

Simultaneous Compliance and Registration Estimation for Robotic Surgery

Siddharth Sanan*, Stephen Tully*, Andrea Bajo†, Nabil Simaan† and Howie Choset*

*The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA

Email: sanan@cmu.edu, stully@ece.cmu.edu, choset@cs.cmu.edu

†Dept. of Mechanical Engineering
Vanderbilt University
Nashville, TN

Email: andrea.bajo@gmail.com, nabil.simaan@vanderbilt.edu

Abstract—Leveraging techniques pioneered by the SLAM community, we present a new filtering approach called simultaneous compliance and registration estimation or CARE. CARE is like SLAM in that it simultaneously determines the pose of a surgical robot while creating a map, but in this case, the map is a compliance map associated with a preoperative model of an organ as opposed to just positional information like landmark locations. The problem assumes that the robot is forcefully contacting and deforming the environment. This palpation has a dual purpose: 1) it provides the necessary geometric information to align or register the robot to *a priori* models, and 2) with palpation at varying forces, the stiffness/compliance of the environment can be computed. By allowing the robot to palpate its environment with varying forces, we create a force balanced spring model within a Kalman filter framework to estimate both tissue and robot position. The probabilistic framework allows for information fusion and computational efficiency. The algorithm is experimentally evaluated using a continuum robot interacting with two benchmark flexible structures.

I. INTRODUCTION

Minimally invasive surgery (MIS) offers patients many benefits including fast recovery, reduced pain and improved cosmesis compared to open surgery. These advantages, however, come at the cost of increased challenges to the surgeons in terms of manipulation dexterity, diminished haptic presence, and limited situational awareness. The limited situational awareness during MIS stems from the loss of visualization due to the extremely narrow field of view imposed by the necessary use of thin endoscopes. Added to the narrow field of view is the lack of haptic sensory immersion which precludes the detection and use of surface and stiffness features to deduce correlations between the intraoperative scene and the preoperative imaging.

Robot-assisted MIS and computer-aided surgery were introduced to assist surgeons in overcoming the limitations of MIS. Robot-assisted MIS allows surgeons to augment their skills with the dexterity and precision of the robots (e.g. [8, 18]). Newer systems allow surgeons to extend their reach deeper into the patient’s anatomy (e.g. [4, 22]). Computer-aided surgery allowed surgeons to achieve intra-operative navigation

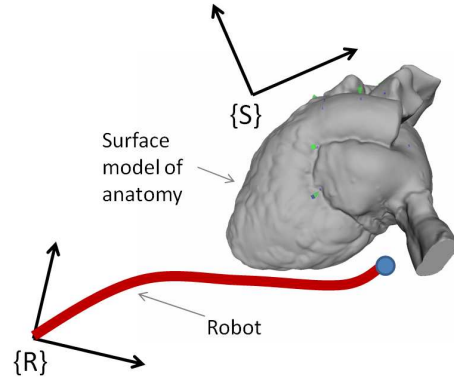


Fig. 1. The registration parameters define the transformation between the coordinate frame S of the surface model and the coordinate frame R of the robot.

and assistance in maintaining correspondence between pre-operative surgical plans and their execution.

Computer-aided image overlay augments surgeon’s understanding of the anatomy and its correspondence with pre-operative images though an overlay of 3D reconstructed pre-operative images on the visible anatomy[16]. Successful image-guided robot-assisted surgery hinges on registering the surgical tool to the coordinate frame of a known *a priori* 3D surface model (as obtained from pre-operative 3D image reconstruction), depicted in Fig. 1.

Registration to rigid anatomy such as bones is easy and does not require intraoperative update due to organ shifting, swelling or gravitational pull. When operating on flexible anatomy this task becomes more challenging. Organ shifting distorts the mapping between intraoperative sensitive anatomy and its corresponding interpretation in image space. This further complicates the task of the surgeon in avoiding critical anatomy such as nerves, arteries, biliary duct, etc.

While there are approaches to overcome organ shift using deformable registration from intraoperative surface scanning (e.g. [10, 13]), these approaches typically work using white-

light based imaging data. In cases where the critical anatomy is obscured by blood or is underneath tissue layers one can use either intraoperative imaging or mechanical palpation data. For example, mechanical stiffness imaging has been shown to allow surgeons to detect hidden features based on probing for changes in relative stiffness [11, 20, 23, 21, 17].

A common approach to achieve registration is to use an iterative closest point (ICP) algorithm [3] which typically assumes the surrounding environment is rigid and does not account for tool-tissue interaction. Unfortunately ICP is too computationally expensive to update registration parameters in real time. Probabilistic filtering based approaches have been developed to perform registration [15, 12]. These techniques primarily use line of sight measurements, i.e., 3D vision or range scanning, to output a map of locations of anatomical structures and landmarks. There has also been prior work in estimating tissue compliance using palpation, without solving the registration problem, using filtering methods [9, 23].

Our approach is to use contact between the tool and the environment to estimate the alignment between the tool and the anatomical structure. If the environment were rigid, estimating registration using geometric contact constraints alone would be possible. However, since we are dealing with flexible environments with unknown elastic properties, the parameters defining the elastic properties must also be estimated. In our formulation, we incorporate contact force measurement to estimate both the environment stiffness and registration parameters simultaneously using an iterated extended Kalman filter (IEKF) based algorithm. In this work, we extend the approach of [19] by relaxing two assumptions made in that work: 1) the environment has uniform stiffness everywhere, 2) the environment stiffness is known. In this work, we initialize our IEKF with a distribution of stiffness and then we use contact and force information to estimate the stiffness map and register the robot frame to the *a-priori* model of the environment.

The contribution of this work is a new method capable of “on the fly” registration of a robot to a flexible environment. We have adopted a probabilistic method that allows us to readily estimate the compliance or stiffness of the environment while simultaneously aligning the robot to an organ. Such information can be used for image-guided surgery and haptic feedback during MIS. The identified environment stiffness can be used to locate tumors and major arteries at the surgical field.

The outline of the remainder of the paper is as follows: in Sec. II we outline the problem statement and assumptions, in Sec. III we introduce our registration and compliance estimation approach, in Sec. IV we test the observability of our problem using a simple test case, and in Sec. V we discuss the experiments we performed. Finally, in Sec. VI, we conclude the paper with a discussion of the contributions we have made and future work.

II. PROBLEM STATEMENT AND ASSUMPTIONS

Problem Statement: Given an *a priori* model of the environment and a set of robot end effector locations with

associated contact force magnitudes do the following: 1) create a set of estimated stiffness normal to the surface, 2) find the registration parameters relating the robot frame to the model of the environment. In the following discussion, we make few simplifying assumptions. The following list provides the simplifying assumptions and the rationale behind making these assumptions:

- 1) Palpation data is obtained using a robot capable of force sensing. For example, we use a continuum robot with force sensing algorithms as demonstrated in [21].
- 2) The position data of the robot tip is accurate. If the robot kinematic model is not accurate then the robot tip is tracked using external measurements such as a magnetic tracker.
- 3) The forces that the robot applies on the environment during palpation are small enough not to cause a large deformation of the surface. Although this assumption limits the ability of palpation to discern deep underlying stiffness features, we are still able to detect stiffness features close to the surface of the environment. A typical case where detection of a shallow underlying stiffness feature is useful is during the final stages of resection around an artery with the aim of avoiding inadvertent cutting of the artery.
- 4) The friction between the robot tip and the environment are negligible. The rationale for this assumption is that the anatomy is typically lubricious and also the robot tip may be equipped with a rolling/spherical bearing in a manner similar to the work of [14].
- 5) The deformation of the model is predominantly in the direction normal to the surface at contact point. This assumption follows from the previous one since lack of friction eliminates tangential forces on the environment.

We also emphasize that the motivation for estimating the environment stiffness from contact detection is for detecting *relative* stiffness changes as the robot probe tip moves along the surface. We are not interested in exact data of stiffness for diagnosis. Instead, relative stiffness changes can be used to detect nodules, cysts, arteries, etc. Examples of works using relative stiffness changes for stiffness imaging include [14, 20, 5]. Since the exact value of the environment stiffness is not important and since we limit the palpation forces to levels associated with small deformations of the tissue we use a simple linear model of the environment stiffness.

III. FILTERING FOR REGISTRATION AND COMPLIANCE ESTIMATION

The goal of the CARE algorithm is to estimate a state vector that encodes the alignment of the robot as well as the compliance of the surface. We therefore define \mathbf{x} as follows,

$$\mathbf{x} = [\mathbf{x}_{k,robot} \ \boldsymbol{\theta}_{k,robot} \ s_k \ \mathbf{x}_{k,tip}]^T, \quad (1)$$

where $\mathbf{x}_{k,robot}$ and $\boldsymbol{\theta}_{k,robot}$ represent the position and orientation vectors, respectively, of the robot base in the coordinate frame of the surface model, s_k is a vector defining the stiffness

in a direction normal to the surface, and $\mathbf{x}_{k,tip}$ represents the position of the robot tip in the base frame.

During data collection for generating the input to the CARE algorithm, the robot may be telemanipulated using hybrid position-force control ensuring a regulated force on the surface subject to a rough estimate of the surface normal. The robot detects that it is in contact with a surface using an implementation similar to [1] or by using a threshold on the force estimated at the tip of the robot. The robot infers the contact force by observing the joint actuation forces and solving the inverse statics problem in a manner similar to [21]. For our implementation, the surface is assumed to be locally planar and the contact force is assumed to be normal to the surface. In addition to measuring force, the robot also measures its tip position in the base frame of the robot.

There are two components of the filtering process: a prediction step that evolves the state estimate based on a motion model and a measurement update step that corrects the state estimate with live sensor data. For our implementation, we have designed the prediction step assuming that the parameters are static, with a large uncertainty added to the tip position due to tool movement.

Measurement Function: We assume a simple linear spring model governs the tool-environment interaction forces. This assumption holds within the surgical context even though tissue is highly non-linear because we are interested in mapping relative changes in stiffness and the absolute estimates of stiffness are not important to the problem. For example, an artery is detected by a change in the stiffness from the surrounding tissue. Since we care only about detecting the change between the tissue and artery, absolute stiffness values are unnecessary and we therefore avoid using complex non-linear models. Additionally, this assumption is valid because we apply only small forces on the environment so as to produce minimal local impingement on the surface without a substantial deformation of the organ. This ensures that we are able to estimate the shape and relative stiffness, and that our measurement technique does not change the environment in a way that makes the problem intractable.

Therefore, the measurement function used to predict the expected measurements (force and tip position) at time-step k given \mathbf{x}_k is as follows,

$$h(x_k) = \begin{pmatrix} -\mathbf{n}^T(\mathbf{x}_{k,tip}^S - \mathbf{p})s(\mathbf{x}_k) \\ \mathbf{x}_{k,tip} \end{pmatrix} \quad (2)$$

where \mathbf{n} is the surface normal, \mathbf{p} is a point on the (relaxed) surface, and $s(\mathbf{x}_k)$ is a function that extracts the appropriate surface stiffness given the state \mathbf{x}_k . $\mathbf{x}_{k,tip}^S$ is the current estimate of the robot tip in the coordinate frame of the surface given by:

$$\mathbf{x}_{k,tip}^S = R(\mathbf{x}_k)\mathbf{x}_{k,tip} + t(\mathbf{x}_k), \quad (3)$$

$R(\mathbf{x}_k)$ and $t(\mathbf{x}_k)$ are the rotation matrix and translation vector respectively, that perform this transformation based on the

robot base position $\mathbf{x}_{k,robot}$ and the robot base orientation $\theta_{k,robot}$.

In (2), the surface normal \mathbf{n} is obtained by finding the normal associated with the plane on the surface that is closest to the current estimate of the robot tip position $\mathbf{x}_{k,tip}^S$. Similarly, the function $s(\mathbf{x}_k)$ extracts the stiffness corresponding to the plane closest to $\mathbf{x}_{k,tip}^S$.

For the measurement update step, we use an iterative algorithm that prevents divergence of the state estimate. The approach is similar to the iterated extended Kalman filter (IEKF) described in [2]. After each new measurement is obtained, the IEKF attempts to update the state mean $\hat{\mathbf{x}}_{k|k}$ to minimize the following cost function:

$$\hat{\mathbf{x}}^+ = \arg \min_{\mathbf{x}} [\hat{\mathbf{x}}_{k|k} - \mathbf{x}]^T \mathbf{P}_{k|k}^{-1} + [\mathbf{z}_k - h(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{z}_k - h(\mathbf{x})] \quad (4)$$

where \mathbf{z}_k is the k -th observation, $\mathbf{P}_{k|k}$ is the state covariance and \mathbf{R} is the measurement noise covariance.

The above optimization problem corresponds to finding the maximum likelihood estimate given the current observation \mathbf{z}_k . IEKF minimizes the cost function (4) using a recursive Gauss-Newton update procedure. The update equations are as follows:

$$\mathbf{x}_0 = \hat{\mathbf{x}}_{k|k} \quad (5)$$

$$\mathbf{K}_i = \mathbf{P}_{k|k} \mathbf{H}_i^T [\mathbf{H}_i \mathbf{P}_{k|k} \mathbf{H}_i^T]^{-1} \quad (6)$$

$$\mathbf{x}_{i+1} = \hat{\mathbf{x}}_{k|k-1} + \quad (7)$$

$$\mathbf{K}_i [\mathbf{z}_k - \mathbf{h}(\hat{\mathbf{x}}_i) - \mathbf{H}_i (\hat{\mathbf{x}}_{k|k-1} - \mathbf{x}_i)],$$

where \mathbf{H}_i is the Jacobian of $\mathbf{h}(\mathbf{x})$ linearized about \mathbf{x}_i . The above recursion is continued until convergence. We say the filter has converged when the $\|\mathbf{x}_{i+1} - \mathbf{x}_i\| < \epsilon$, where $\epsilon \ll 1$. At convergence, the state mean is updated as $\hat{\mathbf{x}}_{k|k}^+ = \mathbf{x}_i$. While we do not have a guarantee of convergence, in our experiments the filter has converged every time. It should be noted that a single iteration of the above procedure corresponds to the standard EKF update step.

Once the mean has been updated using the iterative procedure described above, the covariance matrix is updated using the standard EKF covariance update equation,

$$\mathbf{K}_i = \mathbf{P}_{k|k} \mathbf{H}^T [\mathbf{H} \mathbf{P}_{k|k} \mathbf{H}^T]^{-1} \quad (8)$$

$$\mathbf{P}_{k|k}^+ = \mathbf{P}_{k|k} - \mathbf{K} \mathbf{H} \mathbf{P}_{k|k} \quad (9)$$

where \mathbf{H} is the Jacobian of $\mathbf{h}(\mathbf{x})$ linearized about the updated mean $\hat{\mathbf{x}}_{k|k}^+$.

It should be noted that a key distinguishing feature of our filtering approach to registration is that we update our registration to minimize the residual force (i.e., the difference between the measured force and predicted force) at each contact point instead of minimizing a geometric quantity such as distance between the contact point and surface. We note that we choose to instantiate a number of these Kalman filters, each of which attempts to estimate the unknown registration parameters. When the filter has utilized all the available observations, we take the hypothesis with the highest measurement likelihood as the final result.

IV. OBSERVABILITY

In order to apply the filter described in the previous section, the state must be observable. We test observability through the use of a simple test case, described below. Assume an environment with three planes of different orientation, indexed by $i = (1, 2, 3)$. Each plane is defined by a normal vector \mathbf{n}_i and a plane point \mathbf{p}_i . For this test case, we assume we have knowledge of which plane we are contacting and that both \mathbf{n}_i and \mathbf{p}_i are known. Each plane has an unknown stiffness s_i that is encoded within the state vector for estimation.

The state vector is of length 9, which defines the position and orientation of the robot base in the coordinate frame of the surface model and encodes the stiffness of each of the three planes.

$$\mathbf{x}_k = [x \ y \ z \ r_z \ r_y \ r_x \ s_1 \ s_2 \ s_3]^T$$

We obtain three measurements from palpating the three planes. We next obtain three additional measurements through palpating the same three planes with a different force. The point on the plane where contact occurs does not affect the analysis. Let us assume that all of these force measurements are obtained and collected together. We will use the term $h_j(\mathbf{x}_k)$ to represent the j -th element of the measurement function

$$\mathbf{h}(\mathbf{x}_k) = [h_1(\mathbf{x}_k) \ h_2(\mathbf{x}_k) \ \dots \ h_6(\mathbf{x}_k)]^T$$

and the force measurement model can be written as follows:

$$h_j(\mathbf{x}_k) = -\mathbf{n}_i^T (\mathbf{R}(\mathbf{x}_k)\mathbf{q}_j + \mathbf{t}(\mathbf{x}_k) - \mathbf{p}_i) s_i$$

where, for a measurement index j , the associated plane that is being contacted is i -th plane. \mathbf{q}_j is the robot tip position corresponding to the j -th observation, measured in the robot frame. Using Lie derivatives, we can define two different observability functions as follows,

$$\mathbf{o}^0(\mathbf{x}_k, \mathbf{u}_k) = \begin{pmatrix} (\mathcal{L}_{\mathbf{f}}^0 h_1)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^0 h_2)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^0 h_3)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^0 h_4)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^0 h_5)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^0 h_6)(\mathbf{x}_k, \mathbf{u}_k) \end{pmatrix}, \quad \mathbf{o}^1(\mathbf{x}_k, \mathbf{u}_k) = \begin{pmatrix} (\mathcal{L}_{\mathbf{f}}^1 h_1)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^1 h_2)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^1 h_3)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^1 h_4)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^1 h_5)(\mathbf{x}_k, \mathbf{u}_k) \\ (\mathcal{L}_{\mathbf{f}}^1 h_6)(\mathbf{x}_k, \mathbf{u}_k) \end{pmatrix}$$

where, \mathbf{u}_k is the input at step k . The zero-th Lie derivative of a function is the function itself,

$$(\mathcal{L}_{\mathbf{f}}^0 h_j)(\mathbf{x}_k, \mathbf{u}_k) = h_j(\mathbf{x}_k), \quad (10)$$

and the first Lie derivative of a function is as follows,

$$(\mathcal{L}_{\mathbf{f}}^1 h_j)(\mathbf{x}_k, \mathbf{u}_k) = \frac{\partial h_j(\mathbf{x}_k)}{\partial \mathbf{x}_k} \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k). \quad (11)$$

Substituting Eq. 10 and 11 into the functions $\mathbf{o}^0(\mathbf{x}_k, \mathbf{u}_k)$ and $\mathbf{o}^1(\mathbf{x}_k, \mathbf{u}_k)$, we can write a combined observability function,

$$\mathbf{o}(\mathbf{x}_k, \mathbf{u}_k) = \begin{bmatrix} \mathbf{o}^0(\mathbf{x}_k, \mathbf{u}_k) \\ \mathbf{o}^1(\mathbf{x}_k, \mathbf{u}_k) \end{bmatrix} = \begin{bmatrix} \mathbf{h}(\mathbf{x}_k) \\ \mathbf{H}_k \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) \end{bmatrix},$$

where the function $\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k)$ assumes the state parameters are static, i.e. $\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) = \mathbf{x}_k$. We can then compute the Observability matrix by differentiating the observability function with respect to the Kalman state vector

$$\mathcal{O}(\mathbf{x}_k, \mathbf{u}_k) = \frac{\partial}{\partial \mathbf{x}_k} \mathbf{o}(\mathbf{x}_k, \mathbf{u}_k).$$

To test whether the system is observable, we must determine if the observability matrix is invertible. This can be determined by looking at the rank of the matrix,

$$\text{rank}(\mathcal{O}(\mathbf{x}_k, \mathbf{u}_k)) = 9.$$

We obtained this value through the use of symbolic analysis using Matlab (The Mathworks Inc., Natick, MA). The resulting rank of 9 means that all 9 parameters in the Kalman state vector are observable.

V. EXPERIMENTAL EVALUATION

The experimental evaluation involved using a 5-mm diameter two-segment continuum snake robot. This robot uses several NiTi backbones in each one of its bendable segments. Each segment consists from a central backbone surrounded by several secondary backbones that are pushed or pulled on to bend the segment. Our robot uses three circumferentially distributed NiTi secondary backbones for actuation in a manner similar to [7]. The robot tip is a spherical probe equipped with a magnetic tracker allowing for exact measurement of the robot tip location.

Figure 2 shows the robot with its spherical probe. An Ascension 3D guidance trackSTARTM magnetic tracker with metal immune magnetic transmitter was used for tracking the location of the 0.9 mm magnetic coil embedded in the probe contacting the environment. We assume the magnetic tracker measurement noise is normally distributed with zero mean and standard deviation $\sigma_t = 3$ mm. The robot is also equipped with load cells for measuring the actuation forces on the secondary backbones. Using algorithms such as [21] for estimating the force at the robot tip from measurement of the forces on the load cells, we obtained estimates of the forces at the robot tip. The force sensor noise is assumed to be normally distributed with zero mean and standard deviation $\sigma_f = 1$ g (0.01 N). Evaluation of the force estimation accuracy is reported in [21]. The control environment was implemented on a real-time control system using Matlab xPC. This system was used to palpate the object with varying force levels at multiple points on its surface.

We evaluated our probabilistic filtering-based approach using two experiments. The first experiment utilized a simple planar surface consisting of two different materials. We used this as a proof of concept experiment and to develop an intuitive understanding of our approach. We then applied our algorithm to a more complex setup, where inclusions stiffer than the surrounding environment were embedded in the material. We describe the methodology and results of the two experiments in the following sections.

A. Planar Surface

The experimental model, shown in Fig. 2 was made from silicone. The stiffer inclusions were made from rubber wire with 80 Shore A hardness. The experimental model was palpated with two force levels of 5 g (0.049 N) and 25 g (0.245 N) and tip position data was collected during each palpation.

Dataset and Experiment: We split the palpation data from the experiment into two data sets. The first was used to infer a best approximation to the relaxed position of the surface (shown in Fig. 3 as a gray plane). This data was also used to infer the ground truth stiffness (0.089 N/mm and 0.196 N/mm) for the outer and inner regions, respectively. The second dataset was used for testing CARE. The test data was artificially disturbed with a translation in x and z -directions, resulting in a large registration error with respect to the surface. This was done to create a challenging registration test for our approach. CARE was applied to this data by iterating through all the contact points in the test dataset using equations outlined in Sec. III.




Fig. 2. The flexible planar surface consisting of two regions of different material. The inner region is stiffer than the outer region. The robot contacts the surface with a spherical probe at its end.

The registration result of our probabilistic filtering method, CARE, compared to the initial registration, is shown in Fig. 3. The initially perturbed point cloud of contact points is shown in green and are poorly registered to the surface. After the filter applies the contact force measurements, the algorithm correctly aligns the contact points to the surface such that the points are properly displaced “into” the surface to generate forces that agree with the applied palpation force. Note that points to the right of the surface in Fig. 3 are inside the surface while points to the left of the surface are outside the surface.

We note that, during the experiment, the estimate of the x dimension of the robot registration becomes highly correlated with the estimate of the surface compliance. This occurs because the estimation problem is not observable until the surface is palpated at two different forces. In Fig. 4, an error plot is shown that depicts the error between the estimated compliance and the true compliance of the surface for the two (inner and outer) regions. The data illustrates the sudden drop

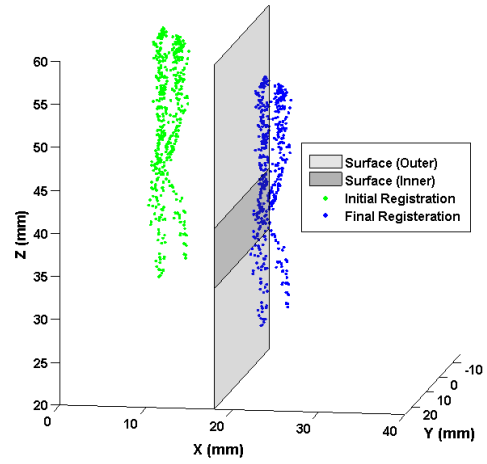


Fig. 3. The contact points using the initial guess registration (shown in green) and after the registration parameters have been estimated using our filtering algorithm (shown in blue). After application of the filter all the tip positions are ‘inside’ the planar surface.

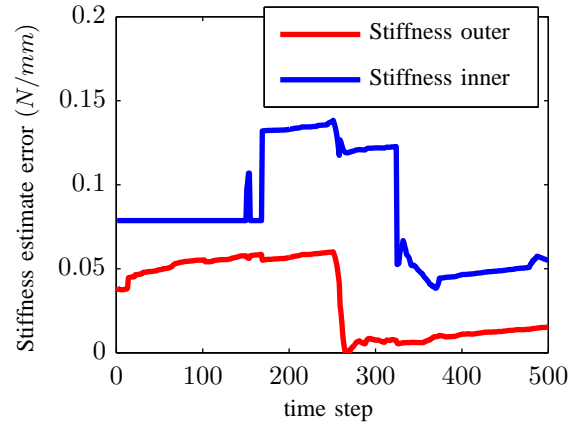


Fig. 4. This figure shows the error history of the stiffness estimates for both the inner and outer regions as measurements are incrementally included in the filtering algorithm.

in error when the surface is palpated at a second contact force at the 250-th time step. The final registration error in terms of aligning the robot to the *a priori* surface was 0.55mm. The stiffness estimation error is larger in the inner region because there are fewer data points in that region.

It should be noted that for the experiment described above it would be impossible for any registration method that utilizes geometric information alone to correctly register the robot frame along the z direction. This is because the problem is ill-posed due to the geometric symmetry that exists in the object along the z -direction. The compliance map breaks this symmetry and therefore allows registration in the z -direction.

For the purpose of comparison, we applied the standard ICP algorithm to the above problem. The initial point cloud obtained via the palpation experiment (green in Fig. 5) has registration error along both the x and z -directions. The

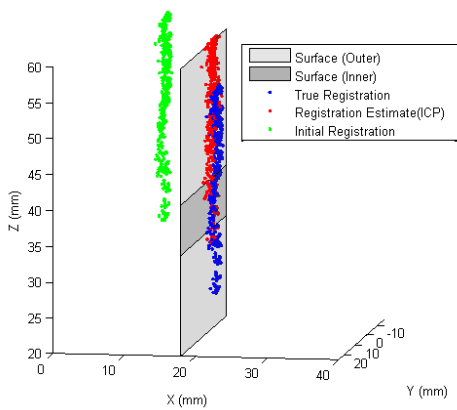


Fig. 5. The point cloud obtained by palpating the surface is registered to the surface using ICP. The green point cloud corresponds to the initial guess registration, the red point cloud corresponds to the ICP registration estimate and the blue point cloud corresponds to the ground truth registration. Registration error in X is reduced while registration error in Y does not decrease.

point cloud obtained after application of ICP is shown in blue. As can be seen in Fig. 5, ICP minimizes the distance between the point cloud and the surface plane leading to an improved registration estimate in the x-direction. However, as surface compliance and/or palpation forces become larger the registration accuracy in terms of aligning the robot and surface frames will decrease because ICP will attempt to align the deformed surface to the undeformed surface. Additionally, registration along the z-direction is not improved because of the inherent geometric symmetry in the problem. An accurate prior compliance map can address this issue by utilizing the palpation forces and compliance information to generate a deformed target surface for registration. However, our algorithm attempts to solve for both the compliance map and registration simultaneously.

B. 3D Soft Volume

The previous experiment involving the planar surface served as a relatively simple test to verify the feasibility of our approach. We also tested our filtering algorithm on a more complex soft object with 3D geometry and embedded hard material, shown in Fig. 6. The surface geometry was obtained by scanning the surface with a magnetically tracked probe dragged over the surface with only the probe self weight acting on the surface (the probe weighed less than 0.5 grams).

Experiment details: After obtaining the surface geometry of the object, the surface was palpated with two force levels of 5 g (0.049 N) and 25 g (0.245 N) using the continuum snake. The tip of the snake was tracked using the magnetic tracker. As the surface scan and palpation were done with the same tracker setup, the robot frame R and surface frame S are coincident and therefore the ground truth registration parameters of the robot are known. Using the known ground truth registration we estimate the ground truth surface stiffness. Figure 7 shows the ground truth stiffness distribution on the pyramid surface. The state vector for this problem is considerably larger as the




Fig. 6. A pyramid shaped silicone specimen with embedded hard material used in the experiments.




Fig. 7. Ground truth stiffness distribution on the pyramid surface. High stiffness is indicated in red and low stiffness in dark blue.

size of the stiffness vector \mathbf{s}_k is equal to the number of facets in the surface model, i.e, each face has an associated stiffness. The total number of facets in our surface model are 236.

For testing, we artificially perturbed the palpation data contact points with a rigid transformation T_{pert} . The corresponding ground truth registration for the perturbed data set is T_{pert}^{-1} . Using CARE we simultaneously estimated the registration of the robot and surface stiffness parameters of the object. We also estimated the registration parameters by registering the perturbed contact point cloud to the surface using ICP.

Figure 8 shows a perturbed contact point cloud based on the initial registration parameters. The initial registration error is 10mm in the z-direction. As a result all initial contact points are above the surface. Once the algorithm has iterated through all the observations in the data set, all contact points are below the relaxed surface, as shown in Fig. 9. Fig. 10 shows the error in the registration and stiffness estimates as the algorithm iterates through the data. Fig. 11 shows the error in the registration and stiffness estimates when an orientation error of 11.5° about the Y-axis is applied to the data set. The stiffness error is computed by taking the Euclidean distance between the ground truth stiffness vector and the estimated stiffness vector. The distribution of the estimated stiffness on the surface model is also shown for both the experiments in




Fig. 8. Data points prior to registration.





Fig. 9. Data points after registration.

Fig. 12 and 13. It can be seen from direct observation that the estimated stiffness distribution is qualitatively similar to the ground truth stiffness distribution.


VI. CONCLUSION AND FUTURE WORK

The results of our experiments show that we can use palpation as an effective strategy for registering the robot “on the fly” to a flexible environment and to infer the surface compliance or stiffness of the environment. Unlike conventional approaches, CARE uses probabilistic filtering with forceful contact, allowing us to estimate the compliance/stiffness of the environment while simultaneously aligning the robot to an organ. The registration and compliance information estimated by CARE can be of significant assistance during the use of image guidance and haptic feedback for minimally invasive surgery.

The significance of our approach is apparent as it relates to a “chicken-and-egg problem”: if the compliance of the environment is known, the expected displacement of each contact point into the surface can be determined, thus simplifying registration. Conversely, if the registration is known, the stiffness can be determined by observing the surface displacement for a given contact force. We believe that solving for both the



(a) Registration error



(b) Stiffness estimation error

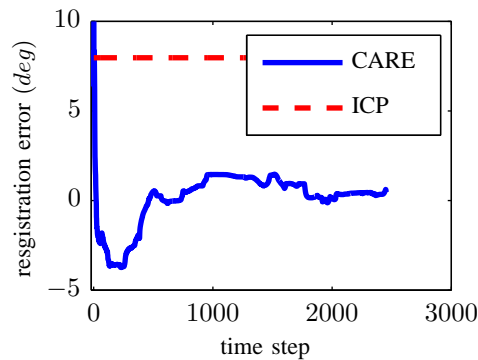
Fig. 10. Registration and stiffness estimation errors when an initial registration error of 10mm in the z-direction is applied to the data set.

registration and compliance requires a probabilistic approach such as the one presented in this paper.


The work we presented relies on the use of probabilistic filtering for the estimation of uncertain data. The real world is noisy and so it is crucial to model and estimate this uncertainty via sensor fusion techniques. Probabilistic filtering methods like CARE allow us to report confidence measures to the surgeon during image-guided interventions for increased situational awareness and more advanced visual feedback. A promising application of our approach is to provide feedback to the surgeon regarding the location of hard and soft regions in the surgical field.

Our approach does not aim to substitute the use of medical imaging for registration but to augment it. Rigid and deformable registration techniques that have been studied extensively in medical imaging can be combined with the approach described in our paper to leverage sub-surface features using palpation force data. It is also worth noting that any probabilistic filtering technique can be adopted within the CARE framework. While we have chosen to use the IEKF in our implementation due to our prior experience and success with it, other filtering techniques such as those based on the unscented Kalman filter can also be easily adopted within the CARE framework.

The limitations of our proposed approach stem partly from



(a) Registration error



(b) Stiffness estimation error

Fig. 11. Registration and stiffness estimation errors when an initial registration error of 11° about the Y-axis is applied to the data set.

our simplifying assumptions. Because we assume that the robot does not deform the environment significantly we are currently limited in our ability to detect stiffness features embedded deep inside the model. Another limitation is that our force controller needs some estimate of the surface normal. In this work we have not attempted online estimation of the surface normal during the implementation of force control.

Future work will include adoption of well known methods for surface normal algorithms during force control (such as [6]). Another possible direction for future research is to not only estimate surface properties as we have done in this paper but to also estimate volumetric properties of the tissue. This will allow us to not only identify regions on the surface of an organ that are soft or hard but to also localize, in 3D, embedded features such as arteries and tumors.

Finally, although in this work we register the robot to the environment using contact force data while not accounting for underlying stiffness features, the eventual goal of our approach will be to also include such information by using a stiffness distribution prior which could be obtained from imaging (e.g. ultrasound elastography).

ACKNOWLEDGMENT

This work has been supported in part by the National Robotics Initiative Award IIS-1327597 and ARO Award




Fig. 12. Estimated stiffness distribution on the pyramid surface for the test case with translation registration error of 10 mm in the z-direction. High stiffness is indicated in red and low stiffness in dark blue




Fig. 13. Estimated stiffness distribution on the pyramid surface for the test case with orientation registration error of 11.5° about the Y axis. High stiffness is indicated in red and low stiffness in dark blue

W911NF1010343.

REFERENCES

- [1] Andrea Bajo and Nabil Simaan. Kinematics-based detection and localization of contacts along multisegment continuum robots. *Robotics, IEEE Transactions on*, 28(2):291–302, 2012.
- [2] B.M. Bell and F.W. Cathey. The iterated kalman filter update as a gauss-newton method. *IEEE Transactions on Automatic Control*, 38(2):294–297, 1993.
- [3] Paul J Besl and Neil D McKay. A method for registration of 3-d shapes. *IEEE Transactions on pattern analysis and machine intelligence*, 14(2):239–256, 1992.
- [4] A. Degani, H. Choset, A. Wolf, and M.A. Zenati. Highly articulated robotic probe for minimally invasive surgery. In *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pages 4167–4172, May 2006.
- [5] Vladimir Egorov and Armen P Sarvazyan. Mechanical imaging of the breast. *IEEE transactions on medical*

- imaging*, 27(9):1275–87, September 2008. doi: 10.1109/TMI.2008.922192.
- [6] A. Fedele, A. Fioretti, C Manes, and G. Ulivi. On-line processing of position and force measures for contour identification and robot control. In *[1993] Proceedings IEEE International Conference on Robotics and Automation*, pages 369–374. IEEE Comput. Soc. Press, 1993. ISBN 0-8186-3450-2. doi: 10.1109/ROBOT.1993.292009.
- [7] R Goldman, A Bajo, L Maclachlan, R Pickens, S Herrell, N Simaan, et al. Design and performance evaluation of a minimally invasive telerobotic platform for transurethral surveillance and intervention. *IEEE transactions on biomedical engineering*, 2012.
- [8] Gary S Guthart and John Kenneth Salisbury Jr. The Intuitive™ telesurgery system: overview and application. In *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, volume 1, pages 618–621. IEEE, 2000.
- [9] Amir Haddadi and Keyvan Hashtrudi-Zaad. A new method for online parameter estimation of hunt-crossley environment dynamic models. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 981–986. IEEE, 2008.
- [10] M Hayashibe and Y Nakamura. Laser-pointing endoscope system for intra-operative 3D geometric registration. In *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, volume 2, pages 1543–1548, Seoul, Korea, Republic of, 2001. IEEE. ISBN 0-7803-6576-3. doi: 10.1109/ROBOT.2001.932830.
- [11] Robert D Howe, William J Peine, DA Kantarinis, and Jae S Son. Remote palpation technology. *Engineering in Medicine and Biology Magazine, IEEE*, 14(3):318–323, 1995.
- [12] A. B. Koolwal, F. Barbagli, C. Carlson, and D. Liang. An ultrasound-based localization algorithm for catheter ablation guidance in the left atrium. *The International Journal of Robotics Research*, 29(6):643–665, May 2010.
- [13] Ray A Lathrop, Douglas M Hackworth, and Robert J Webster. Minimally invasive holographic surface scanning for soft-tissue image registration. *Biomedical Engineering, IEEE Transactions on*, 57(6):1497–1506, 2010.
- [14] Hongbin Liu, David P Noonan, Benjamin J Challacombe, Prokar Dasgupta, Lakmal D Seneviratne, and Kaspar Althoefer. Rolling mechanical imaging for tissue abnormality localization during minimally invasive surgery. *Biomedical Engineering, IEEE Transactions on*, 57(2):404–414, 2010.
- [15] P. Mountney, D. Stoyanov, A. Davison, and G. Z. Yang. Simultaneous stereoscope localization and soft-tissue mapping for minimal invasive surgery. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2006.
- [16] Terry Peters and Kevin Cleary. *Image-Guided Interventions: Technology and Applications*. Springer Science, 2008. ISBN 078-0-387-73856-7.
- [17] Pinyo Puangmali, Kaspar Althoefer, Lakmal D. Seneviratne, Declan Murphy, and Prokar Dasgupta. State-of-the-Art in Force and Tactile Sensing for Minimally Invasive Surgery. *IEEE Sensors Journal*, 8(4):371–381, April 2008. ISSN 1530-437X. doi: 10.1109/JSEN.2008.917481.
- [18] Nabil Simaan, Kai Xu, Wei Wei, Ankur Kapoor, Peter Kazanzides, Russell Taylor, and Paul Flint. Design and integration of a telerobotic system for minimally invasive surgery of the throat. *The International journal of robotics research*, 28(9):1134–1153, 2009.
- [19] Stephen Tully, Andrea Bajo, George Kantor, Howie Choset, and Nabil Simaan. Constrained filtering with contact detection data for the localization and registration of continuum robots in flexible environments. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3388–3394. IEEE, 2012.
- [20] P S Wellman, E P Dalton, D Krag, K A Kern, and R D Howe. Tactile imaging of breast masses: first clinical report. *Archives of surgery (Chicago, Ill. : 1960)*, 136(2):204–8, March 2001. ISSN 0004-0010.
- [21] Kai Xu and Nabil Simaan. Intrinsic wrench estimation and its performance index for multisegment continuum robots. *Robotics, IEEE Transactions on*, 26(3):555–561, 2010.
- [22] Kai Xu, Roger E Goldman, Jienan Ding, Peter K Allen, Dennis L Fowler, and Nabil Simaan. System design of an insertable robotic effector platform for single port access (spa) surgery. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 5546–5552. IEEE, 2009.
- [23] Tomonori Yamamoto, Balazs Vagvolgyi, Kamini Balaji, Louis L Whitcomb, and Allison M Okamura. Tissue property estimation and graphical display for teleoperated robot-assisted surgery. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, pages 4239–4245. IEEE, 2009.