Int J CARS (2010) 5 (Suppl 1):S137–S142
DOI 10.1007/s11548-010-0450-7

Hierarchical Normal Vector Information based registration for surgical navigation by using saliency information

H. Dang[1], L. Gu[2], X. Zhuang[3], Z. Luo[4]
[1]Shanghai Jiao Tong University, Biomedical Engineering, Shanghai, China
[2]Shanghai Jiao Tong University, Med-X Research Institute, Shanghai, China
[3]University College London, Medical Physics & Bioengineering Department, London, United Kingdom
[4]Shanghai Jiao Tong University, School of Software, Shanghai, China

Keywords Normal vector information · Medical image registration · Saliency · Hierarchical

Purpose

Fast and accurate image registration methods have been in great demand in many biomedical applications, especially in image guided surgery and therapy. In specific, when the patient changes pose or has a surgical resection between two samplings (mono-modality) or the subject is imaged in substantially two different ways (multi-modality), those images needs registration before they can be used in surgical navigation systems. Mutual Information (MI) is a popular similarity metric for image registration, especially multi-modality registration, and much improvement on it has been achieved. However, it has a number of well-known drawbacks: MI metric is highly non-convex and has many local minima; it develops from information theory and has a discrete nature; finally it decouples the connection between the gray value and its location information, which is its biggest problem.

To complement the drawbacks of MI metric in image registration, we have proposed an alternative measure which is based on normalized gradients using properties of differential geometry, named Normal Vector Information (NVI) [1]. Normal Vector (NV) is the vector that is perpendicular to the local isosurface of one pixel in one image (Fig. 1.). We have also proved that the average accuracy of the success registrations of NVI is better than MI and NMI methods, while NVI consumes more time because it employs all the pixels in one image to compute their NV values [2]. Thus, in this paper, to speed up NVI towards practical surgical navigation application while preserving its high accuracy, we propose a novel hierarchical model for NVI based registration, which falls into two steps. In the first step, registration starts from feature points carefully selected from image's saliency map to quickly get a coarse transform. In the second step, based on this coarse transform, registration switches from feature points to all pixels and reaches a fine transform finally.

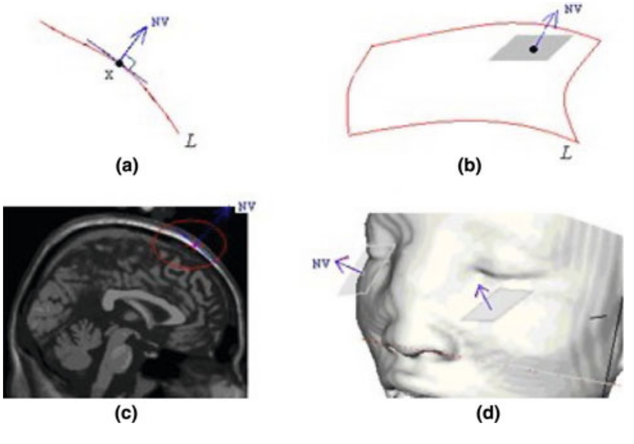

Fig. 1 Normal Vector Information of Image: (**a**), (**c**) NVI in 2D image; (**b**), (**d**) NVI in 3D image

The selection of feature points is very important in this hierarchical model, since it determines to which extent the feature-points-based coarse transform approximates the final one. In fact, registration in second step will cost less time, the closer this approximation is. Here we introduce saliency information on each pixel as criterion for selecting feature points. Saliency theory was first discovered in neuroscience by L. Itti and etc [3]. Different with basic properties in images such as color, texture and shape, saliency provides high level semantic understanding of image so that it has been widely used recently for image analysis. Many methods have been developed to define the saliency of image, i.e., using edge gradient, local phase, and salient regions. Inspired by contrast-based saliency measure [4], we define a two-step scale and rotation invariant saliency operator [5] to extract saliency of each pixel from its surrounding intensity contrast.

Methods

A. Calculation of Normal Vector

In normalized gridded images, Normal Vector can be calculated by Grey Intensity Gradient:

$$G_x = [g(x+1,y,z) - g(x-1,y,z)]/2\mathrm{spacing}_x$$
$$G_y = [g(x,y+1,z) - g(x,y-1,z)]/2\mathrm{spacing}_y$$
$$G_z = [g(x,y,z+1) - g(x,y,z-1)]/2\mathrm{spacingz}$$

Then Normal Vector value in (x, y, z) is given by:

$$N_{x,y,z} = (G_x, G_y, G_z)/\sqrt{\left(G_x^2 + G_y^2 + G_z^2\right)}$$

B. Definition of NVI Metric

Here cosine of the included angle of two NVs is adopted to evaluate the similarity between two NVs. Then the average of all the square of cosine values of the NVI images could be a standard in evaluating how much the two NVI images are in similarity, that is, how much the respective images are aligned. This method can exclude the problem that the two images have information at different scale in some regions such as the central part of the brain. The NVI similarity measure between image A and B is defined as below:

$$S(A,B) = (1/|V|)\sum \cos^2\theta(N_A(X), N_B(X))\ (X \in V)$$

In practice, the calculation of cosine of the included angle is:

$$\cos\theta = NV_A \cdot NV_B = N_{xA}N_{xB} + N_{yA}N_{yB} + N_{zA}N_{zB}$$

C. Contrast-based Saliency Measure

The contrast value $C_{i,j,k}$ on one perception unit (i,j,k) is defined as follows:

$$C_{i,j,k} = \sum d\left(p_{i,j,k}, q\right)(q \in \Theta)$$

where $p_{i,j,k}$ and q denote the stimulus perceived by perception units. Θ is the neighborhood of perception unit (i,j,k). The size of Θ controls the sensitivity of perceive field. d(p,q) is the difference between $p_{i,j,k}$ and q, which may employ any suitable distance measure according to applications.

D. Two-step Scale and Rotation Invariant Saliency Operator

Our implementation of the above contrast-based saliency measure is defined as follows:

$$C_{i,j,k} = \sum \left(I_{i,j,k} - I_q\right)^2 (q \in \Theta)$$

where intensity I is used as stimulus on perceive field, and the difference d(p,q) is computed by the square of intensity difference. In practice, we set Θ as 6 neighbourhood pixels around (i,j,k), 2 pixels in each axis. After all saliency values $C_{i,j,k}$ in one image are computed, a saliency map is formed. By adjusting threshold K, we can select proper number of pixels whose saliency values are larger than K as feature points for coarse registration in the first step.

Results

Our registration is carried out in a surgical navigation system named SurgView developed by our lab. The registration framework employs many ITK (the Insight Toolkit) classes, which provide the implementations for the algorithms of Linear Interpolation and Rigid Transformation interactive with our NVI metric. In order to test the robustness and universality of our method in clinical setting, our source data are 3D clinical patient data including prostate ultrasonic images (354*189*175) and kidney CT images (319*239*239), both from Shanghai Renji Hospital. A P-IV 1.60 GHz PC, 2.5 GB main memory, MS Windows XP desktop computer running Windows-XP was employed to run the registration. All registration experiments are mono-modality since the main concern here is to examine whether and to what extent the hierarchical model with saliency information can speed up NVI based registration. Multi-modality registration experiments will be conducted in later work.

Here we offer US (Ultrasonic) to US and CT to CT registration experiments. The maximal iterations in both cases are set to 200 steps when the learning rate is set to 0.05 for both translation and rotation parameter. X, Y, Z translation scope randomly varies from 0 pixel to 10 pixels, angle of rotation randomly varies from 0 to 10 degree. The crucial threshold is set to be twice as large as the average of all saliency values in one image. The adaptivity of this threshold ensures a proper proportion of salient points in different input images can always be selected. The results of US to US registration are reported in Fig. 2. In these series of experiments, the registration using only NVI demonstrates a good result of image alignment, but costs too much time for calculating. The hierarchical model which employs feature points for registration, costs much less time and produces an equally satisfactory registration result.

Conclusion

In this abstract, we propose a novel hierarchical model for NVI based registration, which starts registration from feature points carefully selected from image's saliency map, inherits coarse transform and finally reaches fine transform on all pixels. A quantitative evaluation under our surgical navigation system shows this hierarchical model largely speeds up NVI based registration while preserving its high accuracy and avoiding being trapped in local minima. To facilitate the application of this hierarchical model in image fusion in surgical navigation systems, later work will emphasize on examining and modifying this hierarchical model in multi-modality setting.

Acknowledgement

This work was supported in part by National natural science foundation of China, Grant No. 60872103.

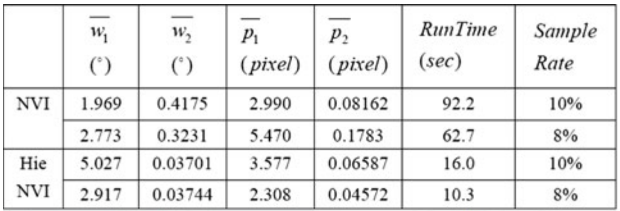

	$\overline{w_1}$ (°)	$\overline{w_2}$ (°)	$\overline{p_1}$ (pixel)	$\overline{p_2}$ (pixel)	RunTime (sec)	Sample Rate
NVI	1.969	0.4175	2.990	0.08162	92.2	10%
	2.773	0.3231	5.470	0.1783	62.7	8%
Hie NVI	5.027	0.03701	3.577	0.06587	16.0	10%
	2.917	0.03744	2.308	0.04572	10.3	8%

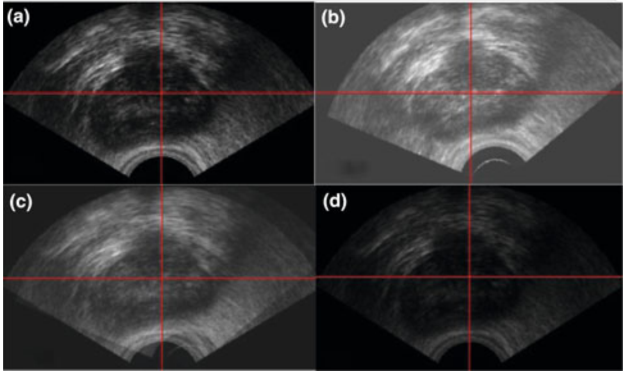

Fig. 2 The table above is the result of ultrasonic to ultrasonic registration. Hie NVI is short for Hierarchical NVI. w_1 and w_2 mean average displacement in translation before and after registration. p_1 and p_2 mean average displacement in rotation before and after registration. Run_Time is the average run time for registration and Sample_Rate gives the sampling rate. The images below are one slice of 3D ultrasonic image in our surgical navigation system. (**a**) The reference image; (**b**) The test image after random transform; (**c**) Overlapping of the reference image and the test image; (**d**) Overlapping of the reference image and the transformed test image after registration using Hierarchical NVI

References

[1] Xiahai Zhuang, Lixu Gu, Jianfeng Xu (2005) Medical image alignment by normal vector information. (Eds.): CIS 2005, Part I, LNAI 3801, pp. 890–895, Springer-Verlag Berlin Heidelberg, 2005

[2] Xiahai Zhuang, Lixu Gu (2006) Normal vector information registration and comparisons with mutual information. Engineering in Medicine and Biology Society, 2006 EMBS ‘06.28th Annual International Conference of the IEEE Aug, 2006: 3827–3830

[3] L. Itti, C. Koch (2001) Computational modeling of visual attention. Nat. Rev. Neurosci. 2(2001), pp. 194–203

[4] Yu-Fei Ma, Hong-Jiang Zhang (2003) Contrast-based image attention analysis by using fuzzy growing, Proceedings of the eleventh ACM international conference on Multimedia, November 02-08, 2003, Berkeley, CA, USA [doi:10.1145/957013.957094]

[5] Binjie Qin, Zhijun Gu, Xianjun Sun, Yisong Lv (2008) Registration of images with outliers using joint saliency map. IEEE Signal Processing Letters, 17(1):91–94, 2010

SSD based real-time tracking of myocardial motion using 3D streaming TEE volumes

B. Ramachandran[1], V. Parthasarathy[1], C. Hatt[2], A. Raval[2], A. Jain[1]
[1]Philips research, IGT, Briarcliff, Manor, NY, United States
[2]University of Wisconsin, Madison, United States

Keywords Cardiac · Motion · 3D ultrasound · Real-time

Purpose

Over 1 million people in the United States and 19 million people worldwide experience sudden cardiac death or myocardial infarction each year. Treatments have significantly evolved over the past decades and image guided interventions have successfully replaced invasive surgical methods in some cases. Image guided interventions benefit from patient specific anatomy, which is typically obtained using high resolution MRI/CT scans. Though these imaging modalities provide high quality images and a detailed description of the cardiac anatomy, it is still a challenge to use them intra-operatively to provide real-time guidance. Real-time 3D trans-esophageal echocardiography (TEE) offers a potential break-though, offering real-time volumetric cardiac tissue visualization at low a cost and without any side-effects. Real-time TEE can compute the cardiac motion and hence compensate the pre-operative plan to match the current location of the anatomy. In this work, we develop an algorithm which tracks a set of points in live 3D TEE image of the heart in real-time to help compensate for cardiac motion during surgery.

Methods

In this work, we present a real-time tracking scheme. Point tracking in TEE images is a challenging task due to the low signal to noise ratio and inherent speckle noise. Many techniques and tracking algorithms have been proposed, most of them applied for 2D tracking. Maximum likelihood techniques have been used to estimate motion in a sequence of B-Mode ultrasound images, wherein the motion field is computed using accurate statistical models of the ultrasound images. Block matching methods are based on comparing intensity of image blocks using norm, and accounting for speckle patterns, but have not been accurate enough to be extended from 2D to 3D. Automatic contour tracking, wherein the surface is treated as a parametric model with a sequential state space is used for tracking the deformation state using a Kalman filter or predictor while minimizing mean square error and supporting the estimation of past present and future states. Alternatives to such a methods study the motion of the heart, understanding the dynamic of the heart. Another major class of tracking algorithms used for tracking in ultrasound images is based on sum square difference (SSD) and on cross-correlation. Complex motions which can be modeled as a combination of translation, rotation and scaling can be tracked in the general framework of object tracking using SSD, including real-time motion of internal organs in 3D ultrasound images. Since the logarithm performed to convert the RF US data to B-mode images converts the multiplicative Raleigh noise to additive noise, SSD based algorithms are usually efficient. SSD and cross-correlation coefficient along with maximum likelihood estimator have also been applied to 3D cardiac ultrasound for applications such as for ASD correction. All aforementioned techniques apply universal description of objects without considering any temporal relationships, which leads to temporal inconsistence between adjacent frames. To overcome this difficulty, forward prediction using motion manifold learning has been evaluated. Other techniques such as tracking deforming objects using particle filtering, fast point tracking algorithm, demon's algorithm and condensation algorithm have been suggested, though more in camera images than in volumetric US volumes. To preserve a high run-time from streaming volumetric datasets, we evaluate a 3D SSD based algorithm for tracking myocardial motion in the TEE images. An effort is made to detect for error or drift while tracking by making use of a weighted squared difference (WSD) value as a metric and correcting for the same by switching the search space and mask. This method (Fig. 1) provides us with an implementation that is both real-time and accurate for our clinical application.

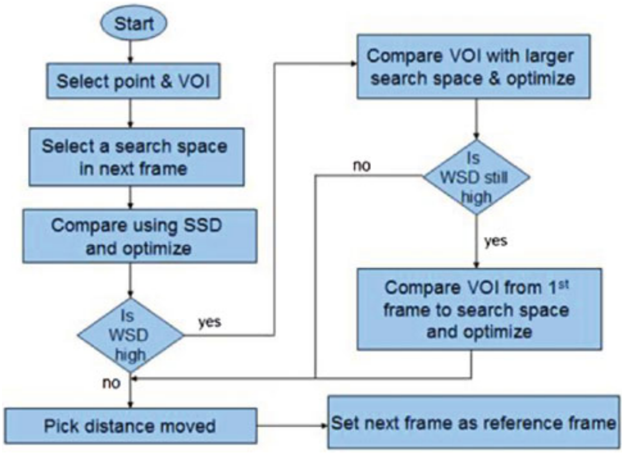

Fig. 1 Flowchart showing the working of the motion tracking algorithm

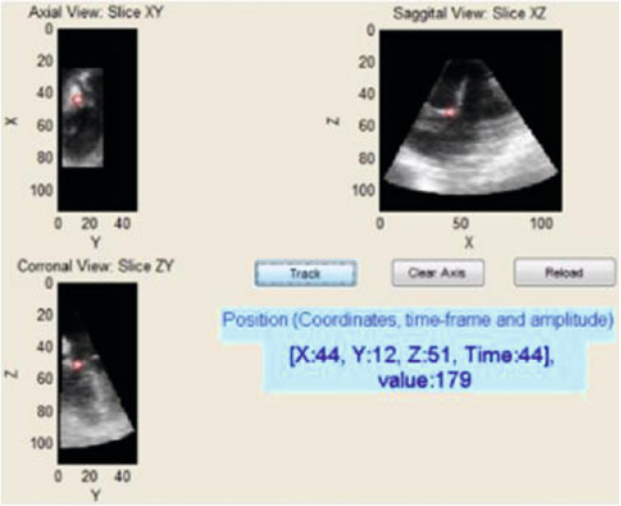

[X:44, Y:12, Z:51, Time:44], value:179

Fig. 2 Snapshot of tracker as it tracks 3D cardiac motion in a pig

Results

The algorithm is validated using simulation, phantom and animal data. The simulation data consisted of rotation, translation and scaling performed to different extents, wherein to-and-fro rotation about the Z-axis and to-and-fro translation in the X and Y dimensions mimicked the myocardial motion. 3D pig datasets at end diastole were taken as the reference to create the moving datasets with known motions. The output given by the tracking algorithm is compared with the true position under different conditions of rotation, translation and scaling undergoing myocardial motion. Results showed that the algorithm offered a tracking error of less than 1 pixel for translational myocardial motion under 20 mm/cardiac-cycle and rotation of 20 degrees/cardiac-cycle. In phantom datasets, a cardiac phantom having motion driven from its apex and made from a poly-vinyl alcohol cryogen was immersed in a water tank and used for the experiments. The Philips iE33 X7-2t TEE probe was used to acquire live 3D ultrasound images, the volumes streamed from the US machine to the workstation. In a similar protocol, pre-clinical data was also collected using pigs (Fig. 2). Since ground truth in phantom and pre-clinical data is not quantitative, condition for success was defined as not losing the tracked feature. In datasets that varied in length from 20–60 seconds, the algorithm was found to be able to continuously track the feature of interest. The algorithm offers a frame rate of 25–100 Hz in a Matlab implementation, depending on the choice of the different parameters.

Conclusions

We present an algorithm to track myocardial motion of a pre-selected point in a live 3D TEE image. The drift detection using WSD and its correction significantly improved the robustness of the algorithm. Tracking was found to be accurate, real-time and robust. However, it was significantly influenced by the quality of the data, namely its spatial and temporal resolution. This understanding will be critical in improving the robustness of tracking during clinical interventions. As future work, an ECG/state-space based correction needs to be incorporated that additionally performs a similarity measure at every end diastole. Finally, these preliminary results show promise towards improving the robustness of motion tracking in these low resolution volume streams (high speckle noise and low spatial/temporal resolution).

3D Shape Analysis of Arthritic Hips: A Preliminary Study

R. Ellis[1], B. Rasquinha[1], G. Wood[1], J. Rudan[1]
[1]Queen's University, Kingston, Ontario, Canada

Keywords Computer assisted surgery · Orthopedic surgery · Hip anatomy · Osteoarthitis

Purpose

A conundrum in the anatomy and kinematics of the hip joint is that, although the hip is ubiquitously labeled as a ball and socket (or spherical) joint in medical textbooks, in many of the same medical textbooks the acetabulum and femoral head are further described as respectively being approximately hemispherical and spheroidal. Mechanically, a ball and socket joint implies a particular kinematic behavior: pure rotation about an isocenter. It also implies a particular geometry for the articulating surfaces of the joint, as the kinematics are due to the congruent spherical shapes purely rotating relative to each other. If the hip articular surfaces are aspherical, then passive hip kinematics cannot be a pure rotation and must include some degree of translation (sliding motion).

The objective of this study was to quantitatively measure the sphericity of the bony articular surfaces of the hip joint - that of the femoral head and of the acetabulum - in order to determine whether or not the arthritic hip has a spherical geometry that allows it to act as the ball and socket joint that it is commonly supposed to be.

Methods

The data were derived from preoperative CT datasets of 16 patients who, prior to this study, underwent computer-assisted hip resurfacing surgery. The CT datasets were acquired with a slice spacing of 1.25 mm and an approximate in-plane spatial resolution of 0.5 mm per pixel. A skilled technical assistant segmented the proximal femur and acetabulum using commercial software (Materialise, Leuven, Belgium) for the purpose of computer-assisted hip resurfacing. These segmentations produced 3D triangulated surface meshes that were used as the primary input for this study.

Each 3D model was further processed to isolate the articular surfaces of the acetabulum and the femoral head. This processing eliminated the fovea, any osteophytes, and all structures distal to the femoral head and surrounding the acetabulum. The resulting data were converted to a set of points, which were the vertices of the STL mesh output by the Materialise software.

The points were fit to an ellipsoid using a least-squares algorithm. The technical difficulty of fitting a partial set of points is not insignificant; if an ordinary least-squares method is used for algebraic optimization, the governing equations can be ambiguous (fitting a general quadric surface such as a paraboloid) and statistically the equations are termed inconsistent. Geometric approaches, such as orthogonal distance regression, are highly sensitive to initialization conditions and to the data's derivations from an ellipsoid. Preliminary uses of these alternative methods on the data were unsatisfactory.

Instead, the adjusted least squares (ALS) algorithm of Markovsky et al. was used. This algorithm relied on an implicit formulation of an ellipsoid that was converted mathematically into a parametric form. A point x lies on an ellipsoid centered at a point c if and only if it satisfies the equation

$$(x - c)^T A(x - c) = 1$$

A point x on an ellipsoid can also be found by (a) scaling a unit sphere to a generic ellipsoid by premultiplying x by a scaling matrix S that is diagonal and positive definite, (b) rotating to an ellipsoid centered on the origin by pre-multiplying by a matrix R, and (c) translating the center by adding a vector c. The Schur decomposition of A can be used to convert the implicit form into a parametric form

$$x = RSP + c$$

where p is a point on the unit sphere.The diagonal scaling matrix S deforms a sphere to an ellipsoid, so its elements can be ordered as

$$S_{11} \geq S_{22} \geq S_{33}$$

In this notation, the major eccentricity of such an ellipse is $E_U = S_{11}/S_{22}$ and the minor eccentricity is

$$E_L = S_{22}/S_{33}.$$

The point data were imported into Matlab (Mathworks, Cambridge, USA). The ALS fit was found using code that Markovsky contributed to Matlab. Custom code was used to decompose the structural matrix in order to find the scaling matrix S.

The values of E_U and E_L were found for each subject, sorted in increasing order, plotted, and grouped into bins for histogram plots. Also, the relationships between femoral and acetabular eccentricity and major and minor eccentricity were examined.

Results

The histograms of eccentricity were produced and analyzed separately for the articular surfaces of the femur and acetabulum, and for major and minor eccentricity (Fig. 1). These histograms were compared to the spherical articular surfaces assumption.

Statistical comparisons were performed using the mean values μ, standard deviations σ, and judging comparisons using a two-sided Student's T-test for significance.

For the femur, the mean major eccentricity had a mean value $\mu = 1.10$ (standard deviation $\sigma = 0.07$) with a minor eccentricity

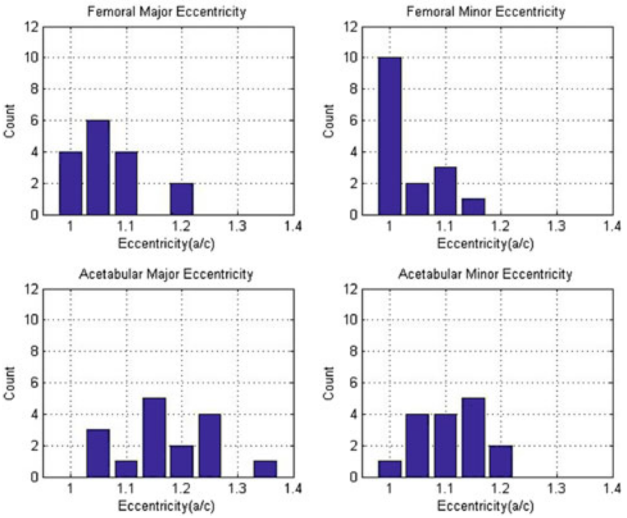

Fig. 1 Histograms indicating the distribution of femoral and acetabular eccentricities

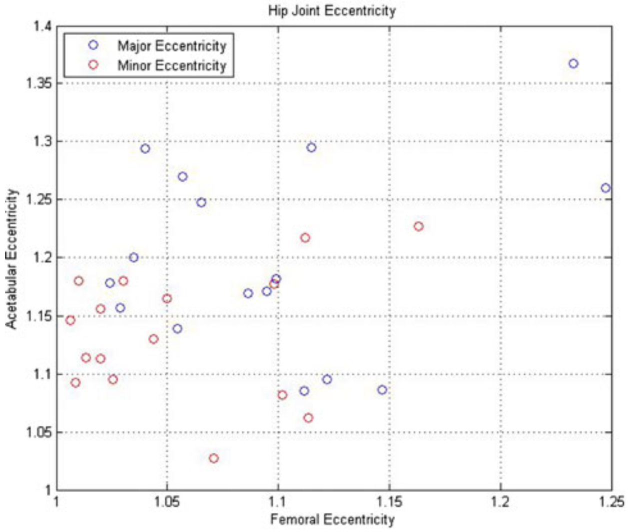

Fig. 2 Scatter plot examining the relationship between femoral and acetabular eccentricities. Each Data point was derived from a femur-acetabulum pair of a single joint

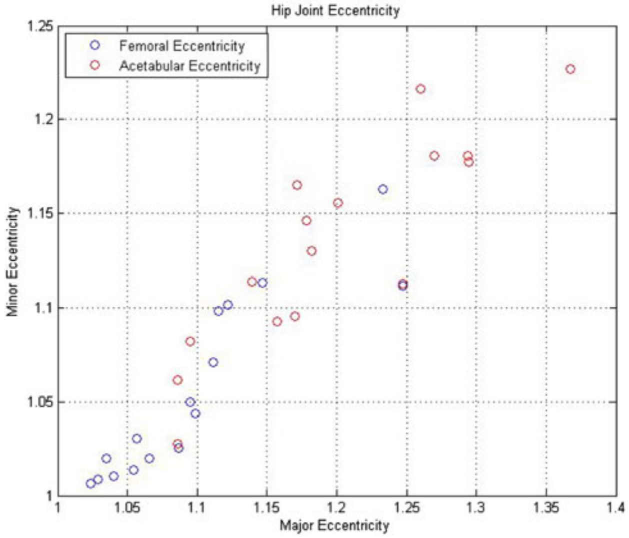

Fig. 3 Scatter plot examining the relationship between major and minor eccentricities. Each data point was derived from a single piece of bony anatomy (femoral head or acetabulum)

mean of $\mu = 1.05$ ($\sigma = 0.05$). For the acetabulum, the mean major eccentricity had a mean value $\mu = 1.20$ ($\sigma = 0.08$) with a minor eccentricity mean of $\mu = 1.14$ ($\sigma = 0.06$). These shapes were highly significantly different from spherical for both the femoral head and acetabular cup $p < 0.001$).

Scatter plots examining the relationships between femoral and acetabular eccentricity and major and minor eccentricity are given in Figs. 2 and 3, respectively. There was a roughly linear relationship between major and minor eccentricity for both the femoral (R2 = 0.833) and acetabular (R2 = 0.767) surfaces of the hip joint. This indicated that the articular surfaces did not take on any general ellipsoidal shape, but rather tended to be ellipsoids from a specific set of shapes governed by the linear relationships. There was no clear relationship between femoral and acetabular eccentricities for either major (R2 = 0.067) or minor (R2 = 0.041) eccentricities.

For the femoral heads, there was no significant difference between the major and minor eccentricities ($p > 0.58$). For the acetabular cups there was a highly significant difference between the major and minor eccentricities ($p < 0.001$). The femoral heads and acetabular cups were, using the major eccentricities, very significantly different ($p < 0.001$).

Geometrically, the femoral heads ranged from nearly spherical to distinctly aspherical. Because the major and minor eccentricities were correlated, the head shapes were best described as oblate spheroids. The acetabular cups had significantly different eccentricities and are best described as general ellipsoids. For this population of arthritic patients the hips were definitely not shaped to form the traditional ball and socket joint.

Conclusion

Given the assumption of a spherical hip joint in both medical texts and gait analysis, the distribution of hip eccentricities (as given in Fig. 1) is perhaps the most significant result. The histogram indicates that not only were none of the articular surfaces spherical, but also that there was no trend towards spherical shape. If spherical articular surfaces were the typical case then one would expect the histograms to show a very positive (or left) skew, with the highest counts at or near the lowest bin with an eccentricity of 1. Figure 1 shows this is not the case, with the plots centering around higher eccentricity values for both the femur and acetabulum. This indicates that the anatomy of the articular surfaces of the hip had a natural variation around some ellipsoidal shape as opposed to a spherical shape.

Because this population had both parts of the hip joints non-spherical in shape, it is highly probable that hip joint motion for these patients was definitely more complex than the idealized mechanical spherical joint that is widely accepted. The incongruence of shape of the bony anatomy of the hip indicates that the motion cannot be purely rotational about an isocenter, but rather some combination of rolling, sliding and translation. While it is presently unknown how closely the motion of the hip resembles that of a spherical joint, it is clear that there are unmodeled behaviors neglected in treating the hip as a spherical joint.

These variations of shape are well known in pediatric orthopedics, especially in Legg-Calvé-Perthes disease. For example, the Stulberg classification and the Mose quotient are both frequently used to classify hips in determining prognoses for patients, although these are 2D shape estimations made from plain radiographs. There has been virtually no work reported on the biomechanical implications of asphericity of the hip. It is known that more spherical 2D classifications have better prognoses, and aspherical or mis-matched shapes have poorer prognoses, but the kinematics of aspherical hips appears to be a fertile area for future work.

The present study was limited by the small number of CT scans available at the time of writing. Work is underway to extend the number of patients, and hence the robustness of the findings. It is important to note that patients selected for hip resurfacing tend to be younger than those selected for total hip replacement, so these findings may not extend to an older population.

In conclusion, if arthritic hips do indeed vary significantly from the ideal spherical shape that is commonly assumed, then this variation may need to be taken into account when planning and performing reconstructive hip surgery for early arthritis.

VURTIGO: Visualization Software for Real-time, Image-guided Interventions

S. Pintilie[1], L. Biswas[1], K. Anderson[1], S. Dick[1,2], G. Wright[1,3], P. Radau[1]
[1]Imaging Research, Sunnybrook Health Sciences Centre, Toronto, Canada
[2]Dept. of Cardiology, Sunnybrook Health Sciences Centre, Toronto, Canada

[3]Dept. of Medical Biophysics, University of Toronto, Toronto, Canada

Keywords Real-time· Cardiovascular· Visualization· Interventional guidance

Purpose

With the development of rapid imaging acquisition protocols, magnetic resonance imaging (MRI) has recently been applied to the guidance of interventional procedures. Guidance by real-time MRI is attractive compared with x-ray fluoroscopy, because MRI has better soft tissue contrast and is not a source of harmful radiation which is a concern for long procedures. The Vurtigo software presented here is designed to enhance the visualization of real-time medical imaging for therapeutic interventions.

Methods

Pipeline Architecture: The communication system between Vurtigo and the MR scanner is composed of several pieces of software that communicate over TCP/IP sockets. The central piece is the Geometry Server that is a synchronized storage location for several types of data including images, image plane orientations, physiological data, and catheter parameters. Multiple clients can send and receive information from the server simultaneously. The scanner control software RTHawk provides a communication link between the MRI scanner and the Geometry Server. Vurtigo and other client applications are connected directly to the Geometry Server and thus are independent of MRI scanner architecture. Vurtigo is able to passively read the scan plane orientation from the server, or drive the location of the scan plane by writing a new orientation.

Vurtigo Design: Vurtigo is written in C++ which has fast execution time when compared to interpreted languages such as Python. Also, adopting C++ frees the developer to utilize a wide selection of C/C++ libraries, such as VTK and the DICOM Toolkit (DCMTK) that have been integrated. For cross-platform compatibility, the Vurtigo project uses the Qt interface and CMake as the build system. The application has been tested on WinXP, Ubuntu Linux and MacOSX Leopard. Another advantage of Vurtigo is its plug-in design which provides a modular and easily extensible framework for developers. The core comprises the interface as well as storage and rendering of a dynamic set of objects, e.g. an MRI volume or a tracked device. Plug-ins add functionality by adding, removing, and processing these objects.

Results

Features: Vurtigo provides the user with a 3D global coordinate system where it can render multiple 2D real-time planes, 3D volumes, 4D (e.g. cine MRI) volumes as well as actively tracked interventional devices such as catheters. Volumes in Vurtigo can be displayed as 3 cut planes or as ray traced composite/MIP renderings. The orientations of the these objects are known from DICOM information so that they can be fused for visualization in approximate spatial alignment. All of the objects may be modified in real-time with rendering updated as required. The latency of sending an MR image through the Geometry Server to Vurtigo was 56 ± 59 msec. Rendering speeds depend on the number of objects being rendered.

Interventional Application: Passive interventional devices, such as catheters, can be difficult to identify on MR images. They can be modified to include one or more coils along their length that can be actively tracked using an MR projection technique. Depending on the number and location of tracked coils, Vurtigo is able to display the device's tip, the tip and direction, or a spline indicating a flexible device's shape. With Vurtigo in drive mode, real-time imaging planes can be attached to a tracked location on or near the device and automatically updated to follow the device's movement. The real-time images show the current state, with their lower resolution and smaller field of view (FOV) compensated by a registered prior volume that provides a high resolution spatial guide. By obtaining the prior volume data and the real-time slices during the same exam, these datasets will be spatially aligned except in cases of subject motion. Rendering the volume, the real-time planes and the interventional device together improves the interventionalist's ability to navigate with respect to the patient's anatomy.

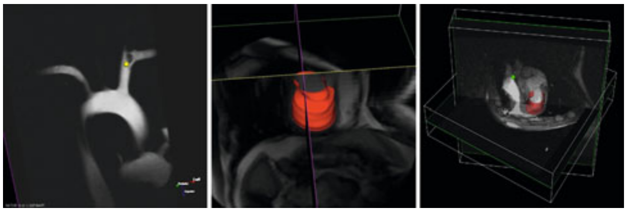

Fig. 1 Vurtigo's interventional capabilities. Left image (**a**): Tracking a catheter tip in an aortic arch phantom. Center image (**b**): Segmented LV endocardial surface from a prior volume. Right image (**c**): tracking a catheter tip in-vivo with two prior volumes as a roadmap

In an initial test a tracking catheter with a single coil at its tip was displayed as a sphere with a 3-Plane rendering of a prior MRI volume acquired from an aortic arch phantom (Fig. 1a). As the catheter was moved through the phantom, the rendered catheter tip moved synchronously in real-time as Vurtigo received streamed catheter coordinates from the Geometry Server.

To demonstrate the capabilities of the system in-vivo we advanced an active catheter in a porcine model to the right ventricle (RV). A surface coil was placed on the body surface and used for large FOV imaging. Two 4D volumes were acquired with the surface coil and rendered in Vurtigo for a static roadmap. One of the volumes was a short axis view of the left ventricle (LV). This volume was segmented and a mesh was produced that displays the endocardial LV surface (Fig. 1b). This is an important feature to support catheter guidance and therapeutic ablation of the surface for ventricular arrhythmias. The rendering produced by Vurtigo is presented in Figure 1c where the green sphere is the catheter tip (coil centroid) and the red surface is the endocardial surface mesh. Intravascular real-time imaging can also be performed in two orthogonal planes, with one offset in front of the catheter to image the vessel cross-section, while the other images the vessel in-plane. The positions of the two planes can be fixed relative to the position of the catheter tip, such that their positions move with the catheter. The visualization can be presented directly to the interventionalist by an MR compatible display beside the scan bed. In cases where fluoroscopic guidance is very difficult due to a total vascular occlusion that blocks contrast diffusion, the value of MRI guidance with Vurtigo visualization would be emphasized.

Conclusion

We have presented an open-source, cross-platform, freely available application that is an advanced visualization system for MRI-guided cardiovascular interventions. The project has reached the stage where it has a vital role in supporting in-vivo experiments. Further information and source code is available from www.vurtigo.ca.