

Tumor Localization using Automated Palpation with Gaussian Process Adaptive Sampling

Animesh Garg^{1,2}, Siddarth Sen², Rishi Kapadia², Yiming Jen²,
Stephen McKinley³, Lauren Miller², 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 haptic palpation has the potential to precisely estimate the geometry of embedded tumors during robot-assisted minimally invasive surgery. We formulate the tumor boundary localization problem in terms of Bayesian optimization along implicit curves defined by a Gaussian Process representation of estimated tissue stiffness. We formulate three palpation algorithms in this context: (1) Expected Variance Reduction (EVR), which emphasizes exploration by minimizing variance, (2) Upper Confidence Bound (UCB), which balances exploration with exploitation using the estimated mean, and Implicit Level Set UCB (ILS-UCB), a variant of UCB that prioritizes sampling near a level set. We compare these algorithms in controlled simulation experiments, varying levels of measurement noise and bias. We find that ILS-UCB significantly outperforms the other algorithms in terms of symmetric difference between tumor boundary estimate and ground truth, reducing error by up to $10\times$. In initial physical experiments with a haptic probe and the Intuitive dVRK surgical robot using multiple sample points and a variant of travelling salesman to plan paths between them, the ILS-UCB algorithm is $4\times$ faster than a raster scan path. Details and supplementary material available at: berkeleyautomation.github.io/gpas

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 allow identification and localization of subsurface structures or tumors based on changes in tissue stiffness relative to the surrounding substrate [27].

Robotic Surgical Assistants (RSAs), such as Intuitive Surgical’s da Vinci system have been shown to facilitate precise minimally invasive surgery [3, 26] 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) [15], 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.

Despite the increasing use of RMIS in cancer treatment [20], the lack of haptic perception during RMIS, as compared to open surgery, has been shown to increase the risk of accidental tissue damage [4] and the likelihood of




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

leaving behind cancerous cells during tumor resection [28]. Moreover, RMIS procedures are primarily controlled by surgeons in a local teleoperation mode (master-slave with negligible time delays). Introducing autonomy of surgical sub-tasks such as tumor localization with palpation has the potential to assist surgeons, reduce fatigue & cognitive load, and facilitate supervised autonomy.

To plan and automate tumor removal, a precise estimate of the tumor boundary is necessary. Cancerous tumors are however often visually indistinguishable and embedded in healthy tissue. Uncertainty in the location of tumor (or cyst) boundary during surgery can result in imprecise incision. A negative margin could result in an excessively conservative estimate (cutting out healthy tissue), while a positive margin could result in breaking the tumor membrane and spreading cancerous tissue.

In this paper, we propose an autonomous tumor local-

¹IEOR, ²EECS, ³ME, University of California, Berkeley, CA USA; {animesh.garg, siddarthsen, rishikapadia, yjen, mckinley, laurenm, goldberg}@berkeley.edu

ization algorithm for palpation in robot-assisted minimally invasive surgery. We formulate the tumor boundary localization problem as Bayesian optimization along implicit curves defined by a Gaussian Process representation of estimated tissue stiffness. Our approach performs adaptive sampling using a sampling criterion that balances exploration and exploitation along estimated level sets of the surface stiffness as opposed to creating a 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 one 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 our method converges to the tumor boundary faster and can accommodate higher levels in measurement noise and bias. In addition to verification 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 [17] on the da Vinci Research Kit (dVRK) [30]. Initial results on dVRK suggest that our approach can localize the tumor boundary up to 4x faster than uniform raster search.

II. BACKGROUND AND RELATED WORK

Palpation sensors are a subclass of haptic and force sensors that mimic the biological sense of cutaneous touch. In RMIS, palpation sensors can estimate relative tissue stiffness and allow the surgeon to adjust force control input for safer tissue manipulation. Haptic feedback can be obtained by using a number of transduction principles [15, 21]. See Girão et al. [8] and Tiwana et al. [25] for surveys of existing haptic and force feedback devices in the context of robotic and biomedical applications respectively. Konstantinova et al. [15] provide a survey of a number of RMIS haptic feedback devices.

Haptic Tumor Localization. A number of recent works have focused on leveraging haptic feedback to automate mapping subsurface stiffness or material variation in soft tissue for RMIS procedures. Nichols et al. [19] present a method for haptic localization of subcutaneous tissue boundaries. Their method uses elastography data to train a classifier for stiffness discrimination between tumors and surrounding tissue, and explores the search space using discrete measurements along a grid, performing local refinement around points that are classified as boundaries of a tumor. In Goldman et al. [9], a method for automated palpation is presented that recursively increases measurement resolution in areas where the measurement passes a certain threshold. Ayvali et al. [1], present an algorithm to automate palpation for stiffness mapping for the purpose of registration of surface geometry to pre-operative data.

Active Search and Mapping. Our approach draws on prior work from adaptively estimating shapes and curves using

noisy measurements. 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, 6, 11]. Dragiev et al. [6] 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. The method by Bjorkman et al. [2] initializes an object shape estimate using stereo vision, and calculates a Gaussian process implicit surface 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. In Hollinger et al. [13], coverage-based inspection paths are planned based on estimated uncertainty over a Gaussian Process Implicit surface (GPIS) model of a ship's hull.

The approach presented in this paper is most closely related to recent work in active level-set estimation using mobile sensors for environmental modelling [10, 12]. In these applications, the idea is that rather than trying to achieve low estimate error everywhere, one should instead look for regions that represent boundaries of level sets, or where a scalar function crosses a specific threshold. Both [10, 12] use a sampling criterion based on the ambiguity of the function value at a particular point being above or below the selected threshold to drive selection of subsequent sampling locations. In [12], a receding horizon path planner is used to plan efficient paths for sensing to reduce uncertainty specifically around a threshold value for plankton level modeling using aquatic robots. In Gotovos et al. [10], a traveling salesman algorithm is used to plan paths that sample a set of new measurement sites chosen using the same information measure.


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

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,



(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 prioritizes sampling near level sets between the max and minimum values, and uncertainty.

C_{GT} . We evaluate the proposed method and compare to other sampling approaches by comparing the symmetric difference between ground truth (C_{GT}) and the algorithm estimate (C). Note that we penalize both under and overestimation equally.

IV. GAUSSIAN PROCESS ADAPTIVE SAMPLING

Algorithm overview. The stiffness map estimate $S(x)$, represented as a Gaussian process, is initialized with measurements collected at randomly selected locations. In the results shown in Section VI and VII, we use 5 initial measurements. 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 dimensionality [22]. Formally, a Gaussian process generates data located throughout a domain such that any finite subset of the range follows a multivariate Gaussian distribution. Gaussian processes (GP) are often used for estimating and modeling continuous spatial data. GP models provide a smooth estimate everywhere, even given sparse sets of training data, allow multi-modal sensor fusion, and provide statistical representation of the estimate useful for active estimate refinement using Bayesian optimization methods. We use a Gaussian process to represent the stiffness $S(x)$ and associated uncertainty from the observed palpation measurements.

A. Overview: Gaussian Process Regression

In general, the input data for a Gaussian process is a set of training data D with observations Y taken at points X , or $D = \{X, Y\} = \{(x_1, y_1), \dots, (x_n, y_n)\}$ for a set of n training samples. The Gaussian Process model treats a function f as 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 $\sigma_{x^+}^2$, i.e.

$$p(f(x^+) | x^+, X, Y) = \mathcal{N}(f(x^+); \mu_{x^+}, \sigma_{x^+}^2),$$

where

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

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

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 sensor. The covariance function (or kernel) K determines the correlation between input locations x_i . We use the squared exponential kernel [22]. And σ_n^2 is the variance corresponding to the additive noise term $\delta(\cdot)$.

The stiffness $S(x)$ is then used to define the tumor boundary contour (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 [7, 5, 16]. We define the tumor boundary (C) as an implicit curve based on a Gaussian process representing the stiffness $S(x)$. We define the tumor boundary as the α -level set for the stiffness $S(x)$, such that:

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

Since $S(x)$ is an estimate with associated uncertainty, we define the implicit curve using the Gaussian process 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

Below, we introduce three palpation algorithms that use different sampling criterion for generating sequential measurements, which we evaluate 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 estimate, the mean $\mu(x)$ and variance $\sigma(x)$ of the Gaussian process.¹

For all three palpation algorithms, we select sampling locations that maximize the sampling criterion over the

¹Sampling criterion are often referred to as *acquisition functions* in Gaussian process optimization.

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, [10], we select a set of samples during each iteration of the algorithm using the estimated mean and the known variance from the past iteration and solve a traveling salesman problem approximately to plan paths between the set of new measurement sites chosen using the acquisition function \mathcal{A} .

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.

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

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

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

Prioritizing high stiffness areas guides sampling toward regions that are likely to be tumor vs. surrounding tissue, and prioritizing high variance regions prioritizing 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 trades off between exploration and exploitation. Rather than seeking to maximize the mean, however, 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:

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

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.

An implicit level set representation means that we prioritize searching the expected tumor boundary, conditional on estimate uncertainty, and do not seek to precisely learn a stiffness map of the entire workspace. Our hypothesis is that 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.

VI. SIMULATION EXPERIMENTS

We compare the three palpation algorithms described in Section V in simulation for boundary 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 $2.5 \times 5 \text{ cm}$ 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 exploration and exploitation well for this application. We use Python package GPy [24] for implementing Gaussian process regression.

We evaluate performance of all three palpation 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*, as presented in our prior work [17],

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

where k represents the slope, Y_{min} and Y_{max} the maximum and minimum measurements, and $(x - C)$ is the distance between points x and the closest point on the tumor boundary. Additive measurement noise is drawn from $\delta \sim \mathcal{N}(0, \sigma)$, and measurement bias $\beta x_1 (Y_{max} - Y_{min})$ that increases linearly 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 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) levels are 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.

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 different tumor models, a circular disk shaped tumor (#1) and a horse-shoe shaped tumor (#2). The error is reported as a percentage of the search space area after 10 iterations (with a batch size of 10, i.e. 100 measurements) and is averaged over five trials in each case. We note that 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

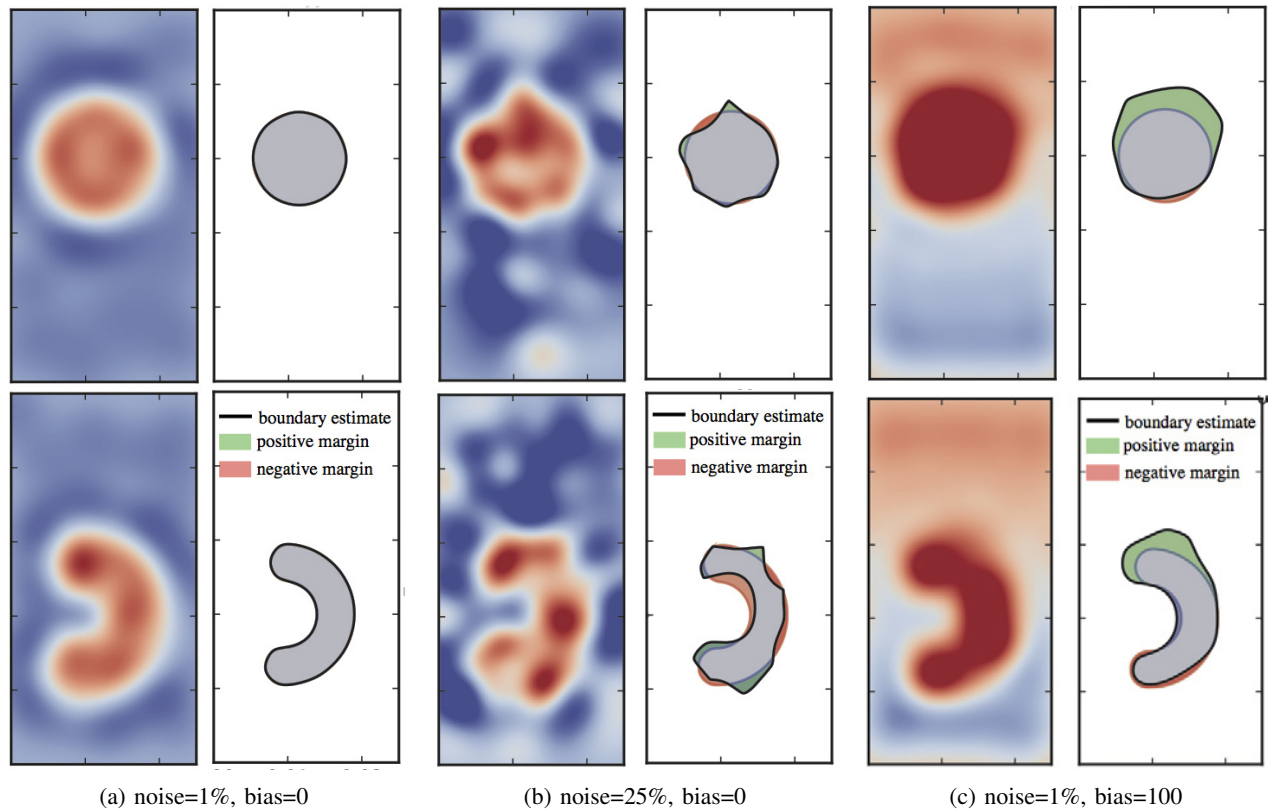


Fig. 3: Estimated stiffness maps and boundary estimates for simulated experiments after 10 iterations (of batch size 10, i.e. 100 measurements) using the ILS-UCB algorithm, for different noise levels and measurement bias. Regions in blue denote surrounding tissue with lower stiffness and regions in red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

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.

VII. PHYSICAL EXPERIMENTS

In this section we present a preliminary demonstration of tumor localization on a physical phantom similar to the simulated phantoms, using the dVRK robot.

dVRK: Hardware and Software. We use the Intuitive Surgical da Vinci Research Kit (dVRK) surgical robot assistant with the 8 mm Needle Drivers, as in [18, 23]. We interface with the dVRK using open-source electronics and software developed by WPI and Johns Hopkins University [14]. 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 placed in a 100 mm long, 50 mm wide, 20 mm deep mold and filled with a softer silicone rubber *Ecoflex 00-20 (Smooth-On)* as subcutaneous tissue. An opaque dermal phantom of 1 mm thickness was created using a stiffer (shore hardness 2A) *DragonSkin 10 Medium* silicone rubber

(*Smooth-On*). The dermal phantom was overlaid on the subcutaneous phantom to create the final tissue phantom setup as shown in Figure 1.

Stiffness measurement with Palpation Probe. The physical experiments in this study use a customized sensor, *PALP*, as presented in our prior work [17]. *PALP* is a low-cost and disposable palpation device designed to fit on a dVRK classic tool-tip and pass through an RMIS trocar port. The *PALP* probe uses a displacement-based contact sensing mechanism. A spherical indenter of 4.5 mm diameter allows point and sliding palpation. The end-effector of the probe has a known spring constant and uses Hall Effect sensing 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 (measured by an incremental Hall Effect magnetic encoder) can be linearly related to a tissue reaction force (F) using Hooke's Law ($F = k\delta_p$), where k is the spring constant ($k = 0.08$ 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 as shown empirically in [17].

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 on the horseshoe-shaped tumor. At each iteration, the probe scans

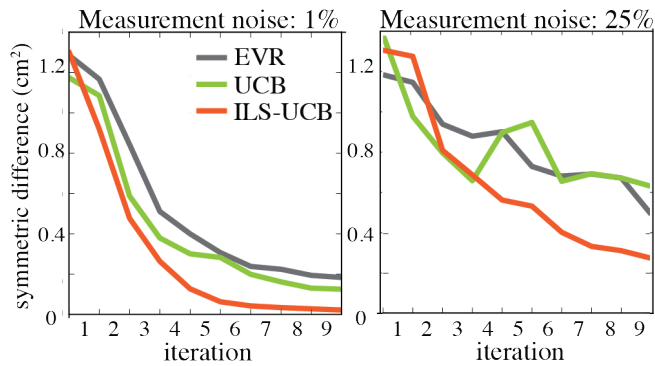


Fig. 4: Simulation Experiments: Convergence of the symmetric difference error (as % of search space area) for palpation algorithms as a function of iteration for two different levels of measurement noise, using a horseshoe shaped tumor in simulation. As the noise increases, there is randomized exploration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

the surface between the selected points, continuously collecting measurements. In the interest of reducing computation time, scanning trajectories on the robot were computed by interpolating between the points selected based on the sampling criterion. Selected points were ordered by approximately solving a Travelling salesman problem at each iteration. Measurements collected between the selected points were incorporated in the GP update (this differs from the simulation experiments, which did not incorporate measurements in between selected points). Future work will explore direct trajectory optimization with respect to the sampling criterion.

The robot moves at a speed of 5mm/s and measurements are collected at a rate of 1 sample/mm. 10 iterations on the robot, including computation and scan time, was completed in 3:07 minutes (averaged across 5 runs). As a comparison, performing a continuous raster scan using the robot at a rate of 5mm/s, with 2.0mm between rows (27800 samples total), takes 12 minutes.

Figure 5 shows a visual comparison of the tumor boundary estimation on dVRK. Figure 5(a) shows a top-down view of the custom molded rubberized tumor (blue) embedded in a silicone matrix (white). Figure 5(b) shows the estimate of the mean stiffness obtained after 10 iterations (in batches of 10) of our localization algorithm with ILS-UCB as the palpation algorithm. We use 50% level set of the mean estimate as the tumor boundary ($\alpha = 0.5$) which is indicated as a black line on the image.

VIII. DISCUSSION AND FUTURE WORK

We study the problem of tumor boundary estimation for robot-assisted minimally invasive surgery. We propose an algorithm for autonomous tumor localization using palpation based on Gaussian process adaptive sampling along implicit curves defined by stiffness measurements. The simulation results show that our algorithm can achieve up to a $10x$ reduction in error in boundary estimate. We have also demonstrated our algorithm on a physical system implemented with a displacement based palpation probe on a da Vinci Research kit (dVRK) and we observe that algorithmic search along

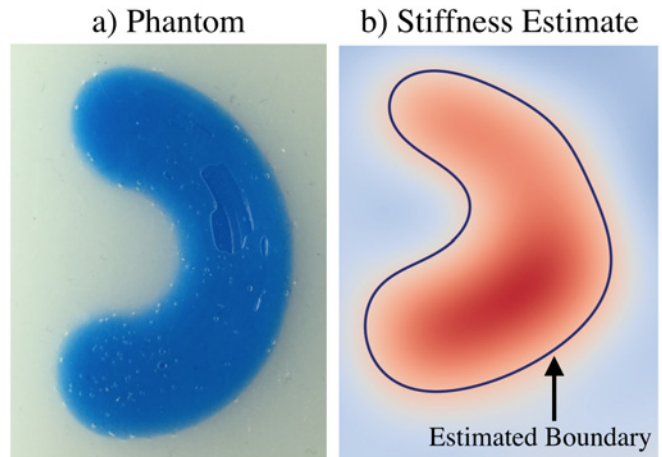


Fig. 5: Tumor boundary estimation on physical system using ILS-UCB with the PALP probe mounted on the dVRK. (a) The picture shows a custom molded rubberized tumor (blue) embedded in a silicone matrix (white). (b) The image shows our algorithm’s estimate of the tumor with the estimate of the boundary mean shown in black contour line obtained after 10 iterations of our localization algorithm with a total of 100 measurements/

implicit curves is on average $4x$ faster than uniform raster scanning.

A limitation of the proposed adaptive sampling method for tumor localization is that an accurate estimate of the surface geometry is required for automating haptic sensing. While using pressure or force-based haptic measurement devices such as *PALP*, controlling the applied force is critical in interpreting measurements and requires surface map which is difficult to reliably obtain. Maintaining a constant indentation and performing measurements along surface normals is necessary for a correct estimate of sub-surface stiffness, and by extension, of the tumor boundary. In the experiments presented in Section VI, we assume the palpation surface is planar, and known a-priori. But as observed in Figure 3(c), if the surface is slightly tilted in an unknown manner, the resulting estimates are subject to high error levels. Future work will focus on extending methods in this paper to boundary estimation on unknown, non-planar phantoms, combining surface estimation with stiffness estimation.

Details and supplementary material available at: berkeleyautomation.github.io/gpas

Acknowledgements: This research was performed with UC Berkeley’s Automation Sciences Lab 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, and 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,” *ArXiv e-prints*, Sept. 2015.

- [2] M. Bjorkman, Y. Bekiroglu, V. Hogman, and D. Kragic, "Enhancing visual perception of shape through tactile glances," in *IEEE Int'l Conf. on Intelligent Robots and Systems (IROS)*, 2013.
- [3] S. A. Darzi and Y. Munz, "The Impact of Minimally Invasive Surgical Techniques," in *Annu Rev Med.*, vol. 55, 2004, pp. 223–237.
- [4] B. Demi, T. Ortmaier, and U. Seibold, "The touch and feel in minimally invasive surgery," in *Haptic Audio Visual Environments and their Applications, 2005. IEEE International Workshop on, 2005*.
- [5] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian process implicit surfaces for shape estimation and grasping," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on.* IEEE, 2011, pp. 2845–2850.
- [6] —, "Uncertainty aware grasping and tactile exploration," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on.* IEEE, 2013, pp. 113–119.
- [7] —, "Re-grasping for gaussian process implicit surface," in *Workshop on Autonomous Grasping at International Conference on Robotics and Automation (ICRA)*, 2014.
- [8] P. S. Girão, P. M. P. Ramos, O. Postolache, and J. M. D. Pereira, "Tactile sensors for robotic applications," *Measurement*, 2013.
- [9] 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.
- [10] A. Gotovos, N. Casati, G. Hitz, and A. Krause, "Active learning for level set estimation," in *IJCAI*, 2013.
- [11] 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.* IEEE, 2013.
- [12] 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 *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, May 2014.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] 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 *Robotics and Automation (ICRA), 2015 IEEE International Conference on.* IEEE, 2015, pp. 4919–4926.
- [17] S. McKinley, A. Garg, S. Sen, R. Kapadia, A. Murali, K. Nichols, S. Lim, S. Patil, P. Abbeel, A. M. Okamura, *et al.*, "A disposable haptic palpation probe for locating subcutaneous blood vessels in robot-assisted minimally invasive surgery," *IEEE CASE*, 2015.
- [18] 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 *Int. Conf. on Robotics and Automation (ICRA)*, 2015.
- [19] K. Nichols, A. M. Okamura, *et al.*, "Methods to segment hard inclusions in soft tissue during autonomous robotic palpation," *Robotics, IEEE Transactions on*, vol. 31, no. 2, pp. 344–354, 2015.
- [20] K. Ohuchida and M. Hashizume, "Robotic surgery for cancer," *The Cancer Journal*, vol. 19, no. 2, pp. 130–132, 2013.
- [21] 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," *Sensors Journal, IEEE*, 2008.
- [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," 2016.
- [24] The GPy authors, "GPy: A gaussian process framework in python," <http://github.com/SheffieldML/GPy>, 2012–2015.
- [25] 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.
- [26] 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.
- [27] S. K. Venkatesh, M. Yin, J. F. Glockner, N. Takahashi, P. A. Araoz, J. A. Talwalkar, and R. L. Ehman, "Magnetic resonance elastography of liver tumors-preliminary results," *AJR. American Journal of Roentgenology*, vol. 190, no. 6, p. 1534, 2008.
- [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.