Image-Based Motion Detection

Using the Concept of Weighted Directional Descriptors

BY EYAL ZADICARIO, SHLOMI RUDICH, GHASSAN HAMARNEH, AND DANIEL COHEN-OR

The use of image guidance in medical applications is constantly growing because of its tremendous impact on the future of health care. Although image-based tissue tracking has been thoroughly explored in the academic literature for years, it has not yet matured to become widely accepted by clinicians. Undetected tissue movements in image-based clinical procedures may cause safety and efficacy difficulties. We introduce an image-based approach for detecting tissue movements during clinical procedures. Our method has been validated in more than 600 true clinical cases. The results show that our algorithm agrees with an expert analysis in 98% of the cases, showing zero events of false alarms and zero events of undetected motion. The results show that the approach provides a clinically ready motion-detection algorithm. These robust results are achieved by introducing the concept of weighted directional descriptors (WDDs). The technique analyzes the directivity and confidence level of each anatomical feature and uses it to weight local inputs resulting in a robust motion vector. The robustness is further increased by a novel preprocess that screens out features that may be misleading or are repeated in the adjacent search zone.

The technique meets the requirements, as defined by our clinicians, and is now integrated in true medical systems. In particular, our approach has been uniquely developed and integrated into a clinical product. ExAblate is the first Food and Drug Administration (FDA)-approved magnetic resonance (MR)-guided noninvasive surgical device using focused ultrasound therapy. It is used in commercial clinics and in leading medical academic research institutions, attesting to the success of our method and its practical clinical value.

Image-Guided Therapy

In recent years, medical imaging is moving from being an off-line diagnostic tool to an intraoperative modality, which together with noninvasive therapy procedures provide a revolutionary concept of image-guided therapy systems [1], [2]. A typical image-guided treatment starts by acquiring a set of pre-operative images used to analyze the anatomy and plan the treatment [3]. However, during the treatment, the anatomy may change shape and shift from its original placement causing the planning images and the treatment plan to become invalid. This is a concern to both safety and efficacy of a treatment that depends heavily on image guidance. Therefore, there is a need for a tool that will detect motion events and issue a warning so that the proper actions can be taken by the physician. A critical success factor of any motion detection technique lies in its reliability, mainly having zero events of undetected motion. In addition, it needs to be robust by minimizing the events of false alarm where no motion actually occurred. Finally, to be integrated in a clinical environment, the method needs to be fast, requires very little interaction, and be robust to image quality and user inputs. Our approach is based on user input to define the area to be monitored. The target region is analyzed to identify key features to be tracked for displacements based on updated images of the same anatomy taken during treatment.

In this work, we develop a method that meets these requirements for adoption in a clinical setting. Our proposed method was developed as part of a clinical product ExAblate. It is the first and only FDA-approved medical device for using focused ultrasound under MR guidance and control. It is commercially available in the United States, Europe, and parts of Asia with more than 5,000 treatments performed worldwide. It is installed in more than 70 clinics and hospitals and in more than 30 top academic centers researching additional potential clinical applications that can benefit from this technology.

Although our approach can be generalized to other clinical applications, in this work, we focus on a medical device using ultrasound to heat the target tissue to a point of irreversible coagulation. Since ultrasound energy can traverse through soft tissue, this can be done completely noninvasively. However, to make it safe and controlled, the energy needs to be guided to the target, and the thermal rise needs to be measured and controlled in real time. To achieve this, the system is integrated inside an MR scanner. Initially, MR images are acquired to describe the anatomy and plan the treatment. During the treatment, the target is continuously scanned to measure the thermal effect and to confirm there was no movement of the anatomy. This is a concern of clinicians that our approach was challenged with. Movement of the target may result in a situation where the treatment is

Digital Object Identifier 10.1109/MEMB.2009.935727

being delivered to nontargeted tissue or that some target tissue remains untreated. Both are undesired and are the key motivation for implementing a motion-detection mechanism in the treatment procedure.

Our approach is not a generic image analysis technique but rather a tailored approach to the clinical application we were challenged with. The device is currently being evaluated and used for treatment of the targets in uterus, bones, brain, breast, and prostate. These are quasi-static organs in which motion is not expected during the treatment.

Background and Related Work

Image-guided therapy of soft tissue organs has been widely explored in recent years in an effort to use technological development in computer-assisted interventions toward improving patients' health care [5], [7]–[9]. Vast research has evolved in elastic deformable models for image-based organ deformation measurements, e.g., [10]–[12]. Despite these efforts, very little of this has matured to be accepted in routine clinical applications. The technological challenge is great since the clinical environment demands a robust solution with little tolerance to errors. Tracking of soft tissue organs presents a special challenge since they undergo nonrigid deformations. Various models for segmentation and tracking of deformable anatomy have been proposed using parametric and geometric active contours [13], [14]. However, the results were insufficiently robust in situations of objects that have poor contrast [15].

Motion detection was also addressed by attempting registration techniques. The registration of medical images is in general a difficult problem with numerous proposed methods [16]. Linear transformations do not have enough degrees of freedom to track the deformable anatomy accurately, and dense deformation fields have too many degrees of freedom requiring long computation times. Most techniques that estimate dense deformation fields require the setting of the amount of spatial regularization to produce smooth results and counteract image degradation. However, setting the amount of regularization is a nontrivial task. Further, the large degrees of freedom renders the method more susceptible to entrapment into local minima of the similarity criteria [22]–[25]. To address these difficulties, various researchers suggested matching a sparse set of points as the basis of image registration or shape matching, e.g., [17] and [18]. The point feature is the simplest form of a feature that can be extended to curves or surfaces. Variations of this concept include elastic registrations based on thin-plate splines, in which a set of corresponding anatomical-point landmarks drive a minimum bending energy warp [17], [19]. However, point-based matching has several drawbacks. The presence of noise and the existences of outliers make it impractical for routine situations [26]. Finally, the geometric transformations may need to incorporate high-dimensional nonrigid mappings to account for deformations of the point sets [17].

A successful and highly adopted approach for extracting distinctive invariant features that can be used for image matching [using scale invariant features transform (SIFT)] was suggested by Lowe (2004) [20], [27], [28]. It has been shown that SIFT descriptors can be used to achieve a fairly robust object detection in still 2-D scalar images. SIFT has been used in medical imaging [4], extended to scalar images of arbitrary dimensionality, and applied to determine point correspondence between pairs of medical images in 3-D [28]. However, the approach has several drawbacks. It ignores key attributes of anatomical structures and may be easily misled by duplicated features within the area of interest. A comprehensive robustness comparison was performed between motion detection based on SIFT and our proposed approach.

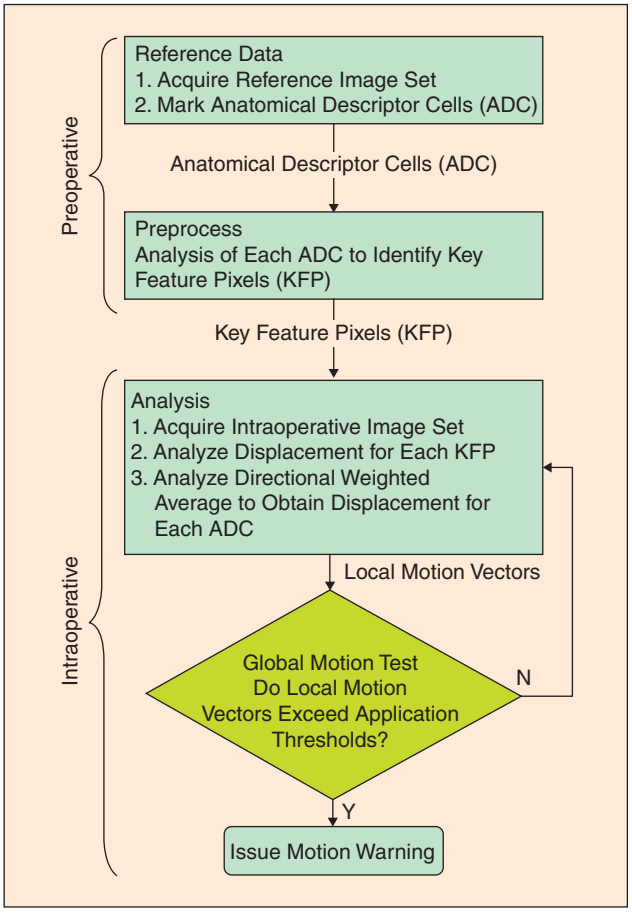

Fig. 1. An overview of algorithm flow method.

Methods

Overview

The proposed method makes use of information available during the existing treatment flow. Planning images are acquired

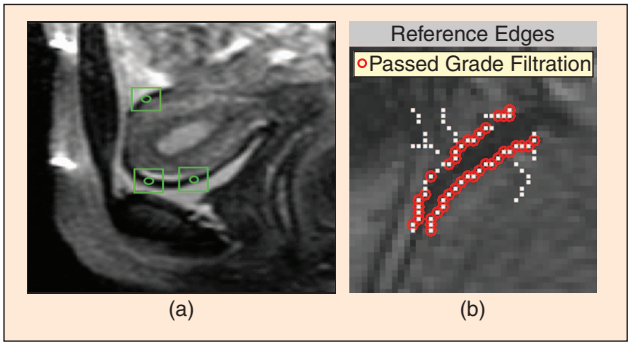

Fig. 2. Motion-detection components. (a) ADC marked by the user (circles) and the patch in the image (square) to be considered in the analysis of the anatomical descriptor. (b) KFP: an example of one patch from (a). Pixels with high-gradient value (white) have been screened to select pixels having autocorrelation above threshold (red).

ExAblate is the first FDA-approved MR-guided noninvasive surgical device using focused ultrasound therapy.

at the beginning of the treatment (e.g., MR T2-weighted image sequences) covering the entire volume of the treatment target. At planned intervals (e.g., every 3 s), updated images of some portion of the target are being rescanned. These data are processed by the algorithm, and in case motion has been estimated to exceed predefined thresholds, a motion warning is issued to the physician.

Our algorithm is separated into four steps (Figure 1). Initially, a baseline reference image set is acquired to include the entire volume of interest. In the second step, the volumetric data are preprocessed to identify landmark features in the anatomy, anatomical descriptor cells (ADCs), and to select and classify the most robust key feature pixels (KFPs) within each cell. The third step is applied during the treatment. The anatomy is periodically rescanned. In the analysis phase, the displacement for each KFP is estimated, contributing to the local motion vector of each cell. Finally, in step four, the motion vectors are compared with predefined criteria, and a motion warning is issued if needed.

Preprocessing the Reference Image Set

A robust motion detection technique needs to distinguish between tissue displacements that are inevitable or have no effect on the planned procedure (e.g., blood vessels) and ones that need to raise a clinical concern. This requires a clinical understanding of the anatomy and procedure. In our approach, gestures are given by the clinical user marking the areas of clinical interest. The automated analysis is limited to cells surrounding each key feature identified by the user, labeled as {Ci}. Figure 2 shows sample cells (green bounding boxes) that have been placed surrounding inputs given by the clinical user. In each ADC, we need to screen the KFPs that may provide reliable evidence for local displacement.

First, we apply an edge detector [6] to identify pixels that have significant local gradient to become candidate KFPs. Figure 2(b) shows the selected pixels for the top cell in Figure 2(a). The following screening steps are based on an autocorrelation matching score (ACMS) to quantify the similarity between two image patches (P_1, P_2) as defined in (1). Repeating this for each pixel results in a 2-D autocorrelation map, as can be seen in Figure 3. The goal of the second step is to distinguish between true features and noisy pixels. We compute an integral autocorrelation value for pixel P by normalized sum of pixels in its autocorrelation map, as defined in (2). A significant anatomical feature has an integral autocorrelation value that is above a threshold whereas a noise pixel does not. Figure 3(a) shows an example of a feature that has low-integral value and (c) a robust feature where the integral value is above the threshold. The third screening method

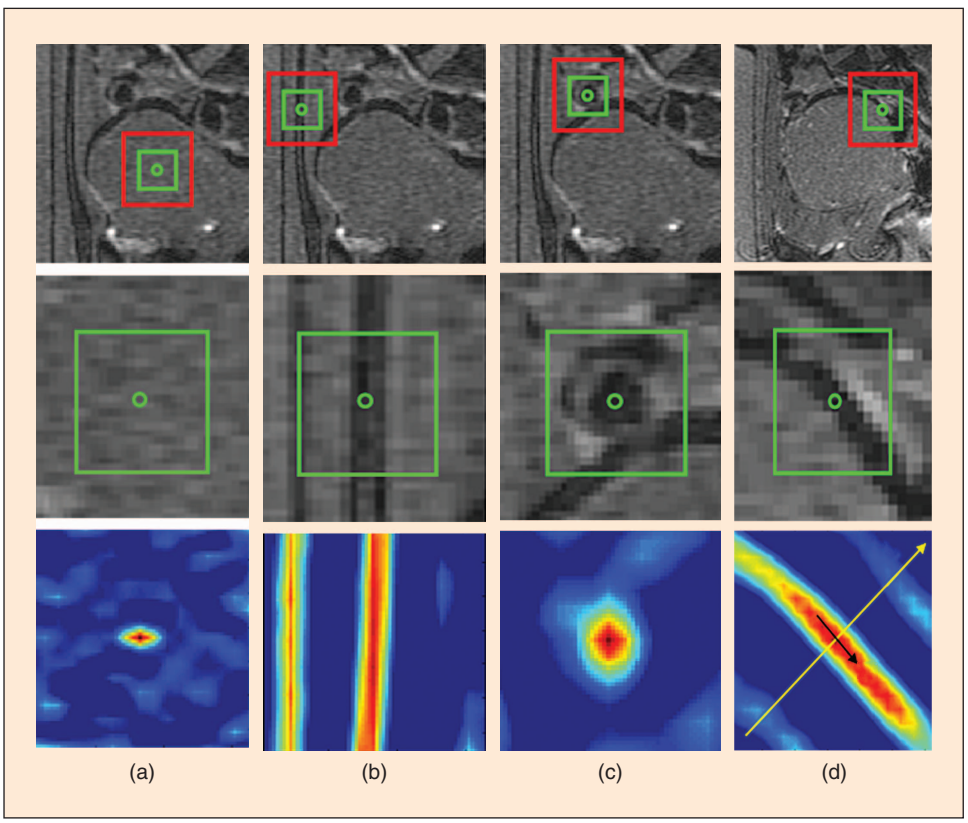

Fig. 3. Autocorrelation sample analysis cases. (a) Poor correlation: pixel with correlation value below the required threshold (CCM = 0.21). (b) Duplicate correlation: Situation of duplicate feature that may later become misleading (CCM = 0.44 and directional grade = 0.85). (c) Nondirectional correlation: sample of pixel with high correlation value with a nondirectional behavior (CCM = 0.62). (d) Directional correlation: directional feature where displacement along S can be very reliable and displacement along W should be ignored (CCM = 0.51 and directional grade = 0.75).

eliminates candidate features that are repeated and can be misleading in the motion analysis phase. The repeated features are identified by applying a labeling algorithm [21] to the auto correlation map and confirming there is not more than a single-connected label. Finally, we characterize each KFP according to its directivity. The motivation for this comes from the understanding that a key feature may have a directional nature and this should be properly accounted for in the analysis phase. As seen in Figure 3(d), some features have a directional nature. We classify the KFPs to be omni [Figure 3(c)] or directional [Figure 3(d)]. The directivity score of a KFP (P) is calculated by a 2-D ray scan on the autocorrelation map. Equation (3) describes how the directional score is computed for each ray. If a certain ray has a significantly higher value than the average, the pixel is considered to have a preferred directional nature. The segment with the high value represents its weak (W) direction and orthogonal to it is the strong direction (S). Figure 3(d) shows an example of a directional KFP where the vectors of W and S are identified.

$$\mathrm{ACMS}(P_1^R, P_2^R, \mathrm{ref}) = \frac{\sum_{i=1}^{n}\sum_{j=1}^{m}\left(\mathrm{ref}(P_1^R)_{ij} - \overline{\mathrm{ref}}(P_1^R)\right)\left(\mathrm{ref}(P_2^R)_{ij} - \overline{\mathrm{ref}}(P_2^R)\right)}{\sqrt{\left(\sum_{i=1}^{n}\sum_{j=1}^{m}\left(\mathrm{ref}(P_1^R)_{ij} - \overline{\mathrm{ref}}(P_1^R)\right)^2\right)\left(\sum_{i=1}^{n}\sum_{j=1}^{m}\left(\mathrm{ref}(P_2^R)_{ij} - \overline{\mathrm{ref}}(P_2^R)\right)^2\right)}}. \quad (1)$$

$$\text{Integral Autocorrelation Value } (P) = \frac{1}{nm}\sum_{x=\lfloor n/2\rfloor}^{\lfloor n/2\rfloor}\sum_{y=\lfloor m/2\rfloor}^{\lfloor m/2\rfloor}\mathrm{ACM}_P(x,y). \quad (2)$$

$$\text{directional grade}_\theta = \frac{1}{2m+1}\left[1 + \sum_{x=1}^{m}\left(\mathrm{ACM_P}(x, y(x,\theta)) + \mathrm{ACM_P}(-x, -y(x,\theta))\right)\right]. \quad (3)$$

where

$$y(x,\theta) = \lfloor \tan(\theta)x + 9(1 - \tan(\theta))\rfloor$$

and

$$\text{directional grades vector} = \{\text{directional grade}_\theta\}_{\theta=0:15:165}.$$

Intraoperative Motion Vector Analysis

During the clinical treatment, the anatomy is periodically scanned, and an updated input image set is acquired. On the basis of this updated information, we analyze each ADC. The cross-correlation computation is repeated for each KFP by comparing the patch from the original reference image, ref(P1), with patches from the input image (P2), as defined in (4), resulting in a cross-correlation score (CCMS). The terms are identical to the one described in (1). Only in (4), we compare two different images: the reference image and input image. A pixel with a high score determines a high correspondence between the two patches. However, the location with highest correlation value is not necessarily the best estimate for the displacement. To enhance robustness, we assume local rigidity within each cell. Equation (5) shows an energy function that combines a measure for how well the proposed location corresponds to the original image data given by (4) and a second term (neighbor's correlation) that reflects the similarity to the displacements of neighboring KFPs, as in (6). The coefficients for the weighted ratio determine how much we enforce local rigidity with respect to variability in local displacement.

$$\mathrm{CCMS}(P_1^R, P_2^I, \mathrm{ref}, \mathrm{input}) = \frac{\sum_{i=1}^{n}\sum_{j=1}^{m}\left(\mathrm{input}(P_1^I)_{ij} - \overline{\mathrm{input}}(P_1^I)\right)\left(\mathrm{ref}(P_2^R)_{ij} - \overline{\mathrm{ref}}(P_2^R)\right)}{\sqrt{\left(\sum_{i=1}^{n}\sum_{j=1}^{m}\left(\mathrm{input}(P_2^I)_{ij} - \overline{\mathrm{input}}(P_1^I)\right)^2\right)\left(\sum_{i=1}^{n}\sum_{j=1}^{m}\left(\mathrm{ref}(P_1^R)_{ij} - \overline{\mathrm{ref}}(P_1^R)\right)^2\right)}}. \quad (4)$$

$$\text{Correlated Displacement Value} = \alpha \times (\text{Max Cross-Correlation Score}) + \beta \times (\text{Neighbors Corrolation}). \quad (5)$$

$$\text{Neighbors Corrolation } (P_j^I) = 1 - \sum_{n=1}^{N}\left(\left\|P_j^I - P_n^I\right\| - \left\|P_j^R - P_n^R\right\|\right)^2. \quad (6)$$

The estimated displacement of each cell Ci is a directional weighted averaging of all its related pixels. The weighted averaging is designed to ensure two things: 1) pixels with higher cross-correlation value will have more contribution to the average displacement and 2) directional pixels contribute only in the appropriate (S) direction [Figure 3(d)] as seen in (7) and (8).

$$\Delta X_{\mathrm{AVG}} = \frac{\sum_{i=1}^{N} dX_i \times \mathrm{CCV}x(Pi)}{\sum_{i=1}^{N}\mathrm{CCV}x(Pi)}$$

$$\Delta Y_{\mathrm{AVG}} = \frac{\sum_{i=1}^{N} dY_i \times \mathrm{CCV}y(Pi)}{\sum_{i=1}^{N}\mathrm{CCV}y(Pi)}. \quad (7)$$

$$\mathrm{CCV}_x(Pi) = \text{Max Cross-Correlation Score}(Pi) \times \cos\theta.$$
$$\mathrm{CCV}_y(Pi) = \text{Max Cross-Correlation Score}(Pi) \times \sin\theta. \quad (8)$$

Figure 5 shows an example of how the directional and nondirectional KFPs are combined to result in the local motion vector. The original location of the cross mark is shown in a dashed outline. The original KFPs are shown in green cross marks (labeled 1–6). A diagonal displacement has been applied (gray arrow). The contribution of nondirectional KFPs is taken in both axes (green vectors), whereas in the directional

Our approach has been uniquely developed and integrated into a clinical product.

case (blue vectors), only its strong component (red vectors) affects the final motion vector (gray vector).

The final part of the algorithm compares the displacements with a predetermined threshold. This threshold is expected to vary according to the clinical procedure, its safety requirements and depending on the treated organ. Since we do not assume the overall displacement to be rigid, there is no reason to assume that different anatomical descriptors, which represent different locations in the anatomy, will agree on the same displacement. Therefore, the analysis is based on comparing each descriptor cell with the motion threshold. Figure 4 shows an example of a global motion test. Figure 4(a) shows the reference image and Figure 4(b) reveals the input image (after the motion occurred) taken during treatment. The green circles show the original location of the anatomic descriptors on both images. The red circles show the estimated location of the descriptor on the input image. Figure 4(c) shows a 2-D plot of the displacement for each descriptor in a sample situation. When a single or a group of anatomical descriptors have a displacement that exceeds the threshold defined for the application, we issue a motion warning.

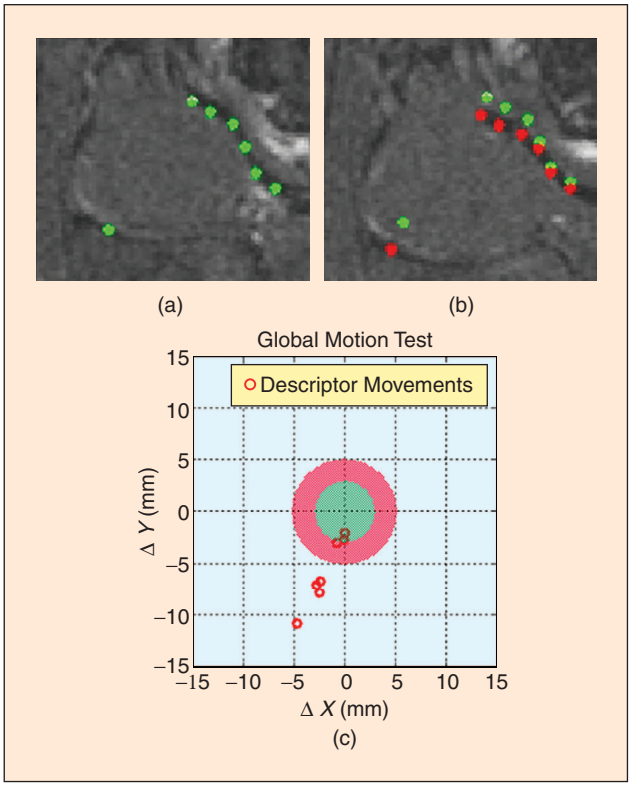

Fig. 4. (a) User-defined anatomical descriptors on the reference image. (b) Analyzed descriptors (red) on the input image. User-defined descriptors (green) are superimposed. (c) Global motion test: the green and red areas mark the allowed and critical displacements, respectively.

Data Sets

To test our approach, we used a clinical database that contains imaging data of hundreds of treatments performed with a noninvasive MR-guided focused ultrasound therapy system. The data sets were anonymized and originated from treatments of uterine fibroids (i.e., benign tumor in the uterus). We used this sample database to test the algorithm with specific thresholds defined by clinicians, setting a displacement of 4 mm (or higher) to be large and an error below 2 mm to be small and insignificant.

MR images are acquired and used to describe and monitor the patients' anatomy during the treatment. Typically, the treatment volume is covered by 20–30 sagittal images acquired for reference at the beginning of the treatment. The annotation on the reference images includes easily noticeable landmarks in and adjacent to the target volume. The intraoperative images are of the same modality and scanner and include one sample slice acquired every 3 s.

The algorithm is implemented using MATLAB infrastructure running on Windows XP. Performance was typically less

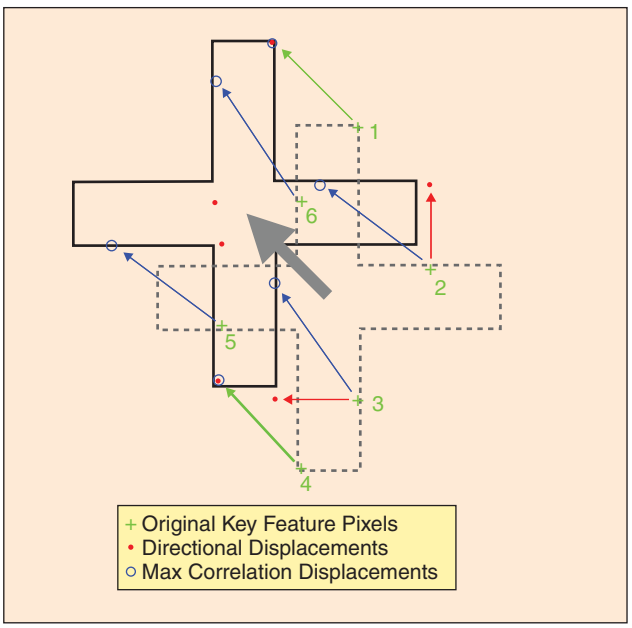

Fig. 5. Local motion vector. Input image with diagonal displacement. Nondirectional contribution (green) and strong component (red) of nondirectional (blue) are combined to result in final local motion vector.

The robust results are achieved due to the weighted directional descriptors approach.

than 3 s to analyze ten features in a pair of images of 256 × 256 running on a 3-GHz PC machine.

Results

Gold Standard Analysis

The most straightforward quantitative test is a comparison to gold standard. Overall, 606 anatomical descriptors were selected by the clinical expert users in 32 randomly selected treatments from the clinical database.

The algorithm results are compared with the results of a manually marked displacement done by expert clinicians (Figure 6). According to our clinical criteria, an error lower than 2 mm is considered a hit (acceptable). An error exceeding 2 mm is considered misplaced. An error exceeding 4 mm is defined as a false alarm. Finally, the category associated with safety aspects and with the highest concern is a case of motion not detected by the algorithm (undetected motion or false negatives). Our approach is robust and has zero events of undetected motion. In Figure 7, it can be seen that 98%, vast majority of the cases, is classified as hits, and only less than 2% of the descriptors is misplaced by the algorithm. From safety aspects, it is important to note that, throughout these tests, there were zero events of undetected motion or false alarms. Both situations, even at low rate, would cause safety or reliability concerns that would lead clinicians to reject the approach.

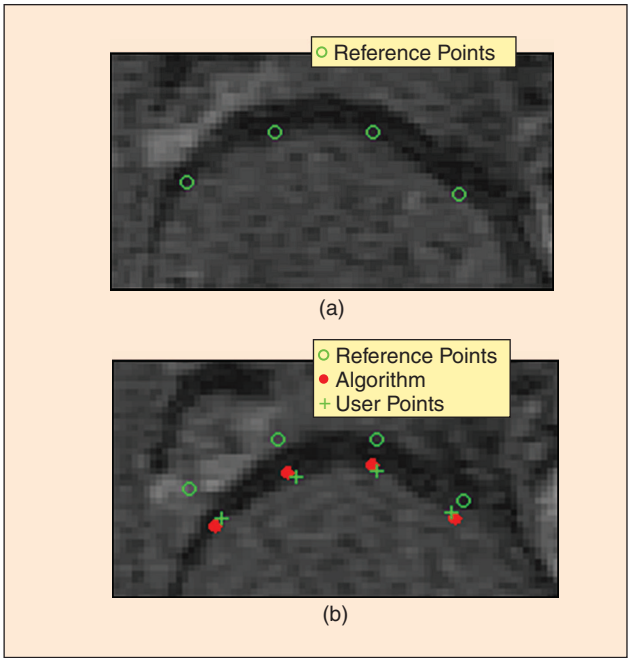

Fig. 6. An example of a gold standard procedure. (a) Anatomical descriptors on reference image set. (b) Input image with original location descriptors (circles) and expert user gold standard placement (cross marks). In red are the algorithm results.

Localization Sensitivity

Our approach relies on user input for marking the anatomical feature that requires tracking. This part cannot be automated since it requires an understanding of the anatomy and the clinical procedure performed. However, our approach does not demand that this input be very precise but rather a mere gesture. To test the sensitivity of the algorithm to the user input, we repeated the same analysis after randomly altering the location of the user input by ±3 mm.

The results show that the average displacement error increased by only 0.15 mm (from 4.08 to 4.23). The rest of the results were the same as in the original test. This is a key attribute in providing an approach that is clinically acceptable. We do not have to impose stringent and stressful constraint on our users by requiring very accurate labeling. The approach is satisfied with a mere gesture of where the target KFPs are, not requiring accurate user input.

Comparison to SIFT

The use of SIFT is very common in computer vision and can be used for motion tracking and registration [20]. To compare the performance of our proposed concept to SIFT, a typical set of

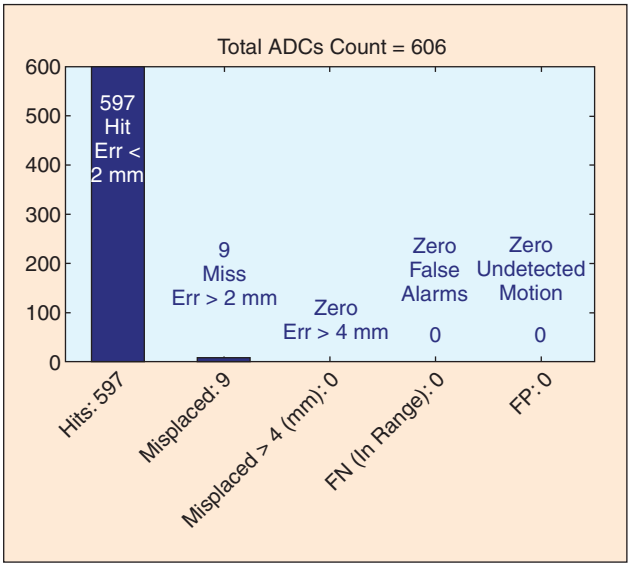

Fig. 7. Results of the gold standard analysis. Histogram shows the result of 606 tests comparing results from our approach with expert user manually analyzed displacement.

The use of scale invariant feature transform is very common in computer vision and can be used for motion tracking and registration.

brain MR images were used. A gold standard displacement was manually marked by expert clinicians. This result was compared with the estimated displacement of SIFT and of the proposed approach. The mean error for SIFT was 17.1 mm compared with 1.76 mm in our approach. SIFT, although a popular approach in academic research, is not robust for clinical use. The error is ten times larger than the proposed approach, which would lead to clinically unacceptable situations.

Discussion

The uniqueness of the approach is that it deals with true intraoperative combination of challenges making it practical for immediate use in a clinical environment, thus translating a huge academic research challenge into true accepted medical practice. Ensuring safety via low rate of false alarms and minimizing events of undetected motion to zero leads to a successful implementation in the clinical arena.

Medical images are taken with a wide field of view to provide the physician with the full view of the relevant anatomy. However, the scanned volume includes structures that are of no interest to the procedure being applied. The notion of excluding the nonrelevant anatomy requires medical understanding of the anatomy and the procedure performed. Ignoring this would result in false alarms that would be not accepted clinically. This is why our approach relies on user input for the designation of the anatomical descriptors that are to be tracked. Although user input is essential, it is important to keep it from being demanding or time consuming. We minimize the user interaction to mere gestures pointing to the features of interest, and we show the approach is resilient to errors and the localization does not have to be precise. The detailed evaluation methods included also a sensitivity analysis to image quality. As could be expected poor image quality affects the results. It leads to more false alarms and an increase in displacement errors. However, it is important to mention that the key safety criterion of zero undetected motion events is still maintained with our approach.

The analysis of anatomical features shows that a unique preprocess is essential to identify the key pixels to be tracked. Features that have a pattern that is repeated in several locations within the search range must be ignored since they may be misleading during the analysis phase. In many situations, a feature has a directional nature (e.g., edges). Such features may provide very reliable displacement in one direction (i.e., normal to the edge), but any other displacement must be ignored [Figure 3(d)]. These realizations are essential to reliably estimate the displacement as well as to avoid false alarms.

Medical images are challenging for many automated approaches that are commonly used elsewhere due to the complexity of the details and the demanding clinical requirements. These factors became evident when comparing our WDDs approach with the notable approach of SIFT, which is very successful in other image processing applications. The SIFT approach was misled by repeated features in the image [Figure 3(b)]. The SIFT approach ignores information given by displacement of edges because of their directional nature. We found that, without introducing these concepts, the results are not robust enough to be accepted in clinical environment.

The approach has been implemented for quasi-static organs where motion is not expected during the treatment (uterus, breast, prostate, and brain). This was a fundamental phase in validating the proof of concept of the image-based motion detection technique. Further development is now being done to adapt the approach to rapidly moving organs such as liver where the treatment target area is constantly moving. The main challenges are in high-temporal resolution and higher performance algorithmic implementation. Since the organ is not static, there are additional algorithmic challenges in estimating the motion and comparing it with a predictive model and updating the treatment plan accordingly.

The development of this novel technique for intraoperative motion detection was challenged by more than just the technical algorithmic difficulties. There was a need to translate clinical needs to technical requirements that would meet the expectations of clinicians. The robustness and performance are key factors in the acceptance of this tool and relying on it in true clinical situations. InSightec is developing additional products for treating prostate cancer, brain tumors, and breast cancer, which will benefit from our experience that will be implemented in these applications in the future. In coming years, clinical studies will be conducted to develop statistical data needed for reimbursement coverage and additional clinical application will be investigated to reach wide use of this technology for improving patient-care and health-care costs.

Conclusions

The use of image guidance in clinical procedures brings huge benefits to patients and to health-care applications. When relying on preoperative imaging, tissue movement can create safety concerns if not detected. Although studied thoroughly in the academic arena, image-based motion detection is a challenge that has not yet met the goal of being widely accepted in the true clinical environment.

The robustness of the proposed algorithm has been validated on true operative data. The encouraging robust results have been shown to be resilient to inaccurate user input and imperfect image quality. The results show accurate motion analysis with very low rate of false alarms and zero events of undetected motion under routine clinical conditions. The approach has matched the clinical requirements and is integrated in an image-based noninvasive therapeutic system under MR imaging system. This technique is bound to have a

huge impact in procedures where image-based tissue motion detection is needed.

The approach has been validated by clinicians and is already in use in medical device products. In future, it will be implemented in additional clinical procedures and will have a significant impact by improving health care delivered to patients.

Eyal Zadicario graduated in aeronautical engineering (cum laude) from the Technion, Israel Institute of Technology and completed his M.Sc. degree (summa cum laude) in computer science from the Tel Aviv University School of Mathematics. He is currently the director of Neuro Programs in InSightec. Previously, he was a part of the software development team responsible for research development of SW applications and techniques for medical device applications.

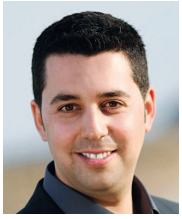

Shlomi Rudich received his B.Sc. degree in biomedical engineering from Tel-Aviv University in 2007 and a Software Developer Certification from the Technion in 2009. He is a clinical applications engineer at InSightec.

Ghassan Hamarneh received his bachelor's degree from Jordan University in 1995 and master's degree, with distinction, in digital communications from Chalmers University, Sweden, in 1997. He completed his doctoral studies at Chalmers University of Technology, Sweden, in 2001 and was a predoctoral research fellow at the University of Toronto, Canada, from 2000 to 2001. He is an associate professor in computing science at Simon Fraser University (SFU), Canada. He is the codirector and cofounder (2003) of the Medical Image Analysis Laboratory (MIAL) at SFU. His research interest includes medical image analysis (segmentation, registration, and anatomical shape modeling and analysis). Before joining SFU in 2003, he was a postdoctoral fellow at the Hospital for Sick Children (Mouse Imaging Center) and the University of Toronto, Canada from 2001 to 2003.

Daniel Cohen-Or received his B.Sc. cum laude in both mathematics and computer science in 1985, an M.Sc. cum laude in computer science in 1986 from Ben-Gurion University, and his Ph.D. degree from the Department of Computer Science, State University of New York, Stony Brook, in 1991. He is a professor at the School of Computer Science. He received the 2005 Eurographics Outstanding Technical Contributions Award. His research interests include computer graphics, in particular, rendering and modeling techniques; image synthesis, motion, and transformations; shapes and surfaces; and surface reconstruction.

Address for Correspondence: Eyal Zadicario, InSightec LTD, 5 Nahum Heth, Tirat Carmel, Israel 39120. E-mail: eyalz@insightec.com.

References

[1] S. Dimaio, T. Kapur, K. Cleary, S. Aylward, P. Kazanzides, K. Vosburgh, R. Ellis, J. Duncan, K. Farahani, H. Lemke, T. Peters, W. Lorensen, D. Gobbi, J. Haller, L. Clarke, S. Pizer, R. Taylor, R. Galloway, Jr., G. Fichtinger, N. Hata, K. Lawson, C. Tempany, R. Kikinis, and F. Jolesz, "Challenges in image-guided therapy system design," *NeuroImage*, vol. 37, Suppl. 1, pp. S144–S151, 2007.
[2] D. Verellen, M. Ridder, and G. Storm, "A (short) history of image-guided radiotherapy," *Radiotherapy Oncol.*, vol. 86, no. 1, pp. 4–13, 2007.
[3] D. Jaffray, P. Kupelian, T. Djemil, and R. M. Macklis, "Review of image-guided radiation therapy," *Expert Rev. Anticancer. Ther.*, vol. 7, no. 1, pp. 89–103, Jan. 2007.
[4] M. Moradi, P. Abolmaesoumi, and P. Mousavi, "Deformable registration using scale space keypoints," *Proc. SPIE*, vol. 6144, pp. 791–798, 2006.
[5] E. R. Mcveigh, "MRI of myocardial function: Motion tracking techniques," *Magn. Reson. Imag.*, vol. 14, no. 2, pp. 137–150, 1996.
[6] Z. Fang, T. Möller, G. Hamarneh, and A. Celler, "Visualization and exploration of time-varying medical image data set," in *Proc. Graphics Interface 2007 (ACM Int. Conf. Proc. Ser., 234)*, 2007, pp. 281–288.
[7] M. J. Murphy, "Tracking moving organs in real time," *Semin. Radiat. Oncol.*, vol. 14, pp. 91–100, 2004.
[8] G. C. Sharp, S. B. Jiang, and H. Shirato, "Prediction of respiratory tumor motion for real-time image-guided radiotherapy," *Phys. Med. Biol.*, vol. 49, pp. 40–425, 2004.
[9] J. Tokuda, S. Morikawa, T. Dohi, and N. Hata, "Motion tracking in MR-guided liver therapy by using navigator echoes and projection profile matching," *Acad. Radiol.*, vol. 11, pp. 111–120, 2004.
[10] B. J. Matuszewski, J. K. Shen, L. K. Shark, and C. J. Moore, "Estimation of internal body deformations using an elastic registration technique," in *Proc. Int. Conf. Medical Information Visualization—BioMedical Visualization*, 2006, pp. 15–20.
[11] J. Cho, "Elastically deformable model-based motion-tracking of left ventricle," in *Proc. 26th Annual Int. Conf. IEEE EMBS 1*, Sept. 2004, vol. 3. pp. 1925–1928.
[12] X. Papademetris, A. J. Sinusas, D. P. Dione, R. T. Constable, and J. S. Duncan, "Estimation of 3-d left ventricular deformation from medical images using biomechanical models," *IEEE Trans. Med. Imag., 21*, vol. 7, pp. 786–800, July 2002.
[13] Y. Chenoune, E. Delechelle, E. Petit, T. Goissen, J. Garot, and A. Rahmouni, "Segmentation of cardiac cine-MR images and myocardial deformation assessment using level set methods," *Computer. Med. Imag. Graph. 29*, vol. 8, pp. 607–616, Sept. 2005.
[14] T. Mcinerney and D. Terzopoulos, "Deformable models in medical image analysis—A survey," *Med. Image Anal. 1*, vol. 2, pp. 91–108, June 1996.
[15] F. Huang and J. Su, "Face contour detection using geometric active contours," in *Proc. 4th World Congress on Intelligent Control and Automation*, Nov. 2002, vol. 3. pp. 2090–2093.
[16] J. B. A. Maintz and M. A. Viergever, "A survey of medical registration," *Med. Image Anal.*, vol. 2, pp. 1–36, Mar. 1998.
[17] H. Chui and A. Rangarajan, "A new point matching algorithm for non-rigid registration," *Comput. Vis. Image Understand.*, pp. 114–141, Feb.–Mar. 2003.
[18] Z. Yefeng and D. David, "Robust point matching for nonrigid shapes by preserving local neighborhood structures," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, pp. 643–649, Apr. 2006.
[19] K. Rohr, H. S. Stiehl, R. Sprengel, W. Beil, T. M. Buzug, J. Weese, and M. H. Kuhn, "Point-based elastic registration of medical image data using approximating thin-plate splines," *Visual. Biomed. Comput.*, pp. 297–306, 1996.
[20] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, pp. 91–110, Nov. 2004.
[21] R. M. Haralick and L. G. Shapiro, *Computer and Robot Vision*. Reading, MA: Addison-Wesley, 1992, pp. 40–48.
[22] M. I. Miller, A. Trouve, and L. Younes, "On the metrics and euler-lagrange equations of computational anatomy," *Annu. Rev. Biomed. Eng.*, vol. 4, pp. 375–405, 2002.
[23] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *Int. J. Comput. Vis.*, vol. 61, no. 2, pp. 139–157, Feb. 2005.
[24] G. E. Christensen, P. Yin, M. W. Vannier, K. S. C. Chao, J. L. Dempsey, and J. F. Williamson, "Large-deformation image registration using fluid landmarks," in *Proc. 4th IEEE Southwest Symp. Image Analysis and Interpretation*, pp. 269–273, 2000.
[25] J.-P. Thirion, "Image matching as a diffusion process: An analogy with Maxwell's demons," *Med. Image Anal.*, vol. 2, no. 3, pp. 243–260, Sept. 1998.
[26] I. K. Sethi and R. Jain, "Finding trajectories of feature points in a monocular image sequence," *IEEE. Trans. Pattern Anal. Mach. Intell.*, vol. 9, no. 1, Jan. 1987.
[27] M. Moradi and P. Abolmaesumi, "Medical image registration based on distinctive image features from scale- invariant (SIFT) keypoints," in *Proc. Computer-Assisted Radiology and Surgery (CARS) Conf.*, 2005, p. 1292.
[28] W. Cheung and G. Hamarneh, "n-SIFT: n-Dimensional scale invariant feature transform," *IEEE. Trans. Image Process.*, vol. 18, no. 9, pp. 2012–2021, 2009.