Utility-Guided Palpation for Locating Tissue Abnormalities

Elif Ayvali[1], Alexander Ansari[1], Long Wang[2], Nabil Simaan[2] and Howie Choset[1]

***Abstract*—Palpation is a key diagnostic aid for physicians when looking for tissue abnormalities. This paper focuses on autonomous robotic palpation for locating regions of interest representing possible tumor locations or underlying anatomy (e.g. a hidden artery). Many approaches direct the robot to exhaustively palpate the entire organ. To reduce exploration time, we define a utility function to guide the search to palpate likely regions of interest and update this function as new palpation data is collected. The search approach, presented in this paper, incorporates prior information from preoperative images, which can provide an estimate for the location of suspicious sites, thereby reducing unnecessary manipulation and exploration of the organ. To generate search trajectories that encode a coverage goal and locate all regions of interest, two planners are adapted. The first planner is based on ergodic coverage and the second planner is based on Bayesian optimization algorithm (BOA). Both planners are evaluated via simulation and experimentally to elucidate their strengths and weaknesses. The results demonstrate a higher efficacy of the BOA planner in detecting all regions of interest while avoiding exhaustive palpation scan of the organ.**

***Index Terms*—Surgical Robotics: Planning, Reactive and Sensor-Based Planning, Probability and Statistical Methods**

I. Introduction

SURGEONS rely on palpation to understand the anatomy, to detect anatomical abnormalities (e.g. possible tumor sites) and to localize sensitive underlying anatomy such as arteries and nerve bundles. While easy to achieve in an open surgery, palpation capabilities are unfortunately lost during robotic minimally invasive surgery. To address this sensory deficiency, there has been a significant focus on developing tactile and force sensing techniques (e.g. [1]–[5]). In addition to works on sensory methods for palpation, there have been some works focused on autonomous probing and palpation. For example, some works direct the robot to exhaustively probe all the vertices of a predefined grid on the organ's surface along a predefined path, such as a raster scan that covers the entire organ (e.g. [6]–[9]). Some of these works [7], [8] use adaptive grid resolution to increase palpation resolution around boundaries of regions of interest marked by high stiffness gradients. Such an exploration strategy can take a significant amount of time. We have recently employed Bayesian optimization algorithm (BOA) [10] to guide exploration and to reduce the number of observations needed to locate a tumor or a hidden artery. However, this approach is also seeded with an initial coverage plan, i.e. a sparse grid uniformly distributed across the organ. Therefore, these prior works generate palpation paths that is distributed across the entire organ's surface and do not incorporate prior information which can reduce unnecessary manipulation and exploration of the organ.

This work investigates methods for enabling efficacious palpation for detecting possible tissue abnormalities. Unlike other works relying on fixed or variable resolution coverage of the organ, we seek to investigate methods that use anatomical priors and online palpation sensory data to guide the path of palpation to likely locations of regions of interests.

In our formulation, we define a utility function that measures the benefit of taking sensor measurements, i.e. force and position, at each point on the organ and is used by a path-planner to guide the palpation to regions where there is a likelihood of observing a tissue abnormality. A Gaussian process (GP) represents the utility function, and GP regression [11] is used to update it efficiently using new palpation data. We assume that regions that are expected to have high stiffness are expected to provide high utility in identifying tissue abnormalities. Thus, GP represents the expected stiffness distribution over the organ.

Our approach differs from prior work [2], [3], [6], [7], [10], [12], [13] in that the search starts from a user-specified region of interest determined from the preoperative images and progressively expands out. To achieve this, Section IV introduces a weighted utility function which combines the utility function represented by the GP with those provided by a user-specified prior. It also incorporates a distance measure between the current and future locations, which allows the robot to favor regions in its immediate vicinity. To adapt the palpation path in closed-loop, we introduce a mechanism to decay the user-specified prior, as the importance of the prior decreases over time, and use new measurements to update the utility function.

We adapt and investigate the performance of two palpation path planners, one based on ergodic coverage [14] and another based on BOA [15]. Although both planners adapted in this work are myopic, i.e. they plan each action sequentially, they were chosen because they encode coverage strategies that locate all regions of interest, and meet the computational demands of an online search process. We provide a comparison

Manuscript received: September, 10, 2016; Revised December, 07, 2016; Accepted January, 03, 2017.

This paper was recommended for publication by Editor Tamim Asfour upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by NRI Large IIS-1426655 and IIS-1327566.

[1]E. Ayvali, A.Ansari, H. Choset are with the Robotics Institute at Carnegie Mellon University, Pittsburgh,PA 15213, USA (eayvali@,raansari1@,choset@) cmu.edu

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

Digital Object Identifier (DOI): see top of this page.

and discuss the strengths and weaknesses of both planners in Section V. The ergodic coverage algorithm introduced by Mathew and Mezić [14] explicitly incorporates a coverage metric, called the metric for ergodicity, as an objective function to distribute the time spent searching regions in a domain with respect to a desired distribution. Relative to the work in [14], where a stationary distribution is defined to guide the coverage, we use a utility function to guide coverage and update it online using new measurements. The planner based on BOA balances the trade-off between exploration and exploitation in search. As exploration initially drives the search to unexplored points of the search space, exploitation narrows the search to points where the utility function has maxima. Relative to global search methods based on BOA [10], [15]–[17], our approach facilitates priors that can initially bias the search towards user-specified regions of interest and allows for gradual expansion of the search from this region.

II. Related Work

This section focuses on previous works on active sensing [18]–[20], which share similar goals and strategies to those presented in this work. Active sensing determines a strategy to support a sensing objective such as optimizing sensor performance, acquiring information for target localization and estimation tasks.

Planning for Information Acquisition: Many approaches to target localization use some measure of utility, e.g. expected information gain, to plan actions or trajectories of a robot [21]–[24]. This requires maintaining a measure of utility [21], [25] over all possible measurements, which can be computationally expensive [26], [27]. In this context, GPs are becoming increasingly popular to represent objective functions in planning and decision-making since a continuous model of the function can be formed with few samples [28]–[31].

Planning long-term optimal trajectories is computationally expensive and in most cases intractable under uncertainty [32], [33]. To keep computations at a manageable level, the community turned to suboptimal approaches to plan locally optimal feasible trajectories [21], [22] or the next best action sequentially [34]–[36]. For example, previous works often use gradient-based methods, where the robot follows the local gradient of expected utility [23], [24]. These methods seek immediate improvement (i.e., driving the agent to the nearest local maximum) rather than maximizing long-term reward. Therefore, they are highly sensitive to local maxima and can suffer if the local gradient information is inaccurate or information is distributed.

Exploration and Exploitation Trade-off: Achieving long-term information gain (reward) necessitates a balance between exploration, visiting unexplored states of the search space, and exploitation, visiting states that are expected to be informative based on the current belief [37]. Previous works commonly use heuristics to encode exploration and exploitation as two objectives through multiple utility functions [37] or switch between these two objectives based on a stopping criterion [38], [39]. In this context, BOA is a powerful sequential sampling strategy that drives the search to a global maximum of an objective function after few samples, i.e. a utility function that is expensive to evaluate [15], [16]. In the robotics community, BOA has been previously used to find the optimal parameters of a trajectory which maximizes some measure of utility [16], [17], and for sequentially selecting sensing locations to find the locations where the sensor measurements take extreme values [10], [40]. Previous works commonly neglect the cost of motion, and do not address how to incorporate prior information, when available, to bias the search to user-specified regions of the domain. Additionally, BOA is a global search method and there is no established means to gradually expand the search space.

Coverage Approaches: Some approaches to search direct the robot to exhaustively pass over all points in the search domain to locate targets of interests [41]–[43]. These, commonly referred to as coverage approaches [41], [44] can take significant amounts of time. Therefore, probabilistic approaches were developed to exploit prior information, when it exists, to direct the search process to take less time. Many of these methods incorporate a coverage metric, called ergodicity [14], as an objective to bias the search with respect to a distribution. Previous work [27] shows that such a coverage strategy performs better than driving the search to the most likely location, i.e. information surfing, in a target localization example.

Mathew and Mezić [14] used the metric for ergodicity as an objective function and derived control laws for robots with first-order and second-order dynamics to cover a domain with respect to a probability density function (pdf) defined over the search domain. To accommodate broader classes of robots with nonlinear dynamics, Miller and Murphey [45] incorporates the metric for ergodicity in a nonlinear trajectory optimization framework and plan *feed-forward* trajectories over a specified time horizon. By considering future time steps, this approach is less myopic than [14]. However, trajectory optimization is a computationally intensive, iterative process and so is generally impractical to implement in an online, model predictive fashion.

III. Methods

The palpation path is driven by a planner that uses a weighted-utility function. Section III-A introduces GP regression which allows efficient update of the utility function using new palpation data. Section III-B gives background on ergodic coverage, followed by Section III-C on BOA, which are used to plan palpation paths.

A. Gaussian Process Regression

Gaussian processes are sample-efficient tools to perform nonparametric regression [11]. Intuitively, one can view GP as a distribution over functions. Given a d-dimensional search domain $X \subset \mathbb{R}^{\mathrm{d}}$, the distribution of function values at a point $\boldsymbol{x} \in X$ is represented by a random variable, y, and has a Gaussian distribution, $N(\mu(\boldsymbol{x}), \sigma^2(\boldsymbol{x}))$. Given a set of n observations of the function $\bar{\boldsymbol{y}} = [y_1, y_2, \ldots, y_n]^T$ at $[\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_n]^T$, GP regression [11] can be used to make predictions on the distribution of function values at a new

point $\boldsymbol{x_*} \in X$ [11]

$$p(y_*|\bar{\boldsymbol{y}}) \sim N(\boldsymbol{K}_*\boldsymbol{K}^{-1}\bar{\boldsymbol{y}}, k_{**} - \boldsymbol{K}_*\boldsymbol{K}^{-1}\boldsymbol{K}_*^{\boldsymbol{T}}), \quad (1)$$

where $\boldsymbol{K}$ is the $n \times n$ covariance matrix whose elements $\boldsymbol{K}_{ij}(i, j \in [1, ...n])$ are calculated using a squared exponential covariance function $k(\boldsymbol{x_i}, \boldsymbol{x_j})$ [11]. Similarly, $\boldsymbol{K}_*$ is a $1 \times n$ vector defined as $\boldsymbol{K}_* = [k(\boldsymbol{x_*}, \boldsymbol{x_1}), \ldots, k(\boldsymbol{x_*}, \boldsymbol{x_n})]$, and finally $k_{**} = k(\boldsymbol{x_*}, \boldsymbol{x_*})$.

In our application, the search domain is $X \subset \mathbb{R}^2$. We choose this representation because every point on the organ's surface can be uniquely mapped onto X. The variable y represents the stiffness estimate obtained from robot's force and position measurements upon probing the organ at $\boldsymbol{x}$[1]. We discretize the search domain and perform GP regression to obtain the mean, $\mu(\boldsymbol{x})$, and variance, $\sigma^2(\boldsymbol{x})$, of the stiffness distribution across the organ.

B. Ergodic coverage

Ergodic theory is the statistical study of time-averaged behavior of dynamic systems [46]. Mathew and Mezić [14] introduced a metric to quantify the ergodicity of a robot's trajectory with respect to a given probability distribution function [14]. For completeness, we will first introduce the metric for ergodicity and then explain how it is used as an objective in planning.

The time-average statistics of a trajectory, $\boldsymbol{q} : (0, t] \to X$, quantifies the fraction of time spent at a point, $\boldsymbol{x} \in X$, where $X \subset \mathbb{R}^{\mathrm{d}}$ is a d-dimensional search domain. Mathew and Mezić [14] define the time-average statistics of a trajectory at a point $\boldsymbol{x}$ as

$$C(\boldsymbol{x}) = \frac{1}{t}\int_0^t \delta(\boldsymbol{x} - \boldsymbol{q}(\tau))d\tau, \quad (2)$$

where δ is the Dirac delta function. Let $\xi(\boldsymbol{x})$ be a pdf – also referred as a coverage distribution in [14] – defined over the search domain. The ergodicity of a robot's trajectory with respect to $\xi(\boldsymbol{x})$ is defined as [14]

$$\Phi(t) = \sum_{k=0}^{m} \lambda_k \left|c_k - \xi_k\right|^2, \quad (3)$$

where c_k and ξ_k are the Fourier coefficients of $C(\boldsymbol{x})$ and $\xi(\boldsymbol{x})$, respectively. The coefficient λ_k places higher weights on the lower frequency components and $m \in \mathbb{Z}^n$ is the number of basis functions. We observe that for $m > 10$, the value of (3) is negligible, therefore we use $m = 10$ basis functions.

The goal of ergodic coverage is to generate optimal controls $\boldsymbol{u}^*(t)$ for a robot, whose dynamics is described by a function $\dot{\boldsymbol{q}}(t) = g(\boldsymbol{q}(t), \boldsymbol{u}(t))$ such that

$$\begin{aligned} \boldsymbol{u}^*(t) &= \arg\min_{\boldsymbol{u}} \Phi(t), \\ \|\boldsymbol{u}(t)\| &\leq u_{max}, \end{aligned} \quad (4)$$

Mathew and Mezić [14] consider first-order, $\dot{\boldsymbol{q}}(t) = \boldsymbol{u}(t)$, and second-order systems, $\ddot{\boldsymbol{q}}(t) = \boldsymbol{u}(t)$, and derive control laws for minimizing Eq. (3) (see [14] for the derivation). As in Mathew and Mezić [14], we solve (4) by discretizing the exploration time and solving for the optimal control input that maximizes the rate of decrease of the ergodic metric at each time-step.

[1]In other applications, y can represent another problem-specific measure of utility. For example, when a parametric sensor model (observation likelihood) is available, information-theoretic functions are commonly used to quantify the utility of a measurement [19].

C. Bayesian Optimization Algorithm

BOA is a sequential sampling strategy for finding the global maxima of black-box functions [16]. A GP is used as a surrogate for the function to be optimized. BOA uses the posterior mean, $\mu(\boldsymbol{x})$, and variance, $\sigma^2(\boldsymbol{x})$, of the GP for all $\boldsymbol{x} \in X$, to sequentially select the next best sample as the point that maximizes an acquisition function such as expected improvement (EI) given by [16]

$$\xi_{EI}(\boldsymbol{x}) = \begin{cases} (\mu(\boldsymbol{x}) - y^+)\Phi(z) + \sigma(\boldsymbol{x})\phi(z) & \text{if } \sigma(\boldsymbol{x}) > 0 \\ 0 & \text{if } \sigma(\boldsymbol{x}) = 0 \end{cases} \quad (5)$$

where $z = \left(\frac{(\mu(\boldsymbol{x}) - y^+)}{\sigma(\boldsymbol{x})}\right)$, y^+ is the current maximum. $\phi(\cdot)$ and $\Phi(\cdot)$ are the probability density function and cumulative distribution function of the standard normal distribution, respectively. As exploratory sampling drives the uncertainty down, EI narrows the search to a global maximum.

IV. Closed-loop Search Process

The following sections describe the closed-loop search process illustrated in Fig. 1. Section IV-A introduces the weighted utility function. Section IV-B describes the process of planning palpation paths using ergodic coverage and BOA.

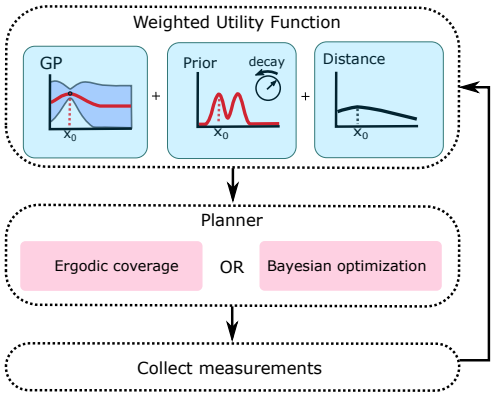

Figure 1. Closed-loop search process is guided by a weighted utility function. The shaded area in the GP represents the uncertainty in the model.

A. Weighted Utility Function

To guide palpation, this work uses a weighted utility function that has three components: 1) a GP component representing a utility function, i.e expected stiffness distribution, 2) a user-specified prior on regions of interests, e.g. possible tumor locations, and 3) a distance measure which allows the search process to favor regions nearby and acts as a surrogate motion model for the BOA-based planner, by guiding the

search to visit points in the immediate vicinity of the robot. The weighted utility function is defined as

$$\xi(\boldsymbol{x}) = \frac{1}{c}\left\{\frac{\xi_{GP}(\boldsymbol{x})}{\int_X \xi_{GP}(\boldsymbol{x})} + \frac{\xi_p(\boldsymbol{x})}{\int_X \xi_p(\boldsymbol{x})} + \epsilon_d\left[1 - \frac{\xi_d(\boldsymbol{x})}{\int_X \xi_d(\boldsymbol{x})}\right]\right\}, \tag{6}$$

where $\xi_{GP}(\boldsymbol{x})$ is a function which assigns each $\boldsymbol{x} \in X$ a scalar value according to the mean and variance of the GP. For BOA-based planner, $\xi_{GP}(\boldsymbol{x})$ is the acquisition function given by (5). For ergodic coverage, $\xi_{GP}(\boldsymbol{x})$ is defined as the upper confidence bound of the GP, $\xi_{GP}(\boldsymbol{x}) = \mu(\boldsymbol{x}) + \sigma(\boldsymbol{x})$, since the goal is to generate a palpation path that visits all possible regions that can correspond to the targets of interest, i.e. a tumor or a hidden artery. The distance measure, $\xi_d(\boldsymbol{x})$, is the Euclidean distance of $\boldsymbol{x}$ from the last location. The user-specified prior $\xi_p(\boldsymbol{x}) = \sum_k \eta^{n_k}\xi_p^k(\boldsymbol{x})$ is represented by a mixture of k Gaussians where the mean of each Gaussian is centered at the expected location of a target of interest, and $\eta < 1$ is the decay rate. Each term in the parenthesis in (6) is normalized and takes a value between $[0, 1]$. The scalar c is a normalizing constant such that the integral of $\xi(\boldsymbol{x})$ over the search domain is 1[2]. The weight of the distance measure, ϵ_d, is a positive number and $\epsilon_d < 1$. Therefore, the search is mostly dominated by $\xi_{GP}(\boldsymbol{x})$ and $\xi_p(\boldsymbol{x})$. The goal of $\xi_d(\boldsymbol{x})$ is to favor regions near the current location.

Initially, $\xi_{GP}(\boldsymbol{x})$ is a uniform distribution, and the distance measure is $\xi_d(\boldsymbol{x}) = 0$. In this case, the user-specified prior information, $\xi_p(\boldsymbol{x})$, guides the search. Hence, the search process starts in a region defined by $\xi_p(\boldsymbol{x})$. If we collect n_k measurements inside the region defined by the variance of the k^{th} Gaussian, the prior information $\xi_p^k(\boldsymbol{x})$ decreases by η^{n_k}. This process allows an effective way to trade off the bias of the user-specified prior and allows the GP to take over the search process.

B. Planning

This work demonstrates two computationally efficient planning methods. The BOA-based planner operates in a relatively straightforward manner. That is, it sequentially decides which point in the search domain to take measurements. Once the robot reaches the planned point, it probes the organ at that point, takes the measurements and updates the utility function. The process is closed-loop in that the GP is updated using the stiffness estimate at the probed point.

In ergodic coverage, we solve a model predictive control problem in the limit as the receding horizon goes to zero for a first-order system of the form, $\dot{\boldsymbol{q}}(t) = \boldsymbol{u}(t)$, subject to the constraint $\|\boldsymbol{u}(t)\|_2 \leq u_{max}$. The optimal feedback is computed at each time step for T seconds as in [14]. Over the T second interval, a set of points are probed along the trajectory at a fixed frequency, f_s, and then (6) is updated using the stiffness estimates at the probed points.[3] Figure 2(b) shows two trajectories generated with different u_{max}.

Ergodic coverage requires one additional parameter to encode the desired search behavior. As mentioned, ergodic coverage distributes the time spent in areas of the search space in proportion to the utility they provide. It can be difficult to coax the process to prioritize nearby regions and expand the search gradually because even after adding a distance measure, (6) results in non-zero utility measure in remote regions of the search-space. As shown in Fig 2(a), the algorithm will still visit regions with small but nonzero utility measure, although the time spent in those regions will be small. Therefore, we pass (6) through a threshold filter such that $\xi(\boldsymbol{x}) = 0$ if $\xi(\boldsymbol{x}) < I$. The parameter study in the next section includes guidelines on how to select I.

[2]Ergodic coverage requires $\xi(\boldsymbol{x})$ to be a pdf.

[3]We select f_s such that the spacing between consecutive probed points is smaller than the size of the smallest feature we want to detect.

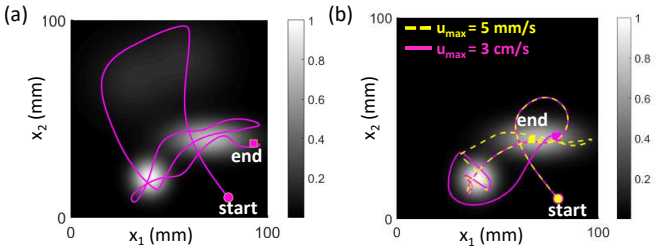

Figure 2. (a) Ergodic coverage of a given pdf. The color bar shows the measure of the distribution. (b) Ergodic coverage using [$u_{max} = 5$mm/s , $T = 60$s] (dashed line), and [$u_{max} = 3$cm/s , $T = 10$s] (solid line). We select T such that the generated trajectory results in over 99% ergodic coverage of (6) as measured by the metric for ergodicity. For higher velocity limits a shorter duration should be used.

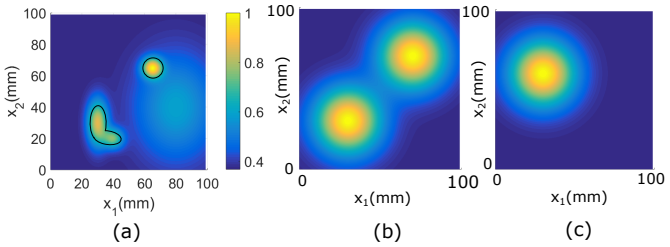

Figure 3. (a) A ground truth stiffness map generated using $100 \times 100 = 10000$ samples. (b) A user-specified good prior, (c) A user-specified inaccurate (poor) prior on the expected location of the tumors.

V. Results

In this section, we first present a parameter study to demonstrate the effects of the parameters used in (6) and then present experimental results using a phantom silicone tissue model with two inclusions representing tumors.

A. Parameter Study

Figure 3(a) shows a representative stiffness map of an organ. Since this is a simulated study, we assume that when the robot palpates (or probes) the organ at x and collects measurements, it obtains the stiffness estimate at x. In the following section, the stiffness will be estimated using actual force and position measurements. As performance measures we report the number of probed points on the organ, the total path length (PL) traveled, and the search area covered. Our goal is to minimize manipulation of the organ and search area. Therefore, in the results we use the number of probed points and the search area covered as the relevant measures

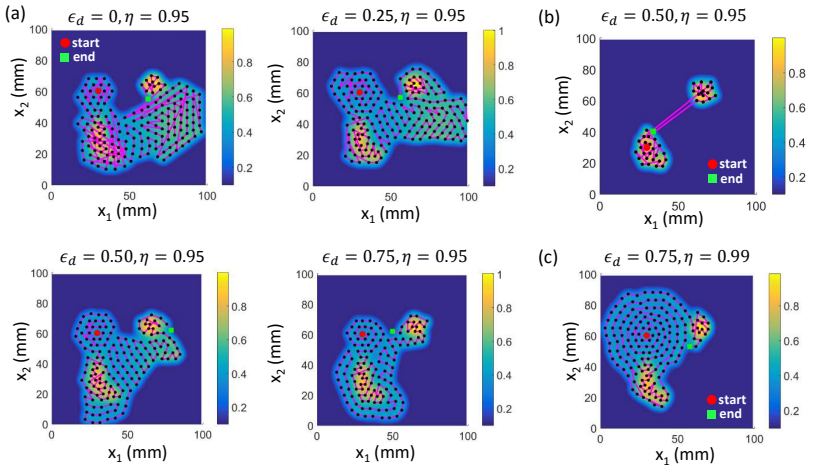

Figure 4. Search trajectories for BOA overlaid on the mean of the GP (i.e. expected stiffness (N/m)). (a) Set of figures showing the effect of ϵ_d when an inaccurate prior (Fig. 3(c)) is used; (b) The targets are quickly identified when the user-specified prior is good; (c) Using a high decay rate, $\eta = 0.99$, results in the search to be strongly biased by the user-specified prior. Black dots are the probed points on the organ.

Table I
(Y): WITH NOISE, (N): NO NOISE, (G): GOOD PRIOR, (P): POOR PRIOR, (PL): PATH LENGTH

Parameters ϵ_d	η	Noise (Y/N)	Prior (G/P)	Probed Points	PL (mm)	Area (%)
0.00	0.99	N	P	145	1919.90	26.40
0.25	**0.99**	**N**	**P**	**152**	**1094.70**	**27.53**
0.50	0.99	N	P	160	964.75	27.97
0.75	0.99	N	P	160	837.59	29.54
0.00	0.95	N	P	180	1613.20	32.41
0.25	0.95	N	P	180	1032.60	32.55
0.50	0.95	N	P	158	817.67	26.38
0.75	**0.95**	**N**	**P**	**135**	**649.27**	**23.55**
0.00	0.95	N	G	50	1163.6	13.50
0.25	0.95	N	G	47	466.74	8.82
0.50	**0.95**	**N**	**G**	**42**	**333.38**	**5.94**
0.75	0.95	N	G	50	359.03	7.32
0.50	0.95	Y	P	154 ± 36	788.27 ± 185.32	26.45 ± 5.94
0.50	**0.95**	**Y**	**G**	**49 ± 1**	**384.39 ± 15.68**	**7.62 ± 0.55**

to assess the performance. To have a fair comparison for all the parameter sets, we assume that the actual locations of the tumors are known, and terminate the search process when the uncertainty of the GP at the locations defined by the tumors in Fig. 3(a) fall below a threshold.

1) BOA: Figure 4 shows the search trajectories generated for different parameters. The colormap in Fig. 4 shows the predicted mean of the stiffness distribution. Table I shows a summary of results. The first two blocks in Table I are meant to provide a comparison of the effects of decay rate of the user-specified prior, η. The second and third blocks show the effects of having a good vs. a poor prior. The fourth block shows the effect of noise and the statistical results of 5 runs with 10% noise added to the stiffness estimate. The best parameter set for each block is shown in bold letters in Table I.

The results show that incorporating the distance measure, $\xi_d(\boldsymbol{x})$, into (6) significantly reduces the total path traveled (compare $\epsilon_d = 0$ vs. $\epsilon_d > 0$ in Table I). Also note that in Fig. 4(a) as ϵ_d increases, the generated trajectory becomes smoother. When η is closer to 1, the algorithm takes longer to forget the user-specified prior, and hence the initial user-specified prior biases the search. This results in exhaustive exploration of the region defined by the user-specified prior as shown in Fig. 4(c). When η is around 0.95, it provides a better trade-off, allowing the user-specified prior to guide search until enough samples have been collected for the GP to take over. We found $0.85 < \eta < 0.95$ to be a good range for the decay rate. The results for $\eta = 0.99$ were added only to explain the effect of the decay rate, and are not considered as they bias the search and completely follow the user-specified prior. Therefore, they are grayed out in Table I. Perhaps unsurprisingly, when the prior is good all performance measures are minimized for each parameter set compared to the case with a poor prior.

2) Ergodic Coverage: The parameter study includes values between $0 < \epsilon_d \leq 0.50$. For the threshold parameter, the study includes values between the range $0.50 \leq I \leq 0.70$. For $I \lesssim 0.50$, it is difficult to coax the search process to prioritize regions in the vicinity of the current location, and for $I \gtrsim 0.70$ search grows very slowly. We fix the velocity limit as $u_{max} = 5$mm/s and set $T = 60$s. We probe 25 points along the trajectory, update (6) based on the stiffness estimates at the probed points and compute a new trajectory.

Figure 5 shows the search trajectories for three different sets of parameters. Table II shows a summary of results. The first and second block show the results for the case with a poor prior and a good prior, respectively. The third block shows the effect of noise. The best parameter set for each block is shown in bold letters. The second block of Table II shows that the performance measures of ergodic coverage is less sensitive to the choice of parameters when initialized with a good prior. This is because the target features are in the region defined by the user-specified prior, and hence the search process is confined in that region (see Fig. 3(b) and Fig. 5 (c)). In the first block, note that increasing ϵ_d reduces the search area covered for all cases. Increasing the distance weight results in favoring regions nearby which in return results in more conservative expansion of the search space. For a given distance weight, ϵ_d, increasing I results in the trajectory being spatially more dense because there is a less area to cover (see Fig. 5 (a) vs. (b)).

3) Discussions: The parameter study shows that both search strategies start from a region defined by the user-specified prior and progressively expand the search space. We observe that BOA is more susceptible to noise, while ergodic coverage exhibits spatially more consistent expansion. The generated trajectories can vary depending on the noise level. Upon locating a region of interest, BOA segments the boundaries of the tumor quickly and covers less area, especially when the user-specified prior is good, which makes it advantageous for this application where accurate segmentation of boundaries and avoiding unnecessary exploration

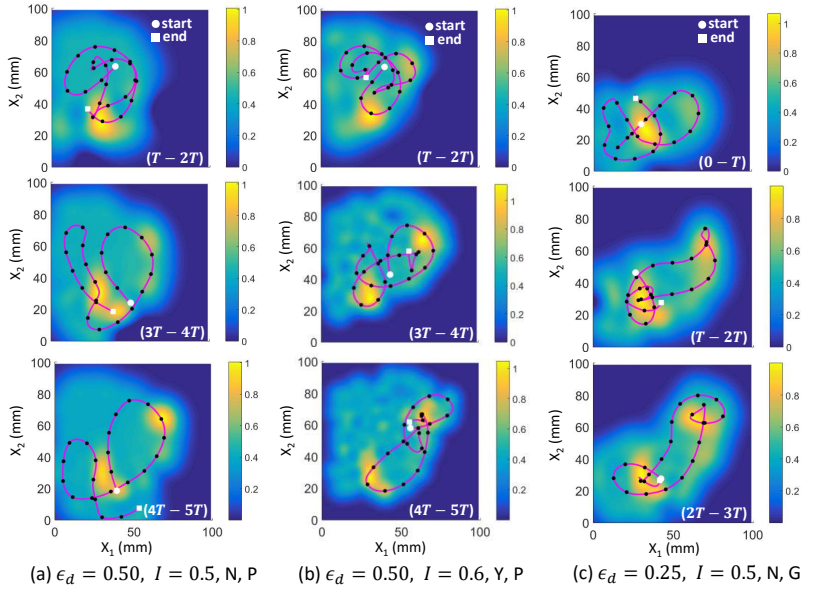

(a) $\epsilon_d = 0.50,\ I = 0.5$, N, P (b) $\epsilon_d = 0.50,\ I = 0.6$, Y, P (c) $\epsilon_d = 0.25,\ I = 0.5$, N, G

Figure 5. Ergodic search trajectories overlaid on the mean of the GP (i.e. estimated stiffness) for different sets of parameters. (c) shows the measurement sequence $(0-T, T-2T, 2T-3T)$ for the first row in the second block with parameters $\epsilon_d = 0.25, I = 0.5$. Due to space limit we only show some of the time intervals. Solid purple lines are the search trajectories and black dots are the probed points on the organ.

Table II
(Y): WITH NOISE, (N): NO NOISE, (G):GOOD PRIOR, (P): POOR PRIOR, (PL): PATH LENGTH

Parameters ϵ_d	I	Noise (Y/N)	Prior (G/P)	Probed Points	PL (mm)	Area (%)
0.25	0.5	N	P	100	1200	42.15
0.50	0.5	N	P	125	1500	38.90
0.25	0.6	N	P	100	1200	34.16
0.50	**0.6**	**N**	**P**	**100**	**1200**	**31.08**
0.25	0.7	N	P	125	1500	41.90
0.50	0.7	N	P	125	1500	33.73
0.25	0.5	N	G	50	600	23.56
0.50	**0.5**	**N**	**G**	**50**	**600**	**23.06**
0.25	0.6	N	G	50	600	30.73
0.50	0.6	N	G	50	600	30.97
0.25	0.7	N	G	50	600	27.27
0.50	0.7	N	G	50	600	27.30
0.50	0.6	Y	P	142 ±14	1700 ±173.21	32.41 ± 5.62
0.50	**0.6**	**Y**	**G**	**67 ± 14**	**800 ± 173.21**	**27.97 ± 3.49**

Table III
COMPARISON OF THE ALGORITHMS

Method	Probed Points	PL (mm)	Area (%)
BOA	**39**	**230.18**	**17.65**
Ergodic Coverage	45	450.00	29.04

are important. BOA is less likely to miss small "hotspots" as the search expands slowly and consecutive probed points are close. On the other hand, ergodic coverage captures the robot constraints thereby offering both planning and control for guiding the robot. Ergodic coverage is less sensitive to noise as a set of observations are used to update the GP. It has more parameters to be tuned, although we found a wide range of parameters to perform well. In simulated tests with different stiffness maps, these observations were consistent. In BOA, it takes less than 0.05s to update (6) on a 100×100 grid and select the next best location to take measurements[4]. In comparison, it takes 0.25s on average to generate an ergodic trajectory and update (6).

B. Experimental Results

The BOA and ergodic coverage algorithms were tested on experimental data obtained using a Cartesian robot as shown in Fig. 6(a). The robot is equipped with an ATI Nano43 6-axis force sensor mounted on the robot end-effector, and it is capable of executing hybrid motion/force tasks running on Matlab xPC Target. Probing a point on the organ involves detecting surface contact and estimating the surface normal followed by probing along the computed surface normal up to 3mm depth. We employ a linear stiffness model and use the slope of the line that best fits the force-displacement profile similar to [10]. Therefore, every point on the organ's surface is associated with a scalar stiffness value. The silicon organ has two stiff inclusions that represent tumors. Fig. 6(b) shows the computer-aided design (CAD) model of the organ in the form of a triangular mesh obtained from CT scans.

In a clinical scenario, the approximate locations of the tumors are determined from MRI or CT scans. Before the surgical procedure, the surgical tool needs to be registered to the surgical site to be able to find correspondences between the preoperative scans and the visible anatomy. We assume 5mm positional error and 0.1rad $(5.73°)$ angular error in registration parameters to represent the inaccuracy in the registration estimate, which is within the reported range of registration errors [47].

We set $T = 30$s, probe 15 points along the trajectory, and then generate a new T second trajectory. Figure 6(e) shows the result for BOA after probing 50 points on the organ and Fig. 6(f) shows ergodic trajectory for the time intervals $(0-T, T-2T, 2T-3T, 3T-4T)$. Both algorithms started searching the regions defined by the registration-specified prior and then expanded out. Table III shows the performance measures. BOA covered less area searching for the target features, obtained more samples around the targets and quickly identified the boundaries. The experimental results are consistent with the parameter study.

VI. CONCLUSIONS

This paper presented algorithms for robot-assisted palpation, which we believe should be useful for future semi-automated robotic palpation systems. We derived a means to use anatomical priors and online palpation sensory data to guide the path of palpation to likely locations of tissue abnormalities, e.g. possible tumor sites. We introduced two methods to generate

[4]The code runs on MATLAB R2015a in Windows 10 on a laptop with i7 CPU and 8 GB RAM. The code was not optimized and the computation time can be substantially improved with a C++ implementation and by taking advantage of matrix sparsity in GP regression.

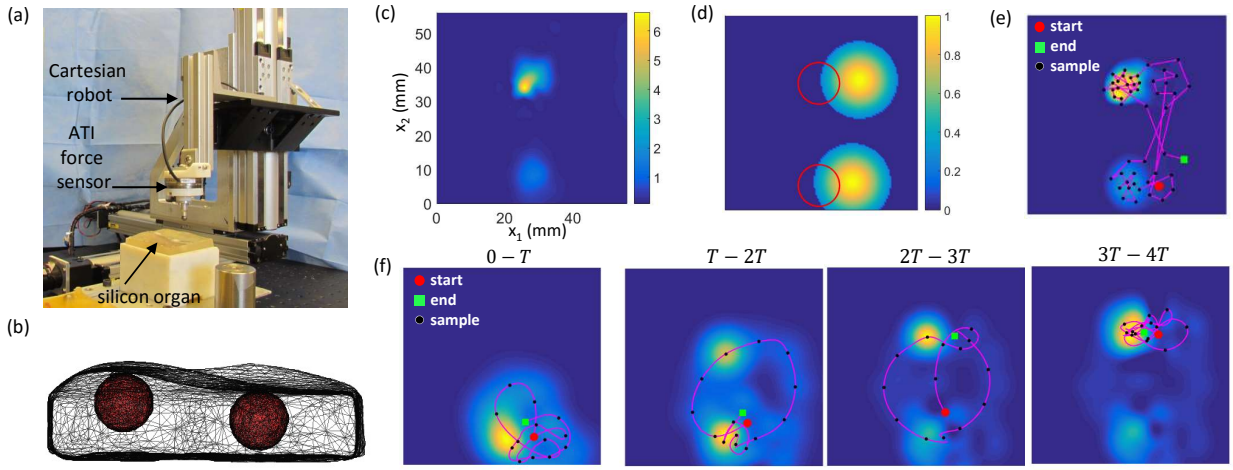

Figure 6. (a) Experimental setup, (b) CAD model of the organ and stiff inclusions, (c) Actual stiffness map, (d) Prior information obtained from registration and red circles showing the actual locations of the inclusions, (e) BOA, (f) Ergodic coverage.

palpation paths, one based on Bayesian optimization another on ergodic coverage, and compared them with real experimental data collected through physical experiments. Both planners encode coverage strategies that locate all regions of interest, and meet the computational demands of an online closed-loop search process.

The weights used in (6) are intuitive and we found them easy to specify in practice. The included parameter study demonstrates a wide range of values tend to produce good results and indicates trends which may be useful in fine-tuning performance. BOA-based planner is more susceptible to noise as it updates the trajectory after every observations. However, it is very effective in quickly identifying tumors when initialized with a good prior. In robotic surgery, there is reasonably a good prior. Therefore, BOA can be preferred as it requires less number of observations (probing) and less search area to be covered. On the other hand, ergodic coverage shows more consistent expansion of the search space even when the algorithm is seeded with inaccurate prior and subjected to noise. Additionally, ergodic coverage offers both planning and model-based control in one package.

There are several directions for future work. In our application, we used a simple experimental setup and an unconstrained environment to evaluate our methods to avoid additional sources of error. In our future work, we will demonstrate the proposed method using a continuum robot [48] and the da Vinci Research Kit [49]. We believe the framework introduced in this work can be used to address robots with uncertain motion models by propagating the uncertainty in robot's state into the utility estimation using GP regression methods which can incorporate input noise [50].

ACKNOWLEDGMENTS

The authors would like to acknowledge the assistance of George Mathew and Igor Mezić in implementation of the ergodic coverage algorithm.

REFERENCES

[1] P. S. Wellman, E. P. Dalton, D. Krag, K. A. Kern, and R. D. Howe, "Tactile imaging of breast masses: first clinical report," *Archives of surgery*, vol. 136, no. 2, pp. 204–208, 2001.

[2] 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.

[3] 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.

[4] I. B. Wanninayake, L. D. Seneviratne, and K. Althoefer, "Estimation of tissue stiffness using a prototype of air-float stiffness probe," in *Int. Conf. on Robotics and Automation (ICRA)*, 2014, pp. 1426–1431.

[5] K. Xu and N. Simaan, "An investigation of the intrinsic force sensing capabilities of continuum robots," *IEEE Transactions on Robotics*, vol. 24, no. 3, pp. 576–587, 2008.

[6] M. Beccani, C. Di Natali, L. J. Sliker, J. Schoen, M. E. Rentschler, and P. Valdastri, "Wireless tissue palpation for intraoperative detection of lumps in the soft tissue," *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 2, pp. 353–361, 2014.

[7] 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.

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

[9] 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)[Submitted]*, 2016.

[10] 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," *The Int. Conf. on Robotics and Automation*, 2016.

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

[12] H. Liu, D. P. Noonan, B. J. Challacombe, P. Dasgupta, L. D. Seneviratne, and K. Althoefer, "Rolling mechanical imaging for tissue abnormality localization during minimally invasive surgery," *IEEE Transactions on Biomedical Engineering*, vol. 57, no. 2, pp. 404–414, 2010.

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

[14] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4, pp. 432–442, 2011.

[15] 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.

[16] 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.

[17] R. Marchant and F. Ramos, "Bayesian optimisation for informative continuous path planning," in *Int. Conf. on Robotics and Automation (ICRA)*, 2014, pp. 6136–6143.

[18] C. Kreucher, K. Kastella, and A. O. Hero Iii, "Sensor management using an active sensing approach," *Signal Processing*, vol. 85, no. 3, pp. 607–624, 2005.

[19] C. Cai, *Information-driven sensor path planning and the treasure hunt problem*. ProQuest, 2008.

[20] L. Mihaylova, T. Lefebvre, H. Bruyninckx, K. Gadeyne, and J. D. Schutter, "Active sensing for robotics - a survey," in *in Proc. 5 th Intl Conf. On Numerical Methods and Applications*, 2002, pp. 316–324.

[21] F. Bourgault, T. Furukawa, and H. F. Durrant-Whyte, "Optimal search for a lost target in a Bayesian world," in *Field and Service Robotics*. Springer, 2003, pp. 209–222.

[22] J. Vander Hook, P. Tokekar, and V. Isler, "Cautious greedy strategy for bearing-based active localization: Experiments and theoretical analysis," in *Int. Conf. on Robotics and Automation (ICRA)*, 2012, pp. 1787–1792.

[23] A. Ryan, "Information-theoretic tracking control based on particle filter estimate," in *Proceedings of the AIAA Guidance, Navigation, and Control Conf., Honolulu, HI, Aug*, 2008, pp. 18–21.

[24] B. Grocholsky, R. Swaminathan, J. Keller, V. Kumar, and G. Pappas, "Information driven coordinated air-ground proactive sensing," in *Proceedings of the IEEE Int. Conf. on Robotics and Automation*, 2005, pp. 2211–2216.

[25] G. Zhang, S. Ferrari, and C. Cai, "A comparison of information functions and search strategies for sensor planning in target classification," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 42, no. 1, pp. 2–16, 2012.

[26] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using rao-blackwellized particle filters." in *Robotics: Science and Systems*, vol. 2, 2005, pp. 65–72.

[27] L. M. Miller, Y. Silverman, M. A. MacIver, and T. D. Murphey, "Ergodic exploration of distributed information," *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 36–52, 2016.

[28] R. Martinez-Cantin, N. de Freitas, E. Brochu, J. Castellanos, and A. Doucet, "A Bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot," *Autonomous Robots*, vol. 27, no. 2, pp. 93–103, 2009.

[29] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," *arXiv preprint arXiv:0912.3995*, 2009.

[30] A. Krause and C. S. Ong, "Contextual Gaussian process bandit optimization," in *Advances in Neural Information Processing Systems*, 2011, pp. 2447–2455.

[31] B. Frank, C. Stachniss, N. Abdo, and W. Burgard, "Using Gaussian process regression for efficient motion planning in environments with deformable objects." in *Automated Action Planning for Autonomous Mobile Robots*, 2011.

[32] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press Cambridge, 1998, vol. 1, no. 1.

[33] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using Rao-Blackwellized particle filters." in *Robotics: Science and Systems*, vol. 2, 2005, pp. 65–72.

[34] F. Amigoni and V. Caglioti, "An information-based exploration strategy for environment mapping with mobile robots," *Robotics and Autonomous Systems*, vol. 58, no. 5, pp. 684–699, 2010.

[35] P.-P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich, "Viewpoint selection using viewpoint entropy." in *VMV*, vol. 1, 2001, pp. 273–280.

[36] X. Liao and L. Carin, "Application of the theory of optimal experiments to adaptive electromagnetic-induction sensing of buried targets," *IEEE Trans. Pattern Anal. Mach. Intellce*, vol. 26, no. 8, pp. 961–972, 2004.

[37] A. Singh, A. Krause, and W. J. Kaiser, "Nonmyopic adaptive informative path planning for multiple robots," in *Proceedings of the 21st Int. Jont Conf. on AI*, 2009, pp. 1843–1850.

[38] A. Krause and C. Guestrin, "Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach," in *International Conference on Machine learning*, 2007, pp. 449–456.

[39] K. H. Low, J. M. Dolan, and P. Khosla, "Adaptive multi-robot wide-area exploration and mapping," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*, 2008, pp. 23–30.

[40] R. Marchant and F. Ramos, "Bayesian optimisation for intelligent environmental monitoring," in *International Conference on Intelligent Robots and Systems*, 2012, pp. 2242–2249.

[41] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013.

[42] S.-W. Ryu, Y.-h. Lee, T.-Y. Kuc, S.-H. Ji, and Y.-S. Moon, "A search and coverage algorithm for mobile robot," in *Ubiquitous Robots and Ambient Intelligence (URAI), 2011 8th International Conference on*. IEEE, 2011, pp. 815–821.

[43] L. Paull, S. Saeedi, M. Seto, and H. Li, "Sensor-driven online coverage planning for autonomous underwater vehicles," *IEEE/ASME Transactions on Mechatronics*, vol. 18, no. 6, pp. 1827–1838, 2013.

[44] H. Choset, "Coverage for robotics–a survey of recent results," *Annals of mathematics and artificial intelligence*, vol. 31, no. 1-4, pp. 113–126, 2001.

[45] L. M. Miller and T. D. Murphey, "Trajectory optimization for continuous ergodic exploration," in *American Control Conf. (ACC), 2013*. IEEE, 2013, pp. 4196–4201.

[46] K. E. Petersen, *Ergodic theory*. Cambridge University Press, 1989, vol. 2.

[47] C. Hoffmann, S. Krause, E. Stoiber, A. Mohr, S. Rieken, O. Schramm, J. Debus, F. Sterzing, R. Bendl, and K. Giske, "Accuracy quantification of a deformable image registration tool applied in a clinical setting," *Journal of Applied Clinical Medical Physics*, vol. 15, no. 1, 2014.

[48] J. Ding, R. E. Goldman, K. Xu, P. K. Allen, D. L. Fowler, and N. Simaan, "Design and coordination kinematics of an insertable robotic effectors platform for single-port access surgery," *IEEE/ASME transactions on mechatronics*, vol. 18, no. 5, pp. 1612–1624, 2013.

[49] P. Kazanzides, Z. Chen, A. Deguet, G. S. Fischer, R. H. Taylor, and S. P. DiMaio, "An open-source research kit for the da vinci® surgical system," in *Int.Conf. on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 6434–6439.

[50] A. McHutchon and C. E. Rasmussen, "Gaussian process training with input noise," in *Advances in Neural Information Processing Systems*, 2011, pp. 1341–1349.