J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*
Copyright © 1999 John Wiley & Sons, Inc.

MEDICAL IMAGE PROCESSING

Medical image processing is a broad and general term used to refer to the application of image processing and analysis techniques to medical data [X rays, computer tomography (*CT*), nuclear medicine, and ultrasound] for visualization or quantification purposes, or both. It is a relatively new field that is evolving in response to the increasing widespread availability of good medical image acquisition systems. Clinicians can now "see inside" the human body. Medical image processing deals with making measurements on these (usually digital) pictures and interpreting the measurements in a clinically meaningful way. Medical image processing now plays a regular part in diagnosis and treatment in many areas of medicine, including breast cancer, cardiac disease, diseases of the brain, and orthopedics, and for automating tasks such as image-guided surgery and surgery planning.

Early work in this area mainly involved the direct application to medical data of two-dimensional (2-D) image-processing algorithms developed for analyzing 2-D static visual images. The limitations of this approach were soon realized, as noise and signal properties of medical images can be quite different from those of visual images. As a result, research is moving towards developing approaches that make use of knowledge of anatomical geometry and physical properties of tissues to interpret medical imagery. Another recent trend is the increased use of three-dimensional (3-D, volumetric) images and image sequences in the medical domain. This presents new challenges in terms of visualization needs and processing speed.

This article is divided into three main sections. The following section consers specific aspects of data processing that are relevant to processing medical images for visualization and display, reconstructing an object from multiple image slices, and medical image enhancement and segmentation. The section after considers the tasks of interpreting and analyzing medical images. Medical image registration (bringing two data sets into spatial alignment), medical shape representations and analysis, and medical image tracking and analysis are discussed. Some medical image application areas are outlined in the final section.

Data Processing

Visualization and Display. Image processing is the focus of this entry. However, one of the most important roles it plays is to preprocess data so that information of interest can be visualized in a clinical meaningful way. This is particularly important for 3-D and 3-D + time data. Here we restrict the discussion to visualizing volumetric data. For an illustration of medical visualization in practice, see, for example, 1.

One of the simplest ways to visualize a 3-D data set is using planar reformatting. This involves treating the scan data as a volume and mathematically defining a cutting plane through the volume. The intensity of the voxels traversing this plane is mapped as pixels into a 2-D image (the slice plane or cross-sectional image). This technique is particularly useful for visualizing anatomy in directions that are more intuitive than those provided by the data acquisition process. Another method of visualizing 3-D data is using a technique known as maximum-intensity projection (*MIP*), which has proved especially popular for visualizing magnetic resonance angiography (*MRA*) data (2). An example is shown in Fig. 1. (3). The main idea is to create an "X-ray"

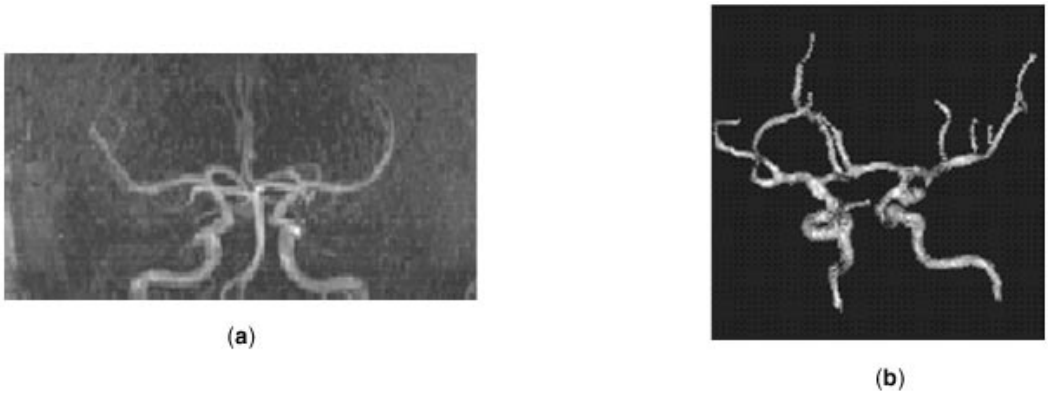

Fig. 1. (a) a maximum-intensity projection (*MIP*) and (b) a surface rendering of an MRA volume of a brain aneurysm (3). Note how the surface visualization gives a better indication of three-dimensionality by the use of shading.

projection image of the 3-D data set, that is, to project a ray from a point source (usually positioned at infinity so that parallel projection geometry is assumed) through the volume. The maximum intensity of all voxels encountered in this traversal is displayed in the resulting 2-D image. MIP images are typically generated from a number of viewing angles about a rotary axis, and the resulting image sequence displayed as a movie. The key advantage of this method is its computational simplicity. Its key disadvantage is the loss of 3-D information in the projection process; also, because only the maximum intensity is retained, in some cases low-intensity structures are not seen. Variants of the standard projection technique have been proposed to deal with this, including depth-sensitive MIP and multifeature-extraction ray tracing (3,4).

The two main techniques used for visualizing medical data in three dimensions are surface rendering and volume rendering (5,6). These methods use spatial information about objects in three dimensions, shadowing, and perspective to create a 3-D impression of the data volume (a *displayable*), which is projected on a 2-D screen.

In *surface rendering* a surface is first extracted from a data set, typically using an intensity threshold operation (isovalue selection), reconstructed as a mathematical model, and then rendered using standard computer graphics techniques such as raycasting (7). One of the classic methods of reconstruction, called the *marching cubes algorithm*, converts the thresholded data set into polygonal surface patches (8). Much of the original information content of the image volume is thus lost, leaving a clear representation of the surface in question. This makes surface rendering especially useful for displaying anatomical data such as CT images of bone. The main advantages of surface rendering include the relatively rapid speed of rendering and reduced storage requirements once the surface has been determined. The disadvantages include that for some data types (e.g., ultrasound data) it may prove difficult to segment a good surface, and this approach can handle small features incorrectly (9,10).

In *volume rendering*, shading and perspective are determined by spatial position and the relative intensity difference in the original data. This is an important difference between surface and volume rendering. In the former case you effectively create a mathematical model of the surface in terms of geometric primitives such as points, lines, and polygons and render the object assuming that it is opaque (no light is transmitted through it). The interior of the object is not described. Volume rendering assumes the object is transparent. It produces an image without generating an intermediate geometric description. Clearly, from the medical visualization perspective this is more powerful, because you can look inside an object and visually inspect different tissue regions. Thus it is typically used for displaying soft tissue structures such as cerebral tumors or other soft tissue lesions. However, introducing transparency makes basic ray-casting (a fundamental step

in all rendering techniques) computationally impractical unless some sort of acceleration method is used to improve performance. This is because the entire 3-D data set must be used for each new viewpoint. This has limited the practical application of volume rendering until quite recently.

The surface and volume rendering techniques described above all aim to give the appearance of three-dimensionality in a planar image (a 2-D display device). True 3-D display or stereo rendering can be achieved by exploiting binocular parallax (11). There are a variety of ways to do this, including generating lenticular photographic prints, using a polarized filter projection system (as used in 3-D movie presentations), and using liquid-crystal shutter glasses or virtual reality (*VR*) heads-up displays (as used for viewing computer-generated images or video games) (5). The basic idea is to generate two images of a scene taken from slightly different viewpoints and to use special optical processing, equipment, or glasses to create a stereoscopic effect, similar to the way in which the human visual system perceives 3-D information. This can be particularly useful to determine the relative depth of close objects in a scene (e.g., vascular structures) and for visualizing textured or translucent objects, such as tumors, which do not have regular shapes or clearly defined boundaries and which are not readily seen in 2-D slice visualization (12).

The advantages of using true 3-D display over using surface or volume rendering in medical applications are strongly dependent on the application and are much debated. For instance, head displays have the advantage of being portable and interactive, but are inconvenient for a surgeon to use in an operating room. For diagnostic needs, clearly, bulkiness is not an issue, but then you might prefer a method that could be readily used for documentation (such as lenticular photographic prints). Limiting factors in the end are cost and image quality.

Reconstruction. Here we define the term *reconstruction* to mean postprocessing of two or more 2-D images acquired from a medical imaging system to generate a 3-D representation of an object or volume in space. Thus we do not consider, for example, how to construct a magnetic resonance image from its k -space representation, or a computed tomography image from radiographs (X-ray projections). See the separate entries on individual medical imaging modalities for discussions on these methods.

With this definition, reconstruction methods can be classified as volumetric-based versus sparse-view-based, or alternatively intensity-based versus feature-based. We consider the former classification.

Volumetric methods involve using a dense set of 2-D scans to reconstruct an object volume. The simplest example of this type of approach involves interpolating between a series of parallel cross-sectional MR or X-ray CT scans. A common variant of this idea is called *spatial compounding*. It involves mapping into a volume the intensities of pixels in a 2-D slice that is located at a known position and orientation in 3-D space. When two (or more) 2-D slices intersect at a voxel location, the intensity is taken as the average intensity of the two values. Spatial compounding is used in 3-D freehand ultrasound imaging, where a number of 2-D B scans (ultrasonic reflection images) are acquired in rapid succession (a sweep), possibly from different look directions (acquisition angles), and processed to give a visualisation of 3-D anatomy (13,14). By imaging from a variety of look directions the effects of noise (speckle) and imaging artifacts can be significantly reduced. However, spatial compounding requires very accurate slice registration, which has proved difficult to achieve with ultrasound data because localizing image features accurately in noisy ultrasound data is difficult. This has limited the clinical application of 3-D freehand ultrasound to date. It may prove that recent advances in multimodality image registration (see "Registration and Matching" in the next section) will provide a solution to this problem in the near future (15).

Sparse-view methods differ from volumetric methods in that they use a small number of 2-D images to reconstruct a 3-D object or scene and the images are typically acquired at a large angular separation. This is the classical stereo vision setup, and many of the methods used in sparse view medical object reconstruction have been borrowed from the computer vision and photogrammetric literature. Sparse view reconstruction has proved especially popular in 3-D reconstruction of coronary arteries from X-ray images (angiograms). In these applications it is typically possible to only acquire four to eight views, so that classical X-ray reconstruction methods based on, for example, filtered backprojection cannot be employed. In the conventional biplane

angiographic imaging system set-up, two images are simultaneously acquired at 90° to each other; the center points of vessels or bifurcations (junctions) are extracted in both views, either manually or using semiautomatic methods; these features are matched between views; and then triangulation is used to find the location of the vessels in three dimensions. Having reconstructed the essential global geometric structure, it is possible to estimate the size and the shape of blood vessels from X-ray intensities and assumptions about vessel shape (15,16,17). This approach has not proved so successful for cerebral vessel reconstruction, due to the increased complexity of and variability in the cerebral vasculature branching pattern, which makes establishing correspondences difficult. Although there has been some progress in this area using model-based techniques that employ a 3-D preoperatively acquired model of vessel structure to guide correspondence matching (18), the general problem of 3-D reconstruction from a few projection angiograms using geometric methods remains unsolved.

Finally, there has been some recent work investigating variants of the iterative algebraic reconstruction technique (*ART*) (19) to reconstruct cerebral vessels and pathologies such as aneurysms and arteriovenious malformations from few-view data (20,21). Classical ART and its multiplicative equivalent (MART) do not work well on limited-view data, because the optimization problem is underconstrained. To overcome this, regularization can be used to favor a smooth reconstruction result. This approach has produced promising results and may find application in interventional neuroradiology procedures and for angiographic examinations performed prior to radiotherapy. In this case, full volumetric CT reconstruction is not feasible, due to acquisition-time constraints or limitations of equipment that cannot rotate at the high speeds required for 3-D acquisition.

Image Enhancement and Segmentation. Much effort in medical image processing is dedicated to enhancing the visibility of structures of interest in an image (such as tumors or vessels) and reducing imaging noise and the effects of imaging artifacts, as a preprocessing step prior to image measurement and interpretation. The process of producing an image where features of interest are more clearly visible is called *image enhancement*. Although general- purpose image-processing filters such as Gaussian and median filtering can improve image appearance, the danger is that the application of such methods can make lesions less easy to locate or, worse still, remove them. Thus the most successful approaches to image enhancement are methods that aim to model the physics of image formation or make use of prior knowledge about object properties such as shape. Thus image enhancement is very much a modality- and application-dependent task. Here we highlight some examples in different domains.

By modeling the mammographic imaging process Highnam et al. (22), show that it is possible to transform a mammographic X-ray image into an image (representation) of breast tissue where principal degradation effects and "uninteresting" structure have been removed. Orkisz et al. (23) developed a matched-filter algorithm for enhancing small vessels in magnetic resonance angiographic images based on anisotropic diffusion (see below), which favors improving the contrast of line-like features. To correct for intrascan intensity inhomogeneities due to RF coils in MR scans, the bias field, Wells et al. (24), developed an adaptive algorithm that classifies pixels according to their most probable tissue type. This approach assumes a parametrized model for the intensities in the image data (in this case they used a Gaussian mixture model) and uses the expectation maximization (*EM*) algorithm (25), to estimate the model parameters (see below for further discussion of the EM algorithm and its application in segmentation). Related ideas have been used, for example, to fit intensity models to multispectral MR images (26), and MRA data (27) so as to identify tissue regions.

The general goal in medical image segmentation is to partition a raw 2-D or 3-D image into regions corresponding to anatomically meaningful structures as a precursor to geometric reasoning or interpretation. For example this might involve identifying vessels in a MRA, classifying tissue in brain MR images, or detecting and characterizing lesions in breast X-ray images. This task differs from image enhancement in that in image enhancement the output is another image. In image segmentation the output is a set of image features (regions, edges, corners), which may subsequently be grouped, prior to being used as input to a measurement or visualization. The segmentation approach adopted is clearly application-dependent and may involve modeling and extracting static or dynamic edges or textured regions, or using local matched filters (e.g., wavelets).

Many of the algorithms in the literature are adaptations of algorithms developed for analyzing visual images. Straightforward methods borrowed from the 2-D image-processing literature—such as thresholding, mathematical morphological operators, or Canny edge detection—can be defined in three dimensions (28), and can be applied to segment sufficiently clean low-textured CT and MRI data. However, these approaches do not always work well on highly textured and noisy data such as ultrasound and brain MR images. This is because in these cases boundaries are not necessarily well modeled by a straight (1-D) step edge [Canny's model (29)], and the assumption that object regions are associated with image regions of constant intensity corrupted with additive noise is not valid.

To overcome these difficulties a number of generic, model-driven medical image segmentation methods have evolved. We mention three of the key types of algorithm here. These all make use of geometric models or models of image physics to help guide image segmentation. Two of these model-based approaches make use of a combination of geometric and intensity knowledge rather than local intensity information only; they are based on deformable models (30), and anisotropic diffusion (31), and make use of a combination of geometric and intensity knowledge. A third, EM segmentation (24), is a statistical approach that has its foundations in statistical stochastic theory.

Deformable model segmentation combines *a priori* knowledge of the shape of a geometric structure (typically represented parametrically as a 2-D contour spline or a 3-D surface) with constraints derived from the image in an energy minimization functional, which can be accomplished iteratively using numerical algorithms based on, for example, finite element methods. For instance, in two dimensions, the energy functional $E(x, y)$ of a deformable contour model could be represented by

$$E(x, y) = S(x, y) + P(x, y) \qquad (1)$$

Here, $S(x, y)$ is the internal energy characterizing the deformation of the contour, and $P(x, y)$ is a scalar potential derived from image measurements. For example, the contour will be attracted to intensity edges in an image $I(x, y)$ if $P(x, y) = c|\text{grad } I(x, y)|$, where c is a weighting factor. Numerous applications of deformable-model-based segmentation have appeared in the medical image analysis literature for automatically or semi-automatically segmenting. 2-D and 3-D anatomical structures (see the extensive review in see the extensive review in Ref. 32). This approach does, however, require a good initialization, and the user is usually left to decide when the process has converged. The traditional approach does not take into consideration changes in topology, which are particularly important for vessel segmentation, but there has been some recent research aimed at extending the deformable model concept to accommodate this need, as in Ref. 33.

A variant of the general idea is a *dynamic* deformable model that replaces Eq. (1) by a functional that describes the shape and motion of an object. This allows the quantification of shape evolution (deformation) over time. See the subsection "Intrepretation" below on motion and tracking.

Anisotropic diffusion was introduced by Perona and Malik as a method for performing edge-preserving smoothing prior to edge detection (31). This is a particularly attractive feature in medical applications because of their low signal-to-noise ratio and weak edges. The original idea was to smooth homogeneous regions while preserving the boundaries according to the anisotropic diffusion equation

$$I_t = \text{div}[c(x, y, t) \text{ grad } I(x, y)]$$

Here I_t represents the image, and $c(x, y, t)$ is a conductance function that is monotonic in the image gradient. The conductance function defines the manner in which diffusion occurs across an edge. For example, one might penalize smoothing at edges with a high gradient (which are assumed significant) but allow smoothing at edges with a low gradient (assumed spurious). Various authors have investigated different forms of the conductance function to incorporate knowledge of expected geometric structure or image physics. For instance, Steen and Olstad (34), integrated anisotropic diffusion with an ultrasound signal model, Gerig et al. applied

Fig. 2. Example of anisotropic diffusion: (a) original image; (b) classic Perona–Malik anisotropic diffusion; (c) knowledge-based anisotropic diffusion (37).

the approach to multichannel MR data (35), Krissian et al. constrained smoothing of 3-D vascular images according to an analysis of local differential structure (36), and Sanchez-Ortiz et al. defined the conductance function as a product of a function of the image gradient and a (cylindrical) shape function to segment cardiac cine MR images (37) (see Fig. 2).

The EM algorithm provides a method for estimating the distribution parameters of a data set given an appropriate model for the distribution by maximizing the likelihood of the distribution. Having estimated the parameters, each data point can subsequently be labeled as belonging to the class for which the conditional probability of belonging to the class is maximized. This approach has been successfully applied to identify cerebral tissues from MR image data with (24), and without (26,38) bias field correction, and to segment vascular structures (27). An attraction of the method is that classification is fully automatic and that different models for the intensity distributions can be readily introduced to accommodate statistical variations in the data.

Finally, due to the wide variety of image types and image quality, it is frequently not possible to automatically segment an image. Thus there is a role for interactive tools for image segmentation where a manually placed approximate contour is adjusted using an iterative minimization process to find an optimal boundary. Snakes or active contour algorithms and live-wire segmentation (39), are examples of this type of approach.

Image Understanding, Analysis, and Interpretation

Although one use of medical image processing is to provide better visualization, its prime use is for quantification, that is, to provide clinical meaningful measurements to improve diagnosis and treatment. Much of the work in this area involves the application of geometric methods to solve general computer vision problems such as registration and matching, shape analysis, and motion and tracking. In some cases this has led to the development of new concepts and new techniques that go beyond the boundaries of computer vision (e.g., in the areas of medical image registration and, increasingly, medical motion analysis).

Registration and Matching. Medical image analysis frequently involves integrating information from multiple modalities. For example, in radiotherapy treatment, it is useful to be able to register a MR scan (which gives better delineation of tumor tissue) with a CT scan (which is needed to estimate radiation dosage). It is also frequently desired to integrate anatomical and functional modalities. Registration of time sequences of images of the same modality is also important, to monitor treatment and disease progression for example. The process of bringing the modalities into spatial alignment is referred to as *registration*. This is one of the most successful areas where medical image processing has been applied in clinical practice to date. Most work has focused on the 2-D–3-D and 3-D–3-D rigid registration of two images (no time involved).

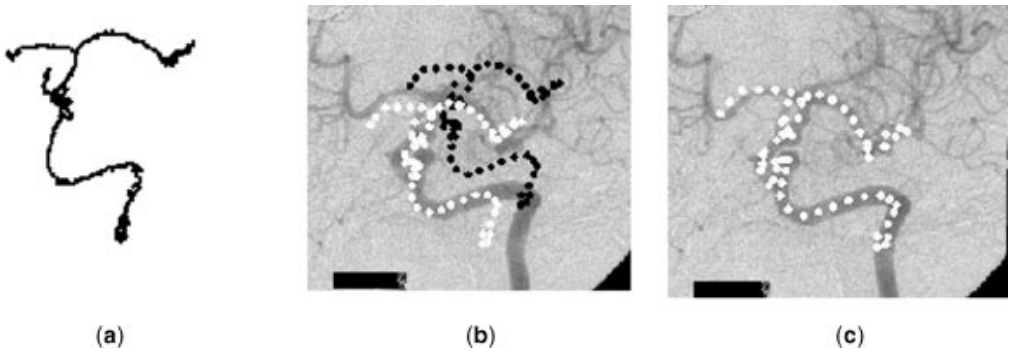

Fig. 3. Example of 3-D–2-D feature-based registration: (a) 3-D skeleton of a 3-D MRA segmentation scan like Fig. 1(b); (b) initial approximate alignment (white dots) of projected 3-D skeleton model (black dots) on 2-D X-ray view; (c) best-fit 3-D–2-D registration of projected model to image data (45).

For recent surveys of registration methods the reader is referred to Refs. 40 and 41. In Ref. 41, registration methods are classified according to nine criteria: dimensionality; nature of registration basis (extrinsic or intrinsic); nature of transformation (rigid, affine, projective, or curved); domain of transformation; degree of interaction; optimization procedure; modalities involved; subject (intrasubject, intersubject, or image-to-atlas); object type (part of the body). We will briefly consider the second criterion to highlight some of the popular state-of-the-art methods.

Extrinsic registration involves placing markers (fiducial objects) on a patient that are designed to be readily visible and detectable in each modality. Correspondence between the data sets is established manually or automatically, and the parameters of the transformation determined using for example linear or nonlinear least-squares fitting. A commonly used fiducial object utilized in neurosurgery is a stereotactic frame that is screwed rigidly to a patient's skull (42). Noninvasive markers glued to the skin are also popular. This type of approach is relatively simple and fast to use, but is often restricted to rigid transformations (translations and rotations)—clearly, the markers cannot move between acquisitions.

Intrinsic methods utilize image content and can be divided into two types: feature-based and voxel-based. Feature-based registration (43,44,45) utilizes salient feature points (landmarks), contours, or surfaces derived via some segmentation process as features for matching (see Fig. 3). These features might for example be anatomically significant points located interactively by a user, or automatically localized points of extrema in curvature or object surface points. Optimization algorithms are then used to optimize measures such as the average distance (L_2 norm) between each landmark and its closest counterpart (the Procrustean metric) or the iterated minimal landmark distance. Optimization of the latter using the iterative closest point (*ICP*) algorithm and its variants has proved especially popular(43,46,47), gradient measures (49), or mutual information (50, 51,52) (see also Ref. 53). These methods are being actively researched at the current time. They tend to be computationally expensive but are automatic and, with the increased speed of computers, are now usable in clinical practice for applications such as rigid or affine 3-D–3-D registration of head images. However, they have not yet been successfully applied to time-constrained tasks such as intraoperative 2-D–3-D registration.

Most of the methods currently in clinical use perform rigid-model registration. However, there is also great interest in using deformable models to deform one image elastically to another. This type of approach finds application in intrasubject matching or atlas registration as well as for tracking objects, such as heart chambers, over time. Typically a template is constructed from one image and is then deformed to fit a segmented structure in the second using an interative process. The demon algorithm (54), and the ICP algorithm (43), have been used for this purpose. The success of the approach depends on the choice of deformation model (elasticity constraint) and auracy of the segmentation process. Registration algorithms using nonpoint features such as

special curves and special points on those curves are also being investigated (55,56). The motivation is that richer feature descriptors enable the quality of match to be defined more robustly, which should lead to more powerful methods of registering medical images. However, the need to extract reproducible geometric features for multimodality registration, and reliable features in noisy data has limited the practical use of this approach to date.

Interpretation.

Shape Representations and Analysis. The representation of 2-D and 3-D shape and expected shape across a population is a very important problem in medical image analysis. For instance, many medical objects do not have any identifying characteristics (such as color or texture) other than shape. Shape parameters (such as area, eccentricity, axes of symmetry) and their variances can provide a natural way to describe properties of an anatomical object for detection or classification purposes. Further, the definition of a parametrized model of shape enables model-based (so-called top-down) segmentation and registration methods to be employed, which are more robust to image measurement error than pure image-based (so-called bottom-up) techniques.

Most work in medical shape analysis deals with the statistical analysis of landmarks, that is, sets of specially selected points (57,58). A prime reason for this is that the statistical analysis of point sets is well understood and that measures such as the Euclidean distance metric between two points and the covariance of a point set can be defined. Higher-level shape primitives, such as boundary points (points with direction normals), curves, medial loci (points on loci of symmetry), and surfaces, appear less frequently in the literature. These provide more powerful features for registration and matching. However, the statistical analysis of nonpoint primitives is still being developed (59), and they are typically more difficult to extract from image data.

A popular method to describe a medical object is as an *eigenshape model* (60,61,62). An eigenshape model is a statistical shape model that represents the variability of a set of examples (the training set) in a compact and statistically robust way. The model is constructed by representing each example in the training set as a vector in a high-dimensional vector space and performing principal component analysis (*PCA*) upon the resulting set of vectors (63). PCA determines the modes of variation of the data as the eigenvectors of the data set with largest eigenvalues. If the variability is well represented by a few modes, then the dimensionality of the data can be reduced. Typically three to eight modes capture 95% of the total variability. When the eigenshape model is constructed from a set of closely spaced points along the boundary of an object, it is sometimes called a *point distribution model* (*PDM*) (62). PDM s and the related concept of an active shape model (*ASM*) which combines a model of shape and gray-level appearance, have been successfully applied to a variety of medical problems, for instance, to locate structures in noisy images (64), to learn normal and abnormal left-ventricle "average"-shape models (65), and to detect abnormal instances of 2-D biological shapes (62). One key issue in eigenspace analysis is how to select suitable corresponding features on the training-set examples. In most cases this is done by hand, which requires expert knowledge of the objects being analyzed and is time-consuming, and the problem of how to do this in three dimensions is unsolved. Further, the individual modes do not typically describe shape variation in a clinically meaningful way. However, in spite of these limitations, eigenshape decomposition remains one of the most powerful and convenient ways to simplify the description and shape analysis of medical objects.

Another way to describe shape is in terms of medial (skeletal) primitives. Here the essential idea is to summarize a shape by the loci of points that lie medially between its boundaries and their distance to the boundary. The work has been largely restricted to two dimensions and has focused on extracting skeletons from objects rather than the applications of the representation. Pizer et al. develop a related medial representation of shape, called the *core*, that is extracted directly from image intensities rather an (already extracted) object boundary. Cores can be defined in three dimensions and, unlike classical skeletonization, are insensitive to small image disturbances (66,67).

A skeleton or core characterizes shape in terms of constituent parts. An alternative view is to characterize a shape via a sequence of transformations (or deformations) from a simple shape. Koenderink's theory on

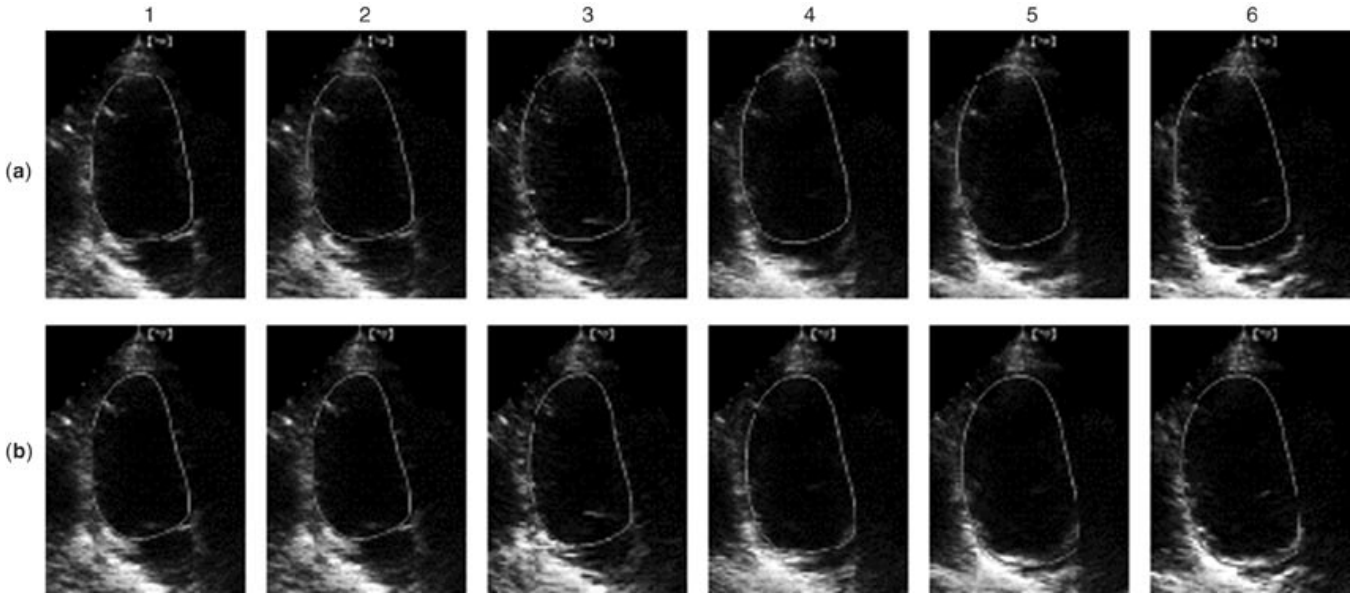

Fig. 4. 2-D echocardiogram model-based tracking: (a) tracking using a simple (affine) motion model; (b) tracking using a motion model learnt from some training data. The latter enables the actual motion to be followed more accurately; this is most clearly seen by comparing frames 4 to 6 in (a) and (b) (75).

dynamic shape (68), and Kimia's *reaction–diffusion-space* theory (69) fall into this category, but neither has yet been shown to be usable for computational image analysis.

Finally, there is a whole subfield of applied statistics, called morphometrics that deals with multivariate analysis of shape (70,71). In this approach the emphasis is on analyzing *landmark sets* (point sets), which are sets of anatomically significant features located on the object of interest. It is assumed that these have been identified in each image instance of the object by some mean (usually manual selection). The key idea is to define a distance (the Procrustes distance), or metric, that is a least-squares formula used to describe the distance between a pair of shapes. This then enables measures such as an average shape to be defined and classical linear multivariate analysis to be performed to test group differences and the like. Bookstein suggested using a thin-plate spline to warp one average shape onto another as a way to visualize shape differences as patterns in the deformation grid. Such techniques are being used, for example, in studies of schizophrenia to understand whether the shape characteristics of a schizophrenic's brain, as measured from MR brain scans, differ statistically from those of a normal human brain.

Motion and Tracking. There has been considerable interest in developing tracking and deformation analysis methods for nonrigid objects, motivated primarily by the need to develop techniques for assessing the dynamical behavior of a heart, and in particular the left ventricle. Much effort has been undertaken to automate the tracking and analysis of heart chamber movement from 2-D echocardiographic image sequences using, for example, methods based on optic-flow computation (72) or snakes (73,74). Heart movement can be complex, and ultrasound measurements are noisy, with the net result that a classical snake can wander from the object of interest. To increase tracking robustness, Cootes et al. propose to constrain the allowable contour deformation to lie within the bounds of a shape model learnt from a large training set (62). Jacob et al. develop a Kalman-filter-based tracker whose system model is a motion model whose parameters have been learnt from observing normal heart motion (75) (see Fig. 4).

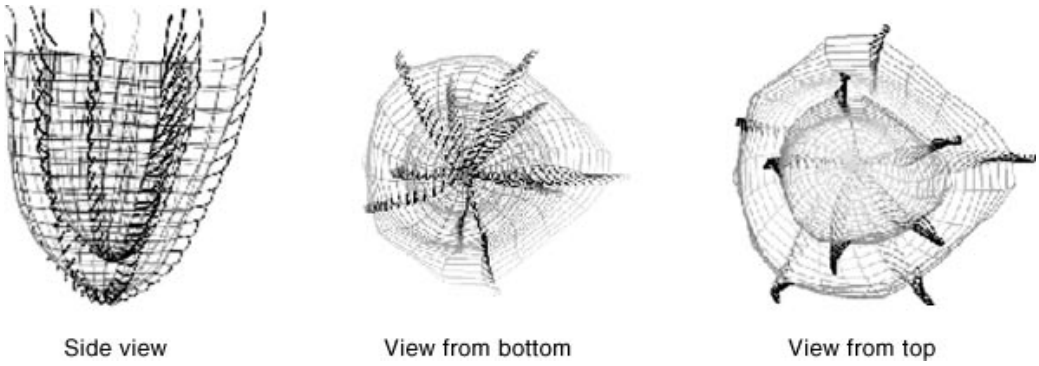

Fig. 5. Wireframe model showing the trajectories of selected points on the left ventricle endocardial and epicardial walls. The model was derived by best-fitting a 4-D (3-D + T) heart model to tagged MR volume data (79).

Two-dimensional ultrasound analysis has the advantages that is relatively cheap and can be performed rapidly and is the most widely accepted method of cardiac image analysis in clinical practice today. It is however limited in that slice analysis does not enable 3-D motion to be characterized—the socalled through-plan problem—with the result that ultrasound scan view directions have to be carefully chosen. With the recent advances in 3-D cardiac imaging technology [especially nuclear-medicine SPECT or positron emission tomography (*PET*) imaging, tagged-MR imaging, and 3-D ultrasound], 3-D heart image analysis may soon become a reality. To date, the focus in 3-D cardiac image analysis has been in two areas: 3-D heart reconstruction and tracking, and shape modeling (76,77,78,79). Two general approaches to tracking have evolved. The first approach, suitable for MR, SPECT or PET data, performs a segmentation and surface fit on each volume frame and then computes the motion field between model fits in pairs of successive frames (76,77). The second approach, used in tagged-MR image analysis, finds features (usually manually extracted points) in each frame, establishes the frame-to-frame correspondence of features, and then fits some model to the features (78,80). The first approach involves extensive volume processing and is too slow for practical use. The principal difficulty in the second approach is reliable feature detection.

The end goal of left ventricle tracking is to provide intuitive parameters that describe characteristic normal and abnormal heart chamber motions, but how model parameters correlate to clinical conditions is still an active area of research. The superquadric model of Ref. 78, which describes basic heart motion in terms of a contraction, elevation, and twist, and the 4-D polar transformation of Ref. 79, which adapts this to incorporate parameters describing regularity and periodicity (Fig. 5), are perhaps the most developed methods at the current time.

Some Applications

Image processing is now utilized in many areas of medicine to help improve diagnosis and treatments, and a complete list of applications falls outside the scope of this review. Here we highlight a few of the areas:

- *Atlas Matching*. This is a special case of registration where a brain image acquired from a patient is compared with an "average" or statistical brain image, or *template*, obtained from imaging of many subjects. Any statistically significance different local or global shape change is attributed to abnormal growth or development. The key issue is how to develop an average representation of neuroanatomy that captures the normal variability and structural and functional differences of the brain observed across a healthy

population; methods based on deformable templates (81,82) and probabilistic techniques (83,84) have been applied to solve this problem. Brain atlases are used to help understand pathological states such as Alzheimer's disease as well as to guide neurosurgical procedures.

- *Orthopedic and Craniofacial Surgery*. Image processing is increasingly being used in the planning and execution of orthopedic surgery, for example, to register real-time intraoperative fluoroscopy images with 3-D preoperative CT images in spine surgery or to guide a robot in hip replacement surgery (85). Here, imaging is used in presurgical planning and to provide in vivo measurements to improve the accuracy and precision of surgical execution. Similar needs have led to the application of 3-D imaging in craniofacial surgical reconstruction (86). 3-D CT-based modeling and visualization can be used to develop a very precise 3-D computer model of a custom implant, typically using a mirror image of the normal side of the face as the template for the prosthesis. During the execution of surgery, the preoperative model can be registered to the patient's skull (using a number of optical beacons placed on the bone-fragments to act as landmarks) and be used to guide cutting and bone-fragment repositioning. Medical image processing is also starting to be used for orthopedic surgery simulation and in related clinical areas (87,88).
- *Neurosurgery*. Neurosurgery provided one of the earliest application domains for 3-D medical imaging in clinical practice and remains one of the dominant successful areas of application of 3-D medical image processing. One of the basic requirements in neurosurgical operations is to be able to register a preoperative 3-D image (typically a CT and/or MR scan) to a patient as they are seen during surgery in order to guide a surgeon to the target of interest. Thus the main problem is one of 3-D-to-2-D (MR or CT to video image) registration. Conventional stereotaxy, in which a rigid frame is attached to the patient's skull before preoperative scanning and left in place for surgery to provide a common frame of reference, was used in early work. It was cumbersome and found to be suitable only for a small of operations in practice. The general trend has been to move to frameless registration which does not use an attached frame or landmarks, but relies on techniques borrowed from visual image object recognition to register a preoperative 3-D model to real-time video images (89,90). For instance, Grimson et al. at MIT and Brigham and Women's Hospital, Boston, have developed a system that performed a 3-D-to-3-D surface registration between a preoperative MR (or CT) scan and a laser range scan of the patient's skull in the operating room, and then determined the relationship between a video camera and the patient by matching video images of known points on the table to the actual 3-D laser data of the same features. The registered 3-D model is then transformed into the frame of reference of the video image, so that the surgeon sees the model superimposed on the video image (89). Colchester et al. at Guy's Hospital, London, have independently developed a similar system, VISLAN. This system initially performs a registration to the skin surface, but has the capability to update this initial registration in the event of patient movement using a point-based method that uses markers affixed to the skull at the beginning of surgery (90). Variants of both systems are now used regularly in clinical practice. More recent research includes work aimed at increasing the accuracy of the guidance process by using multiple video images (91), and the development of registration methods that relax the rigid-body assumption to enable tissue deformations (such as brain swelling) to be accommodated (92,93) in order to maintain the registration during a surgical procedure.
- *Mammography*. Mammographic image processing primarily deals with processing X-ray images of a breast (X-ray mammograms) to diagnose breast cancer. One of the earliest reports on mammographic image processing by 94. Since this time therhas been much research in this area, motivated by the need to support the growing size of national screening programs (25 million woman are screened annually in the EC at a cost of approximately US $3 billion a year) and the desire to introduce computer-assisted methods to reduce the cost and improve the accuracy of diagnosis. Mammographic image processing represents a tough challenge to image processing because the images typically have a low signal-to-noise ratio, in the 5 dB to 6 dB range, and many of the abnormalities that one wants to detect, such as microcalcifications, appear as quite subtle and irregular changes in intensity. Highnam and Brady provide an excellent up-to-date review on the field (95). Research has addressed issues such as understanding noise properties

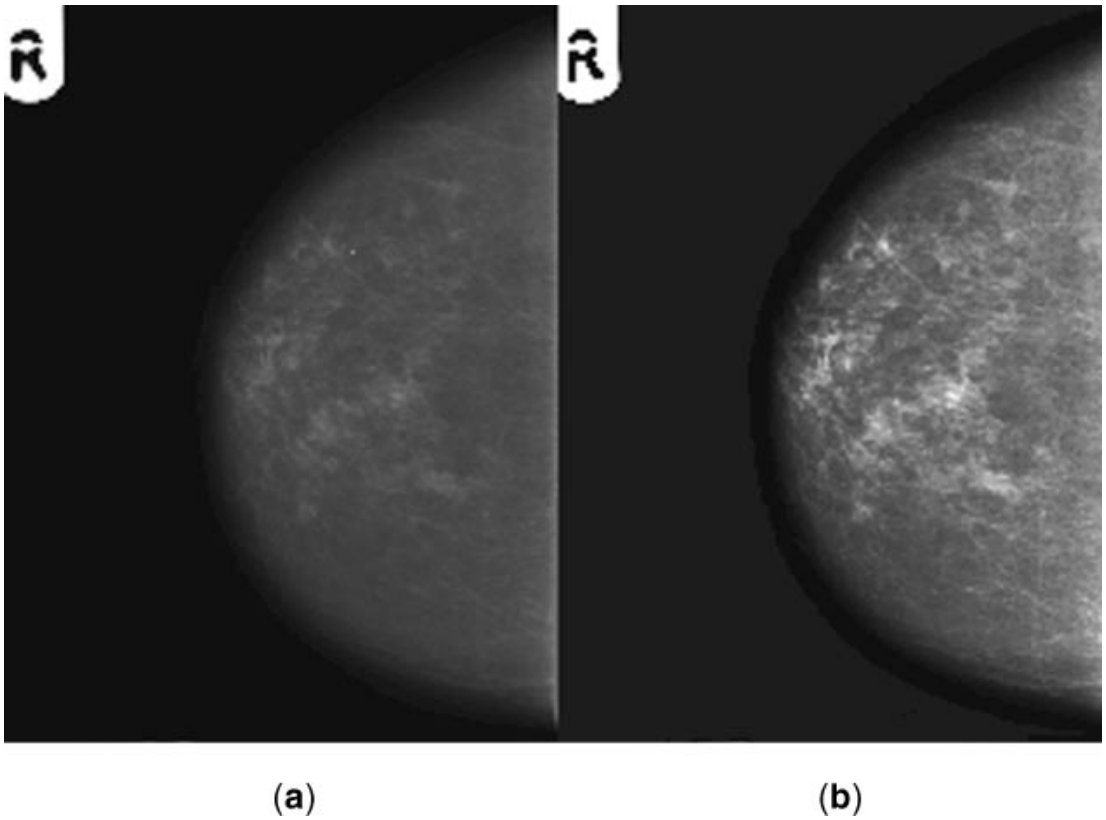

Fig. 6. X-ray mammogram (a) before and (b) after physics-based enhancement (95).

of images (96), image enhancement (22), developing mammographic feature extraction methods (97), and quantification (98). Some work has also been done to match two X-ray views of the same breast to try to infer 3-D information (such as the volume extent of a suspected mass, and more precise locations of microcalcifications) that cannot be inferred from a single X-ray view (99). Finding reliable features turns out to be the limiting factor here that has prevented the successful application of this idea to date. Highnam and Brady (95) argue that successful mammographic image processing has to be firmly based on a physics-based approach that takes into consideration how the image is formed. In their approach they go through an image enhancement step to construct what they called an h_{int} image which is a measure of breast tissue (rather than intensity) at each image location, from a conventional X-ray image. All further image measurement and interpretation are based on the h_{int} image. (see Fig. 6.) Today, most of the successful algorithms that are used in clinical practice are individual techniques designed to assist a radiographer in diagnosis. However, the ultimate goal is to develop fully automatic procedures that could, at a minimum, reduce the large volume of mammograms viewed by a radiographer to a small number of "suspect" cases. Fully automatic microcalcification detection and classification schemes and quantification are active subjects of current research (100,101,102,103). The idea of embedding information provided by image processing into an AI-based clinical decision-making system to aid trained radiographers in the interpretation of mammograms is another interesting area under investigation (104).

- *Cardiology*. Most medical image-processing research in cardiac analysis has focused on the problem of tracking left-ventricle boundary motion to provide global heart function measurements such as the ejection

fraction (ratio of minimum volume to maximum volume) as outlined in the preceding section. As noted there, 2-D echocardiograpy is the most widely used cardiac imaging technique in clinical partice. Commercial echo systems have a relatively limited image-processing capability, consisting perhaps of acoustic boundary detection, simple linear or area measurement, and intensity profiling tools. Automated image analysis systems are still at the laboratory stage of development (75,105,106). Recent advances in ultrasonic imaging technology such as second-harmonic imaging, and the use of contrast agents, together with the introduction of digital and 3-D ultrasound technology, have created renewed interest in this area. 3-D SPECT imaging is used routinely to provide 3-D density maps of blood perfusion. Image processing is used to register stress and rest images and to quantify the extent of diseased tissue (myocardial infarction) (107,108,109). Techniques for measuring 3-D heart motion based on tagged MR are also being actively developed. However, tagged-MR and nuclear-medicine cardiac imaging are currently only available in large research laboratories.

Further Reading

For further information about current research in medical image processing refer to the leading journals in the field, such as *Medical Image Analysis* (Oxford University Press) and the *IEEE Transactions on Medical Imaging*, or to specialized journals on medical imaging modalities, the appropriate clinical fields, or computer vision and image processing journals. The journals referenced below should provide a good starting point.

BIBLIOGRAPHY

1. T. Shiemann, *et al.* Segmentation of the visible human for high quality volume based visualization, *Med. Image Anal.*, **1**: 263–270, 1997.
2. S. W. Atlas, *et al.* Intracranial aneurysms: Depiction on MR angiograms with a multifeature extraction, ray-tracing postprocessing algorithm, *Radiology*, **192**: 129–139, 1994.
3. D. L. Wilson An improved planning protocol for the endovascular treatment of intracranial aneurysms, Ph.D. thesis, Department of Engineering Science, University of Oxford, Oxford, UK, 1998.
4. G. T. Herman Three-dimensional imaging on a CT or MR scanner, *J Comput. Assist. Tomogr.* **12**: 450–458, 1988.
5. W. Schroeder K. Martin B. Lorensen *The Visualization Toolkit: An Object-oriented Approach to 3D Graphics*, Upper Saddle River, NJ: Prentice Hall, 1997.
6. A. Kaufman R. Yagel, D. Cohen, Volume graphics, *IEEE Comput.* **26** (7): 51–64, 1993.
7. J. D. Foley, *et al. Computer Graphics: Principles and Practice* 2nd ed., Reading, MA: Addison-Wesley, 1990.
8. W. E. Lorensen H. E. Cline, Marching cubes: A high-resolution 3D surface construction algorithm, *Comput. Graphics*, **21** (3): 163–169, 1987.
9. H. Fuchs, M. Levoy, S. M. Pizer, Interactive visualization of 3D medical data, *IEEE Comput.* **22** (8): 46–51, 1989.
10. S. L. Wood, Visualization and modeling of 3D structures, *IEEE Eng. Med. Biol.*, **11** (2): 72–79, 1992.
11. L. F. Hodges, Tutorial: Time-multiplexed stereoscopic computer graphics, *IEEE Comput. Graphics Appl.*, **12** (2): 20–30, 1992.
12. A. Hernandez, *et al.* Acquisition and stereoscopic visualization of 3D ultrasound breast data, *IEEE Trans. Ultrason. Ferro. electr. Freq. Control*, **43**: 576–579, 1996.
13. A. Moskalik, *et al.* Registration of three-dimensional compound ultrasound scans of the breast for refraction and motion correction, *Ultrasound Med. Biol.*, **21**: 769–778, 1995.
14. R. N. Rohling, A. H. Gee, L. Berman, *Automatic Registration of 3D Ultrasound Images*, Tech. Rep. CUED/F-INFENG/TR 290, Engineering Department, Cambridge, University, Cambridge, UK: 1997.
15. K. Kitamura, J. M. Tobis, J. Sklansky, Estimating the 3D skeletons and transverse areas of coronary arteries from biplane angiograms, *IEEE Trans. Med. Imaging*, **7**: 173–187, 1988.
16. M. Garreau, *et al.* A knowledge-based approach to 3D reconstruction and labelling of vascular networks from biplane angiographic projections, *IEEE Trans. Med. Imaging*, **10**: 122–131, 1991.

17. L. V. Tran, R. C. Bahn, J. Sklansky, Reconstructing the cross sections of coronary arteries from biplane angiograms, *IEEE Trans. Med. Imaging*, **11**: 517–529, 1992.
18. E. Bullitt, *et al.* Three-dimensional reconstruction of intracranial vessels from biplane projection views, *J. Neurosci. Methods*, **66**: 13–22, 1996.
19. A. C. Kak M. Slaney, *Principles of Computerized Tomographic Imaging*, New York: IEEE Press, 1988.
20. L. Launay, *et al.* 3D reconstruction of cerebral vessels and pathologies from a few biplane digital angiographies, *Proc. Visualization Biomed. Comput. Conf. 1996*, pp. 122–128, 1996.
21. E. Payot, Reconstruction vasculaire tridimensionelle en imagerie pat rayons X, Doctoral thesis, Laboratoire d'Electronique de Technologie et d'Instrumentation (*LETI*), 1996.
22. R. Highnam, *et al.* A representation for mammographic image processing, *Med. Image Anal.*, **1** (1): 1–18, 1996.
23. M. M. Orkisz, *et al.* Improved vessel visualisation in MR angiography by nonlinear anisotropic filtering, *Magn. Reson. Med.*, **37**: 914–919, 1997.
24. W. M. Wells, III, *et al.* Adaptive segmentation of MRI data, *Int. Conf. Comput. Vision, Virtual Reality Robot. Med.*, pp. 59–69, 1995.
25. A. P. Dempster, N. M. Laird, D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *J. R. Stat. Soc.*, **39**: 1–38, 1977.
26. Z. Liang, J. R. MacFall, D. P. Harrington, Parameter estimation and tissue segmentation from multispectral MR images, *IEEE Trans. Med. Imaging*, **13**: 441–449, 1994.
27. D. L. Wilson J. A. Noble, An adaptive segmentation algorithm for extracting arteries and aneurysms from time-of-flight MRA data, *IEEE Trans. Med. Imaging*, to appear October 1999. An earlier version of this work was presented in *Proc. Int. Conf. Image Process. Med. Imaging*, pp. 423–428, 1997.
28. O. Monga, *et al.* Recursive filtering and edge tracking: Two primary tools for 3D edge detection, *Image and Vision Computing*, **9** (4): 203–214, 1991.
29. J. Canny, A computational approach to edge detection, *IEEE Trans. Pattern Anal. Mach. Intell.*, **PAMI-8**: 679–698, 1986.
30. M. Kass, A. Witkin, D. Terzopoulos, Snakes: Active contour models, *Int. J. Comput. Vision*, **1** (4): 321–331, 1988.
31. P. Perona J. Malik, Scale-space and edge detection using anisotropic diffusion, *IEEE Trans. Pattern Anal. Mach. Intell.*, **12**: 629–639, 1990.
32. T. McInerey D. Terzopoulos, Deformable models in medical image analysis: A survey, *Med. Image Anal.*, **1** (2): 91–108, 1996.
33. T. McInerey D. Terzopoulos, Topologically adaptable snakes, *Proc. IEEE Int. Conf. Comput. Vision*, pp. 840–845, 1995.
34. E. Steen B. Olstad, Scale-space and boundary detection in ultrasonic imaging using nonlinear signal-adaptive anisotropic diffusion, *SPIE, Image Process.*, **2167**: 116–127, 1994.
35. G. Gerig, *et al.* Nonlinear anisotropic filtering of MRI data, *IEEE Trans. Med. Imaging*, **MI-11**: 221–232, 1992.
36. K. Krissian, G. Malandain, N. Ayache, Directional anisotropic diffusion applied to segmentation of vessels in 3D images, *INRIA Tech. Rep. 3064*: 1996.
37. G. I. Sanchez-Ortiz, D. Rueckert, P. Burger, Knowledge-based anisotropic diffusion of vector-valued 4-dimensional cardiac MR images, *Proc. Br. Mach. Vision Conf.*, pp. 605–614, 1996.
38. Z. Liang, R.J. Jaszczak, R.E. Coleman, Parameter estimation of finite mixtures using the EM algorithm and information criteria with application to medical image processing, *IEEE Trans. Nucl. Sci.*, **39**: 1126–1133, 1992.
39. W. Barrett E. Mortensen, Interactive live-wire boundary extraction, *Med. Image Anal.*, **1** (4): 331–341, 1996.
40. S. Lavallee, Registration for computer integrated surgery: Methodology state-of-the-art, in R. H. Taylor et al., (eds.), *Computer-Integrated Surgery*, Cambridge, MA: MIT Press, 1996, pp. 115–143.
41. J. B. A. Maintz M. A. Viergever, A survey of medical image registration, *Med. Image Anal.*, **2** (1): 1–36, 1998.
42. L. D. Lunsford, *Modern Stereotactic Neurosurgery*, Boston: Martinus Nijhoff, 1988.
43. J. Feldmar, *et al.* Rigid, affine and locally affine registration of free-form curves and surfaces, *Int., J. Comput. Vision*, **18**: 99–119, 1996.
44. S. Lavallee, R. Szeliski, L. Brunie, Matching 3D smooth surfaces with their 2D projections using 3D distance maps, Proc. *SPIE*, **1507**: 322–336, 1991.
45. Y. Kita, D. L. Wilson, J. A. Noble, Real-time registration of 3D cerebral vessels to X-ray angiograms, Proc. MICCA198, *Lect. Notes Comput. Sci.*, **1496**: Springer Verlag, 1125–1133, 1998.

46. P. J. Besl N. D. McKay, A method for registration of 3D shapes, *IEEE Trans., Pattern Anal. Mach. Intell.*, **PAMI-14**: 239–256, 1992.
47. A. Rangarajan, H. Chui, F. L. Bookstein, The soft assign Procrustes matching algorithm, *Proc. Int. Conf. Image Process. Med. Imaging*, pp. 29–42, 1997.
48. J. Weese, *et al.* An approach to 2D/3D registration of a vertebra in 2D X-ray fluoroscopies with 3D CT images, *Proc. Joint Conf. Comput. Vision, Virtual Reality Robot. Med. (CVRMSD)* and *Med. Robot. Comput. Assisted Surgery (MRCAS)*, 1997, pp. 119–128.
49. L. Lemieux, *et al.* A patient-to-computed-tomography image registration method based on digitally reconstructed radiographs, *Med. Phys.*, **21** (11): 1749–1760, 1994.
50. P. A. Viola, Alignment by maximization of mutual information, Ph.D. thesis, Massachusetts Institute of Technology, AI Tech. Rep. 1548, 1995.
51. A. Collignon,, *et al.* Automated multimodality image registration using information theory, *Proc. Int. Conf. Image Process. Med. Imaging*, pp. 263–274, 1995.
52. C. Studholme, D. L. G. Hill, D. J. Hawkes, Automated 3D registration of MR and CT images of the head, *Med. Image Anal.*, **1** (2): 163–175, 1996.
53. G. P. Penney, *et al.* A comparison of similarity measures for use in 2D–3D medical image registration, *IEEE Trans. Med. Imaging*, **17**: 586–595, 1998.
54. J.-P. Thirion, Image matching as a diffusion process: An analogy with Maxwell's demons, *Med. Image Anal.*, **2** (3): 243–260, 1998.
55. J.-P. Thirion, New feature points based on geometric invariants for 3D registration, *Int. J. Comput. Vision*, **18** (2): 121–137, 1996.
56. A. P. Gueziec, X. Pennec, N. Ayache, Medical image registration using geometric hashing, *IEEE Comput. Sci. Eng.*, **4** (4): 29–41, 1997.
57. S. Small, *The Statistical Theory of Shape*, Berlin: Springer-Verlag, 1996.
58. I. L. Dryden K. V. Mardia, *Statistical Shape Analysis*, New York: Wiley, 1998.
59. X. Pennec N. Ayache, Uniform distribution, distance and expectation problems for geometric features processing, *J. Math. Imaging Vision*, **9**: 49–67, 1998.
60. S. Sclaroff A. P. Pentland, Modal matching for correspondence and recognition, *IEEE Trans. Pattern Anal. Mach. Intell.*, **17**: 545–561, 1995.
61. C. Nastar N. Ayache, Frequency-based nonrigid motion analysis: Application to four dimensional medical images. *IEEE Trans. Pattern Anal. Mach. Intell.*, **18**: 1067–1079, 1996.
62. T. F. Cootes, *et al.* Active shape models—their training and application, *Comput. Vision Image Understanding*, **61**: 38–59, 1995.
63. W. J. Krzanowski, *Principles of Multivariate Analysis*, Oxford, UK: Oxford University Press, 1988.
64. T. F. Cootes, *et al.* The use of active shape models for locating structures in medical images, *Image Vision Comput.*, **12** (6): 276–285, 1994.
65. A. Hill C. J. Taylor, Model based image interptreation using genetic algorithms, *Image Vision Comput.*, **10**: 295–300, 1992.
66. S. M. Pizer, D. Eberly, D. S. Fritsch, Zoom-invariant vision of figural shape: The mathematics of cores, *Comput. Vision Image Understanding*, **69** (1): 55–71, 1998.
67. B. S. Morse, *et al.* Zoom-invariant vision of figural shape, *Comput. Vision Image Understanding*, **69** (1): 72–86, 1998.
68. J. J. Koenderink, *Solid Shape*, Cambridge, MA: MIT Press, 1990.
69. B. B. Kimia, A. R. Tannenbaum, S. W. Zucker, Shapes, shocks and deformations. I: The components of shape and reaction–diffusion space, *Int. J. Comput. Vision*, **15**: 189–224, 1995.
70. F. L. Bookstein, *Morphometric Tools for Landmark Data*, Cambridge, UK: Cambridge University Press, 1991.
71. F. L. Bookstein, Shape and the information in medical images: A decade of the morphometric synthesis, *Comput. Vision Image Understanding*, **66** (2): 97–118, 1997.
72. G. E. Mailloux, *et al.* Computer analysis of heart motion from two-dimensional echocardiograms, *IEEE Trans. Biomed. Eng.*, **34**: 356–364, 1986.
73. N. Ayache, I. Cohen, I. Herlin, Medical image tracking, in A. Blake and A. Yuille (eds.), *Active Vision*, Cambridge, MA: MIT Press, 1992, pp. 285–302.

74. V. Chalana, *et al.* A multiple active contour model for cardiac boundary detection in echocardiographic sequences, *IEEE Trans. Med. Imaging*, **15**: 290–298, 1996.
75. G. Jacob, *et al.* Evaluating a robust contour tracker on echocardiographic sequences, *Med. Image Anal.*, **3** (1): 63–76, 1999.
76. P. Shi, *et al.* A unified framework to assess myocardial function from 4D images, *Proc. Int. Conf. Comput. Vision, Virtual Reality, Robot. Med.*, pp. 327–337, 1995.
77. E. Bardinet, L. Cohen, N. Ayache, Superquadrics and free-form deformations: A global model to fit and track 3D medical data, *Proc. Int. Conf. Comput. Vision, Virtual Reality, Robot. Med.*, pp. 319–326, 1995.
78. J. Park, D. Metaxas, L. Axel, Analysis of left ventricle wall motion based on volumetric deformable models and MRI-SPAMM, *Med. Image Anal.*, **1** (1): 53–71, 1996.
79. J. Declerck, J. Feldmar, N. Ayache, Definition of a 4D continuous planispheric transformation for the tracking and the analysis of LV motion, *Med. Image Anal.*, **4** (1): 1–17, 1998.
80. T. S. Denney J. L. Prince, 3D displacement field reconstruction from planar tagged cardiac MR images, *IEEE Workshop Biomed. Image Anal.*, pp. 51–60, 1994.
81. A. C. Evans, *et al.* Warping of a computerized 3D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis, in *Medical Imaging III, SPIE*, 1991, pp. 264–274.
82. G. E. Christensen, S. C. Joshi, M. I. Miller, Volumetric transformation of brain anatomy, *IEEE Trans. Med. Imaging*, **16**: 864–877, 1997.
83. A. C. Evans, *et al.* An MRI-based probabilistic atlas of neuroanatomy, in S.D. Shorvon et al. (eds.), *Magnetic Resonance Scanning and Epilepsy*, New York: Plenum, 1994, pp. 263–274.
84. P. M. Thompson, A. W. Toga, A surface-based technique for warping 3-dimensional images of the brain, *IEEE Trans. Med. Imaging*, **15**: 1–16, 1996.
85. R. H. Taylor, *et al.* An image-directed robotic system for precise orthopaedic surgery, *IEEE Trans. Robot. Autom.* **10**: 261–275, 1994.
86. C. B. Cutting, F. L. Bookstein, R. H. Taylor, Applications of simulation, morphometrics, and robotics in craniofacial surgery, in R. H. Taylor et al., (eds.), *Computer-Integrated Surgery*, Cambridge, MA: MIT Press, 1996, pp. 641–662.
87. S. Gibson, *et al.* Simulating arthroscopic knee surgery using volumetric object representations, real-time rendering and haptic feedback, *Proc. Comput. Vision, Virtual Reality Robot. Med. Conf.*, Berlin: Springer-Verlag, 1997.
88. S. Cotin, H. Delingette, N. Ayache, Real-time non-linear elastic deformations of soft tissues for surgery simulation, *IEEE Trans. Vis. Comput. Graphics*, **5** (1): 62–73, 1999.
89. W. E. L. Grimson, *et al.* An automatic registration method for frameless stereotaxy, image guided surgery, and enhanced reality visualization, *IEEE Trans. Med. Imaging*, **15**: 129–140, 1996.
90. A. C. Colchester, *et al.* Development and preliminary evaluation of VISLAN, a surgical planning and guidance system using intra-operative video imaging, *Med. Image Anal.*, **1** (1): 73–90, 1996.
91. M. J. Clarkson, *et al.* Registration of multiple video images to pre-operative data for image-guided surgery, *Proc. Med. Image Understanding Anal. 98*, Leeds, UK, pp. 73–76, 1998.
92. P. J. Edwards, D. L. G. Hill, D. J. Hawkes, Image guided interventions using a three component tissue deformation model, *Proc. Med. Image Understanding Anal. Conf.*, Oxford, UK, pp. 33–36, 1997.
93. E. Grimson, *et al.* Clinical experience with a high precision image-guided neurosurgery system, *Proc. Med. Image Computing and Computer-Assisted Interventions, Lect. Notes in Comput. Sci.*, **1496** Springer-Verlag, 1998, pp. 63–73.
94. F. Winsberg, *et al.* Detection of radiographic abnormalities in mammograms by means of optical scanning and computer analysis, *Radiology*, **89**: 211–215, 1967.
95. R. H. Highnam, J. M. Brady, *Mammographic Image Processing*, Kluwer Ser.—Med. Imaging, Boston: Kluwer Academic Publishers, 1999.
96. R. M. Nishikawa, M. J. Yaffe, Signal-to-noise properties of mammographic film–screen systems, *Med. Phys.*, **12**: 32–39, 1985.
97. T. C. Parr, *et al.* Model-based classification of linear structures in digital mammograms, *3rd Int. Workshop on Digital Mammography Int. Congr. Ser, Excerpta Med.*, **1119**: 1996.
98. J. H. Richter, E. Claridge, Extraction of quantitative blur measures for circumscribed lesions in mammograms, *Med. Inf.* **16**: 229–240, 1991.
99. W. Spiesberger, Mammogram inspection by computer, *IEEE Trans. Biomed. Eng.* **26**: 213–219, 1979.

100. N. Karssemeijer, Adaptive noise equalization and recognition of microcalcification clusters in mammograms, *Int. J. Pattern Recognition Artif. Intell.*, **7**: 1357–376, 1993.
101. D. H. Davies, D. R. Dance, Automatic computer detection of clustered calcifications in digital mammograms, *Phys. Med. Biol.*, **35**: 1111–1118, 1990.
102. T. Ema, *et al.* Image feature analysis and computer-aided diagnosis in mammography: Reduction of false-positive clustered microcalcifications using local edge-gradient analysis, *Med. Phys.*, **22**: 161–169, 1995.
103. R. M. Nishikawa, *et al.* Computer-aided detection of clustered microcalcifications on digital mammograms, *Med. Biol. Eng. Comput.* **33**: 174–178, 1995.
104. P. M. Taylor, J. Fox, A. Todd-Pokropek, A model for integrating image processing into decision aids for diagnostic radiology, *Artif. Intell. Med.*, **9**: 205–225, 1997.
105. V. Chalana, Y. Kim, A methodology for evaluation of boundary detection algorithms on medical images, *IEEE Trans. Med. Imaging*, **16**: 642–652, 1997.
106. P. J. Slomka, *et al.* Automated alignment and sizing of myocardial stress and rest scans to three-dimensional normal templates using an image registration algorithm, *J. Nucl. Med.*, **36**: 1115–1122, 1995.
107. A. Giachetti, On-line analysis of echocardiographic image sequences, *Med. Image Anal.*, **2** (3): 261–284, 1998.
108. R. Mullick, N. F. Ezquerra, Automatic determination of LV orientation from SPECT data, *IEEE Trans. Med. Imaging*, **14**: 88–99, 1995.
109. J. Declerck, *et al.* Automatic registration and alignment on a template of cardiac stress and rest reoriented SPECT images, *IEEE Trans. Med. Imaging*, **16**: 727–737, 1997.

J. ALISON NOBLE
University of Oxford