2016 IEEE International Conference on Automation Science and Engineering (CASE)
Fort Worth, TX, USA, August 21-24, 2016

Tumor Localization using Automated Palpation with Gaussian Process Adaptive Sampling

Animesh Garg,[1,2] Siddarth Sen,[2] Rishi Kapadia,[2] Yiming Jen,[2]
Stephen McKinley,[3] Lauren Miller,[2] Ken Goldberg[1,2]

***Abstract*— In surgical tumor removal, inaccurate localization can lead to removal of excessive healthy tissue and failure to completely remove cancerous tissue. Automated palpation with a tactile sensor has the potential to precisely estimate the geometry of embedded tumors during robot-assisted minimally invasive surgery (RMIS). We formulate tumor boundary localization as a Bayesian optimization model along implicit curves over estimated tissue stiffness. We propose a Gaussian Process Adaptive Sampling algorithm called Implicit Level Set Upper Confidence Bound (ILS-UCB), that prioritizes sampling near a level set of the estimate. We compare the ILS-UCB algorithm to two alternative palpation algorithms: (1) Expected Variance Reduction (EVR), which emphasizes exploration by minimizing variance, and (2) Upper Confidence Bound (UCB), which balances exploration with exploitation using only the estimated mean. We compare these algorithms in simulated experiments varying the levels of measurement noise and bias. We find that ILS-UCB significantly outperforms the other two algorithms as measured by the symmetric difference between tumor boundary estimate and ground truth, reducing error by up to $10x$. Physical experiments on a dVRK show that our approach can localize the tumor boundary with approximately the same accuracy as a dense raster scan while requiring at least $10x$ fewer measurements.**

I. Introduction

Palpation—using the sense of touch to examine part of the body or organ—is frequently used during surgery for in-situ assessment and localization of cancerous tissue for diagnosis or tumor resection. During open surgery, a surgeon can directly palpate tissue to identify and localize subsurface structures or tumors based on changes in tissue stiffness relative to the surrounding substrate.

Robotic Surgical Assistants (RSAs), such as Intuitive Surgical's da Vinci system have been shown to facilitate precise minimally invasive surgery [4, 27] by providing increased dexterity and control to the surgeon. While there have been advances in providing haptic feedback for Robot-assisted minimally invasive surgery (RMIS) [16], RSAs used in clinic still largely lack haptic sensing. During an RMIS procedure, the operating surgeon depends primarily on vision for complex tasks such as tumor localization and resection.

In spite of the increasing use of RMIS, the lack of haptic perception as compared to open surgery can potentially increase risk of tissue damage [5] and likelihood of incomplete removal of cancer cells [28]. Moreover, RMIS procedures are primarily controlled by surgeons through local teleoperation (master-slave with negligible time delays). Introducing autonomous surgical sub-tasks such as tumor localization with palpation has the potential to assist surgeons, reduce fatigue & cognitive load, and facilitate supervised autonomy.

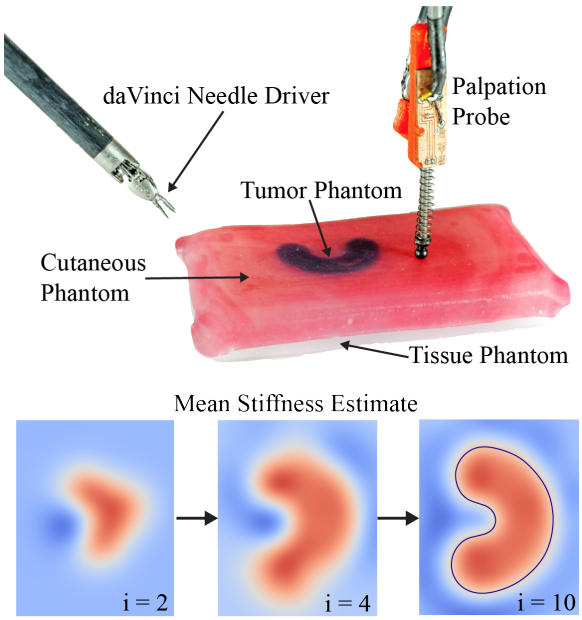

Fig. 1: This figure illustrates autonomous localization of an embedded tumor in a tissue phantom. The top image shows the experimental setup with a palpation probe mounted on a Da Vinci Research Kit (dVRK). The sequence of images on the bottom illustrates the progression of the estimated stiffness over the tissue surface at intermediate stages (i=2, 4) and also the final estimate (i=10), with the estimated tumor boundary shown as a black line.

Automation of tumor resection requires a high confidence estimate of both the organ geometry and the tumor boundary [18]. Organ geometry uncertainty can result in incorrect stiffness estimates and uncertainty in tumor (or cyst) location can result in imprecise incision. A negative margin could result in a conservative estimate (cutting out excessive healthy tissue), while a positive margin could result in undercutting the tumor and spreading the cells.

In this paper, we propose an algorithm for autonomous tumor localization with palpation in RMIS for a given organ geometry. We pose tumor boundary localization as a Bayesian optimization problem along implicit curves defined by a Gaussian Process (GP) representing estimated tissue stiffness. Our approach focuses on reducing uncertainty

All authors are affiliated to AUTOLAB at UC Berkeley. [1] IEOR, [2] EECS, [3] ME, University of California, Berkeley, CA USA;
{animesh.garg, siddarthsen, rishikapadia, yjen, mckinley, laurenm, goldberg}@berkeley.edu

along level sets of the surface stiffness as opposed to creating a high-certainty stiffness map for the entire search area. We compare our approach to two other palpation algorithms for mapping subsurface stiffness: one which prioritizes exploration alone, and the other which balances exploration and exploitation of high stiffness areas, as opposed to along implicit curves.

We evaluate the symmetric difference in the boundary estimate obtained with the ground truth in all three cases as described in Section VI. Simulation results suggest that as compared to EVR and UCB, ILS-UCB converges to the tumor boundary more quickly and can accommodate higher levels in measurement noise and bias (see Section VI). In addition to evaluation in simulation, we demonstrate preliminary results for autonomous segmentation of subcutaneous tumor in soft tissue phantom with an end-effector mounted palpation sensor presented in our prior work [19] on the da Vinci Research Kit (dVRK) [30] in Section VII. Initial results on dVRK suggest that our approach can localize the tumor boundary with approximately the same accuracy while requiring at least $10x$ fewer measurements than uniform raster search.

II. Background and Related Work

Palpation sensors are a subclass of haptic sensors that mimic the biological sense of cutaneous touch. In RMIS, palpation sensors can allow the surgeon to estimate relative tissue stiffness at the surface and adjust forces for safer tissue manipulation. Haptic feedback can be obtained by using a number of transduction principles [16]. See Girão et al. [9] and Tiwana et al. [26] for surveys of existing haptic and tactile feedback devices for robotic and biomedical applications. Konstantinova et al. [16] provide an excellent treatise on haptic feedback devices in RMIS.

Active Search and Mapping. Literature from automating grasping and grasp planning, for example, examines the problem of estimating and refining 2D and 3D estimates of objects using different sensing modalities [2, 7, 12]. Our approach draws on prior work from adaptively estimating shapes and curves using noisy measurements. Dragiev et al. [7] represent the shape of 3D objects using Gaussian Process Implicit Surfaces (GPIS), and their algorithm explores the shape estimate by attempting grasps in areas with highest variance along the implicit surface, using information from failed grasp attempts (missing or unexpected contacts) to refine the GPIS estimate. In Hollinger et al. [14], coverage-based inspection paths are planned based on estimated uncertainty over a GPIS model of a ship's hull. The method by Bjorkman et al. [2] initializes an object shape estimate using stereo vision, and calculates a GPIS representation of object geometry. The estimate is then refined by iteratively collecting haptic measurements at points along the estimated surface with the highest predicted uncertainty.

Tactile Tumor Localization. A number of recent works focus on leveraging haptic/tactile feedback to automate mapping subsurface stiffness or material variation in soft tissue for RMIS procedures. Goldman et al. [10] recursively increase spatial resolution in areas where the measurement passes a certain threshold during stiffness mapping. Similarly, Nichols et al. [21] use elastography data from discrete measurements along a grid, train a classifier for stiffness discrimination between tumors and surrounding tissue, and perform local refinement around points identified as boundary points. Ayvali et al. [1], present an algorithm for registration of surface geometry to pre-operative data.

It is worth noting that uncertainty in the stiffness map depends on the accuracy of the organ surface estimate. A few recent studies have studied tactile surface estimation [3, 24], using a high-definition multi-axis force sensor that is unavailable in the form-factor for RMIS. In this paper, we are assuming that we have the surface estimate *a priori*, and hence we can use a uniform measurement noise in the GP update (see eq. (2)).

The approach presented in this paper is most closely related to recent work in active level-set estimation using mobile sensors for environmental modelling [11, 13]. Hitz et al. [13] use a receding horizon path planner to reduce uncertainty specifically around a threshold value for plankton level modeling using aquatic robots. In Gotovos et al. [11], a traveling salesman algorithm is used to plan paths that sample a set of new measurement sites. Both [11] and [13] use a sampling criterion based on the ambiguity of the function value at a particular point being above (or below) a threshold to select subsequent sampling locations. Similar to these approaches, rather than achieving low estimate error everywhere, the palpation strategy in this paper is to focus on regions representing boundaries of level sets, or where a scalar function crosses a specific threshold.

III. Problem Statement

This paper studies the problem of localizing the boundary of a subcutaneous tumor using a palpation probe that provides a measure of effective surface stiffness.

Assumptions. We assume a single, solid, connected 3D tumor is embedded in a volume of tissue. We assume the tumor is within depth d from the surface, resulting in measurable stiffness differences, and that the difference in stiffness measured at the surface due to the embedded tumor is at least Δk. We assume that we have access to the surface of the tissue for probing, the surface geometry is known, and the boundary of the tumor projected on the surface is smooth with an upper bound on the local curvature (κ).

Expected Output. The goal is to estimate the non-parametric curve $\mathcal{C}$ representing the projection of the subcutaneous tumor(s) $\mathcal{T}$ on the soft tissue surface. The curve is defined by level sets of the stiffness $\mathcal{S}(x)$ over the surface parameterized by $x \in \mathbb{R}^2$, measured using the palpation probe. Our algorithm produces a sequence of locations on the surface for palpation.

Evaluation. In our simulation experiments, we assume we have access to the ground truth of the inclusion boundary, $\mathcal{C}_{GT}$. We evaluate the proposed method and compare to other

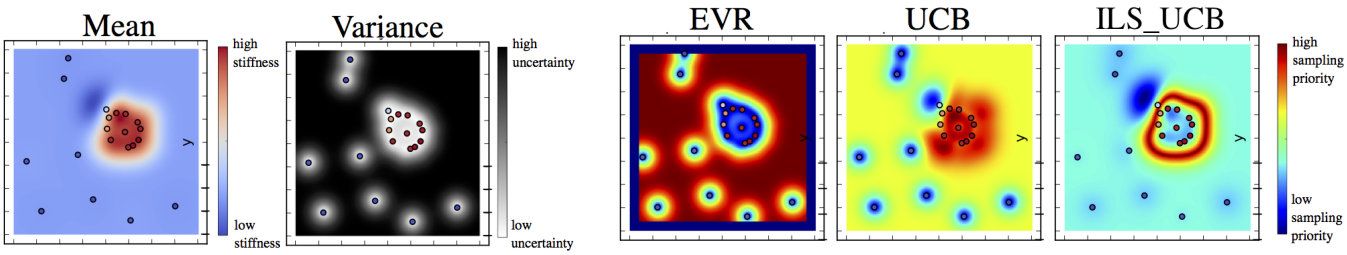

(a) Gaussian Process Mean and Variance Estimate

(b) Three different Palpation Algorithms

Fig. 2: (a) Mean and variance of a circular tumor for Gaussian process estimate after 2 iterations of batch size 10 in simulation. (b) Different palpation algorithms evaluated for the Gaussian Process estimate (a). EVR prioritizes exploration in unsampled regions, UCB prioritizes exploitation of the maximum stiffness areas and uncertainty, and ILS-UCB balances sampling near level sets between the max and minimum values, and uncertainty.

sampling approaches by comparing the symmetric difference between ground truth $\mathcal{C}_{GT}$ and the algorithm estimate $\mathcal{C}$. Note that we penalize both under and overestimation equally.

IV. Gaussian Process Adaptive Sampling

Algorithm overview. The stiffness map estimate $\mathcal{S}(x)$, represented as a Gaussian process (GP), is initialized with measurements collected at randomly selected locations. Based on the current estimate and uncertainty, measurement locations are iteratively selected according to sampling criterion, defined in Section V, to refine the estimate. Measurements are taken, and appended to the set of data used to retrain the Gaussian Process. The updated estimate is then used to select new locations to probe.

Gaussian processes. GPs extend multivariate Gaussian distributions to infinite dimensions [22]. GPs are often used to estimate and model continuous spatial data. GP models provide a smooth estimate everywhere, even given sparse sets of training data, allow multi-modal sensor fusion, and provide a statistical representation of the estimate useful for active refinement using Bayesian optimization methods. We use a GP to represent the stiffness $\mathcal{S}(x)$ and associated uncertainty from the observed palpation measurements.

The input data for a GP is a set of training data D with observations Y taken at states X, or $D = \{X,Y\} = \{(x_1,y_1),...(x_n,y_n)\}$ for a set of n training samples. The GP model assumes a function f is a noisy spatial process $y_i = f(x_i) + \delta$, where $x \in \mathbb{R}^2$, $\delta \sim \mathcal{N}(0,\sigma)$ is additive zero-mean Gaussian measurement noise. Given training data D, the posterior distribution for the function $f(\cdot)$ at new points x^+ is Gaussian with mean μ_{x^+} and variance $\sigma^2_{x^+}$, i.e. $p(f(x^+)|x^+,X,Y) = \mathcal{N}(f(x^+);\mu_{x^+},\sigma^2_{x^+})$, where

$$\mu_{x^+} = k_+^T(K+\sigma_n^2 I)^{-1}Y, \tag{1}$$

$$\sigma^2_{x^+} = k(x^+,x^+) - k(x^+,X)^T(K+\sigma_n^2 I)^{-1}k(x^+,X) \tag{2}$$

Here $k(x^+,X)$ is the $n\times 1$ vector of covariances between x^+ and the n training inputs X, K is the covariance matrix of the inputs X, and σ_n^2 is the noise variance of the additive measurement noise $\delta(.)$. The covariance function (or kernel) K determines the correlation between input locations x_i. We use the squared exponential kernel [22] for the experiments in this paper because they produce very sooomth output, but other kernels such as Matern and Periodic ones may also be used based on the domain.

Tumor Boundaries as Implicit Surfaces. The stiffness $\mathcal{S}(x)$ can be used to define the tumor boundary contour $\mathcal{C}$. Implicit curves are defined by the set of points for which an implicit function—a scalar-valued function defined over $\mathbb{R}^2$—takes on a particular value [8, 6, 17]. We define the tumor boundary ($\mathcal{C}$) as an implicit curve based on the GP representing the stiffness $\mathcal{S}(x)$, i.e. as the α-level set for the stiffness $\mathcal{S}(x)$, such that:

$$\mathcal{S}(x)\begin{cases} = \alpha, & x \text{ on } \mathcal{C} \\ > \alpha, & x \text{ inside } \mathcal{C} \\ < \alpha, & x \text{ outside } \mathcal{C}. \end{cases} \tag{3}$$

Since $S(x)$ is an estimate with associated uncertainty, we define the implicit curve using the GP mean and choose α based on maximum and minimum deflection measurements observed, representing likely tumor boundaries. This relaxes the need to define precise expected stiffness measurements prior to probing.

V. Three Palpation Algorithms

As described in Section IV, an algorithm can balance the exploration and exploitation over the GP representation of the stiffness $\mathcal{S}(x)$. We introduce three palpation algorithms that are evaluated in Section VI. All three iteratively select measurements by optimizing a *sampling criterion* $\mathcal{A}(\mu_{t-1}(x),\sigma_{t-1}(x))$, parametrized by the sufficient statistics of the GP estimate, the mean $\mu(x)$ and variance $\sigma(x)$.

Sampling criterion are often referred to as *acquisition functions* in GP optimization. In all 3 of these cases, we select sampling locations that maximize the criterion over the search space $\mathcal{X}$. Because the variance depends only on the locations of samples (see eq. (2)), not the measurements, one can select sets of sample points that take into account local variance decrease following measurements, prior to taking them. Following [11], we select a set of samples during each iteration of the algorithm using the estimated mean and the known variance from the prior iterations and solve a traveling salesmen problem approximately to plan paths between the selected measurement locations.

1. Expected Variance Reduction (EVR)
The EVR algorithm is a purely *exploratory* approach, selecting sampling points where the variance of the Gaussian process estimate is highest: i.e.

$$x_t = \arg\max_{x \in \mathcal{X}} \quad \sigma_{(t-1)}(x).$$

2. Upper Confidence Bound (UCB)
The UCB palpation algorithm balances exploration, i.e. prioritizing areas with high uncertainty (high GP variance σ) and exploitation i.e. areas where the expected stiffness is high (high GP mean μ):

$$x_t = \arg\max_{x \in \mathcal{X}} \quad \gamma * \mu_{t-1}(x) + (1-\gamma)\sigma_{t-1}(x).$$

Prioritizing high stiffness areas guides sampling toward regions that are likely to be tumor vs. surrounding tissue, and prioritizing high variance regions prioritizes sampling where the confidence bound is very large (there is high uncertainty in the stiffness estimate).

3. Implicit Level Set Upper Confidence Bound (ILS-UCB)
The ILS-UCB algorithm also trades off between exploration and exploitation. This algorithm prioritizes searching the expected tumor boundary, conditional on estimate uncertainty, and does not seek to precisely learn a stiffness map of the entire workspace. Intuitively, by reducing the estimation space to specifically localize the tumor boundary, we can reduce the total number of measurements–and consequently the time–required to achieve an estimate of the boundary.

Hence, rather than prioritizing the both high variance and high mean like UCB, ILS-UCB prioritizes sampling in areas near a level set of the mean represented by the Gaussian process implicit surface, i.e. to minimize the implicit potential defined by $\mu_{t-1}(x) - h_{t-1}$, and where the confidence interval is large:

$$x_t = \arg\max_{x \in \mathcal{X}} \quad (1-\gamma)\sigma_{t-1}(x) - \gamma * |\mu_{t-1}(x) - h_{t-1}|.$$

The level set h_{t-1} is not assumed *a priori*, but is a percentage α of the current estimated mean: $h_{t-1} = \alpha\,(max\mu_{t-1}(x) - min\mu_{t-1}(x))$. Note that the second term in the equation above is negative, as we are trying to sample in locations where the distance to the level set is minimized.

VI. Simulation Experiments

We compare the three palpation algorithms described in Section V in simulation for estimation of two phantoms with known tumor geometry: 1) a circular disk (area $1.23\,cm^2$) and 2) a horseshoe (area $1.26\,cm^2$). The search space is a 2.5×5 cm region. For selection of sampling points, the search area was discretized into a 200 × 200 grid (40,000 points), and γ for the UCB and ILS-UCB algorithms was chosen to be 0.5, which empirically balanced exploitation and exploration well for this application. In the simulations experiments, we use 5 initial measurements and a 50% level set (α) of the stiffness map as the tumor boundary which is indicated as the tumor boundary in Figure 3. We use Python package GPy [25] to implement Gaussian process regression.

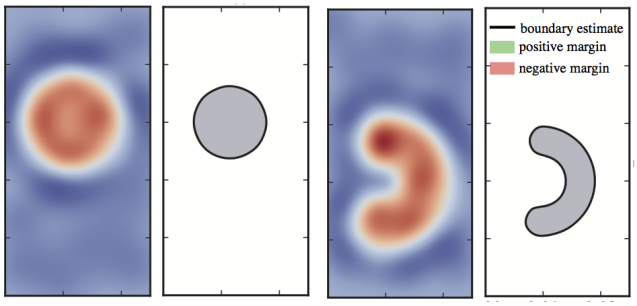

(a) noise=1%, bias=0

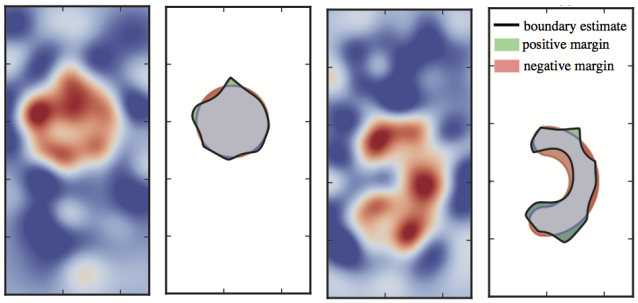

(b) noise=25%, bias=0

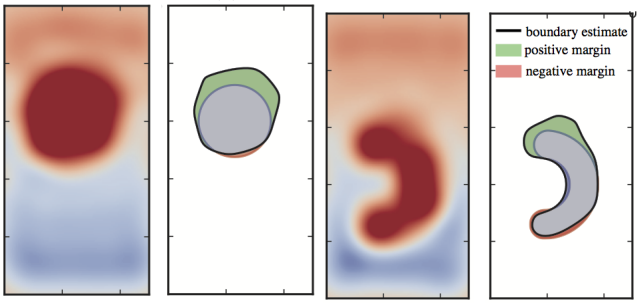

(c) noise=1%, bias=100

Fig. 3: Estimated stiffness maps and boundary estimates for simulated experiments after 10 iterations (of batch size 10, i.e. 100 measurements) using the ILS-UCB algorithm, for different noise levels and measurement bias. Regions in blue denote surrounding tissue with lower stiffness and regions in red denote higher stiffness.

We evaluate performance of all three algorithms varying two possible noise sources: additive measurement noise (σ), and systematic measurement bias (β). The latter arises when, for example, unmodeled deviations in the palpation surface lead to systematic error in the stiffness measurements due to non-constant probe indentation.

Measurements are simulated using a *sigmoidal* model, which approximates the probe measurements made using our a customized sensor, *PALP* [19],

$$Y(x) = Y_{min} + \frac{Y_{max} - Y_{min}}{1 + e^{-k(x-\mathcal{C})}} + \delta + \beta x_1 (Y_{max} - Y_{min}), \qquad (4)$$

where k represents the slope, Y_{min} and Y_{max} the maximum and minimum measurements, and $(x-\mathcal{C})$ is the distance between points x and the closest point on the tumor boundary. δ is additive measurement noise ($\delta \sim \mathcal{N}(0,\sigma)$), and we use a linear measurement bias $\beta x_1 (Y_{max} - Y_{min})$ that increases from 0 along one dimension x_1 (vertical dimension in Figure 3), proportional to a bias constant β.

Table I shows the final boundary error for each palpation

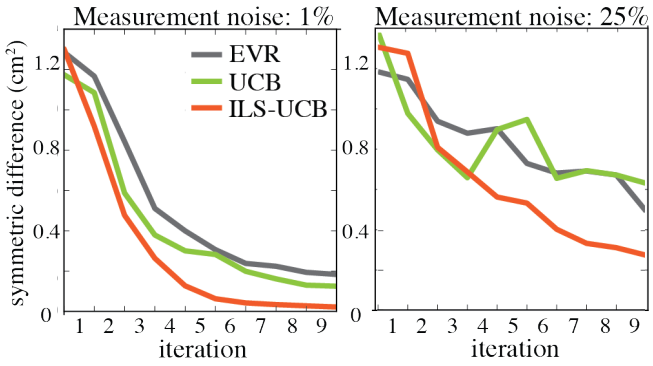

Fig. 4: Simulation Experiments: Convergence of the symmetric difference error (as % of search space area) for palpation algorithms as a function of iteration for two different levels of measurement noise, using a horseshoe shaped tumor in simulation. As the noise increases, there is randomized exploration which results in non-smooth convergence curves in the right graph.

TABLE I: Simulation: Symmetric difference of boundary estimate from the Ground Truth with varying levels of measurement noise and bias in the measurement function (see eq(4)) for each of the three probing algorithms. We use two tumor models, a circular disk shaped tumor (#1) and a horse-shoe shaped tumor (#2). Error is reported as a percentage of the search space area after 10 iterations (with a batch size of 10, i.e. 100 measurements), averaged over five trials. Using the ILS-UCB palpation algorithm outperforms other algorithm in most cases and achieves up to *10x* reduction in error.

Variance (σ)	Tumor#1 Circular Disk			Tumor#2 Horse-Shoe		
	EVR	UCB	ILS-UCB	EVR	UCB	ILS-UCB
1 %	0.840	0.567	**0.056**	1.467	1.001	**0.175**
5 %	0.807	0.672	**0.177**	1.525	1.256	**0.373**
10 %	1.189	0.951	**0.393**	2.155	2.135	**0.749**
25 %	2.610	2.870	**1.314**	3.987	5.116	**2.210**

Bias (β)	Tumor#1 Circular Disk			Tumor#2 Horse-Shoe		
	EVR	UCB	ILS-UCB	EVR	UCB	ILS-UCB
1%	0.759	0.573	**0.060**	1.255	1.093	**0.141**
5%	0.667	0.818	**0.267**	1.460	1.186	**0.305**
10%	5.064	5.085	**3.906**	4.212	4.881	**3.234**
100%	11.084	10.818	**9.810**	**10.084**	10.085	10.091

algorithm and phantom. Each value is the symmetric difference error between the estimated boundary (α is the 50% level set of the mean) and the ground truth tumor boundary, as a percentage of the search area, averaged over five trials. Measurement noise (σ) is shown as a percentage of the difference between the simulated measurement maximum and minimum. Measurement bias constant β is shown as a percentage of the difference between the measurement maximum and minimum.

Figure 3 shows three cases of varying noise and bias levels with corresponding stiffness mean and estimated boundary for both tumors. Figure 4 shows the error for the horse-shoe tumor (tumor 2) as a function of iterations, at two measurement noise values, for all three algorithms. Performance for all three algorithms degrades with increasing noise and measurement bias, and all three algorithms have similarly high error for the highest noise and bias cases.; final error is up to 10% smaller, however, using the ILS-UCB algorithm.

TABLE II: Physical Experiments: Symmetric difference of the Ground Truth from the boundary estimate obtained from Raster Scans and ILS-UCB Adaptive Sampling algorithm. We use two tumor shapes, a circular disk shaped tumor (#1) and a horse-shoe shaped tumor (#2) similar to the Simulation setup. Error is reported as a percentage of the search space area (26.46 cm^2). Raster scan uses approximately 10,000 samples and ILS-UCB uses 200 samples(20 iterations with a batch size of 10). ILS-UCB algorithm achieves the same order of performance with 10*x* fewer samples than Raster scan.

	Tumor#1 Circular Disk	Tumor#2 Horse-shoe
Raster Scan	3.60	3.73
ILS-UCB		
Trial 1	0.92	5.73
Trial 2	2.99	5.68
Trial 3	3.48	6.67
Trial 4	1.01	7.08
Trial 5	0.97	8.65
Mean	**1.87±1.25**	**6.76±1.21**

VII. Physical Experiments

In this section, we evaluate performance of the ILS-UCB algorithm, which outperformed the alternative algorithms in simulation, for tumor localization on a physical phantom using the dVRK robot. We compare the ILS-UCB algorithm to a dense raster scanning strategy, evaluating the total number of measurements made, time, and error between the estimated boundary and the Ground Truth.

dVRK: Hardware and Software. We use the Intuitive Surgical da Vinci Research Kit (dVRK) surgical robot assistant, as in [20, 23]. We interface with the dVRK using open-source electronics and software developed by WPI and Johns Hopkins University [15]. The software system is integrated with ROS and allows direct robot pose control.

Soft Tissue Phantoms. Tumor phantoms were molded from silicone rubber (thickness 4.5 mm; Shore hardness 30A), and are embedded in softer silicon rubber *Ecoflex 00-20* (*Smooth-On*) with a total size of $100 \times 50 \times 20$ mm ($L \times W \times H$) simulating subcutaneous tissue.

Stiffness Measurement with Palpation Probe. We use a customized sensor, *PALP*, presented in our prior work [19]. *PALP* is a low-cost, disposable sensor that mounts on a DVRK classic tool-tip. The PALP probe uses a displacement-based contact sensing mechanism. A spherical indenter of 4.5 mm diameter enables point and sliding palpation. Probe's end-effector has a known spring constant and uses Hall Effect sensor to calculate indentation force by measuring displacement of magnets affixed to the end-effector. The probe-tip displacement (δ_p) relative to the body of the device can be linearly related to a tissue reaction force (F) using Hooke's Law ($F = k\delta_p$), where k is the spring constant ($k = 0.08\,\mathrm{N/mm}$). The sensor provides estimated stiffness values using inverse calculations as in work by Yu et al. [29]. The probe can sense stiffness difference for tumors embedded up to a depth $d = 8$ mm [19].

Tumor Localization. We demonstrate the performance of tumor localization using the ILS-UCB algorithm, which outperformed other methods in simulation in terms of robustness to noise. We perform the localization experiment

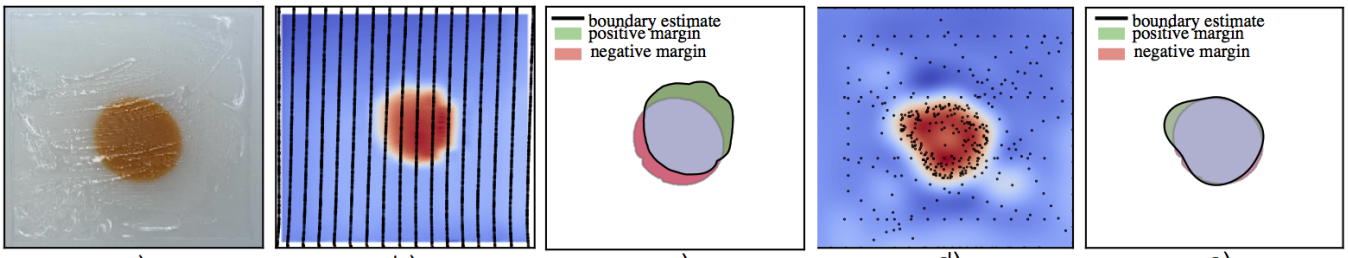

Fig. 5: Tumor boundary estimation for a circular tumor phantom with the PALP probe mounted on the dVRK. (a) shows a top-down view of a molded rubberized tumor (orange) embedded in a silicone matrix (white). (b) shows the estimated stiffness map and the measurement locations for the raster scan (black). (c) shows the estimated boundary (black), with error shown as positive and negative margins for the raster scan. (d) shows the estimated stiffness map measurement locations using ILS-UCB (black), and (e) shows the estimated boundary using ILS-UCB. Raster Scanning in this case takes 9,965 measurements and results in an error of 0.95 cm^2 from Ground Truth. ILS-UCB takes 200 measurements and results in an error of 0.24 cm^2. Note that because ILS-UCB prioritizes sampling along an implicit level set of the stiffness, the measurements in (d) are clustered near the boundary of the tumor, resulting in faster convergence.

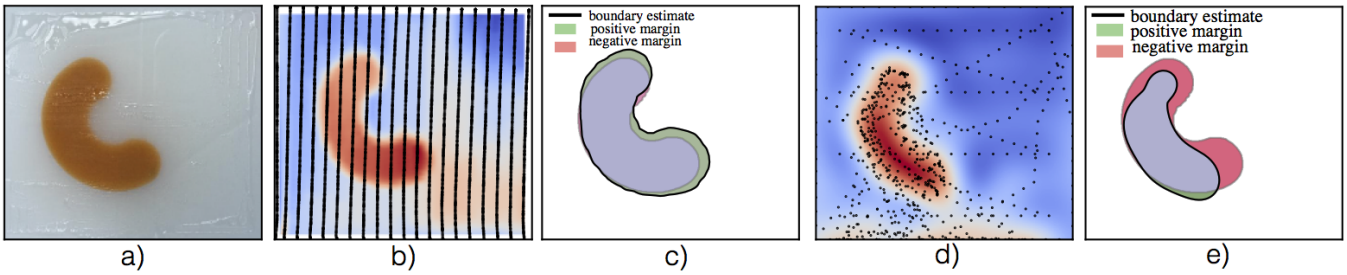

Fig. 6: Tumor boundary estimation for a horseshoe-shaped tumor phantom with the PALP probe mounted on the dVRK. (a) shows a top-down view of a molded rubberized tumor (orange) embedded in a silicone matrix (white). (b) shows the estimated stiffness map and the measurement locations for the raster scan (black). (c) shows the estimated boundary (black), with error shown as positive and negative margins. (d) shows the estimated stiffness map measurement locations using ILS-UCB (black), and (e) shows the estimated boundary using ILS-UCB. Raster Scanning in this case takes 11,774 measurements and results in an error of 0.98 cm^2 from Ground Truth. ILS-UCB takes 200 measurements and results in an error of 1.51 cm^2.

on a circular disk and a horseshoe-shaped tumor similar to the simulation setup. The search space is 5.4×4.9 cm (26.46 cm^2) and the tumor areas are 3.84 cm^2 for the horseshoe and 1.84 cm^2 for the circle.

In these physical experiments, we use 16 samples on a uniform grid across the workspace to initialize the Gaussian process representation. Each tumor localization trial is run for at most 20 iterations with a batch of 10 measurements each iteration. Hence for each trial 200 points are selected for measurement. Selected points were ordered by approximately solving a Travelling Salesman Problem at each iteration. In the interest of reducing computation time, scanning trajectories on the robot were computed by linearly interpolating between the points selected based on the sampling criterion.

At each iteration, the probe moves on the surface between the selected points, continuously collecting measurements. The robot moves at 5mm/s and measurements are collected at 1 sample/mm. Unlike simulation experiments, measurements obtained between the selected points were incorporated in the GP update in physical experiments to speed up convergence.

As a comparison, we perform a continuous raster scan using the robot the same speed with 2.0 mm between rows and obtaining 9,965 measurements for the circle and 11,774 measurements for the horseshoe. Table II details the results from raster scan as well as 5 trials of ILS-UCB on both the circular and horseshoe tumor. As in simulation, the values are the symmetric difference between estimated boundary and ground truth. Ground truth of the tumor boundaries is calculated by registering the search area between the camera image and the robot. The resulting image is then rectified and the tumor is segmented and the area is calculated from the resulting closed contour.

Figures 5a-c and 6a-c show the physical phantom, the raster scan path, and the estimate of the mean stiffness obtained using the raster scanning strategy. Figures 5d-e and 6d-e show the estimate of the mean stiffness obtained after 20 iterations (in batches of 10) using ILS-UCB as the palpation algorithm for the two sets of physical experiments. Estimates of the boundary as well as the positive and negative margins are shown for Figures 5e and 6e. We use the 70% level set of the estimated stiffness map as the tumor boundary ($\alpha = 0.7$) which is indicated as a black line on the image. Note that because ILS-UCB prioritizes sampling along an implicit level set of the stiffness, the measurements in both Figures 5d and 6d are clustered near the boundary of the tumor, resulting in faster convergence.

VIII. Discussion and Future Work

This paper studies tumor boundary estimation for robot-assisted minimally invasive surgery. We propose a Gaussian Process Adaptive Sampling algorithm, ILS-UCB, for

autonomous tumor localization using palpation along implicit curves defined by stiffness measurements. Simulation results show that ILS-UCB can achieve up to a *10x* reduction in boundary estimation error over other methods. We also perform physical experiments using a custom palpation probe on a da Vinci Research kit (dVRK) and observe that algorithmic search along implicit curves requires at least $10x$ fewer measurements than uniform raster scanning.

A limitation of the proposed adaptive sampling method for tumor localization is that a very accurate estimate of the surface geometry is required for a correct estimate of sub-surface stiffness and, by extension, of the tumor boundary. We show in Figure 3(c) that stiffness estimates will be effected by errors in surface estimates. While using contact-based tactile probes such as *PALP*, controlling the applied force is critical in interpreting measurements, in addition to maintaining a constant indentation and performing measurements along surface normals. Our proposed method currently selects points at each iteration and then uses an approximate TSP route through these points. The resulting trajectories visit the selected points but do not optimize for the measurements along the path. Future work will explore direct continuous trajectory optimization with respect to the sampling criterion. We will focus future work on extending the current approach to non-planar surfaces, and to uncertain surface estimates using, e.g., a heteroskedactic GP model to model non-uniform measurement noise.

Acknowledgements: This research was performed with UC Berkeley's AUTOLAB under the UC Berkeley Center for Information Technology in the Interest of Society (CITRIS) "People and Robots" Initiative (robotics.citris-uc.org) The authors were supported in part by the U.S. National Science Foundation under NRI Award IIS-1227536: Multilateral Manipulation by Human-Robot Collaborative Systems, by Google, Cisco, by a major equipment grant from Intuitive Surgical, and by generous donations from Andy Chou and Susan and Deepak Lim.

REFERENCES

[1] E. Ayvali, R. A. Srivatsan, L. Wang, R. Roy, N. Simaan, and H. Choset, "Using bayesian optimization to guide probing of a flexible environment for simultaneous registration and stiffness mapping," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2016.

[2] M. Bjorkman, Y. Bekiroglu, V. Hogman, and D. Kragic, "Enhancing visual perception of shape through tactile glances," in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, 2013.

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

[4] S. A. Darzi and Y. Munz, "The Impact of Minimally Invasive Surgical Techniques," in *Annu Rev Med.*, vol. 55, 2004, pp. 223–237.

[5] B. Demi, T. Ortmaier, and U. Seibold, "The touch and feel in minimally invasive surgery," in *IEEE International Workshop on Haptic Audio Visual Environments and their Applications*, 2005.

[6] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian Process Implicit Surfaces for shape estimation and grasping," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2011.

[7] ——, "Uncertainty aware grasping and tactile exploration," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2013.

[8] ——, "Re-grasping for Gaussian Process Implicit Surface," in *Workshop on Autonomous Grasping at International Conference on Robotics and Automation (ICRA)*, 2014.

[9] P. S. Girão, P. M. P. Ramos, O. Postolache, and J. M. D. Pereira, "Tactile sensors for robotic applications," *Measurement*, 2013.

[10] 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.

[11] A. Gotovos, N. Casati, G. Hitz, and A. Krause, "Active learning for level set estimation," in *IJCAI*, 2013.

[12] P. Hebert, T. Howard, N. Hudson, J. Ma, and J. W. Burdick, "The next best touch for model-based localization," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, 2013.

[13] G. Hitz, A. Gotovos, F. Pomerleau, M.-E. Garneau, C. Pradalier, A. Krause, and R. Siegwart, "Fully autonomous focused exploration for robotic environmental monitoring," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, May 2014.

[14] G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, "Active planning for underwater inspection and the benefit of adaptivity," *International Journal of Robotics Research*, 2012.

[15] P. Kazanzides, Z. Chen, A. Deguet, G. Fischer, R. Taylor, and S. DiMaio, "An Open-Source Research Kit for the da Vinci Surgical System," in *IEEE Int. Conf. Robotics and Automation (ICRA)*, 2014.

[16] J. Konstantinova, A. Jiang, K. Althoefer, P. Dasgupta, and T. Nanayakkara, "Implementation of tactile sensing for palpation in robot-assisted minimally invasive surgery: A review," *Sensors Journal, IEEE*, vol. 14, no. 8, pp. 2490–2501, Aug 2014.

[17] J. Mahler, S. Patil, B. Kehoe, J. van den Berg, M. Ciocarlie, P. Abbeel, and K. Goldberg, "Gp-gpis-opt: Grasp planning with shape uncertainty using gaussian process implicit surfaces and sequential convex programming," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2015.

[18] S. McKinley, A. Garg, S. Sen, D. V. Gealy, J. P. McKinley, Y. Jen, M. Luo, D. Boyd, and K. Goldberg, "An Interchangeable Surgical Instrument System with Application to Supervised Automation of Multilateral Tumor Resection." in *IEEE Conf. on Automation Science and Engineering (CASE)*, 2016.

[19] S. McKinley, A. Garg, S. Sen, R. Kapadia, A. Murali, K. Nichols, S. Lim, S. Patil, P. Abbeel, A. M. Okamura, and K. Goldberg, "A disposable haptic palpation probe for locating subcutaneous blood vessels in robot-assisted minimally invasive surgery," in *IEEE Conf. on Automation Science and Engineering (CASE)*, 2015.

[20] A. Murali, S. Sen, B. Kehoe, A. Garg, S. McFarland, S. Patil, W. Boyd, S. Lim, P. Abbeel, and K. Goldberg, "Learning by Observation for Surgical Subtasks: Multilateral Cutting of 3D Viscoelastic and 2D Orthotropic Tissue Phantoms," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2015.

[21] K. Nichols, A. M. Okamura, *et al.*, "Methods to segment hard inclusions in soft tissue during autonomous robotic palpation," *IEEE Trans. on Robotics*, vol. 31, no. 2, pp. 344–354, 2015.

[22] C. E. Rasmussen, "Gaussian processes for machine learning." MIT Press, 2006.

[23] S. Sen, A. Garg, D. V. Gealy, S. McKinley, Y. Jen, and K. Goldberg, "Automating Multiple-Throw Multilateral Surgical Suturing with a Mechanical Needle Guide and Sequential Convex Optimization," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2016.

[24] 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 Int. Conf. on Robotics and Automation (ICRA)*, 2016.

[25] The GPy authors, "GPy: A gaussian process framework in python," http://github.com/SheffieldML/GPy, 2012–2015.

[26] M. I. Tiwana, S. J. Redmond, and N. H. Lovell, "A review of tactile sensing technologies with applications in biomedical engineering," *Sensors and Actuators A: Physical*, vol. 179, pp. 17–31, 2012.

[27] R. Veldkamp, E. Kuhry, W. Hop, J. Jeekel, G. Kazemier, H. J. Bonjer, E. Haglind, L. Pahlman, M. A. Cuesta, S. Msika, *et al.*, "Laparoscopic surgery versus open surgery for colon cancer: short-term outcomes of a randomised trial," *Lancet Oncol*, vol. 6, no. 7, pp. 477–484, 2005.

[28] S. B. Williams, M.-H. Chen, A. V. D'Amico, A. C. Weinberg, R. Kacker, M. S. Hirsch, J. P. Richie, and J. C. Hu, "Radical retropubic prostatectomy and robotic-assisted laparoscopic prostatectomy: likelihood of positive surgical margin (s)," *Urology*, 2010.

[29] W. Yu, Y. Li, Y. Zheng, N. Lim, M. Lu, and J. Fan, "Softness measurements for open-cell foam materials and human soft tissue," *Measurement Science and Technology*, vol. 17, no. 7, p. 1785, 2006.

[30] Z. Zhang, A. Munawar, and G. Fischer, "Implementation of a Motion Planning Framework for the da Vinci Surgical System Research Kit (dVRK)," in *Hamlyn Symposium on Medical Robotics (HSMR)*, 2014.