BOUNDARY ESTIMATION IN ULTRASOUND IMAGES

W J Lin[1], S M Pizer[1,2,3], and V E Johnson[4]

Departments of [1]Computer Science,
[2]Radiology, and [3]Radiation Oncology
University of North Carolina
Chapel Hill, NC 27599-3175

[4]Institute of Decision Science and Statistics
Duke University
Durham, NC 27706

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. We propose a Bayesian approach to the surface definition problem of ultrasound images, and study this approach in two dimensions. We formulate the problem as the estimation of posterior means and standard deviations of Gibbs distributions for boundary believability and normal direction. A set of filters of directional derivatives of Gaussians are used to measure the edge strength and orientation at multiple scales. The likelihood function is based on the measurement at the smallest scale. The prior distribution reflects shape properties at multiple scales. It uses a pyramid algorithm for contour analysis where the lengths of contours are computed and contour gaps are closed at multiple scales. The outcome of the pyramid algorithm is the length and weight global attributes for each pixel. These attribute values are incorporated into the Gibbs prior using a data augmentation scheme. The design and implementation of such an approach are the subject of this paper.

Keywords

Volume rendering; surface estimation; Bayesian; likelihood function; prior; posterior; Markov random field; Gibbs distribution; data augmentation.

1. INTRODUCTION

Three dimensional arrays of digital data are being 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 a three dimensional volume of ultrasound data in real time (von Ramm et al, 1988), (Shattuck et al, 1984). These volume data represent complex anatomy or functional process under study. Effective visualization of these volume data helps physicians in diagnostic interpretation or treatment 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 classification to obtain a partial opacity for every voxel. The second process performs surface shading at every voxel of volume data with a locally computed surface normal. Non-binary classification increases the likelihood that small or poorly defined features are perserved (Levoy 1988).

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 to 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 two dimensions. We formulate the problem as the estimation of posterior means and standard deviations of Gibbs distributions for boundary 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 edge-related measurements of ultrasound images to which Gibbs-compatible likelihood functions pertain. Then the design of Gibbs priors reflecting shape-related knowledge about object boundaries is presented, followed by a summary of the complete algorithm. We then present results and close with a discussion.

2. METHODS

2.1 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 on a grid 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 image 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 existence of a boundary at two neighboring voxels increases when their normal direction measurements are aligned and conform to the assumed boundary orientation but decreases when the normal directions are not aligned or contradict to the assumed boundary orientation.

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 the normal vector associated with the boundary. 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 show a low confidence. The mean of the boundary value can be used to modulate the opacity of the voxel and the standard deviation of the boundary value