Integrated Bronchoscopic Video Tracking and 3D CT Registration for Virtual Bronchoscopy

W.E. Higgins, J.P. Helferty, and D.R. Padfield

Penn State University, University Park, PA 16802 USA

[JH]Scitec, Princeton, NJ 08540

[DP]General Electric, Schenectady, NY

ABSTRACT

Lung cancer assessment involves an initial evaluation of 3D CT image data followed by interventional bronchoscopy. The physician, with only a mental image inferred from the 3D CT data, must guide the bronchoscope through the bronchial tree to sites of interest. Unfortunately, this procedure depends heavily on the physician's ability to mentally reconstruct the 3D position of the bronchoscope within the airways. In order to assist physicians in performing biopsies of interest, we have developed a method that integrates live bronchoscopic video tracking and 3D CT registration. The proposed method is integrated into a system we have been devising for virtual-bronchoscopic analysis and guidance for lung-cancer assessment. Previously, the system relied on a method that only used registration of the live bronchoscopic video to corresponding virtual endoluminal views derived from the 3D CT data. This procedure only performs the registration at manually selected sites; it does not draw upon the motion information inherent in the bronchoscopic video. Further, the registration procedure is slow. The proposed method has the following advantages: (1) it tracks the 3D motion of the bronchoscope using the bronchoscopic video; (2) it uses the tracked 3D trajectory of the bronchoscope to assist in locating sites in the 3D CT "virtual world" to perform the registration. In addition, the method incorporates techniques to: (1) detect and exclude corrupted video frames (to help make the video tracking more robust); (2) accelerate the computation of the many 3D virtual endoluminal renderings (thus, speeding up the registration process). We have tested the integrated tracking-registration method on a human airway-tree phantom and on real human data.

Keywords: virtual bronchoscopy, 3D imaging, pulmonary imaging, video tracking

1. INTRODUCTION

Modern lung-cancer assessment involves 3D CT assessment followed by bronchoscopy.[1,2] In recent years our research group has been constructing tools for image-guided lung-cancer assessment.[3–8] For a particular patient, our tools are used in two stages:

1. Stage 1 (Off-line 3D CT Assessment): The patient's 3D CT chest scan undergoes airway-tree segmentation,[7] airway centerline analysis,[6,8] and region-of-interest (ROI) definition. The airway-tree segmentation provides data for computing endoluminal airway renderings. The centerline analysis provides path trajectory information for navigating through the major airways. Finally, the ROIs denote suspect mediastinal lymph nodes or peripheral nodules to consider for follow-on bronchoscopic biopsy. These steps result in a case study, which provides guidance data during the next stage.[3,4]

2. Stage 2 (Image-Guided Bronchoscopy): With the computer interfaced to the bronchoscope and the previously computed case study loaded, the system provides guidance to each of the ROIs. The live video stream from the bronchoscope is fed into the computer and used for registration (matching) to the 3D CT "virtual world." It is this registration step that enables image guidance of bronchoscopy.[5]

Previously, this system relied on a method that only used static registration of the live bronchoscopic video to corresponding virtual endoluminal views derived from the 3D CT data.[5] The system only performed the registration at manually selected 3D CT sites. It did not draw upon the motion information inherent in the bronchoscopic video. Further, the registration procedure was slow, as it involved computing a large number of endoluminal renderings. In this paper, we introduce a method that integrates live bronchoscopic video tracking and faster 3D CT registration.

Send correspondence to WEH; email: weh2@psu.edu

An endoscope moving through the bronchial tree can be viewed as a camera moving through a rigid body. Previous research has demonstrated the possibility of calculating the change in camera position in 3D coordinates based on the 2D movement of scene information.[9,10] These methods involve translating the position change obtained from optical flow calculations into a 3D motion vector. Horn and Weldon introduced a method of directly calculating motion parameters by using the optical flow constraint equation.[11] These methods form the basis of the tracking/matching method we have developed.[12]

Mori *et al.* presented a technique to track the motion of endoscopic video frames. Their technique is able to track a significant number of frames, but requires over 20 seconds of processing time per frame.[13,14] Mori *et al.* also described a method for effectively identifying misleading frames resulting from the presence of bubbles. This method, however, relies on the matching and tracking of every video frame and involves a correlation of video and CT data, thus necessitating a large amount of computation.[15] Because this method requires a registration step for every frame in a video sequence, it cannot be applied to our method, which matches only every few frames and employs optical flow parameters for the tracking of those frames between each matching step. Also, the above methods have only been tested on video-taped sequences and are presently too computationally intense for live endoscopy procedures.

Our proposed method has the following advantages: (1) it tracks the 3D motion of the bronchoscope using the bronchoscopic video; (2) it uses the tracked 3D trajectory of the bronchoscope to assist in locating sites in the 3D CT "virtual world" to perform the registration. In addition, the method incorporates techniques to: (1) accelerate the computation of the many 3D virtual endoluminal renderings (thus, speeding up the registration process); and (2) detect and exclude corrupted video frames (to help make the video tracking more robust). A companion paper describes the combined video tracking and matching algorithm.[12] This paper focuses on the details involved in the accelerated endoluminal rendering calculations and in detecting/excluding of corrupted video frames.

2. METHODS

During Stage-2 Bronchoscopy, the virtual-bronchoscopy computer is interfaced to the bronchoscope hardware, so that it receives a live video stream from the bronchoscope. Guidance to a particular planned ROI biopsy site begins by highlighting the precomputed trajectory (path) to the ROI. Initially, the physician manually picks an obvious site, such as the main carina, on this trajectory; this brings the 3D CT "virtual world" to this position. The physician then moves the bronchoscope near this site, and a registration is done between the 3D CT virtual world and the live "real world" of the bronchoscopic video. Once this initial registration is done, the physician picks a new site further down the precomputed trajectory, closer to the biopsy site of interest. The physician then moves the bronchoscope in the direction of this site, based on this visual cue; this action initiates a combined tracking/matching algorithm, depicted in Figure 1 and described previously by Helferty and Higgins.[12]

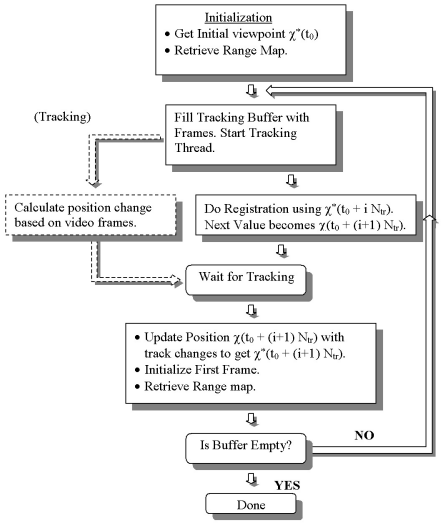

Figure 1. Block diagram of previously proposed algorithm for combined registration and tracking.[12]

The video produced during this movement is fed into a near real-time 3D optical-flow algorithm to track the motion of the bronchoscope. Also, at periodic intervals during this movement, the video (real world) is registered to the 3D CT data (virtual world). The 3D tracking information is used to update the previous registration location before the more precise registration is done. An optimization procedure, based on the mutual information criterion, is used for registration. Tracking only is generally done on roughly 75% of all frames, and full-blown video-CT registration is done on the remainder. Because drift error accumulates fairly rapidly with the optical-flow method, the more robust matching step must be done to ensure a continued alignment of the video and CT data. The matching step also provides an updated range map for the optical-flow method.

Two issues are crucial to the functionality of this integrated tracking-registration procedure. The first issue is the slowness of the virtual endoluminal rendering process. A typical video-CT registration step will require between 50 and 200 different computed endoluminal views, which are tested against the candidate video frame. Previously, we used a brute-force marching-cubes rendering algorithm for all airway surface triangles to compute these views. This results in considerable wasted computation. We have devised a method that predetermines the relevant subset of triangles for each site along a trajectory of interest. Only triangles within the actual field of view of a site are preserved for each site. Also, our method takes into account that for incremental movements, only a small number of the previously used triangles may change for the new site.

The second issue is that misleading video frames can greatly affect the accuracy of the optical-flow tracking calculations. Notice in Figure 1 that no processing is done on the frame to determine whether it is faulty before calculating the optical flow parameters. These misleading frames must be removed from consideration during video tracking. Four types of misleading frames occur in bronchoscopy (Figure 3): (1) drop outs, resulting from electrical interference; (2) wash outs, resulting from water or mucus; (3) wall frames, which occur when the bronchoscope hits a wall; and (4) bubble frames, which result when the patient breathes or coughs. All of these commonly occurring events confuse the tracker because it is not able to recognize these video sequences as faulty and thus attempts to track them. We have devised a series of simple image processing algorithms to detect these frames.

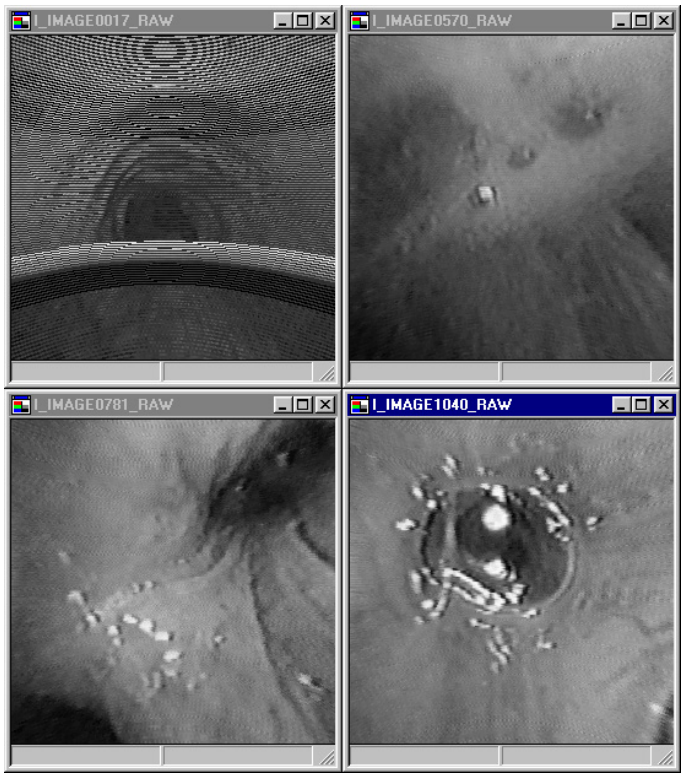

Figure 2. Examples of four types of misleading video frames. The frame in the upper left is a drop-out frame. The upper right frame is a wash-out frame. The lower left frame is a wall frame. The lower right frame is a bubble frame. These frames belong to the 1100-frame sequence from case h005, per Table 3.

The remainder of this section highlights our efforts with these issues.

2.1. Rapid Endoluminal Rendering

Figure 3 depicts the approach for rapidly rendering the required endoluminal views. The basic idea is to only render those airway branches that are within the observer's direct line of sight; all other branches need not be rendered because they cannot be seen. We omit considerable small-level detail in the discussion below. These details can be found in Dirk Padfield's thesis.[16]

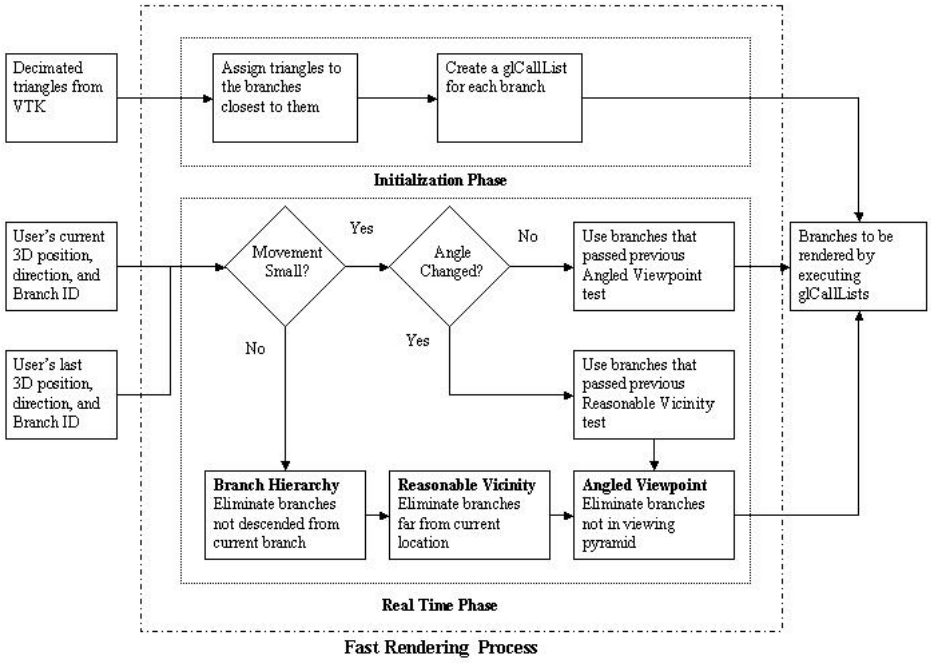

Figure 3. Block diagram for fast endoluminal rendering.

We now outline the steps in this approach. Initially, an off-line calculation is made whereby all precomputed airway-tree triangles are first associated with a particular guidance branch. For a given triangle, this is done by finding the branch that is closest in terms of Euclidean distance. These triangle associations are in practice stored for later use as part of the `glCallLists` used by OpenGL for rendering.[17]

The next series of calculations occur in real-time during live bronchoscope tracking and registration. Assume that the observer's current 3D position within the airway tree is known. In practical terms, this is done by first registering the current video position to the 3D CT virtual world. With a registration done, we can use the current 3D position to determine for this position its current branch (e.g., left main bronchus), parent branch (e.g., trachea), and all branches derived from the current branch. For our system, all path information used for navigating through the major airways is precomputed and stored in the path file.[6,8] The path file, loaded during navigation, contains detailed 3D viewing site, branch, and path data for the airway tree.

Three separate methods are then run to eliminate branches from rendering. First, the Branch Hierarchy Method only includes triangles for branches that lie in the subclass defined above (current branch, parent branch, sibling branches). We point out that this method assumes that the bronchoscope points forward (a reasonable assumption as it cannot look backwards!) and that *all* descendant branches, regardless of their distance from the current position, are considered. Thus, the next two methods attempt to further limit the number of branches considered. The Reasonable Vicinity Method eliminates branches that have no part of their extent within a fixed distance from the current location. Such far-off branches are not deemed to be visible from the current position. The distance we use is the length of the longest branch in the airway tree (typically the trachea), a conservative choice. Finally, the Angled Viewpoint Method eliminates branches that are not within the viewing pyramid (for rendering) defined

by the current 3D location and viewing direction.

Subsequent 3D navigation positions (viewing sites) tend not to differ greatly from the current 3D position. For such circumstances, it is often true that the same, or nearly the same, subset of triangles can be used for rendering new endoluminal views. Thus, the fast rendering method first determines if a small movement was made for the next 3D position. If it is not small, then all three sets of calculations described earlier must be redone. Otherwise, a check is made to see if the viewing direction (angle) changed a sufficiently small amount. If this change is too large, then calculations done earlier up to the Reasonable Vicinity Method can be reused but the Angled Viewpoint Method must be rerun. Otherwise, all triangles for the subset of branches associated with the previous viewing position can be reused for the new position. When the actual renderings are calculated during the registration steps, only the indicated subset of triangles is used.

2.2. Misleading Video-Frame Rejection

The algorithms for video-frame rejection are deliberately designed to be simple, so that the processing speed of the overall video tracking algorithm stays near real time. We point out that it is not necessary to remove all corrupted video frames, since periodic registration steps are done throughout the tracking/matching process. With these preliminary comments, we know introduce the Video-Frame Rejection methods, as shown in Figure 4. Please note that the methods constituting this processing *replace* the block entitled "Calculate position change based on video frames" shown in the previous Helferty-Higgins method of Figure 1.

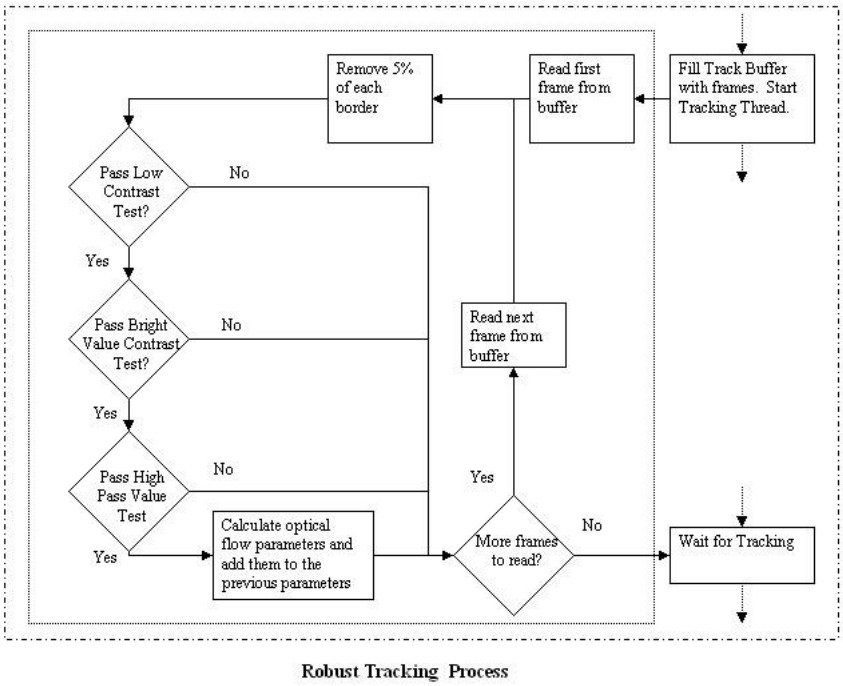

Figure 4. Block diagram for video-frame rejection. This diagram replaces the block delimited by dotted lines and entitled "Calculate position change based on video frames" shown in Figure 1.

At the outset of processing video frames from the tracking buffer, we assume that the first frame is normal. This can easily be guaranteed since the tracking/matching process is triggered manually. Further, the outer 5% of all frames is reomved from processing, since the bronchoscope sometimes incorrectly includes a dark misleading border, arising from distortion (these data can mislead the Bright-Value Contrast test, discussed below).

A succession of three simple tests are then run on each frame, with two preliminary processing steps, as summarized below:

1. Apply a simple window average filter to produce a smooth image, free of noise spikes:

$$I \rightarrow I_{ave}$$

2. Compute the histogram of I_{ave}. This histogram will be used to locate various brightness levels in the image.

3. Low Contrast Test — Reject frames having insufficient (bland) contrast. Such frames occur from wash-outs and drop-outs. In particular, if

$$\text{contrast} = P_{99\%} - P_{1\%} < T_1,$$

then reject the frame. $P_{X\%}$ is defined as the pixel gray-level where $X\%$ of the pixels in I_{ave} are below gray-level $P_{X\%}$ and T_1 is a threshold.

4. Bright-Value Contrast Test — Bubbles appearing in endoscopic video frames often cause bright reflections and produce bright patches accounting for approximately 5% of the image pixels. Some misleading wall frames are also identified by this method. This test appears below:

   (a) If

$$\text{brightness range} = P_{100\%} - P_{95\%} > T_2$$

   then continue with the second test below. Otherwise, pass the frame.

   (b) If the distance between the darkest and brightest clusters of pixels in I_{ave} is small, then reject the frame. Frames exhibiting such close concentrated bright and dark regions are typically bubble or wall frames. See Dirk Padfield's thesis for further detail.[16]

5. High-Pass Filter Test — Various line artifacts and patterns can appear in corrupted frames as a result of electrical interference. Such frames must be rejected. These frames are easily detected by the following test:

   (a) Compute a high-pass filtered image:

$$I_{highpass} = I - I_{ave}$$

   (b) If the sum of the absolute values of all pixels in $I_{highpass}$ exceeds a threshold T_3, then reject the frame. (The frame has too much energy in high-frequency components, which typically produce line artifacts.)

3. RESULTS

We provide results for the proposed methods for rapid endoluminal rendering and misleading video-frame rejection. We also give an overview of combined tracking/matching performance in the Discussion.

3.1. Tests for Rapid Endoluminal Rendering

Figure 5 overviews the locations of viewing sites considered for tests on human case h005. Table 1 gives the results of the branch/triangle elimination methods. Notice that triangles not within the user's field of view are successively eliminated by the methods. The final column of the table gives the percentage of the number of branches and triangles that were not eliminated by any of the methods; these are the triangles that are then rendered. Note that the percentage of branches retained and triangles retained are related but not necessarily the same. This is because different branches have different numbers of triangles; for example, the branch in the trachea normally has many more triangles that make up its wall than other branches because the airway is thicker and longer. As the results make clear, anywhere from 62% to 98% of the triangles were eliminated by the methods, saving on rendering computation.

Figure 6 depicts the rendering for site E of this example. Notice that no triangles are (apparently) missing from this view; this is true for the other sites as well.

It is necessary to also consider the overall increase in rendering speed. (It is possible that the increase in computation necessitated by the elimination methods might outweigh any benefits in decreasing the speed of individual renderings.) Comparisons of the matching times at the various locations for h005 is given in Table 2. As the table shows, the overall rendering speed is reduced significantly.

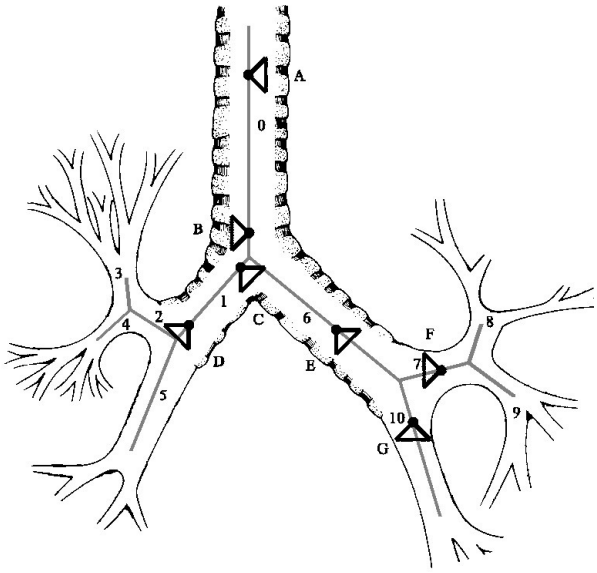

Figure 5. Example bronchial tree with various user locations and directions. The tested locations for human case h005 are marked A through G at points on the outside of the tree closest to their respective locations. Point F is not included in thr results presented here, as it represents an impossible site for a bronchoscope to reach; consult Dirk Padfield's thesis for further detail.[16]

Loc.	Branch Hierarchy		Reasonable Vicinity		Angled Viewpoint		% Retained	
A	126	100,000	15	11,227	1	2,468	1%	2%
B	126	100,000	39	23,020	5	4,921	4%	5%
C	126	100,000	56	35,851	27	20,598	21%	21%
D	54	41,562	35	20,849	35	20,849	28%	21%
E	71	55,969	67	54,251	43	37,703	34%	38%
F	—	—	—	—	—	—	—	—
G	3	4,733	3	4,733	3	4,733	2%	5%

Table 1. Number of branches and triangles retained by the Branch Elimination method for case h005. 126 branches are in this tree and 100,000 triangles were used to represent the bronchial tree. The first column gives the IDs of the locations shown on the bronchial tree of Figure 5. The first number under each method indicates the number of branches that were retained by the method, and the second number indicates the number of triangles retained. The branches retained by each method are passed to the next method. The last column displays the percentage of branches and triangles, respectively, that were retained overall.

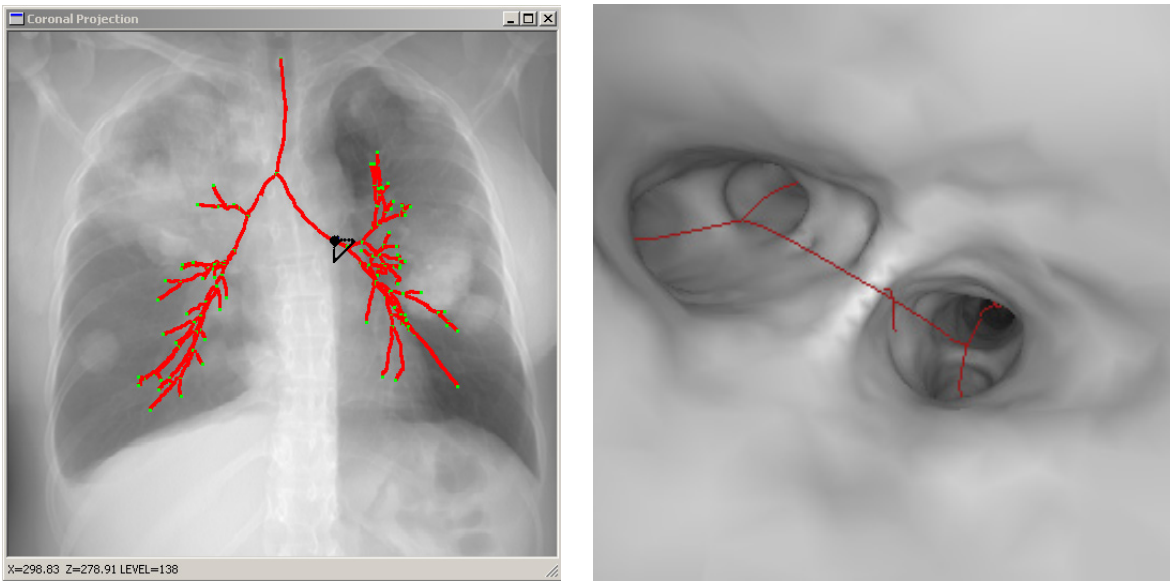

Figure 6. 3D position and direction of location E of the example bronchial tree. The left image, a coronal maximum-intensity projection of 3D CT data set and associated path file [lines], gives the location and viewing direction of the user in the tree, as identified by a large dot attached to a viewing triangle. The right image is the CT-based endoluminal rendering rendering at the location.

Loc.	Old	New	Times Faster
A	10.66	6.11	1.74
B	11.79	7.22	1.63
C	12.19	6.87	1.77
D	12.82	6.23	2.06
E	16.50	8.54	1.93
Avg.	12.79	6.99	1.83

Table 2. Comparison of matching time results for h005. All results are given in seconds. The first column gives the locations on the bronchial tree that were matched. The next two columns compare timing. "Old" refers to the system before the implementation of the branch elimination methods, and "New" refers to the system using the new methods. The last column gives the amount of speed up resulting from the new methods. The last row gives the average of these results over all locations.

Wash-out	Wall	Bubble	Drop-out
9-15	847-853	1026-1054	8
17-19			16
21-23			20
556-600			25-32
613			37-38
623-630			573
634-645			
974-979			
1056-1058			

Table 3. Ground truth table for the 1100-frame h005 video sequence. The first row gives the names of the various types of misleading frames. The numbers below each misleading frame type corresponds to the frame or range of frames of the video sequence that are this type of misleading frame. There are a total of 88 wash-out frames (8% of the sequence), 7 wall frames (1% of the sequence), 35 bubble frames (3% of the sequence), and 14 drop-out frames (1% of the sequence).

Frame Type	Total	Correct Detection		Missed Detection		False Positive
	No.	No.	% of Total	No.	% of Total	No.
Wash-out	88	46	52.27	42	47.73	3
Wall	7	0	0.00	7	100.00	4
Bubble	35	26	74.29	9	25.71	0
Drop-out	14	14	100.00	0	0.00	0
TOTAL	144	86	59.72	58	40.28	7

Table 4. Results of misleading frame detection for the h005 sequence. The first column gives the type of frame being detected. The second column gives the total number of frames of that kind in the sequence based on the ground truth given in 3. The third and fourth columns give the number and percentage of frames correctly detected, the next two columns give the number and percentage of frames that were missed, and the last column gives the number of normal frames that were detected as this type of frame.

3.2. Tests for Misleading Video-Frame Rejection

Tests were conducted on video sequences from three human case studies (h005, h006, and h008).[16] For each video sequence, a ground-truth table was found for the number of wash-out, wall, bubbles, and drop-out frames in the sequence. Table 3 gives an example ground-truth table for case h005. If a frame was identified as more than one type of misleading frame, then precedence was given to it being a wash-out frame first, then a wall frame, a bubble frame, and a drop-out frame. For example, if a frame had characteristics of a wall frame and a drop-out frame, it would be identified as a wall frame. (We did this because of the precendence of the methods used to find these features; if the frame fails one test, it is considered misleading and skips the other tests.) After the sequences were tested by the methods, the number and types of frames eliminated by the Low Contrast, Bright-Value Contrast, and High-Pass Filter tests were determined. The results of these tests for case h005 is summarized in Table 4. Note that a significant percentage of misleading frames are detection, with virtually no false positives. Our results placed a premium on not rejecting good frames (i.e., on not having false positives). By adjusting the thresholds, more frames would conceivably detected, but at a much higher false positive rate.

4. DISCUSSION

We have tested the integrated tracking-registration method on a human airway-tree phantom and on real human data. For live bronchoscopic guidance within the phantom, we have observed that the procedure works well and quickly. Several hundred video frames, representing approximately 6 cm of motion, have been used in the combined video tracking and 3D CT registration; the process took roughly two minutes, more than an order of magnitude faster than previously proposed research. Regarding human tests, we have found that 55-82% of the triangles on

average are rejected during the rendering process (3 separate human cases considered), greatly speeding up the registration time. Regarding video processing, we have found 13%-73% of the video frames are misleading (wash-out and bubble-frame portions of a sequence can last for many seconds); our detection procedure successfully rejected on the order of 60-70% of these frames.

The proposed method offers a promising approach beyond merely registering manually selected discrete sites. It better uses the motion information inherent in the bronchoscopic video, hence enabling a tighter coupling between the real bronchoscopic video world and virtual 3D CT world. Live human bronchoscopy, however, is rife with "uncooperative events," such as coughing, electrical interference, etc. Thus, human intervention is essential for overseeing the process.

ACKNOWLEDGMENTS

This work was partially supported by grants #CA74325 and #CA91534 from the NIH.

REFERENCES

1. H. P. McAdams, P. C. Goodman, and P. Kussin, "Virtual bronchoscopy for directing transbronchial needle aspiration of hilar and mediastinal lymph nodes: a pilot study," *Am. J. Roentgenology*, vol. 170, pp. 1361–1364, May 1998.
2. I. Bricault, G. Ferretti, and P. Cinquin, "Registration of real and CT-derived virtual bronchoscopic images to assist transbronchial biopsy," *IEEE Transactions on Medical Imaging*, vol. 17, no. 5, pp. 703–714, Oct. 1998.
3. A. J. Sherbondy, A. P. Kiraly, A. L. Austin, J. P. Helferty, S. Wan, J. Z. Turlington, E. A. Hoffman, G. McLennan, and W. E. Higgins, "Virtual bronchoscopic system combining 3D CT and endoscopic video," *SPIE Medical Imaging 2000: Physiology and Function from Multidimensional Images*, vol. 3978, pp. 104–115, A. Clough and C.T. Chen, eds., 2000.
4. J. P. Helferty, A. J. Sherbondy, A. P. Kiraly, J. Z. Turlington, E. A. Hoffman, G. McLennan, and W. E. Higgins, "Image-guided endoscopy system for lung cancer assessment," *IEEE International Conference on Image Processing 2001*, pp. 307–310, Oct. 7-10 2001.
5. J. P. Helferty and W. E. Higgins, "Technique for registering 3D virtual CT images to endoscopic video," *IEEE International Conference on Image Processing 2001*, pp. 893–896, Oct. 7-10 2001.
6. R. Swift, A. Kiraly, A. Sherbondy, A. L. Austin, E. A. Hoffman, G. McLennan, and W. E. Higgins, "Automatic axes-generation for virtual bronchoscopic assessment of major airway obstructions," *Computerized Medical Imaging and Graphics*, vol. 26, no. 2, pp. 103–118, March-April 2002.
7. A. P. Kiraly, E. A. Hoffman, G. McLennan, W. E. Higgins, and J. M. Reinhardt, "3D human airway segmentation methods for virtual bronchoscopy," *Academic Radiology*, vol. 9, no. 10, pp. 1153–1168, Oct. 2002.
8. A. Kiraly and W. E. Higgins, "Analysis of branching tubular structures in 3D digital images," *IEEE International Conference on Image Processing 2002*, pp. II–333 — II–336, 2002.
9. K. I. Diamantaras, T. Papadimitriou, M. G. Strintzis, and M. Roumeliotis, "Total least squares 3-D motion estimation," *IEEE International Conference on Image Processing-98*, vol. 1, pp. 923–927, 1998.
10. T. Jebara, A. Azarbayejani, and A. Pentland, "3D structure from 2D motion," *IEEE Signal Processing Magazine*, pp. 66–84, May. 1999.
11. B. K. Horn and R. J. Weldon, Jr., "Direct methods for recovering motion," *International Journal of Computer Vision*, vol. 2, pp. 51–76, 1988.
12. J. P. Helferty and W. E. Higgins, "Combined endoscopic video tracking and virtual 3D CT registration for surgical guidance," *IEEE International Conference on Image Processing 2002*, pp. II–961 — II–964, Sept. 22-25 2002.
13. K. Mori, Y. Suenaga, J. Toriwaki, J. Hasegawa, K. Katada, H. Takabatake, and H. Natori, "A method for tracking camera motion of real endoscope by using virtual endoscopy system," *SPIE Medical Imaging 2000: Physiology and Function from Multidimensional Images, A. Clough and C.T. Chen, eds.*, vol. 3978, pp. 122–133, Feb. 12-17, 2000.
14. H. Shoji, K. Mori, J. Sugiyama, Y. Suenaga, J. Toriwaki, H. Takabatake, and H. Natori, "Camera motion tracking of real endoscope by using virtual endoscopy system and texture information," *SPIE Medical Imaging 2001: Physiology and Function from Multidimensional Images, A. Clough and C.T. Chen, eds.*, vol. 4321, Feb 18-22 2001.
15. D. Deguchi, K. Mori, J. Hasegawa, J. Toriwaki, and H. Natori et al., "Camera motion tracking of real bronchoscope using epipolar geometry analysis and CT derived bronchoscopic images," *SPIE Medical Imaging 2002: Physiol. Func. from Multidim. Images*, vol. A Clough and C. Chen (ed.), v. 4683, pp. 30–41, 2002.
16. D. Padfield, "Integrated image registration and tracking of endoscopic video and 3D CT images," May 2002. M.S. Thesis, Penn State University.
17. R. Fosner, *OpenGL: Programming for Windows 95 and Windows NT*. Reading, MA: Addison Wesley, 1997.