Segmentation of carpal bones from CT images using skeletally coupled deformable models

Thomas B. Sebastian[a], Hüseyin Tek[a], Joseph J. Crisco[b], Benjamin B. Kimia[a,*]

[a]*LEMS, Division of Engineering, Brown University, Providence, RI 02912, USA*
[b]*Department of Orthopedics, Rhode Island Hospital, Providence, RI 02903, USA*

Abstract

The in vivo investigation of joint kinematics in normal and injured wrist requires the segmentation of carpal bones from 3D (CT) images, and their registration over time. The non-uniformity of bone tissue, ranging from dense cortical bone to textured spongy bone, the irregular shape of closely packed carpal bones, small inter-bone spaces compared to the resolution of CT images, along with the presence of blood vessels, and the inherent blurring of CT imaging render the segmentation of carpal bones a challenging task. We review the performance of statistical classification, deformable models (active contours), region growing, region competition, and morphological operations for this application. We then propose a model which combines several of these approaches in a unified framework. Specifically, our approach is to use a curve evolution implementation of region growing from initialized seeds, where growth is modulated by a skeletally-mediated competition between neighboring regions. The inter-seed skeleton, which we interpret as the predicted boundary of collision between two regions, is used to couple the growth of seeds and to mediate long-range competition between them. The implementation requires subpixel representations of each growing region as well as the inter-region skeleton. This method combines the advantages of active contour models, region growing, and both local and global region competition methods. We demonstrate the effectiveness of this approach for our application where many of the difficulties presented above are overcome as illustrated by synthetic and real examples. Since this segmentation method does not rely on domain-specific knowledge, it should be applicable to a range of other medical imaging segmentation tasks.
© 2002 Published by Elsevier Science B.V.

Keywords: Segmentation; Carpal bones; Skeletally coupled deformable model

1. Introduction

Degenerative joint disease is commonly attributed to alterations in joint loading and joint kinematics due to traumatic injury. In the wrist, despite widespread clinical awareness of dynamic and static wrist instability, little is known about the pathoanatomy and kinematics of these conditions. Patients may continue to be incapacitated by pain following stressful activities months after injury, even though radiographs and other static imaging studies appear normal. Characterizing the true 3D kinematics of the carpal bones following these ligament injuries would provide better insight for the development of diagnostic techniques, surgical treatment, rehabilitation, the design of prosthetic devices and more appropriate treatment strategies (Lee and Massear, 1993; Mayfield, 1984; Ruby et al., 1987; Savelberg, 1991; Savelberg et al., 1991, 1993).

Invasive methods for measuring 3D joint motion are common in orthopedic research. While instrumentation ranges from video to stereoradiogrammetry, all of the existing methods use specific landmarks (e.g. infrared reflectors or implanted tantalum balls) on each rigid body segment. Such methods have been used in vivo (Karrholm et al., 1989; Koh et al., 1992; Lafortune et al., 1992) to

*Corresponding author.
E-mail address: kimia@lems.brown.edu (B.B. Kimia).

study the knee and hip, but the small size of the carpal bones of the wrist and the invasive nature of these methods limits the in vivo application. Our general approach to measuring 3D carpal motion in vivo requires registration of bone surfaces extracted from multiple CT volumes (Crisco et al., 1999). The extraction of these surfaces, in turn, requires robust and reliable *segmentation* methods. A key goal of this paper is to develop a segmentation technique suitable to this domain.

Medical image segmentation, however, has proved to be a challenging task. This is true, in particular, for the segmentation (and registration) of carpal bones in the wrist from CT images. While the segmentation of bones in X-ray and CT images is viewed to be a relatively straightforward task, carpal bone segmentation is difficult because the volumetric datasets contain irregularly shaped bones with small inter-bone distances relative to the resolution of CT imaging. A key fact is that bone tissue cannot be characterized uniformly: the outer layer of the bone tissue, or *cortical bone*, is denser than the *spongy bone* it encases. Thus, under CT imaging cortical bone appears brighter and smooth, while spongy bone appears darker and textured. In addition, due to the close spacing of some carpal bones and inherent blurring in CT imaging, the inter-bone space often appears brighter than the background (soft tissue), substantially reducing boundary contrast at these points. Finally, blood vessels resemble the background, creating gaps in the surface of bone images. In the image domain, these characteristics translate into four challenging areas for segmentation techniques (Fig. 1): (i) gaps in the cortical shell; (ii) weak or diffused bone boundaries due to the partial volume effect in CT imaging; (iii) textured areas corresponding to the spongy bone alternating between bone-like and tissue-like intensities; and (iv) the narrow inter-bone regions which tend to be diffused.

While it is possible for medical experts to segment these images by using thresholding or manual seeding, e.g., by using the ANALYZE package[1] or similar tools, this process is highly labor intensive, considering the time required for accurate manual correction. Thus, the development and use of segmentation techniques that minimize user interaction is highly desirable. We have implemented and evaluated several techniques for the segmentation of carpal bones from a sequence of 2D CT images, including global thresholding (Weeks et al., 1984; James et al., 1992), statistical methods (Duda and Hart, 1973), seeded region growing (Adams and Bischof, 1994), deformable models like snakes (Kass et al., 1988), balloons (Cohen and Cohen, 1993), bubbles (Tek and Kimia, 1997), morphological watersheds (Vincent and Soille, 1991), and region competition (Zhu and Yuille, 1996). The experience with the use of these techniques in our domain has prompted us to combine three classes of these approaches in a single framework. Specifically, in one approach taken by Tek and Kimia (1997), numerous seeds are initialized, both inside and outside objects of interest, which then grow by image-dependent forces. The growing seeds merge in the absence of boundaries to form larger seeds, and finally slow down near boundaries, thus trapping the boundary between the inner and outer regions. The success of this technique is dependent on the existence of boundaries with sufficient contrast and symmetric initialization in the case of weak boundaries. This is due to the monotonic nature of growth: once a region has evolved beyond object boundaries it can no longer return to capture it. Region competition (Zhu and Yuille, 1996), on the other hand, also relies on the growth of seeds, but implements a

---

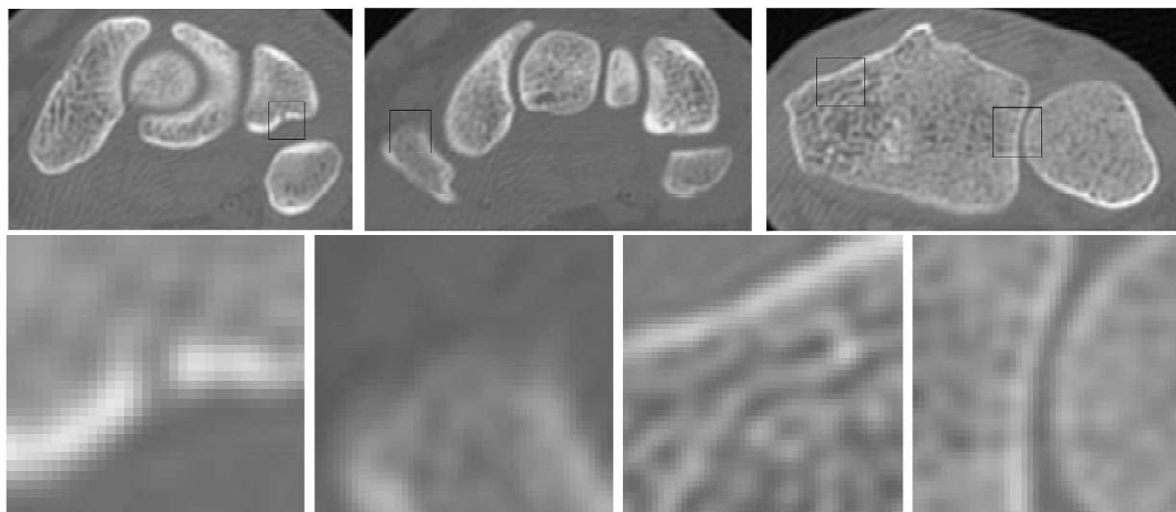[1] ANALYZE is a medical imaging software package developed at the Mayo Foundation, Rochester, MN.

Fig. 1. Three slices from a three-dimensional CT image of carpal bones (top row) demonstrate gaps or weak edges, diffused edges, textured areas, and extremely narrow inter-bone regions that make automatic segmentation difficult. Each window is zoomed in the bottom row detailing the above features, respectively.

local competition between growing seeds, once they have become adjacent. This local back and forth competitive movement of adjacent regions is based on a statistical decision depending on to which of these regions a point is more likely to belong. The central assumption underlying this scheme is that the growth of seeds leads to regions that characterize distinct areas. This assumption fails when the growing seeds in 'waiting' for other regions to arrive, acquire and encompass two statistically distinct domains. Seeded region growing (Adams and Bischof, 1994) avoids this difficulty by implementing a *global* competition among growing regions, but does not implement the local 'back and forth' competition between them, thus also not allowing for recovery from errors.

The approach presented in this paper combines these three ideas, namely, deformable models implemented in the curve evolution framework, local back and forth competition or region competition, and the global competition of seeded region growing, under one framework. The main idea is to rely on the inter-region skeleton as a predictor of boundaries resulting from the growth of current seeds. Assuming current growth conditions continue to hold, this predicted boundary couples two seeds and is used to affect their respective growth process by modulating their deformation speed: if a point on the pair of paths leading to the formation of a skeletal point is more likely to belong to one seed compared to another, then the former region should grow faster at that point to capture it. Region competition then becomes a special case, i.e., when the two regions become adjacent. The idea of global competition in seeded region growing is implemented by the long-distance competition among neighboring seeds, mediated by the inter-region skeleton, with the advantage that it eliminates the irrelevant interaction between the very distant seeds with other seeds in between.

This paper is organized as follows. Section 2 reviews our experience with some current segmentation techniques for carpal bone segmentation. In Section 3, we describe the skeletally coupled deformable model (SCDM). The implementation details are discussed in Section 4. The segmentation results and validation studies are reported in Section 5.

2. Segmentation of carpal bones: current approaches

We have investigated the use of several segmentation techniques for their specific use in the recovery of carpal bone surfaces from CT images. These methods include global thresholding (Weeks et al., 1984; James et al., 1992), statistical classification (Duda and Hart, 1973), seeded region growing (Adams and Bischof, 1994), region competition (Zhu and Yuille, 1996), deformable models like snakes (Kass et al., 1988), balloons (Cohen and Cohen, 1993), their curve evolution counterparts (Malladi et al., 1994; Caselles et al., 1993, Tek and Kimia, 1995), and watershed segmentation (Vincent and Soille, 1991). This section briefly reviews our experience with the application of these techniques to carpal bone segmentation.

2.1. Statistical methods

Global thresholding (Weeks et al., 1984; James et al., 1992) is the simplest statistical segmentation technique, where pixels are classified based on their intensity values. However, choosing the right intensity threshold, which typically varies from one dataset to another, is difficult. Interactive manual selection of the threshold is tedious and operator-sensitive. Even with the optimal threshold, final segmentation based on thresholding creates holes and in some images two adjacent bones merge (Fig. 2). Morphological filters fill in the holes, but often merge distinct bones. The intensity threshold can be automatically found by fitting a mixture of Gaussian distributions to the intensity values, and estimating the distribution parameters using the expectation maximization (EM) algorithm (Dempster et al., 1977; Redner and Walker, 1984). Once the parameters are estimated, classical Bayesian Decision theory (Duda and Hart, 1973) is used to find the decision boundaries, which in the 1D case, is the intensity threshold we want. However, the estimation of intensity threshold for carpal bone segmentation using EM is not reliable as there is an overlap between the intensities of the bone and soft tissue (Fig. 3). Other statistical methods include maximum likelihood methods or non-parametric methods
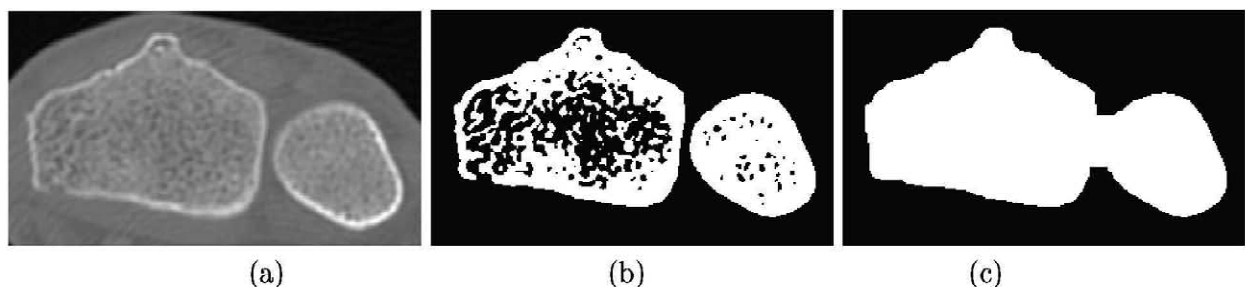

Fig. 2. An example of a 2D slice where global thresholding is used. (a) Original image. (b) Segmentation using the optimal threshold. (c) Result of applying morphological filtering (Serra, 1982) on the segmented image. Observe that there are holes in the bones as well as a gap in the bone contour. While morphological filtering can close the holes and gaps, the inter-bone region will also be closed and the shape of the boundary will be slightly altered.
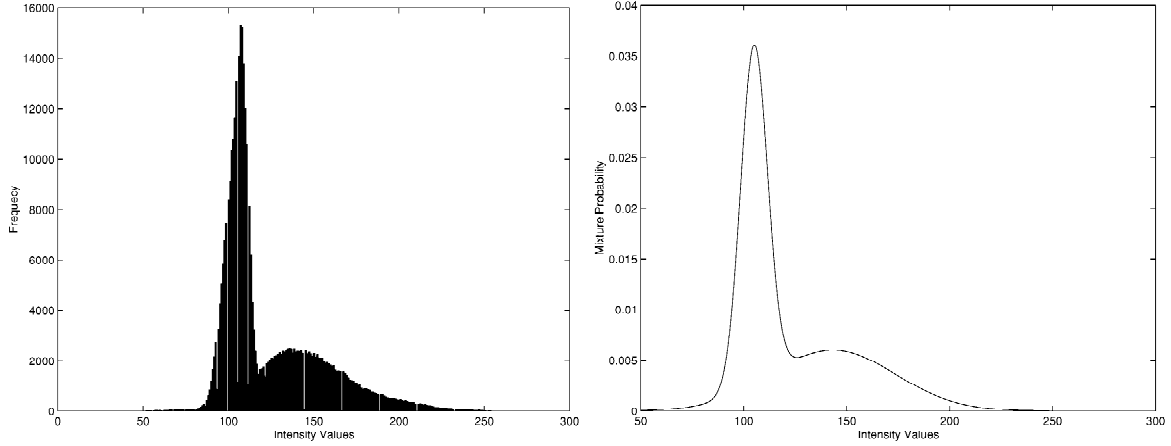
Fig. 3. Left: A histogram of intensities of soft tissue and bones, obtained from a set of 2D slices. Note the significant overlap of the distribution of these two classes. Right: The result of fitting a mixture of Gaussian distributions to the intensity values using the EM algorithm (Dempster et al., 1977; Redner and Walker, 1984).

like k nearest neighbors (kNN) (Duda and Hart, 1973, Cooper et al., 1980).

2.2. Deformable models

Deformable models are geometric descriptions of contours or surfaces which evolve under a suitable energy. These were pioneered by Kass et al. (1988), who introduced *snakes*, energy minimizing splines influenced by image forces and external constraint forces. Formally, let $\mathscr{C}(s) = (x(s), y(s))$ represent the snake, where s is the arc length parameter. The energy functional of a snake is defined as

$$E(\mathscr{C}) = \int_0^1 [E_{\text{int}}(\mathscr{C}(s)) + E_{\text{image}}(\mathscr{C}(s)) + E_{\text{con}}(\mathscr{C}(s))] \, ds, \quad (1)$$

where E_{int} represents the internal energy of the spline that imposes regularity on the curve, E_{image} represents image forces that are responsible for pushing the snake towards salient image features, and E_{con} are the external constraint forces. The energy landscape of the snake E is typically not convex and can have several local minima. Hence, snakes must be initialized close to the boundaries for proper convergence. To overcome the initialization restrictions of the snakes, Cohen and Cohen (1993) proposed the *balloon* model. Balloons grow under a similar energy to the snake, but with an additional constant inflation force to transport the initialized model closer to the edges.

Snakes (Kass et al., 1988) and balloons (Cohen and Cohen, 1993) cannot easily handle topological changes. Caselles et al. (1993) and Malladi et al. (1994) resolve this issue by using the curve evolution approach which models the snake as level sets of surfaces evolving under similar energies, with the additional dimension allowing for the notion of interior/exterior of the snake. Specifically, a curve $C(s)$ is considered to be the zero level set of the surface $\phi(x, y) = 0$, and the zero level set of ϕ is evolved according to

$$\frac{\partial \phi}{\partial t} = S(x, y)(\beta_0 - \beta_1 \kappa(x, y))|\nabla \phi|, \quad (2)$$

where $S(x, y)$ is the image based speed function defined for the zero level set and extended to all other level sets, defined as $\hat{S}(x, y) = 1/(1 + |\nabla G_\sigma * I(x, y)|)$, β_0 and β_1 are constants, and κ is the curvature. While these models handle topological events well, they do not address the initialization and convergence issues. Tek and Kimia (1995) use a random initialization of many seeds, which in the early stages of growth resemble *bubbles*, to address the initialization issue and to include regional characterization of each bubble in the evolution process (Fig. 4). Bubbles are very effective in segmenting images with small structures, e.g., vascular structures, which render manual initialization impractical. However, practical use of bubbles requires that (i) parameters be tuned for simultaneous convergence of bubbles, and (ii) an appropriate domain-dependent stopping time be selected by the user. Another approach to handle topological events has been proposed by McInerney and Terzopoulos (1995) where multiple seeds are explicitly modeled as topologically adaptable snakes. Caselles et al. (1997) and Kichenassamy et al. (1995) proposed *geodesic active contours* to tackle the convergence issue. This is a curve evolution counterpart to snakes, and inherits the initialization difficulties, thus still requiring significant user interaction for segmentation. In summary, for carpal bone segmentation, the deformable models encounter similar problems: (i) not all models converge at weak/diffused boundaries; (ii) the contour smoothing terms do not allow entry into the narrow inter-bone region; (iii) as these models rely only on the local information along the boundary, the texture inside the bone slows snakes down, resulting in poor convergence on the bone boundary; (iv) initialization requirements imply significant and often impractical user interaction.
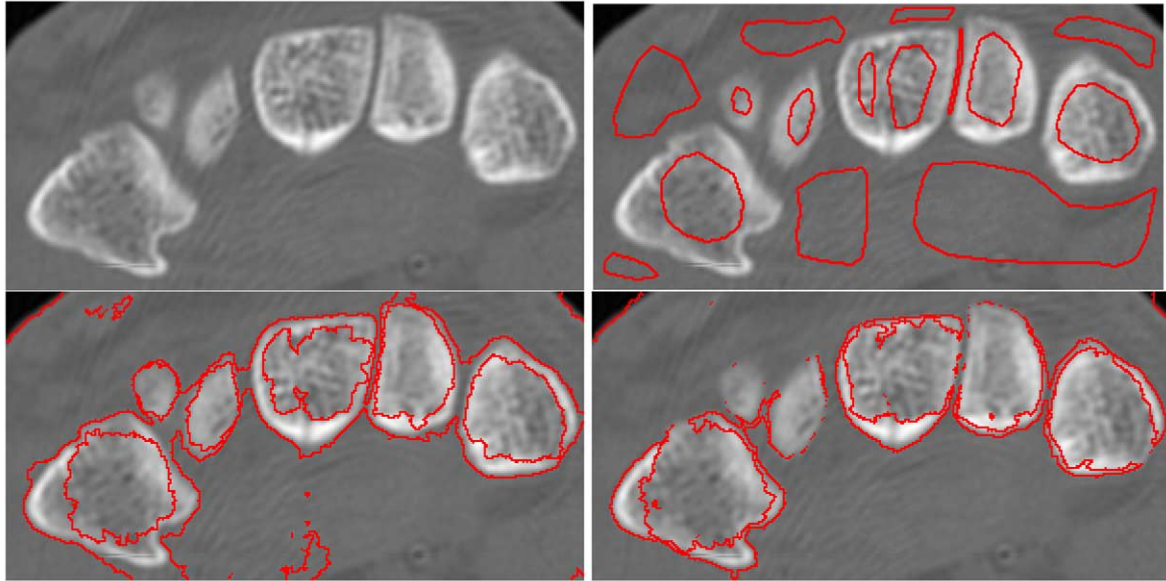
Fig. 4. This figure illustrates a shortcoming of the bubbles (Tek and Kimia, 1995) approach. Top row: Left: Original image. Right: Initialized seeds. Bottom row: Left: Initialized seeds have grown to near bone boundaries; note the convergence of bubbles in the small bone on the top left. Right: In the course of further iterations, which are necessary for convergence in other places, some bubbles leak out, e.g., the bubble inside the small bone evolves to include the inter-bone region.

2.3. Region growing/seeded region growing

Another class of techniques for segmentation is *region growing* and *merging*, where initialized seeds grow by annexing 'similar' pixels, where similarity is defined by a statistical test. The goal of region growing and merging is to divide the domain R of the image I into regions $\{R_i : i = 1 \ldots N\}$, such that $R = \cup_{i=1}^{N} R_i$, $R_i \cap R_j = \phi$, if $i \neq j$, and I satisfies a homogeneity criterion on each R_i (Beveridge et al., 1989). Region growing is typically followed by region merging where small regions are grouped into larger ones using a statistical test. A popular choice is the *Fisher's* test. Let R_i and R_j be two adjacent regions, with n_i, n_j, μ_i, μ_j, σ_i^2 and σ_j^2 the sizes, sample means and the sample variances of R_i and R_j, respectively. Then the regions R_i and R_j are merged if the squared *Fisher distance*, $(n_i + n_j)(\mu_i - \mu_j)^2/(n_i\sigma_i^2 + n_j\sigma_j^2)$ is below a certain threshold. Region growing methods work well in noisy images, but are sensitive to seed initialization, and in general result in jagged boundaries and poorly localized edges.

Seeded region growing (Adams and Bischof, 1994) improves traditional region growing by introducing a 'competition' between growing regions by ordering all candidate 'growth' pixels according to some suitability criteria. Seeded region growing starts off with a set of seeds, and in each step of the algorithm one pixel is added to one of the growing regions, $R_i : i = 1 \ldots N$. Let T be the set of all unallocated pixels which border at least one of the regions, and is given by

$$T = \left\{ x \notin \bigcup_{i=1}^{N} R_i \,\middle|\, N(x) \cap \bigcup_{i=1}^{N} R_i \neq \emptyset \right\}, \qquad (3)$$

where $N(x)$ is the set of 'immediate' neighbors of the pixel x, $\delta(x)^2$ is a measure of how different x is from the adjacent region R_i it may join. If x adjoins two or more regions, the region that minimizes $\delta(x)$ is chosen. In each iteration, the pixel $z = \text{argmin}_{x \in T} [\delta(x)]$ is appended to $R_i(z)$. This process is repeated until all the pixels have been allocated to one of the regions. The global competition ensures that the growth of regions near weak or diffused edges is delayed till other regions have had a chance to reach this area. This simple, fast technique remarkably improves the performance of traditional region growing methods. However, seeded region growing does not incorporate any geometrical information and hence can 'leak' through narrow gaps or weak edges (like the ones seen in the cortical shell (Fig. 5(b)). It also tends to merge bones that are very close to each other (Fig. 5(a)). Another significant issue is that seeded region growing does not allow for recovery from errors, i.e., once a seed 'leaks out' of the region (object) it is supposed to capture, it cannot be pushed back into the region.

2.4. Region competition

Zhu and Yuille propose an interesting model of *region competition* (Zhu and Yuille, 1996) which combines the geometrical features of deformable models and the statistical nature of region growing, by using a combination of statistical and smoothing forces for seed growth. It also introduces a local competition between regions when they

---

[2]The definition of $\delta(x)$ used in (Adams and Bischof, 1994) is $\delta(x) = |g(x) - \text{mean}_{y \in R_{i(x)}} [g(y)]|$, where $g(x)$ is the intensity at the image point x.
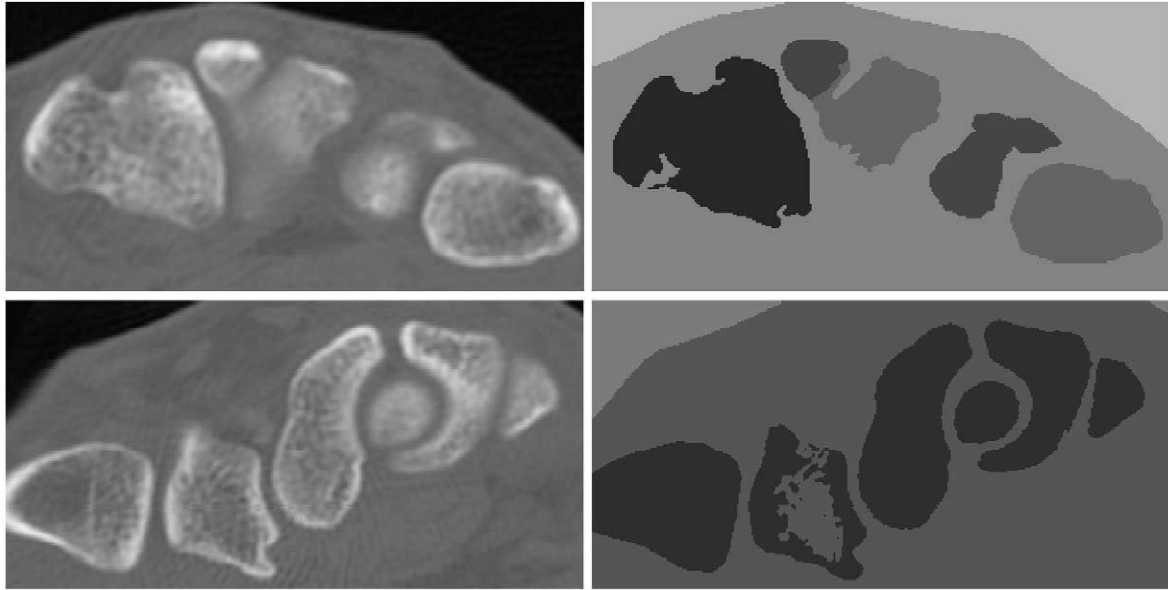
Fig. 5. An application of seeded region growing (Adams and Bischof, 1994) to our domain. Seeded region growing works fairly well in carpal bone segmentation. However, in some cases seeded region growing leaks through gaps and weak edges on the surface of bones.

contact each other, by trading pixels that result in a decrease of energy, thus allowing recovery from errors. Specifically, let R_i, $i = 1 \ldots N + 1$ denote N seed regions (objects) and one background region R_B. Let ∂R_i be the boundary of region R_i, and let the intensity values in region R_i be consistent with having been generated by a probability distribution $P(I|\alpha_i)$, where α_i are the parameters of the distribution. The local deformation at a point C on the boundary of a region R_i consists of a smoothing force $\kappa \vec{N}$ and a statistical force $\log(P(I_C|\alpha_i))\vec{N}$, where κ is the curvature at the boundary of the region, I_C is the image at point C and $\vec{N}$ is the normal to the boundary. As a result of competition between two adjacent regions R_i and R_j, the local deformations of their boundary are based on a single smoothing term for the boundary and a competition between the two statistical forces, leading to

$$\frac{\partial C}{\partial t} = -\beta_1 \kappa \vec{N} + [\log(P(I_C|\alpha_i)) - \log(P(I_C|\alpha_j))]\vec{N}, \quad (4)$$

where C is a point on the common boundary (Fig. 6), α_i

and α_j are parameters of the probability distribution describing the regions R_i and R_j, respectively, and β_1 is a constant. It should be noted that Eq. (4) is derived as the gradient descent for minimizing description length in MDL (Leclerc, 1989), and that the curvature term represents Euclidean curve shortening flow (Grayson, 1987). Region competition implements a *back and forth* competition between adjacent regions, which is continued to convergence. After convergence, two adjacent regions are merged if it leads to a decrease in energy; in this case the competition resumes and is continued until a final convergence is reached resulting in the final segmentation. Region competition is a powerful technique that works well in a wide variety of images, including ones with diffused or weak edges between noisy regions. However, as the authors themselves indicate, this method needs improvement in at least two areas: (i) the lack of effectiveness of smoothing terms leading to jagged boundaries, and (ii) the effective similarity of speeds for statistically different regions. Both drawbacks can be traced to discreti-
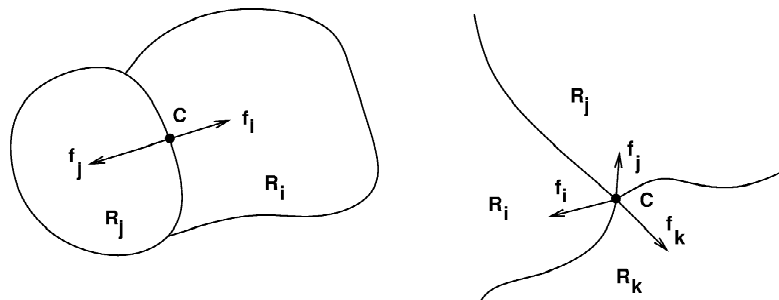

Fig. 6. This figure adapted from (Zhu and Yuille, 1996) illustrates the competition of statistical region forces between adjacent regions used in region competition.

zation effects: curvature computation is not reliable due to the discrete representation of the boundary and boundaries can move either by an entire pixel or not at all. Also, the motion is restricted to four (or possibly eight) directions. Another drawback, which we have found experimentally for carpal bone segmentation, is that while the approach reduces the sensitivity to initial placement of seeds, it is not completely invariant to it (Fig. 7). If seeds are placed 'asymmetrically' with respect to boundaries, i.e., if a seed is placed close to a boundary, and another seed is placed far away, it is possible for the first seed to 'leak through' if the boundary is weak and represents similar regions. If the second seed were to arrive at roughly the same time, the 'region competition' would have reversed the situation and pushed back the extending region. However, due to asymmetric initialization, the first seed has a chance to grow for some time, and loses its statistical characterization, i.e., the parameters specifying the intensity distribution of the seed (mean and variance if a Gaussian model is used) change sufficiently, that it can embody a large number of pixels as well as the competing region can. In this case, recovery is not possible as region competition cannot push back the extruding region (Fig. 7). Finally, observe that while the method implements a merge process, it does not allow for a region to 'split', another significant aspect needed for error recovery.

2.5. Watershed segmentation

Watershed segmentation (Vincent and Soille, 1991) is a morphological gradient-based technique, which can be intuitively described as follows: View the gradient image as a height map, and gradually 'immerse it in water', with water leaking through the minimum gradient points and rising uniformly and globally across the image. Place a 'dam' when two distinct bodies of water (catchment basins) meet and continue the process until water has reached all the points of the image. The dams provide the final segmentation. This can be interpreted in the image domain as the growth of seeds placed on the minima of the image gradient height map at a time proportional to their height, that finally converges on the crest lines of the gradient map. This is a powerful approach especially where local gradients cannot be defined, e.g., diffused edges. However, watershed segmentation cannot be used in our domain due to a well-known and un-resolved *over-segmentation* problem (Fig. 8). Much research is underway, e.g., using marker methods (Vincent and Soille, 1991) to deal with this problem.

As an example of an approach that uses domain-specific knowledge, Tagare et al. (1993) have proposed a dynamic programming based deformable template for carpal bone segmentation. The algorithm is based on the observation that bone boundaries in CT images are locations of strong intensity gradients. A circular or elliptical template is manually placed in the target area using the user's knowledge of carpal bone anatomy. Dynamic programming is used to maximize edge strength (computed using Sobel gradient and the zero crossings of the Marr–Hildreth edge operator) and image intensity along the boundary of the template while keeping it smooth and closed. This 'optimal' boundary is manually adjusted by the user, if
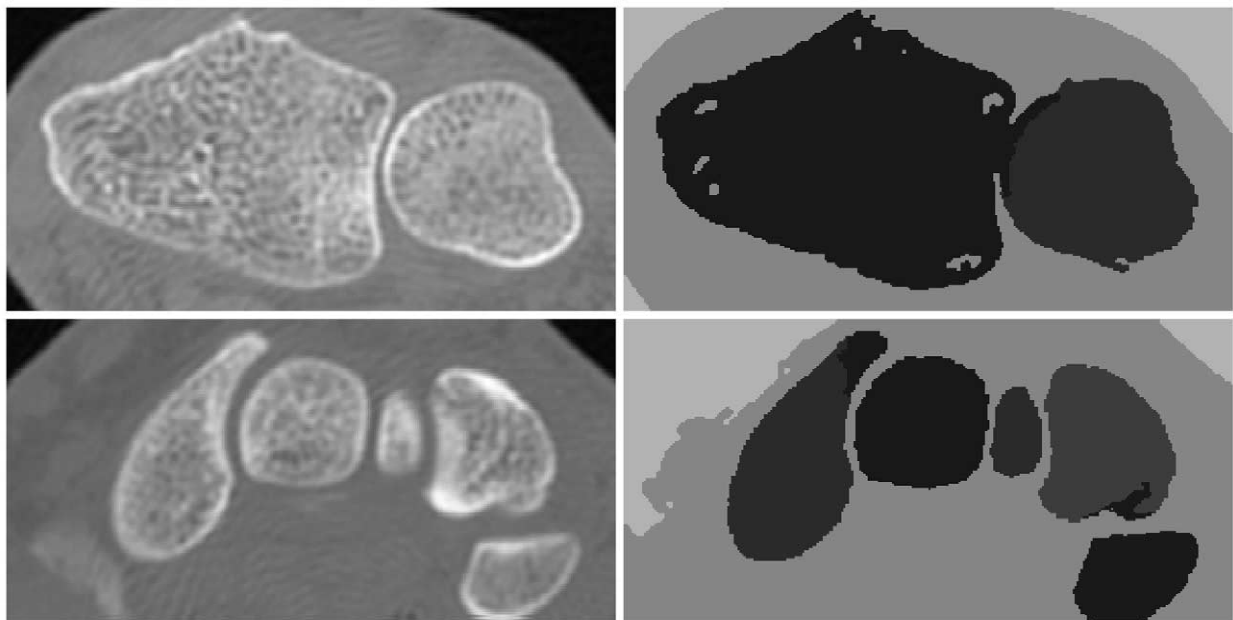

Fig. 7. While region competition (Zhu and Yuille, 1996) frequently captures the correct bone boundaries, in some cases it does not, as illustrated here. In the first case, region competition merges the bones and creates some holes in the captured regions. In the second case region competition fails to segment the bone on the extreme left of the image. Here the seeds that were initialized lost their statistical character, resulting in an incorrect segmentation. The main flaw is that seeds compete only after they have fully grown and become adjacent.
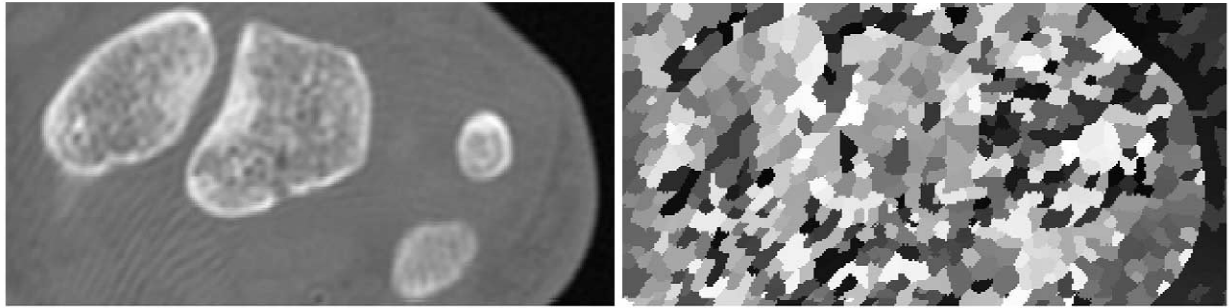
Fig. 8. An example of a 2D slice where watersheds over-segment the image. Much research is underway to correct this problem, such as marker methods (Vincent and Soille, 1991), and this issue currently remains un-resolved.

necessary. While this algorithm is less labor intensive than manual outlining of bone, it still requires significant user interaction.

3. The skeletally coupled deformable model

Our proposed approach, the *skeletally coupled deformable model* (SCDM), can be viewed as a combination of three of the approaches presented above: curve evolution deformable models such as bubbles (Tek and Kimia, 1997), seeded region growing (Adams and Bischof, 1994), and region competition (Zhu and Yuille, 1996). Seeded region growing (Adams and Bischof, 1994) implements a 'global competition' in that every region simultaneously competes with every other. This competition is too global in the sense that it allows two distant regions, which may have many other regions in between and are therefore unrelated, to interact. On the other hand, region competition (Zhu and Yuille, 1996) implements a 'local competition' between adjacent regions, but there is no competition between these regions *before* they have become adjacent. Rather, each region competes with the background. However, observe that delaying competition between regions until they contact each other can cause the seeds to lose their 'character', i.e., their statistical properties, during growth, resulting in an incorrect segmentation (Fig. 7).

SCDM combines the global seeded region growing and the local region competition and brings it under one curve evolution framework by implementing a *long range, predicted competition*, between regions, *before* they become spatially adjacent. This long range coupling is mediated by the *inter-seed skeleton*. The skeleton is interpreted as a predicted location of the final, converged boundaries of growing seeds, and feeds back the statistical desirability of this prediction to the local deformation or growth process.

3.1. Skeletons mediating long-range coupling

To see how long range coupling is mediated by the inter-seed skeleton, consider Fig. 9. Without loss of generality, we restrict our attention first to two regions (seeds), R^+ and R^-, and the background R_B. Let $\mathscr{S}$ denote the *shocks*[3] of R_B as defined in (Kimia et al., 1990, 1995; Siddiqi and Kimia, 1996; Tek et al., 1997), to represent the inter-seed skeleton. Consider a point $A^- \in \partial R^-$, the boundary of R^-, its corresponding shock point $\mathscr{S}(A^-) = A$, and the paired shock point A^+ on the boundary of R^+, $A^+ \in \partial R^+$ such that $\mathscr{S}(A^+) = \mathscr{S}(A^-) = A$. The main idea is that the skeletal point A is interpreted as the predicted point of collision between the boundary points A^+ and A^-, and is used to couple boundary points A^+ and A^-, assuming current growth conditions hold. The proposed competition is based on the prediction that A^- of R^- captures points from the path from A^- to A, and similarly A^+ captures points on the path from A^+ to A, measuring the appropriateness of this prediction, and feeding this measure back into the deformation process. As a first approximation, we compare the relative statistical desirability of each growth path A^-A and A^+A by comparing the two end points of each path. The competition will then be a predicted competition: if the region were to grow as current conditions imply, how would they compete? The process then feeds this competition to modulate growth.

3.2. Local and long-range forces

Specifically, first, consider the local statistical growth force at a point A^- used in typical region growing algorithms (a similar situation holds for A^+),

---

[3]Shocks are the medial axis points (locus of centers of circles that lie completely inside the shape, and are at least bitangent to the boundary) augmented with a notion of dynamics. The distance to the boundary (radius of the circle) at each shock point can be interpreted as the time of formation of the shock in a wave propagation scheme where waves are sent from the boundary. This dynamic interpretation of the shock trajectory allows us to associate a direction of flow, an instantaneous velocity (derivative with respect to radius), and an acceleration to each shock point. The combined geometric (curvature of axis) and dynamics (acceleration along axis) represented in a hierarchical graph is the shock set representation. In summary, shocks are medial axis points augmented with with the notions of speed, direction, type, label, grouping, and a hierarchy of these groups. For details see (Kimia et al., 1990, 1995; Siddiqi and Kimia, 1996; Tek et al., 1997).
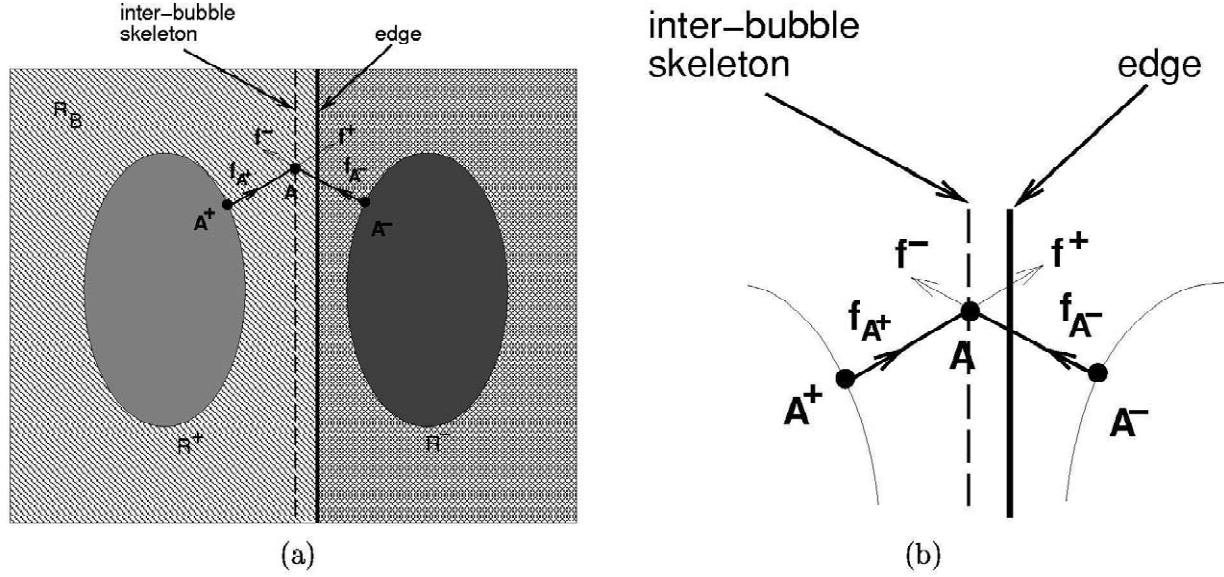
Fig. 9. A sketch of how the skeleton allows for the coupling of forces between nearby seeds, R^+ and R^-, e.g., coupling of A^+ and A^- through A. Note that f_{A^-} and f_{A^+} are the *local* forces at the boundary points A^+ and A^-, respectively, while f^- and f^+ are the *predicted* forces at the skeleton/shock point A due to predicted grown seed regions corresponding to R^+ and R^-. The net force at A^- is computed based on these four forces, Eq. (9).

$$f_{A^-} = \log(P(A^-|\alpha^-)), \tag{5}$$

where P is typically $(1/\sqrt{2\pi\sigma^2})\, e^{-(I(x,y)-\mu)^2/(2\sigma^2)}$, $I(x, y)$ is the intensity at A^-, and μ and σ are the mean and standard deviation of the adjoining region, respectively. Note that the range of the values of f_{A^-} is $(-\infty, 0)$, and is rescaled to ensure numerical stability of the evolution. If the two regions R^+ and R^- are infinitesimally close, then the point A^+ and A^- directly compete for the possession of A, as in region competition, by the net force, $f_{net} = f_{A^-} - f_{A^+} = \log P(A|\alpha^-) - \log P(A|\alpha^+)$. However, when the two regions are distant, they do not interact in the region competition approach (Zhu and Yuille, 1996). Our proposal then is to use the skeleton for (i) a long distance competition by coupling pairs of points; and (ii) at the same time reduce the effect of local competition as the distance between A^+ and A^- is increased. Thus, we define the net force representing competition at a distance by

$$f_{local} = f_{A^-} - \lambda(d)f_{A^+}, \tag{6}$$

where λ is a monotonically decreasing function of the distance between A^- and its competitor A^+, along the path going through the skeletal point A, consisting of two segments A^-A and AA^+ of length $2d(A^-, A)$ (Fig. 9(b)). We use

$$\lambda(d) = \frac{1}{\sqrt{2\pi\sigma_d^2}} e^{-d(A^-, A^+)^2/2\sigma_d^2}. \tag{7}$$

Thus, when the regions are adjacent, f_{local} represents region competition as in (Zhu and Yuille, 1996), but when the regions are moved apart, the local competition is gradually reduced to zero, such that $f_{local} = f_{A^-}$, which

represents traditional region growing, but implemented in a curve evolution framework.

Second, global competition is based on a comparison of the end points of the growth path, which in turn requires a comparison of which region R^+ or R^- should eventually enclose A, namely, a numerical comparison of the statistical forces, $f^- = \log P(A|\alpha^-)$ and $f^+ = \log P(A|\alpha^+)$. This global interaction can only be valid at a distance and thus must be modulated by $(1 - \lambda)$. Thus, the long range competition of the regions R^+ or R^- for A is defined as

$$f_{long\_range} = (1 - \lambda)(f^- - f^+), \tag{8}$$

where λ is defined in Eq. (7). The term $(1 - \lambda)$ is nearly 1 when regions are sufficiently distant (with respect to σ), but as regions draw closer to each other, the effect of long-range coupling is reduced to zero.

The total statistical force is then defined as a combination of the local and long range competition as

$$f_{stat}(A^-) = f_{local} + f_{long\_range}$$
$$= f_{A^-} - \lambda f_{A^+} + (1 - \lambda)(f^- - f^+). \tag{9}$$

In this way, the *local competition* (between local forces f_{A^+} and f_{A^-} on the boundaries of R^+ and R^-) controls the movement of A^- when the regions are close to each other. When the two regions are adjacent, $\lambda = 1$, and the net force is $F_{A^-} = f_{A^-} - f_{A^+}$ as in region competition. On the other hand, the *long range*, *predicted competition* between the statistical force f^- of A belonging to R^- and f^+ of A belonging to R^+ modulates the movement of A^- when the regions are far away. If the skeletal point is more likely to belong to one region compared to the other, based on current region parameters, the former region should grow

faster than the latter to capture the skeletal point. We emphasize that the inter-seed skeleton evolves as the seeds evolve. As an example, in Fig. 9, the skeletal point A is more likely to be part of region R^+ and hence the long range force coupling effectively slows A^- down, allowing A^+ to 'catch up', thus 'symmetrizing' the regions R^+ and R^- with respect to the edge.

In addition to the statistical force, we include a boundary smoothing term as in region competition (Zhu and Yuille, 1996), which is curvature dependent, so that the final force is

$$F = (\alpha F_{\text{stat}} + \beta F_{\text{smooth}}), \tag{10}$$

where $F_{\text{smooth}} = \kappa$, the curvature of the boundary.

3.3. Handling many-to-one skeletal coupling

While generally the coupling between two boundaries is one-to-one, there are cases where this breaks down, and where the evolution is no longer well-defined. The first case is a generic case when three points A^-, A^+, A^{++} couple through a skeletal point (junction) A (Fig. 10(a)), i.e. $\mathscr{S}^{-1}\mathscr{S}(A^-) = \mathscr{S}^{-1}A = \{A^-, \ A^+, \ A^{++}\}$. In this case three boundaries compete to capture this skeletal point. Thus, we use the average of two forces as the competing forces for the third, i.e.,

$$f_{\text{stat}}(A^-) = f_{A^-} - \lambda \left(\frac{f_{A^+} + f_{A^{++}}}{2} \right)$$
$$+ (1 - \lambda) \left(f^- - \frac{f^+ + f^{++}}{2} \right). \tag{11}$$

Note that four boundary points on different seeds do not generically share a shock point (Giblin and Kimia, 1999).

Second, when a segment of one of the boundaries is exactly circular that entire segment collapses to a single skeletal point. This is not a generic occurrence in the space of all curves, but we are making their occurrence generic, since we approximate each boundary segment within a pixel with a circular arc. However, since we are concerned with inter-region skeletons, this is a problem only when the skeletal point happens to be equidistant from another boundary as well (Fig. 10(b)). This is not a generic occurrence. When such cases do occur during the course of the evolution, it is handled by an averaging step similar to the junction case. Specifically, the circular arc segment is sampled using ENO points A_i, $i = 1, \ldots, n$ (see Section 4.1 for details on ENO points) and the forces are computed as

$$f_{\text{stat}}(A^-) = f_{A^-} - \lambda \left(\frac{\sum_{i=0}^{n} f_{A_i}}{n} \right) + (1 - \lambda) \left(f^- - \frac{\sum_{i=0}^{n} f_{A_i}}{n} \right). \tag{12}$$

Observe that since the extent of the circular arcs does not exceed a pixel, forces are generally continuous over such a range, and averaging is meaningful.

Finally, the singularities of one boundary can couple with many points on another boundary (Fig. 10(c)). Since our evolution process uses curvature smoothing, sharp singularities are smoothed out. However, in practical implementation, due to discretization high curvature points act like singularities, and we have to deal with the subsequent many-to-one coupling. When this occurs, we use averaging of forces as before to define the evolution process.

For a formal discussion of the special configurations of the skeleton under a family of deformation see (Giblin and Kimia, 1999). In numerous computer simulations we have observed that the many-to-one mapping except at the generic case of a junction occurs infrequently, and the averaging step described in Eq. (12) is rarely used. In
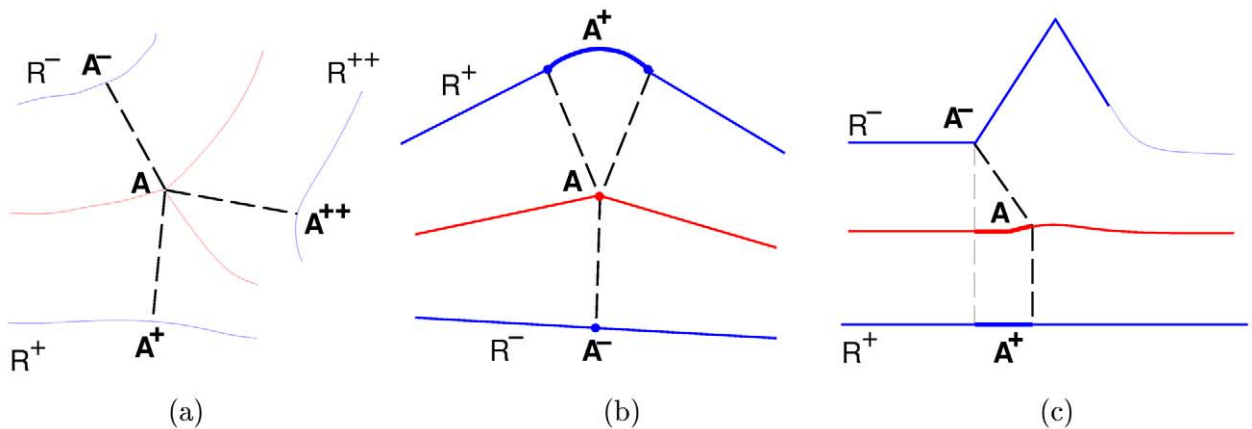

Fig. 10. This sketch shows the three cases where many-to-one mappings can occur in coupling two boundaries (blue) using the inter-seed skeleton (red). (a) Three isolated points A^-, A^+, A^{++} from boundaries of three distinct regions R^-, R^+, R^{++} can give rise to skeletal point (junction). In this case, A^- has to be coupled with two points from two distinct boundaries. (b) A portion of the boundary of region R^+ is circular, and this segment corresponds to a point on the skeleton leading to a many-to-one coupling. (c) In this case the boundary of region R^- has a singularity at point A^- and hence corresponds to many points on the skeleton leading to one-to-many boundary coupling.
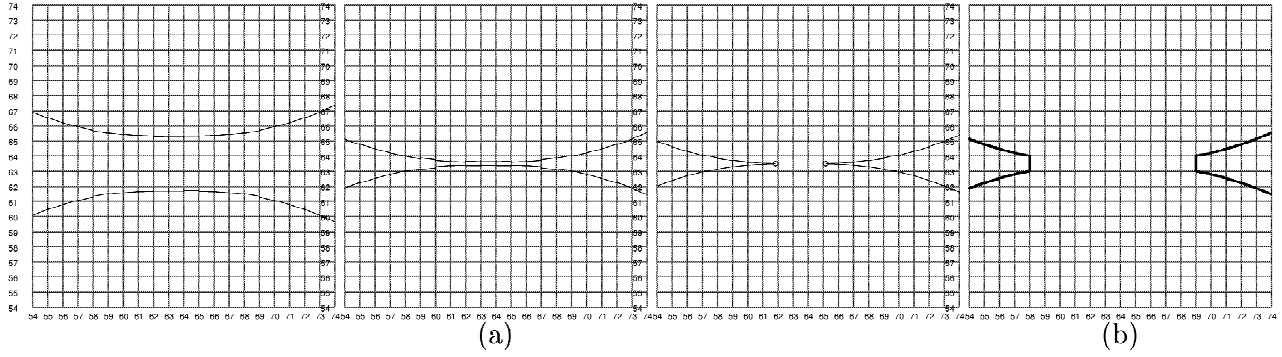
Fig. 11. (a) This figure from (Siddiqi et al., 1997) illustrates the reconstruction of curves in the presence of nearby curves from the discrete grid, using ENO interpolation. (b) In contrast, bilinear interpolation fails in the presence of nearby curves shown more accurately by the middle figure in (a).

addition, the many-to-one mapping in its rare occurrence is not stable over time, and can affect one or a few iterations. Moreover, what happens in the duration of one or a few iterations of the evolution does not typically affect the final results.

4. Implementation of SCDM

We now discuss the details of implementing the framework described in Section 3. Several components are necessary to complete the implementation: (i) initialize seeds, (ii) characterize statistical properties of the seeds, (iii) compute inter-seed skeletons, (iv) couple boundary points through the inter-seed skeletons, (v) evolve seeds, and (vi) compute subpixel forces.

We first note that the evolving boundary must be represented at subpixel resolution. As the authors note in (Zhu and Yuille, 1996), the implementation of region competition algorithm has discretization drawbacks and would benefit greatly from a subpixel implementation. Also, the approach proposed here in modulating growth is critically dependent on the reliable, accurate, subpixel detection and representation of the skeleton. In particular, computation becomes intricate as two seeds approach each other and share a pixel's area (Fig. 11) Observe that a

subpixel implementation of region competition requires regions to be adjacent, thus constantly operating in a mode where more than one curve is present within a pixel (Fig. 12). This requires the reliable identification of the subpixel boundary and computation of the forces at subcell points. For the subpixel boundary within a pixel, we use a Piecewise Circular (PC) approximation (Siddiqi et al., 1997), which can be extracted accurately from the distance transform surface used in the curve evolution approach. The forces are computed at subcell locations by a non-uniform ENO interpolation technique (Osher and Shu, 1991; Siddiqi et al., 1997), Appendix B.

Second, SCDM is implemented in a curve evolution framework which the allows for subpixel movement of the seed boundary, as well as for the existence of multiple boundaries within a pixel, while maintaining a pixel-based computation based on the discrete grid, as described below. In this framework the curve (explicit representation) is embedded as the zero level-set of an evolving surface ϕ (implicit representation). This allows the boundaries of regions to move consistently in a subpixel fashion on a discrete grid. Formally, let the initialized seeds or their growth at any point be represented by $R = \cup_{i=1}^{N} R_i$ where each region R_i is bounded by $C_i = \partial R_i$. Each region evolves by a statistical force F_{stat}, defined in Eq. (9), and smoothing force F_{smooth} as
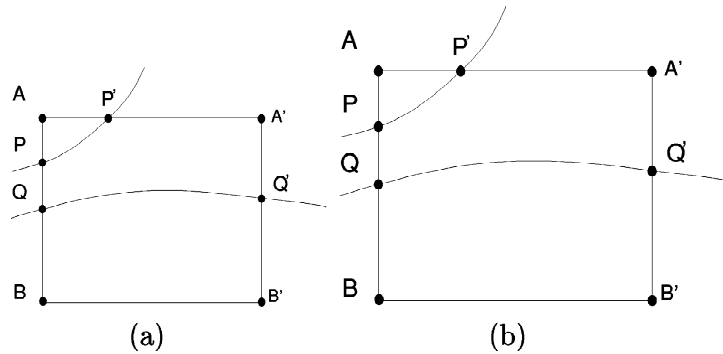

Fig. 12. This figure illustrates two cases where more than one curve is present within a pixel. Using ENO interpolation, the subpixel curve samples P, P', Q, and Q' can be accurately identified. Observe that ENO interpolation implies that P is only affected by the changes in A and Q is affected by changes in B, even if in this geometric arrangement A is closer to both P and Q.

$$\frac{\partial C_i(x, y)}{\partial t} = (\alpha F_{\text{stat}} + \beta F_{\text{smooth}})\vec{N}. \tag{13}$$

This is achieved by evolving the embedding surface ϕ using the Osher–Sethian formulation (Osher and Sethian, 1988)

$$\frac{\partial \phi(x, y)}{\partial t} = (\alpha F_{\text{stat}} + \beta F_{\text{smooth}})|\nabla\phi|, \tag{14}$$

where $\phi(x, y) = 0$ gives the curve C_i; we use the distance transform to generate ϕ from C_i. Table 1 gives the top-level sketch of the algorithm for the implementation of SCDM. The different components of that algorithm are described in the following subsections.

4.1. A subpixel method to construct explicit curve representation from implicit surface representation

As noted above, a subpixel implementation of SCDM requires the reliable identification of the subpixel boundary (zero level set) from the embedding surface. For *isolated* and *smooth* curves, bilinear interpolation of the discrete surface values leads to a reliable identification of the boundary, i.e., the zero level set. However, when *discontinuous* or *multiple* distinct curves exist within a pixel (Fig. 12), standard interpolation smoothes out the discontinuities and merges nearby curves, leading to large computational errors in the recovery of each curve (Siddiqi et al., 1997), (Fig. 11). Specifically, consider two contours that are within the square area defined by four pixels (Fig.

Table 1
Top-level algorithm for the implementing SCDM

Inputs:

    1. Gray-scale intensity image which is to be segmented

    2. Initialized seeds.

1. Initialize:

    1.1. Construct an implicit representation of the initialized seeds via the distance transform.

    1.2. Compute a subpixel interpolation (ENO) of the input gray-scale image. (Appendix B).

2. Find the subpixel forces at each grid point adjacent to evolving boundary: (Section 4.2)

    2.1. Find the explicit representation of evolving boundary from its implicit representation. (Section 4.1).

    2.2. Compute the inter-seed skeletons. (Section 4.3).

    2.3. Compute the statistics (μ, σ^2) for each seed.

    2.4. Find force at each point P on the discrete grid point adjacent to the evolving boundary.

        2.4.1. Find the closest point A^- on the seed boundary to P.

        2.4.2. Couple the boundary point A^- with its corresponding point(s) A^+. (Section 4.3).

            i. Find the shock point A corresponding to A^-.

            ii. Find other boundary point (or points) A^+ corresponding to shock point A.

        2.4.3. Compute forces at A^-.

            i. Compute the image intensity distribution along the normals from A^- and A^+ to skeletal point A. (Section 4.4)

            ii. Compute coupled statistics force. (Equation 9).

            iii. Compute the curvature at the boundary point A^-.

            iv. Compute the total force at A^- and use as force at P. (Equation 10).

3. Solving the PDE: (Section 4.5)

    3.1. Propagate forces computed from the explicit representation (zero level set) to other level sets of the implicit representation.

    3.2. Compute the ENO derivatives of surface.

    3.3. Evolve the surface by updating the PDE. (Equation 15)

4. Convergence: (Section 4.6)

    4.1. If the change in the surface is significant go to step 2, else stop and return boundaries and skeleton.
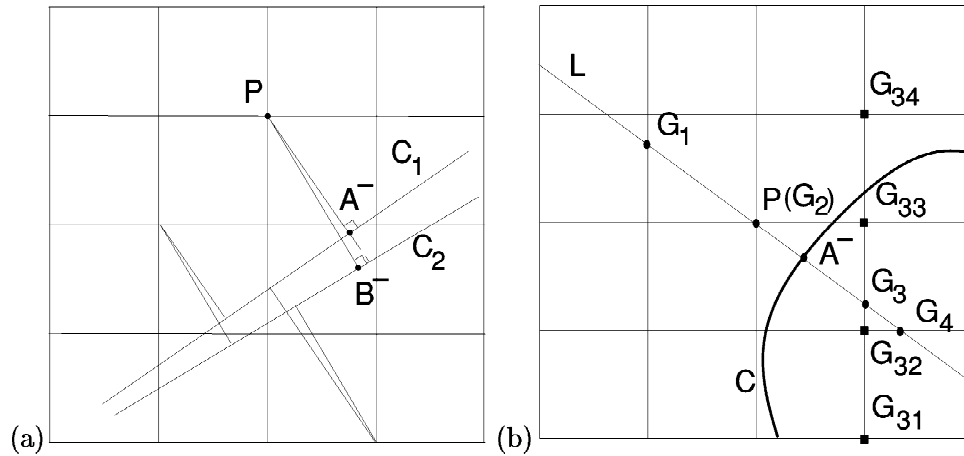
Output:

    1. Final boundaries and their skeleton.

Fig. 13. (a) The subpixel deformation of a curve (C_1 is the original curve and C_2 is the deformed curve) can be accomplished via updating its embedding surface at each grid point P, by the subpixel movement of the closest point on the curve A^-. (b) The subpixel computation of the force at P requires a two stage ENO interpolation; (i) along the line (L) normal to the curve (boundary of the seed) using points $\{G_i\}$ (solid dot) and (ii) the computing the force for each G_i along either the horizontal or vertical directions. For example, the force at G_3 is computed via an ENO interpolation along the vertical grid line using the grid points $\{G_{3i}\}$ (solid squares).

12(a)). When these contours are represented as the zero level set of a single implicit surface, where the surface is positive for the areas within each closed contour, then all the four surrounding grid crossings A, A', B and B' must necessarily be positive. Standard interpolation cannot recover either curve, since the results of the bilinear interpolation for four positive numbers cannot be zero anywhere in the area between them. Siddiqi et al. (1997) proposed a method (GENO, geometric ENO) based on the Essentially Non-Oscillatory (ENO) Interpolation methods (Osher and Shu, 1991), which is summarized in Appendix B. ENO interpolation is effectively a one-sided interpolation, that can recover adjacent contours (within a pixel) accurately from distance transform surface via a three step process: (i) ENO interpolation along horizontal and vertical discrete grid lines to recover curve samples along grid lines. Since this interpolation is one-sided, both curve samples on grid lines P, P', Q, and Q' are recovered provided a third curve is not present in the two pixel vicinity; (ii) tracing these samples by tracking the interior, and (iii) a GENO interpolation of grid crossing samples of the curve to represent the curve inside the pixel using a piecewise circular (PC) model (Banchoff and Giblin, 1994; Siddiqi et al., 1997).

The accuracy of reconstruction of PC curves from their distance transform using this approach is examined in this paper in Appendix A. First, it is shown that an arbitrary isolated line can be recovered exactly using a second order ENO interpolation[4]. Second, it is shown that for an isolated circular curve, the second order ENO reconstruction errors are extremely small. Several simulations show that the reconstruction errors for other shapes are very

small, and the method is suitable to recover the explicit representation of curves from their implicit representation.

4.2. Evolving the implicit surface representation by the desired evolution of explicit curve representation

We now address the issue of whether changes in the discrete implicit surface can lead to reliable, accurate and predictable subpixel changes in the contour, and vice versa. The surface should be transformed in such a way that the modified surface is the distance transform of the modified curve. Appendix A shows that the PC reconstruction from the distance transform satisfies this accuracy requirement. Consider how the surface point at P needs to be updated when the curve changes from C_1 to C_2 (Fig. 13(a)). Let A^- and B^- be the closest points from P on the curves C_1 and C_2, respectively. Hence the surface at P should be updated by $|PB^- - PA^-|$, i.e., the difference in distance from P to the two curves. For small movements this can be approximated by B^-A^-. This requires identifying the closest point A^- on the PC curve to the each pixel P, and computing the force at subcell point A^- (Fig. 13(a)).

The identification of the closest point on the subpixel boundary A^- is straightforward from the CEDT method, Fig. 13(b). We have examined the effectiveness of this update scheme in a few cases where a pre-defined force is specified at each subpixel point. The results are accurate and robust, as illustrated by accurate rotations of a straight line and a circle (Fig. 14(a) and (b)).

4.3. Computing inter-seed shocks and coupling of boundary points using the shocks

The computation of the subpixel force at A^- relies on coupling the boundaries via the intermediate skeleton

---

[4]We typically use second order ENO interpolation, but we are not restricted from using higher order schemes when necessary.
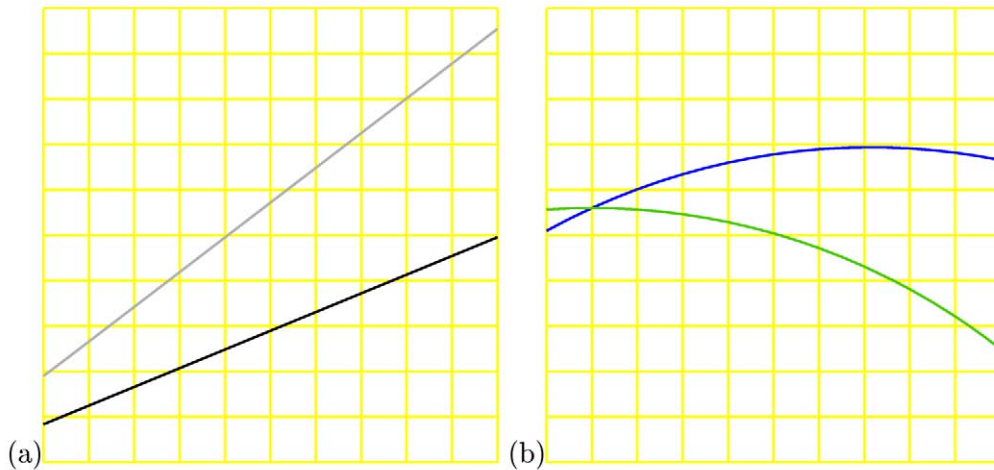
Fig. 14. This figure shows the results of deforming a curve by updating its embedding surface at each grid point by the force at the closest point on the subpixel boundary. In this example, the subpixel force is specified by a rotation field: (a) the rotation of a line (grey line is the original and the black line is the rotated line); (b) the case of rotating an arc (blue is the original and green is the rotated). In each case the exact and the computed curves after rotation are plotted, but the exact curve cannot be viewed as at this resolution they are visually identical. These results illustrate the accuracy of the approach.

(Section 3), and requires: (i) computing the inter-seed skeleton, (ii) identifying the point A on the skeleton, corresponding to A^- (Fig. 15), (iii) identifying the point A^+ on another seed that has the same skeletal point A, Fig. 15, and (iv) accurately computing the image intensity and image derivatives along the normal to the seed boundaries at A^-, A^+ and A, so that forces may be computed.

The inter-seed skeletons (shocks) are computed via the wave-propagation technique based on the contour based distance transform (CEDT) (Tek and Kimia, 1998; Tek et al., 1997). The input to the CEDT is a boundary model, which in our case is the evolving seed boundaries. The CEDT propagates distances and orientation from the boundary models such that it is immediately clear at each

point on the wavefront how far the wave has traveled, what the direction of propagation is, and which point on the original boundary gave rise to it. This allows for the accurate and numerically efficient computation of the inter-seed skeleton. In addition, the approach of computing the skeleton using the CEDT directly identifies the skeletal point corresponding to a boundary point and vice versa. This makes the coupling of the boundary points A^- and A^+ straightforward.

4.4. A two-stage nonlinear interpolation of image intensity

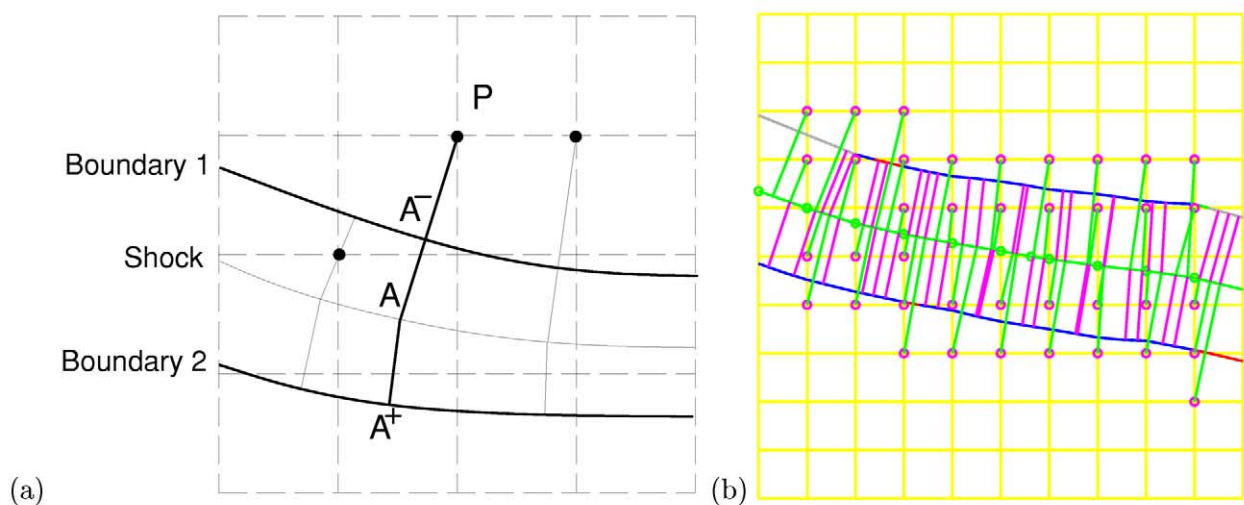We now discuss how the image intensity and image

Fig. 15. (a) An illustration to show how the surface is updated. Dotted lines represent the grid lines whose crossings are the pixel center. The solid lines are the boundaries of each deformable model. To update the surface at P, the closest point on the boundary (from P), A^-, is found. A^- is coupled with the point, A^+ on the boundary of the adjacent region via the inter-boundary skeleton (shock) point A. (b) Actual coupling in a cropped portion of a simulation in a real image. The boundaries of each region represented in implicit form as the zero crossing of a surface represented on the pixels only. Observe (i) the coupling between two deformable contours through the skeleton and (ii) the connection between the pixels (solid dots) and the deformable contours.

derivatives along the normal to the seed boundaries are computed. To avoid the pitfalls of bilinear interpolation, we use a two stage ENO interpolation to compute the image intensity distribution and image derivatives along a line segment with arbitrary orientation. In a pre-processing step, ENO interpolation is used to compute the intensity distribution along every horizontal and vertical gridline of the image. Let L be the line segment along which we want to compute the intensity distribution and image derivatives. In the first stage, the intersections $\{G_i\}$ of this line segment L with the gridlines (Fig. 13(b)) are found. The image intensity values at these points $\{G_i\}$ are computed via a second order ENO interpolation along the corresponding discrete grid line. For example, the image intensity at G_3 in Fig. 13(b) is computed via an ENO interpolation along the vertical grid line using the points $G_{31}, G_{32}, G_{33}, G_{34}$. In the second stage, a second ENO interpolation along the line segment L using the values at the G_i is done to obtain the intensity distribution along the L, which specifies the intensity distribution and derivatives at any point along L. We use this interpolation technique to compute the image intensity at the closest point on the subpixel boundary A^- along the normal to the boundary, which in turn determines the statistical force at P (Eq. (5)).

4.5. Updating the PDE for surface evolution

The surface update equation (Eq. (14)) can be rewritten as

$$\frac{\partial \phi(x, y)}{\partial t} = \alpha F_{\text{stat}} |\nabla \phi| + \beta \kappa |\nabla \phi|$$

$$= \alpha F_{\text{stat}} (\phi_x^2 + \phi_y^2)^{1/2}$$

$$+ \beta \left(-\frac{\phi_{xx} \phi_y^2 - 2\phi_{xy} \phi_x \phi_y + \phi_{yy} \phi_x^2}{(\phi_x^2 + \phi_y^2)^{3/2}} \right) (\phi_x^2 + \phi_y^2)^{1/2}$$

$$= \alpha F_{\text{stat}} (\phi_x^2 + \phi_y^2)^{1/2}$$

$$+ \beta \left(-\frac{\phi_{xx} \phi_y^2 - 2\phi_{xy} \phi_x \phi_y + \phi_{yy} \phi_x^2}{(\phi_x^2 + \phi_y^2)} \right). \quad (15)$$

The embedding surface derivatives ϕ_x, ϕ_y, ϕ_{xx}, ϕ_{yy}, ϕ_{xy} are computed using an upwind scheme using ENO interpolation. As discussed in Section 4.2, the surface is updated by the force at the closest point on the evolving boundary, F_{stat}.

4.6. Convergence of the algorithm

Finally, we address the issue of convergence, a strong point of this algorithm. Lack of convergence is one of the main drawbacks of traditional level-set-based segmentation methods. Recall that when seeds get close to one another in SCDM, the evolution is governed by local force competition (Eq. (6)). In addition, if the seeds are close to the object boundaries, the local forces cancel each other

and these seeds stop moving. Thus, SCDM through the use of skeletal coupling resolves the convergence problems of traditional deformable models (Fig. 4). Eventually, all the seeds capture their respective object boundaries, and the evolution is stopped. In practice, the evolution is stopped when the net forces at all boundary points fall below a certain threshold. We chose this threshold to be 0.1 pixel. We have observed from numerous computer simulations that the method is insensitive to the choice of their threshold. Fig. 16 illustrates the convergence properties of SCDM. Observe that in the intermediate iterations portions of the bone outline has been captured by the evolving seeds. Those converged seeds do not move in the further iterations that are necessary for other portions of the seeds to capture their object boundaries.

5. Results, evaluation and discussion

In this section we present the results of applying the SCDM method to synthetic images (Fig. 17), and to carpal bone segmentation (Figs. 16, 18 and 19), especially to illustrate the performance of SCDM in the problem areas listed in Section 1, namely, gaps and weak edges, diffused edges, bone texture, and narrow inter-bone spaces. In order to segment an image using SCDM, one has to choose the user-specified parameters, namely, σ_d in the definition of λ (Eq. (7)), and parameters α, β in the definition of the net force (Eq. (10)). The values for these parameters were chosen based on experimentation. We chose $\sigma_d = 1$, so that the local force competition dominates when the competing regions are within a few pixels. The parameters α, β are chosen based on numerical constraints and on our experience with this type of PDE (Kimia et al., 1995) to have values $\alpha = 0.3$ and $\beta = 0.1$. All parameters were fixed after an initial fine-tuning stage. We have observed that the performance of the algorithm is not sensitive to moderate changes in these values.

Fig. 17 compares the results of applying seeded region growing, region competition and SCDM for segmenting some synthetic images. Fig. 17(a) illustrates the application to an image with a weak edge. SCDM captures the edge well, as does region competition. However, seeded region growing leaks through the gap as it does not incorporate geometric information. Fig. 17(b) illustrates a case where the seeds are initialized across the boundaries. Since SCDM and region competition allow for recovery from errors, both of them capture the object. However, observe the jagged nature of the final segmentation, due to the discrete nature of the growth in region competition. Seeded region growing does not allow for regions to shrink, and thus fails to capture the object. Fig. 17(c) illustrates the results of applying SCDM when seeds are initialized asymmetrically. SCDM which implements long range competition captures the edge. Region competition delays competition until the seeds contact each other, and fails to
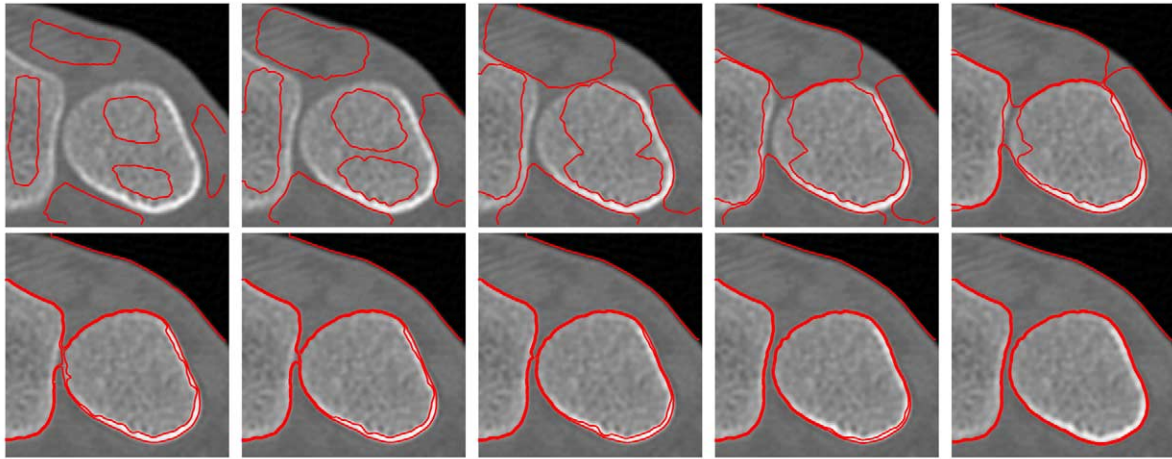
Fig. 16. Several iterations of SCDM on a cropped portion of a CT image of the carpal bones and key advantages of SCDM. First, observe that in the initial iterations, the 'soft-tissue seed' is far away from the inter-bone gap. Hence, the competition for the inter-bone gap is between the two 'bone seeds'. The bone seeds are slowed down considerably, but the seed of the bone on the left leaks out into the soft tissue. However, it is eventually pushed back by the 'soft-tissue seed'. Second, observe how SCDM correctly segments the cortical bone (bright intensity) at the lower right portion of bone on the right. The cortical bone belongs to the tail of distribution for the bone, and hence the evolving 'bone seed' is slowed down, but eventually captures the cortical bone. Finally, observe the convergence properties of the evolving contours. In the intermediate iterations (for example, leftmost image in bottom row), portions of the contours have already captured the bone outline. In further iterations that are needed to capture the inter-bone gap and high intensity cortical bone, these portions of the contours do not move.
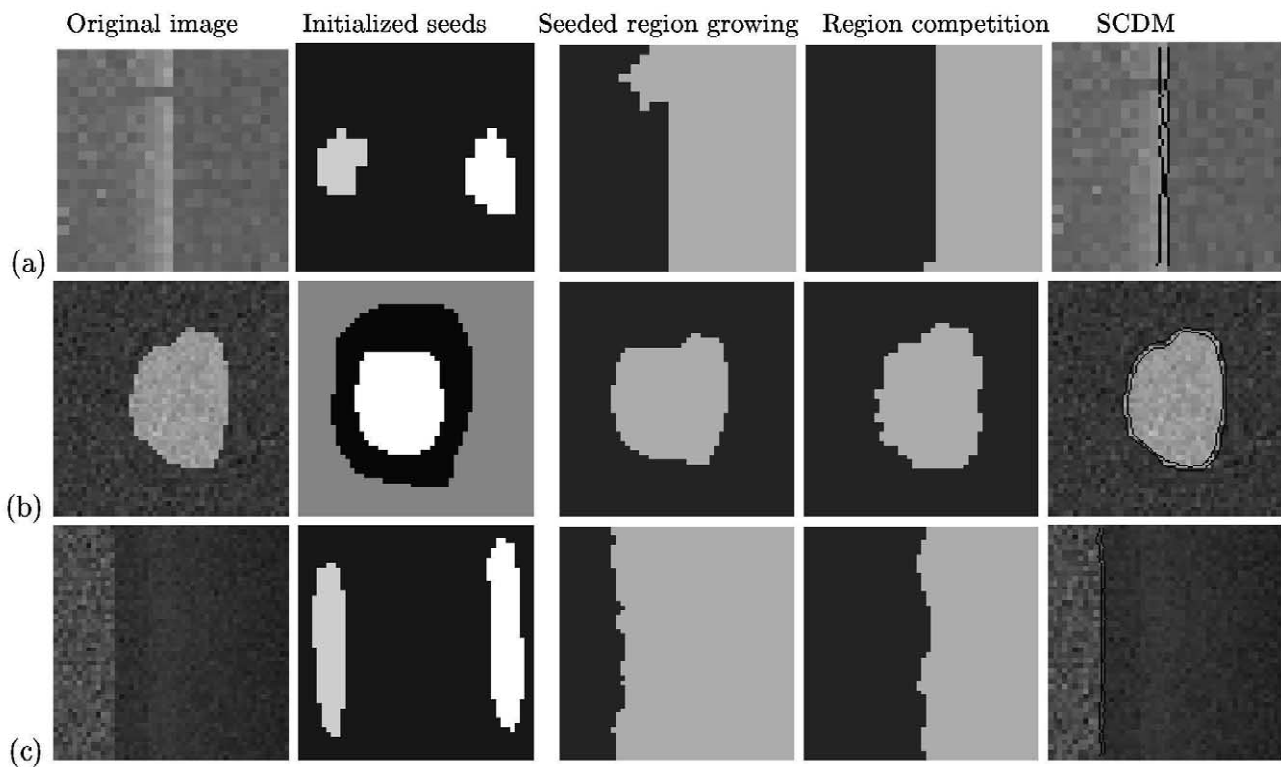

Fig. 17. In all rows left to right: Original image; initialized seeds; segmentation results using seeded region growing, region competition and SCDM. Note that in the examples of SCDM segmentation the converged boundaries for both regions are shown. The skeleton between these boundaries capture edge. (a) Comparison of segmentation results for an image illustrating the effect of a weak edge. Observe that in segmentation using seeded region growing, a seed leaks out through the gap, while region competition and SCDM captures the edge. (b) Compares the effect of seeds initialized on the boundaries. As there is no recovery from errors, seeded region growing fails, while region competition and SCDM captures the edge. However, observe the jagged nature of the final segmentation using region competition. (c) Compares the effect of asymmetrically initialized seeds, i.e., where one seed is one is closer to the edge than the other. While seeded region growing captures the edge it is not localized properly. Region competition fails here, as the seed on the left leaks out and loses its statistical character *before* competition begins. SCDM uses a combination of long-range and local competition and captures the edge.
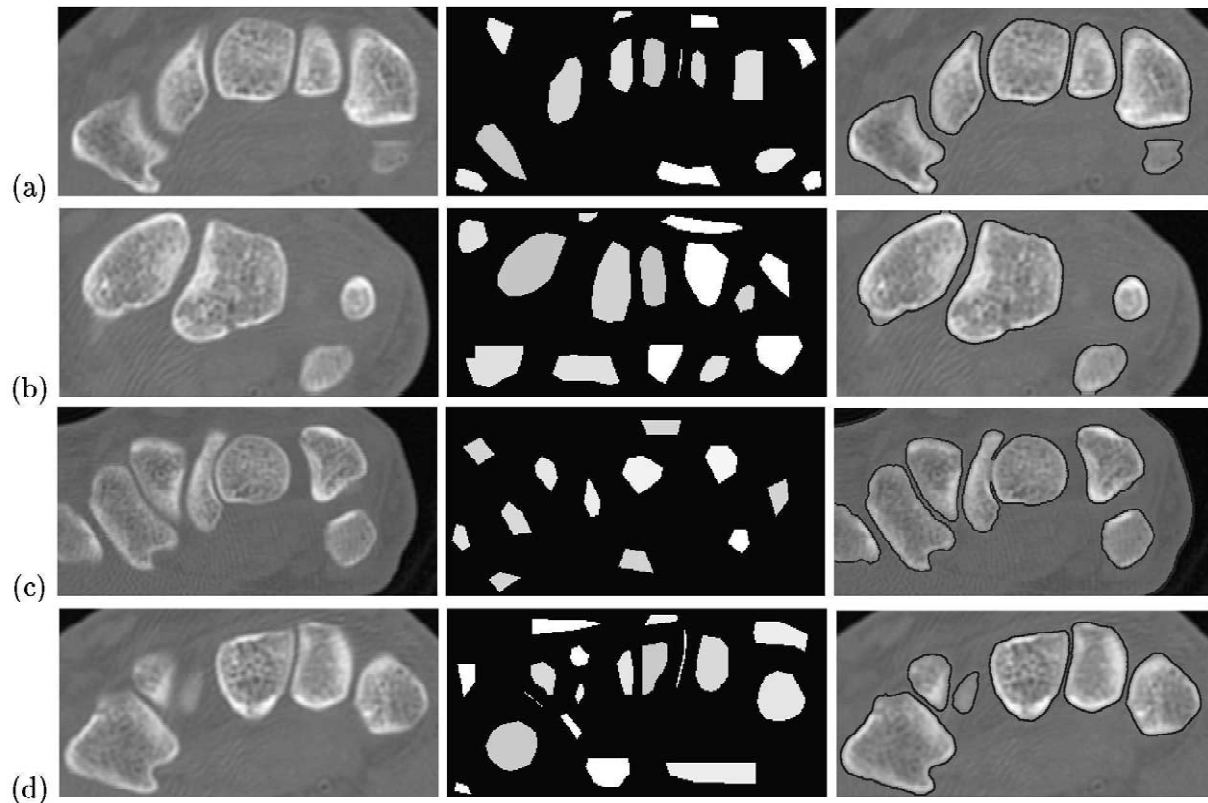
Fig. 18. Examples of carpal bone segmentation using skeletally coupled deformable models. Original image (left column), initialized seeds (middle column) and SCDM segmentation results (right column) are shown.

capture the edge. Seeded region growing captures the edge, but fails to localize it well. Thus, the combined local and global competition in SCDM combines the local competition of region competition and global competition of region growing.

Fig. 16 illustrates some key features of SCDM. Note that the inter-seed skeleton (predicted boundary) may be inaccurate due to the presence of a 'third' region in the middle or due to noise. If there is a third region in the middle, whose seed is remote from inter-seed skeleton of the first two seeds, the two competing seeds are both slowed down, allowing the third seed initialized for the third region to capture the region. This is illustrated in the segmentation of the inter-bone spacing in Fig. 16. Second, if the predicted boundary happens to be positioned on a 'noisy pixel', the forces are generally reduced due to a
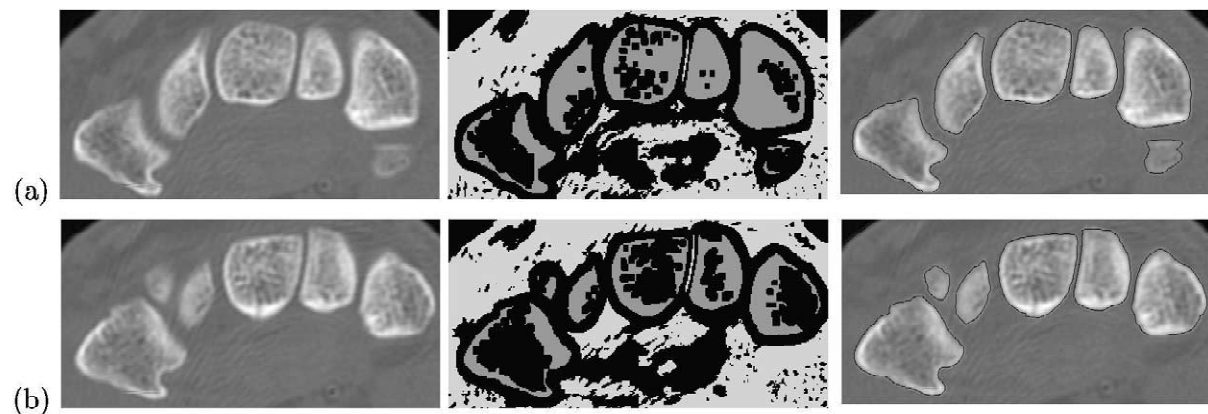

Fig. 19. Use of semi-automated seed initialization for carpal bones segmentation with skeletally coupled deformable models. Original image (left column), initialized seeds (middle column) and SCDM segmentation results (right column) are shown. Seeds in these cases were initialized semi-automatically using the EM algorithm, described in Section 5.

mismatch, thus slowing down each seed, but each by a different extent (generically). This evolves the skeleton out of the noisy region. Once outside this region the evolution will symmetrize the skeleton. Finally, SCDM works well in the presence of object boundaries whose neighboring intensities belong to the 'tail' of its statistical distribution. In this case, the seed for this object is slowed down in approaching the boundary, but eventually does converge on the object outline. This is illustrated in the segmentation of the cortical bone at the lower right portion of bone on the right in Fig. 16.

Fig. 18 shows the results of applying SCDM segmentation to a few representative CT images of the carpal bones. In these examples, seeds were initialized manually, one inside each bone and a few labeled as background in the soft tissue. Observe how SCDM captures the bone contours fairly well in these examples: it works well in images having a gap in the bone contour, e.g., in the lower portion of the second bone from the left in Fig. 18(b). It captures low contrast contours, e.g., the bone on the bottom right in Fig. 18(a). Observe also the growth into some of the narrow inter-bone spaces, e.g., Fig. 18(d). In addition, it does well in case of diffused edges, e.g., top of the first bone in Fig. 18(d). Also, note that the region-based statistical force allows growth in textured areas, as evidenced by all these examples. However, observe that SCDM does not enter some of the thinnest inter-bone spaces whose intensities are diffused. The inherent blurring in CT imaging causes the narrow inter-bone regions to have intensities similar to the bones, rather than soft tissue. This *partial volume effect* does not allow the deformable model representing the soft tissue to enter these diffused inter-bone spaces (Fig. 18(c)). Thus, we are required to initialize seeds at these very thin spaces using a specialized initialization scheme (see below). An explicit modeling of the CT blur such as used in (Laidlaw et al., 1998) could also be employed to resolve this issue.

We have shown that SCDM does well in capturing the bone outlines when seeds were initialized manually. However, initializing multiple seeds in every slice of a 3D volume is tedious. To alleviate this problem, we have examined the use of a semi-automatic statistical classification for seed initialization. Fig. 19 shows the results of SCDM for a few representative slices where we have used this semi-automatic technique for seed initialization. Assuming that the intensity values of the bones and soft tissue come from a mixture of Gaussian distributions, the EM algorithm (Dempster et al., 1977; Redner and Walker, 1984) was used to find the component means and variances. Then, the seeds for the bones and soft tissue were determined using intensity thresholding. In addition to these statistical methods, for narrow inter-bone regions which defy such characterization, the *logical linear n-line operator* (Iverson and Zucker, 1995) was used to initialize seeds in the thin inter-bone regions. Since multiple seeds

can appear in the background, these are identified with a common label so that they may merge.

We conducted a preliminary clinical evaluation of the carpal bone segmentation results of SCDM. This evaluation shows that the carpal bone segmentation by SCDM is clinically meaningful. In addition a comparison was made between the results using SCDM, seeded region growing, and region competition by two hand surgeons from Rhode Island Hospital. Two datasets (101 2D slices) were used for the validation study. All slices were segmented using all three methods, and the hand surgeons chose the best segmentation out of the three (without knowing which segmentation was used in each). The results are: SCDM was chosen as the best segmentation in 82 slices, seeded region growing in 15 slices, and region competition in 4 slices. Fig. 20(a) shows an example of slices where SCDM was chosen to be the best segmentation: SCDM and seeded region growing capture all the narrow inter-bone spaces, whereas region competition merges adjacent bones, and seeded region growing allows the soft tissue seeds to leak through narrow gaps (top of second bone from left, and bottom of fifth bone from left); it also fails to capture the low contrast bone on the bottom right well. Fig. 20(b) shows an example of a slice where seeded region growing was chosen to be the best segmentation. Seeded region growing captured all the bone contours, whereas SCDM merged two adjacent bones, the fourth and fifth bones from the left. Region competition also merged two bones. In addition, it failed to localize the edge properly in the bone at top right. Fig. 20(c) shows an example of a slice where region competition was chosen to be best segmentation. All three methods failed to capture all the inter-bone spaces in this case. However, region competition was rated the best because it best separated the bones on the left. Discussions with the hand surgeons also suggested that in these segmentations, SCDM generally localized the bone contours better. This validation study together with other experiments reveals that the main reason for the occasional failure of SCDM is due to the lack of growth into diffused inter-bone spaces. These occasional cases require modest user interaction but can in future be avoided by explicitly modeling the partial volume effect in SCDM (Laidlaw et al., 1998). We plan to do a rigorous validation study of carpal bone segmentation for SCDM in the future by a quantitative comparison using standard datasets.

We have reconstructed the 3D surface of carpal bones by stacking the segmentation of 2D slices (Fig. 21). The 2D contours of each carpal bone were extracted and all the contours corresponding to a particular bone were identified. The 3D triangulated surface model was then created using the NUAGES software developed by Geiger (1993). Observe that the bone shape as well as the joint spaces (typically one or two pixels and sometimes less than a pixel) are recovered well. The surfaces are then used for registration of bone motion.
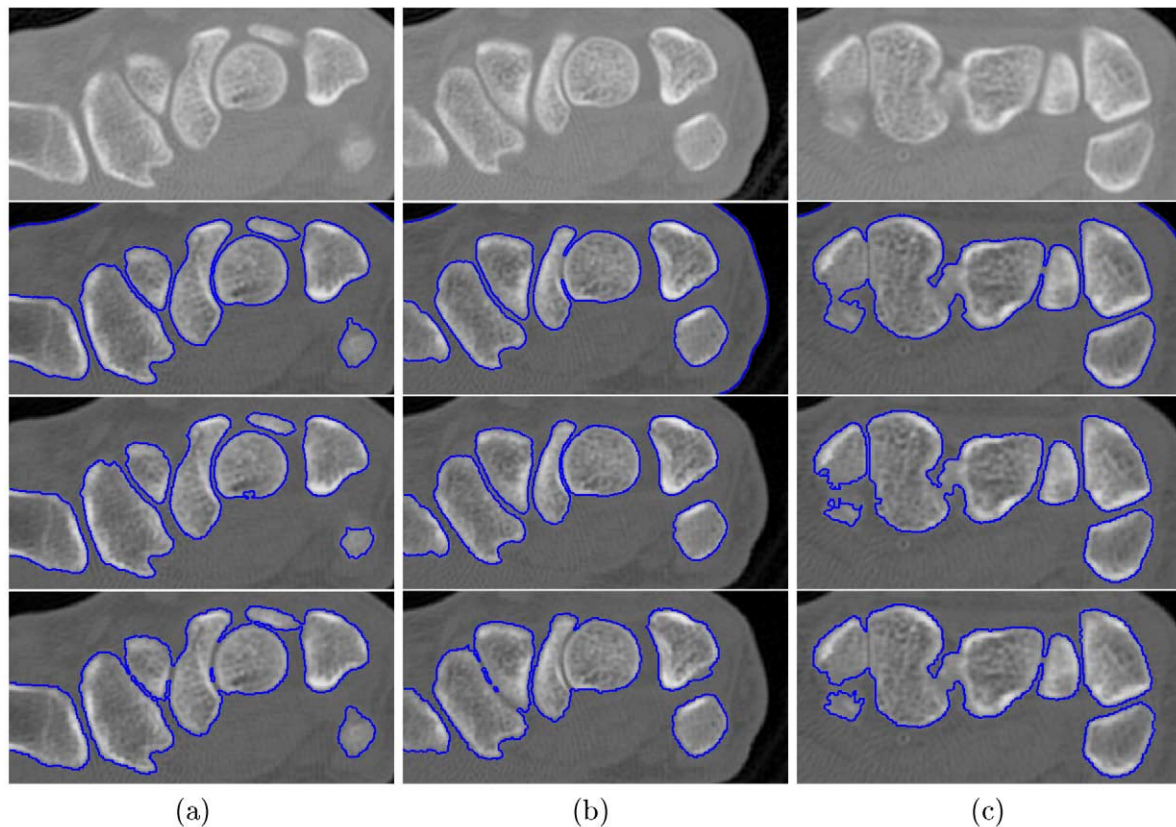
Fig. 20. Examples of CT image segmentations used in the clinical evaluation. Top to bottom rows: Original image, segmentation using SCDM, segmentation using seeded region growing, and segmentation using region competition. (a), (b) and (c) show examples of a slice where the hand surgeons chose the segmentation using SCDM, seeded region growing and region competition, respectively, to be the best.

There are a variety of potential clinical applications for the use of SCDM in segmenting the carpal bones. SCDM segments the carpal bones and the inter-bone joint spaces accurately so that second order differential measurements such as principal curvatures can be reliably made. This allows us to investigate the correlation between the shape of the bones and joint spaces, and the incidence of disease. Specifically, this segmentation allows for an investigation of the correlation between shape of the trapezium (the carpal bone at the base of the thumb), the thumb metacarpal and the thumb carpometacarpal joint space and the incidence of osteoarthritis. Using post-mortem stereophotogrammetry measurements of cadaver bones, Van Mow's group (Xu et al., 1998) have shown that there exist differences in the curvature maps of male and female trapeziums, and hypothesized that the differences in the shape of the trapezium makes the female population predisposed to osteoarthritis in the thumb carpometacarpal joint. Our segmentations are sufficiently accurate to allow for an in vivo measurement of the curvature maps, as indicated by the rather regular appearance of these second order measurements (Fig. 22). In addition, the local curvature differences can be extended to the concept of general shape differences: we intend to use the SCDM

segmentation to compute the 'average shape' of each carpal bone and the joint space, and then characterize the deviations from the average as typical or atypical.

From a technical standpoint, we plan to extend this work along several directions. First, we currently use the inter-seed skeleton as the collision boundary to mediate the long range coupling, effectively assuming that regions grow at constant speed. A better and more intuitive way to determine the collision point is to use a variable speed wave propagation, based on the current speed of regions. Second, the computation of the forces is another area of potential improvement. The basic idea of the approach is to determine the collision point of two growing seeds, compare the paths from the boundary of the seeds up to the predicted collision point, and grow the seeds based on the suitability of these paths. Currently, we use the end points of the paths, i.e., skeletal point and the boundary point, to determine the force. Instead, the forces along the path from the boundary to the predicted collision point can be integrated to determine the force. Due to the volumetric nature of the CT image data and the clinical applications involving 3D anatomical models, a 3D version of SCDM is necessary for its widespread clinical use. Hence, we plan to extend this 2D coupling to 3D, which will enable us to
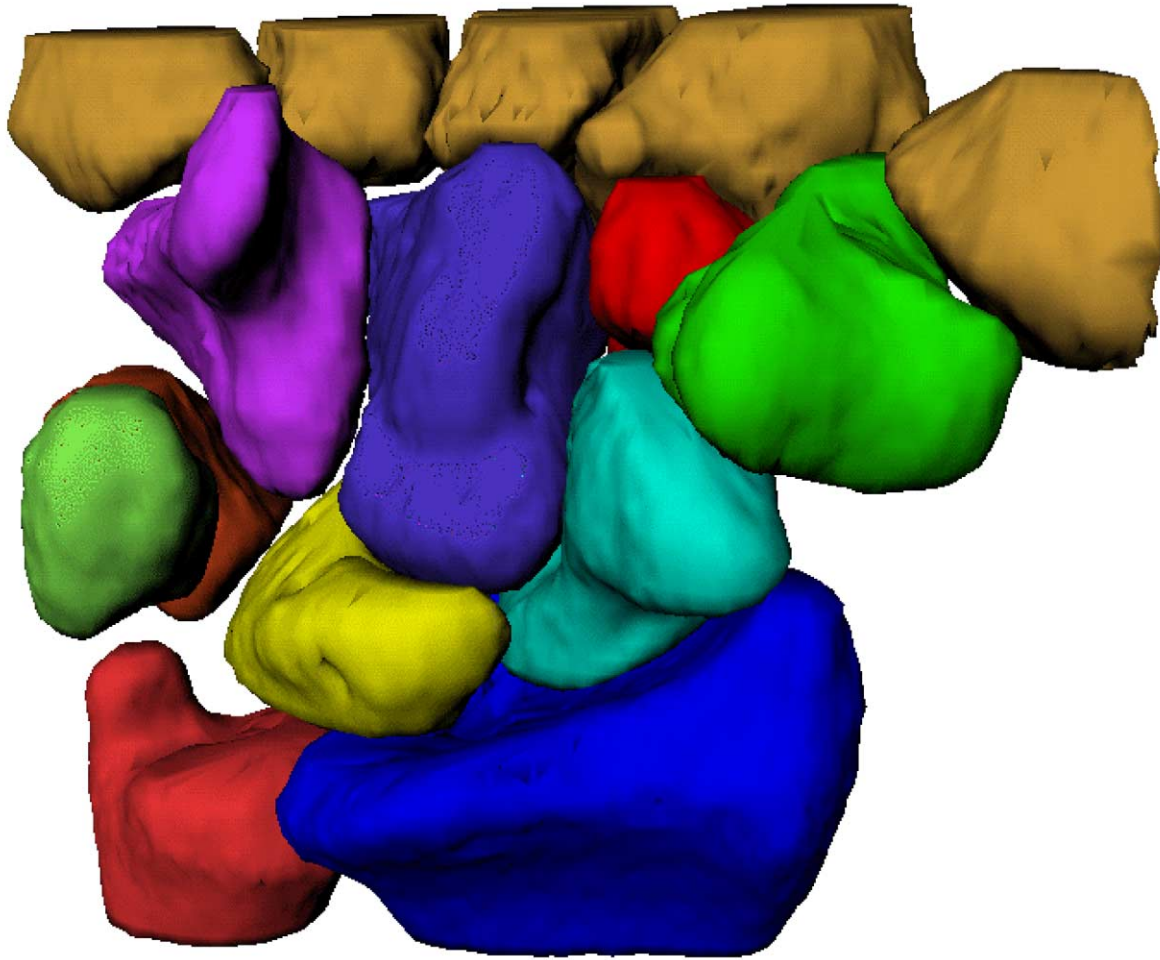
Fig. 21. The 3D visualization of the segmentations using SCDM that shows the relative placement of bones.

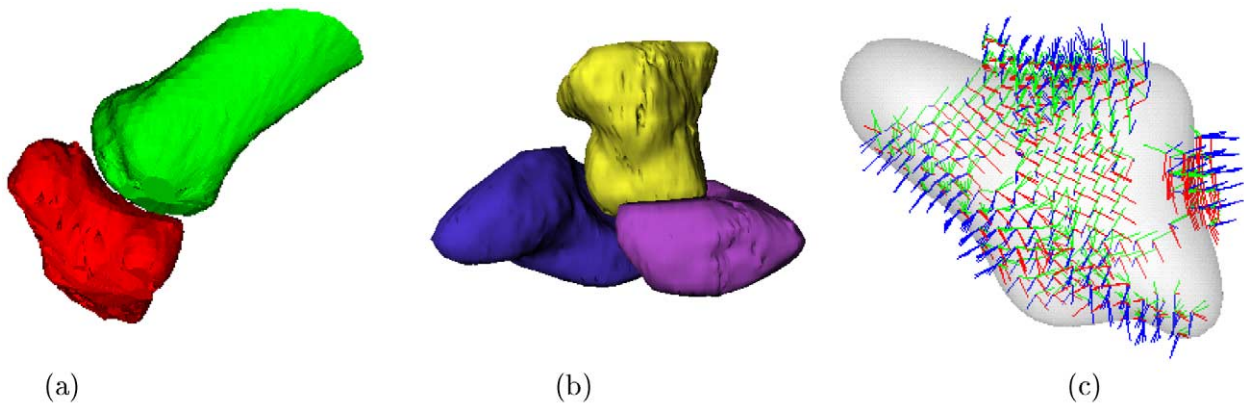(a)   (b)   (c)

Fig. 22. Some potential clinical applications for the use of SCDM in segmenting the carpal bones. Segmentation of (a) carpal bone trapezium and thumb metacarpal, (b) carpal bones scaphoid, lunate and capitate. Note that the bone shape as well as the joint spaces (typically one or two pixels and sometimes less than a pixel) are recovered well. This allows for the investigation of the correlation between bone shape and the occurrence of arthritis (Xu et al., 1998), which established a distinction in the local curvature map between male and female populations. The robust and accurate subpixel recovery of bone shapes allows for the computation of the local curvature map of the bone surface. (c) Principal directions (green and red vectors) and normal (blue vectors) of the carpal bone trapezium computed at points where it contacts other bones. In addition, questions pertaining to the average shape of each carpal bone and typical and atypical deviations can be addressed in the future based on these representations.

take advantage of surface continuity in all three directions. We plan to build on our experience on 3D level set based segmentation (Tek and Kimia, 1997) and 3D curvature-based evolution (Neskovic and Kimia, 1994). However, in addition to components already available or generalized to 3D, namely, (i) distance-transform, (ii) ENO, (iii) subvoxel tracing of the zero level-set, this extension also requires (i) an efficient computation and representation of skeleton in 3D and (ii) a method for coupling two points on the evolving surface. This technology is now emerging (Leymarie and Kimia, 2000, 2001).

6. Conclusion

In this paper, we have presented a segmentation method that combines the advantages of active contour models, region growing, and the global competition in seeded region growing as well as the local competition in region competition. The proposed method (SCDM) uses competition mediated by the inter-seed skeleton to modulate the growth of seeds. The skeletally mediated competition allows for long-range competition thus augmenting the region competition's local competition (Zhu and Yuille, 1996). The subpixel curve evolution based implementation of SCDM avoids the discretization drawbacks of current implementation of region competition (Zhu and Yuille, 1996). SCDM converges when the net forces at all boundary points are negligible. Thus, SCDM avoids the convergence problems of curve evolution methods (Malladi et al., 1994; Caselles et al., 1993; Tek and Kimia, 1995). Also, SCDM is insensitive to the initial placement of the seeds, a drawback of the active contour methods (Kass et al., 1988). However, in carpal bone segmentation, SCDM fails to capture some narrow, diffused inter-bone spaces. This is due to partial volume effect, and can be tackled by modeling the intensity as a continuous function, taking into account the inherent blurring of CT (Laidlaw et al., 1998).

Appendix A

Accuracy of recovery of explicit curve representation from implicit surface representation

The goal of this section is to examine how accurately an explicit representation of the curve $C(s) = (x(s), y(s))$ can be recovered from an implicit representation of it, $f(x, y) = 0$. Specifically, we focus on functions $f(x, y)$ derived as the signed distance transform of a contour, and where the recovered contour is modeled as piecewise circular (PC). We first consider recovery errors for an isolated straight line, and an isolated circular arc, and then consider the recovery error for other curves, e.g., parabola.

First, consider an arbitrary isolated line (Fig. 23(a)). The distance transform (DT) restricted to each horizontal or vertical line is linear, i.e., $A_i \tilde{A}_i$ is linear in i by geometric
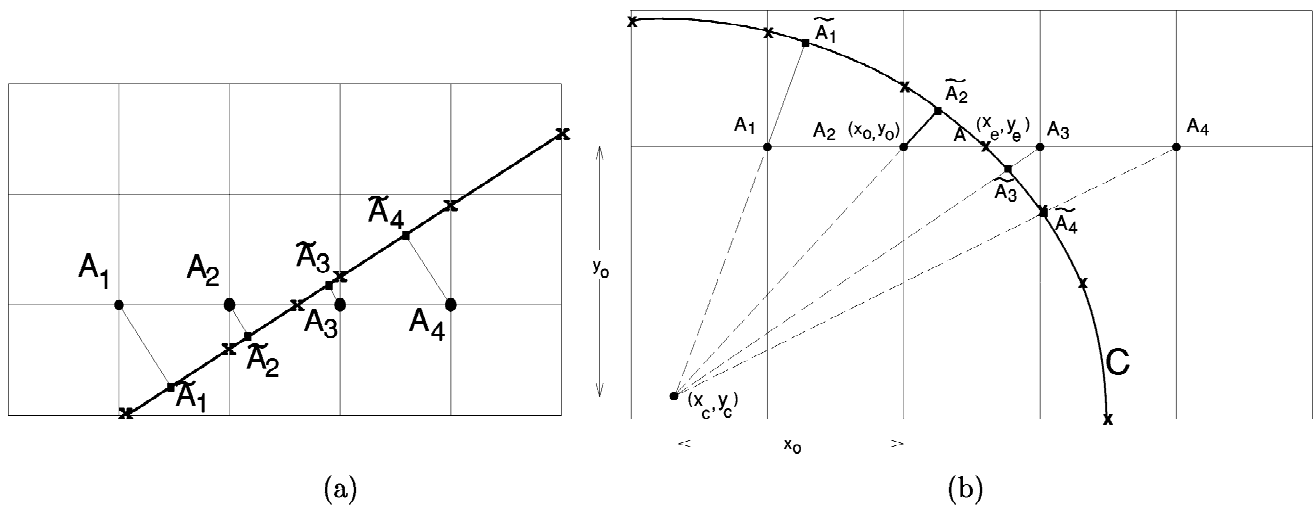

Fig. 23. A profile of the distance transform for an isolated line (a), and circular curve (b), along a horizontal (vertical) grid line to show the effects of interpolation to recover the boundary. The value at the grid crossing A_i is the distance of the grid point from the closest point on the curve, $\tilde{A}_i$ (marked by the black square). The curve (the zero level set of the distance transform) is represented as piecewise circular (PC) curves, using the ENO zero crossings (marked by crosses).
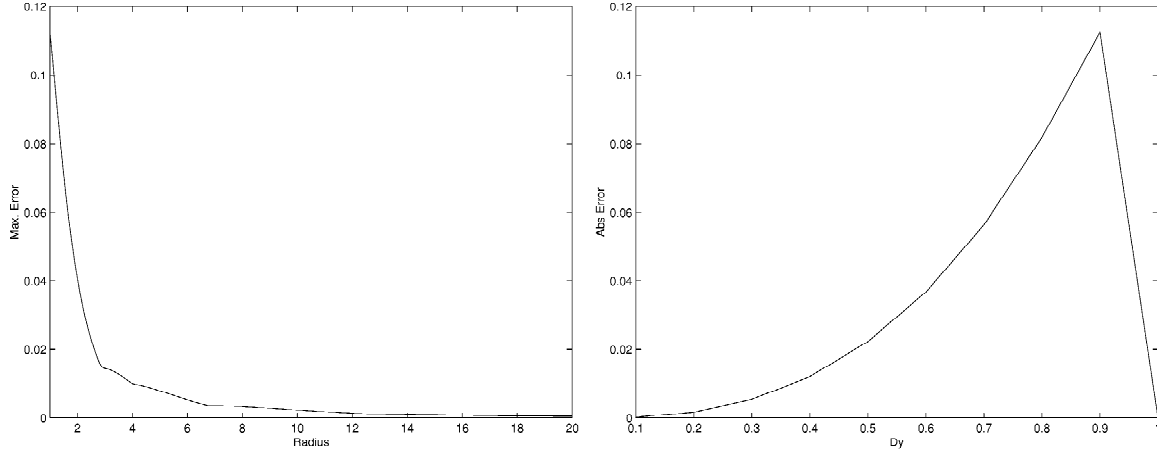
Fig. 24. This figure examines the error in computing the zero crossings of the signed distance transform of an isolated circular curve, i.e., the error in recovering the circular curve. (a) The maximum error for different values of radius R. The maximum error is < 0.12 pixel, for $R = 1$, which is rather low, but which is even smaller if the radius is higher or if the average error is considered, as the plot for the case $R = 1$ (b) shows. D_y is the distance $|y_a - y_c|$.

construction. Then, as the sampled distance transform is first order, using second order ENO interpolation[5] along the horizontal grid line which relies on the values at the grid crossings, namely, $\{A_i \tilde{A}_i\}$, the line is recovered exactly.

Second, consider the recovery of an isolated circular curve, $(x - x_c)^2 + (y - y_c)^2 = R^2$, where (x_c, y_c) is the center of the circle and R is the radius (Fig. 23(b)) from its signed distance transform sampled on the discrete lattice. To recover the curve from the distance transform, all points where the curve intersects the grid lines (marked by crosses in Fig. 23) need to be computed. To simplify calculations, without loss of generality, the origin is shifted to (x_c, y_c). The distance transform, at a horizontal ($y = y_o$) grid crossings $A_i = (x_i, y_o)$ is the distance of the point from the curve is given by $A_i \tilde{A}_i = \sqrt{x_i^2 + y_o^2} - R$ (Fig. 23(b)). Consider now the point $A = (x_e, y_o)$, which lies on this horizontal grid line. To compute x_e, an ENO interpolation along the horizontal grid line is done. If second order ENO interpolation is used, three of the four points A_1, A_2, A_3, A_4 are selected for interpolating the interval $A_2 A_3$. In addition to A_2 and A_3, either A_1 or A_4 (depending on the point that gives the lower curvature) is used for the second order ENO interpolation. Note that in the case of an isolated circle, the point away from the center (A_4 in Fig. 23(b)) is always chosen for second order ENO interpolation. Let (x_0, y_0) denote the coordinates of the grid point interior to the circle and adjacent to the zero crossing, (A_2

in Fig. 23(b)), where $(-R < y_o < R)$ and $x_o = \text{floor}|R^2 - y_0^2|$. We will express all other points in terms of x_0 and y_0, i.e., $A_1 = (x_0 - 1, y_0)$, $A_3 = (x_0 + 1, y_0)$ and $A_4 = (x_0 + 2, y_0)$. The function values at A_1, A_2, A_3, A_4 are given by $\sqrt{(x_0 - 1)^2 + y_o^2}$, $\sqrt{x_0^2 + y_o^2}$, $\sqrt{(x_0 + 1)^2 + y_o^2}$ and $\sqrt{(x_0 + 2)^2 + y_o^2}$, respectively. Hence, the interpolating polynomial for the interval $A_2 A_3$ is given by

$$f(x) = \alpha x^2 + \beta x + \gamma, \tag{A.1}$$

where

$$\alpha = \frac{1}{2}\sqrt{(x_o + 2)^2 + y_0^2} - \sqrt{(x_o + 1)^2 + y_0^2}$$
$$+ \frac{1}{2}\sqrt{x_o^2 + y_0^2}, \tag{A.2}$$

$$\beta = \alpha(2x_o + 1) + \sqrt{(x_o + 1)^2 + y_0^2} - \sqrt{x_o^2 + y_0^2}, \tag{A.3}$$

$$\gamma = \sqrt{x_o^2 + y_0^2} - R - x_o(\sqrt{(x_o + 1)^2 + y_0^2} - \sqrt{x_o^2 + y_0^2})$$
$$+ \alpha x_o(x_o + 1). \tag{A.4}$$

The zero crossings of the distance transform along the horizontal gridline, x_e, is computed by solving the zeros of Eq. (A.1), and choosing the one that lies in the valid range. Thus, we have computed the zero crossing of the distance transform, i.e., the point (x_e, y_o). The actual zero crossing, (x_a, y_a), can be computed analytically by simultaneously solving for the equation of the circle and the horizontal line. Fig. 24(a) shows the maximum error in computing the zero crossings i.e., in recovering the curve from the distance transform, for different values of R, and Fig. 24(b) shows the error for $R = 1$ which even in the worst case is acceptable. We now consider the reconstruction accuracy of some simple shapes, from the distance transform. Fig. 25 shows the error in recovering the curve from the exact distance transform of a parabola and two circles. If the two circles are taken to be representative of two seeds whose

---

[5]To find the polynomial approximation between the grid locations x_j and x_{j+1}, we start by computing a first order polynomial using the values $f(x_j)$ and $f(x_{j+1})$. A second order polynomial is constructed by adding either x_{j-1} or x_{j+2}, whichever produces a smoother polynomial, i.e., lower curvature. Specifically if x_{j+2} is chosen, and $C_1 = f[x_j, x_{j+1}]$ and $C_2 = f[x_j, x_{j+1}, x_{j+2}]$ are the first and second Newton divided differences, second order ENO interpolation for the interval x_j, x_{j+1} is computed as $C_2 x^2 + (C_2(x_j + x_{j+1}) + C_1)x + f(x_j) - (C_1 x_j + C_2 x_{j+1} x_j)$. Further details on computing ENO interpolation can be found in (Siddiqi et al., 1997).
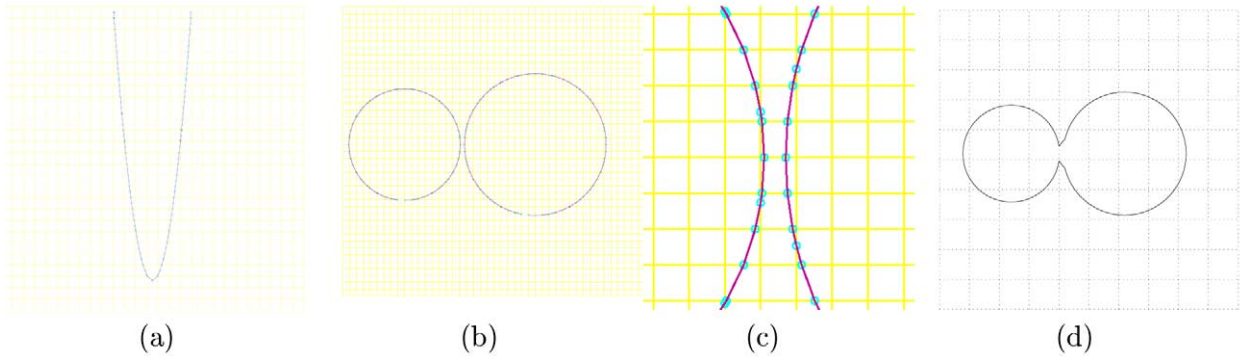
Fig. 25. This figure shows the accuracy in recovering some simple curves from their exact signed distance transform. For the parabola $y = (x - 26)^2 + 3.2$ (a), the maximum error is 0.01 pixel. The two circles, within the same pixel (b) are recovered with a maximum error of 0.005 pixel. A zoomed version of (b) is shown in (c). Bilinear interpolation fails to recover these circular curves (d). This shows that these subpixel methods are essential in implementing the SCDM method.
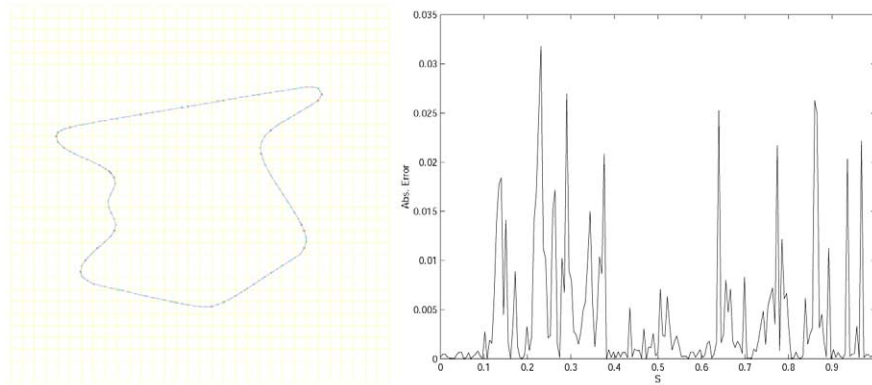

Fig. 26. This figure shows the error in recovery of the zero levelset from the computed distance transform for the shape on the left. The plot on the right shows the absolute error for each zero crossing, all of which are lower than 0.035, a rather acceptable error margin.

growth has made it adjacent, the need for such subpixel methods in implementing SCDM becomes evident. Fig. 26 shows the error in recovering the curve from the computed distance transform using wave propagation for a simple shape. This error is typical and in line with the above analysis, thus rendering the process suitable for implementing SCDM.

Appendix B

ENO interpolation algorithm

This section gives a brief overview of the ENO interpolation algorithm discussed in (Siddiqi et al., 1997). Traditional spline or polynomial interpolation techniques tend to smooth over discontinuities by propagating information across them. ENO interpolation has been proposed to address this problem. The basic idea of ENO interpolation

is to select the data points which give lower variation among two contiguous sets of data points. This means that in intervals neighboring a discontinuity, interpolation is done from the side not containing the discontinuity.

The basic 1D ENO interpolation begins with a first degree polynomial $P_{j+1/2}^{f,1}(x)$ interpolating the function $f(x)$ between the two grid points x_j and x_{j+1}. If we stop here, we obtain the first-order monotone approximation. Whenever a higher order is desired, we add just one point to the existing stencil, chosen from the two immediate neighbors by the sizes of the two relevant divided differences, which measure the local smoothness of the function $f(x)$. Given point values $f(x_j)$, $j = 0, \pm 1, \pm 2, \ldots$ of a (usually piecewise smooth) function $f(x)$ at discrete nodes x_j, we associate an rth degree polynomial $P_{j+1/2}^{f,r}(x)$ with each interval $[x_j, x_{j+1}]$, with the left-most point in the stencil as $x_{k_{\min}^{(r)}}$, constructed inductively as described in Table 2. Note that the ENO interpolation of a 2D function (2D image or embedding surface) is done along the horizontal and vertical gridlines.

Table 2
This table adapted from (Siddiqi et al., 1997) sketches the ENO interpolation algorithm. In the procedure $f[\cdot,\cdots,\cdot]$ are the standard Newton divided differences, $f[x_1, x_2,\cdots,x_{k+1}] = f[x_2,\cdots,x_{k+1}] - f[x_1,\cdots,x_k]/x_{k+1} - x_1$ with $f[x_1] = f(x_1)$

---

1. Find the first-order polynomial interpolation $P_{j+1/2}^{f,1}(x)$ and initialize the first stencil point for the point j, $k_{min}^{(1)}$:

$$P_{j+1/2}^{f,1}(x) = f[x_j] + f[x_j, x_{j+1}](x - x_j)$$
$$k_{min}^{(1)} = j$$

2. If $k_{min}^{(l-1)}$ and $P_{j+1/2}^{f,l-1}(x)$ are both defined, then

$$P_{j+1/2}^{f,l}(x) = P_{j+1/2}^{f,l-1}(x) + c^{(l)} \prod_{i=k_{min}^{(l-1)}}^{k_{min}^{(l-1)}+l-1} (x - x_i),$$

where

$$c^{(l)} = \begin{cases} b^{(l)} & \text{if } |a^{(l)}| \geq |b^{(l)}| \\ a^{(l)} & \text{otherwise} \end{cases}$$

$$k_{min}^{(l)} = \begin{cases} k_{min}^{(l-1)} - 1 & \text{if } |a^{(l)}| \geq |b^{(l)}| \\ k_{min}^{(l-1)} & \text{otherwise} \end{cases}$$

and finally,

$$a^{(l)} = f[x_{k_{min}^{(l-1)}}, \cdots, x_{k_{min}^{(l-1)}+l}]$$
$$b^{(l)} = f[x_{k_{min}^{(l-1)}-1}, \cdots, x_{k_{min}^{(l-1)}+l-1}]$$

References

Adams, R., Bischof, L., 1994. Seeded region growing. IEEE Trans Pattern Analysis and Machine Intelligence 16 (6), 641–647.

Banchoff, T., Giblin, P., 1994. On the geometry of piecewise circular curves. American Mathematical Monthly 101 (5), 403–416.

Beveridge, J.R., Griffith, J., Kohler, R.R., Hanson, A.R., Riseman, E.M., 1989. Segmenting images using localized histograms and region merging. International Journal of Computer Vision 2 (3), 311–352.

Caselles, V., Catte, F., Coll, T., Dibos, F., 1993. A geometric model for active contours. Numerische Mathematik 66, 1–31.

Caselles, V., Kimmel, R., Sapiro, G., 1997. Geodesic active contours. IJCV 22 (1), 61–79.

Cohen, L.D., Cohen, I., 1993. Finite element methods for active contour models and balloons for 2D and 3D images. IEEE Trans. Pattern Analysis and Machine Intelligence 15 (11), 1131–1147.

Cooper, D., Elliott, H., Cohen, F., Reiss, L., Symosek, P., 1980. Stochastic boundary estimation and object recognition. Computer Graphics and Image Processing 12, 326–356.

Crisco, J.J., McGovern, R.D., Wolfe, S.W., 1999. A non-invasive technique for measuring in vivo three-dimensional carpal bone kinematics. Journal of Orthopaedic Research 17 (1), 96–100.

Dempster, A., Laird, N., Rubin, D., 1977. Maximum likelihood from incomplete data via the EM algorithm. Journal of Royal Statistical Society, Series B 39, 1–38.

Duda, R.O., Hart, P.E., 1973. Pattern Classification and Scene Analysis. John Wiley & Sons.

Geiger, B., 1993. Three-dimensional modeling of human organs and its application to diagnosis and surgical planning. Technical Report 2105, INRIA.

Giblin, P.J., Kimia, B.B., 1999. On the local form and transitions of symmetry sets, and medial axes, and shocks in 2D. In: Proceedings of the Fifth International Conference on Computer Vision, KerKyra, Greece. IEEE Computer Society Press, pp. 385–391.

Grayson, M.A., 1987. The heat equation shrinks embedded plane curves to round points. J. Differential Geometry 26, 285–314.

Iverson, L.A., Zucker, S.W., 1995. Logical/linear operators for image curves. IEEE Trans. Pattern Analysis and Machine Intelligence 17 (10), 982–996.

James, S.E., Richards, R., McGrouther, D.A., 1992. Three-dimensional CT imaging of the wrist. Journal of Hand Surgery 17B (5), 504–506.

Karrholm, J., Elmqvist, L., Selvik, G., Hansson, L., 1989. Chronic antrolateral instability of the knee: a roentgen stereophotogrammetric evaluation. American Journal of Sports Medicine 17 (4), 555–563.

Kass, M., Witkin, A., Terzopoulos, D., 1988. Snakes: Active contour models. International Journal of Computer Vision 1 (4), 321–331.

Kichenassamy, S., Kumar, A., Olver, P., Tannenbaum, A., Yezzi, A., 1995. Gradient flows and geometric active contour models. In: Proceedings of the Fifth International Conference on Computer Vision, Cambridge, MA, pp. 810–815.

Kimia, B.B., Tannenbaum, A.R., Zucker, S.W., 1990. Toward a computational theory of shape: an overview. In: Proceedings of the First European Conference on Computer Vision, Antibes, France. Springer Verlag, pp. 402–407.

Kimia, B.B., Tannenbaum, A.R., Zucker, S.W., 1995. Shapes, shocks, and deformations. I: The components of shape and the reaction–diffusion space. IJCV 15 (3), 189–224.

Koh, T.J., Grabiner, M., DeSwart, R., 1992. In vivo tracking of the human patella. Journal of Biomechanics 25 (6), 637–644.

Lafortune, M., Cavanagh, P., Sommer, H., Kalenak, A., 1992. Three-dimensional kinematics of the human knee during walking. Journal of Biomechanics 25 (4), 347–358.

Laidlaw, D., Fleischer, K., Barr, A., 1998. Partial volume Bayesian classification of material mixtures in MR volume data using voxel histograms. IEEE Trans. on Medical Imaging 17 (1), 74–86.

Leclerc, Y.G., 1989. Constructing simple stable descriptions for image partitioning. International Journal of Computer Vision 3 (1), 73–102.

Lee, D., Massear, V., 1993. Comparison of arthroscopic and MRI interpretation of carpal ligamentous injuries: An in vitro study. Personal communication.

Leymarie, F.F., Kimia, B.B., 2000. Discrete 3D wave propagation for computing morphological operations from surface patches and unorganized points. In: Goutsias, J., Vincent, L., Bloomberg, D. (Eds.), Mathematical Morphology and its Applications to Image and Signal Processing. Computational Imaging and Vision, Proceedings of International Symposium on Mathematical Morphology (ISMM), Palo Alto, CA, USA, Vol. 18. Kluwer Academic, pp. 351–360.

Leymarie, F., Kimia, B.B., 2001. The shock scaffold for representing 3D shapes. In: Arcelli, C., Cordella, L., di Baja, G.S. (Eds.), Proceedings of the International Workshop on Visual Form, Capri, Italy. Springer, pp. 216–228.

Malladi, R., Sethian, J.A., Vemuri, B.C., 1994. Evolutionary fronts for topology-independent shape modeling and recovery. In: Proceedings of European Conference on Computer Vision, Stockholm, Sweden, pp. 3–13.

Mayfield, J.K., 1984. Patterns of injury to carpal ligaments – A spectrum. Clinical Orthopaedics 187, 36–42.

McInerney, T., Terzopoulos, D., 1995. Topologically adaptable snakes. In: Proceedings of the Fifth International Conference on Computer Vision, Cambridge, MA, pp. 840–845.

Neskovic, P., Kimia, B.B., 1994. Three-dimensional shape representation from curvature-dependent deformations. In: Proceedings of the IEEE International Conference on Image Processing, Austin, Texas. IEEE Computer Society Press, pp. 6–10.

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

Osher, S., Shu, C.-W., 1991. High-order essentially non-oscillatory schemes for Hamilton-Jacobi equations. SIAM Journal of Numerical Analysis 28, 907–922.

Redner, R.A., Walker, H.F., 1984. Mixture densities, maximum likelihood and the EM algorithm. SIAM Review 26 (2), 195–239.

Ruby, L.K., Cooney, W., An, K., Linscheid, R., Chao, E., 1987. The effect of scapholunate ligament section on scapholunate motion. Journal of Hand Surgery 12A (5), 767–771.

Savelberg, H., 1991. Wrist Joint Kinematics and Ligament Behavior. PhD thesis, Nijmegen University.

Savelberg, H., Kooloos, J.G., de Lange, A., Huiskes, R., Kauer, J.M., 1991. Human carpal ligament recruitment and three-dimensional carpal motion. Journal Orthopaedic Research 9 (5), 693–704.

Savelberg, H., Kooloos, J.G., Otten, J.D.M., Huiskes, R., Kauer, J.M., 1993. Carpal bone kinematics and ligament lengthening studied for the full range of joint movement. Journal of Biomechanics 26 (12), 1389–1402.

Serra, J., 1982. Image Analysis and Mathematical Morphology. Academic Press.

Siddiqi, K., Kimia, B.B., 1996. A shock grammar for recognition. In: Proceedings of the Conference on Computer Vision and Pattern Recognition, pp. 507–513.

Siddiqi, K., Kimia, B.B., Shu, C., 1997. Geometric shock-capturing ENO schemes for subpixel interpolation, computation and curve evolution. Graphical Models and Image Processing 59 (5), 278–301.

Tagare, H.D., Elder, K.W., Stoner, D.M., Patterson, R.M., Nicodemus, C.L., Viegas, S.F., Hillman, R.G., 1993. Location and geometric descriptions of carpal bones in CT images. Annals of Biomedical Engineering 21 (6), 715–726.

Tek, H., Kimia, B.B., 1995. Image segmentation by reaction–diffusion bubbles. In: International Conference of Computer Vision, pp. 156–162.

Tek, H., Kimia, B.B., 1997. Volumetric segmentation of medical images by three-dimensional bubbles. Computer Vision and Image Understanding 64 (2), 246–258.

Tek, H., Kimia, B.B., 1998. Curve evolution, wave propagation, and mathematical morphology. In: Heijmans, H.J., Roerdink, J.B. (Eds.), Mathematical Morphology and its Applications to Image and Signal Processing. Computational Imaging and Vision, Vol. 12. Kluwer Academic, Amsterdam, The Netherlands, pp. 115–126.

Tek, H., Leymarie, F., Kimia, B.B., 1997. Multiple generation shock detection and labeling using CEDT. In: Proceedings of the International Workshop on Visual Form, Capri, Italy. World Scientific, pp. 582–593.

Vincent, L., Soille, P., 1991. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. IEEE Trans. Pattern Analysis and Machine Intelligence 13 (6), 583–598.

Weeks, P.M., Vannier, M.W., Stevens, W.G., Gayou, D., Gilula, L.A., 1984. Three-dimensional imaging of the wrist. Journal of Hand Surgery 10A (1), 32–39.

Xu, L., Strauch, R.J., Ateshian, G.A., Pawluk, R.J., Mow, V.C., Rosenwasser, M.P., 1998. Topography of osteoarthritic thumb carpometacarpal joint and its variations with regard to gender, age, site and osteoarthritic stage. Journal of Hand Surgery 23A (3), 454–464.

Zhu, S.C., Yuille, A.L., 1996. Region competition: unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. IEEE Trans. Pattern Analysis and Machine Intelligence 18 (9), 884–900.