

Using Bayesian Optimization to Guide Probing of a Flexible Environment for Simultaneous Registration and Stiffness Mapping

Elif Ayvali¹, Rangaprasad Arun Srivatsan¹, Long Wang², Rajarshi Roy², Nabil Simaan² and Howie Choset¹

Abstract—One of the goals of computer-aided surgery is to register intraoperative data to preoperative model of the anatomy, and hence add complementary information that can facilitate the task of surgical navigation. In this context, mechanical palpation can reveal critical anatomical features such as arteries and cancerous lumps which are stiffer than the surrounding tissue. This work uses position and force measurements obtained during mechanical palpation for registration and stiffness mapping. Prior approaches, including our own, exhaustively palpated the entire organ to achieve this goal. To overcome the costly palpation of the entire organ, a Bayesian optimization framework is introduced to guide the end effector to palpate stiff regions while simultaneously updating the registration of the end effector to an *a priori* geometric model of the organ, hence enabling the fusion of intraoperative data into the *a priori* model obtained through imaging. This new framework uses Gaussian processes to model the stiffness distribution and Bayesian optimization to direct where to sample next for maximum information gain. The proposed method was evaluated with experimental data obtained using a Cartesian robot interacting with a silicone organ model and an *ex vivo* porcine liver.

I. INTRODUCTION

Surgeons performing minimally invasive surgery (MIS) offer their patients a shorter recovery time and reduced pain at a cost of increased technical difficulty. One of the drawbacks of MIS is the loss of direct sensory feedback. This loss can impede the detection and use of surface and stiffness features which can help the surgeon find correspondences between the intraoperative scene and the preoperative imaging data. Computer-aided surgery (CAS) was introduced to provide surgeons performing MIS with functional and geometric information that can aid intraoperative navigation and execution of the preoperative plans. Mechanical exploration by palpation and tissue manipulation can provide complementary information about geometric constraints [1], tissue characteristics [2] and variation in stiffness throughout the organ [3], thus augmenting the information obtained through conventional image-based CAS. Li et al. provide a recent review on intraoperative tumor localization in robot-assisted minimally invasive surgery which points out to laboratory experiments and in vivo results [4]. The force and tactile sensing technologies designed for MIS are reviewed

in [5]. These works demonstrate that robot-assisted mechanical palpation can facilitate the localization of arteries and other anatomically critical features, which visually cannot be detected, in MIS. Accurate registration of the surgical tool to the coordinate frame of the preoperative model is one of the primary goals in incorporating intraoperative data, such as stiffness, into an *a priori* geometric model of the anatomy. In this paper, we introduce tools for efficient information-guided exploration using mechanical palpation and annotation of an *a priori* geometric model with complementary information.

Several groups previously proposed heuristic algorithms for autonomous exploration [6] and segmentation of stiff features in compliant environments [7]. These methods rely on the robot visiting locations on a discretized grid on the organ's surface. If a rapid change in stiffness is observed between two grid points, the density of the grids around that location is increased to collect more data in an effort to localize a stiff feature. Grid maps comes with the assumption of independence between the grid points, thus ignore the structural dependency in the environment. Additionally, aforementioned previous works on palpation and exploration are not concerned with incorporating the stiffness information into the preoperative geometric data, a task that can help the surgeon in correlating preoperative models, like CT scans, with the current intraoperative scene.

To avoid the costly palpation of the entire organ and enable fusion of the stiffness information into an *a priori* geometric model of the anatomy, we introduce a method based on Bayesian optimization that minimizes the amount of probing required to reveal stiff features and registers the tool to the *a priori* model. The proposed algorithm is advantageous because it only visits regions that bring information gain, contrary to searching a discretized grid on the entire surface of the organ to find stiff features. In the Bayesian formulation, a Gaussian process (GP) [8] is used to define a prior over the unknown stiffness distribution. GP provides a probabilistic description of the stiffness map and captures the variance of the stiffness distribution which helps guide the probing towards unexplored regions. The stiff regions correspond to regions near local maxima of the stiffness distribution and the Bayesian optimization finds the maxima of this unknown stiffness distribution by directing the probing to points that would result in maximum information gain in predicting the stiff regions.

We first introduce GP and Bayesian optimization in Section II, followed by the description of simultaneous registration and stiffness estimation. The experimental evaluations and results are given in Section III.

*This work was supported by NRI Large grant IIS-1426655 and IIS-1327566.

¹E. Ayvali, R. A. Srivatsan, H. Choset are with the Robotics Institute at Carnegie Mellon University, Pittsburgh, PA 15213, USA (eayvali@, rarunsrivatsan@, choset@cs.) cmu.edu

²L. Wang, R. Roy and N. Simaan are with the Dept. of Mechanical Eng., Vanderbilt University, Nashville, TN 37235, USA (long.wang, rajarshi.roy, nabil.simaan) @vanderbilt.edu

II. BACKGROUND

A. Gaussian Processes

A Gaussian process is a collection of random variables, $\{f(x) : x \in \mathcal{X}\}$, indexed by elements from a set \mathcal{X} such that any finite subcollection of random variables has a multivariate Gaussian distribution [8]. A GP, $f(x) \sim \mathcal{GP}(\mu, k)$, is fully specified by its mean function $\mu : \mathcal{X} \rightarrow \mathbb{R}$ and a covariance function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$. Intuitively, we can think of GP as a distribution over functions. The distribution of function values at a point $x_i \in \mathcal{X}$ is represented by a random variable, Y_i . Gaussian processes can be used for regression and to make predictions at a new point $x^* \in \mathcal{X}$ by defining a prior over functions. In machine learning literature, \mathbf{x} is called the training set and \mathbf{x}^* is called the prediction or test set [8]. For simplicity, the prior mean function is generally assumed to be zero mean.

Given a set of n observed inputs $\mathbf{x} = [x_1, x_2, \dots, x_n]^T$ and corresponding outputs $\mathbf{Y} = [Y_1, Y_2, \dots, Y_n]^T$, the random variables \mathbf{Y} are Gaussian distributed with mean $[\mu(x_1), \mu(x_2), \dots, \mu(x_n)]^T$ and covariance matrix \mathbf{K} whose elements are defined by a covariance function, $k(x_i, x_j)$ where $i, j \in [1, \dots, n]$. A commonly used covariance function is the squared exponential kernel defined as

$$k(x_i, x_j) = \sigma_f \exp\left(-\frac{\|x_i - x_j\|^2}{2\ell^2}\right), \quad (1)$$

where σ_f is the variance of the process and ℓ is the length-scale of the kernel. The conditional (predictive) distribution \mathbf{Y}^* can be computed using the conditioning rule of multivariate Gaussian distributions

$$p(\mathbf{Y}^* | \mathbf{Y}) \sim \mathcal{N}(\mathbf{K}_* \mathbf{K}^{-1} \mathbf{Y}, \mathbf{K}_{**} - \mathbf{K}_* \mathbf{K}^{-1} \mathbf{K}_*^T), \quad (2)$$

where \mathbf{K}_* is the $m \times n$ training-prediction set covariance $\mathbf{K}(\mathbf{x}^*, \mathbf{x})$ and \mathbf{K}_{**} is the $m \times m$ prediction set covariance matrix $\mathbf{K}(\mathbf{x}^*, \mathbf{x}^*)$.

We employ GP to model the stiffness distribution of the organ. The values of \mathbf{Y} are the stiffness values associated with the probed points. The position of the points which are probed form the training set, \mathbf{x} , while \mathbf{x}^* is the prediction grid which spans the surface of the organ. Note that GP is continuous and the prediction grid is only used to plot the stiffness distribution for visualization. By using GP, we assume the stiffness distribution changes smoothly across the organ and this smoothness is defined by the choice of the covariance function. In our formulation, we use a local deformation model for stiffness estimation. Therefore, a kernel with a fixed length-scale that is on the same order of the size of the palpation tool's tip is an effective choice. We use the squared exponential kernel with $\sigma_f = 1$ and $\ell = 3\text{mm}$.

B. Exploration and Exploitation with Bayesian Optimization

Bayesian optimization is a powerful framework for global optimization of black-box functions [9]. It is most beneficial when the function does not have a closed-form expression and obtaining observations from the function is expensive.

Bayesian optimization allows for prior beliefs about an unknown function to be updated via a posterior. Stiffness distribution of an organ can be thought of the unknown function we want to optimize whose maxima correspond to the stiff features. In the Bayesian framework, we use GP to define a prior over the stiffness distribution. The sequential nature of the Bayesian optimization can help guide the sampling of the continuous search space. Sequential sampling requires selecting an acquisition function which uses the mean, $\mu(x)$, and variance, $\sigma(x)$, of the predictive GP posterior to compute a function which shows the most likely locations of the global maximum [10].

Acquisition functions such as probability of improvement, expectation improvement (EI) [9], and upper-confidence based methods [11] were proposed to find the global maximum and to balance exploration with exploitation. EI provides global exploration of the search space and local exploitation. The EI acquisition function can be evaluated analytically and is given as [9]

$$EI(x) = \begin{cases} (\mu(x) - Y^+) \Phi(Z) + \sigma(x) \phi(Z) & \text{if } \sigma(x) > 0 \\ 0 & \text{if } \sigma(x) = 0 \end{cases} \quad (3)$$

where

$$Z = \begin{cases} \frac{(\mu(x) - Y^+)}{\sigma(x)} & \text{if } \sigma(x) > 0 \\ 0 & \text{if } \sigma(x) = 0 \end{cases} \quad (4)$$

Y^+ is the output at x^+ which is the current maximum of the sampled points. $\phi(\cdot)$ and $\Phi(\cdot)$ are the PDF and CDF of the standard normal distribution respectively.

Fig. 1 illustrates an expected improvement scenario for a 1D example. Expectation improvement is generally used to find the global maximum of the unknown function and the search is terminated once a desired improvement is achieved. Note that after finding the global maximum, the algorithm still continues exploring other regions in the next iterations in an effort to reduce uncertainty and to find other local maxima. In addition to using EI, we do pure exploration after every 5 samples and select a random point that has 90% uncertainty. Such exploration is advantageous when there are multiple stiff features and the initial samples do not include

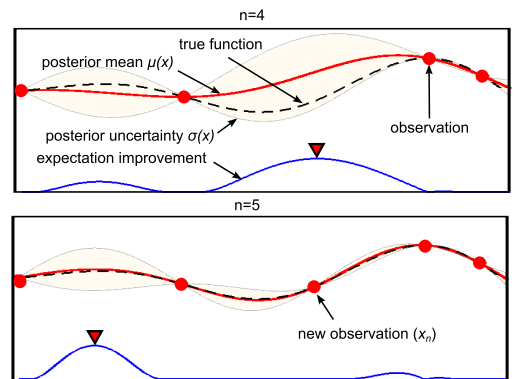


Fig. 1. A 1-D example that starts with an initial training set of 4 points. Red triangle shows the maximum of the EI acquisition function which is the point that should be probed next.

points near stiff regions. Preoperative information such as the size of the tumor or the width of the artery is useful to decide on the density of the samples in the initial set, however this is not the explored in this work. Interested readers can refer to [12] for discussions on the effect of the initial sample size in GP predictions.

Fig. 2 demonstrates the stiffness map obtained using uniform sampling in comparison with the map sampled using EI for two different cases. With EI we can acquire useful information about the stiffness distribution compared to uniformly sampling the surface of the organ for the examples shown.

C. Registration and Stiffness Estimation

Our group has previously developed a method for simultaneously estimating the registration and stiffness distribution over the surface of a flexible environment using a Kalman filtering approach called CARE [13] and a more recent model update method, called Complementary Model Update (CMU), that decouples stiffness estimation from registration, resulting in a more robust implementation [14]. Similar to CARE, the CMU uses the force and position information obtained by interaction of the surgical tool with the organ to estimate the local stiffness and to register the organ to its preoperative model. It is assumed that the local surface deformations are only due to physical interaction of the surgical tool with the organ. We also assume that the surface of the organ is smooth and frictionless, thus the applied force is along the surface normal and increases with depth. Registration is performed by finding the transformation that takes the probed points defined in the tool frame to the corresponding points on the preoperative model of the anatomy. The preoperative model is a computer-aided design (CAD) model in the form of a triangular mesh.

The stiffness at a probed point is estimated using a best line fit between the relative sensed deformation depths and




Fig. 2. Two examples that show the advantage of using EI as a sampling strategy: (a) and (b) show the ground truth stiffness map for a simulated data with an artery and simulated data with multiple stiff features, respectively; (c) and (d) show the maps obtained using uniformly sampling with 100 points; (e) and (f) show the corresponding maps obtained using EI with 100 samples.

sensed forces

$$c_i = L \left(\left\| (\mathbf{p}_\beta^R)_i - (\mathbf{p}_\gamma^R)_i \right\|, ((F_\beta)_i - (F_\gamma)_i) \right), \quad (5)$$

where c_i is the stiffness of the i^{th} probed point, $(\mathbf{p}_\beta^R)_i$ and $(\mathbf{p}_\gamma^R)_i$ are the coordinates of two distinct sensed positions expressed in the tool reference frame, R , corresponding to the i^{th} probed point on the surface of the organ and $(F_\beta)_i, (F_\gamma)_i$ are the corresponding magnitude of sensed forces. The registration estimate is obtained using the estimated stiffness, sensed position and magnitude of the sensed force [14]

$$\mathbf{T} = \operatorname{argmin}_{\mathbf{T}} \sum_{i=1}^n \left\| \mathbf{p}_i^C - \frac{\mathbf{n}_i^C (F_\beta)_i}{c_i} - \mathbf{T} (\mathbf{p}_\beta^R)_i \right\|, \quad (6)$$

where $\mathbf{T} \in SE(3)$ is the homogeneous transformation that is to be estimated, \mathbf{p}_i^C is the location of the probed point in the model's reference frame, given by C , and the surface normal at the probed point is denoted by \mathbf{n}_i^C .

III. EXPERIMENTS AND RESULTS

A. Experimental setup

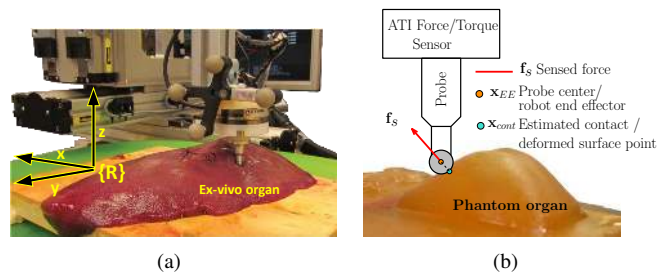


Fig. 3. (a) Experimental setup used for *ex vivo* organ experiments. The tool frame, R , is located at the end effector, (b) Contact location and surface norm estimation

A Cartesian robot with an open architecture controller was used to evaluate the framework proposed in this paper as shown in Fig. 3(a). The robot is equipped with an ATI Nano43 6-axis force sensor at the robot end-effector (EE) and is capable of executing the hybrid motion/force tasks described in Khatib [15]. We assume that the robot end effector has been positioned above the organ. We do not assume any prior knowledge of the local surface normal.

Experiments were conducted with a silicon phantom organ and an *ex vivo* porcine liver. In the silicon phantom experiment the top surface of the organ was lubricated and in the *ex vivo* experiment the organ surface was hydrated to reduce friction during probing. In both experiments, a target region was defined by the user. This region was then used as a reference to generate a uniformly distributed grid map with uniform spacing in the x-y plane of the tool's reference frame, R . For a particular reference probing location, the following procedures were repeated automatically to obtain the force/deformation profile data.

- 1) *Making high force contact*: The robot controller is given a desired probing location \mathbf{x}_p and a force magnitude of 0.5N. The hybrid force/motion controller decouples

the combined commands into compatible (orthogonal) force/motion commands that direct the robot to achieve a desired position in the x-y plane and a stable contact force with the environment along the z axis.

- 2) *Estimating surface norm*: The contact location and surface normal estimation is shown in Fig. 3(b). The surface normal, \mathbf{n} , is computed using the force sensed from the environment, $\mathbf{n} = \mathbf{f}_s / \|\mathbf{f}_s\|$, assuming the surface friction is negligible.
- 3) *Finding low force surface contact point*: In this step, the robot first retreats swiftly away from the surface and then moves towards the surface, along the surface normal this time, to find the zero force intersection. An offset is applied from the robot EE to obtain the estimated contact point as $\mathbf{x}_{\text{cont}} = \mathbf{x}_{\text{EE}} - \mathbf{n}r$ where r is the radius of the robot end effector ball. A 9mm radius probe was used in the silicon examples and a 6.3mm radius probe was used in the *ex vivo* organ experiment.
- 4) *Probing and recording*: The robot is commanded in position control mode along the estimated surface normal direction with 0.3mm increments until 3mm probing depth. Hence, there are 10 position measurements, \mathbf{p} , and 10 force measurements, F , for each point we probe.

B. Bayesian Optimization Guided Probing

To evaluate Bayesian optimization guided probing, we simulated experiments using the experimentally collected data. A block diagram description of the probing method is shown in Fig. 4. Prediction of the stiffness distribution is carried out in the tool frame, R . Initially, the actual registration is unknown, hence we do not know where the probed points correspond to in the preoperative CAD model. We assume palpation is carried out inside a region of interest (ROI). It is assumed that the initial set of samples, \mathbf{x}_0^R , include points on the boundary of the ROI as well as uniformly distributed points inside the ROI. The training set consists of previously probed points, \mathbf{x}_i^R , where $i = 1, 2, \dots, n$. We use gridfit function [16] to interpolate the previously probed points to form a dense grid inside the ROI, \mathbf{x}_j^{*R} where $j = 1, 2, \dots, m$, to make predictions using GP. This grid is used to estimate the stiffness distribution for visualization. The stiffness value at a probed point, c_i , is estimated, in our case by the CMU, and corresponds to Y_i^R at \mathbf{x}_i^R in the GP formulation. Based on the posterior of the predictive distribution given by μ and σ the point at which the expectation improvement takes a maximum value is selected to be the next palpation point, \mathbf{x}_{n+1}^R . As a new point is palpated, it is added to the training set and the prediction grid is regenerated. We use the updated registration estimate, T , to transform the probed points and their associated stiffness values to the corresponding points on the CAD model. This procedure enables displaying experimentally collected stiffness data on the preoperative CAD model. On Matlab, predicting the stiffness distribution and determining the next point to palpate takes 0.11s on average while updating the registration takes 1.2s on average for 100 probed points.

C. Evaluation

The proposed method was evaluated with various stiffness distributions and in the presence of measurement noise. We test the algorithm for four different scenarios:

1. Simulated data of an organ with a multimodal stiffness distribution.
2. Simulated data of an organ with a perturbed multimodal stiffness distribution and sensor noise.
3. Silicon phantom organ with an embedded mock artery.
4. *Ex vivo* porcine liver with a stiff inclusion.

The goal of Example 1 is to demonstrate that expectation improvement is effective in finding all the local maxima and not just the global maximum. We start with an initial set of 19 samples and terminate the palpation after 100 points. Fig. 5(a) shows where we think the position of all the probed points are based on the initial registration guess and their registered position estimated by the CMU algorithm. The registered position of the probed points (deformed points) lie below the surface of the CAD model as expected; validating that the registration estimate is accurate. Fig. 5(b) and (c) show the ground truth stiffness map and the predicted stiffness map, respectively. The predicted stiffness map captures the stiff features present in the ground truth stiffness map.

Example 2 shows the effect of noisy sensor measurements in the prediction of the stiffness distribution. The ground truth stiffness map of Example 2 was obtained by perturbing the stiffness distribution for Example 1. An artificial sensor noise with $(\mu_x, \sigma_x) = (0, 0.3\text{mm})$ was added to the sensed position and $(\mu_F, \sigma_F) = (0, 0.1\text{N})$ was added to the sensed force to simulate a more realistic scenario. Fig. 6(a) shows the registration results. Fig. 6(b) and (c) show the ground truth stiffness map and the predicted stiffness map, respectively. In the presence of noisy sensor measurements, the algorithm still reveals the stiff features and the registration parameters converge to the correct values.

The silicon organ used in Example 3 and the *ex vivo* organ used in Example 4 are shown in Fig. 7(a) and (b), correspondingly. The ground truth stiffness distribution for Example 3 and for the ROI in Example 4 were obtained by interpolating the experimental data at the grid locations shown in Fig. 7(c) and (d) and are shown in Fig. 7(e) and (f), respectively. We emphasize that the organ was discretized only to palpate each grid point with the robot for the purpose of generating a ground truth stiffness map of the organ




Fig. 4. Block diagram description of the Bayesian optimization guided probing

to test our algorithm. Fig. 8(b) shows that sensed force is proportional to the depth of probing for the three locations shown on the CAD model of the liver in Fig. 8(a), validating that the linear stiffness is a valid assumption for 3mm probing depth.

Fig. 9 shows the successful registration and the estimated stiffness map that reveals the mock artery in the silicon organ. Fig. 10 shows the registration result and the position of the embedded triangle overlaid on the estimated stiffness map of the *ex vivo* porcine liver. The actual registration parameters, the estimated registration parameters and the root mean square (RMS) error between the estimated location of all the probed points and their true positions are shown in Table I for all examples. The required clinical accuracy of registration is application dependent. For tumor localization, the registration accuracy can be dictated by the size of the smallest tumor that needs to be resected or ablated [17]. For example, the smallest tumor size detected in livers is 1cm on average [18]. Hence, we conclude that the registration accuracy we achieve is well above the requirements.

TABLE I
REGISTRATION RESULTS

Example	x(mm)	y(mm)	z(mm)	θ_x (deg)	θ_y (deg)	θ_z (deg)	RMS(mm)
Actual 1-3	5	10	-15	11.46	-11.46	5.73	–
Actual 4	5	7	-13	11.45	-5.72	8.59	–
Guess 1-4	0	0	0	0	0	0	–
CMU 1	4.51	9.56	-15.03	12.57	-11.42	5.53	0.91
CMU 2	4.31	9.55	-14.97	12.53	-11.42	5.53	1.04
CMU 3	4.56	9.24	-14.97	12.58	-11.27	5.69	1.14
CMU 4	5.78	6.4	-13.04	11.84	-5.50	8.66	0.74

IV. CONCLUSIONS

This work introduced a probabilistic estimation of the stiffness distribution of the organ using Gaussian processes




Fig. 5. The algorithm starts with an initial set of 19 points: 4 corners and 15 uniformly spaced points. The results are shown for 119 probed points. Example 1: (a) Registration results, (b) Ground truth stiffness map, (c) Estimated stiffness map

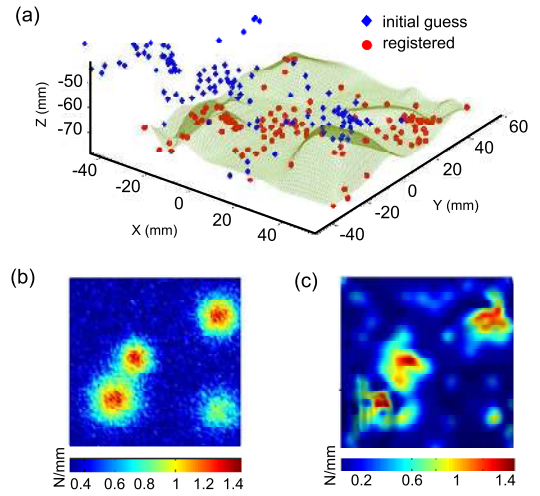


Fig. 6. Example 2: (a) Registration results, (b) Ground truth stiffness map, (c) Estimated stiffness map. The algorithm starts with an initial set of 19 points: 4 corners and 15 uniformly spaced points. The results are shown for 119 probed points.

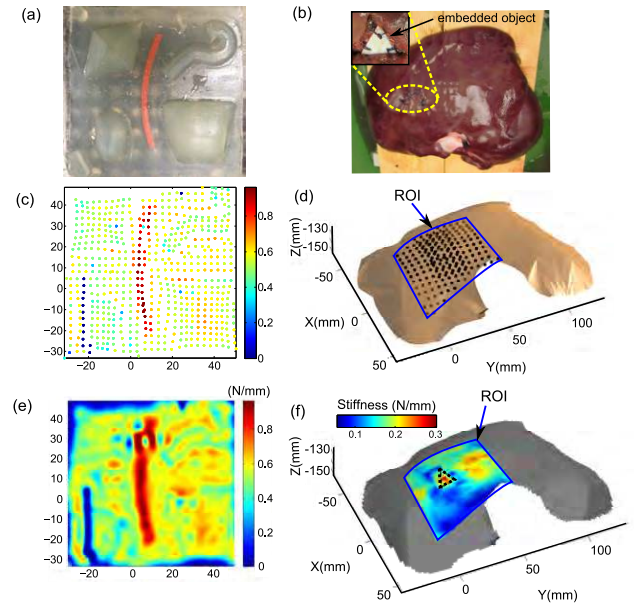


Fig. 7. (a) Silicon phantom organ (b) *ex vivo* porcine liver with an inclusion sutured inside the organ, (c) 619 points were probed on the organ to generate a ground truth for the silicon organ, (d) 196 points were probed on the *ex vivo* organ to generate a ground truth stiffness map, (e) Stiffness map used as the ground truth for the silicon organ, (f) Stiffness map used as the ground truth for the *ex vivo* organ




Fig. 8. (a) CAD model of the organ showing three locations with a square, star and a circle, (b) Force vs. probing depth for the three locations shown on the CAD model.




Fig. 9. (a) Registration results for Example 3, (b) Estimated stiffness map for Example 3 with 119 probed points. The algorithm starts with an initial grid of 19 points: 4 corners and 15 uniformly spaced points.




Fig. 10. (a) Registration results for Example 4, (b) Estimated stiffness map for Example 4 with 79 probed points. The initial set for Example 4 consists of 7 samples that define the boundary of ROI and 12 uniformly distributed samples inside the ROI.

and the use of Bayesian optimization to direct the probing for maximum information gain. We believe fusing intraoperative data into the preoperative model is important to alleviate the limited situational awareness in MIS. The performance of the method was demonstrated by a number of examples and the results show that information-guided probing can avoid probing the entire organ and successfully reveal the stiff regions while registering the tool to the *a priori* geometric model of the organ.

There are several directions for future work. We used a simple experimental setup and an unconstrained environment to evaluate our method to avoid additional sources of error such as workspace limitation and deflection of the robot. In our future work, we will demonstrate the proposed method in real-time using a continuum robot [19] and the da Vinci Research Kit. Another extension of this work is the intraoperative reconstruction of the organ surface and stiffness features as the organ goes through changes during the surgical procedure. We envision that the information-guided probing will enable generation of an updated model of the visible anatomy and reduce the time it takes to reconstruct the intraoperative scene.

ACKNOWLEDGMENT

In a complementary effort, our collaborators are exploring different GP formulations to estimate surface geometry and stiffness using continuous palpation motion [20].

REFERENCES

[1] P. Boonvisut and M. C. Çavuşoğlu, "Identification and active exploration of deformable object boundary constraints through robotic manipulation," *The International Journal of Robotics Research*, vol. 33, no. 11, pp. 1446–1461, 2014.

[2] J. Palacio-Torralba, S. Hammer, D. W. Good, S. A. McNeill, G. D. Stewart, R. L. Reuben, and Y. Chen, "Quantitative diagnostics of soft tissue through viscoelastic characterization using time-based instrumented palpation," *Journal of the Mechanical Behavior of Biomedical Materials*, vol. 41, pp. 149–160, 2015.

[3] A. Talasz and R. V. Patel, "Integration of force reflection with tactile sensing for minimally invasive robotics-assisted tumor localization," *IEEE Transactions on Haptics*, vol. 6, no. 2, pp. 217–228, 2013.

[4] M. Li, H. Liu, A. Jiang, L. D. Seneviratne, P. Dasgupta, K. Althoefer, and H. Wurdemann, "Intra-operative tumour localisation in robot-assisted minimally invasive surgery: A review," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, pp. 509–522, 2014.

[5] P. Puangmali, K. Althoefer, L. D. Seneviratne, D. Murphy, and P. Dasgupta, "State-of-the-Art in Force and Tactile Sensing for Minimally Invasive Surgery," *IEEE Sensors Journal*, vol. 8, no. 4, pp. 371–381, Apr. 2008.

[6] R. E. Goldman, A. Bajo, and N. Simaan, "Algorithms for autonomous exploration and estimation in compliant environments," *Robotica*, vol. 31, no. 01, pp. 71–87, 2013.

[7] K. Nichols, A. M. Okamura, *et al.*, "Autonomous robotic palpation: Machine learning techniques to identify hard inclusions in soft tissues," in *2013 IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 4384–4389.

[8] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[9] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *Journal of Global optimization*, vol. 13, no. 4, pp. 455–492, 1998.

[10] E. Brochu, V. M. Cora, and N. De Freitas, "A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning," *arXiv preprint arXiv:1012.2599*, 2010.

[11] N. Srinivas, A. Krause, M. Seeger, and S. M. Kakade, "Gaussian process optimization in the bandit setting: No regret and experimental design," in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 2010, pp. 1015–1022.

[12] A. Söbester, S. J. Leary, and A. J. Keane, "On the design of optimization strategies based on global response surface approximation models," *Journal of Global Optimization*, vol. 33, no. 1, 2005.

[13] S. Sanan, S. Tully, A. Bajo, N. Simaan, and H. Choset, "Simultaneous compliance and registration estimation for robotic surgery," in *Proceedings of the Robotics: Science and Systems Conference*, 2014.

[14] R. A. Srivatsan, E. Ayvali, L. Wang, R. Roy, N. Simaan, and H. Choset, "Complementary model update: A method for simultaneous registration and stiffness mapping in flexible environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.

[15] O. Khatib, "A unified approach for motion and force control of robot manipulators: The operational space formulation," *IEEE Journal of Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.

[16] J. D'Errico, *Surface Fitting using gridfit*. MATLAB Central File Exchange, 2005, [Retrieved March 2015].

[17] C. Linte, J. Moore, and T. Peters, "How accurate is accurate enough? a brief overview on accuracy considerations in image-guided cardiac interventions," in *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Aug 2010, pp. 2313–2316.

[18] Y. Horiguchi, B. Sekoguchi, H. Imai, T. Suzuki, H. Kubo, H. Itoh, and M. Itoh, "Treatment of choice for unresectable small liver cancer: Percutaneous ethanol injection therapy or transarterial chemoembolization therapy," *Cancer chemotherapy and pharmacology*, vol. 33, no. 1, pp. S111–S114, 1994.

[19] S. Tully, A. Bajo, G. Kantor, H. Choset, and N. Simaan, "Constrained filtering with contact detection data for the localization and registration of continuum robots in flexible environments," in *2012 IEEE International Conference on Robotics and Automation*, 2012, pp. 3388–3394.

[20] P. Chalasani, N. Simaan, L. Wang, R. H. Taylor, R. Roy, and M. Kobilarov, "Concurrent nonparametric estimation of organ geometry and tissue stiffness using continuous adaptive palpation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.