

Copyright © 2006 IEEE. Reprinted from IEEE Transactions on Medical Imaging, vol. 25, no. 4, pp. 435-450, April 2006.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of Cornell University's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

On Measuring the Change in Size of Pulmonary Nodules

Anthony P. Reeves*, *Senior Member, IEEE*, Antoni B. Chan, David F. Yankelevitz, Claudia I. Henschke, Bryan Kressler, and William J. Kostis

Abstract—The pulmonary nodule is the most common manifestation of lung cancer, the most deadly of all cancers. Most small pulmonary nodules are benign, however, and currently the growth rate of the nodule provides for one of the most accurate noninvasive methods of determining malignancy. In this paper, we present methods for measuring the change in nodule size from two computed tomography image scans recorded at different times; from this size change the growth rate may be established. The impact of partial voxels for small nodules is evaluated and isotropic re-sampling is shown to improve measurement accuracy. Methods for nodule location and sizing, pleural segmentation, adaptive thresholding, image registration, and knowledge-based shape matching are presented. The latter three techniques provide for a significant improvement in volume change measurement accuracy by considering both image scans simultaneously. Improvements in segmentation are evaluated by measuring volume changes in benign or slow growing nodules. In the analysis of 50 nodules, the variance in percent volume change was reduced from 11.54% to 9.35% ($p = 0.03$) through the use of registration, adaptive thresholding, and knowledge-based shape matching.

Index Terms—Computed tomography, growth rate estimation, image registration, image segmentation, pulmonary nodules, rule-based segmentation.

I. INTRODUCTION

PULMONARY nodules are approximately spherical regions of relatively higher density that are visible in X-ray images of the lung. The detection and measurement of pulmonary nodules is important in that a nodule may be a manifestation of a malignant cancer. Large nodules, generally defined as being greater than 1 cm in diameter, may be diagnosed by biopsy. Whereas for smaller nodules, that are too small to biopsy, the diagnostic options are more limited. One of the most promising techniques relies on characterizing the nodule based on its growth rate [1]. The growth rate is estimated by measuring the size of the nodule at two points in time and then using an exponential growth rate to model the change in size over the time period. Therefore, for the early diagnosis of lung cancer, it is important to accurately measure the change in size of pulmonary nodules from

two time-separated computed tomography (CT) scans; in this paper we present computer algorithms that have been designed to perform this task.

A. Related Work

Traditional assessment of pulmonary nodules was performed manually on chest radiographs. One of the earliest recognized and best predictors of the likelihood of nodule malignancy was the demonstration of growth [2], [3]. Determinations of nodule size at two points in time can be used to determine the growth rate, or doubling time (DT), of the nodule [4]. Developments in computer-aided analysis of medical images eventually led to automated characterization of pulmonary nodules in chest radiographs. Methods were developed that can be broadly characterized into three classes: detection, size and growth assessment, and shape characterization.

While DT estimates based on the projective two-dimensional (2-D) information contained in chest radiographs are quite useful, it was not until the advent of computed tomography (CT) that we could begin to characterize these lesions radiographically in three dimensions. Not surprisingly, most of the early work in nodule characterization from CT images was 2-D, based on single cross sections of each lesion [5]–[8]. Some of the 2-D metrics studied included nodule diameter, area, compactness, circularity, location within the lung, intensity, and textural information.

With CT scans there is the possibility of three-dimensional (3-D) image analysis; however, early CT image data was highly anisotropic (where the slice thickness in the axial dimension was 10–20 times the size of the in-plane pixel resolution) and the main focus was on the translation of 2-D algorithms that had been developed for both chest radiographs and other imaging modalities analysis [9]–[13]. In addition, much of the work on nodule characterization had been done for the refinement of nodule candidates in detection systems, leading to the traditionally higher slice thickness needed to acquire the whole chest in a single breath-hold [5], [10], [12], [13]. With the advent of multislice CT, it is now possible to image the entire lung volume using 1-mm sections or thinner in a single breath-hold.

Three-dimensional volumetric analyses have been developed for thin-slice CT data that typically involves voxels with an aspect ratio of less than 5. The initial focus of our research group was on 3-D algorithms for size and growth characterization [1], [6], [14], as growth is the parameter most associated with malignancy. An interesting paper by Okada *et al.* describes a Gaussian fitting method that estimates nodule volume without explicit image segmentation [15], and other volumetric methods have also been reported [16], [17]. However, in general, shape characterization has been the major research focus in computer-aided

Manuscript received January 9, 2006; revised January 16, 2006. This work was supported in part by National Institutes of Health (NIH) under Grant R01-CA-63393, Grant R01-CA-78905, and Grant R33-CA-101110. *Asterisk indicates corresponding author.*

*A. P. Reeves is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853 USA (e-mail: reeves@ece.cornell.edu).

A. B. Chan and B. Kressler are with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853 USA.

D. F. Yankelevitz, C. I. Henschke, and W. J. Kostis are with the Joan and Sanford I. Weill Medical College, Cornell University, New York, NY 14853 USA.

Digital Object Identifier 10.1109/TMI.2006.871548

diagnosis rather than growth evaluation, as classifiers based on shape can assess the likelihood of malignancy using a single CT scan, without waiting for a repeat examination. Work has been done on characterization of overall shape and surface irregularity [12], [13], [18]. These 3-D shape descriptors include analogous metrics to those used in 2-D analysis (e.g., 3-D compactness, sphericity), as well new metrics designed for the 3-D domain (e.g., 3-D surface curvature analysis) [19], [20].

Recently, a number of commercial products have been released by CT vendors and others to measure pulmonary nodules and there have been a number of studies conducted to evaluate them. Since the actual volumes of real pulmonary nodules are unknown, such evaluations usually involve radiologists as the gold standard [21], [22] or involve synthetic phantoms for which the volume is known [23]–[25].

B. CT Images of Pulmonary Nodules

The lungs are unique in that they provide high-contrast images for computer analysis where the solid structures such as airways, blood vessels and nodules have a much higher intensity than the surrounding lung parenchyma. Tumors generally appear as solid tissue having a CT density of approximately that of large vessels. There is a second type of lung cancer that grows along the air-containing structures of the lung (alveoli) and have a nonsolid appearance on CT images; these *subsolid nodules* (also known as *ground-glass opacities*) are not considered in this paper.

Cancers are often modeled as growing at a constant rate; therefore, an exponential growth model may be used to characterize the relationship of tumor volume size over time. Tumor growth rates are often expressed by the DT which is the time, in days, required for a tumor to double in volume. Most lung cancers have DTs ranging from 30 days to several hundred days while benign nodules typically have DTs of more than 400 days.

Given two volumetric measurements of a lesion taken Δt days apart, we may estimate the DT using the following expression [1]:

$$DT = \frac{\ln 2 \cdot \Delta t}{\ln \left(\frac{V_2}{V_1} \right)} \quad (1)$$

where V_1 and V_2 are the first and second volumetric measurements respectively.

In order to minimize the time needed to assess the likelihood of malignancy, we need to be able to accurately measure the nodule volume change in two successive CT scans spaced closely in time. In this paper, we propose algorithms to improve the accuracy of volume change measurement, thus reducing the time needed to assess the likelihood of malignancy of a nodule.

C. Organization

In Section II, we characterize the practical image space for nodule measurement and the significance of partial voxels, and show that isotropic interpolation can improve measurement accuracy. In Section III, we consider methods for accurately segmenting a pulmonary nodule from a single CT scan that extend our earlier work on nodule measurement [14]. We describe new methods for adaptive thresholding and for dealing with

nodules that are attached to a large surface. In Section IV, we develop important refinements to the basic method that take into account variations between the different scans. The two scans are registered together in order to more accurately determine the change in the nodule region between them. This refined change in volume is used to compute the growth rate. Finally, in Section V, we describe a method based on stability of non-malignant nodules to determine the improvement in accuracy due to these more complex algorithms.

II. A TWO-LEVEL IMAGE MODEL FOR PULMONARY NODULES

The CT image properties of pulmonary nodules are considered in this section. The two fundamental image properties optical density (image intensity), and image resolution (including the anisotropic CT image space), are considered in turn. Shape issues of pulmonary nodules are considered in Section III.

A pulmonary nodule may be modeled as an approximately spherical region with a density similar to that of other solid tissue. For CT images of the lung we consider that there are two primary materials with very different densities: air and soft tissue. Furthermore, we consider that the value of a CT image voxel is determined by the ratio of these two materials within that spatial region to which it corresponds. Voxels that represent a region containing both air and soft tissue are called partial voxels and the partial volume fraction (PVF) is the fraction of partial voxels to all voxels that are associated with an image region.

The accuracy of volumetric size estimations made by thresholding the nodule region and counting the number of voxels with soft tissue values is inversely related to the PVF (a partial voxel is counted if more than 50% of its area is covered by soft tissue). Furthermore, for regions with generally smooth surfaces, the PVF is dependent on the region size and the image resolution. An error bound for spherical regions has been determined in [14]. In this section, we characterize the practical bounds for the PVF given current CT technology and a clinically significant nodule size range. Furthermore, by using a synthetic image model, we show that the error due to the PVF may be reduced by supersampling the image region.

A. Two-Level Model and Partial Voxels

Fig. 1 shows a typical 12-mm pulmonary nodule. This nodule has a size of 12 mm and the surrounding image region or nodule region-of-interest (ROI) has dimensions of $193 \times 193 \times 21$ voxels or $36.2 \text{ mm} \times 36.2 \text{ mm} \times 21 \text{ mm}$. The normalized histogram for the whole 3-D image ROI is shown in Fig. 2. The histogram bin size is 5 Hounsfield units (HU) and the histogram is plotted on a normalized scale such that the sum of all the bins is one. The width of the two tissue distributions is attributed to a combination of the PVF and the variation due to actual tissue density variation and image noise. Also shown in Fig. 2 are distributions of manually selected regions for both lung parenchyma and nodule soft tissue that did not involve any partial voxels by visual inspection. The distribution, mean \pm standard deviation ($\mu \pm \sigma$), of the lung parenchyma was found to be $-860 \pm 47.8 \text{ HU}$ and the distribution of the nodule tissue was $-14 \pm 43.1 \text{ HU}$.




Fig. 1. Example of a 12-mm nodule. (a) Axial image through the center of the region of interest, and (b) axial, (c) sagittal, and (d) coronal views of the 3-D segmented image using a light shaded rendering model.

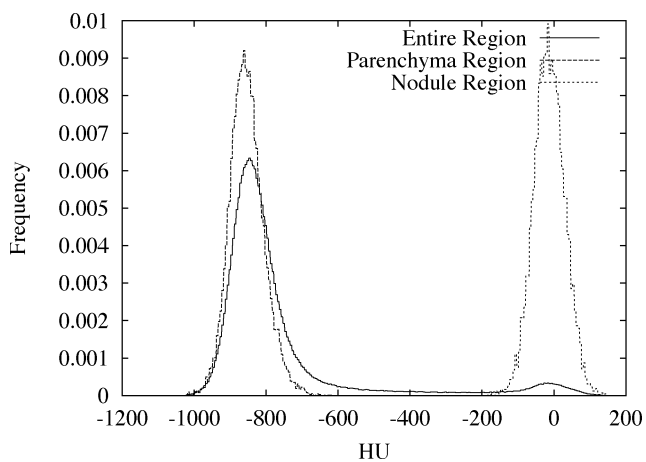


Fig. 2. Normalized density distributions for parenchyma, solid tissue, and the total ROI of the 12-mm nodule in Fig. 1.

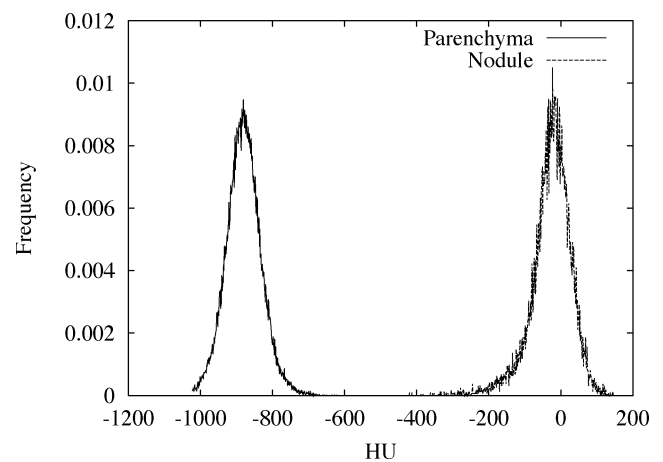


Fig. 3. Normalized density distributions for parenchyma and solid tissue of five representative nodules.

The tissue distribution characteristics can be generalized by estimating the parenchyma and nodule density distributions over a sample set of nodules. Five representative nodules, from 5 mm to 15 mm in size, were selected randomly from the nodule database. A total of 17 870 nodule voxels and 51 468

parenchyma voxels were manually selected for these calculations. The lung parenchyma distribution, mean \pm standard deviation ($\mu \pm \sigma$), was found to be -865 ± 59.7 HU, and the nodule tissue distribution was found to be -12 ± 43.7 HU. The distributions for these tissue types is shown in Fig. 3.

TABLE I
CT IMAGE VOXEL SIZES

Case	XY-size (mm)	Z-size (mm)	Description
A	0.2	1.0	A high resolution diagnostic scan
B	0.2	0.2	Anticipated future CT scanner
C	0.6	2.5	Whole lung scan on a multi-slice CT scanner

The mean value for the nodule tissue is lower than was expected from measurements on large soft tissue regions. However, we found that this was characteristic of our measurements for small lung nodules and, furthermore, the effect was size dependent. For example, the mean pixel value for the nodules of size 10–15 mm was -8 HU while the mean pixel value for nodules of size 5–10 mm was -23 HU. While there may be a significant variation in these distributions between different CT image scans, the graph in Fig. 3 is representative of other scans that we have studied.

B. Nodule Size Considerations

To characterize the effect of scanner reconstruction parameters on nodule size measurement accuracy, three exemplar parameter sets were considered as shown in Table I. Case A in Table I represents a typical high-resolution reconstruction that is used to image a nodule once it has been detected. This corresponds to the highest degree of focus (magnification) possible with the CT scanner; i.e., the finest meaningful reconstruction resolution in the x and y dimensions and the thinnest slice thickness (corresponding to the detector spacing limit and beam collimation limit) of the CT scanner. Case B represents what may be anticipated from a future CT scanner capable of a 0.2-mm slice thickness capability. Such scanner parameters may be anticipated in the near future for scanners that employ higher resolution detectors. Finally, Case C represents the resolution of typical current multi-slice CT scanner for a whole lung scan; The $x - y$ resolution is limited by the need to include the whole body cross section in an single 512×512 pixel image. The axial resolution is limited by the speed of the CT scanner in achieving a whole lung scan in a single breath-hold.

Synthetic 3-D images were created of spherical nodules centered at the origin and with a varying radius r . The value of the voxel V at (x, y, z) in the synthetic image was determined to be

$$V(x, y, z) = (\mu_n - \mu_p) \cdot \alpha + \mu_p \quad (2)$$

where μ_n and μ_p are the mean intensity values for the nodule and parenchyma distributions determined in this section, -12 HU and -865 HU respectively, and α is the percentage of the voxel volume at (x, y, z) that is within the spherical nodule of radius r . By subdividing each voxel volume into 1000 equally sized subvoxels, the value α can be approximated as the percentage of subvoxels with coordinates (x', y', z') that are located inside the spherical nodule of radius r and, thus, satisfy the following equation:

$$(x')^2 + (y')^2 + (z')^2 \leq r^2 \quad (3)$$

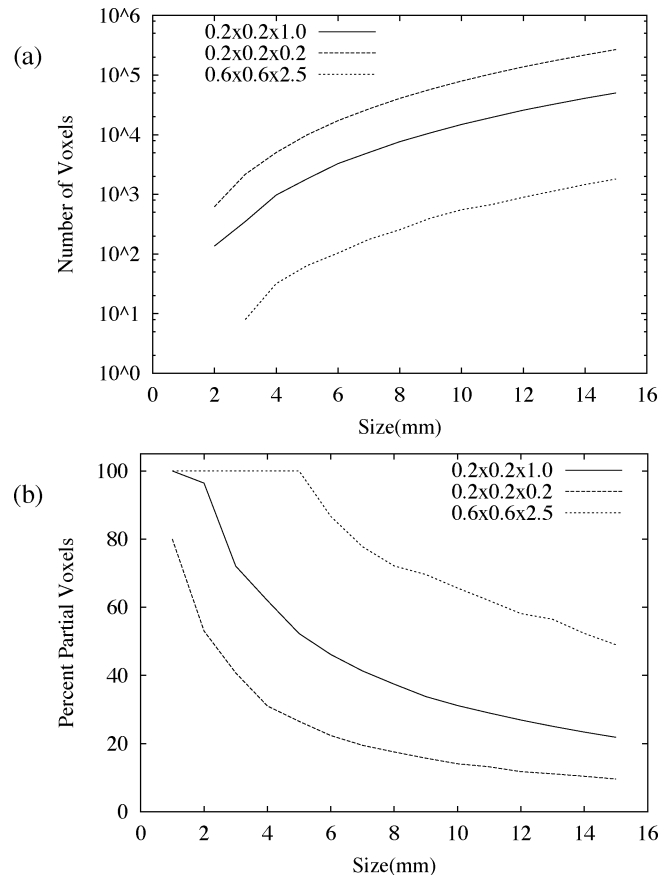


Fig. 4. (a) Number of nodule voxels versus nodule size and (b) percent partial voxels versus nodule size for the scanner parameters given in Table I.

Our analysis involved nodule sizes from 2–15 mm. Nodules smaller than 2 mm are typically too small to routinely detect with current CT scanners. Nodules greater than 10 mm are usually clinically significant by virtue of their size; furthermore, they are large enough to be diagnosed by biopsy or positron emission tomography. Such nodules may also be analyzed by the methods considered in this paper. The number of nodule voxels was measured by counting all the voxels with intensity values greater than the midpoint between the mean parenchyma and mean nodule intensity values ($V > -438.5$ HU). The number of partial voxels was measured by counting all the voxels that were neither the full parenchyma nor the full nodule intensity values ($-865 \text{ HU} < V < -12 \text{ HU}$). Fig. 4(a) shows the number of nodule voxels for the different CT scanner parameter sets defined in Table I. The y-axis is expressed with a logarithmic scale to accommodate the very rapid increase in nodule voxels with nodule size. The PVE expressed as a percentage for each of the spherical nodule sizes is shown in Fig. 4(b). This analysis does not consider the point spread function of the CT scanner or other scanner artifacts which would increase the number of partial voxels.

Two interesting observations may be made from this graph. First, for the whole lung protocol scan where nodules are initially detected, nodules 5 mm and less in diameter have 100% partial voxels. This means that none of the nodule voxels will have the density of solid tissue which directly impairs the detection of such nodules. Second, the percentage of partial voxels is

TABLE II
RESOLUTION PARAMETERS FOR 3-mm NODULES

	Scanner Protocol		
	A	B	C
Number of voxels	352	2176	8
% partial voxels	72.1	40.8	100
Number of complete voxels	98	1288	0

much lower over the size range of interest for the thinner slice protocols. The particular case of 3-mm diameter is considered in detail in Table II. From this table we see that for nodules in the 3-mm size range, size measurement accuracy is highly dependent on the scanner protocol used.

C. Isotropic Resampling

Current Helical CT scanners create images in an anisotropic voxel space. We have used image resampling prior to making nodule measurements [1] to improve measurement accuracy. Also, in some cases, having isotropic data simplifies the implementation of image operations.

In isotropic resampling, the z-direction is resampled (super-sampled) to at least match the resolutions in the x and y dimensions. In some cases, all three dimensions are resampled. Supersampling allows additional shape information to be transferred from the grey-level image to the binary image domain. An alternative approach would be to determine a nodule surface function that is specified to subvoxel accuracy (possibly represented by polygonal mesh) directly from the grey-level values of the anisotropic 3-D image. Subsequent processing, would then be applied to the surface mesh representation. In the limit, both the supersampled binary image and surface representation approaches should lead to similar outcomes.

The improved visual quality due to trilinear isotropic interpolation is shown in [14]. The following experiment using simulated spherical nodules shows that an improvement in measurement consistency can be achieved. Synthetic images of spherical nodule phantoms between 2.0 mm and 8.0 mm in size with a step size of 0.2 mm were generated in a typical high-resolution CT scan resolution space (0.2 mm \times 0.2 mm in-plane and 1.00-mm axial resolution). The spheres were generated using the technique described in the Section II-B. The images were then resampled to an isotropic space (0.2 mm \times 0.2 mm \times 0.2 mm) using trilinear interpolation [26] prior to thresholding and volume estimation. The volume of each sphere was measured for both the resampled isotropic image and the original anisotropic image.

A graph of the percent volume variation versus sphere size for measurements in anisotropic and isotropic voxel-spaces using a size increment of 0.2 mm is shown in Fig. 5. The mean percent volume variation, shown in Table III was calculated over 1-mm size ranges with a size increment of 0.05 mm. In general, as the nodule size increases the percent volume variation decreases, (i.e., the volume measurement is more consistent); this is expected due to the reduction in PVF for larger objects. Table III shows that the mean variation is always less in isotropic space

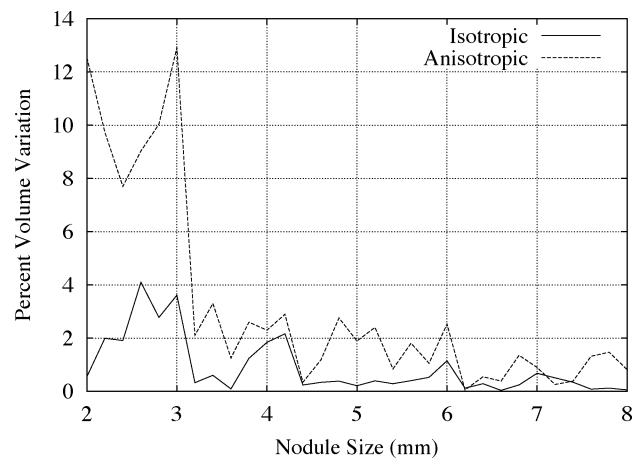


Fig. 5. Percent volume variation versus nodule size for measurements in isotropic and anisotropic voxel-space.

TABLE III
MEAN PERCENT VOLUME VARIATION OF IDEAL SPHERICAL NODULES FOR MEASUREMENTS IN ISOTROPIC AND ANISOTROPIC SPACE

Size Range (mm)	Mean Percent Volume Variation	
	Isotropic	Anisotropic
2 - 3	2.380	7.080
3 - 4	1.368	3.772
4 - 5	1.017	1.941
5 - 6	0.390	1.301
6 - 7	0.386	0.929
7 - 8	0.351	0.651

than in anisotropic space, indicating that volume measurements are more consistent after isotropic resampling. The difference is greatest when measuring small nodules of sizes 2 mm to 4 mm.

In [1], an experiment was conducted to measure the variation of volume measurements from CT images of synthetic spherical phantoms. This experiment was conducted using scanner parameters with a slightly higher x, y resolution than our experiment (0.17 mm versus 0.2 mm). After isotropic resampling, a volume size variation of 0.7% for phantoms of size 3.96 mm was achieved and a variation of 1.3% for phantoms of size 3.2 mm was achieved. These results are in good agreement with the results shown for our simulation shown in Table III.

III. PULMONARY NODULE SEGMENTATION

Segmentation is the most crucial and also most challenging step in the analysis of pulmonary nodules from CT image data. Nodules are frequently attached to other structures, including the local pulmonary vasculature and the pleural surface adjoining the thoracic wall. The geometry of such attachments must be considered in order to successfully segment each type of nodule. We have developed generalized geometric descriptions of the nodules and the attachments, and devised methods for the segmentation of pulmonary nodules in each of these cases.

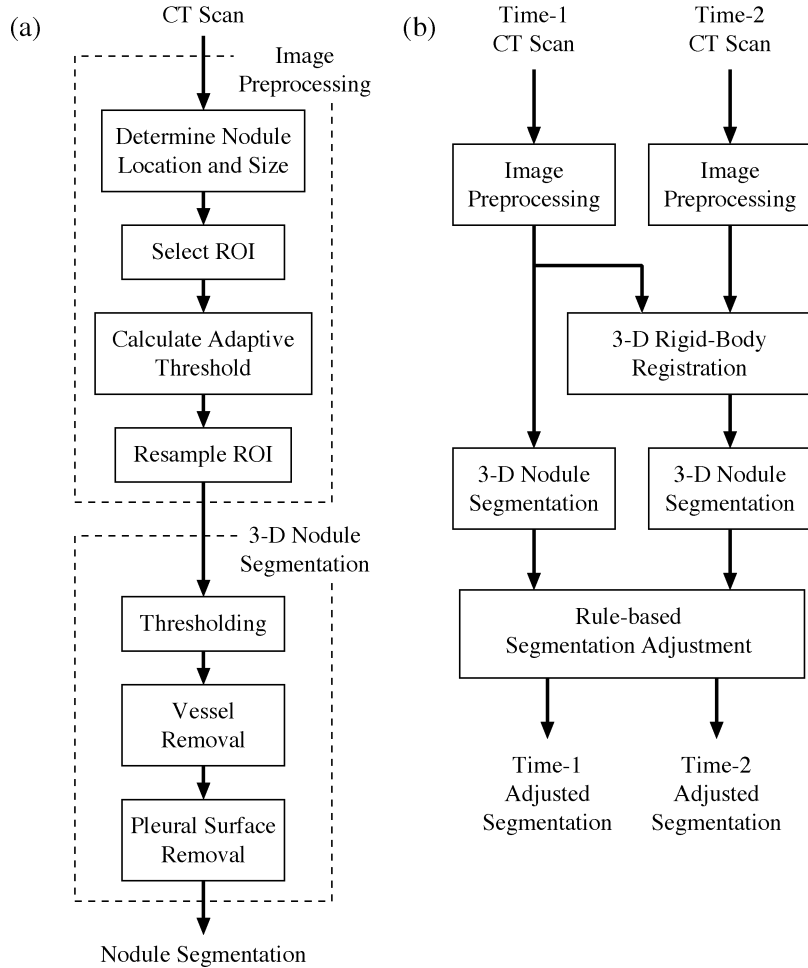


Fig. 6. Segmentation system overview for (a) one nodule and (b) two nodules.

A system overview of the nodule segmentation process is provided in Fig. 6(a). A CT scan first undergoes an image preprocessing stage. In this stage, an approximate location and size of the nodule is determined from the gray level image without segmentation. This information is used to select a region of interest (ROI) around the nodule, an adaptive threshold value is calculated, and the ROI is resampled to isotropic space using trilinear interpolation. Next, the resampled ROI is passed to the 3-D nodule segmentation stage. A threshold is applied to produce a binary image and any attached vasculature structures or pleural surfaces are removed from the nodule.

The following subsections will describe the nodule localization and sizing algorithm, the adaptive threshold calculation, the vessel removal algorithm, and the pleural surface removal algorithm.

A. Nodule Localization and Sizing

The first step in the segmentation process is to determine an approximate location and size of the nodule within the CT image. This information will be used in many of the subsequent stages of the segmentation process.

A method using template functions was developed to estimate the location and size of the nodule in a region without explicit segmentation. Using the two-level model, the nodule is modeled as a high-intensity spherical object surrounded by a

low-intensity background parenchyma. There may also be other high-intensity objects, such as blood vessels or the pleural wall, in proximity to the nodule. We define two template functions for determining the location and size of a nodule, $L_{P,r}(x, y, z)$ and $S_{P,r}(x, y, z)$ respectively. The two template functions are 3-dimensional functions with two parameters, a center point P and a radius r , and the response of a template function is defined as the correlation between the function and the CT image.

The sizing template function $S_{P,r}$ is chosen so that it has a fixed response T when the function is centered over a nodule with radius r and center P . Similarly, the localization template function $L_{P,r}$ is selected so that it has a maximum response when the function is centered over a nodule with radius r and center P . Thus, given an image of a nodule, $I(x, y, z)$, the center P and radius r of the nodule can be determined by selecting the appropriate values for P and r

$$(P, r) = \arg \max_{P, r} \left(\sum_{x, y, z} I(x, y, z) \cdot L_{P, r}(x, y, z) \right) \quad (4)$$

with the size constraint:

$$\sum_{x, y, z} I(x, y, z) \cdot S_{P, r}(x, y, z) = T \quad (5)$$




Fig. 7. Iterations of the nodule localization and sizing algorithm for (a) an isolated nodule, and (b) a pleural nodule. The estimated location is given by a plus, and the estimated radius is represented by the inner circle. The inner circle also represents the positive region of the LOG template function up to the first zero-crossing. The outer circle is twice the radius of the nodule, and the region between the inner and outer circles represents the effective negative region of the LOG template function.

Using the spherical nodule model, the theoretical value T can be determined by finding the response of $S_{P,r}$ of radius r when it is centered over a spherical nodule with the same radius r .

The selection of the two template functions is important and they are chosen to fit the spherical nodule model. The 3-D Gaussian, defined in (7), is selected as the sizing template function because it is normalized with respect to the radius (standard deviation) and, thus, will give a fixed response when the radius of the nodule and the parameter r of the Gaussian are matched

$$S_{P,r}(x, y, z) = G_{P,r}(x, y, z) \quad (6)$$

where

$$G_{P,r}(x, y, z) = \frac{1}{r^3(2\pi)^{\frac{3}{2}}} \exp\left(-\frac{((x-P_x)^2 + (y-P_y)^2 + (z-P_z)^2)}{2r^2}\right). \quad (7)$$

The localization template function, as defined in (9), is the negative Laplacian of the Gaussian (LOG). This function has a positive weight close to its center and a negative weight further out, making it useful in localizing both isolated nodules and nodules on or near the pleural surface. The negatively weighted exterior ensures that the maximum response will occur within the nodule, especially for cases where there is a large solid adjacent structure such as the chest wall

$$L_{P,r}(x, y, z) = -\text{LOG}_{P,r}(x, y, z) \quad (8)$$

$$= \frac{1}{\sigma^2} \left(3 - \frac{x^2}{r^2} - \frac{y^2}{r^2} - \frac{z^2}{r^2}\right) G_{P,r}(x, y, z) \quad (9)$$

The localization and sizing templates are applied to a high-resolution CT image containing a focused region of the lung containing a nodule or nodules. A seed point inside the desired nodule is required as there may be multiple nodules in the scanned region. The response of the sizing function to a nodule of fixed radius decreases monotonically as the function parameter r increases. Using this property, an iterative search scheme was developed to solve for P and r in (4) and (5). Starting at the

seed point inside the nodule and an initial radius of 1, a greedy search is used to perform the maximization in (4) to estimate the nodule center for the given radius. Using the new nodule center estimate, the response to the sizing template function is calculated. If the response is greater than T , then the radius is increased and the process is repeated. If the response is less than T , then the nodule has been successfully localized and sized.

Fig. 7 shows several iterations of this algorithm on an isolated nodule and on a pleural nodule. The nodule location estimate is shown with a plus and the approximate radius is represented by the inner circle. The inner circle also represents the positive region of the LOG template function up to the first zero-crossing. The outer circle is twice the radius of the nodule, and the region between the inner and outer circles represents the effective negative region of the LOG template function. As the radius is increased, the nodule location estimate moves closer to the true center. In these examples, the algorithm was allowed to continue after the termination criteria was met. For the isolated nodule in Fig. 7(a), the localization template function remains over the center of the nodule. However with the pleural nodule in Fig. 7(b), the response of the localization template function has become unstable and the location estimate has moved into the chest wall. This instability occurs when a large portion of the localization template function begins to intersect the boundary of the region of interest. In practice for both cases, the algorithm correctly terminates at the center frame due to the size constraint.

Using the location and size estimate of the nodule, a ROI is selected around the nodule. The ROI is centered over the estimated nodule location with each side equal to 3 times the diameter of the nodule. This ROI of the nodule is used in the rest of the segmentation process.

B. Adaptive Thresholding

The segmentation of nodule tissue from lung parenchyma can be achieved by either a fixed gray-level threshold or by an adaptive threshold. The threshold is selected by choosing a value that best separates the nodule (foreground) from the lung parenchyma (background). In the classical two-level model the intensity histogram will have a bimodal form. The

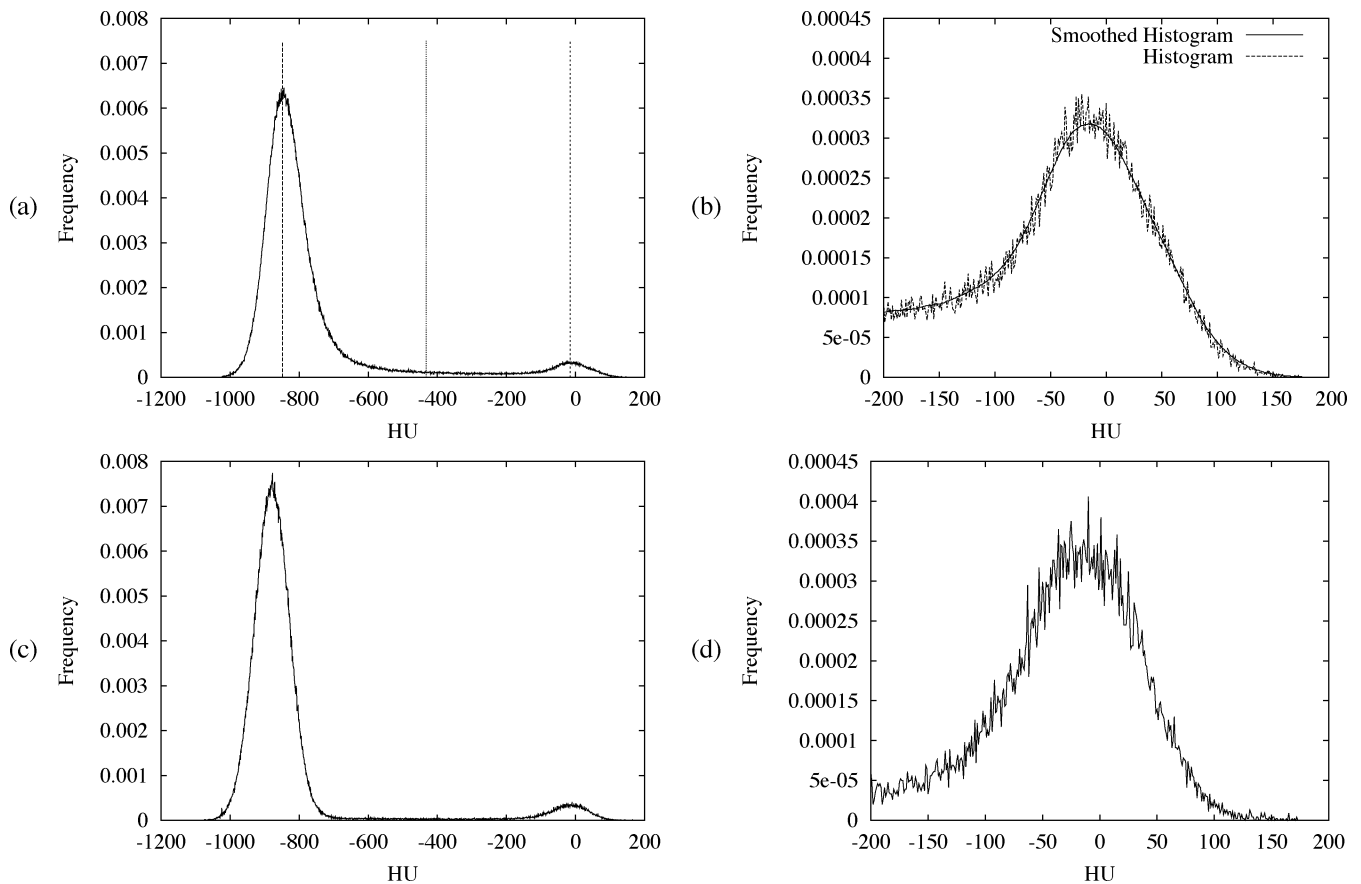


Fig. 8. Histogram of a real 12-mm nodule between (a) -1200 HU and 200 HU, and (b) -200 HU and 200 HU including Gaussian smoothing with ($\sigma = 25$). Histogram over the same ranges, (c) and (d), of a simulated 12-mm spherical nodule.

optimal threshold value that best models the transition between parenchyma and nodule tissues may be determined as the midpoint between the peaks of each mode. In Fig. 8(a), the histogram of the 12-mm nodule shown in Fig. 1 is given with the vertical lines representing the two distribution peaks and the midpoint between them.

A robust fixed threshold may be selected by computing the mean values of lung parenchyma and solid nodule tissue compiled over several cases. Using the mean value of lung parenchyma and solid nodule tissue of -865 HU and -12 HU determined in Section II-A, the fixed threshold was calculated to be -438 HU.

A fixed threshold does not take into account the change in lung parenchyma density due to inspiration of the lungs. In the CT scans we studied, it was found that the density of the parenchyma around the nodule changed on average 9.7 HU \pm 7.0 HU (maximum of 21 HU) between repeat scans of the same patient. Furthermore, it was observed that the lung parenchyma density increases toward the posterior of the lungs (with the patient in the supine position) because of the accumulation of blood due to gravity. In addition, a fixed threshold is not robust to changes in attenuation values due to differences in CT scanner calibration, X-ray dose, and other CT scan parameters.

The goal of the adaptive threshold method is to determine a unique threshold for each scan that compensates for the variations between scans. We define the adaptive threshold for a nodule region as the midpoint between the mean parenchyma

and the mean nodule intensity values. These mean values are approximated by finding the location of the two tissue peaks in the histogram of the region. The histogram is calculated with a bin size of 1 HU and is filtered with a Gaussian ($\sigma = 25$) to reduce the noise caused by the small bin size. The smoothing parameter was determined empirically on the representative nodule data to minimize the noise while maintaining the peak of the nodule distribution [see Fig. 8(b)].

For smaller nodules, the lower number of nodule voxels can cause the solid-tissue distribution peak to be obscured by noise effects including contributions from partial voxels, scanner artifacts and from small vessels. At the limit nodule size of twice the slice thickness, there will be no nonpartial solid voxels at all and, therefore, no peak (this would be a size of 2 mm for the current example of a 1 -mm slice thickness). In practice, the nodule size needs to be larger than this minimum in order to have a statistically significant number of nonpartial voxels (given also the presence of image noise) to produce a well resolved peak value in the histogram. The value of the nodule peak for a given size may be predicted by using the simulated spherical nodule model discussed in Section II. The histogram for this model is shown in Fig. 8(c) and (d).

The histograms of representative nodules of different sizes were examined and the minimum nodule size for the two-peak histogram method was determined empirically to be 6 mm for focused diagnostic scans with a 1 -mm slice thickness. From Fig. 4(b) we see that this limit occurs for a situation in which at

least 60% of the nodules voxels contain only nodule tissue. In cases where the solid-tissue peak value cannot be found through the histogram, the adaptive threshold is calculated as the midpoint between the mean parenchyma value and -12 HU, the average solid-tissue intensity determined in Section II-A.

C. Vessel Removal Using Geometric Constraints

Once a threshold has been determined, the ROI is resampled to isotropic space allowing geometric filters with an isotropic kernel to be applied. The geometric filtering method we have used for removing vessels in the ROI has been used in our earlier studies [1] and is described in detail in [14].

D. Pleural Surface Segmentation

If the nodule is attached to the pleural surface, then the pleural segmentation is performed after the attached vessels have been removed from the nodule. In our earlier work, this was achieved by morphological filtering [14]. An improved method is the following iterative algorithm that separates a nodule from the pleural surface using a clipping plane.

A plane A is constructed that passes through the center of the nodule P and is normal to direction d . The vector d is the direction toward the pleural surface from the center point P , and is approximated as the direction from the center point P to the center of mass of the ROI image. For almost all cases, this approximation is valid because after vessel removal, the largest structure in the nodule ROI is the pleural surface.

The plane A separates the image into a “cut nodule,” which is the connected region on the negative side the plane, and the pleural surface, which is the connected region on the positive side of the plane. The plane is moved in direction d and at each step the change in volume Δ_i is calculated as

$$\Delta_i = V_i - V_{i-1} \quad (10)$$

where V_i is the volume of the “cut nodule” at iteration i . With each iteration, the change in “cut nodule” volume will be relatively constant until the plane intersects the pleural wall, at which point the change in volume will increase dramatically. This event is manifested when the ratio in the change in volume γ , as defined in (11), is greater than a threshold γ_{max}

$$\gamma = \frac{\Delta_i}{\Delta_{i-1}} - 1. \quad (11)$$

When the ratio γ is greater than γ_{max} , a better estimate of the optimal separation plane is calculated by reorienting the plane A about point P to minimize the volume of the “cut nodule.” The ratio γ is recalculated and if the value is less than γ_{max} then the algorithm continues. Otherwise, the plane has reached the pleural wall and the “cut nodule” from the previous iteration is returned as the segmented nodule.

Fig. 9 shows a 2-D example of the algorithm. The cut nodule is shown as dark gray while the rest of the nodule and the pleural surface are light gray. The initial starting point and the initial direction are shown in Fig. 9(a). Fig. 9(b) shows an iteration where the plane does not intersect the pleural wall. Fig. 9(c) shows an iteration where the plane intersects the pleural surface,




Fig. 9. Two-dimensional example of the pleural-wall removal algorithm: the cut nodule is shown as dark gray while the rest of the nodule and the pleural surface are shown in light gray. (a) Given the initial location and direction, (b) the plane is moved in the given direction. (c) The plane eventually intersects the pleural wall causing an increase in the size change of the cut nodule, and (d) the plane is reorientated to minimize the size of the cut nodule. At the next iteration, (e) the plane intersects the pleural wall, but this time (f) reorientating the plane to minimize the cut nodule still results in a large increase in size change. The algorithm terminates, returning the cut nodule from the previous iteration, (d).

causing the change in the nodule size to increase. In Fig. 9(d), the plane is reoriented to minimize the cut nodule, forming a new direction. After another iteration, Fig. 9(e) shows the plane intersects the pleural wall again, but reorientating the plane still leads to a large increase in nodule size, as seen in Fig. 9(f). Thus, the algorithm terminates and returns the cut nodule from the previous iteration, Fig. 9(d).

An example pleural surface segmentation of a nodule is shown in Figs. 10 and 11. A 3-D visualization of the nodule before segmentation and after segmentation is shown in Fig. 10(a) and (b), respectively. The original axial image slices are shown in Fig. 11(a) and (b) shows the resulting segmentation mapped back to the anisotropic axial image space. The white voxels represent the nodule voxels, while the gray voxels represent pleural surface and vessels that were segmented from the nodule.

The pleural segmentation algorithm works best when the surface between the nodule and the pleural wall can be approximated by a plane. Nodules on a convex surface may pose a problem for the segmentation algorithm because the change in




Fig. 10. Example of pleural segmentation: 3-D visualization of the nodule (a) before pleural segmentation, and (b) after segmentation.

volume when the segmentation plane intersects the pleural surface may not be large enough to terminate the algorithm. There are other situations that may cause this algorithm to not work correctly without some assistance. In other work, Mulman [27] showed that for a dataset of 110 pleural nodules the unaided algorithm produced good results for 88.7% of the cases.

IV. IMPROVING SEGMENTATION CONSISTENCY

The accuracy of the change in volume measurement is dependent on the consistency of the segmentations of the nodule in the two images. In the extreme case, a missegmentation of one of the nodules may adversely affect the malignancy predictor by moving its measurement above or below the threshold for malignancy.

There has been some work by Kawata *et al.* [28], [29] on tracking the change of pulmonary nodules in CT images. In [28], the pulmonary nodules are registered together using rigid-body registration and affine registration at two different stages. The nodules are segmented using a 3-D deformable surface model and curvature features are calculated to track the temporal evolution of the nodule. This work was extended in [29], by adding an additional 3-D nonrigid deformable registration stage and the analysis was performed using a displacement field to quantify the areas of nodule growth over time.

The consistency of the nodule segmentation between two scan times can be improved by working with both of the ROIs of the nodule. By registering the two nodule ROIs, we can create a common coordinate space between the regions. With this space, it is now possible to track the change in the nodule in specific areas, along with any inconsistencies in segmentation. These new methods can be incorporated into the existing segmentation procedure to create a new enhanced segmentation system.

The system overview for the segmentation of the same nodule from two CT images is shown in Fig. 6(b). First, each CT scan is passed through the image preprocessing stage, where an ROI is selected, an adaptive threshold is calculated, and the region is resampled to isotropic space. Next, the second nodule region is registered to the first nodule so that both nodules have the same position and orientation in image space. The two nodules are then segmented using the segmentation algorithms discussed previously. Finally, a rule-based segmentation adjustment algorithm is applied to both nodule segmentations. By comparing the nodule segmentations and the thresholded regions, this stage achieves a more consistent measurement of nodule volumes by discounting missegmented nodule voxels using a rule-based system.

The following subsections describe the method used to register two nodules using the 3-D rigid-body transformation and the rule-based segmentation adjustment algorithm.

A. Three-Dimensional Rigid-Body Registration

Given the resampled isotropic ROIs around two nodules, we would like to register the second ROI to the first ROI to facilitate a comparison between the two regions. Because the image regions are of a very focused area of the lung and because the area of interest is composed mainly of solid tissue, we assume that the region will not stretch or skew significantly with different levels of inspiration that may occur during the two CT scans. Thus, the nodule is confined to translation and rotation within the image space, and we can use a rigid-body transformation to register the two nodules. Using a rigid-body model to register the second nodule to the first nodule will preserve the nodule's volume and shape, which is desirable since this is what we wish to measure.

The rigid-body transformation is a function of six parameters: $(t_x, t_y, t_z, r_x, r_y, r_z)$, or translation in x, translation in y, translation in z, rotation about the x-axis, rotation about the y-axis, and rotation about the z-axis, respectively. The rigid-body transformation is a mapping of a point v in 3-D space to a point v' in transformed space defined by the following equation:

$$v' = R_x R_y R_z v + \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix} \quad (12)$$

where R_x , R_y , and R_z are the rotation matrices.

We define the mean-squared difference (MSD) as a similarity metric that quantifies how well two regions are registered together, and is defined as

$$\text{MSD} = \frac{1}{N} \sum_{x,y,z} \sum_{x,y,z} (I_2(x,y,z) - I_1(x,y,z))^2 \quad (13)$$

where I_1 and I_2 are the two regions and N is the number of voxels in I_1 .

The two regions are registered by conducting a search to minimize the MSD over the transform parameters. Powell's method ([30] and [31]), a multi-dimensional direction-set search algorithm, is used as the search strategy because it does not explicitly use derivatives. The initial translation parameter is set to align the center of the second nodule to the center of the first nodule and the rotation parameters are set to zero. After the registration is complete, the second ROI is set to have the same size as the first ROI.




Fig. 11. Example of pleural segmentation: (a) slices from the original image, and (b) the nodule segmentation (white represents nodule, and gray represents segmented pleural surface and vessels).




Fig. 12. Example of rigid-body registration: (a) the first ROI, (b) the second ROI, (c) the second ROI registered to the first ROI, (d) the difference image between the first ROI and the second ROI, (e) the difference image between the first ROI and the registered second ROI.

An example of rigid-body registration on two nodules is shown in Fig. 12. Mid-nodule image slices for the first ROI and the second ROI are shown in Fig. 12(a) and (b), respectively. The registered mid-image is seen in Fig. 12(c). Fig. 12(d) shows the difference image between the first ROI and the second ROI. The difference image shows that the vessels at the top of the image and on the right of the image are not aligned properly. The circular white ring is the growth of the nodule between the two scan times. Fig. 12(e) shows the difference image between the first ROI and the registered second ROI. The vessels are less visible in the second difference image than in the first difference image, meaning that with the registered image the vessels and nodule are aligned better than without registration.

In some cases, the movement of large periphery structures such as vessels or pleural surfaces relative to the nodule may cause misregistration of the nodule. Because the MSD metric equally weighs the alignment of all parts of the image, the larger structure will be registered at the expense of mis-aligning the smaller nodule. This occurs frequently when registering nodules

with a pleural tail (slight attachments to the pleural surface) because differences in lung inspiration may cause the pleural surface to move dramatically relative to the nodule.

A Gaussian-weighted MSD metric can be used to improve registration of nodules with large periphery structures, and is defined as

$$\text{MSD}_g = \frac{1}{N} \sum_{x,y,z} \exp\left(\frac{(g_x - x)^2 + (g_y - y)^2 + (g_z - z)^2}{2g_r^2}\right) \times (I_2(x, y, z) - I_1(x, y, z))^2 \quad (14)$$

where (g_x, g_y, g_z) is the center of the nodule and g_r is the radius of the nodule as determined by the sizing and localization algorithm in Section III-A. An appropriately sized Gaussian function will weigh the alignment of the nodule over the alignment of any of the periphery structures, while still using the attached structures to contribute to the registration of the nodule.

A comparison of registration metrics on a benign pleural-tail nodule is shown in Fig. 13. For a benign nodule, a good nodule




Fig. 13. Comparison of registration metrics on a pleural-tail benign nodule: given the benign nodule in (a) and (b), the second nodule is registered to the first using the MSD metric (c). The difference image (d) shows that the images are misregistered. The second nodule is registered using the Gaussian-weighted MSD (e). The difference image (f) shows that the nodule is registered correctly, even though the high-intensity bone is not.

registration will show little change around the nodule in the difference image. A mid-nodule slice from the original scans is shown in Fig. 13(a) and (b). The second nodule is registered to the first nodule using the MSD metric, and the resulting image is shown in Fig. 13(c). The difference image between the registered image and the first nodule, seen in Fig. 13(d), shows that the nodules are badly misregistered. In contrast, the second nodule is registered using the Gaussian-weighted MSD metric and is shown in Fig. 13(e). In the difference image shown in Fig. 13(f), we see that the nodule has been correctly registered even though the high-intensity bone has not.

B. Rule-Based Segmentation Adjustment

After the registration is complete the two nodules are then segmented using the methods described in the previous section. If the nodule registration is good then it is possible to make a

correspondence between the voxels from the first region to the voxels in the second region. This is a powerful tool in that now we are able to track the nodule region over a time interval with voxel accuracy.

By comparing the corresponding voxels between the segmented nodules and the thresholded images, it is possible to label the voxels in the segmented nodules as nodule repeat voxels, nodule growth voxels, nodule atrophy voxels, or missegmented voxels. Using this knowledge, the nodule segmentations are adjusted by removing the missegmented voxels. This improves the consistency of the segmentations of the nodule in the two different times, thus improving the accuracy of the volume measurements.

Let S_1 be the segmented nodule from the first image (time 1) and T_1 be the threshold image of the first image (before vessel or pleural surface removal). Likewise, let S_2 be the segmented




Fig. 14. Nodule segmentation adjustment: given two thresholded nodule regions, T_1 and T_2 , each ROI is segmented to produce S_1 and S_2 respectively. Active voxels in S_1 are compared with voxels in S_2 and T_2 to mark regions as (a) nodule in time 1, (b) nodule atrophy, or (c) missegmentation in S_1 . Active voxels in S_2 are compared with voxels in S_1 and T_1 to mark regions as (d) nodule in time 2, (e) nodule growth, or (f) missegmentation in S_2 . The missegmented voxels are removed resulting in the adjusted nodule segmentations N_1 and N_2 .

nodule from the second image (time 2) and T_2 be the thresholded second image. A rule-based system is used to mark active voxels in the segmented nodule S_1 as repeat nodule, nodule atrophy, or nodule missegmentation. If an active voxel in the segmented nodule of time 1 corresponds to an active voxel in the segmented nodule of time 2, then that voxel is a repeat nodule voxel because it is present as a nodule in both times. If an active voxel in the segmented nodule of time 1 corresponds to an inactive voxel in the segmented nodule and the thresholded image of time 2, then that voxel is a nodule atrophy voxel because it is present in the first nodule, but disappeared in the second image. Finally, if an active voxel in the segmented nodule of time 1 corresponds to an inactive voxel in the segmented nodule of time 2 and an active voxel in the thresholded region of time 2, then that voxel is a missegmented nodule voxel in time 1 because it was marked as a nodule voxel in time 1 but was also marked as a nonnodule object (vessel or chest wall) in time 2.

An active voxel in the segmented nodule S_2 may be marked as a repeat nodule voxel, a nodule growth voxel, or a nodule missegmented voxel by using a similar set of rules. A summary of the rules for marking a region in S_1 or S_2 are given below.

- (A) Repeat Voxel in Time 1: $R_1 = S_1 \wedge S_2$;
- (B) Atrophy Voxel in Time 1: $A_1 = S_1 \wedge (T_2)'$;
- (C) Missegmented Voxel in Time 1: $M_1 = S_1 \wedge (S_2)' \wedge T_2$;
- (D) Repeat Voxel in Time 2: $R_2 = S_2 \wedge S_1$;
- (E) Growth Voxel in Time 2: $G_2 = S_2 \wedge (T_1)'$;
- (F) Missegmented Voxel in Time 2: $M_2 = S_2 \wedge (S_1)' \wedge T_1$;

where \wedge is the AND operation between binary images, and $(X)'$ is the complement of the binary image X .

Fig. 14 illustrates an example of the segmentation adjustment on a registered vascularized nodule. The thresholded nodule at two different times is shown in Fig. 14(a) and (b), and the segmented nodules are shown in Fig. 14(c) and (d). In the first segmentation, part of the top attached vessel has been missegmented as part of the nodule, and in the second segmentation, part of the bottom attached vessel has been missegmented as part of the second nodule. By using the rules described above, the regions of each segmentation are marked in Fig. 14(e) and (f) as (A) nodule in time 1, (B) nodule atrophy, (C) nodule missegmentation in time 1, (D) nodule in time 2, (E) nodule growth,

or (F) nodule missegmentation in time 2. Finally, the segmentations of both nodules are adjusted in Fig. 14(g) and (h) by removing the missegmented regions (C) and (F). The volumes of the nodules in the adjusted segmentations are more accurate than the volumes in the original segmentations because the vessel attachments have been removed. This leads to a more accurate determination of DT or percent volume change (PVC). Furthermore, regions of growth and atrophy of the nodule can be examined over time.

Fig. 15 shows an example of the rule-based segmentation adjustment on a real nodule. The two nodule regions I_1 and I_2 are thresholded (T_1 and T_2) and segmented (S_1 and S_2). The vessel in S_1 was removed correctly, however part of the vessel adjacent in S_2 has not been successfully segmented from the nodule. Using segmentation adjustment the voxels of the missegmented vessel and the voxels with nodule growth can be identified. Assuming that the two nodule images are registered correctly, the rule-based segmentation adjustment will correct the nodule segmentation if the nodule feature is properly segmented in one of the images. If a nodule feature is missegmented in both images, then no segmentation adjustment will take place; however, in this case, the missegmentation does not greatly affect the change in volume calculation because the missegmented feature will appear to be a stable part of the nodule.

V. CONSISTENCY VALIDATION THROUGH ANALYSIS OF BENIGN NODULES

We have developed several segmentation systems to measure the change in volume of a nodule. We would like to compare the consistency of each system to determine if using the new algorithms are better than before. A consistent algorithm is defined as an algorithm that is capable of measuring the change in volume of a nodule with minimal measurement error.

One method of validating consistency is to use CT scans of phantom nodules. Using phantom nodules is advantageous because the growth rate between repeat scans is zero, making for straightforward calculations of measurement error. However, these phantom nodules are usually simple spheres with no surface features and little or no nodule attachments, but, more importantly, they are only a rough approximation to real *in vivo*




Fig. 15. Example of segmentation adjustment: in regions I1 and I2, a nodule has grown against the adjacent vessel. The two nodule regions I1 and I2 are thresholded (T1 and T2) and segmented (S1 and S2). The vessel in S1 was removed correctly, however part of the vessel adjacent in S2 has not been successfully segmented from the nodule. Using segmentation adjustment the missegmented vessel voxels and the voxels with nodule growth can be identified.

lung nodules. Our algorithms were designed to counter segmentation inconsistencies of nodules with surface features and major attachments. Therefore, little advantage is gained when using our algorithms on phantom nodules with simple surface features and no attachments.

For the comparison of algorithms, we desire a data set consisting of real pulmonary nodules showing little growth between repeat scans. Same-day repeat scans of a nodule would be ideal. However, these repeat scans are usually not available since lung cancer screening protocols are designed to minimize radiation exposure of patients. Given the above constraints, we devised an experiment to evaluate the accuracy of the growth measurement methods based on nodules that were observed to be stable rather than growing at the classical rapid exponential growth curve associated with malignant nodules. We selected a data set consisting of nodules that have been observed to have no significant growth for a period of two years. Next we will show how statistical methods can be used to quantitatively compare volume change measurements from different segmentation algorithms.

A measurement of the change in the volume of a stable nodule is a result of the combination of the actual nodule size change and the measurement error. Consider a data set consisting of N nodules, each with two scans, and a segmentation algorithm that measures the size of a nodule in each of the two scans. The actual growth of the nodule and the error in the growth measurement is modeled as two independent random variables, A and E , respectively. The actual growth of the nodules has a mean of μ_a and a variance of σ_a^2 . The measurement error can be modeled as a Gaussian with mean μ_e and variance σ_e^2

$$E \sim \mathcal{N}(\mu_e, \sigma_e^2). \quad (15)$$

The growth measured by the algorithm is the combination of the actual nodule growth and the measurement error

$$M = A + E. \quad (16)$$

Since A and E are independent random variables, the mean μ_m and variance σ_m^2 of the measured growth can be expressed as

$$\mu_m = \mu_a + \mu_e \quad (17)$$

$$\sigma_m^2 = \sigma_a^2 + \sigma_e^2. \quad (18)$$

Therefore, the variance of the measured growth of the N nodules in the data set is equal to the sum of the variance due to any change in nodule sizes and the variance of the measurement error.

The actual growth distribution is fixed with the selection of the data set. For a given data set, a change in the measured growth between two segmentation algorithms is caused solely by a change in the measurement error. The most accurate algorithm will on average have a lower measurement error, which will result in a reduction of variance in the measured growth distribution. Therefore, different segmentation algorithms can be quantitatively compared for accuracy using the variance (or standard deviation) of the growth measurement over a fixed population of nodules. While there may be systematic errors in the volume measurements of the nodules which are not apparent in the variance measurements, consistency in volume measurement between scans is of greater concern than absolute accuracy because growth is being measured rather than absolute volume.

An experiment was performed to validate the benefits of using adaptive thresholding and rule-based adjustment over the fixed threshold segmentation algorithm. Stable nodules with repeat scans were selected for the experiment.

A. Growth Validation Method and Materials

Fifty stable nodules, either diagnosed as benign through biopsy or observed to have little growth over a period of two years (2-year-no-change or 2YNC nodules), were selected from the Cornell screening databases. Each nodule has two high-resolution CT scans with an axial resolution of 1.00 mm and an in-plane resolution of 0.1875 mm. The scans were obtained with a GE Medical Systems HiSpeed CT scanner with a tube current between 200 mA to 330 mA and a peak voltage of 120 kV to 130 kV. Although the nodules were tracked over two years, the repeat scans were selected such that the interval between scans was the shortest possible given the screening protocol.

Of the 50 nodules, 33 were 2–5 mm nodules, 15 were 5–9 mm nodules, and 2 were larger than 10 mm. The average nodule size was 4.4 mm. Ten nodules had repeat scans within 100 days, while 13 had repeat scans between 100 and 200 days, 11 had repeat scans between 200 and 300 days, 12 had repeat scans

TABLE IV
NODULE DATABASE POPULATION: SIZE AND TIME BETWEEN SCANS
FOR STABLE NODULES

Size (mm)	Number	Time Between Scans (days)	Number
2 - 4	33	0 - 100	10
5 - 9	15	100 - 200	13
10+	2	200 - 300	11
Total	50	300 - 400	12
		400+	4

between 300 and 400 days, and 4 had repeat scans over 400 days. The average separation time between repeat scans was 230 days. Table IV shows the breakdown of the nodule population by size and time between scans.

A seed point in each CT scan was selected by hand for each nodule case. Four variations of the proposed segmentation algorithm were applied to each nodule pair, as listed below:

- 1) segmentation using a fixed threshold;
- 2) segmentation using a fixed threshold and rule-based adjustment;
- 3) segmentation using an adaptive threshold;
- 4) segmentation using an adaptive threshold and rule-based adjustment.

Each segmentation method was compared quantitatively by analyzing the variation of the nodule volume change for each method.

The experiment was run on a 1.5 GHz Pentium 4 with 256 MB RAM using unoptimized research programs. The computation time for a nodule segmentation varies with the size of the nodule. For a 12-mm nodule, the segmentation took 1.5 minutes for the nodule in the first scan and 4 minutes for the nodule in the repeat scan. The segmentation for the repeat scan is slower because rigid-body registration is used. In general, smaller nodules will take less time and larger nodules will take more time.

B. Growth Validation Results

The experiment was run on the 50 nodules and the resulting segmentations were visually inspected for any gross missegmentations, such as failure of the pleural removal algorithm. In the case of a gross missegmentation, the appropriate segmentation parameter was adjusted and the nodule was resegmented.

Traditionally, the DT is used as the measurement of growth rate of pulmonary nodules. However, our experiment was concerned with the amount of measured growth regardless of the time between repeat CT scans. To compare the different segmentation techniques we first evaluated the volume change measured for each nodule. A significant ($P = 0.03$) decrease was seen for the standard deviation of the rule-based adjustment with an adaptive threshold ($\sigma = 29.15$) compared to the fixed threshold method ($\sigma = 38.2$).

For discussion, the PVC was selected as the change metric as it is a normalized measurement that relates to measurement accuracy independent of the nodule size. The PVC is defined as

$$P = 100 \cdot \frac{V_2 - V_1}{V_1} \quad (19)$$

TABLE V
STANDARD DEVIATION OF THE PERCENT VOLUME CHANGE FOR FOUR
SEGMENTATION TECHNIQUES

Segmentation Technique	All Nodules	5 mm	5 mm
Fixed Threshold Segmentation	11.54	8.66	15.69
w/ Rule-Based Adjustment	9.68	7.77	12.58
Adaptive Threshold Segmentation	11.07	7.95	15.39
w/ Rule-Based Adjustment	9.35	7.44	12.22

where V_1 and V_2 are the nodule volumes measured from the time one and time two nodule segmentations, respectively. The standard deviation of the PVC for the 50 nodule pairs was calculated for each segmentation technique and appears in Table V.

For the 50 nodules, the fixed threshold segmentation produced a PVC standard deviation of 11.54, while using rule-based adjustment produced a lower standard deviation of 9.68. Using an adaptive threshold lowered the standard deviation to 11.07, only slightly better than that using a fixed threshold. Using rule-based adjustment with an adaptive threshold produced the lowest standard deviation of 9.35.

The nodules were split into two groups, nodules less than 5 mm in size and nodules greater or equal to 5 mm in size. Rule-based adjustment has a greater impact when the nodules are greater than 5 mm. This happens because larger nodules are more likely to have spicules or significant attachments to vessel structures; therefore, the rule-based adjustment will have a greater affect on the segmentation correction.

For many of the nodules, using the advanced segmentation techniques produced a small improvement. This is because these were nodules without any significant attachments. On the other hand, nodules that showed a large decrease in variation had significant vessel or pleural attachments. Rule-based adjustment is most effective in improving the segmentation consistency when there are significant attachments.

VI. CONCLUSION

The measurement of the growth rate of pulmonary nodules is a very good predictor of malignancy. In this paper, we have presented techniques to improve the accuracy in growth measurement from CT images.

Accurate and robust nodule growth measurement is deceptively challenging both to implement and to validate. However, accurate measurements and knowledge of the accuracy of these measurements are critical to the optimal use of growth measurement for the diagnosis of cancer. Measurement accuracy depends upon many factors including the nodule size, its location with respect to other objects, and the CT scanner parameters, the foremost of which is reconstruction resolution. A second challenge is to acquire two CT scans with exactly the same scanner parameters which can be quite difficult in clinical practice.

Methods have been presented that match two images by density (adaptive thresholding), by location (registration), and by vessel removal consistency (rule-based segmentation). A significant improvement in growth rate measurement accuracy was achieved by using these methods together compared to

the more standard approach of measuring each nodule image individually.

Measurement evaluation is typically performed either by using calibrated phantoms that may not accurately model real nodules or by radiologists subjective evaluations of real nodules for which the accuracy and bias is unknown. Using stable nodules for growth evaluation is a first step involving real nodules that is freer (less biased) than using conventional radiologist measurements as a gold standard. We have achieved promising results with 50 stable nodules; however, more data is necessary to validate the methods for use in a clinical setting. Furthermore a better, more complete, evaluation will require the measurement of growing nodules with well established growth rates in addition to the stable nodules.

REFERENCES

- [1] D. F. Yankelevitz, A. P. Reeves, W. J. Kostis, B. Zhao, and C. I. Henschke, "Small pulmonary nodules: volumetrically determined growth rates based on CT evaluation," *Radiology*, vol. 217, no. 1, pp. 251–256, Oct. 2000.
- [2] V. P. Collins, R. K. Loeffler, and H. Tivey, "Observations on growth rates of human tumors," *Am. J. Roentgenol.*, vol. 76, pp. 988–1000, 1956.
- [3] M. H. Nathan, V. P. Collins, and R. A. Adams, "Differentiation of benign and malignant pulmonary nodules by growth rate," *Radiology*, vol. 79, pp. 221–231, 1962.
- [4] M. Schwartz, "A biomathematical approach to clinical tumor growth," *Cancer*, vol. 14, no. 6, pp. 1272–1294, Nov.–Dec. 1961.
- [5] M. L. Giger, K. T. Bae, and H. MacMahon, "Computerized detection of pulmonary nodules in computed tomography images," *Invest. Radiol.*, vol. 29, no. 4, pp. 459–465, Apr. 1994.
- [6] D. F. Yankelevitz, R. Gupta, B. Zhao, and C. I. Henschke, "Small pulmonary nodules: evaluation with repeat CT—preliminary experience," *Radiology*, vol. 212, no. 2, pp. 561–566, Aug. 1999.
- [7] M. F. McNitt-Gray, E. M. Hart, N. Wyckoff, J. W. Sayre, J. G. Goldin, and D. R. Aberle, "A pattern classification approach to characterizing solitary pulmonary nodules imaged on high resolution CT: preliminary results," *Med. Phys.*, vol. 26, no. 6, pp. 880–888, Jun. 1999.
- [8] D. Cavouras, P. Prassopoulos, and N. Pantelidis, "Image analysis methods for solitary pulmonary nodule characterization by computed tomography," *E. J. Radiol.*, vol. 14, pp. 169–172, 1992.
- [9] M. L. Giger, K. Doi, and H. MacMahon, "Image feature analysis and computer-aided diagnosis in digital radiography. 3. Automated detection of nodules in peripheral lung fields," *Med. Phys.*, vol. 15, no. 2, pp. 158–166, Mar.–Apr. 1988.
- [10] S. Yamamoto, I. Tanaka, M. Senda, Y. Tateno, T. Iinuma, T. Matsumoto, and M. Matsumoto, "Image processing for computer-aided diagnosis of lung cancer by CT (LSCT)," *Syst. Comput. Japan*, vol. 25, no. 2, pp. 67–80, Feb. 1994.
- [11] T. Shimizu, I. Narabayashi, Y. Uesugi, K. Tabuchi, R. Namba, Y. Nakata, Y. Saika, R. Matsui, K. Sueyoshi, and T. Kawai, "Three-dimensional display of pulmonary nodules using helical CT," *Radiat. Med.*, vol. 13, no. 5, pp. 209–216, Sep.–Oct. 1995.
- [12] S. Toshioka, K. Kanazawa, N. Niki, H. Satoh, H. Ohmatsu, K. Eguchi, and N. Moriyama, "Computer aided diagnosis system for lung cancer based on helical CT images," *Proc. SPIE*, vol. 3034, pp. 975–984, Feb. 1997.
- [13] S. G. Armato III, M. L. Giger, C. J. Moran, J. T. Blackburn, K. Doi, and H. MacMahon, "Computerized detection of pulmonary nodules on CT scans," *Radiographics*, vol. 19, no. 5, pp. 1303–1311, Sep.–Oct. 1999.
- [14] W. J. Kostis, A. P. Reeves, D. F. Yankelevitz, and C. I. Henschke, "Three-dimensional segmentation and growth-rate estimation of small pulmonary nodules in helical ct images," *IEEE Trans. Med. Imag.*, vol. 22, no. 10, pp. 1259–1274, Oct. 2003.
- [15] K. Ocada, D. Comanicie, and A. Krishnan, "Robust anisotropic gaussian fitting for colometric characterization of pulmonary nodules in multislice ct," *IEEE Trans. Med. Imag.*, vol. 24, no. 3, pp. 409–423, Mar. 2005.
- [16] E. D. Bomsztyk and H. Rusinek, "Volumetric analysis of lung nodules using a hybrid algorithm," in *Proc. SPIE (Medical Imaging 2004: Image Processing)*, vol. 5370, J. M. Fitzpatrick and M. Sonka, Eds., 2004, pp. 22–29.
- [17] Y. Kawata, N. Niki, H. Ohmatsu, M. Kusumoto, R. Kakinuma, K. Mori, K. Yamada, H. Hishiyama, M. Kaneko, and N. Moriyama, "A computerized approach for estimating pulmonary nodule growth rates in three-dimensional thoracic ct images based on ct density histogram," in *Proc. SPIE (Medical Imaging 2005: Image Processing)*, vol. 5747, J. M. Fitzpatrick and J. Reinhardt, Eds., 2005, pp. 872–882.
- [18] A. P. Reeves, W. J. Kostis, D. F. Yankelevitz, and C. I. Henschke, "Three-dimensional shape characterization of solitary pulmonary nodules from helical CT scans," in *Proc. Computer Assisted Radiology and Surgery (CARS'99)*, Jun. 1999, pp. 83–87.
- [19] Y. Kawata, N. Niki, H. Ohmatsu, K. Eguchi, and N. Moriyama, "Shape analysis of pulmonary nodules based on thin section CT images," *Proc. SPIE*, vol. 3034, pp. 964–974, Feb. 1997.
- [20] W. J. Kostis, A. P. Reeves, D. F. Yankelevitz, and C. I. Henschke, "Three-dimensional curvature analysis of small pulmonary nodules in helical CT scans," *RSNA Scientific Program 2000, suppl. Radiology*, vol. 217P, p. 549, Nov. 2000.
- [21] D. Wormanns, G. Kohl, E. Klotz, A. Marheine, F. Beyer, W. Heindel, and S. Diederich, "Volumetric measurements of pulmonary nodules at multi-row detector CT: *in vivo* reproducibility," *Eur. Radiol.*, vol. 14, pp. 86–92, 2004.
- [22] M. Revel, C. Lefort, A. Bissery, M. Bienvenu, L. Aycard, G. Chatellier, and G. Fria, "Pulmonary nodules: preliminary experience with three-dimensional evaluation," *Radiology*, vol. 231, pp. 459–466, May 2004.
- [23] J. M. Goo, T. Tongdee, R. Tongdee, K. Yeo, C. F. Hildebolt, and K. T. Bae, "Volumetric measurement of synthetic lung nodules with multi-detector row CT: effect of various image reconstruction parameters and segmentation thresholds on measurement accuracy," *Radiology*, vol. 235, no. 3, pp. 850–856, Jun. 2005.
- [24] J. P. Ko, H. Rusinek, E. L. Jacobs, J. S. Babb, M. Betke, G. McGuinness, and D. P. Naidich, "Small pulmonary nodules: volume measurement at chest CT—phantom study," *Radiology*, vol. 228, no. 3, pp. 864–870, Sep. 2003.
- [25] H. Jiang and K. Kelly, "Theoretical prediction of lung nodule measurement accuracy under different acquisition and reconstruction conditions," in *Proc. SPIE (Medical Imaging 2004: Physiology, Function and Structure From Medical Images)*, vol. 5369, A. A. Amininiind and A. Manduca, Eds., 2004, pp. 406–412.
- [26] C. W. Dreveny and L. T. Bruton, "Least-squares trilinear interpolation for digital video coding," in *Proc. 2000 IEEE Int. Symp. Circuits and Systems (ISCAS 2000)*, vol. 4, Geneva, Switzerland, 2000, pp. 701–704.
- [27] Y. D. Mulman, "Segmentation of juxtapleural lung nodules from the pleural surface from CT images using a robust surface estimate," M.Sc. thesis, Cornell Univ., Ithaca, NY, 2003.
- [28] Y. Kawata, N. Niki, H. Ohmatsu, M. Kusumoto, R. Kakinuma, K. Mori, N. Nishiyama, K. Eguchi, M. Kaneko, and N. Moriyama, "Tracking interval changes of pulmonary nodules using a sequence of three-dimensional thoracic images," *Proc. SPIE (Medical Imaging 2000: Image Processing)*, vol. 3979, pp. 86–96, 2000.
- [29] ———, "Analysis of evolving processes in pulmonary nodules using a sequence of three-dimensional thoracic images," in *Medical Imaging 2001: Image Processing, Proceedings of SPIE*, vol. 4322, M. Sonka and K. M. Hanson, Eds., 2001, pp. 1890–1901.
- [30] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Sueten, "Multimodality image registration by maximization of mutual information," *IEEE Trans. Med. Imag.*, vol. 16, no. 2, pp. 187–198, Apr. 1997.
- [31] W. Press, Ed., *Numerical Recipes in C*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1992.