Robust three-dimensional object definition in CT and MRI

Peyton H. Bland[a] and Charles R. Meyer
Department of Radiology, University of Michigan Hospitals, Ann Arbor, Michigan 48109-0553

(Received 21 April 1995; accepted for publication 3 October 1995)

This work describes the application of an object definition algorithm to the medical imaging environment for the task of automated detection of anatomical boundaries in three dimensions in the presence of low spatial frequency nonstationarities. We have chosen the Liou–Jain algorithm and have modified it for use with 3D medical image datasets and extended it by including a recruitment operator that corrects for the algorithm's inherent volume underestimation. The algorithm avoids problems in both traditional statistical segmentation and 2D techniques and elegantly bridges the gap between traditional gradient-based edge finding and regression-based segmentation techniques. Results are shown for MRI datasets from the human abdomen and brain and for a CT dataset of a liver tumor, as well as an MRI scan of a glioma in a rat brain. For comparison, the human abdomen dataset was processed by a multivariate, statistical classifier. The results demonstrate the statistical technique's susceptibility to low spatial frequency nonstationarities due to rf field inhomogeneity; the Liou–Jain algorithm is shown to be immune to this effect. Further, the results show spatial consistency as a result of inherent characteristics of the algorithm. Volumes identified by the algorithm are visualized and assessed qualitatively in three dimensions. Quantitative accuracy of the algorithm's volume estimates is assessed by the use of a phantom. This work demonstrates that this technique is effective in automatically detecting anatomical organ and lesion surfaces in 3D medical datasets that are corrupted by low spatial frequency nonstationarity and in obtaining volume estimates.

Key words: nonstationarity, field inhomogeneity, object definition, segmentation, three-dimensional (3D)

I. INTRODUCTION

Surface detection in medical image datasets arising from CT (computed tomography) and MRI (magnetic resonance imaging) is important in a number of areas. In radiation oncology, accurate, automated delineation of organ and lesion boundaries is vital in developing treatment plans, as well as retrospective assessment of radiation delivery. Chemotherapeutic approaches to cancer therapy depend on accurate estimates of lesion and organ volumes to assess the efficacy of a particular therapeutic regimen. The availability of a three-dimensional anatomical representation may aid in planning and executing certain surgical procedures. All of these applications have the common requirement of identifying organ and lesion boundaries in three dimensions for the purposes of visualization and quantification. This paper describes how we are using a robust object definition algorithm to determine anatomical boundaries noninteractively in medical image datasets. The term "robust algorithm" is borrowed from computer science; it describes an algorithm that performs well in the presence of unusual conditions. In this paper, we use it to characterize the algorithm's inherent ability to define objects in the presence of varying degrees of field nonuniformity as discussed in the following paragraphs.

To date, many segmentation techniques applied to medical imaging have relied on statistical classification.[1–6] Although relatively straightforward in their application and efficient in their computation, these techniques have two serious flaws: They are adversely affected by low spatial frequency nonstationarities in the input image datasets, and spatial relationships are not taken into account unless augmented by other techniques. Bezdek *et al.*[6] have reviewed several segmentation techniques under the general heading of pattern recognition.

In CT images, spatial nonstationarities can arise from cupping artifact,[7] particularly in the head. CT contrast studies may also suffer from spatial intensity variations caused by movement of the contrast bolus through the tissue being imaged relative to the position of the scan plane at a particular time.[8] Spatial inhomogeneity of the rf field coil (especially body and surface coils) in MRI scanners results in low spatial frequency nonstationarities in MRI images.[9–11] (Throughout the remainder of this paper, when terms such as "nonstationarity," "intensity variations," and "nonuniformity" are applied to image datasets, they refer to intensity variations having low spatial frequency content.) Even newer "bird cage" coils for 1.5-T scanners suffer from field inhomogeneities, sometimes to a significant degree.[12] The recent demonstration of high field strength scanners (e.g., in the 4-T range) in producing useful clinical images[13] underscores the need to correct nonstationarities caused by field inhomogeneities and limited rf penetration due to the higher frequencies employed.[14]

These intensity variations cause severe errors in classifying MRI datasets using traditional statistical methods. Figure 1 illustrates the problem in analyzing an MRI exam of the liver using a histogram-based classifier known as HICAP (Histogram Cluster Analysis Procedure).[2] The original gray-scale image [Fig. 1(a)] demonstrates intensity variation be-

 0094-2405/96/23(1)/99/9/$6.00

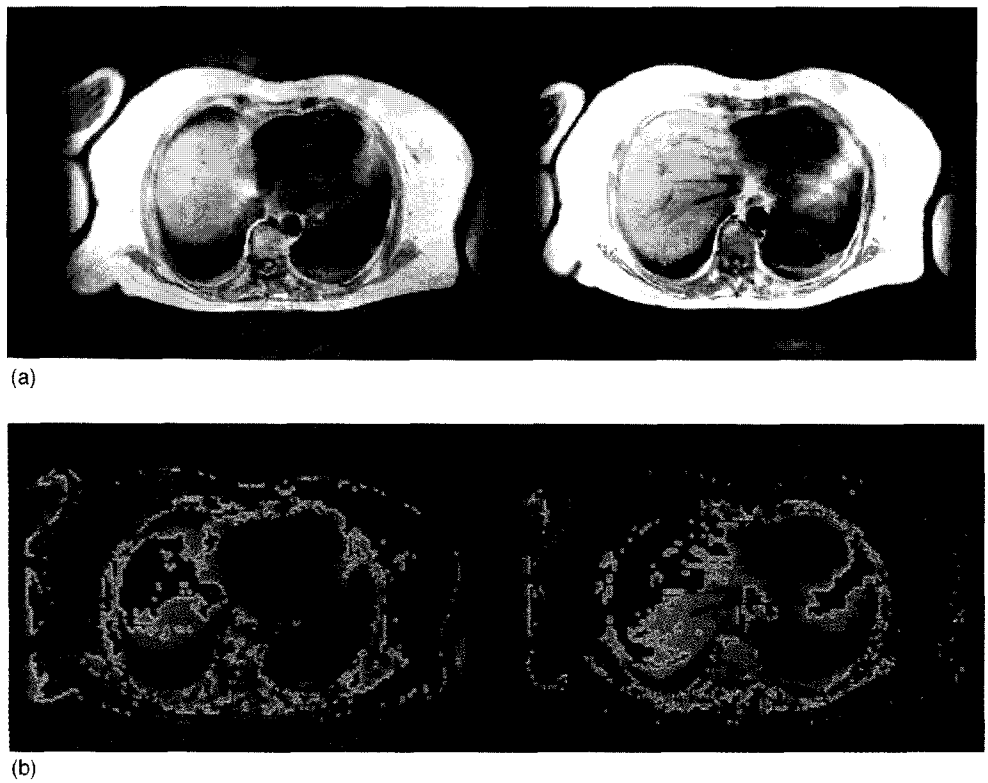

FIG. 1. (a) MRI image of the liver. Note the intensity differences between the left and right sides of the abdomen for subcutaneous fat as well as differences within the liver itself due to field inhomogeneity. (b) Results of statistical classification using HICAP. The clusters illustrated are intended to identify the liver (red and green) and subcutaneous fat (red and blue), but erroneously include other structures as well (e.g., a vertebra, the spinal cord, intercostal and other muscles).

tween the left and right sides of the abdomen due to field inhomogeneity of the scanner. Because of this, the clusters found by HICAP [Fig. 1(b)] do not have a one-to-one correspondence with anatomical structures. For example, in order to define the liver, three clusters must be used. However, in doing so, structures other than the liver are included. Although simple statistical classification is useful in some situations, its application to MRI datasets is complicated since significant manual editing of the clusters or the incorporation of other techniques is necessary to obtain useful results.

Traditional statistical classification algorithms also have the drawback of neglecting spatial information. For example, the formation of a histogram inherently neglects organ shape and anatomical relations between organ cross sections in adjacent slices. Further, some algorithms operate in two dimensions even when the goal is to classify three-dimensional structures. That is, classification is performed on each slice of the dataset independently; clusters from contiguous slices must then be associated to reconstruct a three-dimensional representation. Since spatial relationships are not preserved, clusters in adjacent slices may be spatially inconsistent. Even if the histogram is formed using all of the slices in the three-dimensional dataset, spatial information is still not considered and connectivity must be provided by postprocessing, for example, manual editing[3] or the application of various operators.[5] Morphological operators may also be employed, but even gray-scale morphological operators may not utilize the underlying intensity data appropriately with regard to segmentation, especially in the presence of nonstationarities.

II. MATERIALS AND METHODS

Liou and Jain[15–18] have developed a novel segmentation algorithm that we use as a basic tool to identify 3D volumes in medical image datasets that are corrupted by nonstationarities. It overcomes the fundamental deficiencies in statistical classifiers discussed above. In this section, a brief description of the algorithm (based on Refs. 15–18) is presented. We also discuss dataset preprocessing and modifications made to the algorithm to adapt it for use in the medical imaging environment. Finally, we describe an extension to the algorithm that compensates for volume underestimation, a well-recognized limitation of this segmentation tool.

The input data consists of a series of contiguous two-dimensional images (typically from a tomographic imaging modality such as CT or MRI) that define a three-dimensional rectangular volume dataset. Depending on the scanning parameters, the voxel aspect ratio may be as great as 10:1:1. The resulting anisotropic partial volume effect causes difficulties in computing the gradient magnitude and in generating the volume hypotheses in the Liou–Jain algorithm. Interpolation of additional slices between the original slices offers a partial solution at the expense of increased storage and

computational requirements. (However, interpolated slices retain the same partial volume anisotropy as the original dataset.) Instead, as a preprocessing operation, we perform two-dimensional low-pass filtering on each slice of the original data. The elements in the filter's kernel are calculated as a function of slice thickness and in-plane resolution to filter the data in the *in-plane* dimensions to an extent that approximates the partial volume blurring that occurs during data collection due to slice thickness. Thus the remainder of the algorithm operates on voxels with isotropic partial volume contributions without interpolating along the long axis of the data or subsampling in the in-plane dimensions.

In some datasets, three-dimensional median filtering is performed as discussed later. The kernel dimensions are chosen so that the kernel size expressed in physical coordinates is approximately cubic. Median filtering has the advantage of reducing speckle noise in the data while preserving edges, another important consideration in volume hypothesis generation.

A weighted three-dimensional gradient operator is used to compute the gradient magnitude. The three difference components are scaled by dividing each by its respective voxel dimension. In order to bound and compress the gradient magnitude, an embedding angle transform[18] (inverse tangent of the gradient magnitude) is applied. The transformed gradient magnitude is passed into the segmentation algorithm where it is used as a 3D edge map along with the 3D filtered intensity dataset.

The algorithm begins by partitioning the intensity dataset into a series of volume hypotheses based on discontinuities that typically separate features having different uniformity or smoothness characteristics. Discontinuities are detected by thresholding the gradient magnitude data. However, edge strengths can vary; in medical images this may be caused, for example, by (1) partial volume effects and (2) orientation and profile characteristics of organ–lesion and organ–organ boundaries. Thus, a *set* of thresholds is applied iteratively to locate bounding surfaces in three dimensions. The iteration proceeds from high to low threshold values, generating an independent set of nonoverlapping volume hypotheses at each step. In this manner, the volume hypothesis generated for a particular feature is initially relatively large and becomes smaller as the iteration continues. Each volume hypothesis is tested for smoothness. Volume hypotheses that fail the test are discarded; however, the voxels of these hypotheses are available for consideration at lower thresholds where they may be included in smaller volume hypotheses. Smoothness within a hypothesis is assessed by a regression technique. Intensity values of the voxels within the hypothesis being tested are fit to a polynomial whose independent variables are spatial coordinates in the three-dimensional space of the dataset:

$$\hat{I}(v,\vec{a};x,y,z)=\sum_{i+j+k\leq v} a_{ijk}x^i y^j z^k. \tag{1}$$

The maximum variation order (MVO), v, of the polynomial is chosen to allow a certain variation in the intensities within the volume to be accommodated. In the work presented here, $v=4$ is used. Each volume hypothesis at each threshold is fit separately and independently and is thus characterized by its own set of coefficients $\vec{a}$.

Goodness-of-fit of the regression for each hypothesis is assessed by evaluating the following inequality:

$$\mathbf{s}^2\leq\sigma^2, \tag{2}$$

where $\mathbf{s}^2$ is the residual and is calculated as the sum of squares of the error of the regression fit divided by the degrees of freedom, and σ^2 is estimated from the noise variance of the intensity data using simple statistics over a region of interest. If Eq. (2) is satisfied, the volume hypothesis is accepted as correct. Consider a case, however, where it is not satisfied, and the hypothesis is rejected. In the partitioning scheme described above, the size of the volume hypothesis for a given feature decreases as the threshold iteration proceeds, and $\mathbf{s}^2$ will most likely decrease. Thus a hypothesis that fails to meet the smoothness criterion at a high threshold may satisfy it at a lower one.

As a result of thresholding the gradient magnitude to form the volume hypotheses, volumes are sometimes significantly underestimated. This is because voxels in the gradient magnitude data having values *above* the threshold define the bounding surfaces *between* volume hypotheses and are not included in volume hypotheses themselves. As a result they remain uncommitted (i.e., not included in a volume hypothesis). This causes a reduction of volume at the surfaces of the volumes that is dependent on the profile of the edges. A 3D voxel recruitment operator that functions as an extension of the Liou–Jain algorithm is applied to dilate the volumes selectively by attempting to assign uncommitted voxels to previously identified volumes. This operator may be considered a hybrid morphological operator in that it operates in part as a conventional binary morphological operator but more importantly examines the intensity data (as well as the regression models defined by the coefficients for the identified volumes) in making its decisions regarding the disposition of uncommitted voxels by using a hypothesis testing paradigm.

Figure 2 depicts this three-dimensional operator in two dimensions for simplicity. In this illustration, the kernel of the operator overlays voxels from two volumes identified by the basic algorithm, as well as uncommitted voxels not identified by the algorithm. Each uncommitted voxel is examined to determine whether it should be assigned to a previously identified volume. This determination is based on *both* local morphology and the voxel's intensity value and is distinct from simple statistical or histogram-based techniques. First, the voxel in question must be eight-connected (or 26-connected for the 3D implementation that is actually used in this work) to at least one previously identified volume. This is consistent with connectivity characteristics of the Liou–Jain algorithm and means that voxels are only added to surfaces of existing volumes.

If this connectivity condition is satisfied, the voxel is tested against each of these connected neighboring volumes (A and B in Fig. 2) to determine whether it should be assigned to one of them. This is accomplished by evaluating

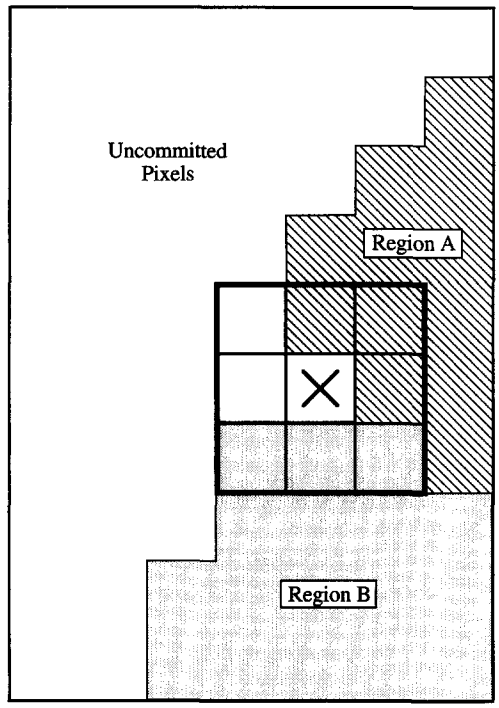

FIG. 2. Two-dimensional example of the recruitment operator used to correct for volume underestimation. The uncommitted pixel marked by the "×" at the center of the 3×3 kernel is eight-connected to both regions A and B. This operator employs hypothesis testing as does the basic Liou–Jain algorithm itself.

the model polynomial (whose coefficients were previously computed during hypothesis testing) of each neighboring volume at the coordinates of the voxel in question according to Eq. (1). The voxel is assigned to the volume whose model offers the best fit as assessed in a manner consistent with the hypothesis testing step of the basic Liou–Jain algorithm. This decision is based on two criteria, both of which must be satisfied in order that the voxel be assigned to an existing volume; otherwise, the voxel remains uncommitted. The goodness-of-fit parameter s^2 is updated by adding the square of the residual of the potentially added voxel to the sum of squares of the residual of the existing volume and dividing by the degrees of freedom; the first criterion is that this updated s^2 satisfy Eq. (2). The second is a more selective criterion based on the residual of the added voxel. It requires that the residual of the added voxel divided by the updated s not exceed a user-specified individual voxel fit parameter which is expressed as the number of sigmas of the residual from zero. That is, if we assume that the residual has a normal distribution with a standard deviation of s and zero mean, the model error of the added point must lie within a user-specified number of sigmas from the mean.

The computations were carried out on a DEC 3000 Model 500X AXP (Digital Equipment Corp., Maynard, MA). Typical execution times on this processor range from approximately 20–60 s, depending on the size of the dataset and the schedule used to iterate through thresholds. The results were visualized in both two and three dimensions using AVS (Application Visualization System from Advanced Visual Systems, Inc., Waltham, MA).

III. RESULTS

The volumes identified by the Liou–Jain algorithm for a MRI exam are presented in Fig. 3 for a representative slice through the volume dataset. This is the same data that was submitted to HICAP analysis (Fig. 1). Figure 3(a) shows the intensity data, while Fig. 3(b) demonstrates the regions identified by the algorithm. The voxel recruitment operator was applied to the volume dataset [Fig. 3(c)]. In Figs. 3(b) and 3(c), note that although some regions may appear to consist of a number of disconnected pieces, they are in fact connected at other locations in the 3D dataset since connectivity in three-space is guaranteed by the algorithm. In fact, the 3D alpha blend/gradient shaded display in Fig. 3(d) demonstrates this.

The algorithm was also applied to a dataset obtained by scanning a rat brain *in vivo* containing a 9L glioma using a SISCO (Spectroscopy Imaging Systems Corp., Sunnyvale, CA) MRI system interfaced to an Oxford (Oxford, England) 7-T magnet.[19] Figure 4 shows the results. In this case, preprocessing includes median filtering to reduce the texture noise in the intensity data which appears in both normal and diseased neural tissue. Without this step, the texture prevents adequate volume hypothesis formation from the gradient magnitude data.

A dataset from a MRI scan of a human head was analyzed, with the results shown in Fig. 5. The slices were collected in the coronal orientation. This analysis demonstrates the performance of the algorithm in a case where fine detail (i.e., the interface between the gray and white matter) is present. Eighteen passes of the recruitment operator were applied over which the volume representing white matter increased from 98.7 to 559 cm^3. The operator is applied repeatedly until the increase in volume between two successive passes drops below a specified fraction. In this case the stopping criterion of 0.1% increase in volume of the white matter was reached between the 17th and 18th passes.

A CT liver exam was analyzed to study the performance of the algorithm in defining a liver lesion, as shown in Fig. 6(a), a case demonstrating a posttransplant lymphoproliferative tumor. The results of the basic segmentation algorithm and voxel recruitment are shown in outline form in Fig. 6(b).

In order to quantify the accuracy of the algorithm, a phantom study was performed. A vinyl glove filled with copper sulfate solution (approximately 15 mM) was imaged by a Picker Vista 0.5-T MRI scanner (Picker International, Cleveland, OH). Sixty slices orthogonal to the long axis of the glove were collected over 30 cm with no gap between slices using a 30-cm field of view; 256×256 images were reconstructed. The volume of the fluid within the glove is calculated as 748.5 cm^3 based on its density and its independently measured weight. As expected, volume underestimation was evident in the result in the form of a gap between the volume identifying the fluid within the glove and the volume identifying the surrounding air, yielding an estimate of 575.8 cm^3

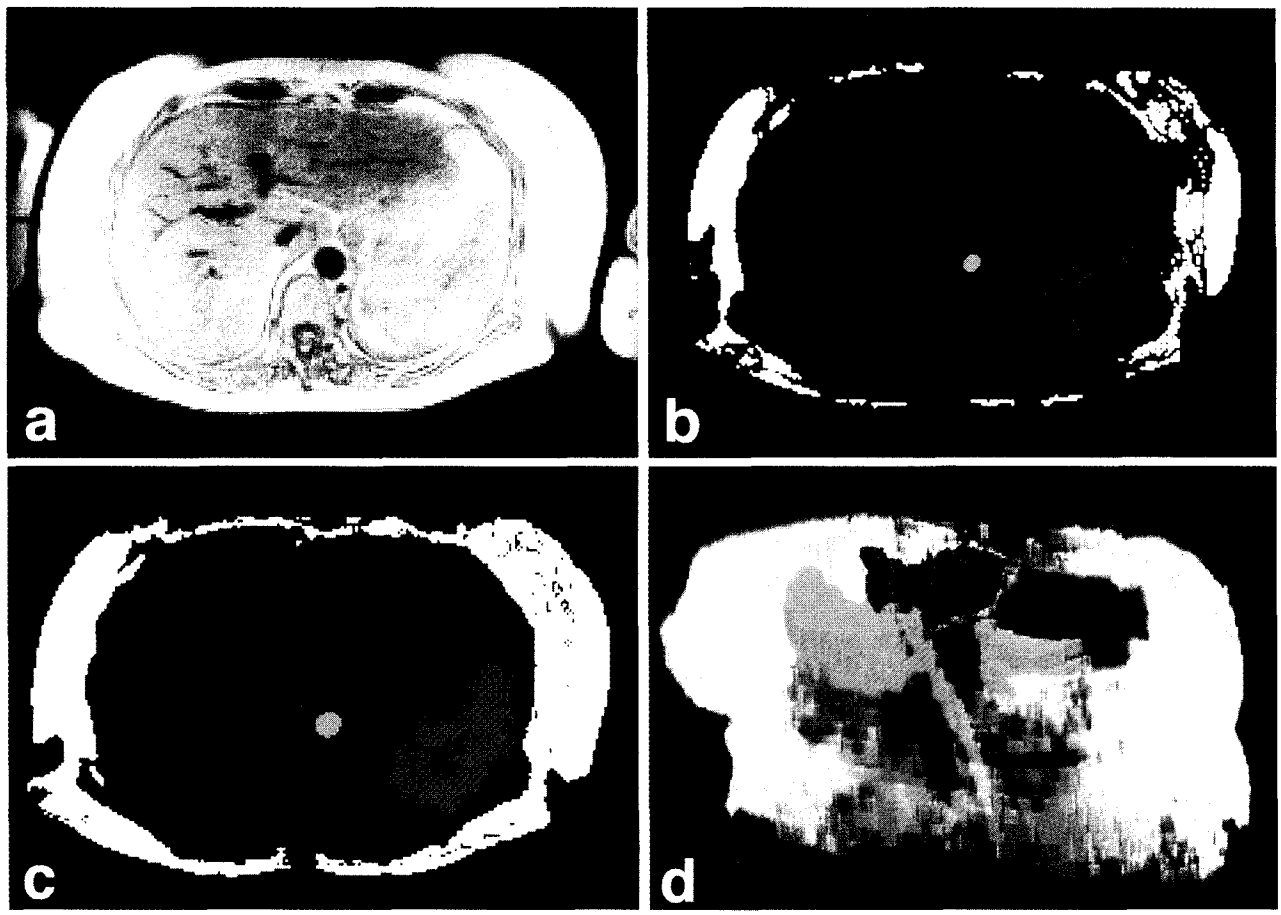

FIG. 3. The same MRI case as presented in Fig. 1 was submitted to the Liou–Jain algorithm. A set of 14 contiguous 2D images forms the volume dataset. The results are displayed using a visualization tool known as AVS. Each volume is indicated by a different color: white—subcutaneous fat, magenta—liver, red—organs of the digestive system, dark orange—spleen, yellow—lungs and aorta (see discussion in the text), green—heart. (a) A representative axial slice of the intensity data. Field inhomogeneity is evident in the intensity variation between the liver's left lobe, its central portion, and its right surface. (b) Segmentation results displayed in two dimensions at the same axial position. Although volumes may appear to be fragmented, they are in fact connected in other axial slices not shown in this figure. (c) Results after application of the recruitment operator. (d) The same data presented using an alpha-blend/gradient shading technique. The view is from the rear of the patient looking down into the abdomen.

for the glove's interior volume. Following recruitment, the estimate of the glove's interior volume is 748.3 cm^3.

IV. DISCUSSION AND CONCLUSION

The difference in performance of the statistical segmentation and the robust object definition procedure in the presence of spatial nonstationarities is well illustrated by Figs. 1 and 3. The nonstationarity resulting from rf field inhomogeneity causes a difference in intensity between different regions of the liver and between the subcutaneous fat on the left and right sides of the images. While this causes serious errors to be committed in the HICAP analysis in detecting the liver, the Liou–Jain algorithm identifies the liver as a single volume. Further, the Liou–Jain algorithm discriminates between the vessels in the liver and liver tissue itself since the vessels are (correctly) not included in the volume hypothesis for the liver. Note that some detail is lost due to volume underestimation, apparent in some of the vessels of the liver; however, some small vessels are well delineated. The alpha-blend/gradient shading display technique used in Fig. 3(d) allows the seven major anatomic features identified by the algorithm to be visualized. It is easy to see an artifactual bridge between the left lung and the descending aorta caused by limited resolution and partial volume averaging in the original scan which causes these two organs to be identified as one. The right lung is identified separately and is presented in yellow along with the left.

The method of thresholding the gradient magnitude field to generate volume hypotheses is simple to implement, efficient to compute, and guarantees completely enclosed volume hypotheses. The surfaces that delineate the volume hypotheses are typically more than one voxel thick and are not included in the volumes themselves. As a result, artifactual links between volume hypotheses are not possible as is the case with some algorithms.[3,20] Another result, however, is that volumes are underestimated. The voxel recruitment operator is used to correct this underestimation. Its operation is consistent with the Liou–Jain algorithm regarding its spatial connectivity characteristics and its use of hypothesis testing in determining whether an uncommitted voxel should be

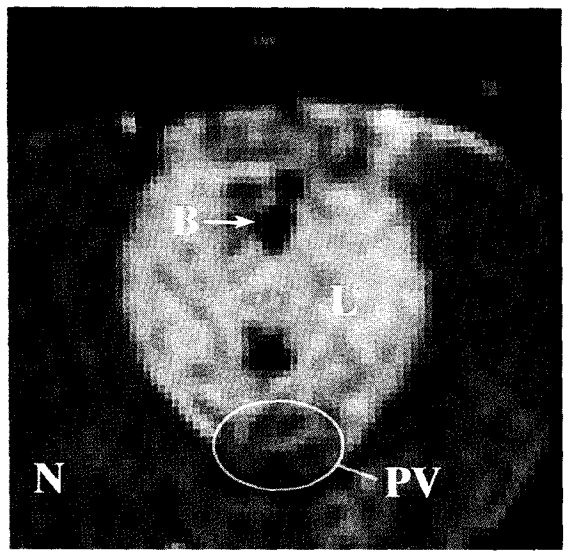

(a)

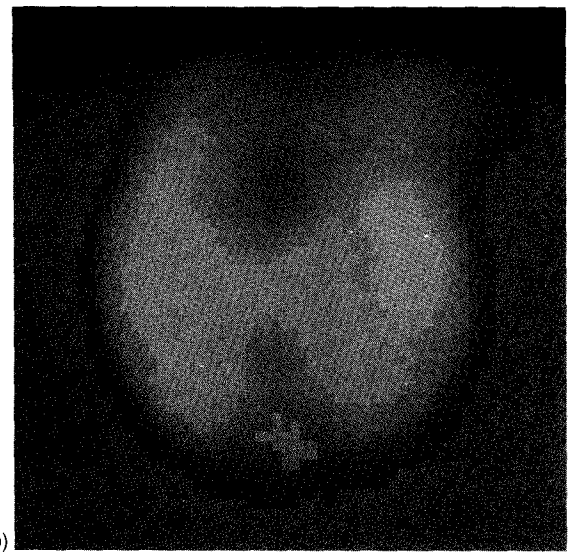

(b)

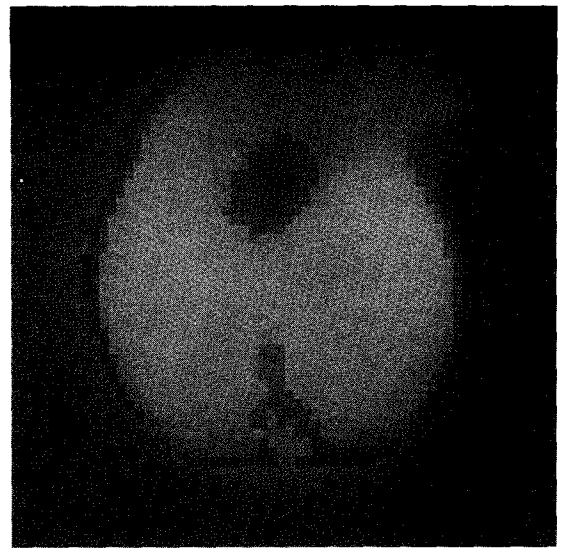

(c)

FIG. 4. Normal tissue and a 9L glioma are identified by the algorithm in a MRI dataset of a rat brain. Note that the original slices (256×256 pixels in-plane) are cropped to 77×77 pixels to include only the tissue of interest. (a) Typical slice of raw data. Note the texture in both the normal tissue (N) and in the lesion (L). Areas of blood (B) are present in the glioma. Partial volume effect (PV) is also evident. (b) The same slice after preprocessing (in-plane low-pass filtering and median filtering) with the results of the segmentation (before voxel recruitment) for this slice superimposed. Two volumes identifying normal brain (blue) and the lesion (red) are indicated. (c) The same slice after recruitment.

added to an existing volume. Since it takes the conservative approach of adding voxels in a single layer to existing volumes, multiple passes of the operator are usually applied. However, even though the algorithm for adding voxels is conservative, dramatic growth in a region can occur even with a single pass of the operator.

The recruitment operator is, in a sense, self-limiting. First, if two existing volumes abut each other, there is, of course, no opportunity for growth at the point of contact. Second, a voxel must satisfy goodness-of-fit criteria in order to be added. Finally, the process is stopped when the increase in volume between successive passes drops below a specified fraction. If a conservative individual voxel fit criterion is used, the fit, as quantified by s [Eq. (2)], often improves slightly with the addition of voxels. The major short-coming of this operator is that if an isolated anatomical feature is too small to have been identified by the basic Liou–Jain algorithm or if the nature of the gradient magnitude field within or near it otherwise prevented appropriate volume hypoth-

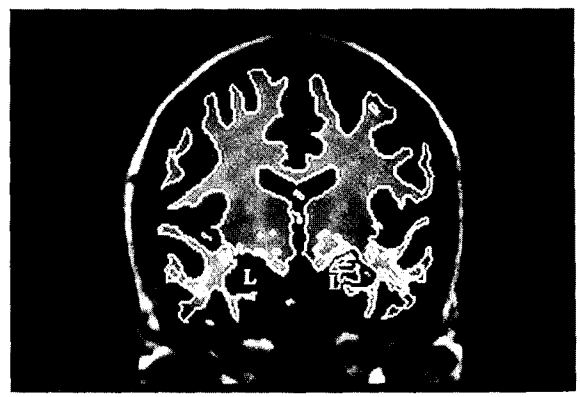

FIG. 5. A MRI scan of a human head. A single slice from the 3D data is shown with the results of the object definition process for the white matter only; the recruitment operator has been applied. (An interior pixel remove operation has been applied to the volume that identifies the white matter so that the white and gray matter can be better visualized in the intensity image.) Note the fidelity with which the interface between white and gray matter is defined. Note also that the algorithm erred in detecting the medial temporal lobes (L) differently on each side.

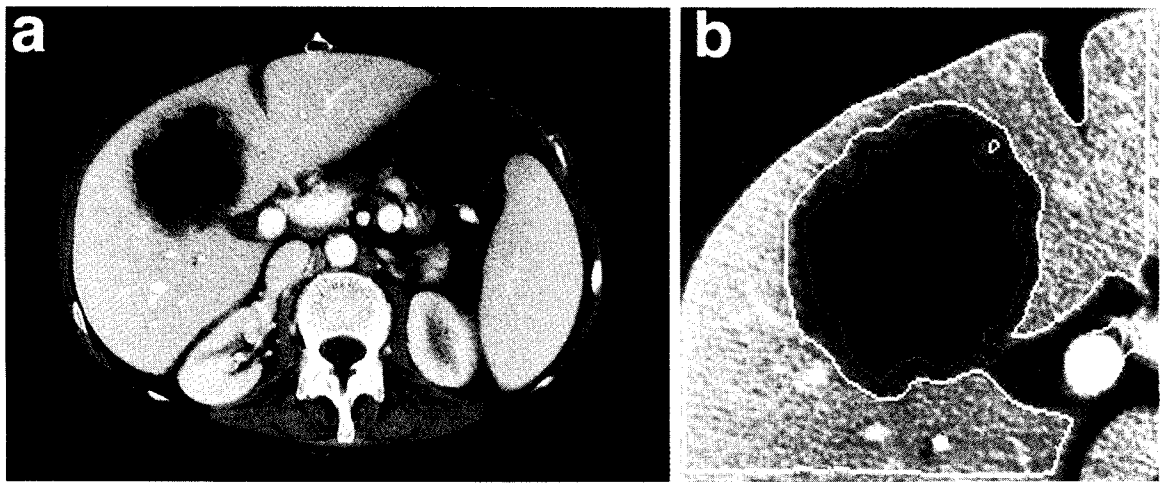

FIG. 6. Definition of a lesion in CT data. (a) A representative axial slice of the intensity data. (b) Results of segmentation and recruitment. The tumor is represented by the red outline and normal liver tissue by the green outline. The region between the identified tumor margin and normal tissue represents uncommitted voxels that remain after applying the recruitment operator.

eses from being formed, then, of course, no recruitment is possible at that location.

Although the MRI head scan (Fig. 5) presents a challenge to the Liou–Jain algorithm, it demonstrates how the recruitment algorithm corrects volume underestimation by propagating the volume surfaces to appropriate anatomical boundaries while preserving small detail. Defining the white matter completely required quite a large amount of growth: The initial output of the Liou–Jain algorithm identified only the central portion of the white matter adequately but almost none of the subcortical white matter immediately beneath the gyri. Even though recruitment was performed by applying many passes with a lenient fit criterion (18 passes using an individual fit parameter of 7.0 sigmas), the operator defined the detailed boundary between gray and white matter without inappropriate growth into neighboring anatomical structures. The initial output of the Liou–Jain algorithm also identified regions that only sparsely sampled the gray matter. However, this was inadequate to allow the recruitment operator to form a volume correctly representing gray matter. Sparse sampling is a result of the volume underestimation characteristic described above. In many locations the thickness of the gray matter layer is comparable to the thickness of the layer lost to volume underestimation, resulting in small or nonexistent hypotheses. The inability of the present method of hypothesis generation to deal with this situation motivates an important piece of future work: designing an operator that forms volume hypotheses with single voxel wide boundaries.

The order of the polynomial used to model the intensity has a bearing on the acceptance of volume hypotheses and is chosen based on the nature of the underlying nonstationarity. (In a practical sense, the order also affects computational speed.) For example, a relatively low-order polynomial may be sufficient to fit a volume containing very smooth, low spatial frequency variations. The order of the polynomial [expressed as "maximum variation order" or MVO, v in Eq. (1)] along with the number of independent variables (three spatial coordinates in our case) determines the form of the polynomial.

To assess the effect of MVO on the residual of the fit, we fit the voxels from the liver of the MRI dataset (Fig. 3) to five different polynomials specified by MVOs of 1, 2, 3, 4, and 5. Figure 7(a) plots the residual variance of the liver volume's fit versus the number of polynomial terms for each

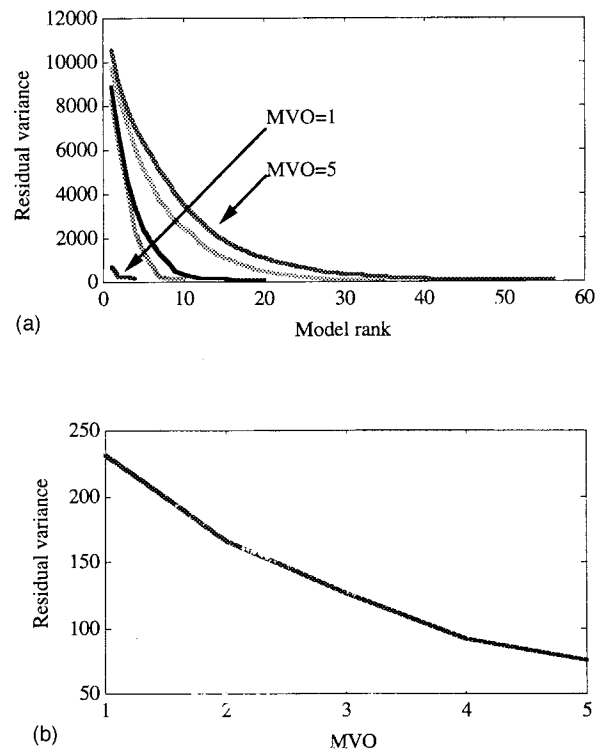

FIG. 7. The effect of polynomial model order on the residual of the model fit using the voxel values of the liver identified in Fig. 3. (a) Variance of the residual vs number of terms in the polynomial for each of five MVOs (1–5). (b) Variance of the residual vs MVO.

of the five MVOs. Figure 7(b) shows a plot of the resulting variance deviation of the residual with all terms considered versus MVO. At a MVO of 4, the standard deviation of the residual is approximately equal to the signal-to-noise ratio measured in this dataset which sets an approximate lower limit on the residual. Further, it is clear that the improvement in the residual becomes less as MVO increases. We conclude that a MVO of 4 is appropriate based on (1) our experience with the performance of the algorithm on MRI datasets (a conclusion also reached by the original authors of the algorithm[16]) and (2) the computational complexity resulting from high MVOs. (Specifying an MVO of 5 results in a polynomial of 56 terms versus 35 terms for a MVO of 4.)

The amplitude of the texture pattern in the rat brain data obtained from the SISCO system is great enough to require the additional step of median filtering during the preprocessing phase. Without texture reduction, volumes identified in the brain and glioma are very sparsely filled because of the difficulty in forming appropriate volume hypotheses by gradient magnitude thresholding. In the case shown, the application of an $11\times11\times3$ kernel (a cube approximately 1.3 mm, on each side in physical coordinates) results in a striking reduction in texture while preserving edges adequately [Fig. 4(b)].

It is noteworthy in Fig. 4(b) that at some locations the edge of the volume identifying *normal* rat brain appears closer to the "true," visually assessed interface with the tumor than does the edge of the *lesion's* volume. Inspection of edge profiles reveals that the shape of the edge of the intensity data results in a gradient magnitude profile that is asymmetrical. When thresholded, the interior portion of the edge (i.e., on the lesion's side) is thicker than the outer portion, causing underestimation of the lesion to be slightly greater than that of normal tissue.

The algorithm is well suited for defining amorphously shaped objects since it makes no assumptions about shape. Thus, one of its prime uses is the identification of lesions. Figure 6 demonstrates the performance of the algorithm in defining a lesion (a posttransplant lymphoproliferative tumor) that appears in a CT liver exam. The results of segmentation and recruitment are shown in outline form in Fig. 6(b). Although field nonuniformity is not a major problem in this case, we use the algorithm because of its 3D characteristics and the ability to use the recruitment operator to refine the edge estimates. The algorithm identified both normal liver tissue and the lesion as shown in outline form in Fig. 6(b). There is a band of uncommitted voxels between them which is consistent with detailed visual analysis of the edge profile. The fact that the edges of the liver and the lesion do not coincide can be used to advantage in object definition for radiation therapy treatment planning. For example, the red outline would correspond to "gross tumor volume" and the green to "clinical tumor volume" which includes regions of suspected microscopic tumor invasion.

Quantitative assessment of an object definition algorithm's performance is often difficult because of the problem of finding an appropriate "gold standard." Even manual tracings made by a trained operator are often subjective (especially in areas of partial volume averaging and low-contrast edges) and are not always reproducible.[21] Quantitative comparisons with automated methods are similarly prone to error. Hence, we rely in part on qualitative assessment of the results.

The vinyl glove phantom is a 3D object having a volume known with sufficient accuracy and precision via independent information (its weight and density). The excellent agreement with the measured volume of the vinyl glove phantom demonstrates that the Liou–Jain segmentation algorithm along with voxel recruitment identifies boundaries correctly in the presence of spatial inhomogeneity.

Counting pre- and postprocessing steps, the elapsed time required by the Liou–Jain algorithm is approximately 20 times greater than that required by HICAP for the MRI dataset presented here. Conclusions drawn from this number must consider the following. First, the time required for the Liou–Jain algorithm is highly dependent on the number of threshold levels used and the number of hypotheses formed and tested at each threshold. Second, we consider the benefits of the Liou–Jain algorithm as presented in this paper to outweigh the computational requirements. Third, note that the nature of the hypothesis generation step allows parallelism to be exploited to a high degree. Separate processors can evaluate hypotheses independently at each threshold, with the results combined into a single dataset. (This high degree of distribution results in some redundant calculations. Implementing an overlap detection scheme via interprocess communication could eliminate this redundancy.) Note that recent work has shown that the Liou–Jain algorithm executes approximately 60 times faster on a RISC processor (DEC 3000 Model 500X AXP) than on our previous generation computer, thus making it more suitable for interactive use.

We have shown that the Liou–Jain algorithm is immune to problems that are inherent in statistical classifiers: It performs well in the presence of low spatial frequency nonstationarities and guarantees connectivity in three-space. The use of this algorithm in identifying structures in the human abdomen and head and in detecting tumors in the human liver and rat brain were demonstrated. Qualitative assessment shows good results in defining anatomical surfaces; quantitative assessment of volume estimates by the use of a phantom shows a high degree of accuracy following the application of a correction for volume underestimation. With attention paid to pre- and postprocessing operations, this algorithm shows promise in defining objects automatically in the medical imaging environment.

ACKNOWLEDGMENTS

The authors are pleased to acknowledge the gracious assistance of Shih-Ping Liou, and Ramesh Jain, in the use of their algorithm for this project, as well as their helpful discussions. In addition, we would like to thank Digital Equipment Corp. for providing access to a DEC 3000 Model 500X AXP used for algorithm development and execution. This work was supported in part by DHHS PHS NIH grants 1R01 CA52709-03A1 and CA59412-01.

[a]Address correspondence to Peyton H. Bland, Ph.D., Department of Radiology, University of Michigan Hospitals, Kresge III, R3300/0553, 1500 East Medical Center Drive, Ann Arbor, MI 48109-0553; Electronic-mail: bland@umich.edu

[1]M. W. Vannier, R. L. Butterfield, D. Jordan, W. A. Murphy, R. G. Levitt, and G. Mokhtar, "Multispectral analysis of magnetic resonance images," Radiology **154**, 221–224 (1985).

[2]S. W. Wharton, "A generalized histogram clustering scheme for multidimensional image data," Pattern Recog. **16**, 193–199 (1983).

[3]L. F. Farjo, D. M. Williams, P. H. Bland, I. R. Francis, and C. R. Meyer, "Liver volume determination from CT images using histogram cluster analysis," J. Comp. Assist. Tomogr. **16**, 674–683 (1992).

[4]M. I. Kohn, N. K. Tanna, G. T. Herman, S. M. Resnick, P. D. Mozley, R. E. Gur, A. Alavi, R. A. Zimmerman, and R. C. Gur, "Analysis of brain and cerebrospinal fluid volumes with MRI imaging," Radiology **178**, 115–122 (1991).

[5]H. E. Cline, W. E. Lorensen, R. Kikinis, and F. Jolesz, "Three-dimensional segmentation of MR images of the head using probability and connectivity," J. Comp. Assist. Tomogr. **14**, 1037–1045 (1990).

[6]J. C. Bezdek, L. O. Hall, and L. P. Clarke, "Review of MR image segmentation techniques using pattern recognition," Med. Phys. **20**, 1033–1048 (1993).

[7]H. H. Barrett and W. Swindell, *Radiological Imaging; The Theory of Image Formation, Detection, and Processing, Vol. 2* (Academic, New York, 1981).

[8]R. C. Nelson, J. H. Moyers, J. L. Chezmer, M. J. Hoel, E. C. Jones, J. E. Peterson, R. D. Cork, and M. E. Bernadino, "Hepatic dynamic sequential CT: section enhancement profiles with a bolus of ionic and nonionic contrast agents," Radiology **178**, 499–502 (1991).

[9]R. B. Lufkin, T. Sharpless, B. Flannigan, and W. Hanafee, "Dynamic-range compression in surface-coil MRI," Am. J. Radiol. **147**, 379–382 (1986).

[10]J. Haselgrove and M. Prammer, "An algorithm for compensation of surface-coil images for sensitivity of the surface coil," Magn. Res. Imag. **4**, 469–472 (1986).

[11]C. E. Hayes and L. Axel, "Noise performance of surface coils for magnetic resonance imaging at 1.5 T," Med. Phys. **12**, 604–607 (1985).

[12]J. Jin, G. Shen, and T. Perkins, "On the field inhomogeneity of a birdcage coil" (abstract), in *Proceedings of the Society of Magnetic Resonance in Medicine*, New York, August, 1993 (Society of Magnetic Resonance in Medicine, Berkeley, CA, 1993).

[13]J. W. Pan, H. P. Hetherington, J. T. Vaughan, and G. M. Prohost, "High resolution neuro-imaging at 4.1T" (abstract), in *Proceedings of the Society of Magnetic Resonance in Medicine*, New York, August, 1993 (Society of Magnetic Resonance in Medicine, Berkeley, CA, 1993).

[14]J. T. Vaughan, J. O. Otu, H. P. Hetherington, P. Noa, and G. M. Prohost, "A high frequency tuned resonator for clinical NMR" (abstract), in *Proceedings of the Society of Magnetic Resonance in Medicine*, Berlin, 1992 (Society of Magnetic Resonance in Medicine, Berkeley, CA, 1992).

[15]S.-P. Liou, *Qualitative Motion Analysis Using a Spatio-Temporal Approach* (dissertation) (University of Michigan, Ann Arbor, 1990).

[16]S.-P. Liou and R. C. Jain, "An approach to three-dimensional image segmentation" (invited paper), Comp. Vision Graphics Imag. Proc.: Image Understand. **53**, 237–252 (1991).

[17]S.-P. Liou, A. H. Chiu, and R. C. Jain, "A parallel technique for signal-level perceptual organization," IEEE Trans. Pattern Anal. Machine Intell. **13**, 317–325 (1991).

[18]S.-P. Liou and R. Jain, "Motion detection in spatio-temporal space," Comp. Vision Graphics Imag. Proc. **45**, 227–250 (1989).

[19]B. Kim, B. D. Ross, S. J. Simerville, and T. L. Chenevert, "MRI determination on intracranial tumor doubling time" (abstract), *Proceedings of the Society of Magnetic Resonance in Medicine*, Berlin, August, 1992 (Society of Magnetic Resonance in Medicine, Berkeley, CA 1992).

[20]M. Joliot and B. M. Mazoyer, "Three-dimensional segmentation and interpolation of magnetic resonance brain images," IEEE Trans. Med. Imag. **12**, 269–277 (1993).

[21]J. M. Henderson, S. B. Heymsfield, J. Horowitz, and M. H. Kutner, "Measurement of liver and spleen volume by computed tomography. Assessment of reproducibility and changes found following a selective distal splenorenal shunt," Radiology **141**, 525–527 (1981).