

Intensity-Based Visual Servoing for Instrument and Tissue Tracking in 3D Ultrasound Volumes

Caroline Nadeau, Hongliang Ren, *Member, IEEE*, Alexandre Krupa, *Member, IEEE*, and Pierre Dupont, *Fellow, IEEE*

Abstract—This paper presents a three dimensional ultrasound (3DUS)-based visual servoing technique for intraoperative tracking of the motion of both surgical instruments and tissue targets. In the proposed approach, visual servoing techniques are used to control the position of a virtual ultrasound probe so as to keep a target centered within the virtual probe's field-of-view. Multiple virtual probes can be servoed in parallel to provide simultaneous tracking of instruments and tissue. The technique is developed in the context of robotic beating-heart intracardiac surgery in which the goal of tracking is to both provide guidance to the operator as well as to provide the means to automate the surgical procedure. To deal with the low signal-to-noise ratio (SNR) of the 3DUS volumes, an intensity-based method is proposed that requires no primitive extraction or image segmentation since it directly utilizes the image intensity information as a visual feature. This approach is computationally efficient and can be applied to a wide range of tissue types and medical instruments. This paper presents the first validation of these techniques through offline robot and tissue tracking using actual *in vivo* cardiac volume sequences from a robotic beating-heart surgery.

Note to Practitioners—This paper is motivated by the practical navigation problem in 3DUS-guided beating-heart robotic surgery. A real-time and robust robot guidance system is of great significance in computer-assisted and computer-navigated robotic surgeries. The current position-based ultrasound guidance systems typically involve sequential feature extraction followed by position estimation, while the proposed intensity-based ultrasound virtual servoing can directly track both tissue and instruments from the raw 3D images. Without necessitating segmentation or feature extraction, the proposed approach is efficient enough to be used in guiding the surgical robot to the target using real-time 3DUS. The implementation is based on virtual probe servoing and can be ex-

tended to other similar volumetric image-based servoing problems in industrial applications. The proposed algorithm can be implemented as a plug-in module, which could be easily integrated into existing image-guided robotic surgery platforms.

Index Terms—Instrument tracking, real-time, robotic surgery, three-dimensional ultrasound (3DUS), visual servoing.

I. INTRODUCTION

Three-dimensional (3D) US-guided robotic beating heart surgery is an emerging computer-assisted surgical technique [1], [2]. As shown in Fig. 1, a robot can enter the beating heart through the vasculature to deploy tools and devices. Substantial challenges to surgeon-guided robot navigation and tissue manipulation arise, however, owing to both the 3D nature of the image sets as well as to image noise. These issues can lead to surgeon disorientation and lack of precision. These challenges can be greatly reduced by automatically tracking the robot tip as well as the surgical tissue targets and using this information to create virtual overlays [3]–[5] or to automate certain surgical motions [6].

While exhibiting a lower signal-to-noise ratio (SNR) compared to other imaging modalities, 3DUS does offer the possibility to extract, intra-operatively and in real time, the motion of multiple targets. Owing to the differences in acoustic impedance of tissue and metal instruments [7], however, prior work has focused on either instrument tracking or tissue tracking as summarized next.

As 3DUS technology is emerging recently in practice, there is a modest amount of prior work on this topic and algorithms typically utilize knowledge of the instrument's geometry to perform tracking [8], [9]. Most of these techniques are dedicated to straight tool shaft detection based on linear discrimination algorithms including Principle Component Analysis (PCA) [10], Hough transforms [11], or Radon transforms [8]. In [12] and [13], a new decoupled method was proposed to detect robots of known curvature, however, it was limited to objects of a specific geometry. In contrast, the method presented here is an intensity-based visual servoing approach that does not require *a priori* knowledge of the target shape and can therefore be applied to a wide range of surgical instruments and robots.

Real-time tissue tracking in 3DUS volumes has been investigated in several contexts. For example, tracking of the blood-tissue boundary inside the heart has been performed in order to maintain a given offset distance between an instrument tip and the tissue [14]. Tissue tracking in this approach is facilitated by tracking the instrument axis. Other proposed approaches for US-based tissue tracking predominantly use 2D images. In [15], five feature extraction methods, based on image similarity measure or on contour segmentation, are compared to track an anatomical point corresponding to the center of an artery as a 2DUS probe is moved robotically along the vessel. A segmentation step is also required in [16] in order to compute the 3D position of a kidney stone in a lithotripsy procedure and in [17], where six geometric features are used for the tracking of an anatomic target. Contour segmentation proves to be effective with convex targets, but is generally time-consuming and not adapted for complex organ shapes, nor to the use of noisy 3DUS volumes provided by a matrix-array 3D probe.

Manuscript received April 02, 2014; accepted June 12, 2014. This paper was recommended for publication by Associate Editor W. Sheng and Editor K. Lynch upon evaluation of the reviewers' comments. This work was supported in part by the National Institutes of Health under USA Grant R01HL073647 and Grant R01HL087797, in part by the Singapore Academic Research Grant R-397-000-139-133, and in part by the US-Comp project of the French National Research Agency.

C. Nadeau was with the Lagadic Research Group, IRISA, Inria Rennes-Bretagne Atlantique, Rennes 35042, France. He is now with CEA LIST, Saclay, Orsay 91400, France (e-mail: cnadeau85@gmail.com).

H. Ren was with Boston Children's Hospital and Harvard Medical School, Boston, MA 02115 USA. He is now with the Department of Biomedical Engineering, National University of Singapore, 117575 Singapore (e-mail: ren@nus.edu.sg).

A. Krupa is with the Lagadic Research Group, IRISA, Inria Rennes-Bretagne Atlantique, Rennes 35042, France (e-mail: alexandre.krupa@inria.fr).

P. Dupont is with Boston Children's Hospital and Harvard Medical School, Boston, MA 02115 USA (e-mail: pierre.dupont@childrens.harvard.edu).

This paper has supplementary downloadable multimedia material available at <http://ieeexplore.ieee.org> provided by the authors. The Supplementary Material includes validation of intensity-based visual servoing through offline robot and tissue tracking using ultrasound imaging sequences acquired during *in-vivo* robotic beating-heart surgery in a porcine model. This video shows results of both intensity-based tracking experiments presented in our paper. This material is 12.3 Mo in size.

Digital Object Identifier 10.1109/TASE.2014.2343652




Fig. 1. Schematic of beating heart closure of a Patent Foramen Ovale (PFO) using real-time 3DUS guidance. Robot is comprised of three telescoping sections. Section 1N is for navigation through the jugular vein into the right atrium. Sections 2M and 3M are used to position and orient the robot tip within the right atrium.




Fig. 2. Image-based visual servoing applied to the virtual probe.

In the proposed approach, visual servoing algorithms are utilized to perform both instrument and tissue tracking without requiring any assumption on object geometry or any segmentation step. This is a new method compared with conventional 3DUS-based target motion tracking or imageless motion tracking as surveyed in [12]. Previous work [18] in ultrasound-based visual servoing involved physically moving a robot-mounted probe to maintain a tissue target within the probe's field-of-view and so to compensate for organ motion. In contrast, the approach proposed here is to track the target motion using a virtual probe's field-of-view defined by a 3D region of interest (ROI). The ROI is automatically moved within the 3DUS volume to compensate for both target and probe motions. With respect to our previous work [18], we present here the first results of the intensity-based approach using 3D information provided by a matrix array probe. We are then able to demonstrate the robustness of the approach for actual surgical volume sets that include fast disturbance motion arising from the heartbeat. Moreover, the visual servoing is applied here for tracking not only the cardiac motion but also the insertion of a surgical tool into the heart.

The structure of this paper is as follows. Section II defines the proposed visual servoing technique, it includes subsections describing the use of image intensity as a visual feature and also provides the model for the interaction matrix that relates intensity variations to motion. Section III describes the collection of surgical image data and offline robot and tissue tracking performed with the surgical data. The conclusion appears in Section IV.

II. TRACKING USING VIRTUAL PROBE VISUAL SERVOING

The principle of the image-based visual servoing applied to the virtual probe is illustrated in Fig. 2. A closed-loop approach is considered where at each time t , a set of current visual features $s(t)$ extracted from the image acquired by the visual sensor is compared to the set of desired visual features s^* , characterizing the target.

The control velocity v_p applied to the virtual US probe is computed so as to minimize the visual error $e(t) = s(t) - s^*$. To obtain an

exponential decrease in the visual error, the classical control law [20] is used

$$v_p = -\lambda \widehat{\mathbf{L}}_s^+ (s(t) - s^*) \quad (1)$$

where λ is the controller gain and $\widehat{\mathbf{L}}_s^+$ is the pseudo-inverse of an estimate of the interaction matrix \mathbf{L}_s that relates the variation of the visual features to the velocity v_p .

In the proposed approach, v_p defines the velocity of a virtual 3DUS probe (defined by a 3D ROI) that moves with respect to the actual ultrasound probe so as to maintain the visual features within the virtual probe's field-of-view. In other words, v_p is applied to the ROI within the actual US probe's field-of-view so as to track the motion of the considered target.

A. Image Intensity as Visual Feature

With visual servoing techniques, the efficiency of the control law is highly dependent of the choice of visual features. Typically, geometric features, which can be extracted by various methods based on image similarity measures, contour segmentation or image processing (threshold and morphological operators), are well adapted for controlling the motion of the visual sensor. However, the processing of US images is often more complex than with other imaging modalities due to the speckle generated by the propagation of the US waves in soft tissues. To avoid the computational complexity of such approaches, an alternate technique is adopted here in which the intensity values of the voxels of the 3D ROI are used as the visual features s ,

$$s = (I_{1,1,1}, \dots, I_{u,v,w}, \dots, I_{M,N,L}) \quad (2)$$

where M , N and L are, respectively, the width, the height, and the depth of the ROI, $I_{u,v,w}$ is the intensity, expressed in gray level and (u, v, w) are the 3D voxel coordinates in the US volume. In consequence, the image processing step of the original ultrasound volume reduces to extracting the intensities of the voxels contained in the ROI. Computing ROI velocity from the intensity-based feature vector depends on formulating the interaction matrix \mathbf{L}_s , as described below, that links the variation of these image features to the motion of the virtual 3DUS probe.

B. Formulation of the Interaction Matrix

The interaction matrix involved in the control law (1) can be derived from the time derivative of the visual features vector $s(t)$. The objective is then to determine an analytic form for the time variation of a single voxel intensity in the 3DUS image as a function of the virtual probe velocity v_p , that is to find $\mathbf{L}_{I_{u,v,w}}$ such as

$$\dot{I}_{u,v,w}(t) = \mathbf{L}_{I_{u,v,w}} v_p.$$

At time t , $I_{u,v,w}(t)$ corresponds to the amplitude $I_{US}(\mathbf{x}(t), t)$ of the US echo reflected by the physical point P with coordinates $\mathbf{x}(t) = [x(t); y(t); z(t)]$ in the US probe frame. Given a fixed 3DUS probe, these coordinates are time varying due to the motion of the scene (physiological motion of the patient or motion of a surgical instrument) and the total derivative of the amplitude $I_{US}(\mathbf{x}(t), t)$ is

$$\dot{I}_{US}(\mathbf{x}(t), t) = \left[\frac{\partial I_{US}(\mathbf{x}(t), t)}{\partial \mathbf{x}} \right]^T \dot{\mathbf{x}}(t) + \frac{\partial I_{US}(\mathbf{x}(t), t)}{\partial t}. \quad (3)$$

The interaction matrix estimation relies on the hypothesis of the constancy of the US wave reflection from a physical 3D point. While this is a strong hypothesis since the intensity of the US echo reflected by an organ interface is dependent on the relative orientation of the US probe to this interface, this conservation equation has been empirically

validated in [18], where a 2D intensity-based approach was applied to move a real 2DUS probe interacting with an abdominal phantom.

The time invariance of the intensity of a 3D point is written as $(\partial I_{US}(\mathbf{x}(t), t))/(\partial t) = 0$, which allows the reduction of (3) as follows:

$$\dot{I}_{US}(\mathbf{x}(t), t) \approx \left[\frac{\partial I_{US}(\mathbf{x}(t), t)}{\partial \mathbf{x}} \right]^T \dot{\mathbf{x}}(t). \quad (4)$$

Here, the interaction matrix associated with a 3D point \mathbf{L}_x can be introduced and defined as

$$\dot{\mathbf{x}}(t) = \mathbf{L}_x \mathbf{v}_p. \quad (5)$$

According to Varignon's formula for velocity composition in a solid [19], this interaction matrix is computed with the skew-symmetric matrix of the vector cross product $[\mathbf{x}]_\times$, that only depends of the coordinates of the 3D point

$$\mathbf{L}_x = [\mathbf{I}_3 - [\mathbf{x}]_\times] = \begin{bmatrix} 1 & 0 & 0 & 0 & z & -y \\ 0 & 1 & 0 & -z & 0 & x \\ 0 & 0 & 1 & y & -x & 0 \end{bmatrix}. \quad (6)$$

Combining (4), (5), and (6) results in

$$\dot{I}_{US}(\mathbf{x}(t), t) \approx \nabla I_{u,v,w} [\mathbf{I}_3 - [\mathbf{x}]_\times] \mathbf{v}_p \quad (7)$$

with $\nabla I_{u,v,w} = [\nabla I_x \nabla I_y \nabla I_z]$ defined as the 3D image gradient associated with the voxel (u, v, w) . It is comprised of three components, $\nabla I_x = (\partial I_{u,v,w})/(\partial x)$, $\nabla I_y = (\partial I_{u,v,w})/(\partial y)$ and $\nabla I_z = (\partial I_{u,v,w})/(\partial z)$ that describe the intensity variation of the voxel (u, v, w) along the three orthogonal axes of the US image frame. With a 3D probe, these three components are directly obtained with three 3D derivative filters applied to the current 3DUS image [18].

The interaction matrix $L_{I_{u,v,w}} \in \mathbb{R}^{1 \times 6}$ associated with the visual feature $I_{u,v,w}$ can now be written as

$$L_{I_{u,v,w}} = \begin{bmatrix} \nabla I_x & \nabla I_y & \nabla I_z \\ y \nabla I_z - z \nabla I_y & -x \nabla I_z + z \nabla I_x & x \nabla I_y - y \nabla I_x \end{bmatrix} \quad (8)$$

in which the coordinates of the 3D point are deduced from the coordinates of the voxel (u, v, w) in the ROI and from the intrinsic parameters of the 3D probe

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} sz_x(u - u_0) \\ sz_y(v - v_0) \\ sz_z(w - w_0) \end{pmatrix} \quad (9)$$

with (sz_x, sz_y, sz_z) the size of one voxel and (u_0, v_0, w_0) the coordinates of the ROI center.

The complete interaction matrix \mathbf{L}_s , built by stacking the $L \times M \times N$ matrices $L_{I_{u,v,w}}$, can then be computed at each iteration of the algorithm using information entirely contained in the current 3DUS image

$$\mathbf{L}_s = \begin{pmatrix} L_{I_{1,1,1}} \\ \vdots \\ L_{I_{M,N,L}} \end{pmatrix}. \quad (10)$$

III. EXPERIMENTS

To evaluate the tracking algorithm, 3DUS image sequences were collected at the maximum rate of the scanner (25 Hz) during a robotic beating-heart Patent Foramen Ovale (PFO) closure procedure in a porcine model, as shown schematically in Fig. 1 (see [1] and [2]). As




Fig. 3. Beating-heart intracardiac surgery with 3DUS probe manually positioned on the epicardial surface.




Fig. 4. Visualization of the tracking results with a simulator developed with the VTK library [22]. (left) The US volume loaded in the simulator is displayed with two orthogonal planes. (right) Two different ROI, displayed in orange in purple, can be tracked using the proposed visual servoing.

described below, these sequences were then imported into a software simulation program that implemented the proposed tracking algorithm at the same 25 Hz rate used for data collection. Yorkshire pigs (50–70 kg) were used in these experiments to approximate human procedures and their use and treatment followed a protocol approved by the Animal Care and Use Committee of the host institution (BCH).

During the procedure, a concentric tube robot [21] comprised of three telescoping curved sections was inserted into the jugular vein and navigated to the right atrium of the heart under 3DUS guidance. The robot sections include a long proximal section for Navigation through the jugular vein (segment 1N in Fig. 1) and two distal Manipulation sections (segments 2M and 3M in Fig. 1) for tissue manipulation and device deployment inside the atrium.

As shown in Fig. 3, the matrix-array 3D probe was positioned epicardially on the porcine heart to image the robotic instrument tip as well as the tissue inside the right atrium. Note that the position of the porcine lung precluded the use of a transesophageal probe as would be employed in a human procedure.

The 3D image acquisition system was an IE33 system (Philips Medical Systems) which was used to acquire US volumes at a frame rate of 25 Hz for a duration of 10 s. Standard settings of the imaging parameters were used during image generation, including 50% overall gain, 50% compression rate, high density scan line spacing, 6 cm image depth, and 0 dB power level. The resulting voxels have anisotropic spacing of 0.290, 0.342, and 0.456 mm in the x , y , and z directions, respectively.

A tracking and rendering software package was implemented as shown in Fig. 4 that could load 3DUS imaging sequences and perform visual-servoing-based tracking on user-defined ROIs. The surgical volume sequence consisted of 250 US volumes that could be played back at the real-time 25 Hz rate in the simulator. All image processing to execute the visual-servoing-based tissue tracking is performed at the real-time rate (25 Hz) corresponding to the acquisition time of a US volume by the 3D probe.

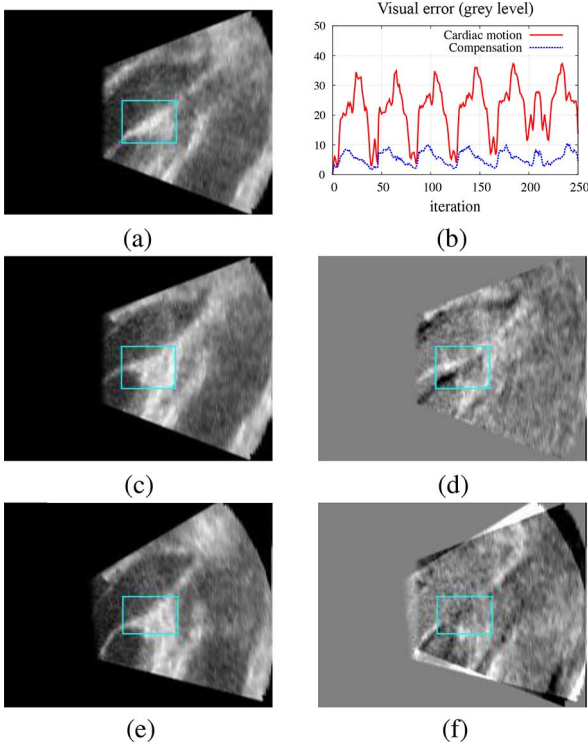


Fig. 5. Tracking results for cardiac tissue. (a) Desired anatomic target is selected by defining the ROI. (b) Visual error with and without compensation. (c) Uncompensated cardiac motion leads to target tissue moving out of ROI. (d) Difference image without compensation. (e) With compensation target tissue remains within ROI. (f) Difference image with compensation.

Heart motion is visualized using two orthogonal images (with green and red contours as shown in Fig. 4, left) as well as an additional view displaying the virtual probe controlled by the proposed intensity-based visual servoing. The first image volume of the data set is used to graphically allow the user to delineate the desired anatomic structure by a bounding box (see Fig. 4, right) that is then used as the 3D ROI by the tracking control law.

In the current implementation, the tracking law assumes rigid motions. While this applies directly to the robot tip, the ROI for tissue tracking must be selected to be sufficiently small so that its motion is approximately rigid. For the tested data set, this corresponds to a ROI size of $40 \times 25 \times 10$ voxels for tissue tracking and of $25 \times 25 \times 10$ voxels for robot tracking.

A. Tissue Target Tracking

For PFO closure, a device must be inserted into the septal ridge above the PFO channel. In order to track the anatomic motion of this ridge, we define it as our tissue target and include it in an ROI. This ROI is a volume of about $1.2 \text{ cm} \times 0.9 \text{ cm} \times 0.5 \text{ cm}$ that is manually defined in the first US image of the acquired sequence (see cyan rectangle in Fig. 5(a), which corresponds to a central slice of the ROI). The visual features vector involved in the control law is then constituted of all the voxel intensities belonging to the ROI. During the subsequent tracking, the velocity computed by visual servoing (1) is applied to this ROI. The gain λ is manually adjusted to ensure the best dynamics of the system, while keeping it stable and the visual control law is applied several times on the same US volume until a new one is available. This strategy ensures a good tracking accuracy and is possible since the control loop period is about 10 ms, while the image acquisition time is 40 ms. In this example, the motion compensated through tracking is




Fig. 6. Robot tracking results. (a) Without compensation, robot tip moves wrt the ROI (displayed in cyan). (b) With compensation, robot tip is kept centered in the ROI. (c) ROI position and orientation wrt the initial image frame when the robot tip is tracked. The $\theta \mathbf{u}$ representation of orientation is considered, where $\mathbf{u} = (u_x \ u_y \ u_z)^\top$ is a unit vector representing the rotation axis and θ is the rotation angle.

predominantly due to beating of the heart. Fig. 5 provides a visual validation of the tracking performance by comparing the ultrasound images of an ROI fixed with respect to the 3D volume [Fig. 5(c)] with an ROI that is updated by tracking compensation [Fig. 5(e)].

The difference image, I_{diff} , between the current image I and the desired one I^* is defined by

$$I_{\text{diff}} = \frac{(I - I^*) + 255}{2}$$

and depicted for the cases of no compensation [Fig. 5(d)] and compensation [Fig. 5(f)]. The effectiveness of the compensation can be seen by comparing the uniformly gray ROI in Fig. 5(f) that results from compensation with the strong gradients occurring in the uncompensated difference image ROI of Fig. 5(d). Image tracking error can also be defined quantitatively as the Euclidean norm of the visual vector normalized by the number of voxels N_{vox} in the considered ROI, as given in (11) and as appears in Fig. 5(b). It can be seen that compensation reduces the norm of the error vector by a factor of about 3.5

$$C = \sqrt{\frac{(\mathbf{s} - \mathbf{s}^*)^\top (\mathbf{s} - \mathbf{s}^*)}{N_{\text{vox}}}}. \quad (11)$$

B. Robot Motion Tracking

Experiments were also performed to evaluate the intensity-based visual servoing for tracking robot motion during surgery. In this case, the desired ROI was defined to include the robot tip and once again the visual features vector was constructed with the intensities of every voxel belonging to this ROI (see cyan rectangle in Fig. 6). The image volume sequence analyzed corresponded to the operator moving the tip inside

the right atrium of the beating porcine heart. The motion corresponds to a forward and backward translation along the direction of the tool shaft.

Robot tip motion is automatically tracked by the proposed intensity approach during its movement, as shown in Fig. 6. In order to most effectively visualize the performance of the tracking algorithm, the 2DUS slice passing through the center of the defined ROI is displayed for both noncompensated [Fig. 6(a)] and compensated [Fig. 6(b)] robot motion. When the tracking is performed, the robot tip remains centered in the ROI in each of the acquired 3DUS volumes.

Fig. 6(c) depicts motion of the ROI with respect to (wrt) a fixed frame attached to the initial US volume during compensation. This computed motion corresponds to the estimated translational motion of insertion (from iteration 20 to iteration 140) and removal (from iteration 140 to iteration 250) of the robot. Indeed, the detected displacement axis, shown as the red dotted line in Fig. 6(a), can be seen to be approximately given by the the 3D vector $(1 - 1 0)^T$.

IV. CONCLUSION

The new direct visual servoing method proposed in this paper provides a real-time estimation and compensation of rigid motions in 3DUS volume sequences. Its effectiveness has been illustrated using data from an actual robotic beating-heart surgery. These experiments provide validation of the proposed approach both for tracking the rigid motion of robotic tools and also for tracking the quasi-rigid motion of heart tissue.

As a next step, the methods will be extended to investigate non-rigid tissue motions using a deformable 3D grid instead of a rigid ROI. Deformation will be accommodated using parametric model based, for example, on 3D thin-plate splines. The objective will be to estimate the variation of the deformable model parameters directly from intensity difference as measured in successive volumes. In future applications, the proposed target tracking approach will be employed to automatically reach and track a target of interest inside the beating heart.

REFERENCES

- [1] A. Gosline, N. Vasilyev, E. Butler, C. Folk, A. Cohen, R. Chen, N. Lang, P. del Nido, and P. Dupont, "Percutaneous intracardiac beating-heart surgery using metal MEMS tissue approximation tools," *Int. J. Robot. Res.*, vol. 31, pp. 1081–1093, 2012.
- [2] N. Vasilyev, A. Gosline, E. B. N. Lang, P. Codd, H. Yamauchi, E. Feins, C. Folk, A. Cohen, R. Chen, P. del Nido, and P. Dupont, "Percutaneous steerable robotic tool delivery platform and metal MEMS device for tissue manipulation and approximation: Initial experience with closure of patent foramen ovale," *Circulation: Cardiovascular Interventions*, vol. 6, pp. 468–475, 2013.
- [3] T. Peters, "Image-guidance for surgical procedures," *Phys. Med. Biol.*, vol. 51, no. 14, p. R505, 2006.
- [4] R. H. Taylor and D. Stoianovici, "Medical robotics in computer-integrated surgery," *IEEE Trans. Robot. Autom.*, vol. 19, no. 5, pp. 765–781, Oct. 2003.
- [5] N. Hungr, M. Baumann, J.-A. Long, and J. Troccaz, "A 3-D ultrasound robotic prostate brachytherapy system with prostate motion tracking," *IEEE Trans. Robot.*, vol. 28, no. 6, pp. 1382–1397, Dec. 2012.
- [6] N. Padoy and G. Hager, "Human-machine collaborative surgery using learned models," in *Proc. IEEE Int. Conf. Robot. Autom.*, vol. 2011, pp. 5285–5292.
- [7] J. Huang, J. Triedman, N. Vasilyev, Y. Suematsu, R. Cleveland, and P. Dupont, "Imaging artifacts of medical instruments in ultrasound-guided interventions," *J. Ultrasound in Med.*, vol. 26, pp. 1303–1322, 2007.
- [8] P. Novotny, J. Stoll, N. Vasilyev, P. Del Nido, P. Dupont, T. Zickler, and R. Howe, "GPU based real-time instrument tracking with three dimensional ultrasound," *Med. Image Anal.*, vol. 11, no. 5, pp. 458–464, 2007.
- [9] S. Yuen, D. Kettler, P. Novotny, R. Plowes, and R. Howe, "Robotic motion compensation for beating heart intracardiac surgery," *Int. J. Robot. Res.*, vol. 28, no. 10, pp. 1355–1372, 2009.
- [10] J. Stoll, H. Ren, and P. Dupont, "Passive markers for tracking surgical instruments in real-time 3D ultrasound imaging," *IEEE Trans. Med. Imaging*, vol. 31, pp. 563–575, Mar. 2012.
- [11] M. Ding, H. N. Cardinal, W. Guan, and A. Fenster, S. K. Mun, Ed., "Automatic needle segmentation in 3D ultrasound images," *SPIE*, vol. 4681, no. 1, pp. 65–76, 2002.
- [12] H. Ren, N. V. Vasilyev, and P. E. Dupont, "Detection of curved robots using 3D ultrasound," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2011, pp. 2083–2089.
- [13] H. Ren and P. Dupont, "Tubular enhanced geodesic active contours for continuum robot detection using 3D ultrasound," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2012, pp. 2907–2912.
- [14] S. G. Yuen, N. V. Vasilyev, P. J. del Nido, and R. D. Howe, "Robotic tissue tracking for beating heart mitral valve surgery," *Med. Image Anal.*, vol. 17, pp. 1236–1242, 2013.
- [15] P. Abolmaesumi, S. Salcudean, W. Zhu, M. Sirouspour, and S. DiMaio, "Image-guided control of a robot for medical ultrasound," *IEEE Trans. Robot.*, vol. 18, no. 1, pp. 11–23, Feb. 2002.
- [16] N. Koizumi, J. Seo, D. Lee, T. Funamoto, Y. Itagaki, A. Nomiya, A. Ishikawa, H. Tsukihara, K. Yoshinaka, N. Sugita, H. Homma, Y. Matsumoto, and M. Mitsuishi, "Construction methodology for NIUTS bed servoing system for body targets," *J. Robot. Mechatronics*, vol. 25, no. 6, pp. 1088–1096, 2013.
- [17] R. Mebarki, A. Krupa, and F. Chaumette, "2D ultrasound probe complete guidance by visual servoing using image moments," *IEEE Trans. Robot.*, vol. 26, no. 2, pp. 296–306, Apr. 2010.
- [18] C. Nadeau and A. Krupa, "Intensity-based ultrasound visual servoing: Modeling and validation with 2D and 3D probes," *IEEE Trans. Robot.*, vol. 29, no. 4, pp. 1003–1015, Aug. 2013.
- [19] M. Hadwiger, J. M. Kniss, C. Rezk-Salama, D. Weiskopf, and K. Engel, *Real-Time Volume Graphics*, pp. 112–114, 2006.
- [20] B. Espiau, F. Chaumette, and P. Rives, "A new approach to visual servoing in robotics," *IEEE Trans. Robot.*, vol. 8, no. 3, pp. 313–326, Jun. 1992.
- [21] P. Dupont, J. Lock, B. Itkowitz, and E. Butler, "Design and control of concentric-tube robots," *IEEE Trans. Robot.*, vol. 26, no. 2, pp. 209–225, Apr. 2010.
- [22] W. Schroeder, K. Martin, and B. Lorensen, *The Visualization Toolkit: An Object Oriented Approach to 3D Graphics*, 3rd ed. Clifton Park, NY, USA: Kitware, Inc., 2003, ISBN-1930934076.