where $\mathbf{H}$ is a symmetric definite positive - often called a bandwidth matrix - that controls the width of the kernel around each sample point $\mathbf{x}_i$. The fixed bandwidth approach often produces an under-smoothing in areas with sparse observations and over-smoothing in the opposite case. Usefulness of varying bandwidths is widely acknowledged to estimate long-tailed or multi-modal density functions with kernel methods.

In the literature, Kernel density estimation methods that do rely on such varying bandwidths are generally referred to as adaptive kernel density estimation methods [56]. An adaptive kernel approach adapts to the sparseness of the data by using a broader kernel over observations located in regions of low density. Two useful state of-the-art variable bandwidth kernels consists of the *sample point estimator* and the *balloon estimator*.

The first one refers to a covariance matrix depending on the repartition of the points constituting the sample :

$$\hat{f}_S(\mathbf{x}) = \frac{1}{M}\sum_{i=1}^{M} \frac{1}{\|\mathbf{H}(\mathbf{x}_i)\|^{1/2}} \mathrm{K}\left(\mathbf{H}(\mathbf{x}_i)^{-1/2}(\mathbf{x}-\mathbf{x}_i)\right)$$

where a common selection of $\mathbf{H}$ refers to

$$\mathbf{H}(\mathbf{x}_i) = h(\mathbf{x}_i)\cdot\mathbf{I}$$

with $h(\mathbf{x}_i)$ being the distance of point $\mathbf{x}_i$ from the k^{th} nearest point. One can consider various alternatives to determine the bandwidth function. Such estimator may be directly used with the uncertainties estimates $\mathbf{H}(\mathbf{x}_i) = \mu\Sigma_{\Theta_i}$ as proposed in [8].

Our registration method assumes an estimation of the uncertainty on the point to be evaluated. In order to make use of this information we first introduce another standard variable bandwidth kernel method knows as *balloon estimator*. It adapts the measures to the point of estimation depending on the shape of the sampled data according to:

$$\hat{f}_B(\mathbf{x}) = \frac{1}{M}\sum_{i=1}^{M} \frac{1}{\|\mathbf{H}(\mathbf{x})\|^{1/2}} \mathrm{K}\left(\mathbf{H}(\mathbf{x})^{-1/2}(\mathbf{x}-\mathbf{x}_i)\right)$$

with $\mathbf{H}(\mathbf{x})$ may be chosen with the same model as *sample point estimator*. Such function may be seen as the average of a density associated to the estimation point

$\mathbf{x}$ on all the sample points $\mathbf{x}_i$. One should point out that such a process could lead to estimates on $\hat{f}(\mathbf{x})$ that do not refer to density function in terms of discontinuity, integration to infinity, etc.

Let us consider $\{\mathbf{x}_i\}_{i=1}^{M}$ a multi-variate set of measurements where each sample $\mathbf{x}_i$ exhibits uncertainties in the form of a covariance matrix Σ_i. Our objective can be stated as follows: estimate the probability of a new measurement $\mathbf{x}$ that is associated with covariance matrix Σ.

Let $\mathbf{X}$ be the random variable associated to the training set and assume a density function f. f may be estimated with $\hat{f}$ in a similar fashion to *sample point* estimator. Therefore f may be expressed in the form $f = \sum f_i$ where f_i are densities associated to a single kernel $\mathbf{x}_i$. Let $\mathbf{Y}$ be the be a random variable for the new sample with density g.

Then one can claim that in order to estimate the probability of the new sample, one should first determine for all possible $\mathbf{u} \in \mathbb{R}^N$ their *distance* from the existing kernels of the training set $\mathbf{X}$, $f(\mathbf{u})$ in a similar fashion as *sample point estimator* and weight them according to their fit with the density function of $\mathbf{Y}$:

$$\begin{aligned} f(\mathbf{x}) &= \int f(\mathbf{u}) g(\mathbf{u}) d\mathbf{u} \\ &= \int \left[\sum_{i=1}^{M} f_i(t) \right] g(t) dt = \sum_{i=1}^{M} \left[\int f_i(t) g(t) dt \right] \end{aligned}$$

The present concept could be relaxed to address the case of non-Gaussian kernels according to a *hybrid* estimator that is considered:

$$\hat{f}_H(\mathbf{x}) = \frac{1}{M} \sum_{i=1}^{M} \frac{1}{\|\mathbf{H}(\Sigma, \Sigma_i)\|^{1/2}} \mathrm{K}(\mathbf{H}(\Sigma, \Sigma_i)^{1/2}(\mathbf{x} - \mathbf{x}_i))$$

Such a density estimator takes into account the uncertainty estimates both on the sample points themselves as well as on the estimation of point $\mathbf{x}$ as introduced in [29]. The outcome of this estimator may be seen as the average of the probabilities that estimation measurement is equal to the sample measurement, calculated over all sample measurements. Consequently, in directions of important uncertainties the density estimation decreases more slowly when compared to the other directions.

This metric can now be used to assess for a new sample the probability of being part of the training set through a process that evaluates the probability for each of the examples in the training set. The resulting approach can account for the non-parametric form of the observed density while the limitation of being time consuming since the cost is linear to the number of samples in the training set. Therefore, there is an eminent need on decreasing the cardinality of the set of retained kernels.

The maximum likelihood criterion expresses the quality of approximation from the model to the data. Consider a set $\mathcal{Z}_\mathrm{K} = \{X_1, X_2, \ldots, X_\mathrm{K}\}$ of kernels extracted from the training set with mean and uncertainties estimates $\{\mathbf{x}_i \Sigma_i\}_{i=1}^{|\mathcal{Z}_\mathrm{K}|}$. Then the probability of any registered shape with associated kernel Y has the

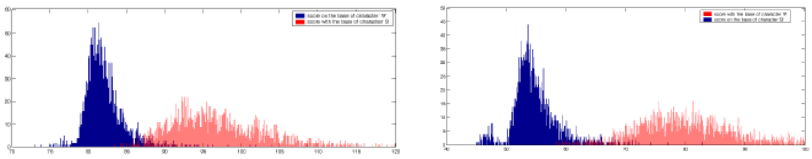

(a) (b)

FIGURE 10. *(a) Distribution of the distance of the training set from the kernel based model build for 3 in logarithmic scale. (b) Distribution of the distance of the training set from the kernel based model build for 9 in logarithmic scale.*

form :

$$P_{\mathcal{Z}_{\mathrm{K}}}(Y) = \frac{1}{|\mathcal{Z}_{\mathrm{K}}|} \sum_{X \in \mathcal{Z}_{\mathrm{K}}} \mathrm{K}(X, Y)$$

and K(X,Y) correspond to the calculation of the hybrid kernel estimator. For such a selection of kernels, one can evaluate the log-likelihood for the entire training set with the associated kernels $\{Y_i\}_{i=1}^{N}$:

$$C_{\mathrm{K}} = \sum_{i=1}^{N} \log(P_{\mathcal{Z}_{\mathrm{K}}}(Y_i))$$

We use an efficient sub-optimal iterative algorithm to update the set $\mathcal{Z}_{\mathrm{K}}$. A new kernel Y is extracted from the training set as the one maximizing the quantity $C_{\mathrm{K}+1}$ associated to $\mathcal{Z}_{\mathrm{K}+1}$ with : $\mathcal{Z}_{\mathrm{K}+1} = \mathcal{Z}_{\mathrm{K}} \bigcup Y$. One kernel may be chosen several times in order to preserve a decreasing order of C_K when adding new kernels. Consequently the selected kernels Y_i used to evaluate the global density probability have prior weight.

The proposed method is indented to provide efficient models for family of shapes with important variation. Digits, is an example where the shape of the characters varies along individuals and therefore one can claim important variability on the training set. Based on this observation and using an important training set from the database, we have considered two digits (random variables of 2000 samples each) that have a quite similar structure, the 3 and 9. Upon intra-class registration two models have been built of 100 kernels each according to the maximum likelihood principle. Then, a cross validation task was performed where for all samples of the database (3 & 9) the probability of being part of the classes 3 & 9 was estimated according to the presented variable bandwidth density function. In [FIGURE (10.a)] one can see in a logarithmic scale the performance of the method using the model built for 3 and applied also to the samples of 9 while the opposite case is presented in [FIGURE (10.b)]. In both cases one can see a clear separation of the two classes and a substantial difference in terms of probabilities between the true and the non-true case. It is important to note that the presented method is not indented for such an application. However, given this validation we can claim that such a model can capture samples of increasing complexity and the use of deformations along with uncertainties provide efficient density estimators.

5.3 Knowledge-based Object Extraction Using Distance Transforms and Level Set Methods

The idea of global as well as local registration can be explored to impose prior knowledge within a level set [30, 31] segmentation process [10, 14, 26, 39, 53]. Such a method is based on the propagation of an initial contour (in practice its implicit function) toward the desired image characteristics [5, 6, 27, 32, 38]. Often, the signed distance function is considered to represent the evolving contour. In the most general case, one can assume a simplistic average shape model [39] $\phi_{\mathcal{M}}$. Then, in order to constrain the segmentation process, one can force the evolving curve ϕ to look like the prior, or:

$$E_{shape}(\phi,(\mathbf{s},\mathbf{R},\mathbf{T})) = \iint\limits_{\Omega} \left(\mathbf{s}\,\phi_{\mathcal{M}}(\mathbf{x}) - \phi(\mathcal{A}(\mathbf{x}))\right)^2 d\mathbf{x},$$

One can further assume that the image refers to a bi-modal partition $(\mathcal{R}, \Omega - \mathcal{R})$ with where the distributions of the visual properties of the object $p_{obj} : \mathcal{N}(\mu_{obj}, \sigma_{obj})$ and the background $p_{bcg} : \mathcal{N}(\mu_{bcg}, \sigma_{bcg})$ are assumed to be Gaussians. Under the assumption of independence between hypotheses and across pixels, the maximum posterior can be used to determine the optimal segmentation results [33], or

$$E_{data}(\phi,\mathcal{N}(\mu_{obj},\sigma_{obj}),\mathcal{N}(\mu_{bcg},\sigma_{bcg})) = \\ -\iint\limits_{\mathcal{R}} \log\left(p_{obj}(\mathcal{I}(\mathbf{x}))\,d\mathbf{x} - \iint\limits_{\Omega-\mathcal{R}} \log\left(p_{bcg}(\mathcal{I}(\mathbf{x}))\,d\mathbf{x}\right.\right.$$

Within the level set framework, one can use the ϕ function to describe this partition [60] according to:

$$E_{data}(\phi,\mathcal{N}(\mu_{obj},\sigma_{obj}),\mathcal{N}(\mu_{bcg},\sigma_{bcg})) = \\ -\iint\limits_{\Omega} \mathcal{H}_\alpha(\phi)\log\left(p_{obj}(\mathcal{I}(\mathbf{x}))\,d\mathbf{x} - \iint\limits_{\Omega} (1-\mathcal{H}_\alpha(\phi))\log\left(p_{bcg}(\mathcal{I}(\mathbf{x}))\,d\mathbf{x}\right.\right.$$

where $\mathcal{H}$ is the Heaviside function

$$\mathcal{H}_\alpha(\phi) = \begin{cases} 1 & , \quad \phi > \alpha \\ 0 & , \quad \phi < -\alpha \\ \frac{1}{2}\left(1 + \frac{\phi}{\alpha} + \frac{1}{\pi}\sin\left(\frac{\pi\phi}{\alpha}\right)\right) & , \quad |\phi| < \alpha \end{cases}$$

Last, but not least toward smooth segmentation results one can impose a minimal length curve constraint according to:

$$E_{smooth}(\phi) = \iint\limits_{\Omega} \delta_\alpha(\phi(\mathbf{x}))\,|\nabla\phi(\mathbf{x})|\,d\mathbf{x}$$

where δ_α is the Dirac function

$$\delta_\alpha(\phi) = \begin{cases} 0 & , \quad |\phi| > \alpha \\ \frac{1}{2\alpha}\left(1 + \cos\left(\frac{\pi\phi}{\alpha}\right)\right) & , \quad |\phi| < \alpha \end{cases}$$

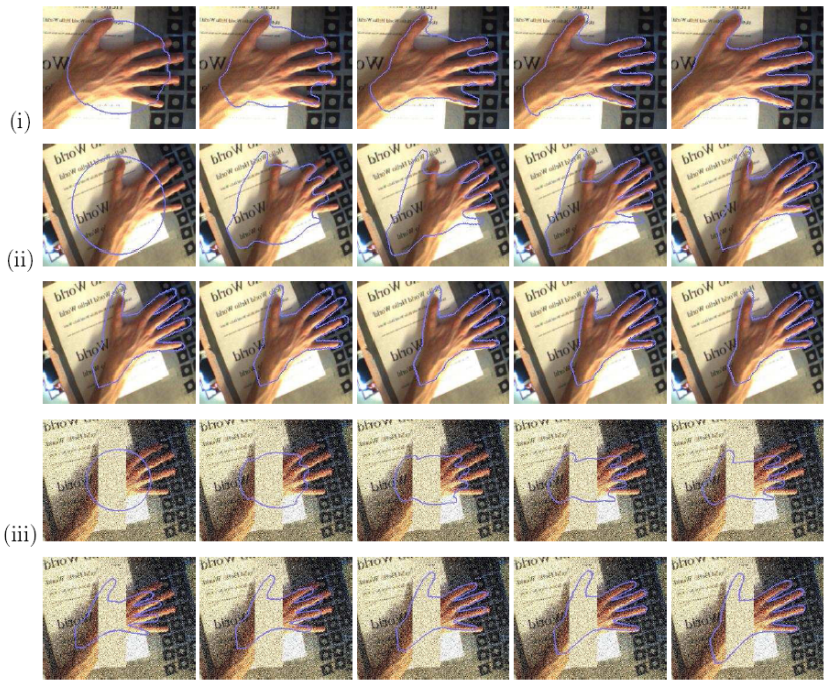

FIGURE 11. *Implicit Representations, Prior Knowledge and Object Extraction under occlusions. The evolution of the contour is presented in a raster scan format. (i) original image from where the prior was extracted, (ii) changes of pose for the object to be recovered, (iii) image with changes of scale, pose, illumination, noise and missing parts.*

with $\left[\frac{\partial}{\partial\phi}\mathcal{H}_\alpha(\phi) = \delta_\alpha(\phi)\right]$. One can now integrate smoothness, data-driven and shape-driven terms toward object extraction according to:

$$E(\phi, (\mathcal{N}(\mu_{obj}, \sigma_{obj}), \mathcal{N}(\mu_{bcg}, \sigma_{bcg}), (\mathbf{s}, \mathbf{R}, \mathbf{T})) =$$
$$w_1 E_{shape}(\phi, (\mathbf{s}, \mathbf{R}, \mathbf{T})) + w_2 E_{data}(\phi, \mathcal{N}(\mu_{obj}, \sigma_{obj}), \mathcal{N}(\mu_{bcg}, \sigma_{bcg})) + w_3 E_{smooth}(\phi)$$

where the object position, its visual properties as well as the transformation between the object and the average model are to be recovered. The calculus of variations with respect to the evolving interface (level set), its projection to mean model as well as the statistical properties of appearance of the object and the background can be considered to recover the lowest potential of the designed cost function. One can refer to the section 2 regarding the derivation with respect to the pose parameters while the level set function evolves according to:

$$\frac{\partial\phi}{\partial t}(\mathbf{x}) = a\,\delta(\phi(\mathbf{x}))\,\left(w_2 \mathrm{div}\frac{\nabla\phi(\mathbf{x})}{|\nabla\phi(\mathbf{x})|} + w_3\,\log\frac{p_{obj}\left(\mathcal{I}(\mathbf{x})\right)}{p_{bcg}\left(\mathcal{I}(\mathbf{x})\right)}\right)$$
$$-\,2w_1\,\mathbf{s}\left(\mathbf{s}\,\phi_{\mathcal{M}}(\mathbf{x}) - \phi(\mathcal{A}(\mathbf{x}))\right)$$

One can also take the partial derivatives of the cost function with respect to the mean and the standard deviation of the normal distributions describing the object and the background intensity properties [38]. Such a method is demonstrated in [FIGURE (11)] to address knowledge-based object extraction and can be quite efficient when seeking objects of limited variations. On the other hand, more advanced models that are based on the same principle can be considered to address segmentation for cases of important deformations [40].

6 Discussion

In this chapter we have studied shape representations of implicit forms, in particular distance transforms. We have demonstrated that such representations can be quite efficient to global and local one-to-one registration. To this end, simple similarity invariant (SSD) and more advanced registration metrics able to account with various global transformations (mutual information) were considered in the space of implicit functions. Local registration was addressed through free form deformations and cubic splines in the space of distance transforms. Furthermore, we have introduced the notion of uncertainties on the registration process to the selected representation space.

Validation of the representation itself as well as the registration methods was done through parametric modeling of shapes and segmentation, non-parametric modeling and recognition and shape-driven level set based object extraction with promising results.

Acknowledgments

The authors would like to thank M.-P. Jolly and R. Ramaraj from Siemens Corporate Research for their input in the ultrasound segmentation part, J.-Y. Audibert from Ecole Nationale des Ponts et Chaussees, France for fruitful discussion related with the selection of most appropriate kernels toward reduction of the dimensionality of the non-parametric pdf used for the recognition of shapes and R. Deriche from I.N.R.I.A. for his contribution on the level-set knowledge-based segmentation part.

7 REFERENCES

[1] S. Belongie, J. Malik, and J. Puzicha. Matching Shapes. In *IEEE International Conference in Computer Vision*, pages 456–461, 2001.

[2] P. J. Besl and N. D. McKay. A Method for Registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[3] G. Borgefors. Hierarchical chamfer matching: A parametric edge matching algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10:849–865, 1988.

[4] A. Can, C. Stewart, B. Roysam, and H. Tanenbaum. A Feature-Based, Robust, Hierarchical Algorithm for Registering Pairs of Images of the Curved Human Retina. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24:347–364, 2002.

[5] V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert. Minimal surfaces based object segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:394–398, 1997.

[6] T. Chan and L. Vese. Active Contours without Edges. *IEEE Transactions on Image Processing*, 10:266–277, 2001.

[7] C. Chefd'Hotel, G. Hermosillo, and O. Faugeras. A Variational Approach to Multi-Modal Image Matching. In *IEEE Workshop in Variational and Level Set Methods*, pages 21–28, 2001.

[8] H. Chen and P. Meer. Robust Computer Vision through Kernel Density Estimation. In *European Conference on Computer Vision*, pages 236–250, 2002.

[9] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Image and Vision Computing*, 10:145–155, 1992.

[10] Y. Chen, H. Thiruvenkadam, H. Tagare, F. Huang, and D. Wilson. On the Incorporation of Shape Priors int Geometric Active Contours. In *IEEE Workshop in Variational and Level Set Methods*, pages 145–152, 2001.

[11] H. Chui and A. Rangarajan. A New Algorithm for Non-Rigid Point Matching. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages II: 44–51, 2000.

[12] A. Collignon, F. Maes, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multimodality image registration using information theory. In *Information Processing in Medical Imaging*, pages 263–274, 1995.

[13] 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.

[14] D. Cremers and S. Soatto. A pseudo distance for shape priors in level set segmentation. In N. Paragios O. Faugeras, editor, *2nd IEEE Intl. Workshop on Variational, Geometric and Level Set Methods*, pages 169–176, 2003.

[15] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor. 3D Statistical Shape Models using Direct Optimization of Description Length. In *European Conference on Computer Vision*, pages 3–20, 2002.

[16] A. Dervieux and F. Thomasset. A finite element method for the simulation of Rayleigh-Taylor instability. *Lecture Notes in Mathematics*, 771:145–159, 1979.

[17] A. Dervieux and F. Thomasset. Multifluid incompressible flows by a finite element method. In W. Reynolds and R.W. MacCormack, editors, *Seventh International Conference on Numerical Methods in Fluid Dynamics*, volume 141 of *Lecture Notes in Physics*, pages 158–163, June 1980.

[18] 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.

[19] J. Feldmar and N. Ayache. Rigid, affine and locally affine registration of free-form surfaces. *International Journal of Computer Vision*, 18:99–119, 1996.

[20] D. Forsey and R. Bartels. Hierachical B-spline Refinement. *ACM Transactions on Computer Graphics*, 22:205–212, 1988.

[21] X. Huang, N. Paragios, and D. Metaxas. Establishing local correspondences towards compact representations of anatomical structures. In *6th Annual International Conf. on Medical Image Computing & Computer Assisted Intervention*, volume 2879 of *LNCS*, pages 926–934, November 2003.

[22] X. Huang, N. Paragios, and D. Metaxas. Registration of structures in arbitrary dimension: Implicit representations, mutual information and free form deformations. In *Technical Report DCS-TR-520, Division of Computer and Information Sciences, Rutgers University*, April 2003.

[23] 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.

[24] X. Huang, S. Zhang, Y. Wang, D. Metaxas, and D. Samaras. A hierarchical framework for high resolution facial expression tracking. In *The Third IEEE Workshop on Articulated and Nonrigid Motion, in conjunction with CVPR'04*, July 2004.

[25] K. Kanatani. Uncertainty modeling and model selection for geometric inference. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(10):1307–1319, 2004.

[26] M. Leventon, E. Grimson, and O. Faugeras. Statistical Shape Influence in Geodesic Active Contours. In *Proceedings of the IEEE CVPR*, pages 316–323, Hilton Head Island, South Carolina, June 2000. IEEE Computer Society.

[27] R. Malladi, J.A. Sethian, and B.C. Vemuri. Evolutionary fronts for topology-independent shape modeling and recovery. In J-O. Eklundh, editor, *Proceedings of the 3rd European Conference on Computer Vision*, volume 800 of *Lecture Notes in Computer Science*, pages 3–13, Stockholm, Sweden, May 1994. Springer Verlag.

[28] D. Metaxas. *Physics-Based Deformable Models*. Kluwer Academic Publishers, 1996.

[29] A. Mittal and N. Paragios. Motion-based background substraction using adaptive kernel density estimation. In *Computer Vision and Pattern Recognition*, volume 2, pages 302–309, 2004.

[30] S. Osher and N. Paragios. *Geometric Level Set Methods in Imaging, Vision and Graphics*. Springer Verlag, 2003.

[31] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed : Algorithms based on the Hamilton-Jacobi formulation. *Journal of Computational Physics*, 79:12–49, 1988.

[32] N. Paragios and R. Deriche. Geodesic Active Contours and Level Sets for the Detection and Tracking of Moving Objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22:266–280, 2000.

[33] N. Paragios and R. Deriche. Geodesic Active Regions Level Set Methods for Supervised Texture Segmentation. *International Journal of Computer Vision*, 46(3):223–247, 2002.

[34] N. Paragios, M. Rousson, and V. Ramesh. Matching Distance Functions: A Shape-to-Area Variational Approach for Global-to-Local Registration. In *European Conference on Computer Vision*, pages II:775–790, 2002.

[35] N. Paragios, M. Rousson, and V. Ramesh. Non-Rigid Registration Using Distance Functions. *Computer Vision and Image Understanding*, 2003. to appear.

[36] X. Pennec and J.-P. Thirion. A framework for uncertainty and validation of 3-d registration methods based on points and frames. *Int. J. Comput. Vision*, 25(3):203–229, 1997.

[37] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever. Mutual Information Based Registration of Medical Images: A Survey. *IEEE Transactions on Medical Imaging*, 22(8):986–1004, 2003.

[38] M. Rousson and R. Deriche. A variational framework for active and adaptative segmentation of vector valued images. In *Proc. IEEE Workshop on Motion and Video Computing*, pages 56–62, Orlando, Florida, December 2002.

[39] M. Rousson and N. Paragios. Shape Priors for Level Set Representations. In *European Conference on Computer Vision*, pages II:78–93, Copenhangen, Denmark, 2002.

[40] M. Rousson, N. Paragios, and R. Deriche. Implicit Active Shape Models for 3D Segmentation in MRI Imaging. In *Medical Imaging Copmuting and Computer-Assisted Intervention*, pages 209–216, 2004.

[41] D. Rueckert, L.I. Sonda, C. Hayes, D. Hill, M. Leach, and D. Hawkes. Non-rigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18:712–721, 1999.

[42] D. Rueckert, L. Sonoda, C. Hayes, D. Hill, M. Leach, and D. Hawkes. Non-rigid Registration Using Free-Form Deformations: Application to Breast MR Images. *IEEE Transactions on Medical Imaging*, 8:712–721, 1999.

[43] T. Sebastian, P. Klein, and B. Kimia. Alignment-based recognition of shape outlines. *Lecture Notes in Computer Science*, LNCS-2059:606–618, 2001.

[44] T. Sebastian, P. Klein, and B. Kimia. Recognition of Shapes by Editting Shock Graphs. In *IEEE International Conference in Computer Vision*, pages 755–762, 2001.

[45] T. Sederberg and S. Parry. Free-Form Deformation of Solid Geometric Models. In *ACM SIGGRAPH*, pages 151–160, 1986.

[46] J. Sethian. Fast Marching Methods. *SIAM Review*, 41:199–235, 1999.

[47] E P Simoncelli. Bayesian multi-scale differential optical flow. In B Jähne, H Haussecker, and P Geissler, editors, *Handbook of Computer Vision and Applications*, volume 2, chapter 14, pages 397–422. Academic Press, San Diego, April 1999.

[48] L. Staib and S. Duncan. Boundary finding with parametrically deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:1061–1075, 1992.

[49] C. Stewart, C.-L. Tsai, and B. Roysam. The dual bootstrap iterative closest point algorithm with application to retinal image registration. *IEEE Trans. Med. Img.*, 22:1379–1394, 2003.

[50] A. Stoddart, S. Lemke, and A. Hilton. Estimating pose uncertainty for surface registration. *Image and Vision Computing*, 16:111–120, 1998.

[51] M. Sussman, P. Smereka, and S. Osher. A Level Set Approach for Computing Solutions to Incompressible Two-Phase Flow. *J. Computational Physics*, 114:146–159, 1994.

[52] M. Taron, N. Paragios, and M.-P. Jolly. Introducing error estimation in the shape learning framework : Spline based registration and non-parametric density estimator in the space of higher order polynomials. Technical Report CERTIS-05-08, Ecole Nationale des Ponts et Chaussees, 2005. http://www.enpc.fr/certis/Papers/05certis07.pdf.

[53] A. Tsai, A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, A. Grimson, and A. Willsky. Model-based Curve Evolution Technique for Image Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume I, pages 463–468, 2001.

[54] R. Veltkamp and M. Hagedoorn. State-of-the-art in Shape Matching. Technical Report UU-CS-1999-27, Utrecht University, 1999. http://webdoc.gwdg.de/ebook/ah/2000/techrep/CS-1997/1997-27.pdf.

[55] P. Viola and W. Wells. Aligment by Maximization of Mutual Information. In *IEEE International Conference in Computer Vision*, pages 16–23, 1995.

[56] M. Wand and M. Jones. *Kernel Smoothing*, volume 60. Chapman and Hall, London, 1995.

[57] Z. Xie and G. Farin. Deformation with hierarchical b-splines. *Mathematical Methods for Curves and Surfaces: Oslo 2000*, pages 545–554, 2001.

[58] A. Yezzi, L. Zollei, and T. Kapur. A Variational Framework for Joint Segmentation and Registration. In *IEEE Mathematical Methods in Biomedical Image Analysis*, pages 44–51, 2001.

[59] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *International Journal of Computer Vision*, 13(2):119–152, 1994.

[60] H-K. Zhao, T. Chan, B. Merriman, and S. Osher. A variational Level Set Approach to Multiphase Motion. *Journal of Computational Physics*, 127:179–195, 1996.

[61] W. Zhu and T. Chan. Stability for Shape Comparison Model. Technical Report 0308, UCLA-CAM, 2003.