NIH Public Access
Author Manuscript
Rep U S. Author manuscript; available in PMC 2014 December 31.

Published in final edited form as:
Rep U S. 2013 December 31; 2012: 3792–3797. doi:10.1109/IROS.2012.6386009.

Space-Time Localization and Registration on the Beating Heart

Nathan A. Wood[1], **Kevin Waugh**[2], **Tian Yu Tommy Liu**[1], **Marco A. Zenati**[3], and **Cameron N. Riviere**[1]

Nathan A. Wood: nwood@andrew.cmu.edu; Kevin Waugh: waugh@cs.cmu.edu; Tian Yu Tommy Liu: tianyul@andrew.cmu.edu; Cameron N. Riviere: camr@ri.cmu.edu

[1]The Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

[2]Computer Science Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

[3]BHS Department of Cardiothoracic Surgery, Harvard Medical School, West Roxbury, Massachusetts 02132

Abstract

This paper presents a framework for localizing a miniature epicardial crawling robot, HeartLander, on the beating heart using only 6-degree-of-freedom position measurements from an electromagnetic position tracker and a dynamic surface model of the heart. Using only this information, motion and observation models of the system are developed such that a particle filter can accurately estimate not only the location of the robot on the surface of the heart, but also the pose of the heart in the world coordinate frame as well as the current physiological phase of the heart. The presented framework is then demonstrated in simulation on a dynamic 3-D model of the human heart and a robot motion model which accurately mimics the behavior of the HeartLander robot.

I. Introduction

Due largely to the improvement in patient outcomes, minimally invasive cardiac therapies have become increasingly appealing in comparison to standard invasive cardiac surgeries. Although these methods provide many advantages for the patient, they possess a number of challenges due to the types of instruments and access points used. With no direct line of sight to the operation field, real-time medical imaging technologies, including magnetic resonance imaging (MRI) [1], fluoroscopy [2], and ultrasound [3], are often used to provide visual feedback for navigation.

Image guided surgery, which uses pre-operative medical images to provide a virtual view of the operating site, is also often used for visual feedback. In this framework the surgical device is often tracked using an electromagnetic position sensor, localized in and registered to the virtual model, and displayed in the visualization [4]–[6]. These methods, while possessing considerable power, can often be negatively affected by the dynamic nature of the heart. While the model of the heart is treated as a static body, deformations of almost 30 mm occur due to the physiological cycles of heartbeat and respiration [7]. For tools and surgical robots that move relatively freely in the cardiothoracic cavity, or inside the heart, accounting for this periodic motion poses a significant challenge due to the changing contact constraints. However, for a robot which adheres to the surface of the heart, this motion may be leveraged to improve localization and registration.

HeartLander, shown in Fig. 1, provides therapies to the heart by adhering to and moving over the epicardial surface. The robot is a miniature inchworm-style robot that adheres to the epicardial surface, inside the pericardium, using suction, and moves by extending and retracting drive wires connecting the two feet while alternating suction. Access to the heart

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

is gained through a subxiphoid skin incision and a small incision in the pericardium at the apex of the heart. Previous work has successfully demonstrated the ability to access the pericardium, move over the surface of the heart, and reach targets accurately [8].

The current methods used for localizing HeartLander on the surface of the heart use several approximations which limit accuracy. Position measurements of the robot come from a 6-degree-of-freedom electromagnetic tracking sensor (microBIRD, Ascension Technology) embedded in the front foot of the robot. The surface of the heart, as previously mentioned, is a dynamic environment which undergoes periodic deformation due to both the heartbeat and respiration cycles. Because the system currently uses a static model of the heart generated from pre-operative CT images, these deformations are treated as noise and filtered out to estimate of the mean location of the robot [8]. This mean location is treated as the position of the robot on the surface of the static heart model. Registration between the map frame and measurement frame is found using markers placed on the chest wall which are identified in each frame. Transformations between the frames are then found using least squares methods.

Although HeartLander has shown considerable success in live animal testing, there remains the possibility of improving the accuracy of robot positioning. If, instead of rejecting the periodic deformations of the surface of the heart as noise, these motions, which vary over the surface of the heart, are treated as features which yield information about the current robot position on the heart, we may be able to improve localization accuracy. Assuming that we possess a map which fully defines how the surface of the heart moves through the physiological phases, this work presents a method for using the motion of the surface to localize on the surface. The concept was demonstrated in 2-D in [9]. In the present work, it is implemented and demonstrated in simulation on a 3-D surface model of the heart.

II. Related Work

Methods for registering and localizing in and around the heart are highly dependent upon the instrument or robot being used. Tully *et al.* use an electromagnetic tracker located at the distal end of a highly articulated snake-like robot to first the shape of the robot [4], and then use inequality constrained kalman filtering to correct registration and localization parameters when the model is found to violate geometric constraints, such as intersecting the pre-operative surface model [5].

Therapies which target the endocardium using intracardiac echo (ICE) catheters use ultrasound images to localize. One method generates a point cloud by extracting heart surface points from 2-dimensional ultrasound images at the catheter tip by rotating the catheter about its longitudinal axis. After sufficient points have been collected the point cloud is registered to a pre-operative model of the left atrium using an iterative closest point (ICP) method [10]. Another method uses a particle filter to recursively estimate the pose of the ICE catheter in the left atrium [6]. In this work the probability of a catheter tip pose is calculated by comparing virtual ultrasound images constructed from a pre-operative model with the actual ultrasound images. The predicted catheter pose is then determined as the weighted sum of the particles, and the registration parameters are then calculated using this pose estimate and measurements from an electromagnetic position tracker.

Several factors differentiate our current work from those presented. First, and most obvious, is that they all work from a static map of the environment so that the periodic motion of the heart is either rejected as noise or accounted for in the registration parameters. While the use of dynamic maps of the entire cardiothoracic cavity may be infeasible at the present time, considerable work has been done on generating models of the deforming heart including representations using splines [11], [12], finite element models [13], and statistical models [14].

A second difference is the measurements, or constraints, available in each system. In [4] the entire shape of the robot is used to constrain the robot with respect to the map. In [6], [10] two separate measurements are used to localize and estimate registration parameters. First the ultrasound images are used to localize within the map, then the position measurements from the electromagnetic tracker are used to register the model to the real world.

In our case, the only measurement available is 6-DOF pose from the electromagnetic tracker. Using only this measurement, along with the constraint that the robot must be on the surface of the heart, we estimate localization and registration parameters simultaneously.

III. Methods

A. Heart Surface Model

The work presented relies on possessing complete maps which describe the periodic motion of a surface. For our purposes, a map of a surface takes the form:

$$M=[\overrightarrow{x}(\varphi),\overrightarrow{n}(\varphi)], \quad (1)$$

where $\phi \in (0, 1]$ is the phase, $\vec{x}$ are the Cartesian coordinates in map frame, and $\vec{n}$ are the surface normals.

In order to develop and test the following methods used for localizing on such surfaces, a surface model of a beating heart was derived from a spline-based model of the epicardium of a human subject [11]. The heart surface model is shown at four different phases of the cardiac cycle in Fig. 2.

B. Simulated System

The simulated system is represented by a given map, M, and the following state vector:

$$s_t=\begin{bmatrix} \overrightarrow{x}_r^m(\varphi) & \overrightarrow{q}_r^m(\varphi) & \varphi & \overrightarrow{x}_m^w & \overrightarrow{q}_m^w \end{bmatrix}^T, \quad (2)$$

where $\overrightarrow{x}_r^m(\varphi)$ is the location of the robot on the surface in map coordinates, $\overrightarrow{q}_r^m(\varphi)$, is the quaternion orientation of the robot in map coordinates, ϕ is the current phase, $\overrightarrow{x}_m^w$, is the location of the map in world coordinates, and $\overrightarrow{q}_m^w$ is the quaternion orientation of the map in world coordinates. Using this representation, the pose of the robot in the world frame is then:

$$\overrightarrow{x}_r^w(\varphi)=\overrightarrow{q}_m^w\overrightarrow{x}_r^m(\varphi)\overrightarrow{q}_m^{w^{-1}}+\overrightarrow{x}_m^w \quad (3)$$

$$\overrightarrow{q}_r^w(\varphi)=\overrightarrow{q}_m^w\overrightarrow{q}_r^m(\varphi) \quad (4)$$

The phase of the system is advanced by:

$$\varphi_t=\varphi_{t-1}=\omega dt, \quad (5)$$

where the velocity, ω, is assumed to be constant. The location and orientation of the map in world coordinates are also assumed to be constant.

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

Control inputs to the robot, $u_t = (\theta_t, d_t)$, rotates the robot about the surface normal through an angle, θ_t, and moves the robot along the surface a distance d_t in the robot's x-direction. Using this framework, we wish to estimate the current state vector, s_t, using a particle filter.

C. Particle Filter Overview

This section gives a brief overview of the particle filter algorithm implemented in this work. A more in-depth treatment of the algorithm and related topics an be found in [15]. The particle filter is a nonparametric Bayes filter which represents the posterior distribution by a set of random samples drawn from the posterior. These samples of the posterior, or particles, are represented as:

$$S_t = [s_t^1, s_t^2, \ldots, s_t^n], \quad (6)$$

where each particle, s_t^i, is a hypothesis of the true state of the system at time t. The set of particles, S_t, then approximates the belief state, where the likelihood for each state hypothesis is proportional to its Bayes posterior:

$$bel(s_t) \sim p(s_t | z_{1:t}, u_{1:t}), \quad (7)$$

where $z_{1:t}$ are all past measurements, and $u_{1:t}$ are all past control inputs. Because of this distribution the more probable regions of the state space will be more densely populated by particles. The particle filter, being a Bayes filter, recursively constructs the belief of the current state from the previous belief state. The general Bayes filter consists of two steps: prediction and correction. In the prediction step the predicted belief, $\overline{bel}(s_t)$, is determined by combining the transition probability using the current control input, u_t, and the prior belief, $bel(s_{t-1})$.

$$\overline{bel}(s_t) = \sum_{s_{t-1}} p(s_t | u_t, s_{t-1}) \, bel(s_{t-1}) \quad (8)$$

The predicted belief is then updated by incorporating the current measurement, z_t.

$$bel(s_t) = \eta p(z_t | s_t) \overline{bel}(s_t) \quad (9)$$

Because the particle filter represents the belief distribution by random samples from the posterior it differs slightly in form from the general Bayes filter, yet it is based on similar principles. In the prediction step, each particle, s_{t-1}^i, $i \in (1, N)$, where N is the number of particles, is advanced using the motion model. The advanced particle is drawn from the state transition distribution.

$$s_i^t \sim p(s_t | u_t, s_{t-1}^i) \quad (10)$$

The set of particles generated by incorporating the state transitions is the filter's representation of the predicted belief distribution, $\overline{bel}(s_t)$. In order to incorporate the current measurement, z_t, importance factors, or weights w_t^i, for each particle is calculated as:

$$w_t^i = w_{t-1}^i p(z_t | s_t^i) \quad (11)$$

NIH-PA Author Manuscript NIH-PA Author Manuscript NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

The updated particle set, $\overline{S}_t$, consists of each particle along with their respective weight. This set of particles represents the predicted distribution $\overline{bel}(s_t)$.

$$\overline{S}_t = \overline{S}_t + \langle s_t^i, w_t^i \rangle \quad (12)$$

In order to force the particle set from the predicted distribution, $\overline{bel}(s_t)$, to the posterior distribution, $bel(s_t)$, the algorithm uses importance sampling. In importance sampling, N particles are drawn with replacement from the predicted set, where the probability of of drawing each sample is proportional to its weight, w_t^i. After resampling each of the particles weight is set to one, and the particle set S_t is distributed according to the posterior, $bel(s_t)$.

D. Particle Filter Implementation

1) Particle Initialization—In order to reduce the size of the space spanned by the state vector, which reduces the number of particles to sufficiently cover this space, it is assumed that at initialization an estimate of the position and orientation of the heart, as well as a single measurement of the pose of the robot in the world frame, are available. Using this information, along with confidence bounds on the position and orientation of the heart, a fixed number of particles, N, are generated such that each particle's registration parameters lie within the confidence bounds of the initial estimate and the location and orientation of the robot on the surface of the heart, when transformed to world coordinates, would produce the initial measurement.

2) Motion Model—Incorporation of the control inputs in the state transition distribution is achieved through use of a robot motion model. As previously described, the control input $u_t = (\theta_t, d_t)$, rotates the robot about the surface normal through an angle, θ_t, and moves the robot along the surface a distance d_t in the robot's x-direction. In order to return a sample from the distribution $p(s_t | u_t, s_{t-1}^i)$, noise is injected into the motion model. The angle and distance each particle moves on the surface is:

$$\theta_t^i = \theta_t + \delta_\theta^i, \quad (13)$$

$$d_t^i = d_t + \delta_u^i, \quad (14)$$

where δ_θ^i and δ_d^i are random samples drawn from $\mathcal{N}(0, \sigma_\theta^2)$ and $\mathcal{N}(0, \sigma_d^2)$ respectively. Also, the same method is used to advance the phase of each particle.

$$\varphi_t^i = \varphi_{t-1}^i + (\omega + \delta_\omega^i) dt, \quad (15)$$

where δ_ω^i is a random sample from $\mathcal{N}(0, \sigma_\omega^2)$. The velocity of the cardiac phase, ϕ, is assumed to be known as this information can be measured in the clinical setting using an electrocardiogram (ECG).

The registration parameters, $\overrightarrow{x}_m^w$ and $\overrightarrow{q}_m^m$, although stationary, also have noise added to ensure coverage of the space. The translational component of the registration parameter is advanced by:

$$\overrightarrow{x}_{m_t}^{w^i} = \overrightarrow{x}_{m_{t-1}}^{w^i} + \delta_{x_m^w}^i \overrightarrow{v}_x^i, \quad (16)$$

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

where $\overrightarrow{v}^{i}_{x}$ is a unit vector uniformly sampled from $\mathbb{R}^3$, and $\delta^{i}_{x^{w}_{m}}$ is a random sample drawn from $\mathcal{N}(0, \sigma^{2}_{x^{w}_{m}})$. While the rotational registration component is advanced by:

$$\overrightarrow{q}^{w^i}_{m_t} = \overrightarrow{q}^{i}_{v_q} \overrightarrow{q}^{w^i}_{m_{t-1}}, \quad (17)$$

$$\overrightarrow{q}^{i}_{v_q} = \begin{bmatrix} \cos\frac{\alpha^i}{2} \\ \sin\frac{\alpha^i}{2} \overrightarrow{v}^{i}_{q} \end{bmatrix}, \quad (18)$$

where $\overrightarrow{v}^{i}_{q}$ is a unit vector uniformly sampled from $\mathbb{R}^3$, and α^i is a random sample drawn from $\mathcal{N}(0, \sigma^{2}_{\alpha})$.

3) Measurement Model—The measurement assumed in this work is the 6-DOF pose of the robot in world frame. This measurement, z_t, is composed of a position vector, $\vec{x_z}$, and a quaternion orientation, $\vec{q_z}$, which are constructed from the ground truth robot state vector as in (3)–(4) and corrupted with Gaussian noise. The weight of each particle is calculated as the probability of the measurement given the particle state. The particle weight is given by:

$$w^{i}_{t} = w^{i}_{t-1} p(z_t | s^{i}_{t}) = w^{i}_{t-1} p(\overrightarrow{x}_z | s^{i}_{t}) p(\overrightarrow{q}_z | s^{i}_{t}) \quad (19)$$

$$p(\overrightarrow{x}_z | s^{i}_{t}) = \eta_x \exp\left(\frac{\left|\left|\overrightarrow{x}_z - \overrightarrow{x}^{w^i}_{r}\right|\right|}{2\sigma^{2}_{x}}\right) \quad (20)$$

$$p(\overrightarrow{q}_z | s^{i}_{t}) = \eta_q \exp\left(\frac{\beta^i}{2\sigma^{2}_{\beta}}\right) \quad (21)$$

$$\beta^i = 2 \arccos\left(\left|\overrightarrow{q}_z \cdot \overrightarrow{q}^{w^i}_{r}\right|\right), \quad (22)$$

where β^i is the angle between the two rotations and a proper distance metric on $SO(3)$ [16].

4) Resampling—The resampling procedure implemented in the system differs in two ways from the general resampling procedure previously mentioned. First, in order to decrease the risk of losing particle diversity resampling only occurs when the variance of the particle weights is sufficiently large, namely when the following holds:

$$\frac{var(w_t)}{\overline{w}_t} > \gamma, \quad (23)$$

where $\bar{w_t}$ is the mean particle weight, and γ is a threshold value. The second strategy used to reduce the sampling error is known as low variance sampling [15]. Instead of just drawing independent samples based on each particle's weight, this method ensures that any particle which has a weight greater than $\frac{1}{N}$, where N is the number of particles, is guaranteed to survive.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

NIH-PA Author Manuscript

IV. Experiments

In order to test the feasibility of using the previously described particle filter framework to estimate localization and registration parameters 100 trials were run in the simulated system. For the trials the heart rate was set to 1.3 Hz. For each simulation a ground truth particle was randomly placed on the surface of the heart with randomly generated registration parameters. At each timestep the ground truth particle was commanded a random motion input such that the maximum robot velocity and path curvature were within the normal HeartLander operating parameters. Simulations were run for 500 time steps, or approximately 15 cardiac cycles.

In the experiment the number of particles was set to be N=2000. The particles were initialized such that each particles registration estimate was within 25 mm and 30° of the initialization estimate, and the phases were uniformly sampled. The motion and observation models as specified in section III-D were used with parameters set to: $\sigma_\theta^2=0.002\,\text{radians}, \sigma_d^2=0.03\,\text{mm}, \sigma_\omega^2=0.01\,\text{radians}$ per second, $\sigma_{x_m^w}^2=0.25\,\text{mm},\ \sigma_\alpha^2=0.002\,\text{radians}$. The observation model parameters are set to: η_x = η_q = 1, $\sigma_x^2=1.4\,\text{mm},\ \sigma_{beta}^2=0.009\,\text{radians}$. The variance resampling threshold was set to be γ = 0.01.

Fig. 3 illustrates the progression of the particle filter through a typical run. The ground truth location of the robot is shown by the large black dot in each image. Fig. 3(a) shows the initialization of the filter, and demonstrates the randomized coverage of our state space. Fig. 3(b)–(g) shows the particle filter after from iterations 10–90, as the particles begins to form clusters around states with high likelihoods, and the most likely clusters become more dense. The reason a few separate clusters form is due to certain geometric symmetries in the surface. Certain distinct regions of the surface under different registration and phase can produce similar world-frame motion, and keeping a diverse set of hypothesis is the correct thing to do for the particle filter. Fig. 3(h) shows the particle filter after 115 iterations. At this point, we have gathered enough observations to greatly reduce the uncertainly in our location, and only a single cluster remain, achieving convergence. From this point onward, the particle filter retains the correct localization until the end of the run.

To make a prediction for the map-frame location of the ground truth using the particle filter at any time during the run, we take the weighted average of all particles. The reason we use an average of particles rather than simply outputting the single highest-weighted particle is for stability: as the weight on each particle is updated after every observation, the highest-weighted particle is likely to change often, and our localization output will be rather noisy.

We recorded the state estimation error made by the particle filter prediction output at each iteration of a run, and averaged the error over 100 runs to obtain a statistically significant result. Each of the 100 runs are completely independent with independent initializations. The results are shown in Fig 4(a)–(d). The error means and standard deviations observed at the end of the trials were phase error of 0.01 ± 0.01, position error of 0.40 ± 0.21 mm, and registration errors of 0.33 ± 0.14 mm and 0.38 ± 0.16°.

V. Discussion

While the results from the simulations are quite promising, there is significant work remaining in order to advance the presented work into real-world use. Our initialization procedure, which provided an initial estimate of the registration parameters, significantly reduced the number of particles required to adequately cover the state space, yet a large number of particles are still required to ensure convergence. Providing a more accurate

NIH-PA Author Manuscript

estimate of the registration parameters could decrease the number of particles, as could providing bounds on the initial position of the robot on the surface of the heart. This is likely feasible in procedures as HeartLander is generally placed on the anterior surface near the apex of the heart. Further reductions may be achieved by incorporating cardiac phase measurements from an electrocardiogram (ECG).

Improvements to the algorithm may also show improvement over the existing framework. Inherent in the system is an ambiguity in the cause of error. It is not possible to determine if the error is due to an error in the location of the robot on the surface of the heart or in the registration parameters. The particle filter relies on the noise in the motion model to cover the state space and settle at a global minimum. A Rao-Blackwellized particle filter, which updates the registration parameters in a deterministic step, similar to the registration calculation in [6], [10], will likely improve performance.

The major hurdle to overcome is the construction of accurate dynamic pre-operative maps for use in actual interventions. As previously mentioned, there is considerable work in the medical imaging community addressing this and the authors remain confident that such maps will be more readily available in the near future.

VI. Conclusions

The presented work acts as a proof-of-concept for localizing on the beating heart using only a 6-DOF position measurement and a map of the heart. Although implementation in the operating room will require further work, the results produced show that not only is localization of mobile robots on a periodically deforming surface feasible, a well-tuned particle filter provides low-error estimates of the true location within a few hundred iterations, and within a few periods of the cyclical deformation. The motion model and observation model can be modeled with the limited data on currently available robot technologies such as simple 6-degree-of-freedom pose.

Acknowledgments

This work was supported in part by the U.S. National Institutes of Health under Grant nos. R01 HL078839 and R01 HL105911.

References

1. Omary RA, Green JD, Schirf BE, Li Y, Finn JP, Li D. Real-time magnetic resonance imaging-guided coronary catheterization in swine. Circulation. 2003; 107(21):2656–2659.
2. Thiagalingam A, Manszke R, D'Avila A, Ho I, Locke AH, Ruskin JN, Chan RC, Reddy VY. Intraprocedural volume imaging of the left atrium and pulmonary veins with rotational x-ray angiography: Implications for catheter ablation of atrial fibrillation. Journal of Cardiovascular Electrophysiology. 2008; 19(3):293–300. [PubMed: 18005027]
3. Novotny, P.; Stoll, J.; Dupont, P.; Howe, R. Real-time visual servoing of a robot using three-dimensional ultrasound; IEEE Int Conf on Robotics and Automation; Apr. 2007 p. 2655-2660.
4. Tully, S.; Kantor, G.; Zenati, MA.; Choset, H. Shape estimation for image-guided surgery with a highly articulated snake robot; IEEE/RSJ Int Conf on Intelligent Robots and Systems; Sep. 2011 p. 1353-1358.
5. Tully, S.; Kantor, G.; Choset, H. Inequality constrained Kalman filtering for the localization and registration of a surgical robot. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems; Sep. 2011
6. Brij Koolwal A, Barbagli F, Carlson C, Liang D. An ultrasound-based localization algorithm for catheter ablation guidance in the left atrium. Int J Robot Res. 2010; 29(6):643–665.
7. Shechter G, Resar J, McVeigh E. Displacement and velocity of the coronary arteries: cardiac and respiratory motion. IEEE Trans Med Imag. 2006; 25(3):369–375.

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

8. Patronik N, Ota T, Zenati M, Riviere C. A miniature mobile robot for navigation and positioning on the beating heart. IEEE Trans Robot. Oct; 2009 25(5):1109–1124.
9. Wood, N.; Liu, TYT.; Waugh, K.; Zenati, M.; Riviere, C. Towards localizing on the surface of the beating heart. IEEE Int. Conf. Engineering in Medicine and Biology Society; submitted
10. Zhong, H.; Kanade, T.; Schwartzman, D. Virtual touch: An efficient registration method for catheter navigation in left atrium. Int. Conf. Medical Image Computing and Computer Assisted Intervention; 2006. p. 437-444.
11. Malchenko S, Vedru J. Model of heart shape cyclic variation for Foucault cardiography simulations. Int J Bioelectromagnetism. 2003; 5(1):318–319.
12. Chen SY, Guan Q. Parametric shape representation by a deformable NURBS model for cardiac functional measurements. IEEE Trans Biomed Eng. Mar; 2011 58(3):480–487.
13. Park, K.; Metaxas, D.; Axel, L. A finite element model for functional analysis of 4D cardiac-tagged MR images. Int. Conf. Medical Image Computing and Computer Assisted Intervention, ser. Lecture Notes in Computer Science; Berlin/Heidelberg: Springer; 2003. p. 491-498.
14. Suinesiaputra A, Frangi A, Kaandorp T, Lamb H, Bax J, Reiber J, Lelieveldt B. Automated detection of regional wall motion abnormalities based on a statistical model applied to multislice short-axis cardiac MR images. IEEE Trans Med Imag. Apr; 2009 28(4):595–607.
15. Thrun, S.; Burgard, W.; Fox, D. *Probabilistic Robotics*, ser. Intelligent robotics and autonomous agents. The MIT Press; Aug. 2005
16. Huynh DQ. Metrics for 3D rotations: Comparison and analysis. J Math Imaging Vis. Oct; 2009 35(2):155–164.

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

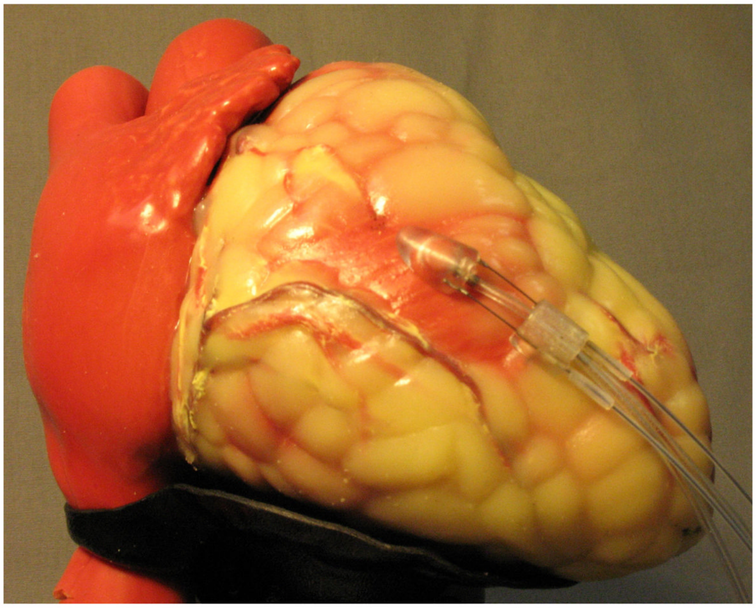

Fig. 1.
The HeartLander robot.

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

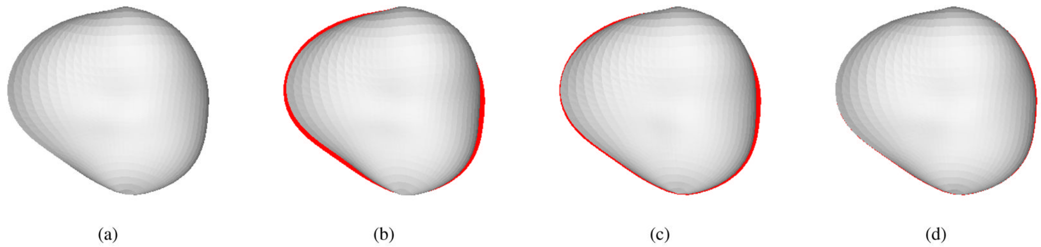

Fig. 2.

Model of the epicardial surface at cardiac phases, ϕ, of (a) $\phi = 0$, (b) $\phi = 0.25$, (c) $\phi = 0.5$, and (d) $\phi = 0.75$. The red background in (b)–(d) denotes the shape of the heart at $\phi = 0$.

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

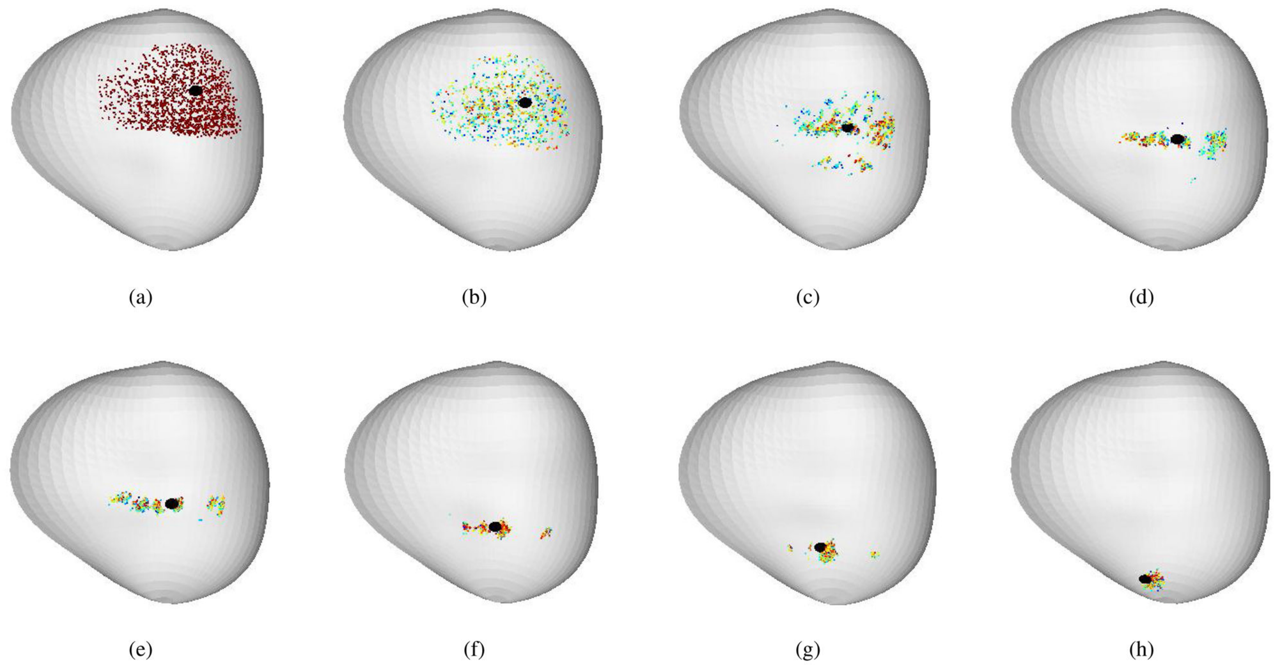

Fig. 3.

Representative results of localizing on the surface of a beating heart using a particle filter. The ground truth position of the robot is shown by the large black dot. Each particle is represented by a small dot whose color denotes the particle weight. Low weights correspond to blue and high to red, with the color scale spanning the weights of the current particles. Plots correspond to (a) Filter initialization with particles randomly distributed over the surface, (b) after 10 iterations, (c) after 30 iterations, (d) after 40 iterations, (e) after 50 iterations, (f) after 70 iterations, (g) after 90 iterations, and (h) after 115 iterations

NIH-PA Author Manuscript
NIH-PA Author Manuscript
NIH-PA Author Manuscript

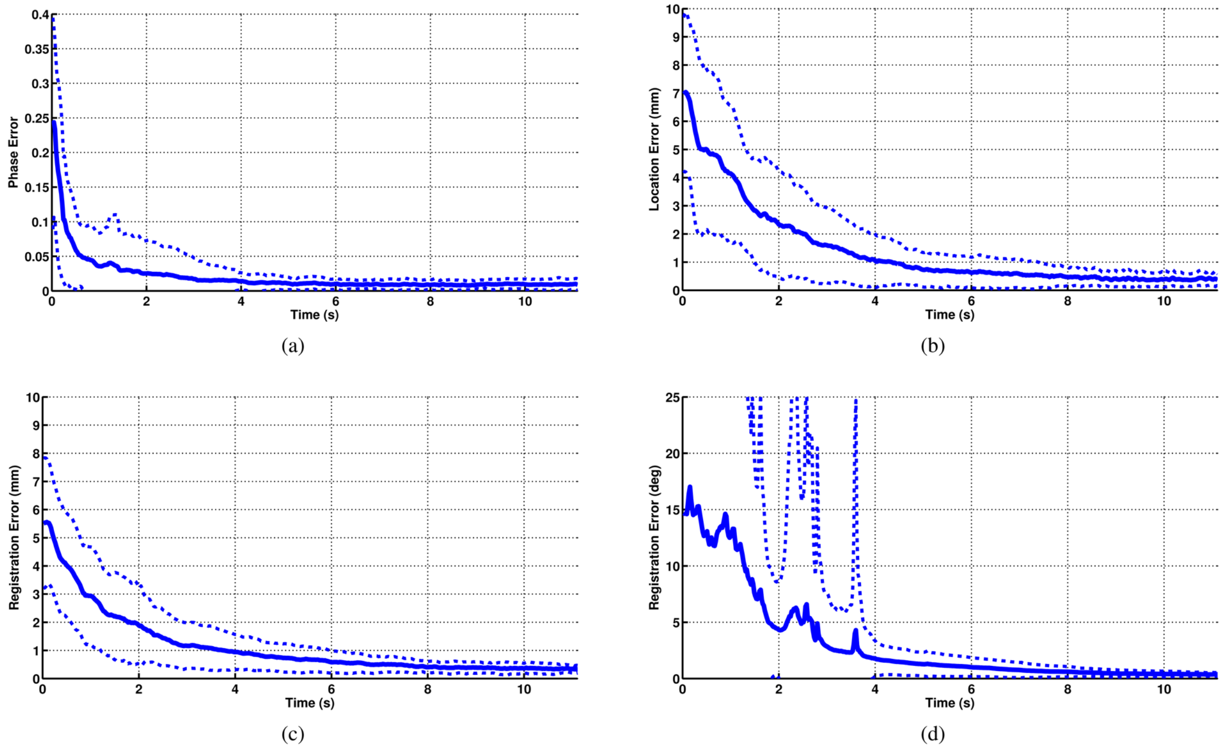

Fig. 4.

Results of localization averaged over 100 runs for (a) phase error, (b) map frame position error, (c) registration position error, and (d) registration orientation error. Mean errors are denoted by solid lines and one standard deviation by dashed lines. Errors for each run are calculated between the ground truth and the weighted average of particles. Position errors are calculated using the euclidean distance, phase and angle errors are absolute differences.