# Surface Estimation in Ultrasound Images

TR91-007
January, 1991

Wei-Jyh Lin
Stephen M. Pizer
Valen E. Johnson

Medical Image Display Group
Department of Computer Science
The University of North Carolina
Chapel Hill, NC 27599-7510

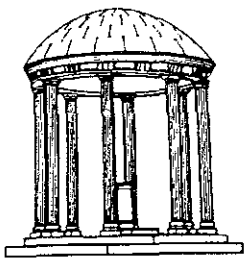

*The research reported herein was carried out with the partial support of NIH grant number P01CA47982 and Duke ERC Grant.*
*To appear in Proceeding Information Processing in Medical Image (IPMI XII), 1991.*

# Surface Estimation in Ultrasound Images

Wei-Jyh Lin
Department of Computer Science
University of North Carolina
Chapel Hill, NC 27599-3175

Stephen M. Pizer
Departments of Computer Science,
Radiology, and Radiation Onchology
University of North Carolina
Chapel Hill, NC 27599-3175

Valen E. Johnson
Institute of Decision Science and Statistics
Duke University
Durham, NC 27706

January 27, 1991

**Abstract**

Surface definition, a process of defining three dimensional surface from volume data, is essential in three dimensional volume data rendering. The traditional method applies a three dimensional gradient operator to the volume data to estimate the strength and orientation of surface present. Applying this method to ultrasound volume data does not produce satisfactory results due to noisy nature of the images and the sensitivity of certain signals to the direction of insonation. A Bayesian approach is proposed here for surface definition of noisy images in general. We formulate the problem as the estimation of posterior means and standard deviations of Gibbs distributions for surface believability and normal direction. The prior distribution reflects shape properties at multiple scales. The design and implementation of such an approach and its application on ultrasound images are the subject of this paper.

## Introduction

Three dimensional arrays of digital data have been generated in many areas of medical imaging in ever increasing number. Multiple 2D slices of computed tomography (CT), magnetic resonance (MR), and single-photon emission computed tomography (SPECT) create volume data. A research project currently conducted in the Duke/UNC Engineering Research Center on Emerging Cardiovascular Technologies includes building a new generation transducer that can capture three dimensional volume of ultrasound data in real time [vonRamm88] [Shattuck84]. These volume data represent complex anatomy or functional process under study. Effective visualization of these volume data helps physicians in diagnostic interpretation or surgical planning.

Volume rendering, a method of direct rendering of volume data, has been successfully applied in visualizing volume data of CT, MRI, and PET images. The operational principle of volume rendering is to render the volume data directly instead of fitting geometric primitives and then rendering the primitives. This direct rendering is done by compositing images from the results of two separate and parallel processes. The first process performs surface shading at every voxel of volume data with a locally computed surface normal. The second process performs surface classification to obtain a partial opacity for every voxel. Independence of shading and classification calculations insures an undistorted visualization of 3D shapes. Non-binary classification increases the likelihood that small or poorly defined features are perserved [Levoy88].

Successful application of volume rendering depends heavily on the estimation of the local surface normal and surface classification. Usually one tries to take advantage of knowledge about the relationship of voxel values with surfaces and their normals. For example, in CT image studies one can make the assumption that CT numbers represent the percentages of material contained in voxels; hence the surface normal and surface classification can be obtained from the local gradient of an approximate percentage measure. Unfortunately this simple technique can not be satisfactorily applied to ultrasound volume data since ultrasound images suffer from serious speckle phenomena due to the coherent radiation source. Common speckle phenomena include random speckles spots from within soft tissues and broken contours on organ boundaries. Applying the above simple classification technique serves to pick up boundaries of random speckle spots as surfaces and miss boundaries at contour gaps. The locally computed surface normals tend to be incorrect.

In this paper we study the surface classification and normal estimation problem from a Bayesian perspective in general. We formulate the problem as the estimation of posterior means and standard deviations of Gibbs distributions for surface believability and normal direction. We show that the Gibbs distribution can be extended to model global structures by using a data augmentation scheme. We apply this method on ultrasound images in two dimensions and show the results. The remainder of this paper is organized as follows. The next section presents the Bayesian approach. The following two sections describe our filter design for producing ultrasound image edge-related measurements to which Gibbs-compatible likelihood functions pertain. Then the design of Gibbs priors reflecting shape-related knowledge about the variables to be extracted is presented, followed by a summary of the complete algorithm. We then present results and close with a discussion.

## Approach

Given an observed image we can apply measurements on the voxel values to determine whether a given voxel is on a boundary and, if it is, the associated normal direction. A representation of the target boundaries and normal directions can then be determined from

the measurements. However, this representation will not always be the true representation of the boundaries and normal directions of underlying targets due to the fact that the image is contaminated by noise. Instead, there is a certain strength of conviction or believability associated with this representation.

The believability is obviously dependent on the measurements and is generally different in different parts of an image. We usually have strong conviction as to the representation of a part of an image when the measurements on that part show relatively high values and weak conviction when the measurement values are low. For example, the output of an edge detector affects our conviction about the presence of an edge – the higher the output value the stronger the conviction. This conviction is also affected by the measurements in the neighboring voxels. For example, the believability concerning the existance of a boundary at two neighboring voxels increases when their normal direction measurements confirm each other but decreases when the normal directions contradict each other.

Putting this dependency of believability on measurements and consistency into a mathematical form produces a Bayesian formulation for the conviction. In other words, the believability is really the posterior odds of a representation given measurements. It depends on a likelihood function, which is a distribution of measurements conditioned on a boundary representation, and a prior, which models the consistency in a representation, that is, the geometry of the situation.

Under the Bayesian framework we can define a random field with the same dimensions as the input image. Each site of the random field has two components: a *boundary value*, whose value is either zero or one, representing whether the corresponding voxel is on a boundary or not, and a *normal* which is a unit vector in the half unit-sphere of positive z component. A representation of object boundary and normal direction is thus a sample drawn from the random field. In regard to this representation there is an associated posterior odds between zero and one, and the value is monotonic with the believability of this representation.

We compute the posterior means and standard deviations of boundary value and normal directions for each voxel. This produces a summary representation of the ensemble of all possible representations. This summary representation provides information useful for volume-rendering. For example, the mean of the normal direction of a voxel can be used as the surface normal for shading. The standard deviation of normal direction can be used to make the mean normal direction fuzzy before it is used as the surface normal. As a result a shiny surface might show a high confidence about the local surface normals while a dull surface showed a low confidence. The mean of the boundary value could be used to modulate the opacity of the voxel, and the standard deviation of the boundary value could be used

to modulate the color. Here the color of a surface would show the relative confidence about our judgement of the presence of a local surface. Ideally this creates an image portraying the target shapes and with additional visual information of confidence about the shapes.

Later we will show that under certain assumptions this random field has the Markovian property. By the equivalence theorem of Markov random fields and Gibbs distributions due to Geman and Geman [Geman84], the posterior distributions of this random field is a Gibbs distribution. The Gibbs sampler can then be used to draw samples from the random field repeatedly. We thus can compute the ensemble posterior mean and standard deviation.

## Filter Design for Boundary Measurement

We take as our likelihood function the distributions of local image features conditioned on the presence or absence of boundaries and, in the case of boundary presence, the true normal direction. The features could be just the image intensity; or more often measurements from some filters. In any case, we wish to select the features that can help better discriminate true boundaries from false ones and select the true normal directions. In other words, we wish to have well separated feature distributions between the cases of boundary presence and boundary absence, with the feature distribution peaking at the true normal direction.

The selection of these features, and in turn the design of the filters, apparently will be application dependent. In the rest of this section and the next section we shall concentrate on ultrasound images in two dimensions. We shall discuss briefly the properties of ultrasound image noise and then introduce a set of filters for detecting boundary and associated normal directions based on those properties. We shall then present our measured distribution of the filter outputs and approximations to the filter output distributions.

Ultrasound imaging is an unusual imaging modality in that it uses a coherent radiation source. A sound beam is produced at a transducer and directed into targets. Echoes are generated at interfaces of media with different acoustic impedances along the path of sonic transmission. The incident beam is partially reflected back toward the transducer with the angle of reflection equal to the angle of incidence [HykesS85]. This type of reflection is responsible for major organ outlines seen in diagnostic ultrasound examinations. These reflections usually appear to be very bright in the image and they form mostly continuous boundaries. We shall call this kind of boundary the *specular boundary.* Specular boundaries tend to be broken where the boundary orientation is parallel to the beam direction and the reflection is directed away from the transducer, a phenomenon called *echo drop-outs.* Specular boundaries also appear to be broken where the random scatterers around the boundary happen to produce a destructive component echo which is out of phase with the specular reflection.

Strong echoes can also originate from within homogeneous organ tissues, like livers,

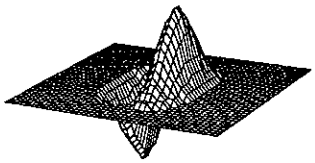

Figure 1: Filter shape

in the absence of apparent acoustic mismatch. These echoes appear in the image as a texture of random bright spots, called *speckles.* In fact, these echoes originate from the supporting fibers in these tissues because they have elasticity much different from their neighboring cells [Fields73]. The intensity distribution of these echoes usually approximates Rayleigh distributions with different mean intensities for different tissue types. The interface of different tissues also forms a boundary and sometimes can be visually detected in the image. We shall call this kind of boundary the *texture boundary.* We are interested in both specular boundaries and texture boundaries.

We wish to design a set of filters which provide accurate edge response and orientation for both specular and texture boundaries. We have selected directional derivatives of anisotropic gaussians as the filter kernels [Canny86], [Korn88]. The rationale for this is as follows. Suppose we know the local boundary orientation. We can compute the average intensities on the two sides of the hypothesized boundary. The absolute value of the difference of the two averages then is a good indication of the presence or absence of an edge. We would like to get as many samples as possible on either side to get reliable statistical averages. But at the same time we would like to avoid the problem of mixing intensities from other nearby image structures. The directional derivatives of anisotropic gaussians with the elongated side oriented along boundary serve these constraints. In our implementation we use such filters of fixed widths but increasing lengths at multiple scales. The purpose of more elongated filters, which will be detailed in the section describing the prior specificaton, is to aid closing broken boundaries of different gap sizes. The shape of such a filter is shown in figure 1.

Notice that these filters work for specular boundaries as well as texture boundaries. What are the correct sizes for the filter to be used in the likelihood function? For texture boundaries we argue that the short side should be at least two speckles wide and the long side should be several speckles in length in order to get good statistical averages. For

specular boundaries smaller filters can be used to get better localization of edges. In current implementation we used the same filter sizes for both kinds of boundaries, which is about 2 speckles for the short side and 2.6 speckles for the elongated side. (Strictly speaking, the filter sizes should also change with the orientations of edges since the speckles are themselves anisotropic and always lie perpendicular to the beam direction. )

In reality we do not know the orientations of boundaries, which is required for the filters to work correctly. We address this problem by applying the filters in several orientations and averaging the responses from all orientations. In the current implementation we use four orientations at 0, 45, 90, and 135 degrees. The response from each filter is taken to form a vector in the polar coordinate system with its corresponding angle. These vectors are transformed to a new vector by doubling their angles, averaging by vector summation, and then halving the angle of resulting vector. This angle doubling and halving before and after the vector summation averaging is necessary because we do not distinguish inner normals from outer normals.

## Empirical Approximation of the Likelihood Function

The likelihood function that we need to estimate is $p(f_b(i,j), f_n(i,j)|B(i,j), N(i,j))$, i.e. the filter response at voxel $(i,j)$ given the true boundary situation at voxel $(i,j)$. $f_b$ and $f_n$ are the magnitude and phase of the smallest scale filter response, $B(i,j)$ is either 1 or 0 representing boundary presence or not respectively, and $N(i,j)$ is the normal direction in the case of boundary presence. Both $f_n$ and $N(i,j)$ are modulated by 180 degrees and quantized to 16 discrete orientations in our implementation. Here we shall assume that $f_b(i,j)$ is independent of $N(i,j)$ and $f_n(i,j)$ is independent of $B(i,j)$. Furthermore, we shall assume that $f_n$ and $f_b$ are independent of each other. It then follows that we need only estimate $p(f_b(i,j)|B(i,j))$ and $p(f_n(i,j)|N(i,j))$ since $p(f_b(i,j), f_n(i,j)|B(i,j), N(i,j)) = p(f_b(i,j)|B(i,j))p(f_n(i,j)|N(i,j))$.

Three test images are used as the training set for determining these conditional distributions. These three images are a slice through a cone phantom simulating a diffuse target, a slice through a liver, and a slice of of a baby doll hanging in a water tank, as shown in figure 2 . Contours in these images have been manually drawn as reference ground truth.

The histograms of $f_b$ on the contours are computed and plotted in figure 3. Regions in the images containing no boundaries are randomly selected to gather samples for computing the histograms of $f_b$ in the absence of boundaries. These regions are as marked in figure 2 and the histograms are plotted in figure 3. These histograms can be reasonably fitted by normal distributions after a cube root transform on $f_b$. Figure 4 shows the transformed histograms and the fitting normal distributions. In these plots $f_b$ is linearly scaled to between zero and one by normalizing by the maximal $f_b$ in each image.

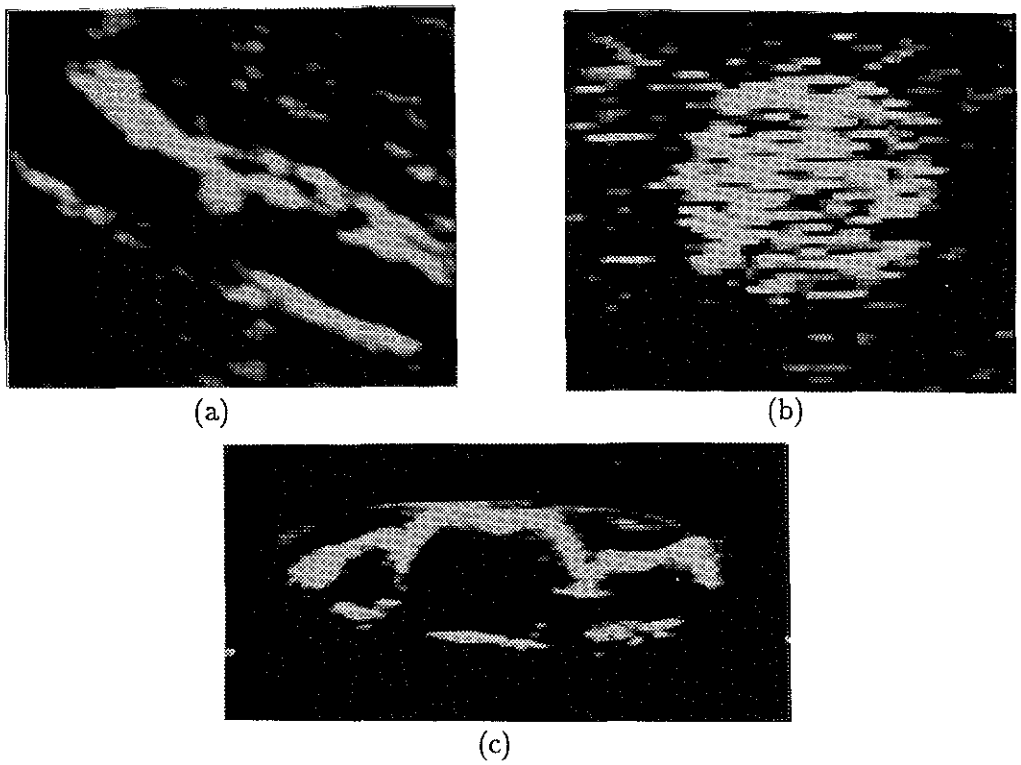


Figure 2: Test images: (a) Cone phantom, (b) Liver, and (c) Baby doll

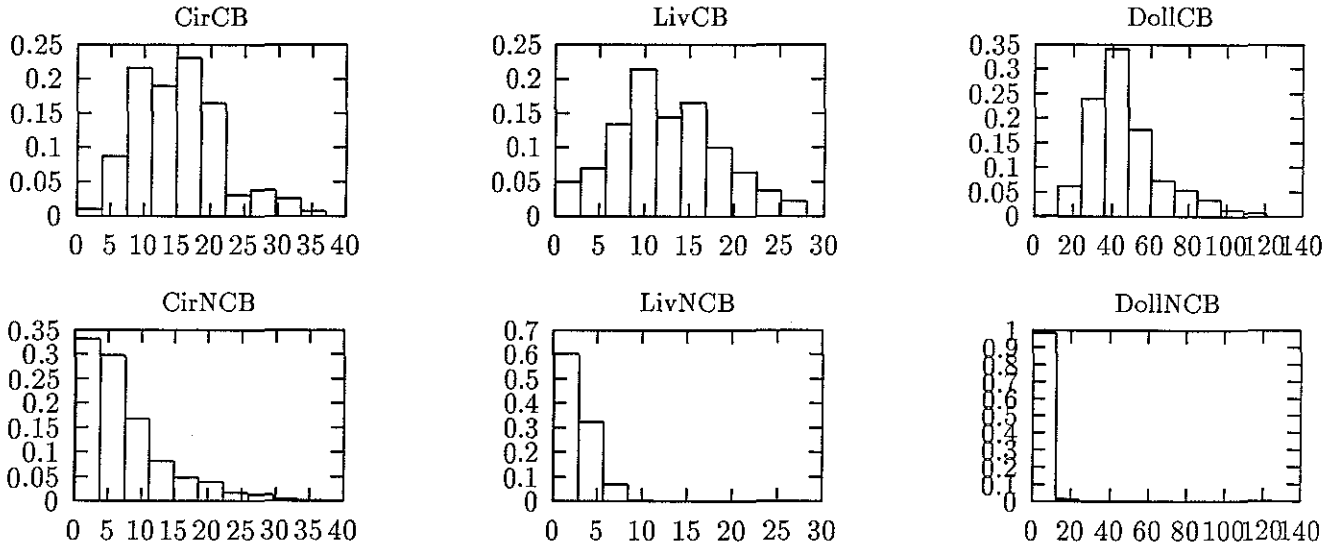


Figure 3: (a), (b), and (c): Histograms of $f_b$ measured on boundaries of cone phantom, liver, and baby doll. (d), (e), and (f): Histograms of $f_b$ measured on no boundary regions of cone phantom, liver, and baby doll

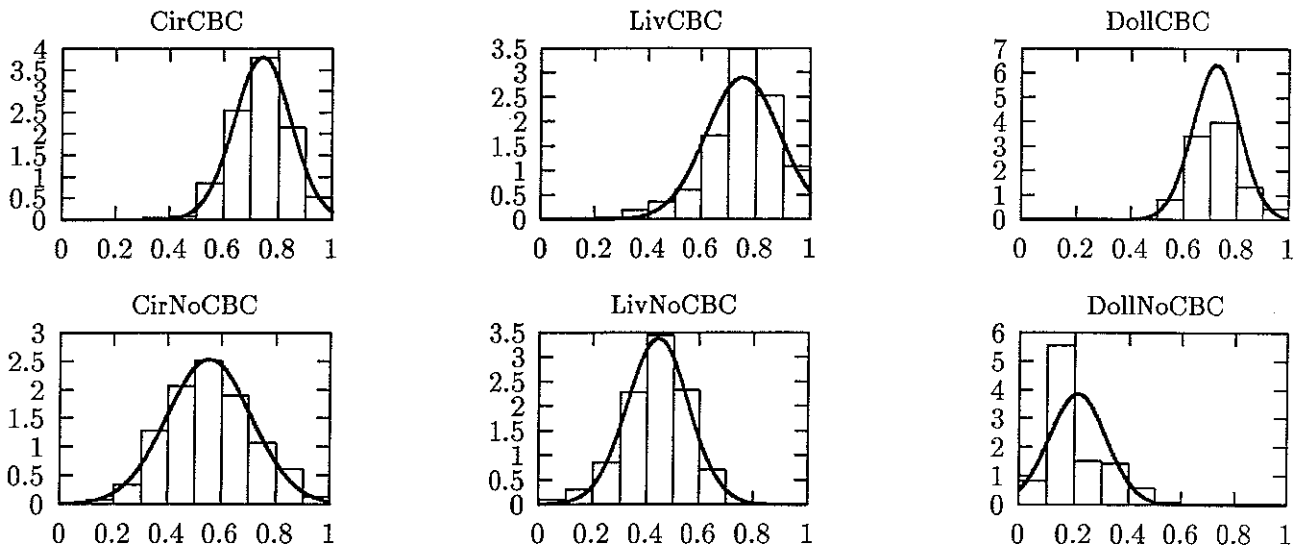


Figure 4: Histograms of $f_b$ after cube root transform and normalization and the fitting normal distributions. Graphs order is the same as in figure 3.

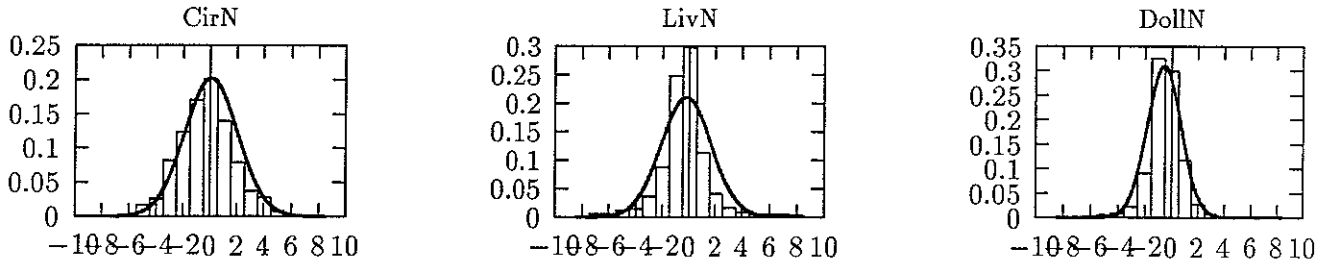


Figure 5: Histograms of $f_n$

$P(f_n(i,j)|N(i,j))$ is estimated by the difference of $f_n(i,j)$ and $N(i,j)$, i.e., $p(h(f_n(i,j),N(i,j)))$, where $h(\cdot,\cdot)$ is a distance measure of normal angles. N(i,j) is computed from the drawn boundaries by fitting a straight line through five neighboring boundary voxels with least-square error. Figure 5 shows the histograms of $h(f_n(i,j),N(i,j))$. These histograms all peak at zero, indicating that our filters give approximately unbiased normal directions. These histograms can also be approximated by normal distributions. In summary:

$$p(f_b(i,j),f_n(i,j)|B(i,j),N(i,j)) =$$
$$\begin{cases} \frac{1}{2\pi\sigma_{b1}\sigma_n}e^{-V_{c1}} & \text{where } V_{c1}=(f_b-\mu_{b1})^2/2\sigma_{b1}^2+h(f_n(i,j),N(i,j))^2/2\sigma_n^2 \quad \text{if } B(i,j)=1 \\ \frac{1}{(2\pi)^{1/2}\sigma_{b2}}e^{-V_{c1}} & \text{where } V_{c1}=(f_b-\mu_{b2})^2/2\sigma_{b2}^2 \quad \text{if } B(i,j)=0 \end{cases} \quad (1)$$

## Prior Specification

The prior defines the distribution of object shapes. Due to the enormous possibilities of object shapes projecting onto the random field it is difficult to directly specify the joint distribution of the large number of random components in the random field. However, since local object shapes do not have absolute correlation with distant parts of global shapes, it seems reasonable to assume a Markovian property for the random field. In other words, we assume that the values of a given site depend only on values of nearby sites.

Random fields having Markovian property are called Markov Random Fields (MRFs). Geman and Geman [Geman84] have shown that for a MRF an equivalent Gibbs distribution can be defined on the field. Let $S=\{s_1,s_2,\cdots,s_N\}$ be the set of grid points on a MRF. A *neighborhood system*, $\mathcal{G}=\{\mathcal{G}_s,s\in S\}$ for S, is any collection of subsets of $S$ for which 1) $s\notin\mathcal{G}_s$ and 2) $s\in\mathcal{G}_r\Leftrightarrow r\in\mathcal{G}_s$, where $\mathcal{G}_s$ is the set of neighbors of s. A subset $C\subseteq S$ is a *clique* if every pair of distinct sites in $C$ are neighbors. A *Gibbs distribution* relative to $\{S,\mathcal{G}\}$ is a probability measure $\pi$ on $\Omega$, the state space of the random field, with the following representation:

$$\pi(\omega)=\frac{1}{Z}e^{-U(\omega)} \quad (2)$$

where $\omega$ is an outcome of the random field, $Z$ is a constant and $U$, called the *energy function*, is of the form

$$U(\omega)=\sum_{C\in\mathcal{C}}V_C(\omega) \quad (3)$$

where $\mathcal{C}$ is the set of cliques for $\mathcal{G}$. Each $V_C$ is a function on $\Omega$ with the property that $V_C(\omega)$ depends only on those components $x_s$ of $\omega$ for which $s\in C$. Such a family $\{V_C,C\in\mathcal{C}\}$ is called a *potential*. $Z$ is the normalizing constant:

$$Z\triangleq\int_\omega e^{-U(\omega)} \quad (4)$$

and is called a *partition function*. For details of MRF and Gibbs distribution we refer readers to [Geman84].

We used a nearest-pair neighborhood system in modeling the local shape probability. There are four kinds of cliques in this neighborhood system: $(s_{i,j}, s_{i+1,j})$, $(s_{i,j}, s_{i,j+1})$, $(s_{i,j}, s_{i+1,j+1})$, and $(s_{i,j}, s_{i+1,j-1})$. For each kind of pair there are three combinations for the boundary values on the two grid points: both on boundaries, exactly one on a boundary, or neither on a boundary. For each case there are constraints on the normal directions for the points on boundaries: if both points are on boundaries the normal direction should change smoothly, while if only one point is on a boundary its normal direction should be oriented such that the assumed boundary direction does not point to the other point. The latter condition effectively penalizes blindly ending boundaries. Specifically, we have the following potentials for horizontal cliques:

$$\begin{aligned} V_{c2} &= B(i,j) * B(i+1,j) * h(N(i,j), N(i+1,j))^2/\sigma_{h1} \\ &+ (1 - B(i,j)) * B(i+1,j) * h(N(i+1,j), 0)^2/\sigma_{h2} \\ &+ B(i,j) * (1 - B(i+1,j)) * h(N(i,j), 0)^2/\sigma_{h2} \end{aligned} \tag{5}$$

where $B(i,j)$ and $N(i,j)$ are the boundary value and normal direction at location $(i,j)$, and $h(\cdot,\cdot)$ is the distance measure of two angles modulated by 180 degrees. $\sigma_{h1}$ and $\sigma_{h2}$ are scale parameters, which control how strongly the object shapes conform to the corresponding shape constraints. Potential functions for the other three cliques are similarly defined.

This set of cliques can effectively capture the smooth property of surface normals and the unlikeliness of blindly ending boundaries. Nevertheless, it is insufficient for distinguishing true object surfaces from ones induced from erroneous filter response of random speckles. Furthermore, the broken surface speckle phenomenon may cause the filter to miss boundaries due to destructive echoes received at transducers. In either case we need to look beyond local neighboring properties to correctly determine the presence of true boundaries.

These problems can be overcome by analyzing object shapes in multiple scales to find support from increasing neighborhoods. Local boundary points suggested by the smallest scale filter are not to be regarded as true boundaries unless enough support is found from extended neighborhoods. Local non-boundary points are not to be regarded as non-boundaries if the global shape structures in some extended neighborhood strongly suggest they are on broken boundaries.

The prior is designed with a component to capture global shape structures in multiple scales. It uses a parallel algorithm for the multiscale analysis and infers two global shape attributes, *length* and *weight*, for all points. The length attribute is the length of the boundary fragment of a boundary point. Boundary points induced from random speckles should have rather short length attributes while those from true object boundaries should have longer length attributes. The weight attribute is designed to address the destructive echo problem on continuous surfaces. The weight of a non-boundary point suggests the chances that point

is actually on a boundary. This weight depends on the presence of two ending boundaries on either side of a point, the compatability of the tangents at the two end points, and the filter responses at different scales.

The algorithm computes length and weight by building a pyramid of cells of 4 by 4 overlapping receptive fields [Shneier81] [Meer90]. At each level in the pyramid, boundary fragments in child cells are joined together in the parent cell and their lengths are summed. A stochastic decision is then made to close the gap between a pair of boundary fragments based on the compatibility condition and the filter response at that scale. A weight inversely proportional to the gap size is then assigned to all points along the trajectory of the closing gaps. This algorithm then makes a downsweep of the pyramid, tracing down the exact boundary locations while propagating the length and weight attributes at the same time. See [Lin91] for details of this algorithm.

With the values of the length and weight attributes, $L(i,j)$ and $W(i,j)$, we define the following potential function:

$$V_{c5} = B(i,j) * \left( \frac{A_{L1}}{L(i,j) + C_L} + \frac{A_{W1}}{W(i,j) + C_W} \right) + (1 - B(i,j)) * \left( \frac{L(i,j)}{A_{L2}} + \frac{W(i,j)}{A_{W2}} \right) \quad (6)$$

where $A_{L1}, A_{W1}, A_{L2}, A_{W2}$ are scaling factors, and $C_W, C_L$ are positive constants introduced to prevent infinite potentials when the attribute values are zeros. Although this potential function looks like a single element clique, it is actually derived from extended neighborhoods. The $L$ and $W$ in this potential are global shape attributes computed from the extended neighborhood, and serve as latent data in our data augmentation scheme to be described in detail in the following section. By incorporating global attributes in a data augmentation scheme we have extended the use of Gibbs distribution to modeling global structures. Currently the parameters in the prior are determined empirically. Formal ways of determining these parameters are under investigation.

## Architecture of the Solution Algorithm

Let $X$ be the random field $X = \{(x_{ij} = (B(i,j), N(i,j)), \forall i,j\}$. Let $Y$ be the output of the filters from the input image, i.e. $Y = \{y_{ij} = (f_b(i,j), f_n(i,j)), \forall i,j\}$. We shall assume the conditional distribution $p(y_{ij}|X)$ of observing filter output $y_{ij}$ at voxel $(i,j)$ given $X$ depends only on $x_{ij}$ and that the values of $y_{ij}$ are independent of each other. It follows from our prior specification and likelihood function that the augmented posterior $p(X|Y,L,W)$ is a Gibbs distribution with the energy function [Lin91]:

$$U(\omega) = \sum_{C_i \in C} \sum_{i=1,6} V_{c_i}(\omega) \quad (7)$$

where $V_{c1}$ comes from the likelihood function, and $V_{c2}, V_{c3}, V_{c4}, V_{c5}$ come from the local shape constraints and $V_{c6}$ comes from the global shape attributes. But what we are really interested in is $p(X|Y)$. We use a data augmentation scheme, which is similar to an

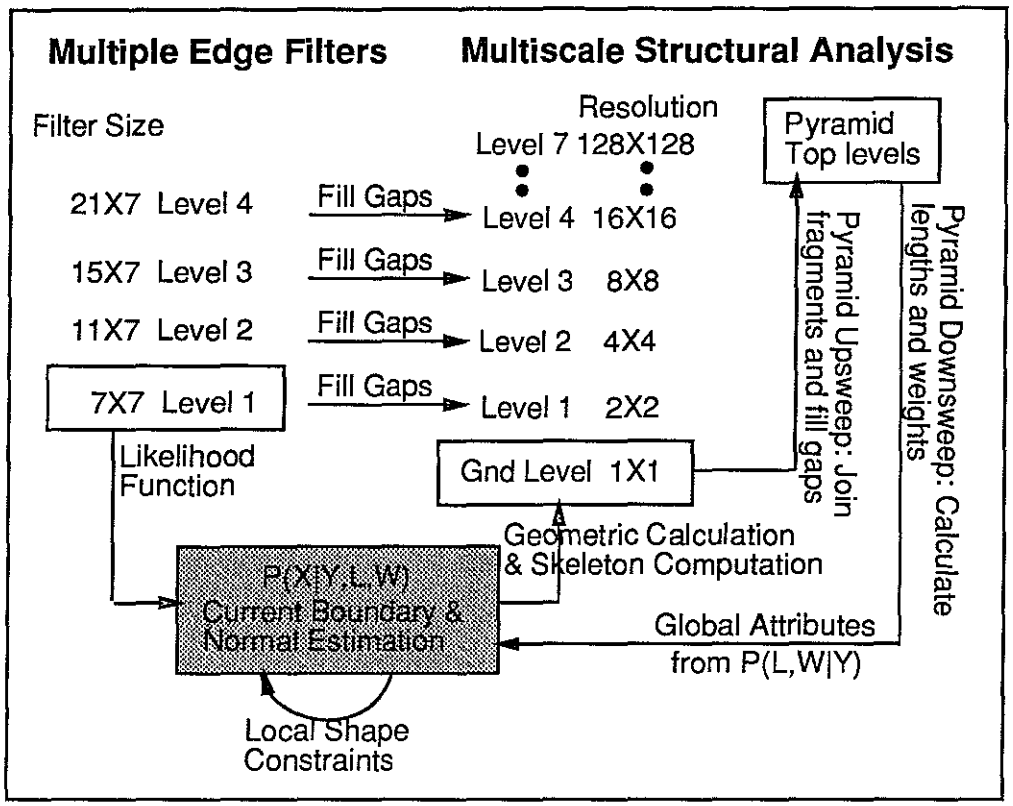


Figure 6: Operation of the Solution Algorithm

EM algorithm [Dempster77], to iteratively calculate $p(X|Y)$ from the augmented posterior $p(X|Y,L,W)$ and the predictive distribution $p(L,W|Y)$ by alternatively sampling these two distribution [Tanner87] [Gelfand90]. The predictive distribution $p(L,W|Y)$ is the integration of $p(L,W|X,Y)P(X|Y)$ over all possible $X$.

The operation of the complete algorithm is shown in figure 6. The current boundary values and normal directions to be estimated are shown in the shaded area. The initial boundary values are obtained by applying a threshold on $f_b(i,j)$, and the initial values of normal directions are $f_n(i,j)$. The algorithm then makes repeated updates on the whole field. Each iteration includes successive updates to each voxel. A sample from $p(L,W|Y)$ is drawn by performing the pyramid operation on the the current configuration. The Gibbs sampler is used to produce a new sample $x_{ij}$ for each voxel. It involves first computing the augmented marginal posterior $p(x_{ij}|X_{S/ij},Y,L,W)$, where $X_{S/ij}$ denotes the neighborhood of $x_{ij}$. Each $p(x_{ij}|X_{S/ij},Y,L,W)$ is computed from equations (2), (4), and (7) with only those cliques including voxel $(i,j)$. A sample is then randomly drawn accordingly and replaces the old $x_{ij}$.

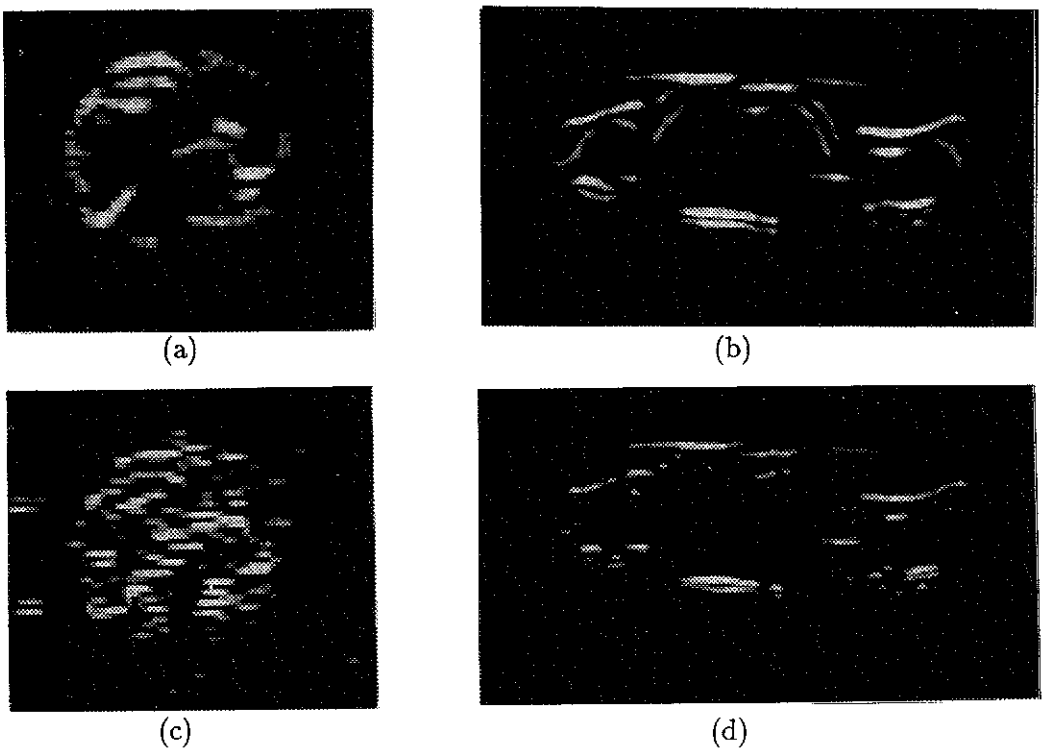


Figure 7: Our filter outputs: (a) Cone phantom, (b) Baby doll. Sobel filter outputs: (c) Cone phantom, (d) Baby doll.

In the current implementation we compute the pyramid only up to the sixth level. At this level each cell covers a 32 by 32 region. This size is considered large enough to compute reliable global attributes to reveal false boundaries. The values of $B(i,j), N(i,j)$ and squared values of them are accumulated during each iteration. They serve to compute the means and variances of the boundary and normal estimations.

## Results

We applied our filters on the baby doll image and the cone phantom image. Figure 7 show the outputs images from our filters, contrasted with those from a Sobel filter, which is a typical filter used in volume rendering. These pictures show our filter gives more accurate boundary measurements on the boundaries of the diffuse phantom and smooths out broken boundaries of the baby doll image.

## Discussion

We have described a Bayesian framework for estimating surfaces in noisy images. This framework produces a reference image of surface representations which can be further mapped to opacity for volume rendering. This reference image is a summary of all possible reconstructions of target shapes. This is a general idea.

We have applied this approach to two dimensional ultrasound images in particular and have demonstrated its usefulness in producing meaningful boundary representations. This boundary representation can not only be used in visualization but can also be the basis of image segmentation. We believe segmenting via this boundary representation will produce better results than segmenting the original ultrasound images based on the Sobel filter. We have implemented an active contour model using Gibbs distributions and are ready to test this idea using this model [Lin91].

This Bayesian framework is general enough to be applied to many applications, provided the likelihood function is appropriately adapted to the particular noise properties. In the case of ultrasound images we use larger elongated filters to handle the large speckle size. The Gibbs prior has also been extended to incorporate the important property of global shapes via multiscale analysis. By incorporating the global attributes in a data augmentation scheme we have extended the use of Gibbs distribution on modeling global structures.

Due to the large size of our filter, the localization of boundaries and corners are not optimal. In future reasearch we wish to use a multiscale likelihood function in which filter outputs of different scales are weighted by types of boundary and boundary curvatures. This should help better locate boundary locations. Another direction of future work is to extend our implementation to three dimensions to truly test this 3D ultrasound visualization approach.

## Acknowledgements

The authors wish to thank Dr. Olaf von Ramm of the Biomedical Engineering Department of Duke University, Mr. Antoine Collet-Billon of Philips-LEP in France, and Dr. Marc Levoy of the Computer Science Department of Stanford University for helpful discussions. Mr. Ryutarou Ohbuchi provided the baby doll image, and Philips-LEP provided the liver and phantom images. This research is supported by NSF grant number CDR-8622201 and NIH grant number P01 CA47982.

## References


[Canny86] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, November 1986.

[Dempster77] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society*, Ser. B(39):1–38, 1977.

[Fields73] Scott Fields and Floyd Dunn. Correlation of echographic visualizability of tissue with biological composition and physiological state. *The Journal of the Acoustical Society of America*, 54(3):809–812, 1973.

[Geman84] Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, 1984.

[Gelfand90] Alan E. Gelfand and Adrian F. M. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409, 1990.

[HykesS85] David Hykes, Wayne R. Hedrick, and Dale E. Starchman. *Ultrasound Physics and Instrumentation.* Churchhill Livingstone, New York, N.Y., 1985.

[Korn88] Axel F. Korn. Toward a symbolic representation of intensity changes in images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-10(5):610–625, September 1988.

[Levoy88] Marc Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8(3):29–37, May 1988.

[Lin91] Wei-Jyh Lin. *Surface Estimation in Ultrasound Images.* PhD thesis, to appear, U. of North Carolina, Chapel Hill, NC, 1991.

[Meer90] Peter Meer, C. Allen Sher, and Azriel Rosenfeld. The chain pyramid: Hierarchical contour processing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:363–376, 1990.

[Shneier81] Michael Shneier. Two hierarchical linear feature representations: Edge pyramids and edge quadtrees. *Computer Graphics and Image Processing*, 17:211–224, 1981.

[Shattuck84] David P. Shattuck, Marc D. Weinshenker, Stephen W. Smith, and Olaf T. von Ramm. Explososcan: A parallel processing technique for high speed ultrasound imaging with linear phased arrays. *Journal of Acoustic Society of America*, 75(4):1273–1282, April 1984.

[Tanner87] Martin A. Tanner and Wing Hung Wong. The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association*, 82(398):528–541, 1987.

[vonRamm88] O. T. von Ramm, S. W. Smith, K. H. Sheikh, and J. Kisslo. Real-time, three-dimensional echocardiography. *Circulation*, 78, 1988. Supplement II.