HHS Public Access

Author manuscript

Eng Comput (Swansea). Author manuscript; available in PMC 2015 March 06.

Published in final edited form as:

Eng Comput (Swansea). 2014 ; 31(3): 530–566. doi:10.1108/EC-06-2012-0145.

Techniques to derive geometries for image-based Eulerian computations

Seth Dillard,
Mechanical and Industrial Engineering, University of Iowa, Iowa City, Iowa, USA

James Buchholz,
Mechanical and Industrial Engineering, University of Iowa, Iowa City, Iowa, USA

Sarah Vigmostad,
Biomedical Engineering, University of Iowa, Iowa City, Iowa, USA

Hyunggun Kim, and
Internal Medicine, Division of Cardiology, University of Texas Health Science Center at Houston, Houston, Texas, USA

H.S. Udaykumar
Mechanical and Industrial Engineering, University of Iowa, Iowa City, Iowa, USA

Abstract

Purpose—The performance of three frequently used level set-based segmentation methods is examined for the purpose of defining features and boundary conditions for image-based Eulerian fluid and solid mechanics models. The focus of the evaluation is to identify an approach that produces the best geometric representation from a computational fluid/solid modeling point of view. In particular, extraction of geometries from a wide variety of imaging modalities and noise intensities, to supply to an immersed boundary approach, is targeted.

Design/methodology/approach—Two- and three-dimensional images, acquired from optical, X-ray CT, and ultrasound imaging modalities, are segmented with active contours, k-means, and adaptive clustering methods. Segmentation contours are converted to level sets and smoothed as necessary for use in fluid/solid simulations. Results produced by the three approaches are compared visually and with contrast ratio, signal-to-noise ratio, and contrast-to-noise ratio measures.

Findings—While the active contours method possesses built-in smoothing and regularization and produces continuous contours, the clustering methods (*k*-means and adaptive clustering) produce discrete (pixelated) contours that require smoothing using speckle-reducing anisotropic diffusion (SRAD). Thus, for images with high contrast and low to moderate noise, active contours are generally preferable. However, adaptive clustering is found to be far superior to the other two methods for images possessing high levels of noise and global intensity variations, due to its more sophisticated use of local pixel/voxel intensity statistics.

Dr Seth Dillard is the corresponding author and can be contacted at: seth-dillard@uiowa.edu.

Originality/value—It is often difficult to know a priori which segmentation will perform best for a given image type, particularly when geometric modeling is the ultimate goal. This work offers insight to the algorithm selection process, as well as outlining a practical framework for generating useful geometric surfaces in an Eulerian setting.

Keywords

Segmentation; De-noising; Eulerian fluid and solid computation; Image-based modeling; Level sets

1. Introduction

Computational modeling of flows in the presence of solid geometries relies on representations of these geometries in various forms. The emerging trend in many applications, such as in biomedical systems (Dubini *et al*., 2011; Steinman *et al*., 2003; Antiga *et al*., 2008), micro-flows in porous media (Lin *et al*., 2010; Al-Omari and Masad, 2004; Yue *et al*., 2003) or in materials modeling (Baer, 2002; Ghosh, 2004) is to reconstruct the geometries to be modeled from image data (acquired from *in vivo* patient data or *in situ* geological sampling, etc.). Imaging modalities can be varied, spanning optical imaging (Tytell and Lauder, 2004), ultrasound (Boukerroui *et al*., 2003; Hovda *et al*., 2008; Kao *et al*., 1997), MRI (van Pelt *et al*., 2012; Makowski *et al*., 2002), and X-ray CT (Cnudde *et al*., 2006; Vignoles, 2001). Figures 1 and 2 illustrate several examples of image/video data that can be employed to compute fluid/solid dynamics. When 3D or 4D (3D + time) images are drawn from these varied sources, a robust pre-processing capability is required to cast the geometries in a form that can be used in computational modeling. In particular, images are typically polluted with noise of different levels and types (Zhou and Bovik, 2002; Wang *et al*., 2004) depending on the imaging modality. To render these images suitable for computations of flows, solid mechanics, or thermal transport, they need to be pre-processed before interfacing with a computational mechanics code that is based on finite difference/ finite element or other framework. While the field of processing images to derive geometries is extensive (Chan and Shen, 2005), this paper examines available approaches from a particular viewpoint, i.e., to choose an approach that delivers the best result that can be utilized by an Eulerian fluid/solid mechanics code. As will be shown later, a suitable choice depends on a number of factors, including image quality (noise intensity, contrast, complexity of the background), desired smoothness of the derived geometry and whether the image is stationary or moving.

In broad terms, there are two approaches to computation of flows in or around the presented geometries. In the first type the computational mesh conforms to the surfaces defined by the modeled geometries; in such body-fitted mesh approaches (Zhang, 2006; Tai *et al*., 2005) volumetric mesh generation is required to fit to the surfaces of geometries (such as fluid flow domains, solid objects, etc.). In the second approach, the computational mesh is independent of the geometries through or around which flows ensue; this latter case corresponds to immersed or embedded boundary approaches (Mittal and Iaccarino, 2005). The present paper builds from previous work in which techniques for solving transport phenomena in complex domains (with stationary or moving boundaries) were developed in a

fixed Eulerian grid framework (Vigmostad *et al*., 2010; Udaykumar *et al*., 2009; Marella *et al*., 2005b). The primary advantage of this framework is that the grid on which PDEs are solved to calculate transport phenomena is decoupled from the embedded geometries, saving the user the rather tedious and onerous task of volumetric mesh generation. The challenge of developing techniques for treating the embedded boundaries as sharp entities has been addressed in several previous works (Liu *et al*., 2005; Marella *et al*., 2005b; Sambasivan and Udaykumar, 2011; Udaykumar *et al*., 2009). In the Eulerian framework it is advantageous to represent the embedded geometries in Eulerian fashion as well. This is achieved by means of a level set field (Osher and Fedkiw, 2001; Osher and Sethian, 1988; Sethian, 2001), which is a signed-normal-distance field that resides on the grid. The zero-contour of this field is the embedded boundary. This blend of level set geometry representation and fixed grid transport solvers can be advantageously exploited when geometries are delivered from acquired images (or video), thus opening up the possibility of circumventing grid generation to conform to complex geometries.

Image-based modeling is based on segmentation, which is the process of partitioning an image into different regions (Brice and Fennema, 1970; Pal and Pal, 1993). This paper examines the application of image segmentation techniques to obtain geometries that can be employed in simulations of flows interacting with complex geometries in mechanics, materials modeling, and biomedical engineering. The ability to accurately describe material properties at the micro- and meso-scales and model the response of the materials to imposed loads is desirable in many engineering and scientific disciplines (Yue *et al*., 2003; Rollett *et al*., 2007). While it is often sufficient to assume homogeneity in materials or assign average multiphase properties (as in the mixture formulation approaches) based upon empirical observations, the mechanical response of a material to an applied force can depend strongly on microstructural details that are missed when such simplifications are made. Examples include the initiation of failure at heterogeneities through crack propagation in metals and other materials (Zhang and Feng, 2011; Liu *et al*., 2010; Chen *et al*., 2006; Duarte *et al*., 2001) and the initiation of detonation via localized heat release due to void collapse (Tran and Udaykumar, 2006a, b). In these cases, while the micro- and meso-structures are critical to the appearance of macro-scale phenomena, the geometric complexity of microstructures makes describing them (not to mention computationally modeling their response) difficult, particularly in three dimensions. However, emerging image analysis techniques such as X-ray CT (Ketcham and Carlson, 2001; Al-Omari and Masad, 2004) have opened a path to 3D material modeling that requires no a priori assumptions about the shapes and orientations of small compositional structures. In biomedical sciences, modeling flows in patient-specific geometries (Steinman *et al*., 2003; Antiga *et al*., 2008) has emerged as an important effort, partly in recognition of the critical role played by biological variability in the response of systems. The wide variety of modalities for imaging *in vivo* geometries makes it challenging to deliver these geometries to a computational code for fluid dynamic or structural dynamic computations. However, it would be attractive (and useful) to develop a technique to seamlessly bridge images to computational geometry, under the unifying framework of level set segmentation, and an Eulerian solver for computational modeling of the mechanics of fluid and solid materials.

With this in mind, we seek to develop a framework that achieves image segmentation with level set geometric representation (Osher and Sethian, 1988), allowing for complex geometries to be modeled in a simple manner as implicit surfaces within a regular Cartesian domain. A significant hurdle to cross in establishing this facility is that images from different modalities contain different levels (and types) of noise. Figures 1, 2, 3 and 4 illustrate the types of images that will be processed using the techniques developed in the work, showing the variety of noise levels and patterns. Arguably, the “cleanest” images are obtained with visible light sources; such images (including video) are typically polluted with additive white noise, which can be removed in a fairly straightforward way. On the other end of the spectrum, images obtained from ultrasound are imbued with large amplitude multiplicative (speckle) noise (Yu and Acton, 2002; Sun *et al*., 2004; Kremkau and Taylor, 1986), and X-ray CT images may contain significant scattering, streaking, or ring artifacts (Barrett and Keat, 2004). Obtaining satisfactory segmentations of features in the presence of various types of noise and processing images to extract geometries that can be used in computational modeling thus requires an evaluation of techniques from the arsenal of tools developed for image processing. Often, these techniques are specialized to one specific modality, or they may stop short of delivering the type and quality of information that is useful to modelers. For example, in biomedical applications it may be sufficient to delineate a lesion (Karmonik *et al*., 2009; Kakar and Olsen, 2009; Pollo *et al*., 2005) or tumor (Venel *et al*., 2008) with limited accuracy measures, such as comparison with “ground truth” (Crevier, 2008); typically this latter is a visual segmentation by an expert. For computational modeling, however, the geometries reconstructed from the image must possess smooth, regular surfaces, and a clear separation of noise from features. This work, therefore, assesses state-of-the-art techniques for feature extraction from noisy images provided by various modalities, and evaluates these techniques for their ability to deliver geometries suitable for computational modeling in an Eulerian framework. Although Lagrangian or body-fitted approaches are not considered in this paper, deriving surface descriptions using tessellations or NURBS descriptions of surfaces may be considered as extensions of the present work (Lorensen and Cline, 1987).

2. Image segmentation

Segmentation is the practice of delineating an imaged scene into distinct groups of objects, and is one of the essences of image-based modeling (Oh *et al*., 2001; Antiga *et al*., 2008). Pixels (or voxels, in 3D) in a digital image are grouped together based on properties like brightness intensity or texture, and are partitioned by different labels or by contours computed on the image. Many algorithms have been developed to accomplish this, including statistical methods (Besag, 1986; Pieczynski, 1992), edge detection (Canny, 1986; Ma and Manjunath, 1997; Bellon and Silva, 2002), and region growing (Mumford and Shah, 1989; Adams and Bischof, 1994; Chan and Vese, 2001) methods. Here we consider three frequently used approaches: active contour evolution (Vese and Chan, 2002), k-means clustering (Kanungo *et al*., 2002; Gibou and Fedkiw, 2005), and adaptive clustering (Pappas, 1992; Ashton and Parker, 1995). These three approaches are well suited to delivering geometries in the form of level set information drawn from image intensity values and have been employed in the image processing community to segment images obtained from

various modalities. Their relative merits and the quality of geometries delivered needs examination, however, to assess which approach works best under what circumstances; in this work we examine all three approaches for images acquired from visible light video, X-ray CT imaging of meso-scale features in different solid materials, and 2D (still images) and 3D ultrasound images. By critically examining these three approaches, specific recommendations are presented on the suitability of the methods for devising a route from image to computation.

2.1 Active contour segmentation

Active contours are based on early applications of elastic deformation theory to computer vision (Terzopoulos *et al*., 1987). In this approach modeled surfaces are assumed to naturally tend toward shapes that minimize the energy of mismatch between the surface contour and an underlying topology. Mumford and Shah (1989) applied this notion to the segmentation problem by treating images as approximate representations of physical objects. They proposed an energy functional that decomposes an image domain O into piecewise smooth regions delineated by a segmentation curve C:

$$E(f,C)=\mu L(C)+\lambda\int_{\Omega}(I_0-I)^2 d\Omega+\int_{\Omega\setminus C}|\nabla I|^2 d\Omega \quad (1)$$

In Equation (1), which is commonly referred to as the Mumford-Shah (M-S) functional, $L(C)$ is segmentation curve length, and $\lambda>0$ and $\mu\geq 0$ are adjustable weighting parameters. The first (curve length, or regularization) term ensures that segmentation curves follow object boundaries smoothly even in the presence of noise. The second term forces the representative function I to approximate the original image intensity I_o, and the third term constrains the spatial variation of I to be smooth in each segment. This way, an image can be decomposed into segments that are discontinuous across their boundaries while maintaining smoothness within each region.

Chan and Vese (2001) later incorporated the M-S functional into a framework that segments an image into two piecewise constant regions of average intensity rather than piecewise smooth regions of slowly varying intensity, eliminating the smoothness constraint in Equation (1). They also recast the energy minimization problem in terms of level sets, so that segmentation curves could be represented implicitly as zero-level isocontours of a signed distance-level set field φ embedded in a Cartesian mesh. The level set formulation also made possible the addition of a second regularization term restricting the area of a segment, along with the original curve length restriction, giving the Chan-Vese (C-V) energy functional:

$$E(\phi,c_1,c_2)=\mu\int_{\Omega}|\nabla\mathscr{H}(\phi)|d\Omega+v\int_{\Omega}\mathscr{H}(\phi)\,d\Omega+\lambda_1\int_{\Omega}(I_0-c_1)^2\mathscr{H}(\phi)\,d\Omega+\lambda_2\int_{\Omega}(I_0-c_2)^2(1-\mathscr{H}(\phi))\,d\Omega \quad (2)$$

In Equation (2), μ, v, and λ_i are weighting parameters and (φ) is the Heaviside function of level set field φ, which C-V approximate in regularized form on a Cartesian mesh with grid spacing Δ as:

$$\mathscr{H}_{\in}(\phi)=\frac{1}{2}\left[1+\frac{2}{\pi}arc\,tan\left(\frac{\phi}{\Delta}\right)\right] \quad (3)$$

The constants c_i in the C-V functional represent the average brightness intensity in each of the two segmentation regions *i*, and are defined as:

$$c_1=\frac{\int_\Omega I_0\mathscr{H}_{\in}(\phi)\,d\Omega}{\int_\Omega \mathscr{H}_{\in}(\phi)\,d\Omega}, \quad c_2\frac{\int_\Omega I_0\,(1-\mathscr{H}_{\in}(\phi))\,d\Omega}{\int_\Omega (1-\mathscr{H}_{\in}(\phi))\,d\Omega} \quad (4)$$

Euler-Lagrange minimization of Equation (2) gives the level set evolution equation:

$$\frac{d\phi}{dt}=\delta_{\in}(\phi)\left[\mu\nabla\cdot\left(\frac{\nabla\phi}{|\nabla\phi|}\right)-v-\lambda_1(I_0-c_1)^2+\lambda_2(I_0-c_2)^2\right] \quad (5)$$

with regularized delta function:

$$\delta_{\in}(\phi)=\mathscr{H}'_{\in}(\phi)=\frac{1}{\pi}\frac{\Delta}{\Delta^2+\phi^2} \quad (6)$$

The presence of the delta function restricts level set evolution to the zero-level segmentation curve, where $\varphi = 0$, though Chan and Vese mention that extension to all level sets of φ can be achieved by replacing $\delta_{\in}(\varphi)$ with $|\nabla\varphi|$. Evolution of Equation (5) proceeds from some initial zero level, usually defined as a closed contour with an interior region possessing level set values $\varphi \geq 0$ and exterior region having $\varphi > 0$.

2.2 k-Means clustering followed by image smoothing

Gibou and Fedkiw (2005) observed that active contour methods could be made significantly simpler and more computationally efficient by pre-processing images with a diffusion operation prior to segmentation so that regularization terms could be dropped from the segmentation algorithm; furthermore, setting $\delta_{\in}(\varphi) = 1$ applies the evolution equation to all level sets of φ. Extension of the evolution equation to all level sets is attractive because it helps to prevent the evolving front from becoming trapped in a local minimum solution, and thus reduces sensitivity to initial placement of the segmentation contour. This leads to the ODE:

$$\frac{d\phi}{dt}=-\lambda_1(I_0-c_1)^2+\lambda_2(I_0-c_2)^2 \quad (7)$$

Gibou and Fedkiw also noted that the evolution of φ is itself not of much interest in the case of image segmentation; rather, the final segmented (steady-state) result is what is of primary importance. It is thus desirable to segment the image in as few steps as possible so that the final result can be quickly obtained. To this end they cast Equation (7) within a *k*-means clustering framework (Lloyd, 1982; MacQueen, 1967; Kanungo *et al.*, 2002) by suggesting a labeling scheme in which the segmentation field is iteratively updated as:

$$\phi^{n+1}=1 \cdot sign\left[\left(\frac{d\phi}{dt}\right)^{n}\right] \quad (8)$$

By setting $\varphi = \pm 1$ in this way, Equation (7) is equivalent to *k*-Means, which minimizes the sum-of-squares error in each of the segmentation clusters *k*:

$$E=\sum_{i=1}^{k}\sum_{I_j \in C_i}|I_j - c_i|^2 \quad (9)$$

Labeling the pixels/voxels $\varphi = \pm 1$ also allows for rapid computation of the constants c_i as:

$$c_1=\frac{\sum I_0\,(\phi+1)}{\sum(\phi+1)}, \quad c_2=\frac{\sum I_0\,(\phi-1)}{\sum(\phi-1)} \quad (10)$$

While the method proposed by Gibou and Fedkiw was outlined for the specific case of two segmented regions, the labeling scheme could be easily modified to allow for extension to multiple clusters.

The greatest advantages of the *k*-Means level set approach are its simplicity and its speed of execution. Evolving the C-V formulation given in Equation (5) is restricted by stability criteria, so that small time steps on the order of $\Delta t = \Delta x / (d\varphi / dt)_{\max}$ must be taken in updating the evolving contour. *k*-Means suffers no such restrictions, and we have found that in practice the update Equation (8) generally converges within two or three iterations, which is consistent with Gibou and Fedkiw's observations.

The main disadvantages are that, due to the absence of built-in regularization, the method can produce noisy segments when faced with even a slightly noisy image, and pixel/voxel labeling results in pixelated (i.e. “stair-stepped”) regions in place of the smooth segments that are produced by active contours. Thus, the image must first be smoothed before the algorithm can proceed, and then the segmentation contours must be smoothed once again in a post-processing step to eliminate pixelation. This second step is not necessary if the intent is to simply process an image for visualization, but if the segmented geometries are intended for computation of flows the surfaces must be smooth (at least smooth enough to be resolved by the flow mesh). Thus, when *k*-Means clustering is used for deriving level sets, denoising becomes an essential complement to segmentation in order to deliver geometries for computation.

In a comparison of various denoising algorithms, including edge-preserving diffusion (Yu and Acton, 2002; Sun *et al*., 2004; Perona and Malik, 1990) and wavelet-based methods (Zhong and Sun, 2008), the speckle-reducing anisotropic diffusion (SRAD) method of Yu and Acton (2002), and later Sun *et al*. (2004), was found to produce the best results on the present class of images of interest to flow computations. Generally, diffusion-based image denoising methods are of the form:

$$\frac{\partial I\,(x,y,z;t)}{\partial t}=\nabla \cdot \left[c\,(x,y,z;t)\,\nabla \mathrm{I}\,(x,y,z;t)\right] \quad (11)$$

with a diffusion coefficient *c* that can either be constant (isotropic or Gaussian smoothing) or a function of some image property such as intensity gradient (anisotropic diffusion). In SRAD, *c* is a function of both the image intensity gradient and the Laplacian of intensity, so that regions of large gradient are preserved – unless they are also regions of high curvature, in which case they are strongly diffused. The SRAD diffusion coefficient is written in three dimensions as:

$$c(q) = \frac{1}{1 + [q^2(x,y,z;t) - q_o^2(t)] / [q_o^2(t)(1+q_o^2(t))]} \quad (12)$$

with instantaneous coefficient of variation:

$$q(x,y,z;t) = \sqrt{\frac{(1/3)(|\nabla I|/I)^2 - (1/6)^2(\nabla^2 \mathrm{I}/\mathrm{I})^2}{(1 + (1/6)(\nabla^2 \mathrm{I}/\mathrm{I}))^2}} \quad (13)$$

$q_o(t)$ is known as the speckle scale function, and is computed as the weighted average instantaneous coefficient of variation over the image domain:

$$q_o(t) = \alpha \frac{\int_V q(x,y,z;t)\, dV}{\int_V dV} \quad (14)$$

In our implementation of SRAD, discretization and updating following Sun *et al*. (2004), along with their use of symmetric boundary conditions. Even though SRAD was originally developed specifically for ultrasound images, which tend to be corrupted by large amounts of multiplicative (speckle) noise, we have found it to be quite effective at smoothing while preserving important details in images captured using optical and X-ray modalities. It can also be applied as a way to smooth pixelated contours generated by the clustering methods, which will be demonstrated in Section 3 when we segment various test images. For each case in which SRAD was used, diffusion was applied for 100-time steps, which was sufficient to approach near steady-state behavior (RMS change in image intensity below 10^{-4}) for all of the images tested.

2.3 Adaptive clustering

Pappas (1992) made significant improvements to clustering-based segmentation by proposing a multi-level adaptive clustering framework, which we have extended to 3D for this study. The method applies Bayes' theorem to an image signal *I* composed of *x* segments, so that:

$$p(x|I) \propto p(I|x)\, p(x) \quad (15)$$

The a priori probability *p(x)* indicates the likelihood of a pixel/voxel *s* belonging to the same segmentation region as its neighbors, and the conditional density *p(I|x)* compares local pixel/voxel characteristics with those of the different segmentation regions distributed in the observed image. By employing a Gibbs Random Field model, Pappas' algorithm estimates the a priori probability as:

$$p(x) = \frac{1}{Z} exp\left[-\sum_{q \in \mathcal{N}_s} V_C(x_q)\right] \quad (16)$$

where Z acts as a normalizing constant. $V_C(x_q)$ are known as clique potentials, which are defined by:

$$V_C(x_q) = \begin{cases} -\beta & \text{if} \quad x_s = x_q \\ +\beta & \text{if} \quad x_s \neq x_q \end{cases} \quad (17)$$

and summed over the nine-point (2D) or 27-point (3D) neighborhood $\mathcal{N}_s$ of each pixel/voxel s. Taken together, these clique potentials are simply a measure of the number of pixels/voxels in a local neighborhood belonging to the same cluster x, weighted by a parameter β.

The conditional probability density $p(I|x)$ is calculated in a manner similar to k-Means, with local values compared to regional characteristics by way of a sum-of-squares error:

$$p(I|x) \ \alpha \ exp\left\{-\sum_s \frac{1}{2\sigma^2}[I_s - \mu_s^{x_s}]^2\right\} \quad (18)$$

Here the variance of brightness intensity σ^2 weights the error between local brightness I_s and the mean of region x_s. Unlike k-Means, however, the mean intensity of each segment x_s is considered to be a smoothly varying mean field μ_s, calculated using an interpolation window approach outlined in (Pappas, 1992). This smoothly varying field allows for pixels/voxels to be compared to the statistics of local regions, making the method more robust to global intensity variations. The combined distribution is then:

$$p(x|I) \ \alpha \ exp\left\{-\sum_s \frac{1}{2\sigma^2}[I_s - \mu_s^{x_s}]^2 - \sum_{q \in \mathcal{N}_s} V_C(x_q)\right\} \quad (19)$$

Pappas' algorithm was later modified by Ashton and Parker (1995) to improve the reliability of segmentation results on images that are heavily polluted by speckle noise, such as those acquired via ultrasound, by allowing for the intensity variance σ^2 to be calculated adaptively in each region x_s, for each window at all levels of resolution:

$$p(x|I) \ \alpha \prod_s \left(\frac{1}{\sigma_{x_s}}\right) exp\left\{-\sum_s \frac{1}{2\sigma_{x_s}^2}[I_s - \mu_s^{x_s}]^2 - \sum_{q \in \mathcal{N}_s} V_C(x_q)\right\} \quad (20)$$

Maximizing this probability function is equivalent to maximizing its natural logarithm, so that it can be viewed alternatively in terms of an energy minimization function:

$$E(x|I) = \sum_s \left\{ln(\sigma_{x_s}) + \frac{1}{2\sigma_{x_s}^2}[I_s - \mu_s^{x_s}]^2\right\} + \sum_{q \in \mathcal{N}_s} V_C(x_q) \quad (21)$$

$E(x|I)$ is minimized using the iterated conditional modes (ICM) approach put forth by Besag (1986) and within the adaptive framework presented by Pappas (1992). The complete algorithm consists of two iterative loops; an outer loop, which computes σ_{xs} and $\mu_s^{x_s}$ given the current configuration of regions x_s in the image, and an inner loop, which executes a raster scan update of the pixels/voxels in the image, whereby the regional x_s label assigned to each pixel/voxel is the one which minimizes the energy function (21). The inner raster scan loop is considered converged when the number of updates is less than some inner threshold ($T_{inner} = W/2^{nlev-1}$ is suggested in (Pappas, 1992)), with adaptive clustering window width W and image coarsening level $nlev$ ($nlev = 1$ is the original image). The outer iterative loop is converged when the number of raster scan updates in the inner loop becomes smaller than an outer threshold (e.g. $T_{outer} = W/10$). In (Pappas, 1992) parameter β is set to the same value for all levels, but we usually found results to be improved when adapting β to increase by some increment $\Delta\beta$ for each level of image refinement as suggested by Ashton and Parker (1995). We also examined texture-based and weighted adaptive clustering approaches(e.g. Boukerroui *et al*., 1998, 2003), but found them to be computationally expensive without significantly improving performance on the image data sets evaluated here.

3. Application of the methods to real image data

The methods described in Section 2 have been applied to images acquired using optical cameras, X-ray CT, and 2D and 3D ultrasound. It is difficult to know a priori which segmentation method will work best for a given data set given that different modalities will produce images with different clarity, contrast, and noise characteristics; thus, the primary objective of this work concerns comparing the performance of the different segmentation algorithms on the image data we wish to use for modeling.

3.1 2D and 3D imaging data sets

3.1.1 Two-dimensional optical imaging—Optical images, such as those captured by CCD- or CMOS-based cameras are used extensively for scientific imaging in applications such as particle image velocimetry (PIV) and video imagery of organisms for elucidation of the mechanics of their complex motions (Tytell and Lauder, 2004). Visible light images are advantageous in that they are intuitive to create and interpret, and the equipment can be quite inexpensive compared with that of other imaging modalities. They are also relatively easy to control in a manner that yields desirable results from a modeling standpoint. While at present most visible light images are limited to two spatial dimensions, emerging technologies may produce 3D optical imaging capabilities (Lynch *et al*., 2012; Belden *et al*., 2010; Zhang, 2010).

To test the methods on visible light images, we segmented image frames from video files captured during two different experiments shown in Figure 1. The first video sequence was of an American eel (*Anguilla rostrata*) swimming in a water flume (courtesy of Dr Eric Tytell) (Tytell and Lauder, 2004), and the second was of a flexible panel plunging in a low-turbulence uniform flow. The eel video was a reduced-resolution AVI file (260 × 54 pixels), with high contrast between the eel and background, but containing significant noise. The

panel was cast from polydimethylsiloxane with embedded spanwise ribs to ensure two-dimensional deformation, and the edge was impregnated with fluorescent paint. Illumination with ultraviolet light provided a high-contrast image of the edge of the panel to provide boundary conditions for a computational fluid dynamics simulation. The panel video, acquired by a consumer video camera in 1080p, was output as a compressed AVI file (360 × 240 pixels) for segmentation. (The image was cropped to 300 × 120 pixels in Figure 1 to highlight the region of interest.)

3.1.2 X-ray CT imaging—The clear advantage X-ray CT holds over visible light imagery is its ability to produce true 3D data sets that render otherwise optically inaccessible internal structures. In addition to its widespread use in health care, CT scanning has also proved useful for geostructural analysis, non-destructive inspection, and other materials-related applications (Cnudde *et al*., 2006). In this work, we consider CT images of three different materials composed of microstructural details that would be useful to model in an Eulerian approach to examine the effects of shock loading on heterogeneous materials (Baer, 2002): an aluminum foam material, concrete, and packed sugar crystals that resemble the material matrix found in, e.g., solid combustible propellants and heterogeneous explosives. While the three images tested share a common acquisition method, they are not equally amenable to segmentation. For instance, the aluminum foam image slices are highly contrasted – it is easy to separate material regions from voids via image intensity differences, which is evident by visual inspection – but the geometry is also complex and filled with narrow filaments that can potentially make segmentation difficult. The concrete images, on the other hand, feature rocky structures that are well defined and generally well separated from each other, but that are also poorly contrasted from the surrounding material. Finally, the sugar crystal images have good contrast and are visually straightforward to interpret, but they are noisy and contain crystalline structures that are closely packed together; in fact, the variation in pixel intensity due to noise is on the same order as the average intensity difference between the crystals and their surroundings, and the noisy patches share the same spatial scale with the narrow gaps between crystal structures. As such, the level of difficulty in segmenting these images increases in the order in which they are presented, as will become clear from the following.

3.1.3 2D and 3D ultrasound images—Portable, relatively inexpensive equipment and the facility to acquire 4D (3D + time) data make ultrasound a promising modality for image-based modeling – particularly in the context of moving boundary problems. Unfortunately, ultrasound images also suffer large amounts of speckle noise and other artifacts (e.g. attenuation, reverberation) (Kremkau and Taylor, 1986), and thus present their own unique challenges to segmentation efforts. In this section we compare the performance of the active contour, *k*-Means, and adaptive clustering methods on two sets of ultrasound images: The first is a 2D image of a liver, from the Philips ultrasound image library (Figure 3), and the second is a 3D transesophageal echocardiographic (TEE) data set featuring a mitral (bicuspid) valve of the heart acquired by one of the authors (HK), provided in 18 slices arranged in cylindrical polar coordinates (Figure 4).

For the mitral valve data set, each slice (samples of which are shown in Figure 4) represents a cross-section of the geometry in cylindrical polar coordinates, and 18 slices separated by 10° increments in the azimuthal (θ) direction comprise the entire image volume. Each slice was obtained from 3D TEE data of a patient. (The Committee for the Protection of Human Subjects at The University of Texas Health Science Center at Houston has approved this study.) The 3D mitral valve geometry (anterior/posterior leaflets and annulus) was obtained using a 3D TEE transducer installed on a Philips iE33 ultrasound unit (Philips Medical Systems, Bothell, WA). The standard 3D TEE echo protocol was performed to acquire the mitral valve geometry.

3.2 Application of image segmentation methods

To save computational time and facilitate analysis, 3D image sets were reduced to 64 × 64 × 64 voxels prior to treatment by selecting 64 adjacent slices and cropping them identically so that a 64 × 64 pixel region remained in the center of each slice. 2D images were either left in their original form, or cropped as indicated to isolate regions of interest. All computations were performed on a Linux desktop workstation with a single Intel Core i5-2500 3.3 GHz CPU and 8 GB RAM. Run times are given in CPU seconds (wall clock time).

Each of the image data sets was segmented using the three methods outlined in Section 2, with SRAD smoothing applied to images prior to segmentation as needed. Level set contours were initialized as spheres (3D) or circles (2D) centered in the image domain, each with a diameter of half the image width:

$$\phi_i = \begin{cases} |x_i - x_c| - r, & \text{if continuous} \\ 1 \cdot sign\left[|x_i - x_c| - r\right], & \text{if pixelated} \end{cases} \quad (22)$$

In Equation (22), x_c denotes the coordinates of the pixel/voxel nearest to the center of the image domain, x_i is the coordinate of the pixel/voxel for which φ_i is being calculated, and r is the radius of the initialization surface or contour. When testing the active contour approach, the C-V evolution Equation (5) was evolved from the initialization given by Equation (22) on each of the image data sets, with central differencing used to compute derivatives. Time steps were incremented as $\Delta t = \Delta x / (d\varphi / dt)_{\max}$. Equation (5) was evolved to convergence, which was defined the RMS change in φ over the image domain falling below the tolerance limit $\in = 10^{-5}$.

k-Means testing followed the straightforward implementation presented in (Gibou and Fedkiw, 2005) with updating given by the labeling scheme (8) and convergence defined the same way it was using active contours, with tolerance set to $\in = 10^{-5}$.

Adaptive clustering was tested following the method proposed by Pappas (1992), coupled with the automatic variance calculation suggested by Ashton and Parker (1995) when segmenting noisy CT and ultrasound images. (This afforded more accurate localized results while eliminating the need to assume an intensity variance for the images a priori.) The algorithm otherwise proceeded along the lines outlined in (Pappas, 1992): Each image was coarsened three levels (decimated by powers of 2) and a minimum overlapping window size of $W = 8$ pixels or voxels in each direction was used. *k*-Means clustering was performed on

the coarsest image level to initialize the segmentation field as the first step. Then the domain was divided into overlapping windows of width: $W=\frac{1}{2}\,min\,(n_x, n_y)$ (or $W=\frac{1}{2}\,min\,(n_x, n_y, n_z)$ in 3D), with n_i representing the image dimensions at each level, and the variance and mean were computed for each segmentation region x_i in each window. The mean was then interpolated onto the image pixel/voxel domain to produce a slowly varying field. The image size was next refined one level and the process repeated, starting with the final result obtained at the coarser-level interpolated onto the refined grid. The minimum window size of $W = 8$ from any coarse level corresponds to $W = 16$ at the next refined level, so that the iterative update procedure takes place on the same number of windows at each level once the initial *k*-Means solution is obtained. Iteration followed the ICM procedure proposed by Besag (1986).

Successful segmentation can be sensitive to the choice of parameters for each of the methods evaluated. It also depends upon the parameters chosen for SRAD smoothing when it is applied to the image intensity or to the segmentation contour as a secondary step. As such, varying parameters to achieve desirable results was a large part of the work performed for this study.

Active contour evolution is driven by the relative weights of four different parameters; λ_1 and λ_2, which bias the contour toward the average brightness value of the inside or outside segmentation region, respectively; μ, which weights the curve length penalty applied to the boundary; and v, which weights the area occupied by the segmented region. Of these parameters, the λ_i have the greatest influence on the final segmentation result, and have values of (1). Small changes in λ_i were sometimes found to have a large effect on the final segment, pushing the contour several pixels/voxels toward the inside or outside of an imaged boundary depending on their relative values. The curve length minimizer μ has significantly less influence, usually taking values of (10)-(100) before noticeably affecting the result. Nonetheless, when given a large enough weight, μ has the effect of minimizing curvature in the segmentation contour. This acts to remove tortuosity from the segment, and generally further biases closed contours toward the outside of regions they are segmenting in an effort to decrease curvature. The second regularization term, which simply adds a constant value of v to the evolution formula (and thus speeds up convergence in some cases), was found to have little or no effect on the final result and was therefore neglected.

Parameter selection for the *k*-Means clustering method was straightforward, as there were only the two λ_i brightness weights to choose from. *k*-Means is essentially a simple threshold method, and the brightness coefficients together determine which two average image intensity values that threshold will lie between. Like the active contours brightness weights, λ_i were of (1) and small variations sometimes led to significant differences in the final segmentation result.

The adaptive clustering approach also required selection of two parameters: the clique potential weight β and the incremental change of that weight, $\Delta\beta$, at each grid resolution level. However, unlike parameters in the other two approaches, optimal adaptive clustering parameters varied considerably from one image to the next; values for β ranged from (0.01)

to (1). Furthermore, the method was sometimes found to be highly sensitive to the value selected, with small deviations giving much different results.

Pre-smoothing images with SRAD and smoothing segmentation contours as a secondary step further required selection of the variational coefficient weight α in Equation (14), which determined how strongly image features would be diffused. In all of the cases outlined in this paper, the segmentation results given are those that best matched visual expectations after balancing the parameters in several different combinations for each segmentation method. The parameters selected for each case are provided in Table I.

After segmentation, a set of points was constructed containing the cell-center locations of pixels/voxels adjacent to each segmented interface. In the case of pixelated segments (the clustering methods) cell centers were simply assigned distance values of $\pm\Delta/2$. For the smooth, continuous segments generated by active contours, cell center locations were found by triangulation (2D) or marching cubes (3D) (Lorensen and Cline, 1987). The fast marching method (Sethian, 1999) was used to construct a narrow-band level set representation of each segmentation contour by solving the Eikonal equation:

$$|\nabla\phi| = F \quad (23)$$

with speed function $F = 1$ and "marching" the solution outward from the initial set of interface points. For the clustering approaches, SRAD was applied to the φ field in a post-processing step to eliminate pixelation. In addition to smoothing, this also had the useful effect of removing small isolated zero-level contours resulting from noise that was not fully eliminated during the initial image denoising step.

A flowchart of the entire image denoising and segmentation process is provided in Figure 5. With these implementations in place the three approaches were subjected to tests as detailed below.

3.2.1 Analysis of the optical image data—. Each video sequence was segmented frame-by-frame with the parameter selections provided in Table I. For the heaving panel frames, active contour and *k*-Means results were found to be relatively insensitive to parameter variation due to image smoothness and sharp transition between the panel and the background. For the eel video frames, however, the brightness weights were set to bias the segmentation contour toward the background's average brightness value. This had the effect of pushing the contour outward toward the edge of the gradient region defining the eel's boundary, where the pixel values were found to vary more smoothly compared with those that lay near the center of the boundary gradient. For active contour evolution, the smoothness weight was given a sufficiently large value to mitigate the effect of pixelated blocks of image intensity along the eel's boundary, while taking care to not make it so large that it would remove desired high-curvature features from the segment (such as the end of the eel's tail).

Adaptive clustering segmentation of the optical images needed to be handled differently than in other cases, following the original formulation proposed by Pappas (Equation (19)) rather than the modified formulation of Ashton and Parker (Equation (21)). This is because

the adaptive variance calculation suggested in Ashton and Parker was designed specifically for noisy images, which have non-zero local variance values by virtue of the noise they contain. However, each of these video examples features a single distinct object and a nearly homogeneous background, so variances in many local regions of the image turned out to be very small. In fact, the variance is zero in the background of the flexible panel, which leads to the first two terms in Equation (21) taking on undefined or very large-magnitude values there. Thus for the American eel and flexible panel image frames, an a priori variance had to be assigned in order to produce a reasonable segment.

For the active contours and *k*-Means cases, the eel frames were smoothed with SRAD prior to segmenting, and the segmentation contours produced by the cluster-based methods were treated again in a secondary smoothing step with SRAD.

3.2.2 Analysis of the X-ray CT data—For this study, 64 slices were chosen from each X-ray CT image data set and cropped to 64 × 64 pixels centered within each slice as illustrated in Figure 2, then segmented with active contours, *k*-Means, and adaptive clustering. Segmentation results for each of the three image volumes are presented in two different ways: Figure 3 shows where segmentation contours intersect an *x*-*y* plane near the center of the data set making up each image (corresponding to slice 32), and Figure 7 gives segmentation results on the faces of each data set volume. These results were obtained using the three approaches, namely, active contours (Figures 6(a) and 7(a)); *k*-Means clustering (Figures 6(b) and 7(b)); and adaptive clustering (Figures 6(c) and 7(c)).

3.2.3 Analysis of the 2D and 3D ultrasound data—Testing on the 2D liver ultrasound image was performed within a square-cropped 128 × 128 pixel view of a large vascular feature (Figure 3), in order to facilitate comparison between the segmentation methods and to eliminate contributions to the final segment from the black background surrounding the ultrasound viewing area. (This isolation from the background could have also been accomplished by restricting segmentation calculations to pixels within the desired viewing area.) SRAD was applied to the image before segmenting with active contours and *k*-Means, effectively removing most of the speckle noise. Applying SRAD prior to segmenting with adaptive clustering led to problems similar to those encountered with the optical images in the previous section, where homogeneity over large portions of the image necessitated an a priori variance estimate. Results obtained by applying SRAD and assuming a variance value were actually worse than those obtained by segmenting the original image using Ashton and Parker's method, thus the latter was chosen. SRAD was, however, applied to the segmentation level set contour to eliminate pixelation in both clustering cases.

In order to prepare the mitral valve data set for smoothing and segmentation, the slices were first cropped to 64 × 64 pixel windows that were centered horizontally about a registration reference line (the green vertical line visible in each slice in Figure 4), so that 32 columns of pixels lay to each side of the centerline of the mitral valve geometry. These slices were then stored in an array with the z-coordinate representing the vertical component of each image slice and the *r*-coordinate representing the horizontal distance of each pixel's center from the mitral valve's centerline. The result was a set of 36 slices separated by $\theta = 10°$, with a height of 64 pixels and a radius of 32 pixels (Figure 8(a)). This cylindrical polar intensity

information was then mapped onto a regular Cartesian mesh of size 64 × 64 × 64 voxels using fourth-order Lagrange interpolation in the r and θ directions, and with linear extrapolation in the r-direction used to fill in the necessary information beyond the radial range of the cylinder in Cartesian space (Figure 8(b)).

After mapping the image into regular Cartesian coordinates, it was processed in the same manner as the other 3D data sets in Section 3.2.2. The image was denoised using 3D SRAD, removing much of the speckle noise (as illustrated by isocontours of average gray-level intensity before and after applying SRAD in Figure 9). Finally, the image was segmented using each of the three methods being evaluated presently. Like in previous cases, segmentation contours produced by the clustering methods were smoothed in a secondary step using SRAD. However, unlike in previous cases, secondary smoothing was also applied to the active contours result. This was done to give a better direct comparison between the three methods; capturing undesirable noise was unavoidable with both active contours and k-Means (Figure 10(a) and (b)). Only adaptive clustering was able to filter excess noise during the segmentation process itself (Figure 10(c)), with SRAD only necessary to smooth out the pixelated segmentation contour.

4. Results

Results produced by the three different segmentation methods for each of the test images were generally quite different from each other. On one hand this is to be expected, as the images all have different levels of contrast and noise, as well as varying degrees of complexity within the imaged scene, and the methods all handle these circumstances differently. On the other hand, some images that visually appeared as though they should be simple segmentation problems for any of the methods turned out not to be straightforward for one or more of them, while other, much noisier images gave usable results.

In an effort to understand the types of image properties that may lead to one segmentation method or the other producing favorable results, we have characterized the images in terms of several different measures relating to their brightness patterns and level of noise pollution. Each of the images has been separated into two regions – a brighter region x_1 with mean brightness intensity μ_1, and a darker region x_2 with mean brightness intensity μ_2. In each case, the two regions are delineated by the segmentation contour visually judged to have given the best results. The quality of each image was then measured in terms of three ratios:

1. the contrast ratio μ_1/μ_2;
2. a signal-to-noise ratio defined by μ_i/σ_i in each region x_i; and
3. an overall contrast-to-noise ratio defined by $|\mu_1 \mu_2|/\sigma$.

These measures are summarized in Table II.

4.1 Images captured in visible light

Both of the optical video sequences feature simple shapes that are well contrasted from their backgrounds and free of large amounts of noise, so segmentation was fairly straightforward

regardless of the method chosen in that the segments described their intended shape without capturing any noise. This was particularly true for the panel frame (Figure 1(e)), which had the highest contrast ratio of any of the images evaluated. Both the eel and panel images also possessed high contrast-to-noise ratios (Table II). Segmentation results are given in Figure 1 for both geometries.

In both the eel and flexible panel cases the active contour approach gave a segment that closely matches visual expectations without any further manipulation, due to its built-in smoothing and regularization. However, because the videos were compressed AVI files containing a large amount of pixelation, both the *k*-Means and the adaptive clustering method yielded a coarse representation of boundaries that could not easily be smoothed without over-diffusing the edges and eliminating legitimate regions of high curvature from the segments, e.g. the end of the eel's tail (Figure 1(c) and (d)). At the other extreme, the slight curvature in the flexible panel was resolved well with active contours, but the clustering methods represented it as a series of “steps” (Figure 1(g) and (h), highlighted by arrows) from which it would be impossible to accurately recover the underlying imaged geometry. Thus for both of these video image frames, active contours provided the most usable results.

Run times for the three segmentation methods were comparable for the eel video, taking ~2*s* per frame to complete. Run times were somewhat longer for the panel video, taking ~5*s* per frame for *k*-Means and ~10*s* per frame for the other methods.

4.2 X-ray CT data

4.2.1 Aluminum foam—The aluminum foam image was also relatively straightforward to segment, likely due to its high contrast ratio. The active contour approach segmented the aluminum foam volume (Figures 6(a) and 7(a), top row) well from a qualitative perspective, giving surface segments that appear similar to the original intensity and segmentation contours that closely trace the edges of bright regions on the interior slice. There are a few material regions where bright patches that should have been separated were grouped together, and a few regions that were missed altogether (indicated by the arrows in Figure 6), but the overall characteristics of the material were captured. Applying the rapid *k*-Means approach (Figures 6(b) and 7(b), top row) gave segmentation results similar to those obtained by active contours, also likely due to the excellent contrast exhibited by the aluminum foam image. While some of the contours do not match the object boundaries as closely as active contours, *k*-Means did manage to keep separated some of the bright patches that had coalesced under active contour evolution. The results of the adaptive clustering approach applied to the aluminum foam are given in Figures 6(c) and 7(c), top row. The results are similar to those found with the active contour model, though it was difficult to match the smoothness of the active contour boundaries without removing some features in the secondary level set smoothing step. Nonetheless, the segments closely match the original data set and appear to be the most useful for constructing a representative material model.

Run times for active contours and adaptive clustering were ~20 s, while *k*-Means took ~15 s to segment.

4.2.2 Concrete—The relatively poor contrast of the concrete volume proved to be more problematic (Figures 6(a) and 7(a), middle row) for the active contour approach than in the case of the aluminum foam. In fact, the concrete CT has the lowest contrast ratio of any of the images evaluated. While some regions were segmented with accuracy, regions near the slice edges in particular were not reliably outlined. This is likely due to the fact that the overall contrast degrades gradually in the radial direction from the *z*-axis, which can be seen in the original image slice shown in Figure 2(b). This degradation in contrast leads to a global intensity variation, which affects the location of the threshold between the average intensities of the two segmentation regions. The problem is compounded by the fact that small changes in the brightness threshold lead to large changes in the segmentation contour's position in regions where the level of local contrast is not sufficient to clearly distinguish between objects. Nonetheless, the active contour model may be a good choice if the contrast could be enhanced without introducing too much noise. The *k*-Means method exhibited considerable difficulty when applied to the concrete image (Figures 6(b) and 7(b), middle row). In the absence of any regularization terms, the combination of noise and lack of contrast produced several visible segmentation errors, particularly near *x*-*z* and *y*-*z* faces as indicated by the arrows in Figure 7. The poor contrast appears to have had a particular impact in this case, causing segmentation contours to follow the threshold between constant region averages c_i even where there are no object boundaries visible. On the other hand, segments produced by adaptively clustering the concrete image (Figures 6(c) and 7(c), middle row) are markedly better than those produced by the other two methods (This is most apparent in Figure 7). While a few features were not resolved as well as in the high-contrast aluminum foam case, the spatially adaptive nature of the algorithm made it much more robust to local degradation of contrast in the already low-contrast image. This is especially evident on the *x*-*z* and *y*-*z* faces of the image volume, where the other methods were corrupted by noise and unable to distinguish many of the features. This is also apparent when comparing 3D surface isocontours in the interior of the concrete image in Figure 11; the adaptive clustering model (Figure 5(c)) is composed entirely of individual rock structures in that region, without the noise or coalescence of structures seen with the other two methods (Figure 11(a) and (b)). Adaptive clustering was clearly superior to the other segmentation methods in this case.

Run times for this case were nearly identical to those encountered previously with the aluminum foam data set.

4.2.3 Sugar crystals—The sugar crystal data set was the most difficult of the three to segment adequately using any of the approaches. The image volume is corrupted by large amounts of noise (note the low overall contrast-to-noise ratio combined with relatively low signal-to-noise ratios in each region of the image, Table II), and the gaps between crystalline structures share the same spatial scale as the noise variation. It was thus found to be impossible to segment individual crystal structures in parts of the image without some of them being grouped together in the same segment. Active contours showed the most difficulty in this setting (Figures 6(a) and 7(a), bottom row), isolating many small regions bounded by the tenuous average intensity threshold. *k*-Means clustering with SRAD applied to the level set field performed better (Figures 6(b) and 7(b), bottom row), mainly because

the secondary smoothing served to remove much of the noise that remained after segmentation was complete. Otherwise, the results would have looked much like a pixelated version of the active contour result. Most notable here was again the adaptive clustering segmentation result, which presented a marked improvement over the other methods for the sugar crystal image (Figures 6(c) and 7(c), bottom row). Some details have been lost on the surface of the image volume, corresponding to slice edges, but interior segments very closely match boundaries in the intensity data. This is especially evident in the overlay of the segmentation contour onto the image slice at the center of the volume in Figure 6; the noisy segments and poorly matched regions have been eliminated altogether. However, the grouping of very closely spaced crystal structures still unfortunately proved to be unavoidable.

It should be noted here that an effort was made to improve the active contour and *k*-Means results by first applying SRAD to eliminate noise in the CT images prior to segmentation. While it did in fact serve to remove much of the noise effectively, particularly that seen in the sugar crystals results shown in the bottom row of Figure 6(a) and (b), it also led to more individual sugar crystals coalescing in the final result. Overall, the adaptive clustering algorithm provided the most useful results out of the three methods applied to the 3D CT images, whether directly or after pre-processing with SRAD.

Run times for the three methods were again ~20 s for active contours and adaptive clustering, and ~15 s for *k*-Means.

4.3 2D and 3D ultrasound

4.3.1 2D ultrasound of the liver—Figure 12 shows the results achieved by each of the three methods overlaid on the original unprocessed image of the liver. Like the visible light images captured in the video sequences in Section 3.2.1, the 2D liver ultrasound image features fairly simple geometries that are highly contrasted from their surroundings (especially after applying SRAD). The contrast ratio is second only to that of the panel video frame, and SRAD smoothing gave the image a favorable signal-to-noise ratio in the lighter speckled region as it was designed to. Also like the visible light images, the liver image gave the best segmentation results when active contours were used. Active contours (Figure 12(a)) successfully segmented each of the four primary features of interest – the large branching vessel and the three smaller vascular cross-sections – much better than the cluster-based approaches, with smooth contours that visually appear to follow object boundaries closely. By comparison, *k*-Means clustering missed (Figure 12(b)) two of the three small vascular features (indicated by arrows in the figure), and adaptive clustering missed one of them. In addition, both *k*-Means and adaptive clustering segmented regions that are most likely attributable to noise as opposed to being desirable features to segment, and adaptive clustering segmented one region that should not have been included at all. (This is likely due to the homogeneity produced by SRAD smoothing in that region giving small variance magnitudes. However, assigning an assumed global variance as in Section 3.2.1 made the results worse in this case.) On the other hand, local adaptive clustering statistics (Figure 12(c)) did help to eliminate some of the tortuosity that was produced by active contours in the presence of noise along the boundary of the large branching vessel. Even so, active

contours settled on a shape that matches visual expectations more closely than the other methods overall. In particular, active contours picked out small features missed by the other methods, and it did not segment anything unexpected. Run times for all three methods were ~1*s*.

Based on previous and forthcoming 3D results it is possible that the results would be different if the liver image were a 2D slice of a three-dimensional set, as the extra brightness data provided by adjacent image slices helps to statistically isolate regions of interest and diminish segmentation errors through the addition of available information. This is particularly true when taking the adaptive clustering approach, since it effectively triples the number of data points available to calculate the clique potentials used in the solution process. Thus if brightness patterns correlate strongly between two or more image slices, they are likely to be interpreted as coherent objects; otherwise they are more likely to be interpreted as noise, as will be seen next.

4.3.2 3D ultrasound of the mitral valve—Figure 10 illustrates the segmentation contours produced by each method applied to the 3D TEE image of the mitral valve, prior to the application of any secondary smoothing operation. It is clear that in this case the adaptive clustering method produced the cleanest segmentation due to its built-in ability to distinguish coherent structures in the ultrasound image and remove noise using local multi-level statistical considerations. The three-dimensionality of the image certainly helped this, as it provided a 27-point neighborhood for calculating clique potentials in place of a nine-point neighborhood, and thus gave a good indication of which set a particular voxel belonged to.

Predictably, the other two methods produced segmentation results that closely followed the isocontours of average intensity shown in Figure 9(b) – just as they are designed to do. In the case of the present ultrasound images, with their large levels of noise and poor contrast, the active contour and *k*-Means approaches fail quite dramatically when compared to the adaptive clustering approach.

As in previous cases, SRAD was applied to the level set field after it was constructed from segmentation contours produced by the clustering-based approaches. In addition, we also applied the same secondary smoothing to the active contour result this time, to give a more direct comparison. The results for each are illustrated in Figure 13. For both the *k*-Means and active contour results, secondary smoothing removed the noisy patches and provided continuous surface descriptions that could be useful from a modeling standpoint. However, the independence adaptive clustering gains by considering local statistical averages and clique potentials likely makes it more trustworthy in this situation, where the boundary of the object of interest is heavily obscured by large-scale speckle noise. Run times for all three methods was ~25*s*.

Figure 14 provides a comparison between the mitral valve segmentation given by adaptive clustering and a hand-segmentation of the same cylindrical polar image set. (Note that the coordinate directions differ between the adaptive clustering segmentation in Figure 14(a) and (b), and the hand segmentation in Figure 14(c) and (d); the *z*-axis is inverted and the *x*-*y*

plane is rotated 90°. Coordinates were therefore adjusted in the figures to maintain similarity between the two cases.)

Segmentation by hand was performed by fitting smooth lines through regions visually identified as being part of the valve leaflet structure, resulting in the zero-thickness membrane shown in Figure 14(c) and (d). As such, one of the biggest differences between the hand-segmentation and the adaptive clustering results is the presence of structural volume in the geometry given by adaptive clustering, which is more accurately representative of the physiological valve geometry. It is therefore hoped that mechanical simulations based upon realistic geometries of this type will help lend insight into complex phenomena that may otherwise be difficult to capture using current techniques.

5. Discussion and conclusions

In this paper we have compared specific implementations of three frequently used segmentation methods in the interest of incorporating information from image data into a level set framework for generating 3D geometries that can be employed in level set-based Eulerian fluid and solid mechanics solvers (Udaykumar *et al*., 2001, 2002; Marella *et al*., 2005a). The different methods evaluated clearly possess their own individual strengths and weaknesses – some intrinsic and some dependent upon the properties of the image being segmented – and conducting this study has helped lend some insight into their usefulness when faced with the task of accurately reconstructing surfaces for the purpose of solid and fluid mechanics modeling.

The set of test images used in this study contained varying levels of image contrast, noise, length scales of the salient features, and global intensity variation. A high level of apparent contrast does not automatically imply greater amenability to segmentation. For example, the sugar crystals CT image seems at first glance as though it should be straightforward to segment: the crystals in each slice appear bright, and the background appears comparatively dark. However, this turned out to be one of the most difficult images to segment out of any we tested. Close inspection reveals that the noise level is rather extreme – many of the voxels corrupted by noise in each of the regions possess intensity values that are very near to the average intensity value of the opposing region to which they do not belong. (This is reflected by the signal-to-noise and contrast-to-noise measures in Table II, which are lower than in other images evaluated.) As a result, the voxels fall on the “wrong side” of the average intensity threshold and are therefore likely to be segmented incorrectly. Unlike true speckle noise, which is multiplicative and thus scales with the overall brightness in a given region, the noise in this image suffers large fluctuations regardless of the local intensity values, making it less responsive to smoothing with SRAD.

5.1 Active contours

Of the three methods investigated here, active contour evolution performs best when given an image that exhibits a good contrast ratio and/or contrast-to-noise ratio. Its primary strength in comparison to the other two methods is that, since the resulting segmentations are continuous in nature, it provides surfaces that are smooth and accurate without further manipulation that can lead to errors when the length scale of the image noise is comparable

to imaged physical features. Active contours can also handle moderate amounts of noise by making use of curve length penalties. This makes it particularly suitable for images captured with optical methods where high-quality images can be acquired by refined imaging hardware and the illumination can generally be well controlled.

However, piecewise constant methods (active contours and *k*-Means) work best when it is desirable for the segmentation contours to follow closely along the isocontour of average intensity in an image; the piecewise constant assumption (which in the end amounts to a simple threshold) that underlies active contour evolution is thus limited when faced with less ideal test images. For instance, when image contrast is poor (e.g. in the concrete data set), the similarity in brightness intensity between an object of interest and its background makes the ease with which the two can be distinguished highly sensitive to any noise that might be present. Small variations in the signal are therefore often likely to be assigned to the wrong segmentation grouping, leading to a result that is at odds with what is visually expected. Global intensity variations present in the domain that are independent of the brightness of objects being imaged will similarly tend to bias segmentation contours away from their expected place, which was also illustrated by the concrete CT results.

5.2 k-Means

For highly contrasted images free of noise, *k*-Means was attractive because of its speed; segments were generally achieved within a couple of iterations. However, it suffered similar limitations to active contours, brought about by the piecewise constant approximation, and its usefulness was even further limited by the fact that it produced a pixelated result that needed to be smoothed in a secondary diffusion operation. In addition it was found to be highly sensitive to noise and to contrast variations, and while it sometimes produced reasonable results when given a pre-smoothed image to segment, we found that it was often difficult to achieve a noise-free result without over-diffusing important details and altering the underlying geometry so that it no longer matched the physical characteristics being imaged. Thus for modeling purposes *k*-Means did not generally produce adequate segments given the test images contained herein.

5.3 Adaptive clustering

Adaptive clustering performed the best of all the methods when faced with a noisy 3D environment. Its consideration of local statistics makes it robust to noise, and global contrast variation is handled by a slowly varying mean field in place of an assumption of constancy within each segmentation region. The performance of adaptive clustering was particularly distinguished when applied to the sugar crystals images and the 3D ultrasound data. Large brightness variations due to noise were effectively ignored, giving smooth boundaries that closely followed visual expectations. Unfortunately, the sugar crystals were too closely spaced in the image to segment correctly with any of the methods; they are in fact difficult to distinguish even by visual inspection in some places.

One of the limitations of adaptive clustering is that, like *k*-Means, it gives pixelated edges that must be smoothed prior to use with computational physics algorithms. Adaptive clustering also required on the order of ten inner raster scan passes for each of about ten

outer iterations to reach convergence on every window at each level, and the updated variance and smooth mean field had to be calculated for each of those iterations. All of this makes the process relatively expensive. It is also the most sensitive to the choice of user-input parameters out of the three methods, sometimes leading to a somewhat lengthy trial-and-error process. Nevertheless, since adaptive clustering was still able to provide useful geometries from noisy images that could not be reliably segmented with active contours, the extra trouble was a necessary compromise in those cases.

It should also be pointed out that one major advantage of the clustering methods in general over active contours is the ease with which they can be extended to include more than two segmentation regions. For example, in Figure 2(b) it can be seen that the concrete material is composed of three different regions: the rock structures, which are light gray in color; the cement matrix, which is dark gray; and voids, which are black. In this work we have only considered segmenting into two regions, but clustering affords the ability to segment into an arbitrary number of regions by selecting an appropriate labeling scheme. This would be especially useful for modeling from images like the concrete, where multiple properties need to be assigned to fully capture the physical characteristics of the material.

Ongoing work is focussed on extending the methodology to include more than two material types, and then embedding the 3D microstructure level sets into a high-resolution Cartesian computational domain and performing simulations of shock loading of these microstructures to determine the material response at high strain rates. Work is also underway to simulate incompressible fluid flows in 2D around image-based level set geometries derived from video. Optical flow (Brox *et al*., 2004) calculations are used to set boundary conditions, so that CFD simulations can proceed directly from segmented video frames without the need for any tedious surface mesh generation procedures. In this way we hope to accurately depict the kinematics of organisms such as the swimming American eel to better understand their performance characteristics, as well as directly replicate the motion of experimental devices such as the plunging flexible panel. In the case of the flexible panel, PIV measurements can be taken concurrently with video, offering a basis for direct comparison between numerical simulations and experimental results. Finally, we ultimately aim to simulate flows through a moving mitral valve, by applying the modeling methods outlined here to 4D echocardiograph data. Optical flow will again supply velocity boundary conditions, giving rise to a completely Eulerian framework that can simulate flows through complex image-derived moving geometries on a regular Cartesian mesh.

Acknowledgments

This work was in part supported by grants from the AFOSR Computational Mathematics program (Program Manager: Dr Fariba Fahroo) and Flow Interactions and Control program (Program Manager: Dr Doug Smith), from the AFRL-RWPC (Computational Mechanics Branch, Eglin AFB, Program Manager: Dr Michael E. Nixon), and from the National Institutes of Health (R01 HL109597). The authors would also like to express their gratitude to Dr Michael Nixon, Dr Martin Schmidt and Dr Joel Stewart from AFRL-RWPC for providing the X-ray CT microstructure images used in this work, to Dr Eric Tytell for providing the swimming American eel video sequence, to Jim Akkala for providing the flexible panel video sequence, and to Dr Yonghoon Rim for providing the echocardiographic image of the mitral valve. Without their generosity this study would not have been possible.

Biography

Dr Seth Dillard (PhD, Mechanical Engineering, University of Iowa, 2011) is an IIHR Post-doctoral Research Scholar in the Department of Mechanical and Industrial Engineering at the University of Iowa.

Dr James Buchholz (PhD, Mechanical and Aerospace Engineering, Princeton University, 2006) is an Assistant Professor in the Department of Mechanical and Industrial Engineering at the University of Iowa.

Dr Sarah Vigmostad (PhD, Biomedical Engineering, University of Iowa, 2007) is an Assistant Professor in the Department of Biomedical Engineering at the University of Iowa.

Dr Hyunggun Kim (PhD, Biomedical Engineering, University of Iowa, 2005) is an Assistant Professor in the Department of Internal Medicine, Division of Cardiovascular Medicine at the University of the Texas Health Science Center at Houston.

Dr H.S. Udaykumar (PhD, Aerospace Engineering, University of Florida, 1994) is a Professor in the Department of Mechanical and Industrial Engineering at the University of Iowa.

References

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

Al-omari A, Masad E. Three dimensional simulation of fluid flow in X-ray CT images of porous media. International Journal for Numerical and Analytical Methods in Geomechanics. 2004; 28(13): 1327–1360.

Antiga L, Piccinelli M, Botti L, Ene-iordache B, Remuzzi A, Steinman DA. An image-based modeling framework for patient-specific computational hemodynamics. Med Biol Eng Comput. 2008; 46(11): 1097–1112. [PubMed: 19002516]

Ashton EA, Parker KJ. Multiple resolution Bayesian segmentation of ultrasound images. Ultrasonic Imaging. 1995; 17(4):291–304. [PubMed: 8677563]

Baer MR. Modeling heterogeneous energetic materials at the mesoscale. Thermochimica Acta. 2002; 384(1):351–367.

Barrett JF, Keat N. Artifacts in CT: recognition and avoidance. Radiographics. 2004; 24(6):1679–1691. [PubMed: 15537976]

Belden J, Truscott TT, Axiak MC, Techet AH. Three-dimensional synthetic aperture particle image velocimetry. Measurement Science and Technology. 2010; 21(12):1–21.

Bellon ORP, Silva L. New improvements to range image segmentation by edge detection. IEEE Signal Processing Letters. 2002; 9(2):43–45.

Besag J. On the statistical-analysis of dirty pictures. Journal of the Royal Statistical Society Series B-Methodological. 1986; 48(3):259–302.

Boukerroui D, Baskurt A, Noble JA, Basset O. Segmentation of ultrasound images – multiresolution 2D and 3D algorithm based on global and local statistics. Pattern Recognition Letters. 2003; 24(4): 779–790.

Boukerroui D, Basset O, Guerin N, Baskurt A. Multiresolution texture based adaptive clustering algorithm for breast lesion segmentation. Eur J Ultrasound. 1998; 8(2):135–144. [PubMed: 9845797]

Brice CR, Fennema CL. Scene analysis using regions. Artificial Intelligence. 1970; 1(3):205–226.

Brox, T.; Bruhn, A.; Papenberg, N.; Weickert, J. High accuracy optical flow estimation based on a theory for warping. 8th European Conference on Computer Vision; Prague, Czech Republic. May 11-14; 2004.

Canny J. A computational approach to edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1986; PAMI-8(6):679–698. [PubMed: 21869365]

Chan, TF.; Shen, JH. Image Processing and Analysis: Variational, PDE, Wavelet, and Stochastic Methods, SIAM (Society for Industrial and Applied Mathematics). Philadelphia, PA: 2005. p. 1-29.

Chan TF, Vese LA. Active contours without edges. IEEE Transactions on Image Processing. 2001; 10(2):266–277. [PubMed: 18249617]

Chen YP, Eskandarian A, Oskard M, Lee JD. Meshless simulation of crack propagation in multiphase materials. Theoretical and Applied Fracture Mechanics. 2006; 45(1):13–17.

Cnudde V, Masschaele B, Dierick M, Vlassenbroeck J, Hoorebeke LV, Jacobs P. Recent progress in X-ray CT as a geosciences tool. Applied Geochemistry. 2006; 21(5):826–832.

Crevier D. Image segmentation algorithm development using ground truth image data sets. Computer Vision and Image Understanding. 2008; 112(2):143–159.

Duarte CA, Hamzeh ON, Liszka TJ, Tworzydlo WW. A generalized finite element method for the simulation of three-dimensional dynamic crack propagation. Computer Methods in Applied Mechanics and Engineering. 2001; 190(15):2227–2262.

Dubini G, Ambrosi D, Bagnoli P, Boschetti F, Caiani EG, Chiastra C, Conti CA, Corsini C, Costantino ML, D'Angelo C, Formaggia L, Fumero R, Gastaldi D, Migliavacca F, Morlacchi S, Nobile F, Pennati G, Petrini L, Quarteroni A, Redaelli A, Stevanella M, Veneziani A, Vergara C, Votta E, Wu W, Zunino P. Trends in biomedical engineering: focus on patient specific modeling and life support systems. Journal of Applied Biomaterials & Biomechanics. 2011; 9(2):109–117. [PubMed: 22065388]

Ghosh S. Computational material modeling: a current perspective. CMES-Computer Modeling in Engineering & Sciences. 2004; 5(1):1–3.

Gibou, F.; Fedkiw, R. A fast hybrid k-means level set algorithm for segmentation. 4th Annual Hawaii International Conference on Statistics, Mathematics, and Related Fields; Honolulu, HI. January 9-11; 2005.

Hovda S, Rue H, Olstad B. New echocardiographic imaging method based on the bandwidth of the ultrasound Doppler signal with applications in blood/tissue segmentation in the left ventricle. Computer Methods and Programs in Biomedicine. 2008; 92(3):279–288. [PubMed: 18471927]

Kakar M, Olsen DR. Automatic segmentation and recognition of lungs and lesion from CT scans of thorax. Computerized Medical Imaging and Graphics. 2009; 33(1):72–82. [PubMed: 19059759]

Kanungo T, Mount DM, Netanyahu NS, Piatko CD, Silverman R, Wu AY. An efficient k-means clustering algorithm: analysis and implementation. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2002; 24(7):881–892.

Kao C, Pan X, Chen C, Hiller EM, Newmark GM, Lipton MJ. Accurate edge extraction and its application in automatic segmentation of ultrasound images of blood vessels and the heart. 1997; 205(Radiology):901–901.

Karmonik C, Basto P, Vickers K, Martin K, Reardon MJ, Lawrie GM, Morrisett JD. Quantitative segmentation of principal carotid atherosclerotic lesion components by feature space analysis based on multicontrast MRI at 1.5 T. IEEE Transactions on Biomedical Engineering. 2009; 56(2): 352–360. [PubMed: 19272944]

Ketcham RA, Carlson WD. Acquisition, optimization and interpretation of X-ray computed tomographic imagery: applications to the geosciences. Computers & Geosciences. 2001; 27(4): 381–400.

Kremkau FW, Taylor KJ. Artifacts in ultrasound imaging. J Ultrasound Med. 1986; 5(4):227–237. [PubMed: 3514956]

Lin CL, Videla AR, Miller JD. Advanced three-dimensional multiphase flow simulation in porous media reconstructed from X-ray microtomography using the He-Chen-Zhang Lattice Boltzmann Model. Flow Measurement and Instrumentation. 2010; 21(3):255–261.

Liu H, Krishnan S, Marella S, Udaykumar HS. Sharp interface Cartesian grid method II: a technique for simulating droplet interactions with surfaces of arbitrary shape. Journal of Computational Physics. 2005; 210(1):32–54.

Liu HL, Lin CL, Sun MT, Chang YH. 3D micro-crack propagation simulation at enamel/adhesive interface using FE submodeling and element death techniques. Annals of Biomedical Engineering. 2010; 38(6):2004–2012. [PubMed: 20414728]

Lloyd SP. Least-squares quantization in Pcm. IEEE Transactions on Information Theory. 1982; 28(2): 129–137.

Lorensen WE, Cline HE. Marching cubes: a high resolution 3D surface construction algorithm. Computer Graphics. 1987; 21(4):163–169.

Lynch, K.; Fahringer, T.; Thurow, B. Three-dimensional particle image velocimetry using a plenoptic camera. American Institute of Aeronautics and Astronautics; Conference Proceedings published by Curran Associates Inc Red Hook; New York, NY. January 9-12; Reston, VA: 2012.

Ma WY, Manjunath BS. Edge flow: a framework of boundary detection and image segmentation. Proceedings, 1997 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. 1997; (June 17-19):744–749.

Macqueen JB. Some methods for classification and analysis of multivariate observations. Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability, I: Statistics. 1967:281–297.

Makowski P, Sorensen TS, Therkildsen SV, Materka A, Stodkilde-jorgensen H, Pedersen EM. Two-phase active contour method for semiautomatic segmentation of the heart and blood vessels from MRI images for 3D visualization. Computerized Medical Imaging and Graphics. 2002; 26(1):9–17. [PubMed: 11734369]

Marella S, Krishnan S, Liu H, Udaykumar H. Sharp interface Cartesian grid method I: an easily implemented technique for 3D moving boundary computations. Journal of Computational Physics. 2005a; 210(1):1–31.

Marella S, Krishnan S, Liu H, Udaykumar HS. Sharp interface Cartesian grid method I: an easily implemented technique for 3D moving boundary computations. Journal of Computational Physics. 2005b; 210(1):1–31.

Mittal R, Iaccarino G. Immersed boundary methods. 2005; 37(Annual Review of Fluid Mechanics): 239–261.

Mumford D, Shah J. Optimal approximations by piecewise smooth functions and associated variational-problems. Communications on Pure and Applied Mathematics. 1989; 42(5):577–685.

Oh, BM.; Chen, M.; Dorsey, J.; Durand, F. Image-based modeling and photo editing. Proceedings of the 28th annual conference on Computer graphics and interactive techniques; Los Angeles, CA. August 12-17; ACM; 2001.

Osher S, Fedkiw RP. Level set methods: an overview and some recent results. Journal of Computational Physics. 2001; 169(2):463–502.

Osher S, Sethian JA. Fronts propagating with curvature-dependent speed – algorithms based on Hamilton-Jacobi formulations. Journal of Computational Physics. 1988; 79(1):12–49.

Pal NR, Pal SK. A review on image segmentation techniques. Pattern Recognition. 1993; 26(9):1277–1294.

Pappas TN. An adaptive clustering-algorithm for image segmentation. IEEE Transactions on Signal Processing. 1992; 40(4):901–914.

Perona P, Malik J. Scale-space and edge-detection using anisotropic diffusion. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1990; 12(7):629–639.

Pieczynski W. Statistical image segmentation. Machine Graphics and Vision. 1992; 1(1-2):261–268.

Pollo C, Cuadra MB, Cuisenaire O, Villemure JG, Thiran JP. Segmentation of brain structures in presence of a space-occupying lesion. Neuroimage. 2005; 24(4):990–996. [PubMed: 15670676]

Rollett AD, Lee SB, Campman R, Rohrer GS. Three-dimensional characterization of microstructure by electron back-scatter diffraction. Annual Review of Materials Research. 2007; 37:627–658.

Sambasivan SK, Udaykumar HS. A sharp interface method for high-speed multi-material flows: strong shocks and arbitrary material pairs. International Journal of Computational Fluid Dynamics. 2011; 25(3):139–162.

Sethian, J. Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science. Cambridge University Press; New York, NY.: 1999.

Sethian JA. Evolution, implementation, and application of level set and fast marching methods for advancing fronts. Journal of Computational Physics. 2001; 169(2):503–555.

Steinman DA, Milner JS, Norley CJ, Lownie SP, Holdsworth DW. Image-based computational simulation of flow dynamics in a giant intracranial aneurysm. AJNR Am J Neuroradiol. 2003; 24(4):559–566. [PubMed: 12695182]

Sun QL, Hossack JA, Tang JS, Acton ST. Speckle reducing anisotropic diffusion for 3D ultrasound images. Computerized Medical Imaging and Graphics. 2004; 28(8):461–470. [PubMed: 15541953]

Tai CH, Bals B, Zhao Y, Liew KM. An efficient parallel computation of unsteady incompressible viscous flow with elastic moving and compliant boundaries on unstructured grids. International Journal for Numerical Methods in Engineering. 2005; 64(15):2072–2104.

Terzopoulos D, Platt J, Barr A, Fleischer K. Elastically deformable models. Computer Graphics. 1987; 21(4):205–214.

Tran L, Udaykumar HS. Simulation of void collapse in an energetic material, Part 1: inert case. Journal of Propulsion and Power. 2006a; 22(5):947–958.

Tran L, Udaykumar HS. Simulation of void collapse in an energetic material, Part 2: reactive case. Journal of Propulsion and Power. 2006b; 22(5):959–974.

Tytell ED, Lauder GV. The hydrodynamics of eel swimming: 1. wake structure. Journal of Experimental Biology. 2004; 207(11):1825–1841. [PubMed: 15107438]

Udaykumar H, Mittal R, Rampunggoon P. Interface tracking finite volume method for complex solid-fluid interactions on fixed meshes. Communications in Numerical Methods in Engineering. 2002; 18(2):89–97.

Udaykumar H, Mittal R, Rampunggoon P, Khanna A. A sharp interface cartesian grid method for simulating flows with complex moving boundaries. Journal of Computational Physics. 2001; 174(1):345–380.

Udaykumar HS, Krishnan S, Marella S. Adaptively refined, parallelised sharp interface Cartesian grid method for three-dimensional moving boundary problems. International Journal of Computational Fluid Dynamics. 2009; 23(1):1–24.

Van Pelt R, Nguyen H, Romeny BT, Vilanova A. Automated segmentation of blood-flow regions in large thoracic arteries using 3D-cine PC-MRI measurements. International Journal of Computer Assisted Radiology and Surgery. 2012; 7(2):217–224. [PubMed: 21779767]

Venel Y, Garhi H, De Muret A, Baulieu JL, Barillot I, Prunier-Aesch C. Comparison of six methods of segmentation of tumor volume on the F-18-FDG PETscan with reference histological volume in nonsmall cell bronchopulmonary cancers. Medecine Nucleaire-Imagerie Fonctionnelle Et Metabolique. 2008; 32(6):339–353.

Vese LA, Chan TF. A multiphase level set framework for image segmentation using the Mumford and Shah model. International Journal of Computer Vision. 2002; 50(3):271–293.

Vigmostad SC, Udaykumar HS, Lu J, Chandran KB. Fluid-structure interaction methods in biological flows with special emphasis on heart valve dynamics. International Journal for Numerical Methods in Biomedical Engineering. 2010; 26(Nos 3-4):435–470.

Vignoles GL. Image segmentation for phase-contrast hard X-ray CMT of C/C composites. Carbon. 2001; 39(2):167–173.

Wang Z, Bovik AC, Sheikh HR, Simoncelli EP. Image quality assessment: from error visibility to structural similarity. IEEE Transactions on Image Processing. 2004; 13(4):600–612. [PubMed: 15376593]

Yu YJ, Acton ST. Speckle reducing anisotropic diffusion. IEEE Transactions on Image Processing. 2002; 11(11):1260–1270. [PubMed: 18249696]

Yue ZQ, Chen S, Tham LG. Finite element modeling of geomaterials using digital image processing. Computers and Geotechnics. 2003; 30(5):375–397.

Zhang S. Recent progresses on real-time 3D shape measurement using digital fringe projection techniques. Optics and Lasers in Engineering. 2010; 48(2):149–158.

Zhang X. Computation of viscous incompressible flow using pressure correction method on unstructured Chimera grid. International Journal of Computational Fluid Dynamics. 2006; 20(9): 637–650.

Zhang YL, Feng XT. Extended finite element simulation of crack propagation in fractured rock masses. Materials Research Innovations. 2011; 15(1):S594–S596.

Zhong JM, Sun HF. Wavelet-based multiscale anisotropic diffusion with adaptive statistical analysis for image restoration. IEEE Transactions on Circuits and Systems I-Regular Papers. 2008; 55(9): 2716–2725.

Zhou W, Bovik AC. A universal image quality index. IEEE Signal Processing Letters. 2002; 9(3):81–84.

Web reference

[20 October 2008] Philips Ultrasound Image Library. available at: www3.medical.philips.com/en-us/secure/images_site/largeImage.asp?size=blowup&classcode=03&appcode=a&imagename=0084-HD11-C5-2-ABD&systemcode=c&div=ultra

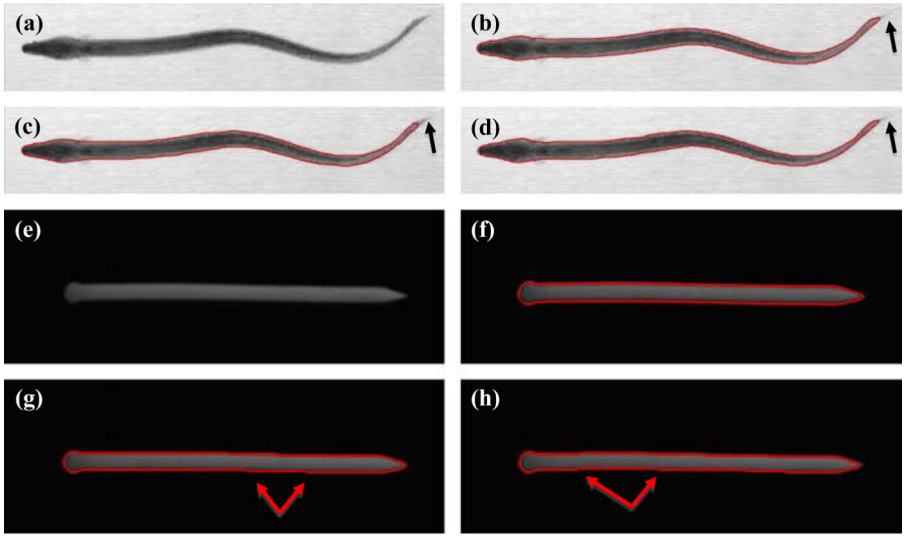

Figure 1.

Top two rows: a single frame of the American eel video. Bottom two rows: a single frame of the flexible plunging panel video

Notes: (a) Active contours segmentation; (b) *k*-means; (c) adaptive clustering; (d) arrows indicate portions of the tail that were not properly segmented due to low contrast and resolution. Note that adaptive clustering resulted in a narrower segmentation due to the lack of biasing toward one side of the body's outline or the other. Bottom two rows: A single frame of the flexible plunging panel video; (e) active contours segmentation; (f) *k*-means; (g) and adaptive clustering; and (h) arrows indicate “stair-stepping,” where the clustering methods were unable to smoothly follow the slight curvature of the panel

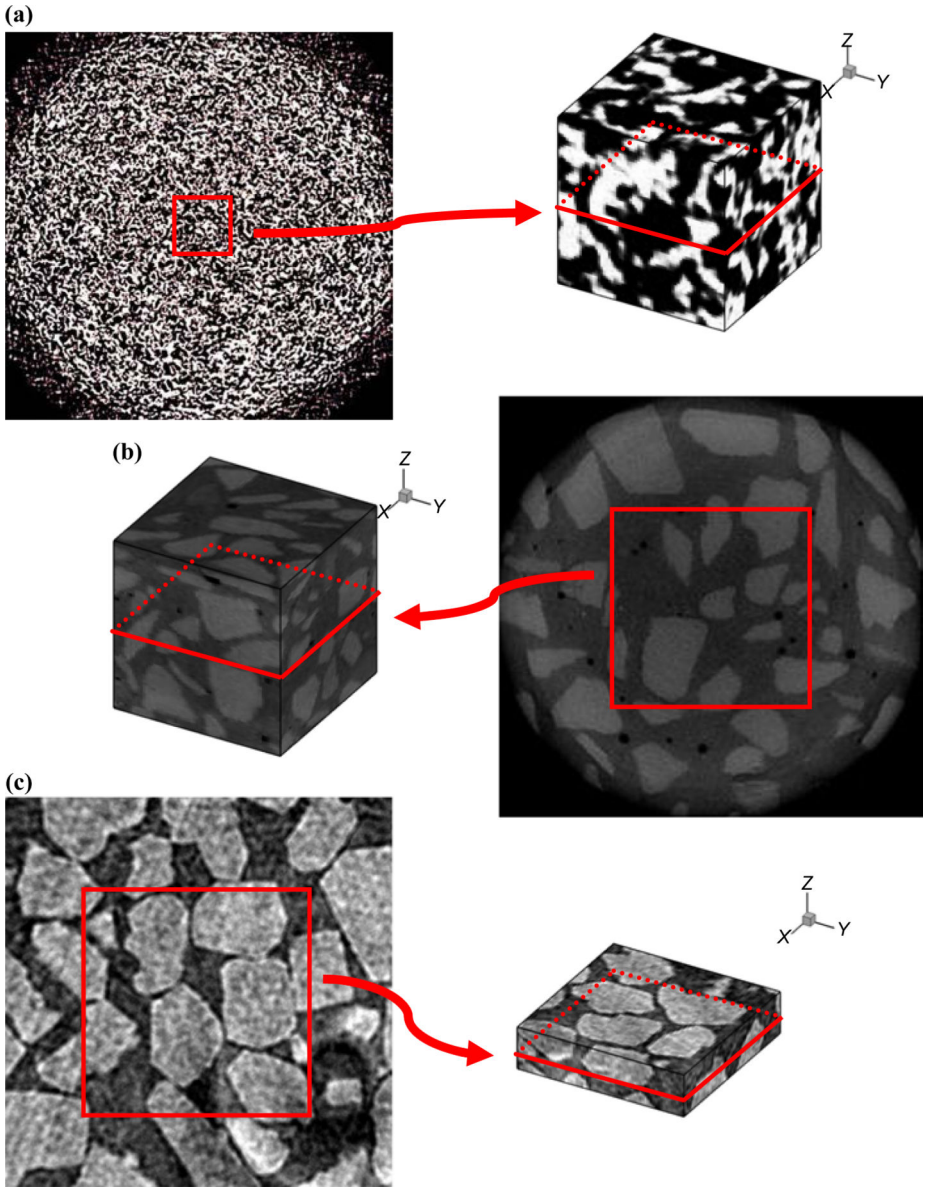

Figure 2.

Representative original image slices and cropped volumes from the 3D microstructure data sets provided

Notes: (a) Aluminum foam; (b) concrete; and (c) sugar crystals

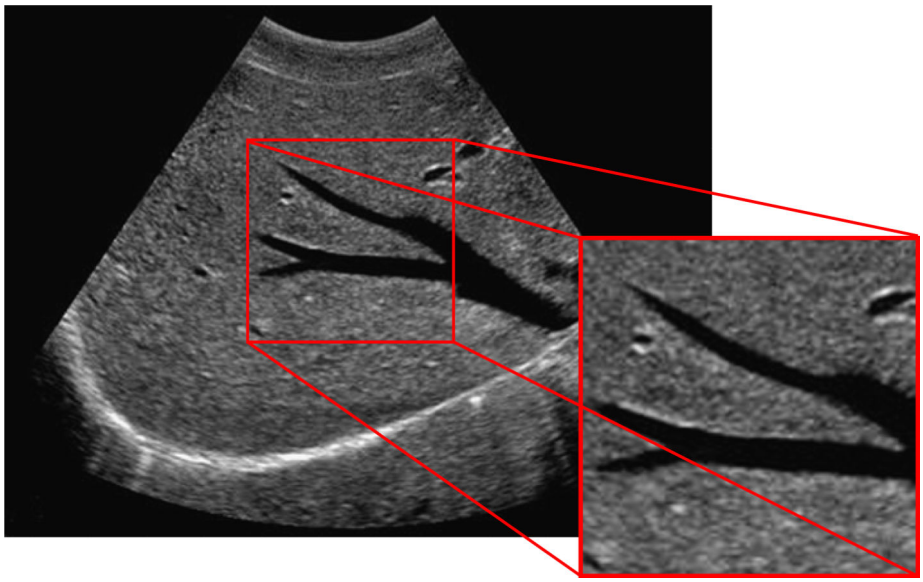

Figure 3.
The liver ultrasound image was cropped to a 128 × 128 pixel region prior to segmentation

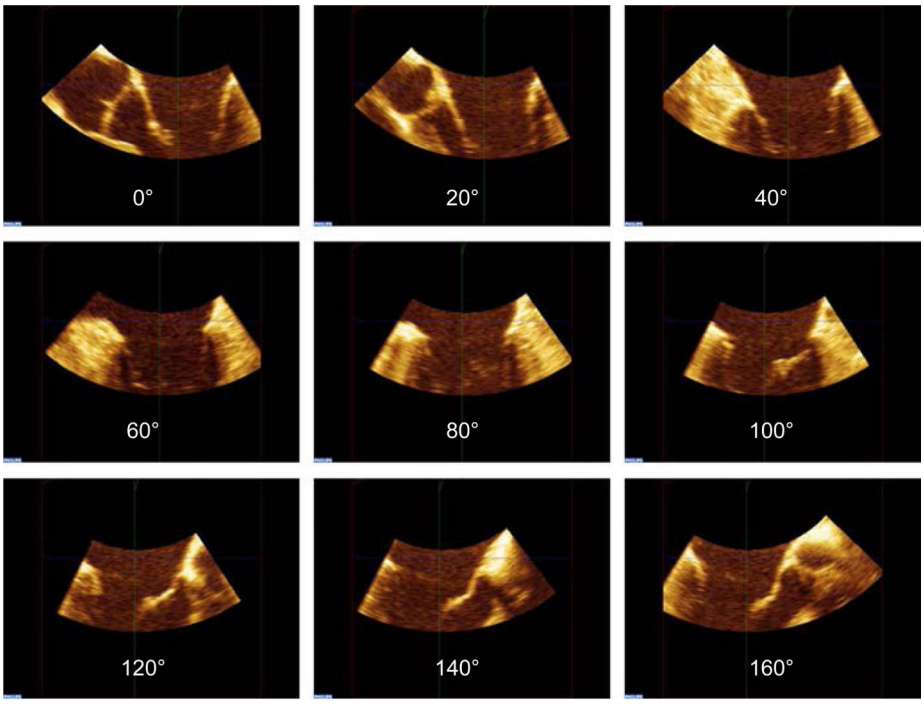

Figure 4.

A sampling of ultrasound image slices in cylindrical polar coordinates

Note: 18 slices in 10° increments make up the entire set

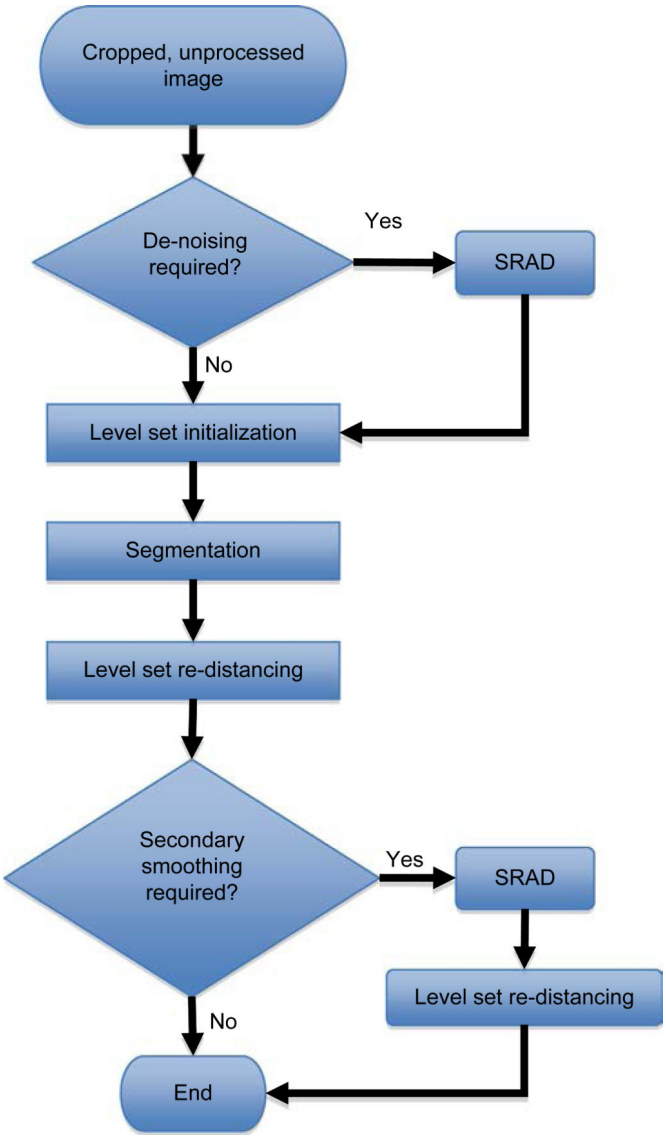

Figure 5.
Flowchart illustrating the image denoising and segmentation process

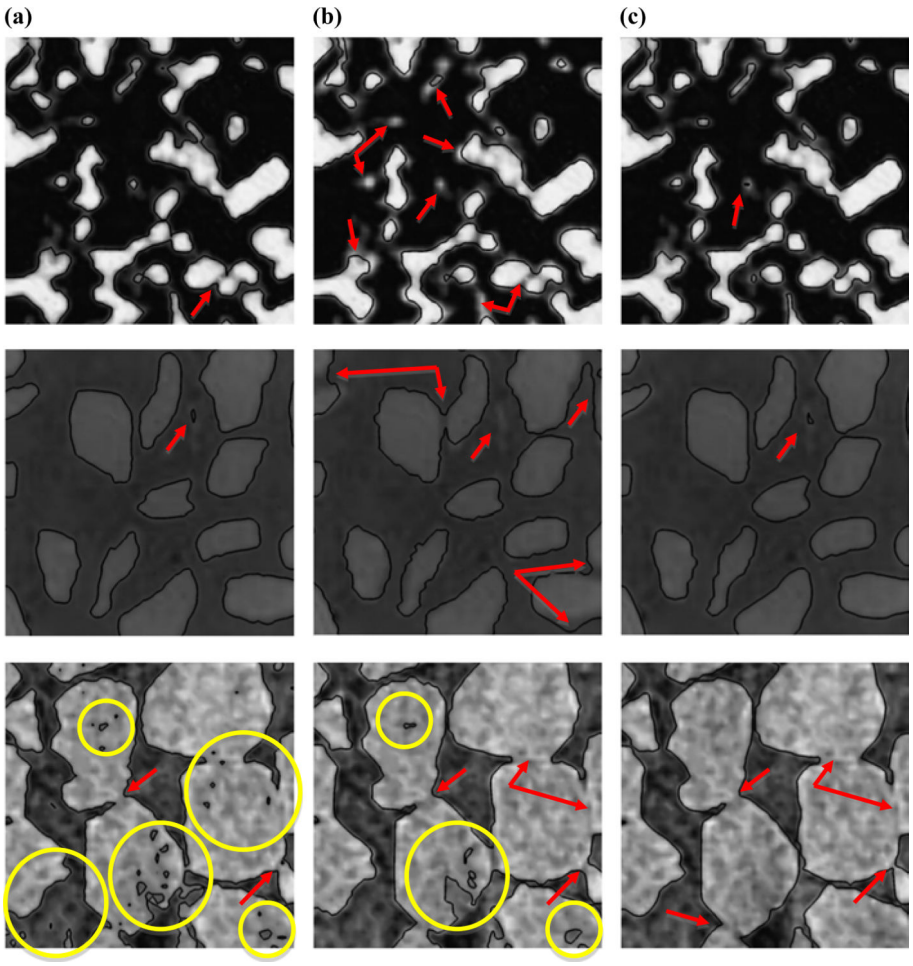

Figure 6.

Segmentation results for aluminum foam (top row), concrete (middle row), and sugar crystals (bottom row)

Notes: (a) Active contours: aluminum foam (μ=0.1, v=0.0, λ_1=1.0, λ_2=1.1), concrete (μ=1.0, v=0.0, λ_1=1.0, λ_2=1.1), and sugar crystals (μ=5.0, v=0.0, λ_1=1.0, λ_2=1.25); (b) *k*-means clustering: aluminum foam (λ_1=1.0, λ_2=1.2, post-smoothing SRAD α=0.15), concrete (λ_1=1.0, λ_2=1.2, post-smoothing SRAD α=0.1), and sugar crystals (λ_1=1.0, λ_2=1.1, post-smoothing SRAD α=0.2); (c) adaptive clustering: aluminum foam (α=0.05, $\Delta\beta$=0.01, post-smoothing SRAD α=0.1), concrete (β=0.001, $\Delta\beta$=0.001, post-smoothing SRAD α=0.2), and sugar crystals (β=0.1, $\Delta\beta$=0.001, post-smoothing SRAD α=0.3). Arrows and circles highlight errors present in each segmentation result

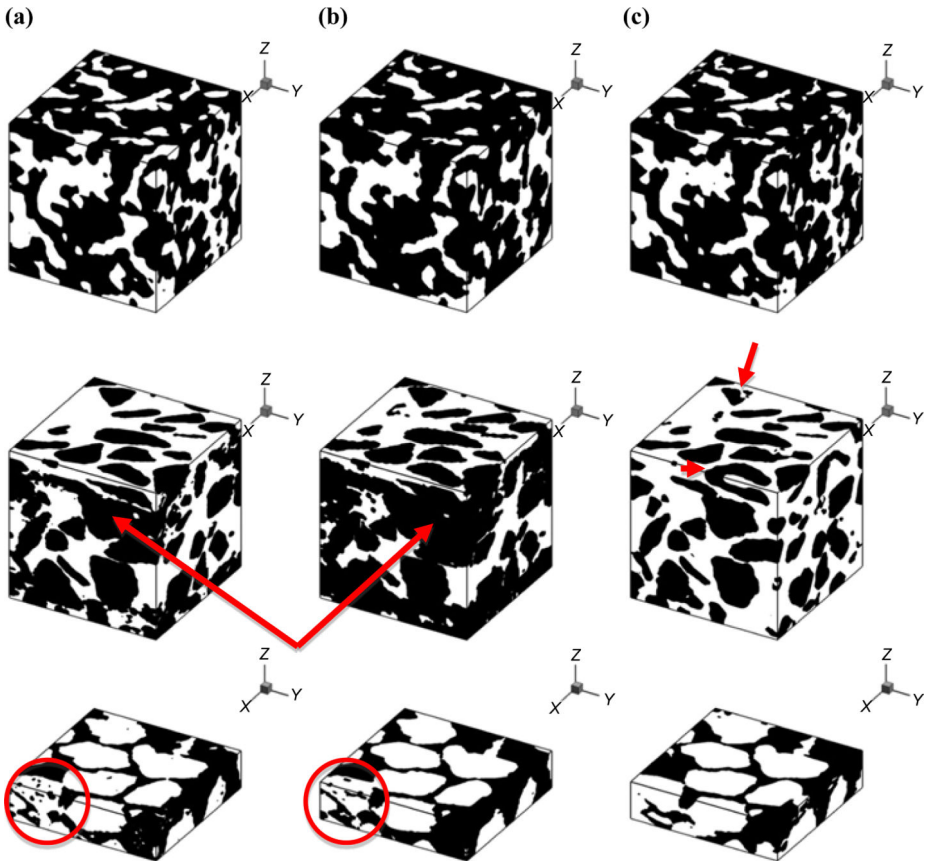

Figure 7.

Segmentation results illustrated on the faces of the 3D microstructure image volumes

Notes: Aluminum foam (top row), concrete (middle row), and sugar crystals (bottom row). Each was segmented by (a) active contours; (b) *k*-means; and (c) adaptive clustering. Arrows and circles highlight segmentation errors

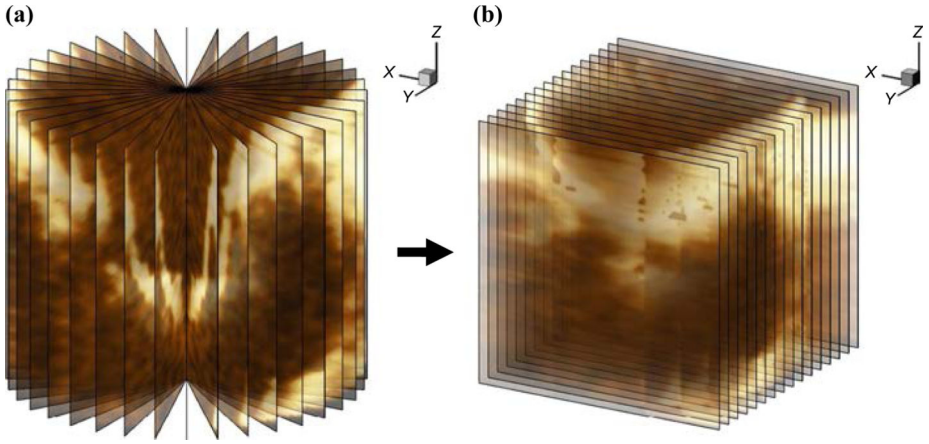

Figure 8.
Cropped image slices in cylindrical polar coordinates
Notes: (a) Are interpolated onto a regular Cartesian mesh; (b) for segmentation

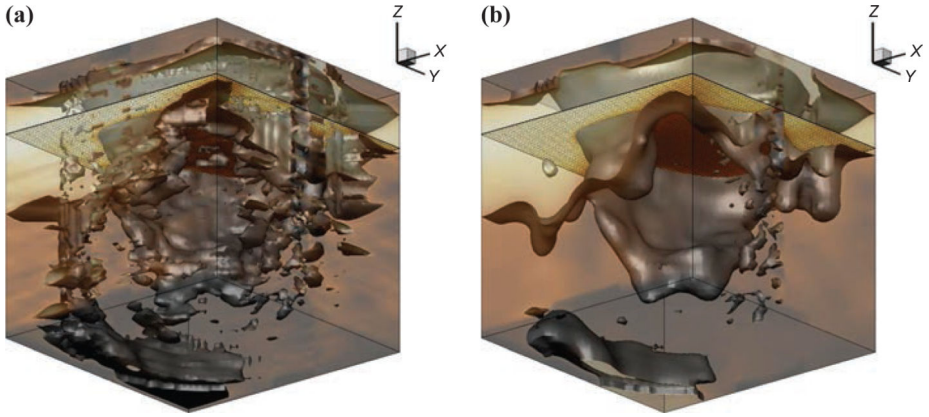

Figure 9.

Noise in the original ultrasound image, illustrated with the average image intensity isocontour

Notes: (a) Was reduced using SRAD; (b) prior to segmentation

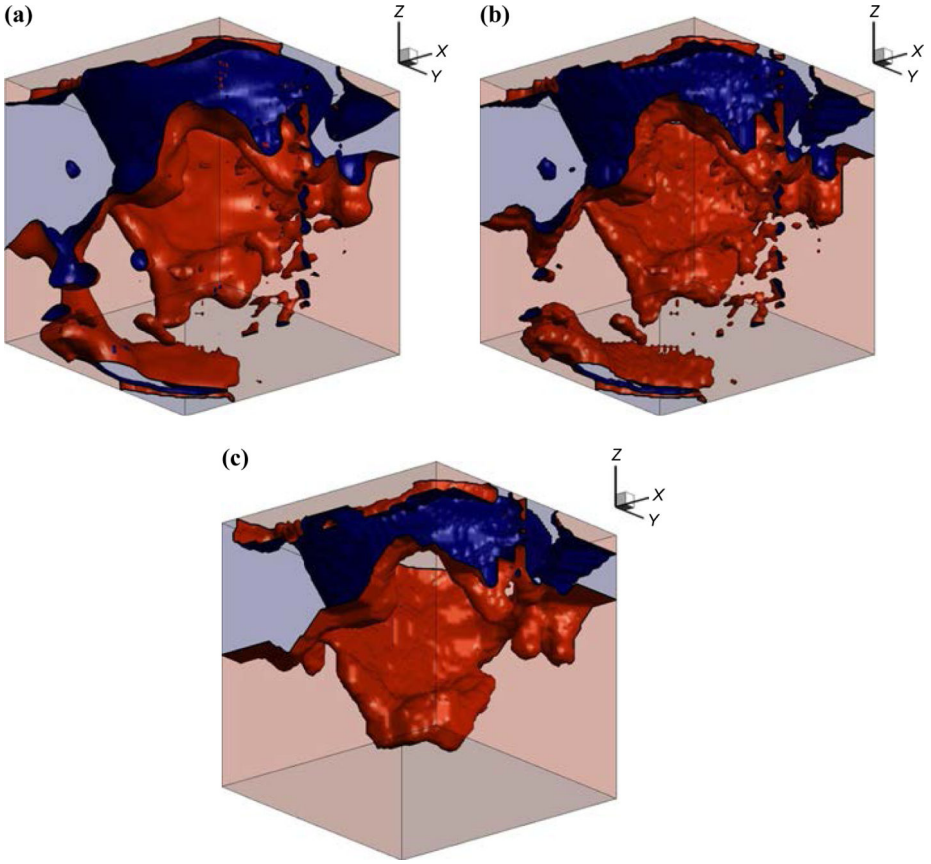

Figure 10.

3D ultrasound mitral valve segmentation results prior to secondary smoothing with SRAD

Notes: (a) Active contours; (b) *k*-means; and (c) adaptive clustering

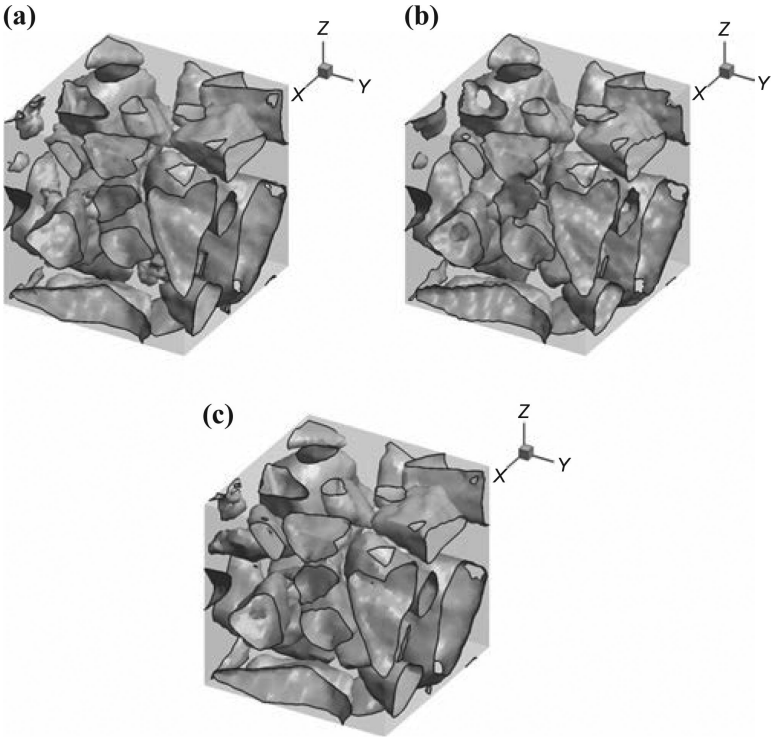

Figure 11.

Close-up views of the interior region of the concrete data set

Notes: (a) Active contours; (b) *k*-means; and (c) adaptive clustering

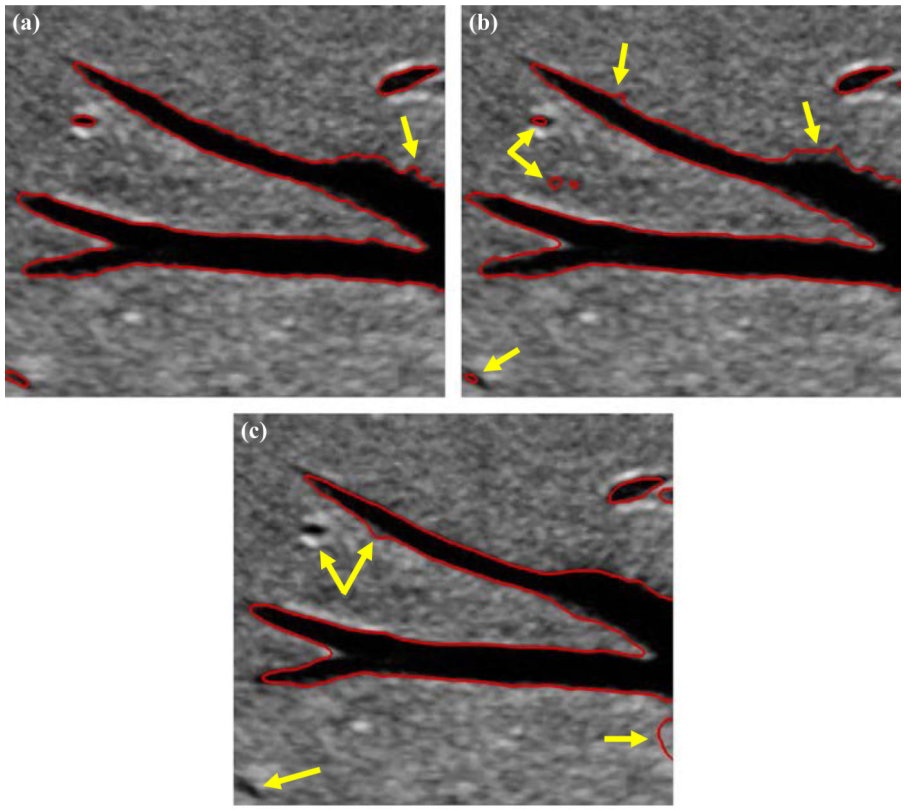

Figure 12.
2D liver ultrasound segmentation results
Notes: (a) Active contours; (b) *k*-means; (c) adaptive clustering. Arrows highlight segmentation errors

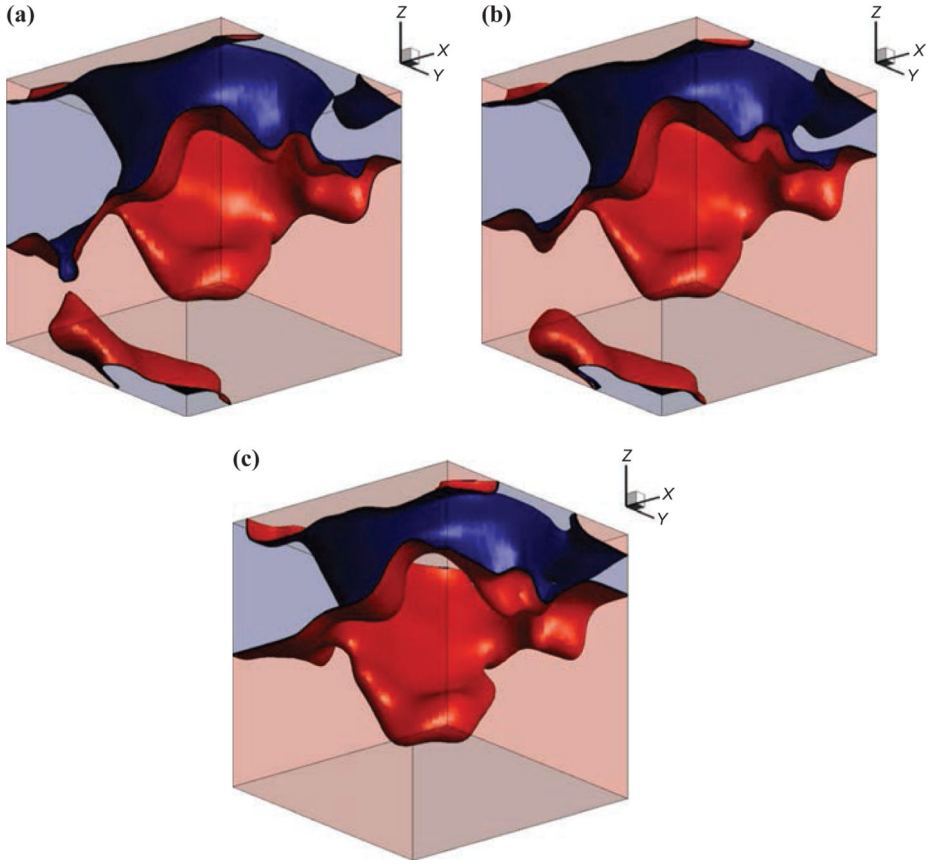

Figure 13.

3D ultrasound mitral valve segmentation results after secondary smoothing with SRAD

Notes: (a) Active contours; (b) *k*-means; and (c) adaptive clustering

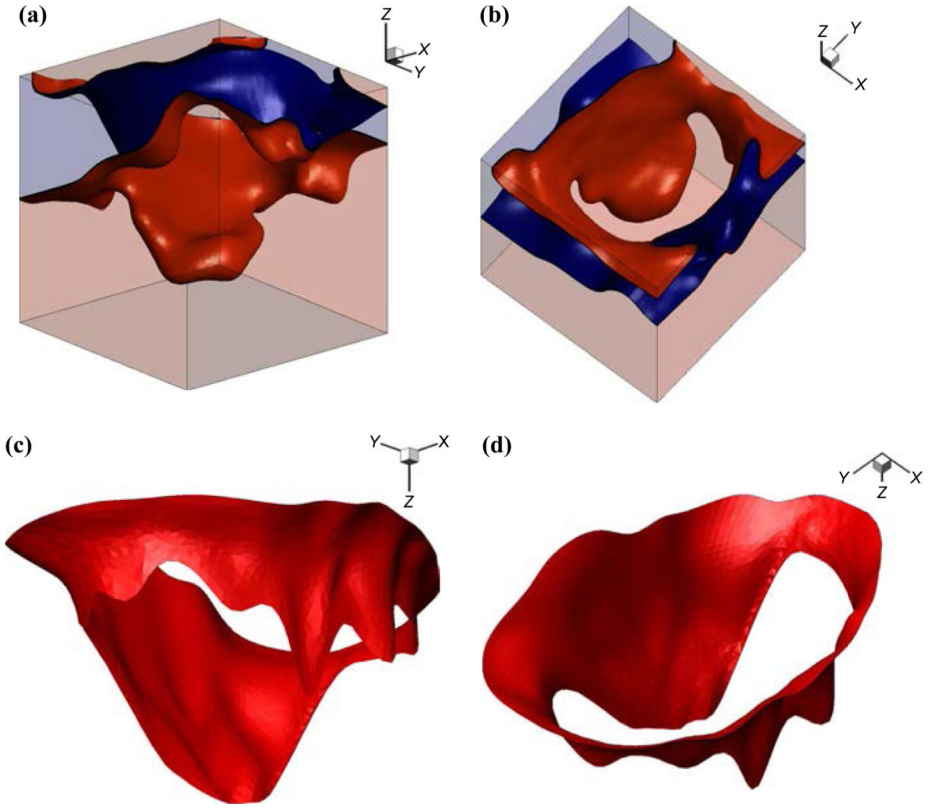

Figure 14.
Two views of the mitral valve segmented with adaptive clustering (a, b), compared with a finite element surface constructed by hand-segmentation (c, d)

Table I

Parameters selected for each image segmentation case

Image	Active contours	*k*-Means	Adaptive clustering	Pre-smoothing
Optical				
Eel	$\lambda_1 = 1, \lambda_2 = 1.1$ $\mu = 10$	$\lambda_1 = 1, \lambda_2 = 1.1$ $\alpha = 3.5$	$\sigma = 16$ $\alpha = 4$	$\alpha = 3.5$
Panel	$\lambda_1 = \lambda_2 = 1$ $\mu = 10$	$\lambda_1 = \lambda_2 = 1$ $\alpha = 8$	$\sigma = 8$ $\alpha = 8$	–
X-ray CT				
Al foam	$\lambda_1 = 1, \lambda_2 = 1.1$ $\mu = 0.1$	$\lambda_1 = 1, \lambda_2 = 1.2$	$\beta = 0.05$, $\Delta\beta = 0.01$	–
Concrete	$\lambda_1 = 1, \lambda_2 = 1$ $\mu = 1$	$\lambda_1 = 1, \lambda_2 = 1.2$ $\alpha = 0.1$	$\beta = 0.001$, $\Delta\beta = 0.001$ $\alpha = 0.2$	$\alpha = 0.25$
Sugar crystals	$\lambda_1 = 1, \lambda_2 = 1$ $\mu = 5$	$\lambda_1 = 1, \lambda_2 = 1.1$ $\alpha = 0.25$	$\beta = 0.1$, $\Delta\beta = 0.01$ $\alpha = 0.3$	$\alpha = 0.25$
Ultrasound				
Liver	$\lambda_1 = 1, \lambda_2 = 1.5$ $\mu = 100$	$\lambda_1 = \lambda_2 = 1$ $\alpha = 1$	$\beta = 0.9$, $\Delta\beta = 0.2$ $\alpha = 1$	$\alpha = 1.5$
Mitral valve	$\lambda_1 = \lambda_2 = 10$ $\mu = 3$	$\lambda_1 = \lambda_2 = 1$ $\alpha = 3$	$\beta = 0.3$, $\Delta\beta = 0.05$ $\alpha = 3$	$\alpha = 3$

Notes: Pre-smoothing with SRAD was only applied when segmenting with active contours or *k*-Means, except in the 3D ultrasound image of the mitral valve, where it was applied for all three methods. Secondary smoothing was also applied for all three methods in the mitral valve case

Table II

Characterization of image quality

Image	μ_1	μ_2	μ_1/μ_2	Without SRAD μ/σ x_1	Without SRAD μ/σ x_2	Without SRAD $\frl{\lvert \mu_1 - \mu_2 \rvert}{\sigma}$	With SRAD μ/σ x_1	With SRAD μ/σ x_2	With SRAD $\frl{\lvert \mu_1 - \mu_2 \rvert}{\sigma}$
Optical									
Eel	220.17	124.27	1.77	31.62	5.02	6.88	38.36	5.08	7.11
Panel	71.96	4.23	17.01	3.40	4.97	5.13	–	–	–
X-ray CT									
Al foam	223.17	27.60	8.09	5.45	0.60	1.90	–	–	–
Concrete	89.11	59.40	1.50	8.88	5.34	1.68	–	–	–
Sugar crystals	163.03	69.31	2.35	5.17	2.32	1.71	–	–	–
Ultrasound									
Liver	106.04	9.71	10.92	4.74	0.60	2.27	8.47	0.63	2.46
Mitral valve	197.28	90.78	2.17	4.64	2.71	1.96	5.31	3.65	1.99

Notes: Each image is separated into two regions – a brighter region x_1 with mean brightness intensity μ_1, and a darker region x_2 with mean brightness intensity μ_2. The two regions are delineated by the segmentation contour judged to give the best results in each case: active contours with SRAD pre-smoothing for the optical images and 2D liver ultrasound image, and adaptive clustering for the others. (SRAD pre-smoothing was also used for the 3D mitral valve ultrasound image.) Image quality is measured in terms of contrast ratio μ_1/μ_2, the signal-to-noise ratio μ/σ in each region, and the overall contrast-to-noise ratio $|\mu_1 - \mu_2|/\sigma$