ELSEVIER

Medical Image Analysis 4 (2000) 201–217

MEDICAL IMAGE ANALYSIS

www.elsevier.com/locate/media

An algorithmic overview of surface registration techniques for medical imaging

Michel A. Audette[a,*], Frank P. Ferrie[b], Terry M. Peters[c]

[a]*Montreal Neurological Institute (McGill University), 3801 University, Montreal, Quebec, Canada H3A 2B4*
[b]*McGill Center for Intelligent Machines (McGill University), Montreal, Canada*
[c]*The John P. Robarts Research Institute, London, Ontario, Canada*

Received 24 September 1998; received in revised form 1 July 1999; accepted 27 August 1999

Abstract

This paper presents a literature survey of automatic 3D surface registration techniques emphasizing the mathematical and algorithmic underpinnings of the subject. The relevance of surface registration to medical imaging is that there is much useful anatomical information in the form of collected surface points which originate from complimentary modalities and which must be reconciled. Surface registration can be roughly partitioned into three issues: choice of transformation, elaboration of surface representation and similarity criterion, and matching and global optimization. The first issue concerns the assumptions made about the nature of relationships between the two modalities, e.g. whether a rigid-body assumption applies, and if not, what type and how general a relation optimally maps one modality onto the other. The second issue determines what type of information we extract from the 3D surfaces, which typically characterizes their local or global shape, and how we organize this information into a representation of the surface which will lead to improved efficiency and robustness in the last stage. The last issue pertains to how we exploit this information to estimate the transformation which best aligns local primitives in a globally consistent manner or which maximizes a measure of the similarity in global shape of two surfaces. Within this framework, this paper discusses in detail each surface registration issue and reviews the state-of-the-art among existing techniques.

Keywords: Registration, Feature, Free-form surface, Surface model, Appearance

1. Introduction

The registration of 3D surfaces is dealt with extensively in machine vision and medical imaging literature. Its applications vary from building terrain maps, in the context of providing autonomy to a planetary rover (Hebert et al., 1989), and depth maps of a sea floor for oceanographic studies (Kamgar-Parsi et al., 1991), to the recognition of objects from a CAD database (Fan et al., 1989), and of course, to reconciling various imaging modalities in biomedical imaging (Collignon et al., 1993). The goal of this paper is to provide a detailed overview of surface registration techniques which have been, or could potentially be, applied to anatomical surfaces.

This problem is a subset of the general medical image registration problem, as surveyed recently by Maintz and Viergever (1998), who also discuss landmark and volume registration, but we emphasize algorithmic details, with a view to providing some motivation for each technique. Surfaces provide more redundancy than landmarks, and this redundancy may be particularly advantageous for characterizing non-rigid motion. Moreover, we can make a distinction between landmarks automatically extracted from surfaces, which can be seen as feature point-based surface registration and are included in this survey, and manually identified landmarks, which may be tedious to determine and less repeatable than the former. Furthermore, a surface-based approach is likely to be less affected than volumes if the two modalities of interest cover parts of the anatomy which overlap only partially, for example if one modality represents a small subset of the anatomy

*Corresponding author.
E-mail address: maudette@bic.mni.mcgill.ca (M.A. Audette).

1361-8415/00/$ – see front matter
PII: S1361-8415(00)00014-1

which appears in the second modality. In other words, we can usually register a subpatch with a larger surface patch, on the basis of local and global surface shape, as will be seen later. Finally, in medical imaging literature, anatomical surfaces are usually explicitly identified within tomographic data such as MRI and CT and are often closed (Herman and Liu, 1979; Udupa, 1982). We expand on this definition by including range images of anatomical structures, such as those obtained by laser-based triangulation, which have a particular relevance to image-guided surgery (Audette and Peters, 1999; Simon et al., 1994b; Kikinis et al., 1994) and which are typically open.

Registration, between modalities *A* and *B*, is the *estimation of a mapping between coordinate systems* Ref_A and Ref_B associated with each modality:

$$\boldsymbol{x}_B = \boldsymbol{T}(\boldsymbol{x}_A), \tag{1}$$

where $\boldsymbol{x}_A = (x_A, y_A, z_A)$ and $\boldsymbol{x}_B = (x_B, y_B, z_B)$ are points in coordinate systems Ref_A and Ref_B respectively which correspond to the same anatomical point, and where the quality of this mapping can be quantified by a global measurement based on fitting residuals. In an ideal, noise- and distortion-free environment where the same anatomy is imaged by two modalities of like scale, the computation of the transformation from point pairs matched on the basis *local* information would produce a relation which is also *globally* consistent. In other words, the resulting transformation would exactly align all pairs of homologous points. In practice, the data contain noise and distortion, and the anatomy itself may distort between images. Therefore the optimal relation (especially if a rigid transformation assumption is maintained) is that which reconciles local homologous point alignment and global consistency in some optimal manner.

Surface registration can be roughly partitioned into three stages, as illustrated in Fig. 1: **choice of transformation**, elaboration of **surface representation** and **similarity criterion**, and **matching and global optimization**. The first issue concerns the assumptions made about the *nature of relationships* between the two modalities, e.g. whether a rigid-body assumption applies, and if not, what type and how general a relation optimally maps one modality onto the other. The second issue determines *what type of information* we extract from the 3D surfaces, which typically characterize their local or global shape, and *how we organize this information* into a representation of the surface which will lead to improved efficiency and robustness in the last stage. The last issue pertains to *how we exploit this information* to estimate the transformation which best aligns local primitives in a globally consistent manner or which maximizes a measure of the similarity in global shape of two surfaces. Within this framework, this paper discusses each surface registration issue in detail and reviews the state-of-the-art among existing techniques.

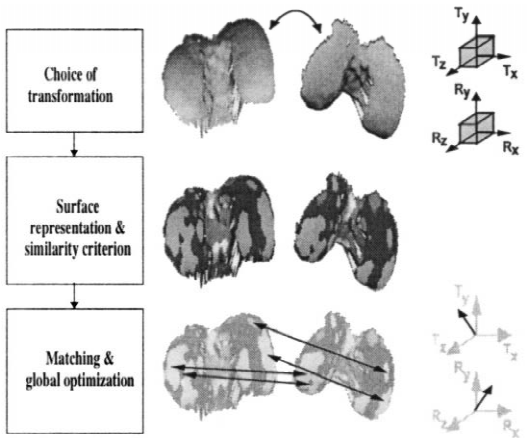

Fig. 1. 3D surface registration framework, featuring the choice of transformation between two anatomical surfaces (and of the search space for transformation parameters); the representation of the surface in terms of matching primitives (e.g. regions of consistent curvature); and finally, the matching of surface primitives and computation of the ''best'' transformation.

2. Choice of transformation

The first stage is the formalization of the assumptions about the type of relation *T* between the two 3D surfaces which is appropriate for mapping points $\boldsymbol{x}_A$ onto $\boldsymbol{x}_B$. In most registration problems, *T* is a transformation between the same anatomy imaged either by different modalities or by one modality at different times. In this context, a **rigid-body transformation** is applicable provided that the deformations sustained by the anatomy are negligible compared with the required accuracy of the transformation. If the deformations between surfaces are significant, and especially if these deformations are caused by factors other than noise and distortion within the modality, then a **nonrigid transformation** must apply. Moreover, one can further classify nonrigid transformations based on whether they are specified by a global or piecewise local fitting.

2.1. Rigid-body transformation

A general rigid-body transformation can be expressed as combination of a rotation and a translation:

$$\boldsymbol{x}_B = \boldsymbol{R}_{AB}\boldsymbol{x}_A + \boldsymbol{t}_{AB}. \tag{2}$$

Consequently, rigid-body registration typically seeks the values of $\boldsymbol{R}$ and $\boldsymbol{t}$ which minimize

$$\min_{\boldsymbol{R},\boldsymbol{t}} \sum_{i=1}^{N} \|\boldsymbol{x}_{B_i} - (\boldsymbol{R}\boldsymbol{x}_{A_i} + \boldsymbol{t})\|^2, \tag{3}$$

given 3D point correspondences $\boldsymbol{x}_{A_i}$ and $\boldsymbol{x}_{B_i}$. The problem can be reformulated in a manner which decouples the computation of $\boldsymbol{t}$ from that of $\boldsymbol{R}$ by referring the coordinates *to the respective centroids of each point set*, leading to the minimization

$$\min_{\boldsymbol{R},\boldsymbol{t}} \sum_{i=1}^{N} \|\boldsymbol{x}'_{B_i} - (\boldsymbol{R}\boldsymbol{x}'_{A_i})\|^2$$

$$\text{where } \boldsymbol{x}'_{A_i} = \boldsymbol{x}_{A_i} - \frac{1}{N}\sum_{j=1}^{N} \boldsymbol{x}_{A_j}$$

$$\text{and } \boldsymbol{x}'_{B_i} = \boldsymbol{x}_{B_i} - \frac{1}{N}\sum_{j=1}^{N} \boldsymbol{x}_{B_j} \tag{4}$$

The translation is given by the difference of centroids:

$$\boldsymbol{t} = \frac{1}{N}\sum_{j=1}^{N} \boldsymbol{x}_{B_j} - \boldsymbol{R}\frac{1}{N}\sum_{j=1}^{N} \boldsymbol{x}_{A_j}. \tag{5}$$

We review four common representations for rotation, as well as a fifth model which represents both translation and rotation. The relevance of choosing one particular representation is that it may lead to more efficient and/or numerically stable estimation of its parameters than others, or may be better suited to a particular surface representation. It is worth noting that these techniques are applicable not only to the registration of surfaces, but also to any set of explicit point pairs. Furthermore, while they may not have been published in a medical imaging context, they are still applicable to anatomical data.

The **orthonormal matrix** representation consists of a 3×3 matrix, which can be viewed as a mapping from reference frame A to frame B, once the translation between their origins is compensated, where each element R_{ij} is a *direction cosine* [i.e. the projection of one axis of reference frame A onto one axis of reference frame B (Craig, 1989)]. Arun et al. (1987) obtain rotation by first computing the *singular value decomposition* (Press et al., 1992) of the matrix $\boldsymbol{H} = \Sigma_{i=1}^{N} \boldsymbol{x}'_{B_i} \boldsymbol{x}'^{T}_{A_i}$ determined from centroid-referred coordinates:

$$\boldsymbol{H} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T, \tag{6}$$

where $\boldsymbol{D}$ is diagonal and $\boldsymbol{U}$ and $\boldsymbol{V}$ are orthonormal. The rotation is given by the expression $\boldsymbol{R} = \boldsymbol{V}\boldsymbol{U}^T$.

We briefly address two interesting representations, the **Euler angles** and **axis-angle** models, although we emphasize less the techniques that employ them, because they are iterative rather than closed-form. We can express rotation as the product of three successive rotations (γ, β, α) of a predefined *fixed* coordinate system about axes $\hat{x}$, $\hat{y}$ and $\hat{z}$, or equivalently, as a succession of rotations about $\hat{z}$, $\hat{y}$ and $\hat{x}$ *moving* axes (Craig, 1989). Huang et al. (1986) use this Euler angle representation to design a 3D iterative motion estimation scheme that is a sequence of well-behaved 2D minimizations involving the projections of (partially rotated) points on the $x-z$, $y-z$ and $x-y$ planes. Moreover, a rotation can also be completely specified by a unique vector whose *direction is the rotation axis* and whose *norm is the rotation angle about this axis* (Ayache, 1991). Lin et al. (1986) adopt the axis-angle representation for a Fourier space approach to rigid-body motion estimation which does not require explicit correspondences.

A widely used representation of rotation is based on **quaternions** (Faugeras and Hebert, 1986; Horn, 1987). A quaternion can be thought of as a *generalization of a complex number*, with *a real part* and *three imaginary parts*, or as *a composite of a 3-vector in* $\mathbb{R}^3$ *and a scalar in* $\mathbb{R}$. Moreover, the rotation quaternion can also be interpreted in terms of the axis-angle representation by the *Euler Symmetric Parameters* (Walker et al., 1991): $\dot{\boldsymbol{q}} = [\sin\theta/2\boldsymbol{n}, \cos\theta/2]^T$. In other words, the orientation of the 3-vector component specifies the axis of rotation, and the norm of the 3-vector and the scalar component are related to the rotation angle about this axis. Horn casts the search for the optimal rotation parameters as a *maximization* based on quaternion components (in contrast to the minimization of Faugeras and Hebert (1986)). His objective function is optimized with respect to rotation by finding the eigenvector corresponding to the largest positive eigenvalue of a matrix $\boldsymbol{N}$ (see Horn, 1987) determined from centroid-referred point coordinates.

The motivation for the **dual quaternion**[1] rigid transformation estimation technique of Walker et al. (1991) is that other rigid transformation estimation techniques first determine optimal orientation and then use this solution to obtain the translation (e.g., Arun et al., 1987), resulting in the accumulation of error in this computation. The dual quaternion technique solves for both relative orientation and position by minimizing a single cost function. The underlying model views the transformation between two coordinate frames as a *translation of the original coordinate frame along a direction* $\boldsymbol{n}$ *by a distance d, followed by a rotation by an angle* θ *with respect to a line having* $\boldsymbol{n}$ *as its direction and passing through a point* $\boldsymbol{p}$, as illustrated in Fig. 2. Walker reports similar accuracy to Arun's SVD technique for estimating rotation, but improved accuracy for estimating translation, across identical sets of point correspondences.

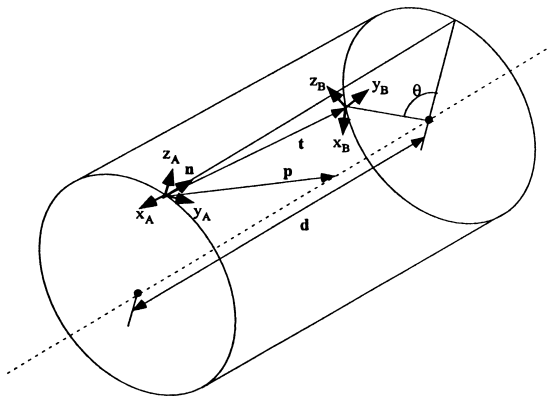

Fig. 2. Illustration of the dual quaternion model for rigid transformation (Walker et al., 1991) (reproduced with the permission of R.A. Volz, copyright Academic Press, 1991).

[1] A *dual number* $\hat{a} = a + \varepsilon b$ is defined as a combination of two real numbers a and b, with a special multiplication rule given for ε given by $\varepsilon^2 = 0$, such that, e.g. $(a + \varepsilon b)(c + \varepsilon d) = ac + \varepsilon(ad + bc)$.

2.2. Nonrigid transformations

In the event that surface deformations must be accounted for in the transformation, a nonrigid approach must be adopted. The simplest classes of nonrigid transformations are basically *generalizations of the rigid body transformation*. A more general formulation is one where a *global polynomial function* maps surface *A* to surface *B*. To make the mapping even more general, a global polynomial can be replaced by *piecewise polynomial relations*, or *splines*, and such functions can be seen as either *interpolating*, whereby the transformation they express is directly determined by the motion between two sets of primitives (and as such is sensitive to errors in the data), or *approximating*, whereby the agreement between the local motion of any pair of matches is reconciled with global consistency, in some optimal manner.

2.2.1. Generalization of rigid body motion: the affine transformation

The generalization of the rigid body transformation relevant to 3D surface matching is the *affine* transformation. The general **affine** class of transformations is characterized by the expression

$$\boldsymbol{x}_B = \boldsymbol{A}_{3\times 3}\boldsymbol{x}_A + \boldsymbol{b}_{3\times 1}, \tag{7}$$

where there is no orthogonality constraint on the elements a_{ij} of matrix $\boldsymbol{A}$ as there is in Eq. (2). Affine transformations do not in general preserve angles or lengths, but parallel lines remain parallel (Foley et al., 1990). The affine transform relating two 3D surfaces is solved by Feldmar and Ayache (1994a), whose method matches ''closest'' points of compatible shape. This ensures that the solution to the minimization $\min_{\boldsymbol{A},\boldsymbol{b}}\ (1/N)\ \Sigma_{i=1}^{N}\ \|\boldsymbol{A}\boldsymbol{x}_{A_i} + \boldsymbol{b} - \boldsymbol{x}_{B_i}\|^2$ tends to a stable solution. Moreover, Henri et al. (1991) determine the translation, rotation and nonuniform scaling which maps stereotactic frame coordinates to corresponding CT or MR voxel values, prior to superimposing stereoscopic DSA images on the equivalent projections of the volumetric scan.

2.2.2. Global polynomial functions

A more general formulation is one where a global polynomial function, typically of order 2–5, maps surface *A* to surface *B* (Lavallée, 1996). Global methods use matched point pairs to generate a single optimal transformation, based on a sufficient number of points to (over)determine the parameters of the transformation, via either approximation or interpolation. Polynomial transformations are typically expressed in either cartesian or spherical coordinates. Moreover, global polynomial transformations are only useful to account for low-frequency distortions because of their unpredictable behaviour when the degree of the polynomial is high (Brown, 1992). This component of inter-surface motion is typically computed after an initial rigid alignment.

Approximation is the search for transformation parameters which map matched points or features *as closely as possible globally, but not necessarily perfectly individually*, whereas interpolation finds the transformation which maps two 3D surfaces so that *matched control points are exactly satisfied*. In the former case we assume that some noise or unwanted distortion exists and should not be accounted for in the transformation. For large numbers of control points, this choice makes sense because matches likely include inaccuracies, but taken together they contain sufficient statistical data to make the transformation reliable. Interpolation is more appropriate for a few accurately-matched control points, since it involves an independent parameter for each control point match, which may result in unexpected undulation for a high-order fit.

For a *cartesian* formulation, the global polynomial transformation in three dimensions can be stated as follows:

$$\begin{aligned} x_B &= \sum_{ijk} a_{ijk} x_A^i y_A^j z_A^k, \\ y_B &= \sum_{ijk} b_{ijk} x_A^i y_A^j z_A^k, \\ z_B &= \sum_{ijk} c_{ijk} x_A^i y_A^j z_A^k, \end{aligned} \tag{8}$$

where a_{ijk}, b_{ijk} and c_{ijk} are the constant polynomial coefficients to be determined. If interpolation is used, these coefficients express a system of $3N$ unknowns which can be determined by N control points. In a least squares approximation, the sum over all matched feature/point pairs of the squared difference between the left and right-hand side of these equations is minimized, for example by setting the partial derivatives associated with these equations to zero.

Jacq and Roux (1993) implement a trilinear interpolation (i.e. where the summation indices *i, j* and *k* each go from 0 to 1, and where $i + j + k \leqslant 3$), determined by eight reference distortion values which span the volume to be warped. Subsol et al. (1994) register skulls and cortical surfaces to build a 3D atlas, by first rigidly matching crestlines and then using an iterative closest point algorithm (both techniques are discussed in Section 3) on crestline points to determine a global second order polynomial to describe the relation between iteratively transformed surface *A* and surface *B*. This results in a 2^n-order polynomial transformation, based on n iterations.

For closed surfaces that can be modelled as functions on a sphere, some authors prefer to work in spherical coordinates (Coppini et al., 1987; Chen et al., 1994). For example, **spherical harmonic** surfaces are closed surfaces on a sphere that can be decomposed into a set of orthogonal functions. To represent an arbitrary shape, the radius $r(\theta, \phi)$ in the spherical coordinate system (centered on the centroid) can be written as a linear sum of spherical harmonic basis functions:

$$r(\theta, \phi) \approx \sum_{n=1}^{N} \sum_{m=0}^{n} [A_{nm} U_{nm}(\phi, \theta) + B_{nm} V_{nm}(\phi, \theta)], \quad (9)$$

where A_{nm} and B_{nm} are basis coefficients computed from the data points from a 3D object or surface and N represents the order of the fitting. The basis functions are:

$$U_{nm}(\theta, \phi) = \cos(m\phi) P_{n,m}(\cos(\theta))$$
$$\text{and } V_{nm}(\theta, \phi) = \sin(m\phi) P_{n,m}(\cos(\theta)), \quad (10)$$

where $P_{n,m}(x) = (1 - x^2)^{m/2} \, \mathrm{d}^m / \mathrm{d}x^m \, P_n(x)$ and $P_n(\cdot)$ is the Legendre polynomial of degree n (Press et al., 1992). In other words, surfaces S_A and S_B, and the displacement between them, can be seen as functions on a sphere, represented by a set of real coefficients A_{nm} and B_{nm}. Because the shape is approximated as a sum of different harmonics, in theory this representation can reconstruct high-frequency surface detail.

Coppini et al. (1987) model the epicardial stretch tensor, based on tracked vascular bifurcations, by performing a third-order spherical harmonic fitting over individual displacements $r_A(\theta_i, \phi_i) - r_B(\theta_i, \phi_i)$, after correction for translation and rotation. Chen et al. (1994) adopt a similar method, but first characterize global shape with **superquadric** surfaces (Barr, 1981; 1984; Bajcsy and Solina, 1987), prior to characterizing local shape variation with spherical harmonics.

2.2.3. Local nonrigid transformations: piecewise polynomials

Global mapping functions do not always adequately capture deformations of anatomical structures, which are often intrinsically local (Lavallée, 1996; Bookstein, 1989). *Piecewise polynomial functions* produce a more general relation. In general, the relative density and reliability of the data determine whether an interpolation or an approximation scheme is used.

One well-documented interpolating scheme is the **thin-plate spline** (Duchon, 1976; Bookstein, 1989). The thin-plate spline over a 2D domain can be expressed as $z(x, y) = U(r) = r^2 \log r^2$, where $r = \sqrt{x^2 + y^2}$, and as such U is a fundamental solution of the biharmonic equation. The interpolant $f(\boldsymbol{p})$ is optimal in that it has *minimum bending energy* amongst all functions which pass exactly through points $\boldsymbol{x}_i = (\boldsymbol{p}_i, z(\boldsymbol{p}_i))$; i.e.

$$\int\int_{\mathbb{R}} \left[\left(\frac{\partial^2 f}{\partial x^2} \right)^2 + \left(\frac{\partial^2 f}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 f}{\partial y^2} \right)^2 \right] \mathrm{d}x\mathrm{d}y \quad (11)$$

is minimized. This type of interpolation function can then be applied to modelling nonrigid motion. Instead of having $f(\boldsymbol{p})$ represent a displacement in the z-direction over the $\boldsymbol{p} = (x, y)$ domain, it can express the x-component of a deformation. Likewise functions $g(\boldsymbol{p})$ and $h(\boldsymbol{p})$ can express the y- and z-components of a deformation (Bookstein, 1989; Evans et al., 1991), where now $\boldsymbol{p} = (x, y, z)$. This method presupposes the use of relatively sparse point matches in order to determine a set of local polynomials or spline functions exactly, when in fact it may be useful to consider denser displacement information to describe a fundamentally underdetermined relationship. Moreover, the notion of minimum bending is better suited to deformation over a 2D, rather than 3D, domain.

Another useful piecewise polynomial representation is the **B-spline**, particularly in the context of *smoothing* (regularizing) and *least-squares* spline approximation (Dierckx, 1995). Given a set of displacement data $\boldsymbol{d}(x_r, y_r, z_r)$, and three sets of knots λ_i, $i = 0, \ldots, f + 1$, μ_j, $j = 0, \ldots, g + 1$, and ν_k, $k = 0, \ldots, h + 1$, we can compute the least-squares pth order volumetric spline

$$\boldsymbol{u}(x, y, z) = \sum_{i,j,k} \boldsymbol{u}_{i,j,k} B_i(x) B_j(y) B_k(z), \text{ where, e.g.}$$
$$B_i(x) = \sum_{l=i}^{i+p+1} \frac{(\lambda_l - x)^p}{\lambda_l - \lambda_i} \quad \text{if } x \in [\lambda_i \lambda_l] \text{ and}$$
$$B_i(x) = 0 \quad \text{elsewhere}, \quad (12)$$

and such that $\delta = \Sigma_r \, (w_r(\|\boldsymbol{d}_r - \boldsymbol{u}(x_r, y_r, z_r)\|))^2$ is minimized. Here, the determination of B-spline coefficients of $\boldsymbol{u}_{i,j,k}$ is by the least-squares solution of an overdetermined linear system. For a smoothing spline approach, the problem is to find the function $\boldsymbol{u}_p(x, y, z)$ minimizing a smoothing norm which is a function of the B-spline coefficients, subject to $\delta < S$. Szeliski and Lavallée (1996) model the nonrigid transformation between two anatomical surfaces as a first order spline in x, y and z, which is constrained by zeroth and first order stabilizers. This approach penalizes large variations of the spline coefficients, while also enforcing agreement with displacement data. We (Audette and Peters, 1999) have recently demonstrated the use of 2D *recursive* splines (Unser et al., 1993a, 1993b) to efficiently characterize nonrigid cortical motion undergone during brain surgery, as captured by a time sequence of range images.

Lastly, a few other ways of characterizing nonrigid motion appear in the literature. Goldgof and Mishra classify the nonrigid motion of surfaces in terms of how it affects their mean and Gaussian curvature properties, namely as *rigid*, *isometric*, *homothetic*, *conformal* and *general nonrigid*[2] (Goldgof et al., 1988a; Mishra et al., 1991). Moreover, Feldmar and Ayache (1994b) determine *locally* affine transformations for individual surface points $(\boldsymbol{A}_{A,i}, \boldsymbol{b}_{A,i})$ by a weighted sum of the rigid transformation

[2]Conformal motion is characterized by *proportionality* of the coefficients of the first fundamental form (do Carmo, 1976):

$$\frac{E_A}{E_B} = \frac{F_A}{F_B} = \frac{G_A}{G_B} = \eta(u, v).$$

This function $\eta(u, v)$ becomes a *constant* over (u, v) for homothetic motion, and *identity* for isometric motion.

parameters of locally neighbouring matched surface points $(\boldsymbol{R}_{A,k}, \boldsymbol{t}_{A,k})$, where $k \neq i$.

3. Surface representation and similarity criterion

The second stage of registration consists of computing a surface representation and defining a matching criterion based on it. In general, the surface representation should be stable over the two modalities or over the time sequence considered. It should afford a similarity criterion which is sufficiently discriminating to associate homologous points unambiguously and efficiently, if the application dictates that algorithmic performance is an issue. There are four approaches to representing a surface for the sake of registration: **feature**, **point**, and **model-based** methods as well as techniques based on **global similarity**. The definition of the similarity criterion is typically closely related to the choice of matching primitive. Furthermore, the criteria for selecting a particular primitive are application-specific, depending for example on whether the anatomy of interest is smooth, such as the cranium, or highly angular, such as a vertebra. Other factors which influence this choice include the size of the transformation to be computed (i.e. whether the two surfaces are separated by an arbitrarily large transformation or roughly aligned), and whether the transformation is rigid or non-rigid.

The feature-based method attempts to express surface morphology as a set of features which are extracted by a preprocessing step. Such features provide a *compact description of the surface shape* (at the expense of losing information), which is quantifiable by stable, discriminating scalar measurements. The similarity criterion is then an outgrowth of this feature characterization: it consists of a comparison of scalar measurements. The point and model-based methods do not attempt to reduce the surface representation to a more compact description, rather they use *all, or a large subset of all, points*. Generally, for the point-based method, the primitive used is often the surface point itself, and the similarity criterion is a distance to be minimized between a pair of surface points. For model-based approaches, often an implicit criterion is used, such as an external force or halting condition driven by two sets of image data, with which an evolving deformable surface model must be reconciled. Finally, a new class of registration methods matches surfaces typically on the basis of their *global* similarity. While there are currently few anatomical applications, the relevance of these methods stems from the feasibility of precomputing a number of training views of a surface generated from a patient's tomographic data, and from the desirability of computing arbitrarily large transformations for smooth, relatively featureless surfaces.

The feature-based and global approaches are potentially more discriminating than point- or model-based techniques, and can therefore resolve a large motion or transformation. On the other hand, most point and model-based methods may be attractive in the case of small or iteratively estimated motion, because they exploit large redundancy of information, which is especially useful for estimating locally nonrigid transformations, as these are inherently underdetermined.

3.1. Feature-based representations

Feature-based matching is largely founded on the use of differential geometry to describe local surface shape (do Carmo, 1976; Besl and Jain, 1986). According to Bonnet's fundamental existence and uniqueness theorem, if two surfaces S_A and S_B possess equivalent fundamental forms I and II (or equivalent Gaussian and Mean curvatures K and H), then there exists an appropriate translation and rotation such that S_A and S_B coincide exactly (Besl and Jain, 1986). In general, feature-based matching is applied to computing *rigid* transformations.

Features used for surface registration fall into three categories: (sparse) *point features* (distinct from dense point-based schemes, also referred to as free-form surface registration), *curves* and *regions*. Point features are salient, well-localized, sparse loci of important geometric significance, such as extrema of curvature: local peaks, pits, saddle points where the two principal curvatures are most pronounced, or where K is at a local minimum or maximum. The second type of feature corresponds to contiguous lines or curves, consisting typically of differential structures such as ridges or boundaries between regions. Regions, in turn, are areas possessing some homogeneous characteristic, such as consistent curvature sign. Each feature in surface S_A can be matched with its homolog in S_B by first characterizing each feature in either surface by parameters expressing its respective topology, and looking for a compatible vector of parameters in the other surface. If there is more than one suitable candidate for a given feature, the match can be disambiguated by assuming that neighbouring features on the same surface which are matched one-to-one should undergo motion consistent with the ambiguous candidate. False candidates can then be eliminated on the basis of motion inconsistency.

Moreover, *accuracy* issues related to feature extraction include the use of the neighbourhood information around each surface point to stabilize the computation of its differential properties (Sander and Zucker, 1990; Ferrie et al., 1993) and the sensitivity to noise statistics of a given numerical method for estimating derivatives and surface curvature (Flynn and Jain, 1989; Abdelmalek, 1990; Roth and Levine, 1993). Finally, another relevant area of research is the application of *recursive* infinite impulse response filters (Proakis, 1996), particularly near-Gaussian exponential filters (Shen and Castan, 1986; Deriche, 1987; Deriche, 1990) to speed up the smoothing stage prior to feature extraction (Monga et al., 1992; Thirion, 1994).

Point features can be matched on the basis of their intrinsic information, such as surface curvature values, as well as position relative to *neighbouring* point features. One such feature is the *extremum of principal curvatures*, computed by Thirion (1994) within volumetric data by detecting the zero-crossings of two ''extremality'' functions. The first function is discussed later in this section, and is used to detect *ridge lines*, or lines of locally maximal higher principal curvature, while the second function finds a local extremum of the lesser principal curvature within each ridge line. The scalar measurements used in the matching include the type of extremal point, which depends on the sign of the principal curvatures, the response of the extremality functions, the principal curvature values, and the distances and orientations of vectors to neighbouring points. An alternate representation, used by Goldgof et al. (1988b) in terrain matching but nonetheless applicable to anatomical surfaces if they possess sharp corners, is the *extremum of Gaussian curvature* (K), detected on the basis of a threshold on the value of K.

A somewhat less compact way of representing surface shape is as a collection of **curves**. One particularly relevant curvilinear feature in medical images is the *ridge or crest line* (e.g. the trough of a sulcus or the maximum height of a gyrus). As pointed out by Maintz et al. (1996), there are many ways to define a ridge, and consequently many ways to detect one. Monga et al. (1992) and Monga and Benayoun (1995) look for a contiguous set of loci of a surface where the largest principal curvature κ_1 is locally maximal. These correspond to the zero-crossings of the extremality function $e_1 = \nabla\kappa_1 \cdot \boldsymbol{t}_1$, where $\nabla\kappa_1$ is the directional derivative of the largest principal curvature, and $\boldsymbol{t}_1$ is the principal direction corresponding to κ_1.

Maintz et al. (1996) propose two operators of their own which stem from the consideration of the gradient $\boldsymbol{w}$ of a smoothed surface and its right-handed normal $\boldsymbol{v}$ in the proximity of a ridge. The gradient at any non-ridge position *points towards* the ridge, but at a ridge position the gradient is *aligned with* the ridge. A consequence of this geometry is that the directional derivative along $\boldsymbol{v}$ of a surface is characterized by a highly concave profile at a ridge point in comparison with other points. Maintz then detects a ridge point either as a minimum of the *second derivative of the surface*, or as a maximum of the derivative of the *direction of the gradient*, along $\boldsymbol{v}$. An alternative to a ridge-based curve is the *distance contour* (Radack and Badler, 1989), which could be applied to anatomical data, and which is the set of points of constant distance from a highly salient point.

Guéziec and Ayache (1994) present an elegant technique for characterizing and matching curves, which in turn is based on the curve matching algorithm of Kishon et al. (1990). Kishon addresses the problem of finding the longest matching subcurve appearing in two curves. This uses local, rotationally and translationally invariant, stable shape signatures, namely *curvature* $\kappa(s)$ and *torsion* $\tau(s)$, of smoothed curves[3]. Each shape signature is represented as a *hash table* (see Section 4), where entries are associated with pairs of curvature–torsion values (κ, τ). The improvements introduced by Guéziec and Ayache relate to the approximation of the curves, to the hash table and to the statistical analysis of various invariants for matching, particularly in directly determining a 3D transformation from homologous points on matched curves.

Regions constitute an even denser feature-based representation. They are typically characterized either by the homogeneity of their local surface shape or as being circumscribed by some boundary. One advantage is that the notion of neighbourhood between regions is natural (Toriwaki and Yokoi, 1988), and leads to the characterization of a surface as an adjacency graph. Matching is then carried out as an exercise in finding the maximal *clique* of compatible subgraphs (see (Radig, 1984) and the references therein). Here, groups of regions are matched based on neighbourhood topology as well as local region characteristics. One definition of homogeneity used to segment surfaces into regions or surface patches is the K/H *sign combination* (Besl and Jain, 1986; Kehtarnavaz and Mohan, 1989). Surfaces can be subdivided into patches according to surface type: elliptic and outwardly bulging ($K > 0/H < 0$), elliptic and inwardly bulging ($K > 0/H > 0$), parabolic and outwardly bulging ($K = 0/H < 0$), parabolic and inwardly bulging ($K = 0/H > 0$), planar ($K = 0/H = 0$) and hyperbolic or saddle shaped ($K < 0/H \neq 0$). In practice, a small nonzero threshold is used to determine sign. An illustration of a K/H surface representation for a range image of a femur epiphysis appears in Fig. 3.

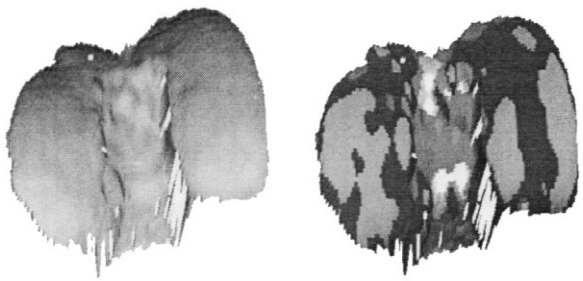

Fig. 3. (a) Range image of femur epiphysis; (b) regions based on K/H sign combination. Both are courtesy of Gilbert Soucy of the McGill Centre for Intelligent Machines.

[3]For a curve $\boldsymbol{\alpha}$: $I = (a, b) \rightarrow \mathbb{R}^3$, parametrized by arclength s, *curvature* $\kappa(s)$ is the scalar $\|\alpha''(s)\|$. It is a measure of how rapidly the curve pulls away from the *tangent* line $\boldsymbol{t}(s)$. The direction of $\alpha''(s)$ is given by the unit *normal* $\boldsymbol{n}(s)$, and the unit vector $\boldsymbol{b}(s) = \boldsymbol{n}(s) \times \boldsymbol{t}(s)$ is called the *binormal* vector. Torsion $\tau(s)$ corresponds to $\|\boldsymbol{b}'(s)\|$, and it is a measure of how quickly the curve pulls away from the *osculating plane* at s spanned by $\boldsymbol{n}$ and $\boldsymbol{t}$ (do Carmo, 1976).

3.2. Point-based methods

Point-based methods register surfaces on the basis of relatively dense point sets brought into correspondence, where these point sets constitute all, or a significant subset of, the available surface point samples. An alternate name for surface matching based on dense point sets is **free-form** surface matching (Besl and McKay, 1992; Zhang, 1994). The two point sets are generally assumed to be relatively close to being aligned, and are usually registered by *iteratively* minimizing a global function such as the sum of squared distances between mutually closest points between (possibly transformed) surface S_A and surface S_B, as expressed by Eq. (3). Differences between many of these methods exist strictly at the level of the choice of distance metric and of the methods of optimally finding a match based on this metric (described in detail in Section 4). An illustration of this technique appears in Fig. 4.

One common distance metric is the distance from a point $\boldsymbol{x}_B$ in set $X_B = \{\boldsymbol{x}_B\} \subseteq S_B$ to the (transformed) point set $X_A = \{\boldsymbol{x}_A\} \subseteq S_A$:

$$d(\boldsymbol{x}_B, \boldsymbol{R}_k \boldsymbol{x}_{A,min} + \boldsymbol{t}_k) = \min_{\boldsymbol{x}_A \in X_A} d(\boldsymbol{x}_B, \boldsymbol{R}_k \boldsymbol{x}_A + \boldsymbol{t}_k). \quad (13)$$

Besl and McKay (1992) propose the *Iterative Closest Point* (ICP) method to determine the closest point pairs according to Eq. (13), then compute the transformation from these pairs with a quaternion technique. The positions of the surface points S_A (''data shape'') are then updated: $X_{A,k} = \boldsymbol{R} X_{A,k-1} + \boldsymbol{t}$ and the process iterates until the mean-square distance, or point matching error, stabilizes to within some tolerance. An accelerated variant of the ICP method (discussed in Section 4) is also proposed. The method is better adapted to registering comparable patches, but a subpatch can also be put into correspondence with a larger patch, at the cost of considering several ''initial translation states''.

Zhang (1994) and Lavallée and Szeliski (1995) adopt an objective function identical to that proposed by Besl, except for weighting factors to accomodate measurement noise, and in the case of Zhang, to exclude outliers. They also introduce execution efficiencies: Zhang uses *k-D trees* to make the search for closest points more efficient, and Lavallée dispenses with this search altogether by pre-computing a *distance map*, which unfortunately does not provide explicit point pairs and thus still entails a search for optimal transformation parameters. Simon et al. (1994a, 1994b) apply this technique to intrasurgical registration of a range image of the head with a surface extracted from tomographic data and address the issue of sensitivity to the perturbation of individual points. We (Audette and Peters, 1999) also use a free-form surface technique to characterize intrasurgical non-rigid cortical motion, where we refine the distance map proposed by Lavallée: a *closest-point map* is produced, which provides explicit point pairs suitable for closed-form transformation computation.

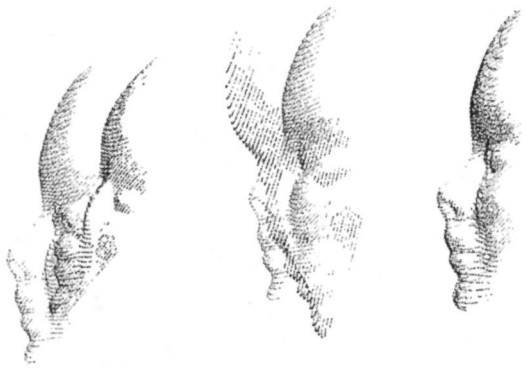

Fig. 4. Iterative point-based registration of phantom face range data (courtesy of S. Lavallée, R. Szeliski and L. Brunie, copyright MIT Press, 1986 [Lavallée et al., 1996]).

Other point-based techniques, which differ from those discussed above in terms of the distance metric which is minimized, have been proposed by Chen and Medioni (1991), by Pelizzari et al. (1989) by Soucy and Ferrie (1997), and by Rangarajan et al. (1997). Chen uses a subset of *control points* in relatively smooth areas, leading to an iterative technique with very good convergence properties. For each surface normal $\boldsymbol{n}_{A,i}$ defined at a control point $\boldsymbol{x}_{A,i}$, its intersection $\boldsymbol{x}_{B,i}$ with surface S_B is found. Next, defining $\boldsymbol{s}_{B,i}$ the plane tangent to S_B at $\boldsymbol{x}_{B,i}$, the global transformation which minimizes the sum of squared distances between the set of transformed $\boldsymbol{x}_{A,i}$ with their corresponding *tangent planes* $\boldsymbol{s}_{B,i}$ on S_B is computed. Pelizzari fits a ''hat'' or external surface, consisting of relatively sparse points from the scalp as imaged by the modality of lesser resolution or coverage, to a ''head'' constituted by the set of 2D contours extracted on a slice-by-slice basis from the higher coverage/resolution modality. The residual which is minimized is the sum of distances from each hat point to the head surface, along a direction from the former point to the head centroid. Soucy proposes an iterative technique which takes local surface shape into account. It matches small surface patches (as small as 3×3 pixels) by minimizing a similarity functional which enforces compatible local shape and piecewise-smooth motion. Rangarajan matches points with a technique (demonstrated on 2D autoradiograph slices but equally applicable to 3D surfaces) which imbeds the search for match pairs and for optimal transformation parameters, as well as the explicit exclusion of outliers, into *one* elegant minimization, which makes the method more robust to initial transformation estimation than the ICP technique.

Finally, Feldmar and Ayache (1994a,b) do not attempt strictly to find closest points, but to find closest **feature vectors**. This approach is a comparison of 8 parameters, namely the coordinates of each point (x, y, z), the components of its normal (n_x, n_y, n_z) and the principal curvatures κ_1 and κ_2 of the surface at that point. They

determine a *globally affine* transformation and then possibly a set of *locally affine* transformations, by extended Kalman filtering (see Section 4), based on point-pairs which minimize the following expression:

$$\begin{aligned} d(\boldsymbol{v}_{A,i}, \boldsymbol{v}_{B,j}) &= \alpha_1(x_{A,i} - x_{B,j})^2 + \alpha_2(y_{A,i} - y_{B,j})^2 \\ &+ \alpha_3(z_{A,i} - z_{B,j})^2 + \alpha_4(n_{xA,i} - n_{xB,j})^2 \\ &+ \alpha_5(n_{yA,i} - n_{yB,j})^2 + \alpha_6(n_{zA,i} - n_{zB,j})^2 \\ &+ \alpha_7(\kappa_{1A,i} - \kappa_{1B,j})^2 + \alpha_8(\kappa_{2A,i} - \kappa_{2B,j})^2, \end{aligned} \tag{14}$$

where α_i are weights determined by minimum and maximum data values. In other words, the closest point to $x_{A,i}$ on S_B is one which best fulfills a compromise between 3D distance, difference in normal orientation, and difference of principal curvatures.

3.3. Model-based representations

Deformable surface modelling consists of expressing surface identification (segmentation) in one volume or surface tracking over a volume sequence, as a model which reconciles the likely shape and/or dynamic behaviour of the surface, according to some **physically-based** or **surface evolution** expression, with raw image data. With some exceptions (Amini and Duncan, 1992; Sclaroff and Pentland, 1993), such methods generally compute curve and surface motion not by explicit matching, but by implicit consideration of image motion in the form of virtual forces that tend to make the model agree with shifts in strong image gradients. Because the thrust of active surface models is mostly on segmentation, only influential techniques and those which emphasize registration are described in detail. For more on active contour and surface models, the reader is referred to McInerney and Terzopoulos (1996) for physically-based models and to Sethian (1996b) and to Kimmel et al. (1997) for surface evolution models. One important issue in segmentation which has an impact on the accuracy of the registration is **anisotropic filtering**, whereby smoothing which is required for surface extraction is carried out tangentially to the surface, without washing out image features in the direction normal to it (Perona and Malik, 1990; Alvarez et al., 1992; Kimia and Siddiqi, 1996).

For physically-based deformable surface models, the basic idea is to model an object which is tracked over time as being in, or quickly reaching, an equilibrium between internal and external virtual forces. **Internal forces** include inertia, damping and strain, and are typically determined by the current state which is assumed for a 3D object. **External forces** are typically determined by image data. The relative motion undergone by the various parts of a surface is a result of the interaction of these two sets of forces, according to formal physical principles.

Important early research includes Kass' *snakes* paper (Kass et al., 1987), Terzopoulos' symmetry-seeking 3D shape model (Terzopoulos et al., 1988) and Bajscy's multiresolution elastic matching (Bacjsy and Kovacic, 1989). Kass casts the 2D contour detection and tracking problem of intensity images as a minimization of the following energy functional:

$$E_{\text{snake}} = \int_0^1 E_{\text{int}}(\boldsymbol{v}(s)) + E_{\text{image}}(\boldsymbol{v}(s)) + E_{\text{con}}(\boldsymbol{v}(s))\, \text{ds}, \tag{15}$$

which is a controlled-continuity spline under the influence of *image* forces and *external constraint* forces. Terzopoulos' model is expressed as

$$\mu \frac{\partial^2 \nu}{\partial t^2} + \gamma \frac{\partial \nu}{\partial t} + \frac{\delta \mathscr{E}(\nu)}{\delta \nu} = \boldsymbol{f}(\nu), \tag{16}$$

and consists of a deformable sheet of elastic material, which is rolled to form a tube, through which passes a deformable spine, also of elastic material. *Coupling* forces try to make the shape retain its axial symmetry, and *extrinsic* forces constrain the shape to be consistent with one or more 2D image projections. The Bajscy model is

$$\mu \nabla^2 u_i + (\lambda + \mu) \frac{\partial \theta}{\partial x_i} + F_i = 0, \quad i = 1, 2, 3, \tag{17}$$

where θ is the *dilatation*[4] at a point on the body, and μ and λ are *Lamé's constants*, which define the elastic properties of the model. The *external forces* $\boldsymbol{F} = (F_1, F_2, F_3)^T$ bring similar regions of two 3D objects into correspondence by enforcing grey-level value correlation and edge alignment between blurred volumes (in a multiresolution framework).

Research conducted separately by Pentland (Pentland and Sclaroff, 1993; Essa et al., 1993) and by Terzopoulos (Terzopoulos and Metaxas, 1991; McInerney and Terzopoulos, 1995) proposes a finite-element model approach to the numerical solution of deformable surface models, and consequently these surface models are based on the *finite-element equilibrium equation*, which has the form

$$\boldsymbol{M}\ddot{\boldsymbol{U}} + \boldsymbol{C}\dot{\boldsymbol{U}} + \boldsymbol{K}\boldsymbol{U} = \boldsymbol{R}, \tag{18}$$

where $\boldsymbol{M}$, $\boldsymbol{C}$ and $\boldsymbol{K}$ are virtual *mass*, *damping* and *stiffness* matrices, $\boldsymbol{U}$ is the displacement of the FEM nodes, and $\boldsymbol{R}$ is the sum of external forces, determined by image data. Pentland and Terzopoulos use a hybrid representation featuring a superquadric ellipsoid upon which is grafted a displacement function, which in turn is estimated by finite-element modelling. In a static segmentation context, this displacement function represents the difference between the simple superquadric shape and the final, more general shape which is more in keeping with image forces. In a surface tracking context, a general shape at t_k is used as a

[4] This is the change in volume per unit initial volume under small strains and sum of principal strains in general (Malvern, 1969).

first estimate towards the final shape at t_{k+1}, and the displacement is just $\boldsymbol{U}_{k+1} - \boldsymbol{U}_k$.

Significant improvements to the static model are suggested by Cohen et al. (Cohen and Cohen, 1990; Cohen, 1991; Cohen et al., 1992b), and by Metaxas et al. (Metaxas and Kakadiaris, 1996; DeCarlo and Metaxas, 1996). If the model is not initialized close enough to the desired surface, short-range image forces may be unable to attract it. To alleviate this problem, Cohen (Cohen and Cohen, 1990; Cohen, 1991) uses a *pressure force* to inflate the model towards the object surface. In (Cohen et al., 1992b), the surface produced by a physical model is characterized in terms of its differential structure, and the model is then expanded to enforce the agreement of the surface's *normal orientation* with that estimated by a Monga–Deriche edge detector (Monga et al., 1991). Metaxas devises methods for adaptively estimating *virtual material properties* (Metaxas and Kakadiaris, 1996), and for using *blending functions* in order to merge two simple (e.g. superquadric) shapes to express more complex surfaces (DeCarlo and Metaxas, 1996).

Improvements relating to surface tracking are proposed by Nastar and Ayache (1993) and by Sclaroff and Pentland (1993). If physically based models can be viewed as virtual masses on the surface with *springs* relating them to their neighbours, this viewpoint can be extended over time to a spring between a boundary point at t_1 and its closest neighbour at t_2, as pointed out by Nastar. Moreover, salient feature points such as curvature extrema are used to anchor the dynamic behaviour of the model, by attaching particularly stiff springs between feature point pairs. Sclaroff suggests that FEM numerical estimation can benefit from a change of basis from *nodal* to *modal* displacements ϕ, where $\boldsymbol{U} = \phi \sin[\omega(t - t_0)]$ and ω represents a virtual frequency of vibration, based on object shape, corresponding to each mode (see also the section in (Bathe, 1982) on mode superposition). This change of basis is justified not on the grounds of expected periodic motion, but because it decouples the system of Eq. (18), which is now expressed as the eigenproblem (neglecting damping by taking $\boldsymbol{C} = \boldsymbol{0}$):

$$\boldsymbol{K\Phi} = \boldsymbol{M\Phi\Omega}^2, \quad \text{where } \Phi = [\phi_1 | \cdots | \phi_p]$$
$$\text{and } \Omega^2 = \begin{bmatrix} \omega_1^2 & & \\ & \ddots & \\ & & \omega_p^2 \end{bmatrix}. \tag{19}$$

The matrix Φ has an interesting interpretation: its vector entries ϕ_i can be ordered according to increasing corresponding eigenvalue ω_i^2. In this case Φ is an orthogonal, frequency-ordered description of on object's shape and its natural deformations, somewhat like a Fourier series. An important consequence of this representation is that it leads to an *explicit* matching algorithm, after an initial rigid alignment, where low-order nonrigid modes ϕ_i are used as a coordinate system, and two shapes are nonrigidly registered on the basis of modal feature vectors.

In contrast to the preceding models, those developed independently by Malladi et al. (1995) and by Caselles et al. (1992) have a geometric, rather than physical, interpretation, and are generally based on a *curve or surface evolution equation*, a partial differential equation of the Hamilton-Jacobi type (Levesque, 1992, Kimmel et al., 1997). These techniques are attractive for registration because they inherently compute contour or surface shape parameters, and possess other advantages, such as robustness to initialization as well as the capacity for a contour or surface to split or to merge. Here, a contour or surface is viewed as the zero level of a higher-dimensional function. For example, to identify a contour in 2D, a 3D hypersurface is initially placed completely inside (or outside) such a shape, and then made to flow outward (or inward) with a speed dependent on the strength of the image gradient $g(I)$, on surface curvature $\mathrm{div}(\nabla u / \|\nabla u\|)$ and on a constant advection term ν which acts as an inflation (or deflation) force:

$$\frac{\partial u}{\partial t} = g(I)\|\nabla u\| \left[\nu + \mathrm{div}\left(\frac{\nabla u}{\|\nabla u\|} \right) \right]. \tag{20}$$

The factor $\mathrm{div}(\nabla u / \|\nabla u\|)$ is the curvature κ of the level-set contour in 2D. In 3D, it is the mean curvature H of the level-set surface. The image gradient-based factor acts as a halting criterion which binds the level-set surface to intensity discontinuities. Recent improvements include a *doublet* term (Caselles et al., 1995, 1997; Kichenassamy et al., 1995) which prevents the surface from overshooting past image gradients, and the application of the *fast marching* level sets algorithm (Sethian, 1996a; Malladi and Sethian, 1998).

Lastly, the model of Amini and Duncan (1992) uses 3D surface points as its raw data, and views the local surface shape as a bending energy from an idealized thin flat plate $\varepsilon_{\mathrm{be}}(u, v) = \kappa_1^2 + \kappa_2^2$. It seeks to match surface points of consistent energy and therefore of consistent principal curvature. Moreover, their model also has a stretching energy term which penalizes *non-conformal* motion (i.e. motion where the proportionality of the various coefficients of the first fundamental form, E, F and G, are not maintained). Consequently, the following energy measure is minimized over two surfaces:

$$\lambda_{\mathrm{be}}\{(\kappa_{1A} - \kappa_{1B})^2 + (\kappa_{1A} - \kappa_{1B})^2\} + \lambda_{\mathrm{st}} \left\{ \left(\frac{E_A}{E_B} - \frac{F_A}{F_B} \right)^2 + \left(\frac{F_A}{F_B} - \frac{G_A}{G_B} \right)^2 + \left(\frac{E_A}{E_B} - \frac{G_A}{G_B} \right)^2 \right\}. \tag{21}$$

The tracking model of Cohen et al. (1992a) is similar to Amini's in that it is characterized by a functional that minimizes curvature differences, between two 2D contours at comparable arclength parameters, while enforcing smooth motion along arclength. The zero arclength value

on each contour corresponds to a matched feature point. They propose an extension to 3D surface tracking in the same paper.

3.4. Techniques based on global shape

The feature, point and model-based algorithms discussed so far can be broadly described as relying on local information to register surfaces. However, there are a few recently-published algorithms which register surfaces on the basis of global surface geometry, that do not rely on a rough prior estimation of the transformation and that may be able to deal with relatively featureless patches. We are alluding to the **spinmap** representation of Johnson and Hebert (1998) and to the **eigenshape** or **appearance-based** methods for registering 3D surfaces (Campbell and Flynn, 1999).

The spinmap representation is a set of 2D footprints conveying global surface shape in the neighbourhood of selected *oriented points* (Johnson and Hebert, 1998), as illustrated in Fig. 5. A 3D oriented point consists of a point $\boldsymbol{p}$ and a surface normal $\boldsymbol{n}$ estimated at that point. In order to describe its neighbouring topology, a 2D basis is formed. To define this basis, Johnson first defines a line L through $\boldsymbol{p}$ whose direction is along $\boldsymbol{n}$, and a tangent plane $\boldsymbol{P}$ through $\boldsymbol{p}$. The two coordinates of the bases are α, the perpendicular distance of a neighbouring point to line L, and β, the perpendicular distance of this point to plane P. A spinmap S_O is the function which maps 3D neighbouring points $\boldsymbol{x}$ to the 2D coordinates of a basis corresponding to an oriented point O:

$$S_O(\boldsymbol{x}) \rightarrow (\alpha, \beta) = \left(\sqrt{\|\boldsymbol{x}-\boldsymbol{p}\|^2 - (\boldsymbol{n}\cdot(\boldsymbol{x}-\boldsymbol{p}))^2}, (\boldsymbol{n}\cdot(\boldsymbol{x}-\boldsymbol{p}))\right). \tag{22}$$

The term spinmap comes from the cylindrical symmetry of the oriented point basis. Given this basis, the 3D shape of the neighbouring surface points can be reduced to a 2D snapshot of global shape, which is a representation invariant to a rotation about its central axis L. 3D surface point correspondences can then be established on the basis of a *2D correlation* between their respective spin maps. The applicability of the technique to anatomical surfaces is demonstrated with skull and bone data sets. The use of points matched by spinmap correlation to determine the transformation between two surfaces is comparable to feature-point matching, as addressed in Section 4.

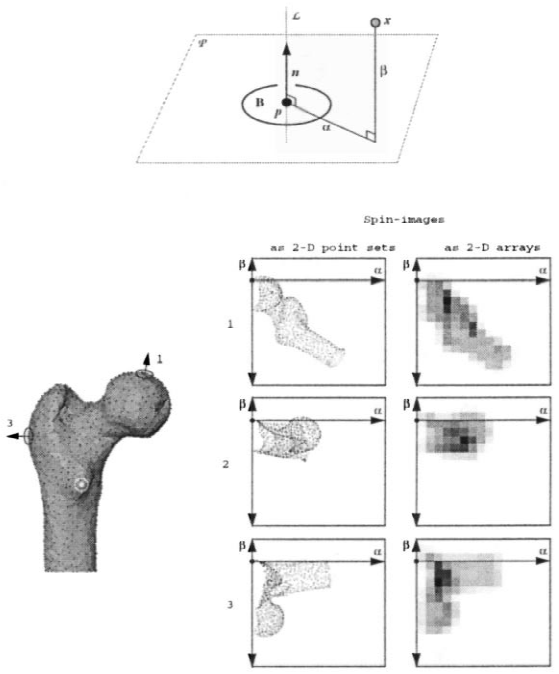

Fig. 5. (a) An illustration spinmap geometry; (b) points distributed on the surface of a femur bone, shown with three point bases and their corresponding spinmaps, appearing as point sets and grey-scale images (courtesy of A.E. Johnson).

Early work on appearance-based techniques mostly deals with face recognition in 2D intensity images, and is based on projecting face images onto a feature space that spans the significant variations among known face images (Turk and Pentland, 1991). Murase and Nayar (1995) apply a similar framework to recognize 3D objects, whose model is stored in a CAD database, from 2D images, as well as estimate the transformation between the two. This in turn is the basis for the eigenshape technique of Campbell and Flynn (1999) for object recognition and pose estimation of range images.

In Murase's application, the basic idea is to take several training views of an object, which rests on a turntable. The object pose is varied by rotating the turntable by an angle θ in increments of a few degrees. Each $N \times N$ image $I(i,j)$ corresponding to a training view is represented as an N^2-vector

$$\hat{\boldsymbol{x}} = [I(1,1)\cdots I(1,N)I(2,1)\cdots I(2,N)\cdots I(N,N)]^T, \tag{23}$$

obtained by concatenating the rows of the image and taking the transpose. An image is then equivalent to a point in a huge N^2-dimensional space. Fortunately, most of the variation in the images can be accounted for by a subspace spanned by just a few vectors, as in **principal component analysis** (or Karhunen-Loeve expansion) (Fukunaga, 1990). Each vector corresponding to a distinct training view constitutes a column of an *object image set*

$$\boldsymbol{X}_{N^2\times m} = [\hat{\boldsymbol{x}}_1, \ldots, \hat{\boldsymbol{x}}_m], \tag{24}$$

where m is the number of training views. A covariance matrix is then defined from the image set:

$$\boldsymbol{Q}_{N^2\times N^2} = \boldsymbol{X}\boldsymbol{X}^T. \tag{25}$$

The eigenvectors $[\boldsymbol{e}_1 \cdots \boldsymbol{e}_k]$ corresponding to the k largest eigenvalues are computed (Murase lists three algorithms for doing this efficiently), where $\lambda_i \boldsymbol{e}_i = \boldsymbol{Q}\boldsymbol{e}_i$, comprising a *parametric eigenspace*.

Thereafter, the set of training views projected onto the parametric eigenspace constitutes samples of a *smoothly varying manifold*. If $k = 3$ this manifold is a closed curve whose arclength is a one-to-one function of θ in the 3D space spanned by $[\boldsymbol{e}_1\boldsymbol{e}_2\boldsymbol{e}_3]$, which in turn can be fitted with cubic splines. Given a new image, its pose can be estimated by projecting it onto the parametric eigenspace. For example, if $k = 3$, this can be achieved by finding the position of its projection within the cube spanned by $[\boldsymbol{e}_1\boldsymbol{e}_2\boldsymbol{e}_3]$) and by finding the closest point on the manifold to this projection. The pose can then be estimated by interpolation along arclength between the two closest discrete poses.

Now, for surface registration (Campbell and Flynn, 1999), we are dealing with more than one pose parameter (i.e. not just θ corresponding to turntable rotation), but the framework is the same (and this could applied to registering an open anatomical surface with a closed surface extracted offline from tomographic data). The steps are (the first three can be implemented offline):

- compute training views of the anatomy at sufficient increments of each pose parameter;
- compute the parametric subspace which captures enough of the shape variation across all poses;
- project back the training views to this subspace to generate a manifold (which may then be smoothed);
- and for a new open surface, project it also onto the parametric subspace and obtain its transformation by interpolating the projections of discrete training poses.

4. Matching, optimization and transformation computation

The third stage comprises the search for corresponding points or feature pairs, based on the surface representation of the second stage, and the computation of the optimal transformation as idealized in the first stage. The search for a match can be either a *succession of comparisons* of discrete candidates, as is frequently the case for feature pairs, or an *iterative minimization* of an objective function, as usually occurs with point- and model-based schemes. The subsequent transformation computation depends on the assumptions made in the first stage. It comprises a closed-form computation or an iterative search for the six or more parameters which best align feature pairs. Alternately, a *Kalman filtering* approach presents advantages to tracking surface points over an extended time sequence, as this estimation of transformation parameters is optimal for the noise characteristics of the whole sequence yet is recursive over time. As for the techniques based on global shape, points whose spinmaps correlate well are used to (over-)determine a transformation in the same manner as discrete feature points (discussed below). The special case of appearance-based matching was fully described in the previous section.

4.1. Discrete feature matching and transformation computation

Feature matching generally determines a rigid transformation, due the relative sparseness of features, and involves a comparison in terms of shape parameters, such as surface curvature values at extrema, curvature and torsion values of curvilinear features, and shape type, averaged properties and size for regions. In the event a particular matching technique does not imbed a transformation computation, most of the rigid-body motion representations reviewed in Section 2.1 lead to a closed-form expression for estimating rotation and translation. Early techniques include **pose clustering** (also known as generalized Hough transform), **sequential hypothesis and test** (also known as prediction and verification and alignment), and **geometric hashing**. We also review the **eigenvector** technique of Shapiro and Brady (1992) for matching point features.

Pose clustering matches like simple or compound structures and derives a transformation from *each* such correspondence. This approach involves quantizing the space of relevant transformations and using it as an accumulator in which each match increments a corresponding cell. A globally acceptable transformation is detected as a cluster in the space of all such candidate transformations (Stockman, 1987).

Sequential hypothesis and test consists of picking a set of feature pairs which are consistent and which determine a transformation, validating these hypothesized matches based on how other features from the two spaces agree with the putative transformation, then possibly backtracking and proceding anew with a new set of consistent feature pairs (Bolles and Horaud, 1986; Chen and Huang, 1988).

Geometric hashing involves precomputing local matching information, which is rotationally and translationally invariant, and storing it in the form of a *hash table*, where each entry is associated with a simple or compound feature to which a local coordinate system (or basis) can be unambiguously assigned, for surface A. Subsequently, a feature or set of features in surface B is similarly characterized. Finally, given the transformation between the two local bases, the consistency of the mapped non-basis features is evaluated, where consistent feature pairs validate this transformation by voting for it (Lamdan and Wolfson, 1988; Kishon et al., 1990).

Shapiro and Brady (1992) match feature points on the basis of consistent same-space distances by an elegant eigenanalysis technique, following the inter-image distance-based matching technique of Scott and Longuet-Higgins (1991). Shapiro suggests accounting for global structure, and proposes that each image or surface be represented by a proximity matrix $\boldsymbol{H}$ of *intra-image* (or intra-surface) distances. The eigenvalues of each matrix are then computed, resulting in a *modal matrix* $\boldsymbol{V}$ whose

columns correspond to the eigenvectors of $\boldsymbol{H}$. Each row of $\boldsymbol{V}$ can be thought of as a feature vector $\boldsymbol{F}_i$ containing the modal coordinates of a feature i. The final stage consists of correlating the two sets of feature vectors $\boldsymbol{F}_{i,A}$ and $\boldsymbol{F}_{j,B}$.

In general, techniques involving some sort of voting or accumulation, such as pose clustering and geometric hashing, are likely to be the most robust, particularly to missing information (such as a partial overlap between different modalities, i.e. registering a sub-patch with a larger surface patch).

4.2. Closest point finding and global optimization

Free-form surface matching involves a closest-point finding stage, which may or may not benefit from some form of preprocessing (k-D trees, distance maps), and frequently a global optimization stage which is used either to compute the best current transformation $\boldsymbol{T}_k$ based on the latest closest-point pairs, or to accelerate the convergence of these fundamentally iterative techniques towards a definitive result.

Classical techniques can be used to implement the minimization of Eq. (13), and generally are of the **unconstrained nonlinear optimization** type (Luenberger, 1984). These approaches are based on the conditions that, for a minimum, the objective function must have a null gradient and a positive semidefinite Hessian matrix. Starting at some initial point, one determines, according to a fixed rule, a direction of movement in the domain, then one moves in that direction to a relative minimum of the objective function along that line. At that point, a new direction is determined and the process is repeated, until some termination condition is met. These techniques include the *method of Steepest Descent*, *multivariate Newton's*, *Conjugate-Gradient* and *Quasi-Newton methods*, and related methods (Luenberger, 1984).

These methods are also used, once point-matches are determined, to iteratively *compute optimal transformation parameters*. Lavallée and Szeliski use a Levenberg-Marquardt technique, which is a hybrid of the Steepest Descent and multivariate Newton's Method (Luenberger, 1984), to iteratively compute both rigid (Lavallée and Szeliski, 1995) and non-rigid (Szeliski and Lavallée, 1996) transformations. Besl and McKay's accelerated ICP technique involves a series of line searches in a seven-parameter space spanned by rotation quaternion and translation (Besl and McKay, 1992). The unaccelerated stage finds closest point pairs and uses a quaternion technique to find the current iteration's least-squares registration. Each such iteration is equivalent to a small step in 7-space, and the accelerated stage fits an interpolant in the direction of this step (from the last 7-vector), whereupon the convergence can be improved by a line search for the minimum. Special switching logic determines whether a linear or parabolic interpolant is used.

The **k-D tree** (Zhang, 1994) is a sequence of bisections in a space of k dimensions (here $k=3$). Each successive cut plane passes through a point chosen such that it divides the remaining points into clouds of roughly equal numbers of points, producing a left and a right son. Each son is split into two grandsons by choosing the appropriate plane parallel to $x-z$, followed by a plane parallel to $x-y$, and so on, while alternating the cut plane orientation, until sufficient divisions occur such that no resulting rectangular parallelopiped contains a data point. Such a cut is termed a ''leaf''. Each node ν of the 3-D tree is characterized by a point $\boldsymbol{x}_{A,i}(\nu)$ through which it passes, and a parameter $t(\nu)$ indicating the orientation of the cut plane. The search for the closest point via a 3-D tree calls a recursive procedure that begins at the root of the 3-D tree assigned to X_A, and exploits the tree structure, branching off into one son or the other depending on the signs of the distance components, to zero-in on the best candidate.

A **distance map** can be precomputed off-line, to determine for each voxel in a volume containing a surface S_B, the closest point on that surface and the distance to it. Thereafter, for a surface S_A falling within this mapped volume, each of its points inherits the closest point precomputed for the voxel on which it falls. In the case of an *octree spline* implementation (Lavallée and Szeliski, 1995), the representation is carried out by a classic octree subdivision (Foley et al., 1990). For each corner vertex of all terminal cubes, the signed distance to its closest point on the surface is computed. The determination of the distance to the closest point on S_B is achieved by interpolation. Alternate techniques for computing the distance map of a surface are by *fast marching level sets* (Sethian, 1996a; Kimmel et al., 1996), where distance is the arrival time of a moving front starting from the initial surface, and by various local mask-based *distance transforms* (Borgefors, 1984, 1986; Paglieroni, 1992).

Alternately, each match pair $\hat{\boldsymbol{x}}_k \equiv (\boldsymbol{x}_{A,k}, \boldsymbol{x}_{B,k})$ can be viewed as a measure $\hat{\boldsymbol{x}}_k$ of a match $\boldsymbol{x}_k$ resulting from the application of the true transformation relating the two surfaces, corrupted by a random error $\boldsymbol{v}_k$. This leads to an **extended Kalman filtering** (EKF) (Kalman, 1960; Sorenson, 1980) formalism for recursively estimating the optimal (with respect to noise statistics) transformation parameters. The application of EKF to tracking surface points over time is an issue of expressing a relationship between positions or feature vectors of the matched pairs and the transformation parameters as a *measurement equation* that can be linearized, and for which the linear Kalman filter can recursively compute the best transformation estimate. This is the technique used by Feldmar and Ayache (1994c) for estimating the global affine transformation relating matched feature vector pairs. Because it is recursive, it presents advantages for tracking surface points over time. Only the new measurements and the statistics and transformation parameters of the previous iteration need to be considered, in computing the parameters which are optimal for the noise statistics of the whole sequence.

4.3. Model-based motion estimation

The method of estimating motion from surface models depends on the nature of the model. As seen in the discussion of Section 3.3, many physically-based models segment the surface explicitly, but surface motion is often just the difference between consecutive shapes at a given surface coordinate (u, v), particularly when the shape at time t_k is used to initialize the shape estimation at t_{k+1} (Terzopoulos et al., 1988; McInerney and Terzopoulos, 1995). The application of surface evolution models to motion estimation is still in its infancy, but as emphasized by Audette and Peters (1999), this framework can be exploited not only to extract a surface reliably, but also to alleviate an explicit search for matches in a subsequent point-based registration scheme.

Finite element modelling is the most prevalent technique for computing physically based models (Essa et al., 1993; Terzopoulos and Metaxas, 1991). Here, the displacement FEM represents a surface as a mesh of 2D simple polygonal elements whose global behaviour (where a surface reaches equilibrium between internal and external forces) can be characterized by an element-by-element analysis of the dynamics involved. The global mass, damping and stiffness matrices are summations of the corresponding element matrices, as are non-concentrated load matrices. In practice, the integration of expression (18) is simplified by neglecting either mass (Terzopoulos and Metaxas, 1991), leading to the *explicit first-order Euler integration*[5] $\boldsymbol{U}^{(t+\delta t)} = \boldsymbol{U}^{(t)} + \delta t\left(\boldsymbol{C}^{(t)}\right)^{-1}\left(\boldsymbol{R}^{(t)} - \boldsymbol{K}\boldsymbol{U}^{(t)}\right)$, or damping, leading to the change of basis to modal coordinates and the expression (19), which can also be integrated numerically.

The Amini surface tracking technique finds closest points in a manner similar to free-form surface registration techniques (Amini and Duncan, 1992). For each point of surface A, the point on surface B that minimizes expression (21) is selected. This ensures that the match is most similar in local surface shape and that its relative movement is nearest to being conformal. The result is a 3D flow field, which may be noisy, defined over the first surface. Thereafter, a *vector smoothing* technique (Horn and Schunck, 1981) is applied to components of each flow vector, and flow estimates are propagated over other regions of the surface.

Other formulations have also been proposed for computing or tracking deformable model points. The change of basis of expression (19) allows Sclaroff and Pentland (1993) to adapt Shapiro's *modal matching* technique to explicitly match FEM surface nodes. Pentland and Horowitz (1991), Metaxas and Terzopoulos (1993) and Terzopoulos and Szeliski (1992) also apply a *Kalman filter* formalism to estimate 3D motion from a physical model. Cohen (1991) suggests a *finite-differences* based technique for finding the energy-minimizing contour of his balloon model. Moreover, surface evolution models are implemented with a combination of central and *upwind* (Osher and Sethian, 1988; Levesque, 1992) finite-differences (Malladi et al., 1995; Kimmel et al., 1997). For alternative approaches to estimating deformable contours and surfaces, see (Terzopoulos and Szeliski, 1992; McInerney and Terzopoulos, 1996).

[5]Based on virtual time steps which need not related to the time lapse separating two surfaces.

5. Conclusion

This paper has presented a survey of surface registration techniques, particularly those which apply to anatomical surfaces, with an emphasis on their mathematical or algorithmic foundations. We chose to represent the process of registration as the succession of three stages: choice of transformation representation, choice of surface representation and similarity criterion, and matching per se and global optimization. According to Section 2, transformations can be categorized as rigid and non-rigid. Several different representations for rigid-body transformations have been surveyed, while non-rigid transformations can be further categorized into global and local polynomial representations, and according to the choice of coordinates (cartesian or otherwise), which is best suited for the problem.

We saw in Section 3 that a feature, point, global shape or model-based approach can be employed to represent a surface for the purpose of registration, and that each has certain advantages. Feature-based and global shape-based techniques can determine an arbitrarily large rigid transformation. A point-based technique is best suited for bringing two surfaces into very close alignment, given a good starting point for the final transformation, particularly if we wish to quantify non-rigid motion, while a model-based approach is advantageous if identifying the surface within a volume must be accomplished prior to registration. Furthermore, we discussed three categories of features: sparse points, curves and regions. We demonstrated that dense point-based techniques vary according to their definition of proximity between points and according to their means of estimating the closest point to a surface. Finally, we pointed out that models for identifying and possibly registering surfaces are either based on physical or surface evolution equations. We identified which techniques compute surface motion implicitly, from the evolution over time of a 3D model, and which explicitly establish correspondences between surface points.

Next, in Section 4 we offered a survey of the main numerical estimation schemes or algorithms for matching features (and points correlated on the basis of spinmaps),

points and model coordinates (matching on the basis of appearance was concluded in the discussion on representation). Features are typically matched by discrete comparisons, in a manner which may imbed the computation of the resulting transformation in the matching process. Since the transformation specified by these matches is usually rigid-body and the robustness of each match is typically very good (due to the discriminating ability of features), the global registration parameters can then be obtained by a closed-form expression. In contrast, matching and registration for point-based techniques are imbedded into a highly iterative process, with the matching characterized either by a classical optimization technique or by a more efficient procedure requiring offline preprocessing, such as a k-D tree or distance map. The registration may also require an optimization or Kalman filtering procedure for finding the best-fitting transformation parameters in the global sense. Lastly, for physical model-based approaches, surface motion computation frequently amounts to identifying the surface over consecutive volumetric scans. This identification is usually implemented with the finite-element model, where a surface at t_1 initializes that at t_2. In contrast, surface evolution models are less prevalent in registration, and typically estimated by finite differences.

In conclusion, the design choices, namely how to model the transformation, whether to use a feature, point, global shape or model-based technique (as well as how to preprocess the data) and how to optimally compute the transformation, are very context-specific. There are instances, e.g. when dealing with rigid tissue, when a rigid-body transformation will suffice, and others, such as tracking soft tissue deformations, where a spline-based non-rigid transformation model is necessary. In between these two poles, there is a whole spectrum of situations where a rigid model may be too constrictive, but an approach based on splines may also be inapplicable, by being underdetermined or inappropriate for a given noise distribution, for example. Likewise, the choice of surface representation is influenced by the size of the transformation (e.g. whether surfaces are already roughly aligned or not), by the smoothness, noise distribution, and mutual overlap of the data, as well as by time constraints. Moreover, establishing correspondences between two sets of primitives is basically an outgrowth of the formalization of the transformation and the choice of primitives, and consequently is also highly related to the situation at hand.

Acknowledgements

The authors wish to thank E.A. Johnson, S. Lavallée, R. Szeliski, L. Brunie, G. Soucy and R.A. Volz, for contributing images and for granting permission to duplicate illustrations for this paper, as well as K. Siddiqi, G. Dudek, S. Zucker and R. Kimmel for stimulating discussion.

References

Abdelmalek, N., 1990. Algebraic error analysis for surface curvatures of 3-D range images obtained by different methods. In: IEEE Proc. ICPR, pp. 529–534.

Alvarez, L., Lions, P.L., Morel, J.-M., 1992. Image selective smoothing and edge detection by nonlinear diffusion. II. SIAM J. Numer. Anal. 29 (3), 845–866.

Amini, A.A., Duncan, J.S., 1992. Differential geometry for characterizing 3D shape change. Proc. SPIE Math. Methods in Med. Imag. 1768, 170–181.

Arun, K.S., Huang, T.S., Blostein, S.D., 1987. Least-squares fitting of two 3-D point sets. IEEE Trans. PAMI 9 (5), 698–700.

Audette, M.A., Peters, T.M., 1999. Level-set surface segmentation and registration for computing intrasurgical deformations. Proc. SPIE 3661 Medical Imaging, to appear.

Ayache, N., 1991. Artificial Vision for Mobile Robots: Stereo Vision and Multisensory Perception. MIT Press, Cambridge, MA.

Bacjsy, R., Kovacic, S., 1989. Multiresolution elastic matching. CVGIP 46, 1–21.

Bajcsy, R., Solina, F., 1987. Three dimensional object representation revisited. In: Proc. IEEE Int. Conf. Comp. Vision. pp. 231–240.

Barr, A.H., 1981. Superquadrics and angle-preserving transformations. IEEE Comp. Graph. Appl. 1, 11–23.

Barr, A.H., 1984. Global and local deformations of solid primitives. Comp. Graph. 18 (3), 21–30.

Bathe, K.-J., 1982. Finite Element Procedures in Engineering Analysis. Prentice-Hall, Englewood Cliffs, NJ.

Besl, P.J., Jain, R.C., 1986. Invariant surface characteristics for 3D object recognition in range images. CVGIP 33, 33–80.

Besl, P.J., McKay, N.D., 1992. A method for registration of 3-D shapes. IEEE Trans. PAMI 14 (2), 239–256.

Bolles, R.C., Horaud, P., 1986. 3DPO: a three-dimensional part orientation system. Int. J. Rob. Res. 5 (3), 3–26.

Bookstein, F.L., 1989. Principal warps: thin-plate splines and the decomposition of deformations. IEEE Trans. PAMI 11 (6), 567–585.

Borgefors, G., 1984. Distance transformations in arbitrary dimensions. CVGIP 27, 321–345.

Borgefors, G., 1986. Distance transformations in digital images. CVGIP 34, 344–371.

Brown, L.G., 1992. A survey of image registration techniques. ACM Comp. Surveys 24 (4), 326–376.

Campbell, R.J., Flynn, P.J., 1999. Eigenshapes for 3D object recognition in range data. In: Proc. IEEE Conf. Computer Vision and Pattern Recognition.

Caselles, V. et al., 1992. A geometric model for active contours in image processing. CEREMADE report no. 9210.

Caselles, V., Kimmel, R., Sapiro, G., 1995. Geodesic active contours. In: IEEE Proc. Int. Conf. Comp. Vision. pp. 694–699.

Caselles, V. et al., 1997. Minimal surfaces: a three dimensional segmentation approach. Numer. Math. 77 (4), 423–451.

Chen, C.W., Huang, T.S., Arrott, M., 1994. Modeling analysis, and visualization of left ventricle shape and motion by hierarchical decomposition. IEEE Trans. PAMI 16 (4), 342–356.

Chen, H.H., Huang, T.S., 1988. Maximal matching of 3-D points for multiple-object motion estimation. Patt. Rec. 21 (2), 75–90.

Chen, Y., Medioni, G., 1991. Object modelling by registration of multiple range images. In: IEEE Proc. Conf. Rob. and Auto, pp. 2724–2729.

Cohen, L.D., 1991. On active contour models and balloons. CVGIP: Image Under. 53 (2), 211–218.

Cohen, L.D., Cohen, I., 1990. A finite element method applied to new active contour models and 3D reconstruction from cross sections. In: IEEE Proc. Int. Conf. Comp. Vision. pp. 587–591.

Cohen, I., Ayache, N., Sulger, P., 1992a. Tracking points on deformable objects using curvature information. INRIA Tech. Report No. 1595.

Cohen, I., Cohen, L.D., Ayache, N., 1992b. Using deformable surfaces to

segment 3D images and infer differential structures. CVGIP: Image Under. 56 (2), 242–263.
Coppini, G. Demi, M., d'Urso, G., L'Abbate, A., Valli, G., 1987. Tensor description of 3D time varying surfaces using scattered landmarks: an application to heart motion. In: Cappellini, V. (Ed.), Time-Varying Image Processing and Moving Object Recognition. North-Holland, Amsterdam, pp. 158–163.
Craig, J.J., 1989. Introduction to Robotics: Mechanics and Control. Addison-Wesley, Reading, MA.
Collignon, A., Vandermeulen, D., Suetens, P., Marchal, G., 1993. Surface based registration of 3D medical images. Proc. SPIE 1898: Med. Imag., pp. 32–42.
DeCarlo, D., Metaxas, D., 1996. Blended deformable models. IEEE Trans. PAMI 18 (4), 443–448.
Deriche, R., 1987. Using Canny's criteria to derive a recursively implemented optimal edge detector. Int. J. Comp. Vision 1 (2), 167–187.
Deriche, R., 1990. Fast algorithms for low-level vision. IEEE Trans. PAMI 12 (1), 78–87.
Dierckx, P., 1995. Curve and Surface Fitting with Splines. Clarendon, Oxford.
do Carmo, M.P., 1976. Differential Geometry of Curves and Surfaces. Prentice-Hall, Englewood Cliffs, NJ.
Duchon, J., 1976. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. R.A.I.R.O Analyse Numérique 10, 5–12.
Essa I., Sclaroff, S., Pentland, A., 1993. Physically based modelling for graphics and vision. In: Martin, R. (Ed.), Directions in Geometric Computing. Information Geometers, Winchester, UK, pp. 161–218.
Evans, A.C., Dai, W., Collins, L., Neelin, P., Marrett, S., 1991. Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis. Proc. SPIE 1445 Med. Imag. V: Image Proc., 236–246.
Fan, T.J., Medioni, G., Nevatia, R., 1989. Recognizing 3-D objects using surface descriptions. IEEE Trans. PAMI 11 (11), 1140–1157.
Faugeras, O.D., Hebert, M., 1986. The representation, recognition, and locating of 3-D objects. Int. J. Rob. Res. 5 (3), 27–52.
Feldmar, J., Ayache, N., 1994a. Rigid and affine registration of free-form surfaces, using differential properties. In: Proc. Euro. Conf. Comp. Vision. pp. 397–406.
Feldmar, J., Ayache, N., 1994b. Locally affine registration of free-form surfaces. In: IEEE Proc. Conf. Comp. Vision and Patt. Rec. pp. 496–501.
Feldmar, J., Ayache, N., 1994c. Rigid, affine and locally affine registration of smooth surfaces. INRIA Technical Report No. 2220.
Ferrie, F.P., Mathur, S., Soucy, G., 1993. Feature extraction for 3-D model building and object recognition. In: Jain, A.K., Flynn, P.J. (Eds.), Three-Dimensional Object Recognition Systems. Elsevier, Amsterdam, pp. 57–88.
Flynn, P.J., Jain, A.K., 1989. On reliable curvature estimation. In: IEEE Proc. Conf. Comp. Vision and Patt. Rec. pp. 110–116.
Foley, J.D., van Dam, A., Feiner, S.K., Hughes, J.F., 1990. Computer Graphics – Principles and Practice. Addison-Wesley, Reading, MA.
Fukunaga, K., 1990. Introduction to Statistical Pattern Recognition, 2nd Edition. Academic Press, Boston.
Goldgof, D.B., Lee, H., Huang, T.S., 1988a. Motion analysis of nonrigid surfaces. In: Proc. Conf. Comp. Vision and Patt. Rec. pp. 375–380.
Goldgof, D.B., Lee, H., Huang, T.S., 1988b. Feature extraction and terrain matching. In: Proc. Conf. Comp. Vision and Patt. Rec. pp. 899–904.
Guéziec, A., Ayache, N., 1994. Smoothing and matching of 3-D space curves. Int. J. Comp. Vision 12 (1), 79–104.
Hebert, M., Caillas, C., Krotkov, E., Kweon, I.S., Kanade, T., 1989. Terrain mapping for a roving planetary explorer. IEEE Proc. Int. Conf. Rob. and Auto. pp. 997–1002.
Henri, C.J., Collins, D.L., Peters, T.M., 1991. Multimodality image integration for stereotactic surgical planning. Med. Phys. 18 (2), 167–177.
Herman, G.T., Liu, H.K., 1979. Three-dimensional display of human organs from computed tomograms. Comp. Graph. and Image Proc. 9 (1), 1–21.
Horn, B.K.P., 1987. Closed-form solution of absolute orientation using unit quaternions. J. Opt. Soc. Am. A 4 (4), 629–642.
Horn, B.K.P., Schunck, B.G., 1981. Determining optical flow. Art. Intel. 17, 185–203.
Huang, T.S., Blostein, S.D., Margerum, E.A., 1986. Least-squares estimation of motion parameters from 3-D point correspondences. In: IEEE Proc. Conf. Comp. Vision and Patt. Rec. pp. 198–201.
Jacq, J.J., Roux, C., 1993. Automatic registration of 3D images using a simple genetic algorithm with a stochastic performance function. In: IEEE Proc. Eng. Med. and Biol. Soc. pp. 126–127.
Johnson, A.E., Hebert, M., 1998. Surface matching for object recognition in complex three-dimensional scenes. Image and Vision Computing 16, 635–651.
Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. Trans. AMSE J. Basic Eng. 820, 35–45.
Kamgar-Parsi, B., Jones, J.L., Rosenfeld, A., 1991. Registration of multiple overlapping range images: scenes without distinctive features. IEEE Trans. PAMI 13 (9), 857–871.
Kass, M., Witkin, A., Terzopoulos, D., 1987. Snakes: active contour models. Int. J. Comp. Vision 1 (4), 321–331.
Kehtarnavaz, N., Mohan, S., 1989. A framework for estimation of motion parameters from range images. CVGIP 45, 88–105.
Kichenassamy, S., Kumar, A., Olver, P., Tonnenbaum, A., Yezzi, A., 1995. Gradient flows and geometric active contour models. In: Proc. IEEE Int. Conf. Comp. Vision. pp. 810–815.
Kikinis, R., Gleason, P.L., Lorensen, W., Wells, W., Grimson, W.E.L., Lozanc-Perez, T., Ettinger, G., White, S., Jolesz, F., 1994. Image guidance techniques for neurosurgery. In: Proc. SPIE 2359: Visualization in Biomedical Computing. pp. 537–540.
Kimia, B.B., Siddiqi, K., 1996. Geometric heat equation and nonlinear diffusion of shapes and images. Computer Vision Image Understanding 64 (3), 305–322.
Kimmel, R., Kiryati, N., Bruckstein, A.M., 1996. Distance maps and weighted distance transforms. J. Math. Imag. Vision 6, 223–233.
Kimmel, R., Kiryati, N., Bruckstein, A.M., 1997. Analyzing and synthesizing images by evolving curves with the Osher-Sethian method. Int. J. Comp. Vision 24 (1), 37–56.
Kishon, E., Hastie, T., Wolfson, H., 1990. 3-D curve matching using splines. In: Proc. Euro. Conf. Comp. Vision. pp. 589–591.
Lamdan, Y., Wolfson, H.J., 1988. Geometric hashing: a general and efficient model-based recognition scheme. In: IEEE Proc. Int. Conf. Comp. Vision. pp. 238–249.
Lavallée, S., 1996. Registration for computer-integrated surgery: methodology, state of the art. In: Taylor, R.H., Lavallee, S., Burdea, G.C., Mosges, R. (Eds.), Computer-Integrated Surgery – Technology and Clinical Applications. MIT Press, Cambridge, MA, pp. 77–97.
Lavallée, S., Szeliski, R., 1995. Recovering the position and orientation of free-form objects from image contours using 3D distance maps. IEEE Trans. PAMI 17 (4), 378–390.
Lavallée, S., Szeliski, R., Brunie, L., 1996. Anatomy-based registration of three-dimensional medical images, range images, X-ray projections, and three-dimensional models using octree-splines. In: Taylor, R.H., Lavallee, S., Burdea, G.C., Mosges, R. (Eds.), Computer-Integrated Surgery – Technology and Clinical Applications. MIT Press, Cambridge, MA, pp. 115–143.
Levesque, R.J., 1992. Numerical Methods for Conservation Laws, 2nd Edition. Birkhäuser Verlag, Basel.
Lin, Z.C., Lee, H., Huang, T.S., 1986. A frequency-domain algorithm for determining motion of a rigid object from range data without correspondences. In: IEEE Proc. Conf. Comp. Vision and Patt. Rec., pp. 194–198.
Luenberger, D.G., 1984. Linear and Nonlinear Programming, 2nd Edition. Addison-Wesley, Reading, MA.
Maintz, J.B.A., Viergever, M.A., 1998. A survey of medical image registration. Med. Image Anal. 2 (1), 1–36.

Maintz, J.B.A., van den Elsen, P.A., Viergever, M.A., 1996. Evaluation of ridge seeking operators for multimodality medical image matching. IEEE Trans. PAMI 18 (4), 353–365.

Malladi, R., Sethian, J.A., 1998. A real-time algorithm for medical shape recovery. In: Proc. IEEE Int. Conf. Comp. Vision. pp. 304–310.

Malladi, R., Sethian, J.A., Vemuri, B.C., 1995. Shape modeling with front propagation: a level set approach. IEEE Trans. PAMI 17 (2), 158–175.

Malvern, L.E., 1969. Introduction to the Mechanics of a Continuous Medium. Prentice-Hall, Englewood Cliffs, NJ.

McInerney, T., Terzopoulos, D., 1995. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4D image analysis. J. Comp. Med. Imag. and Graph. 19 (1), 69–83.

McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. Med. Image Anal. 1 (2), 91–108.

Metaxas, D., Kakadiaris, I.A., 1996. Elastically adaptive deformable models. In: Proc. Euro. Conf. Comp. Vision. pp. 550–559.

Metaxas, D., Terzopoulos, D., 1993. Shape and nonrigid motion estimation through physics-based synthesis. IEEE Trans. PAMI 15 (6), 580–591.

Mishra, S.K., Goldgof, D.B., Huang, T.S., 1991. Motion analysis and epicardial deformation estimation from angiography data. In: IEEE Proc. Conf. Comp. Vision and Patt. Rec. pp. 331–336.

Monga, O., Benayoun, S., 1995. Using partial derivatives of 3D images to extract typical surface features. Comp. Vision and Image Under. 61 (2), 171–189.

Monga, O., Deriche, R., Rocchisiani, J.-M., 1991. 3D edge detection using recursive filtering: application to scanner images. CVGIP: Image Under. 53 (1), 76–87.

Monga, O., Benayoun, S., Faugeras, O.D., 1992. From partial derivatives of 3D density images to ridge lines. In: IEEE Proc. Conf. Comp. Vision and Patt. Rec. pp. 354–359.

Murase, H., Nayar, S.K., 1995. Visual learning and recognition of 3-D objects from appearance. Int. J. Computer Vision 14, 5–24.

Nastar, C., Ayache, N., 1993. Fast segmentation, tracking, and analysis of deformable objects. In: IEEE Proc. Int. Conf. Comp. Vision. pp. 275–279.

Osher, S., Sethian, J.A., 1988. Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. J. Comp. Phys. 79, 12–49.

Paglieroni, D.W., 1992. Distance transforms: properties and machine vision applications. CVGIP: Graph. Models and Image Proc. 54 (1), 56–74.

Pelizzari, C.A. et al., 1989. Accurate three-dimensional registration of CT, PET, and/or MR images of the brain. J. Comp. Ass. Tomog. 13 (1), 20–26.

Pentland, A., Horowitz, B., 1991. Recovery of nonrigid motion and structure. IEEE Trans. PAMI 13 (7), 730–742.

Pentland, A., Sclaroff, S., 1993. Closed-form solutions for physically based shape modelling and recognition. IEEE Trans. PAMI 13 (7), 715–729.

Perona, P., Malik, J., 1990. Scale-space and edge detection using anisotropic diffusion. IEEE Trans. PAMI 12 (7), 629–639.

Press, W.H., Chui, H., Mjolsness, E., Pappin, S., Daviachi, L., Goldman-Rakic, P., Duncan, J., 1992. Numerical Recipes in C: The Art of Scientific Computing, 2nd Edition. Cambridge University Press, Cambridge.

Proakis, J.G., 1996. Digital Signal Processing: Principles, Algorithms, and Applications. Prentice Hall, Upper Saddle River, NJ.

Radack, G.M., Badler, N.I., 1989. Local matching of surfaces using a boundary-centered radial decomposition. CVGIP 45, 380–396.

Radig, B., 1984. Image sequence analysis using relational structures. Patt. Rec. 17 (1), 161–167.

Rangarajan, Flannery, B.P., Teukolsky, S.A., Vetterling, W.T., 1997. A robust point-matching algorithm for autoradiograph alignment. Medical Image Analysis 1(4), 379–398.

Roth, G., Levine, M.D., 1993. Extracting geometric primitives. CVGIP: Image Understanding 48 (1), 1–22.

Sander, P.T., Zucker, S.W., 1990. Inferring surface trace and differential structure from 3-D images. IEEE Trans. PAMI 12 (9), 833–854.

Sclaroff, S., Pentland, A., 1993. Modal matching for correspondence and recognition. M.I.T. Media Lab. Perceptual Computing Section Tech. Rep. No. 201.

Scott, G.L., Longuet-Higgins, H.C., 1991. An algorithm for associating the features of two images. Proc. Royal Soc. London B 244, 21–26.

Sethian, J.A., 1996a. A fast marching level set method for monotonically advancing fronts. Proc. Natl. Acad. Sci. USA 93, 1591–1595.

Sethian, J.A., 1996b. A review of the theory, algorithms, and applications of level set methods for propagating interfaces. Acta Numerica 5, 309–395.

Shapiro, L.S., Brady, J.M., 1992. Feature-based correspondence: an eigenvector approach. Image Vision Comput. 10 (5), 283–288.

Shen, J., Castan, S., 1986. An optimal linear operator for edge detection. In: IEEE Proc. Conf. Comp. Vision and Patt. Rec. pp. 109–114.

Simon, D.A., Hebert, M., Kanade, T., 1994a. Real-time 3-D pose estimation using a high-speed range sensor. In: Proc. IEEE Int. Conf. Rob. and Auto. pp. 2235–2241.

Simon, D.A., Hebert, M., Kanade, T., 1994b. Techniques for fast and accurate intra-surgical registration. In: Proc. Int. Symp. Med. Rob. and Comp. Assis. Surg. pp. 90–97.

Sorenson, H.W., 1980. Parameter Estimation: Principles and Problems. Marcel Dekker, New York.

Soucy, G., Ferrie, F.P., 1997. Surface recovery from range images using curvature and motion consistency. Comp. Vision and Image Under. 65 (1), 1–18.

Stockman, G., 1987. Object recognition and localization via pose clustering. CVGIP 40, 361–387.

Subsol, G., Thirion, J.Ph., Ayache, N., 1994. Non rigid registration for building 3D anatomical atlases. In: IEEE Proc. IEEE Int. Conf. Patt. Rec. pp. 576–578.

Szeliski, R., Lavallée, S., 1996. Matching 3-D anatomical surfaces with non-rigid deformations using octree-splines. Int. J. Comp. Vision 18 (2), 171–186.

Terzopoulos, D., Metaxas, D., 1991. Dynamic 3D models with local and global deformations: deformable superquadrics. IEEE Trans. PAMI 13 (7), 703–714.

Terzopoulos, D., Szeliski, R., 1992. Tracking with Kalman snakes. In: Blake, A., Yuille, A. (Eds.), Active Vision. MIT Press, Cambridge, MA, pp. 3–20.

Terzopoulos, D., Witkin, A., Kass, M., 1988. Constraints on deformable models: recovering 3D shape and nonrigid motion. Artificial Intelligence 36, 91–123.

Thirion, J.-Ph., 1994. Extremal points: definition and application for 3D image registration. In: IEEE Proc. Conf. Comp. Vision and Patt. Rec. pp. 587–592.

Toriwaki, F.I., Yokoi, S., 1988. Voronoi and related neighbors on digitized two-dimensional space with applications to texture analysis. In: Toussaint, G.T. (Ed.), Computational Morphology. North-Holland, Amsterdam.

Turk, M., Pentland, A., 1991. Eigenfaces for recognition. J. Cognitive Neuroscience 3 (1), 71–86.

Udupa, J.K., 1982. Interactive segmentation and boundary surface formation for 3-D digital images. Comp. Graph. and Image Proc. 18 (3), 213–235.

Unser, M., Aldroubi, A., Eden, M., 1993a. B-spline signal processing: Part I – theory. IEEE Trans. Signal Proc. 41 (2), 821–833.

Unser, M., Aldroubi, A., Eden, M., 1993b. B-spline signal processing: Part II – efficient design and applications. IEEE Trans. Signal Proc. 41 (2), 834–848.

Walker, M.W., Shao, L., Volz, R.A., 1991. Estimating 3-D location parameters using dual number quaternions. CVGIP: Image Understanding 54 (3), 358–367.

Zhang, Z., 1994. Iterative point matching for registration of free-form curves and surfaces. Int. J. Comp. Vision 13 (2), 119–152.