Tumor Detection in the Bladder Wall with a Measurement of Abnormal Thickness in CT Scans

Sylvain Jaume*, Matthieu Ferrant, Benoît Macq, *Senior Member, IEEE*, Lennox Hoyte, Julia R. Fielding, Andreas Schreyer, Ron Kikinis, and Simon K. Warfield, *Associate Member, IEEE*

***Abstract*—Virtual cystoscopy is a developing technique for bladder cancer screening. In a conventional cystoscopy, an optical probe is inserted into the bladder and an expert reviews the appearance of the bladder wall. Physical limitations of the probe place restrictions on the examination of the bladder wall. In virtual cystoscopy, a computed tomography (CT) scan of the bladder is acquired and an expert reviews the appearance of the bladder wall as shown by the CT. The task of identifying tumors in the bladder wall has often been done without extensive computational aid to the expert. We have developed an image processing algorithm that aids the expert in the detection of bladder tumors. Compared with an expert observer reading the CT, our algorithm achieves 89% sensitivity, 88% specificity, 48% positive predictive value, and 98% negative predictive value.**

***Index Terms*—Anatomical atlas, bladder wall, computer-aided diagnosis, thickness measurement, tumor detection, virtual cystoscopy, virtual endoscopy.**

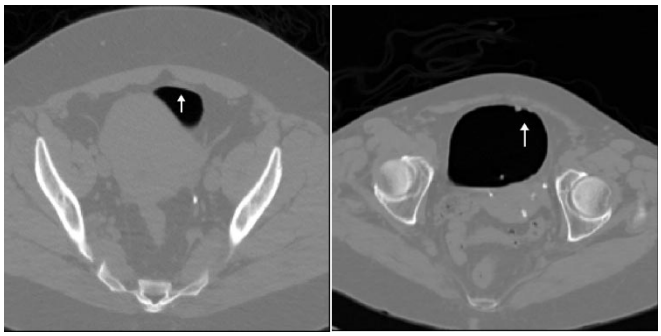

Fig. 1. CT scan of the abdomen shows the bladder cavity in black, and the bladder wall in light gray (thin contour around the cavity). (*left*) Slice through an abdomen CT scan for one patient. (*right*) Slice through an abdomen CT scan for a second patient. The arrows point to wall regions with abnormal thickness.

I. Introduction

TUMORS of the bladder constitute 7% of all malignancies in men and 4% in women [1]. Every year, bladder tumor is diagnosed in more than 50 000 subjects. If doctors detect the tumor when it is still small, they can treat it with a resection, or chemotherapeutic agents, or a combination of both. However, in 60% of the cases, another bladder tumor will appear [2], [3]. Therefore, the patients need regular monitoring to avoid the recurrence of the disease.

Cystoscopy is the conventional procedure for monitoring the bladder [4]. The clinician inserts a probe into the patient's bladder, and orients it to observe on a monitor the surface of the bladder cavity. The clinician has to distinguish the surface deformations due to tumors, from the surface deformations due to the neighboring structures. As a routine exam, conventional cystoscopy exhibits four major problems. First, it is painful; hence, many patients refuse the exam [5]. Second, it increases the risk of urinary tract infection [6]. Third, it does not easily show the diverticula and the bladder neck, especially in males [7]. Fourth, it only shows the inner surface of the bladder wall; the interior of the wall, and the neighboring structures are not visible.

To observe the interior of the bladder wall and the neighboring structures, one can acquire a three-dimensional (3-D) scan of the patient's abdomen. Computed tomography (CT) is the preferred modality of acquisition (see Fig. 1). Other modalities can highlight some useful characteristics, but the images generally exhibit a lower contrast. From the CT scan, a 3-D grid of voxels (i.e. volume elements) is constructed typically with an inslice spacing of 0.5 mm and an interslice spacing of 1–5 mm. To locate potential tumors in the bladder, the clinician reviews the 3-D grid images in the plane of acquisition, or in an orthogonal plane. With this technique, the bladder cavity and neighboring structures are visible in addition to the bladder wall. However, the clinician loses the 3-D visualization and the intuitive navigation the probe offers. Even an experienced radiologist can find it difficult to fully appreciate the 3-D shape of the bladder from the two-dimensional (2-D) images.

II. Related Work

Vining *et al.* [8] segment, and render the surface of the bladder cavity for interactive 3-D visualization from CT scan. In the

Manuscript received December 4, 2001; revised November 1, 2002. This work was supported in part by the Belgian Science Foundation (FRIA and FNRS) and in part by the National Cancer Institute under Grant NIH P41 RR13218, Grant NIH P01 CA67165, Grant NIH R01 RR11747, Grant R01EB000304, and Grant R21-CA80945-01. *Asterisk indicates corresponding author.*

* S. Jaume is with the Telecommunications Laboratory, Université Catholique de Louvain, 1348 Louvain-la-Neuve, Belgium (e-mail: jaume@tele.ucl.ac.be).

M. Ferrant is with General Electric Medical Systems Europe, 78530 Buc, France.

B. Macq is with the Telecommunications Laboratory, Université Catholique de Louvain, 1348 Louvain-la-Neuve, Belgium.

L. Hoyte, A. Schreyer, R. Kikinis, and S. K. Warfield are with the Brigham and Women's Hospital, Harvard Medical School, Boston, MA 02115 USA.

J. R. Fielding is with the University of North Carolina at Chapel Hill, Chapel Hill, NC 27599 USA.

Digital Object Identifier 10.1109/TBME.2003.808828

same way, the clinician manipulates the probe in the conventional cystoscopy, she or he manipulates the computer mouse to navigate inside the bladder cavity. Hence, this scan-based technique is named *virtual cystoscopy* or *CT cystoscopy*. To construct a surface model of the bladder, the CT scan must be segmented. First, the inner boundary of the bladder is contoured in every slice. Then the 3-D surface of the bladder is reconstructed from the pile of segmented slices. Various algorithms can automate the contouring of the bladder cavity since the air regions and the tissue regions appear with very different intensities in the CT scan. The air and tissue classification results in a 3-D image with only two scalar values. An isosurface extraction algorithm, such as the Marching Cubes [9], automatically triangulates the surface between the two values. Finally, the triangulated surface, called a *mesh*, is rendered on the clinician's screen. Using an interactive software (e.g., [10]), the clinician rotates, and cuts the bladder surface, and zooms in the regions of interest.

For 26 patients, Song *et al.* [11] compare conventional cystoscopy, and a combination of transversal CT images and virtual cystoscopy. In their study, 10% of 40 bladder lesions were detected by conventional cystoscopy, but undetected by CT images and virtual cystoscopy. These lesions were all smaller than 5 mm. However the cystoscopic exam is not concerned with detecting every small lesion and irregularity. The purpose of a cystoscopic exam is to judge whether or not the patient will require surgery to remove a tumor. The presence of a single confirmed tumor is enough to ensure the patient will go for surgery.

To enhance the tumors in virtual cystoscopy, Schreyer *et al.* [12] render colored bladder meshes based on the bladder wall thickness. First, they threshold the CT scan to segment the bladder cavity. Second, they edit the contour in every slice to segment the bladder wall. Third, for every voxel on the inner boundary of the bladder wall, they measure the distance to the closest voxel on the outer boundary. This distance locally estimates the bladder wall thickness. Fourth, they triangulate the inner surface with the Marching Cubes algorithm, and map the thickness values onto the mesh vertices. Finally, the clinician can visualize the inner surface, or the outer surface, or both. Colors from red to blue represent thicknesses from 0 to 15 mm.

Fielding *et al.* [13] compare the sensitivity and specificity of three tumor detection methods on 31 patients. The methods are conventional cystoscopy, review of axial CT images, and virtual cystoscopy. They create surface renderings of the bladders based on the thickness representation of Schreyer *et al.*. In a comparison with conventional cystoscopy, they report 83% sensitivity, 36% specificity, 42% positive predictive value, and 71% negative predictive value.

Our contribution is to detect bladder tumors based on a comparison with a normal thickness atlas. During the training, our software needs a set of segmented bladder volumes. For testing a new patient, it inputs a segmented scan of the patient's bladder, and outputs the location of the tumors in the scan, if any. The expert's intervention is limited to the segmentation of the 3-D scan, and the review of the tumor regions the software detects.

The remainder of this paper is structured in four sections. Section III describes the acquisition and segmentation of the image, the construction of the thickness atlas, and its use to detect tumors. Section IV presents our experiments on 26 patients. Section V explains our results and proposes potential improvements. Finally, Section VI concludes the paper.

III. Method

We manually segmented the bladder wall in the CT scan, and then automatically detected the tumors in the bladder wall. Our method was based on five steps: the acquisition of the CT scan, the segmentation of the bladder wall in the scan, the estimation of the thickness of the bladder wall, the construction of a thickness atlas, and the measurement of how much the thickness differs from the atlas. A large difference allowed for the localization of a tumor. These five steps are described below.

A. Image Acquisition

For 26 patients with a history of bladder cancer, we acquired a CT scan of the abdomen. A 12-French Foley catheter was inserted into the bladder, the residual urine drained, and the bladder insufflated with approximately 300 cc room air or to tolerance. The balloon of the catheter was inflated with 5 cc saline. To make urine produced during the scan discernible from the bladder wall, two cubic centimeters diatrizoate meglumine and diatrizoate sodium were instilled through the Foley catheter into the bladder. A 3-D CT image of the bladder was then obtained with the patient in the supine position. The resolution of the 3-D image was $1 \times 1 \times 1.5$ cubic millimeters.

B. Segmentation Method

To segment the bladder wall, we thresholded the bladder scans, and manually outlined the outer boundary of the bladder wall. Note that the contrast in CT scans between air and tissue allows for our automated segmentation of the inner boundary of the bladder wall. However, the contrast between wall tissue and other tissues is much smaller, and their intensities vary over the 3-D CT scan. Therefore, a simple thresholding would not represent the bladder wall correctly. Thus, we used a manual segmentation for the outer boundary of the bladder wall.

We stored the segmentation results in two sets of black (intensity 0) and white (intensity 1) images. In the first set, only the bladder cavity appears in white. In the second set, both the bladder cavity and the bladder wall appear in white. Fig. 2(a) and (b) shows one slice through the first set Fig. 2(a) and the second set Fig. 2(b). We then applied the Marching Cubes algorithm [9] with an intensity of 0.5 on each set of images to create a pair of triangular surface meshes. The mesh extracted from the first set of images models the inner surface of the bladder wall; while the mesh extracted from the second set of images models the outer surface of the bladder wall. A pair of such meshes are rendered in Fig. 2(c) and (d). The left mesh models the inner bladder surface, while the right mesh models the outer bladder surface.

The appearance of the surface is coarse, but its accuracy is below the voxel size. Extracting the isosurface from the CT scan without segmentation would create a smoother surface. However, noise in the CT scan would create many handles, and the surface of the bladder would not be correctly defined.

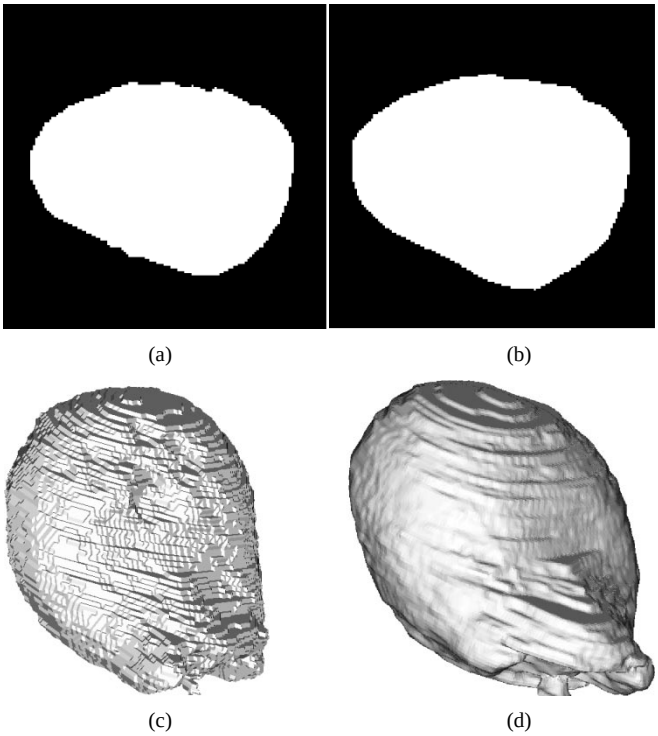

Fig. 2. We model the bladder with two meshes extracted from the segmented volume. (a) One slice of the CT scan after segmentation of the bladder cavity. (b) Same image where the white area includes the bladder wall. (c) Mesh of the inner surface of the bladder. (d) Mesh of the outer surface of the bladder. Both meshes are extracted using the Marching Cubes algorithm on the segmented volume.

C. Thickness Estimation

To estimate the bladder wall thickness, we computed for every vertex of the inner mesh its distance to the closest vertex of the outer mesh. The mesh discretization resulted in mesh vertices being separated by distances similar to the bladder wall thickness. In order to form an accurate estimate of the bladder wall thickness, the triangle meshes were upsampled by a factor of eight by subdividing every triangle of the outer mesh into 64 coplanar triangles. Fig. 3 shows the thickness map for the bladder of two patients. Thin regions of the bladder wall are color-coded in blue, while thick regions are color-coded in red.

D. Atlas Construction

Bladder tumors are characterized by a thickening of the bladder wall. However, the thickness distribution varies over the bladder surface. Our idea was to compare the thick regions with a normal thickness map (*atlas*) to distinguish pathological thickness from normal thickness.

To build the atlas, we used a leave one out cross validation. We left one patient out of our set of 26 patients, and built the atlas for the other 25 patients. The first patient is called the *testing* patient, while the other patients are called the *training* patients. For every testing patient, we computed one atlas with the average, one atlas with the standard deviation, and one atlas with the coefficient of thickness variation. The coefficient of variation is the ratio of the standard deviation over the average

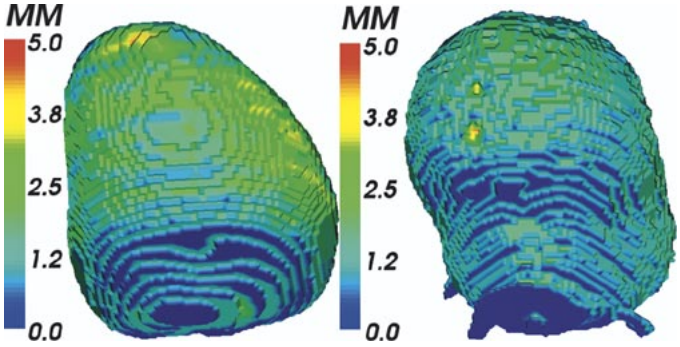

Fig. 3. Distance from inner mesh vertices to the closest vertex on the outer mesh gives a thickness map of the bladder wall. For the two patients shown, the thickness ranges from 0 to 5 mm. We measure a 0-mm thickness when the wall is not visible in the CT scan. Note a thick spot at the top of the first bladder, and two thick spots at the center of the second bladder.

of the thickness; it indicates the relative variability of our thickness estimation.

To build one normal thickness atlas, we averaged the thickness values across the set of 25 training patients. For every patient in the set, we mapped a grid of 600 nodes by 600 nodes on the top half and on the bottom half of the bladder surface. To map these two grids, we used the remeshing algorithm of [14]. We linearly interpolated the thickness values of the bladder mesh to associate values with the new grid nodes. Finally, for every node of the grid, we averaged the thickness across the training patients. The thickness average created the normal thickness atlas for the set of training patients.

The aforementioned procedure would give a correct atlas if no patient in the training set had tumors. Since this was not the case, the tumors would bias our atlas, and decrease the sensitivity of our tumor detection method. To reduce the bias, when one vertex of a training mesh had a thickness value larger than 5 mm, we discarded that value for the atlas computation. We chose a 5-mm threshold since we were interested in small tumors, clinically defined as tumors smaller than 5 mm. Note that the other thickness values for the same training mesh were still inspected. Fig. 4 shows the average atlases for two patients. The values of every atlas are colored on the mesh of the testing patient. We observe that the thickness is 1-mm larger at the top than at the bottom of the bladder.

To build the standard deviation atlas, we again considered every node of the grid for every patient in the set of training patients. This time, we used both the thickness value, and the previously computed average to compute the standard deviation. To build the coefficent of variation atlas, we only considered the nodes of the standard deviation atlas, and the nodes of the average atlas. For every node, we computed the ratio of the standard deviation, and the average. Fig. 5 shows the variation atlases for the two patients selected in Fig. 4. The first row shows the atlases for the standard deviation of the thickness. Note that our thickness estimation has a slightly larger variation at the bottom than at the top of the bladder. The last row shows the atlases for the coefficient of variation. The images confirm that the thickness estimation varies more at the bottom than at the top of the bladder.

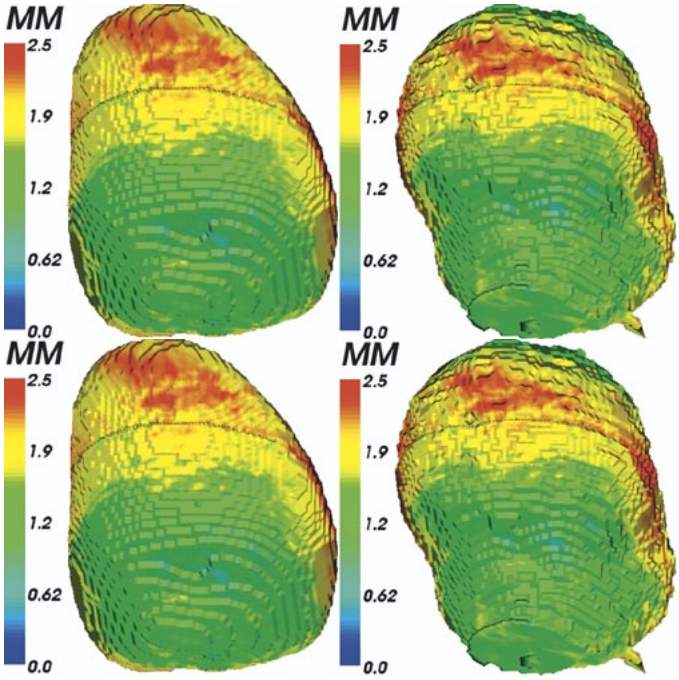

Fig. 4. Average thickness atlases show the distribution of the wall thickness over the bladder surface. The figure shows two atlases computed from a different set of bladder scans. Both atlases indicate that the thickness is 1 mm larger in average at the top of the bladder than at the bottom.

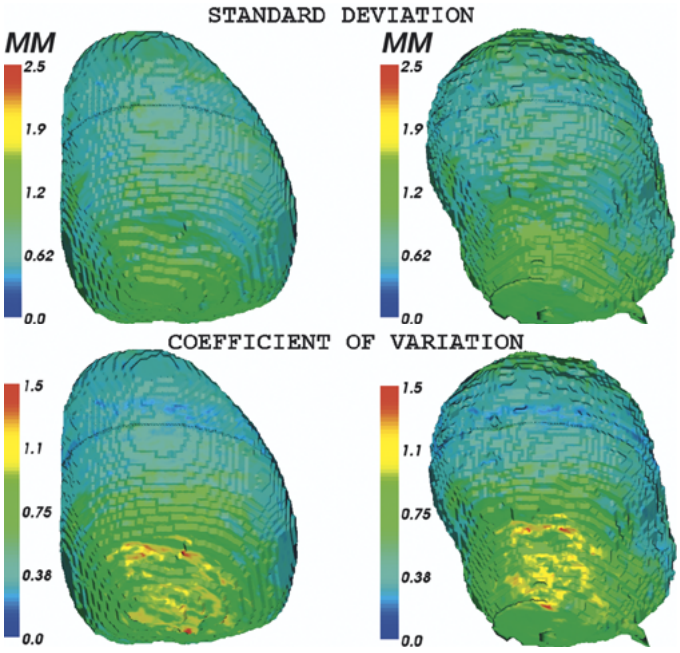

Fig. 5. Variation of the atlas thickness for the sets of patients used in Fig. 4. The first row shows the standard deviation, and the second row the coefficient of variation, i.e., the ratio of the standard deviation over the average. The variation in the atlas appears slightly higher at the bottom than at the top (the standard deviation is 0.6 mm larger).

E. Measurement of Abnormality

We used a Z score to compare the bladder wall thickness for the testing patient to the normal thickness in the atlas. The Z score for a thickness value t is the ratio of its difference to the average μ over the standard devation σ, as shown in (1). This ratio without units characterizes the abnormality of the thickness t. We wanted to find the optimal Z-score threshold for bladder tumor detection. Regions where the Z score was below the threshold would not be detected as diseased; regions above the threshold would be detected as diseased. We evaluated several Z score thresholds for some arbitrarily selected cases, and found that $Z = 2.5$ allowed for a good detection of tumors for those cases

$$Z(t) = \frac{t - \mu}{\sigma}. \tag{1}$$

Positive Z scores characterize thickness larger than the average, while negative Z scores characterize thickness smaller than the average. Since tumors create a thickening of the bladder wall, we only considered the positive Z scores. The larger the Z score, the more atypically thick the bladder wall. Fig. 6 shows the Z score computed for two testing patients. Most of the bladder area has a Z score smaller than one (*blue in the figure*). The thickness in this area is smaller than or close to the average thickness. The top of the first bladder (*top left image*) exhibits a region with a large Z score (*red in the figure*). Since the Z score exceeded 2.5, our software detected this region as diseased. We observe two orange-red spots in the second bladder (*bottom left image*). Again, our software detected these regions as diseased since they exceed the 2.5-Z score threshold. The right column shows one slice through the CT scan for each patient. For every region our software detected as diseased, an arrow points to the region in both the rendering and the CT slice.

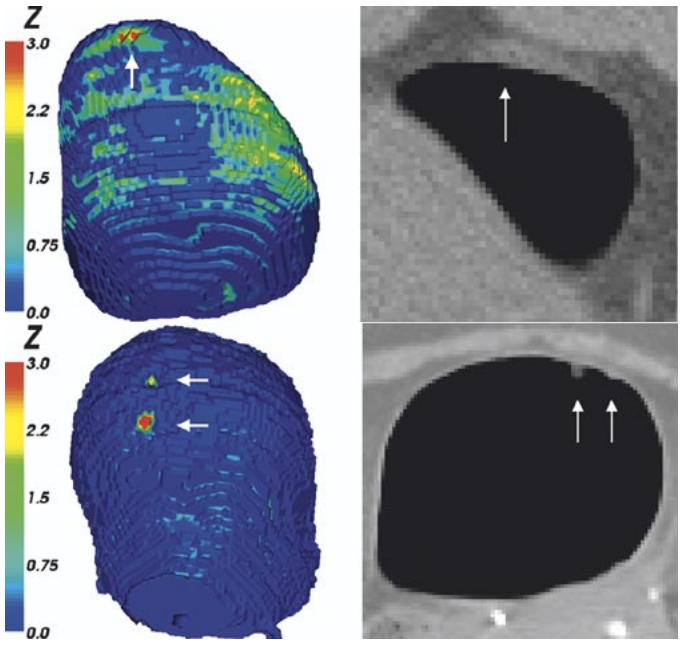

Fig. 6. Relative thickness difference between the patient bladder and the atlas, measured by the Z score, indicates the region of abnormal thickness. (*top left*) Z score for one patient. (*top right*) Zoom in the CT scan for this patient. Similarly, the bottom row shows Z score and CT scan for a second patient. At the top of the first bladder, the Z score reveals a region of abnormal thickness (red). At the center of the second bladder, the Z score reveals two spots of abnormal thickness (red).

IV. Results

To validate our method, we performed a leave one out cross validation. We successively picked one bladder scan out of our set; this bladder scan was our testing data. Our software built an atlas based on the other bladder scans; these bladder scans were our training data. Our software detected the tumors on the testing data. Initial readings made by a radiologist experienced in radiology of the pelvis served as the ground truth. We, thus, compared the outputs of our software with the expert read CT scans. The expert reviewed every bladder scan, and detected the tumors in the scan. Since we had 26 bladder scans in our set, we performed this experiment 26 times with a training set of 25 scans, and a testing set of one scan.

To compare the location of the tumors detected by the software, with the location of the tumors detected by the expert, we separated the bladder wall in six anatomical zones: the top zone, the bottom zone, the anterior zone, the posterior zone, the left zone, and the right zone. For every zone, and for every bladder scan, we compared what the software detected, with what the expert detected. If they both detected a tumor, we counted one true positive. Similarly, if they both detected no tumor, we counted one true negative. If the software detected a tumor, and the expert did not, we counted a false positive (FP). Conversely, if the software detected no tumor, and the expert detected one, we counted a false negative (FN).

Note that we counted the number of diseased zones, not the number of tumors. The detection of several tumors in the same zone gave the same output as the detection of only one tumor in this zone. As an automated screening exam, our software is not intended to accurately locate every tumor, but to detect whether the bladder is diseased or not, and to approximately locate the potential tumors. The detection of one tumor in a zone is enough for sending the patient for further inspection. The software skims out the healthy cases, and only sends the suspicious cases to the clinician. The load of work for the clinician can be reduced, or the number of patients can be increased.

We summed the number of TP occurences together. Similarly, we summed the TN occurences, the FP occurences, and the FN occurences. Our experiments with 26 patients gave $\mathrm{TP} = 16$, $\mathrm{FP} = 17$, $\mathrm{FN} = 2$, and $\mathrm{TN} = 123$. We report these figures in Table I. The meaning of these numbers is the following. Our software correctly detected 16 diseased zones; it incorrectly detected 17 other zones as diseased; it missed two diseased zones; it correctly classified 123 zones as healthy. Note that $\mathrm{TP} + \mathrm{FP} + \mathrm{TN} = 156$, since our data set contained 26 patients, and every patients had six bladder zones. Based on the occurences in Table I, we computed the medical statistics formulated in (2)–(5)

$$\text{Sensitivity} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}} = 89\% \qquad (2)$$

$$\text{Specificity} = \frac{\mathrm{TN}}{\mathrm{TN}+\mathrm{FP}} = 88\% \qquad (3)$$

$$\text{Positive Predictive Value} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FP}} = 48\% \qquad (4)$$

$$\text{Negative Predictive Value} = \frac{\mathrm{TN}}{\mathrm{TN}+\mathrm{FN}} = 98\%. \qquad (5)$$

TABLE I
Comparison of Tumor Detection by an Expert Observer and by our Automated Tumor Detection

	expert detects a tumor	expert detects no tumor
software detects a tumor	TP = 16	FP = 17
software detects no tumor	FN = 2	TN = 123

For every zone and for every bladder, the expert and the software classify the zone as either diseased or healthy. TP is the number of zones detected as diseased by both the expert and the software. FP is the number of zones detected as healthy by the expert and as diseased by the software. FN is the number of zones detected as diseased by the expert and as healthy by the software. TN is the number of zones detected as healthy by both the expert and the software.

Ideally, a detection software should have a 100% score for all these statistics. A 100% sensitivity would indicate that it detects all diseased zones. A 100% specificity would indicate that it never detects a tumor in a healthy zone. A 100% positive predictive value would indicate that, when the software detects a diseased zone, the zone is actually diseased. A 100% negative predictive value would indicate that, when the software detects no diseased zone for a patient, the patient is actually healthy.

V. Discussion

In this section, we comment the results of our validation, we discuss the choice of the ground truth, we compare our method with standard cystoscopy, and we propose applications to other medical problems.

A. Results of the Validation

Table I reports that of the 18 diseased zones detected by the expert our algorithm detected 16 zones and missed the other two. Of the 140 healthy zones, our software correctly classified 123 zones as healthy, and only misclassified 17 as diseased. The resulting high sensitivity (89%) and specificity (88%) indicate that our software could be incorporated in a screening procedure.

The low positive predictive value (48%) means that the clinician needs to review as many healthy patients as actual diseased patients. When our software detected a tumor, the probability the patient actually had a tumor was roughly as high as the probability the patient was healthy. A method with a higher positive predictive value (PPV) would further reduce the work load of the clinician.

The very high negative predictive value (98%) allows the patient to trust our software when it does not detect a tumor. Negative predictive value is the important clinical figure. Virtual CT cystoscopy, when negative, implies that no tumor is present. The patient does not need a standard cystoscopy. It may be useful for screening the patient population with hematuria.

Compared with the method of Fielding *et al.* [13], our method improved the sensitivity from 83% to 89%. This improvement is too small to change the clinical outcome. Besides, they chose conventional cystoscopy as the ground truth, while we chose expert read CT scans as our ground truth. New validation experiments with the same ground truth are needed to accurately compare both methods.

B. Choice of the Ground Truth

An expert performing conventional cystoscopy served as the ground truth in the experiments of [13], while our validation experiments used an expert reading the CT scans. In both cases, the tumor detection is biased toward the tumors detected by the expert. The expert may misclassify a zone using either the scope, or the CT scans. These errors will affect the validation. Since our validation and the validation of Fielding *et al.* use a different ground truth, both methods are not comparable. Nevertheless, our goal is different as we propose a method with an automated reading of the renderings.

Comparison with expert clinical reading is required to validate a tumor detection software before its clinical use. However, such a validation cannot reveal that the software could do better than the expert. The advantage of the software is that it eliminates the variability on the experience and the fatigue of the expert.

C. Comparison with Standard Cystoscopy

Compared with standard cystoscopy, virtual cystoscopy offers a better visibility of the bladder neck. In a standard cystoscopic exam, the clinician may miss tumors in the bladder neck due to the limited flexibility of the scope near the entrance. Combined with high-resolution scans, we expect that our virtual cystoscopy approach will be more accurate in the bladder neck than conventional cystoscopy.

We used interactive segmentation to locate the bladder wall surfaces in our experiments. Interactively segmenting the bladder wall in every slice is time consuming. The widespread clinical implementation of virtual cystoscopy screening will require automated segmentation of the bladder wall. We believe future CT technology will offer higher resolution and better quality images. This will make the automated and accurate segmentation of the bladder wall easier. Combining our method with an automated segmentation software would provide for an entirely automated screening of bladder scans.

D. Potential Applications

We believe our method could be adapted to address problems with other organs, such as the brain. Our method could be advantageous for assessing changes in cortical thickness. The comparison of the cortex of a patient with an atlas of typical cortical thickness could aid in the detection and characterization of neurodegenerative disorders. As reported by several authors [15]–[17], the main requirement for constructing cortical thickness estimates is to accurately segment the cortical surface, which is difficult due to the complex folded structure of the brain and limited spatial resolution of conventional magnetic resonance imaging. The application of our technique would require an accurate registration method capable of aligning corresponding cortical features between subjects.

VI. Conclusion

Our tumor detection method aligns a bladder wall thickness atlas, containing the local mean wall thickness and standard deviation, to a subject. It then identifies tumors as the sections of the patient bladder wall that exceed the mean thickness by more than 2.5 standard deviations. Compared with expert-read CT scans, our tumor detection exhibits 89% sensitivity, 88% specificity, 48% PPV, and 98% negative predictive value. The high negative predictive value makes our software suitable for bladder cancer screening. Our software accurately detects healthy regions, allowing the clinician to focus on the suspicious bladder scans. It evaluates the bladder neck, a significant advantage over standard cystoscopy. The limitation of our current implementation is the interactive segmentation, which is time consuming and susceptible to error. We expect that the development of an automated bladder wall segmentation algorithm would remove this limitation.

Acknowledgment

The authors would like to thank S. Timoner for helpful comments and discussions during the preparation of this paper, S. Haker for explanations on curvature measurements, A. Bharatha for an introduction to medical statistics, I. Friedel for his help with the remeshing code, and the reviewers for suggesting many improvements to this paper.

References

[1] S. Parker, T. Tong, and S. Bolden *et al.*, "Cancer statistics," *CA Cancer J. Clinic.*, vol. 46, no. 5, p. 27, 1996.

[2] M. Cookson, H. Herr, and Z. Zhang *et al.*, "The treated natural history of high risk superficial bladder cancer: 15-year outcome," *J. Urol.*, vol. 158, no. 62, 1997.

[3] R. Hurle, A. Losa, and A. Manzetti *et al.*, "Upper urinary tract tumors developing after treatment of superficial bladder cancer: 7-year follow-up of 591 consecutive patients," *J. Urology*, vol. 53, no. 1144, 1999.

[4] M. Raitanen, M. Leppilahti, K. Tuhkanen, T. Forssel, P. Nylund, and T. Tammela, "Routine follow-up cystoscopy in detection of recurrence in patients being monitored for bladder cancer," *Ann. Chirurgiae et Gynaecol.*, vol. 90, no. 4, pp. 261–265, 2001.

[5] D. Dryhurst and C. Fowler, "A new small-calibre diagnostic flexible cystoscope," *BJU Int.*, vol. 89, no. 3, pp. 194–196, Feb. 2002.

[6] D. Burke, D. Shackley, and P. O'Reilly, "The community-based morbidity of flexible cystoscopy," *BJU Int.*, vol. 89, no. 4, pp. 347–349, Mar. 2002.

[7] T. Bernhardt and U. Rapp-Bernhardt, "Virtual cystoscopy of the bladder based on CT and MRI data," *Abdom. Imag.*, vol. 26, no. 3, pp. 325–332, May/June 2001.

[8] D. J. Vining, R. J. Zagoria, K. Liu, and D. Stelts, "CT cystoscopy: An innovation in bladder imaging," *AJR*, vol. 166, pp. 409–410, 1996.

[9] W. Lorensen and H. Cline, "Marching cubes: a high resolution 3-D surface construction algorithm," in *Proc. Computer Graphics SIGGRAPH*, Anaheim, CA, July 1987, pp. 163–169.

[10] D. Nain, S. Haker, R. Kikinis, and E. Grimson, "An interactive virtual endoscopy tool," presented at the *IMIVA, Workshop Medical Image Computing and Computer-Assisted Intervention*, Utrecht, The Netherlands, Oct. 2001.

[11] J. Song, I. Francis, J. Platt, R. Cohan, J. Mohsin, S. Kielb, M. Korobkin, and J. Montie, "Bladder tumor detection at virtual cystoscopy," *Radiol.*, vol. 218, no. 1, pp. 95–100, Jan. 2001.

[12] A. Schreyer, J. R. Fielding, S. K. Warfield, J. H. Lee, K. Loughlin, H. Dumanli, F. Jolesz, and R. Kikinis, "Virtual CT cystoscopy: color mapping of bladder wall thickness," *Investigat. Radiol.*, vol. 35, no. 5, pp. 331–334, May 2000.

[13] J. R. Fielding, L. Hoyte, S. Okon, A. Schreyer, J. Lee, K. H. Zou, S. K. Warfield, J. P. Richie, K. R. Loughlin, M. P. O'Leary, C. J. Doyle, and R. Kikinis, "Tumor detection by virtual cystoscopy with color mapping of bladder wall thickness," *J. Urol.*, vol. 167, pp. 559–562, 2002.
[14] P. Alliez, M. Meyer, and M. Desbrun, "Interactive geometry remeshing," in *Proc. Computer Graphics*, San Antonio, TX, July 21–26, 2002, pp. 347–354.
[15] B. Fischl and A. Dale, "Measuring the thickness of the human cerebral cortex from magnetic resonance images," in *Proc. Nat. Acad. Sci.*, vol. 97, Sept. 2000, pp. 11 050–11 055.
[16] N. Kabani, G. Le Goualher, D. MacDonald, and A. C. Evans, "Measurement of cortical thickness using an automated 3-D algorithm: a validation study," *NeuroImage*, vol. 13, no. 2, pp. 375–380, Feb. 2001.
[17] A. Yezzi and J. Prince, "A PDE approach for measuring tissue thickness," presented at the *IEEE Computer Society Conf. Computer Vision Pattern Recognition*, Kauai, HI, Dec. 2001.

Sylvain Jaume received the M.Sc. degree in electrical engineering from the Université Catholique de Louvain, Louvain-la-Neuve, Belgium, in 1999. He is currently working towards the Ph.D. degree with scholarships from the Belgian Science Foundation (FRIA and FNRS) at the Université Catholique de Louvain, with research concerning surface modeling from tridimensional scans and anatomical shape analysis.

From March until August 2001, he was a Visiting Student at the Brigham and Women's Hospital and Harvard Medical School, Boston, MA. From October 2001 until September 2002, he was a Visiting Student at the California Institute of Technology, Pasadena, in the Computer Graphics group of Prof. Schröder and Prof. Barr.

Matthieu Ferrant was born in Islington, U.K., in 1974. He received the M.Sc. and Ph.D. degrees in electrical engineering from the Université Catholique de Louvain, Louvain-la-Neuve, Belgium, in 1997 and 2001, respectively.

He was a Visiting Student at the Institute for Perception Research (IPO), the Eindhoven University of Technology, Eindhoven, The Netherlands, within the Erasmus Exchange Program from February until May 1997. From September 1998 until June 1999, he was a Visiting Student at the Massachusetts Institute of Technology, Cambridge, with Prof. F. Dewey within the International Consortium for Medical Imaging Technology. During that period, he was also with the Surgical Planning Laboratory at the Brigham and Women's Hospital, Harvard Medical School, Boston MA, working on medical applications with Prof. Warfield. His main research interests concern deformable surface matching, rigid and nonrigid registration, finite-element modeling, and multiresolution surface representations. He now works for General Electric Medical Systems, Buc, France.

Benoît Macq (S'83–M'84–SM'01) was born in 1961. He received the M.Sc. and the Ph.D. degrees in electrical engineering from the Université Catholique de Louvain, Louvain-la-Neuve, Belgium, in 1984 and 1989, respectively. He performed his Ph.D. degree work on perceptual coding for digital TV under the supervision of Prof. P. Delogne.

Since 1996, he has been a Professor with the Telecommunications Laboratory, the Université Catholique de Louvain. From 1992 to 1996, he was a Senior Researcher with the Belgian Science Foundation (FNRS) at the Université Catholique de Louvain, and an Invited Assistant Professor in the Telecommunications Laboratory. In 1990 and 1991, he worked on network planning for Tractebel S.A., Brussels, Belgium. He was a Visiting Scientist at Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, in 1995, and at the Massachusetts Institute of Technology, Cambridge, during the summer of 1999. He has been a Visiting Professor at the Ecole Nationale Supérieure des Télécommunications, ENST-Paris, France, and at the Université de Nice, Sophia-Antipolis, France, from 1999 to 2000. His main research interests are image compression, image watermarking, and image analysis for medical and immersive communications.

Dr. Macq served as a Guest Editor for the PROCEEDINGS OF THE IEEE, as well as for the *Signal Processing Journal*, and is a Member of the Program Committee of several IEEE and SPIE conferences.

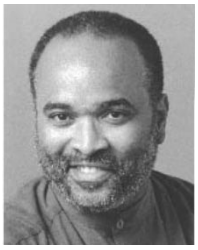

Lennox Hoyte received the M.S. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, and worked in computer hardware design and design automation for several years. He later received the M.D. degree from Stanford University, Stanford, CA, in 1995, and completed his residency training in obstetrics (OB) and gynecology (GYN) at the Brigham and Women's/Massachussetts General Hospital Joint Residency Program, Boston, MA.

Since 1998, he has been with the Surgical Planning Laboratory, Harvard Medical School, Cambridge, MA, where he is currently an Assistant Professor in the Departments of Radiology and OB/Gyn.

Dr. Hoyte is currently a Board Certified Obstetrician/Gynecologist, and a Fellow of the American College of Obstetrics and Gynecology. He is an attending Obstetrician/Gynecologist at the Brigham and Women's Hospital.

Julia R. Fielding received the B.S. degree in chemistry from the University of Michigan, Ann Arbor, and the M.D. degree from the University of Pittsburgh School of Medicine, Pittsburgh, PA, in 1986.

She is currently Associate Professor of Radiology and Chief of Abdominal Imaging at the University of North Carolina, Chapel Hill. After an internship in surgery performed at the Western Pennsylvania Hospital, Pittsburgh, she completed a four-year residency in diagnostic radiology at Boston University, Boston, MA. While a Fellow in abdominal imaging at Brigham and Women's Hospital, Boston, MA, she became interested in genitourinary radiology, including bladder cancer, renal stones, and gynecologic cancers. Collaborating with members of the Surgical Planning Laboratory, she published several scientific papers on the uses of 3-D imaging techniques over the last eight years. She is currently working on developing the use of spectroscopy for identification of extracapsular disease in patients with prostate cancer.

Andreas Schreyer received the M.D. degree from the University of Erlangen/Nuernberg and Regensburg, Germany. In his thesis at the Department of Radiology, University of Regensburg, he examined the limitations and possibilities of maximum intensity projection (MIP) with spiral CT angiography using phantoms.

After graduation, he was trained in abdominal surgery at the Neuperlach/Munich Hospital (Ludwig-Maximillian-University, Munich, Germany). As a MRI Research Fellow at the Brigham and Women's Hospital and Harvard Medical School, Boston, MA, for two years, he got involved in several advanced medical 3-D applications like volumetric measurement of the prostate, virtual endoscopy, and 3-D modeling of the chest for simulated cardiac defibrillation. He is now finishing his residency program in radiology at the University Hospital of Regensburg, Germany. His particular interest is 3-D image processing and virtual endoscopy. Additionally, he focuses on advanced image processing applications in patients with inflammatory bowel disease.

Ron Kikinis received the M.D. degree from the University of Zurich, Zurich, Switzerland, in 1982.

He has worked as a Researcher at the ETH, Zurich, Switzerland, and as a Resident at the University Hospital, Zurich, Switzerland. He is currently the Director of the Surgical Planning Laboratory of the Department of Radiology, Brigham and Women's Hospital and Harvard Medical School, Boston, MA, as well as an Adjoint Professor of biomedical engineering at Boston University, Boston, MA. His research interests include the development of clinical applications for image processing, computer vision, and interactive rendering methods. He is currently concentrating on developing fully automated segmentation methods and introducing computer graphics into the operating room.

Simon K. Warfield (M'99–A'01) received the B.Sc. degree in computer science, the B.E. degree in electrical engineering with first class honors, and the Ph.D. degree in computer science and engineering from the University of New South Wales, in Sydney, Australia, in 1991, 1993, and 1997, respectively.

He is currently an Assistant Professor of Radiology at Harvard Medical School and Director of the Computational Radiology Laboratory at Brigham and Women's Hospital, Boston, MA. Since joining Brigham and Women's Hospital in 1996, his research has focused upon methods for quantitative image analysis through novel segmentation and registration approaches, and in real-time image analysis, enabled by high-performance computing technology, in support of surgery.