

ORIGINAL ARTICLE

Action Augmentation of Tactile Perception for Soft-Body Palpation

Luca Scimeca,¹ Josie Hughes,¹ Perla Maiolino,² Liang He,³ Thrishantha Nanayakkara,³ and Fumiya Iida¹

Abstract

Medical palpation is a diagnostic technique in which physicians use the sense of touch to manipulate the soft human tissue. This can be done to enable the diagnosis of possibly life-threatening conditions, such as cancer. Palpation is still poorly understood because of the complex interaction dynamics between the practitioners' hands and the soft human body. To understand this complex of soft body interactions, we explore robotic palpation for the purpose of diagnosing the presence of abnormal inclusions, or tumors. Using a Bayesian framework for training and classification, we show that the exploration of soft bodies requires complex, multi-axis, palpation trajectories. We also find that this probabilistic approach is capable of rapidly searching the large action space of the robot. This work progresses "robotic" palpation, and it provides frameworks for understanding and exploiting soft body interactions.

Keywords: soft tissue palpation, soft haptics, soft tactile perception

Introduction

THE PALPATION OF SOFT BODIES is a complex medical procedure where physicians palpate the human body for the diagnosis of abnormalities.^{1,2} Practitioners use their hands to explore and feel for abnormalities within the soft tissue of the patient's body, exploiting the physical structure and the sensing capabilities of the human hand.³ This action is widely used for the initial detection and screening of abnormalities within the body, aiding the diagnosis of conditions, including cancer,⁴ abdominal aortic aneurysm,⁵ appendicitis,⁶ and others.⁷⁻⁹

The complexity of this important procedure arises from the complex motions of the practitioner's hand that are in contact with the interacting layers of soft tissues of the human body, which can have many (or infinite) degrees of freedom. In the past, there have been notable attempts to better understand palpation by using robotics technologies.¹⁰⁻¹³ One of the pioneering works in this area, the WAPRO-4 system, is capable of performing simple breast palpation to identify

relatively large inclusions.¹⁴ The use of probes with variable mechanical impedance has been found to improve lumps and tumor detection,¹⁵ and the importance of sensory-motor coordination has also been shown across a number of medical applications.¹⁶⁻¹⁸ One of the key enabling technologies to improve robotic palpation capabilities is tactile sensing,¹⁹ which has led to the in-depth study of the use of tactile sensors for tumor localization.²⁰⁻²⁶ Robotics research has also attempted the development of technologies for medical teleoperation²⁷⁻³⁰ and medical training, such as haptic palpation training systems,³¹⁻³⁴ and virtual reality training systems.^{7,35,36} Finally, efforts have been made in applying machine learning for tumor localization and classification.³⁷⁻⁴¹

Within this body of existing work, there has been limited investigation of the impact of introducing diversity and complexity into the trajectory of the robot hand/probe during palpation. Previous work has only examined robotics palpation systems with simple one-axis vertical displacements,^{24,25,42,43} or horizontal sliding trajectories.^{25,34,44} In contrast, medical practitioners use complex

¹Bio-Inspired Robotics Laboratory, Department of Engineering, University of Cambridge, Cambridge, United Kingdom.

²Oxford Robotic Institute, University of Oxford, Oxford, United Kingdom.

³Morphological Computation and Learning Lab, Dyson School of Design Engineering, Imperial College London, London, United Kingdom.

examinations techniques, including rotations, twists, and percussions that are dependent on the specific body part under investigation.⁸

In the context of palpation, the quality of the tactile information, and hence the ability to make accuracy diagnosis, depends on the quality of these soft interaction as well as the tactile information arising from them. In this article, we hypothesize that the tactile information gained through the interactions between a sensor and the soft human body is improved by introducing complexity into the robot actions. The robot actions can enhance the richness of the physical stimuli arising from the soft interactions between the robot and the soft body to palpate, assisting classification of inclusions and hence diagnosis. As such, the challenge addressed in this article is the optimization of complex palpation trajectories to enable more accurate classification of abnormalities in soft bodies.

We use a six Degree of Freedom robot arm with a sensorized end-effector. To efficiently search the high dimensional action space, we utilize Bayesian inference (in the form of Bayesian Exploration). Bayesian approaches can leverage the cumulative past experiences to rapidly search motion trajectory parameters, and they allow for efficient search of high-dimensional action spaces. This search can enable the robot to select effective trajectories for accurate classification of hard inclusions in soft tissues.

In this article, the Materials and Methods section briefly reports the physical set-up for the experiments, before outlining the Bayesian framework developed for this work. In the Materials and Methods section, we report the results. The Exploring Action Complexity in Robot Medical Palpation section shows the complex relationship between robot palpation trajectory and the ability to perform accurate diagnosis. Bayesian Approaches for Confident Abnormality The Detection section focuses on the use of our Bayesian framework to perform confident diagnosis; however, we show the ability of our framework to find optimal palpation strategies efficiently. The discussion and conclusion are finally reported in the Discussion and Conclusion section.

Materials and Methods

The palpation experiments are performed by using a Robotic Arm with a sensing probe equipped with a capacitive tactile sensor array (Fig. 1). Although alternative sensor technologies could be used, the sensory technology chosen has a number of key advantages for use in palpation. The sensor provides pressure information from seven distributed “taxel” locations on the sensor surface, providing key spatial information with a high sensitivity, which is in line with that required for palpation.⁴⁵ The taxels respond with a bell-shaped curve and their receptive fields overlap,⁴⁶ allowing the detection of abnormal inclusions that are as small as 5 mm in diameter.

As this work focuses on the classification of hard inclusions, as opposed to their localization, we focus on point-based palpation trajectories, which revolve around a predefined point. This is in contrast to existing work, where the localization of inclusions was performed by using sliding trajectories.^{25,34,44} Each palpation experiment lasts 3 s, and, relative to the end-effector’s initial position, the trajectory varies in depth axis (Z), rotation around the x -axis (R_x), and rotation around the y -axis (R_y). As such, each 3D trajectory can be described by six constant motion parameters (A_{rx} , A_{ry} , A_z , ω_{rx} , ω_{ry} , and ω_z) (Supplementary Movie S1 and Supplementary Fig. S3).

The experiments are performed on three phantoms: two flat training phantoms, and a more human-like abdominal phantom (Fig. 1). Inspired by medical palpation of the liver, the Abdominal Phantom is a silicone phantom of a human liver embedded in a cross-sectional replica of a human torso. The replica introduces higher levels of complexity than the flat training phantoms, including a curved surface, skin, and tissues. All phantoms include stiff spherical inclusions of diameter 5, 10, or 15 mm, at a depth of 5 or 12 mm from the surface, as summarized in Figure 1b and c and Table 2. These sizes and depths mirror conditions in which inclusions are typically detected through palpation. More information about phantom development can be found in the Phantom Development section in Supplementary Data (Supplementary Fig. S1).

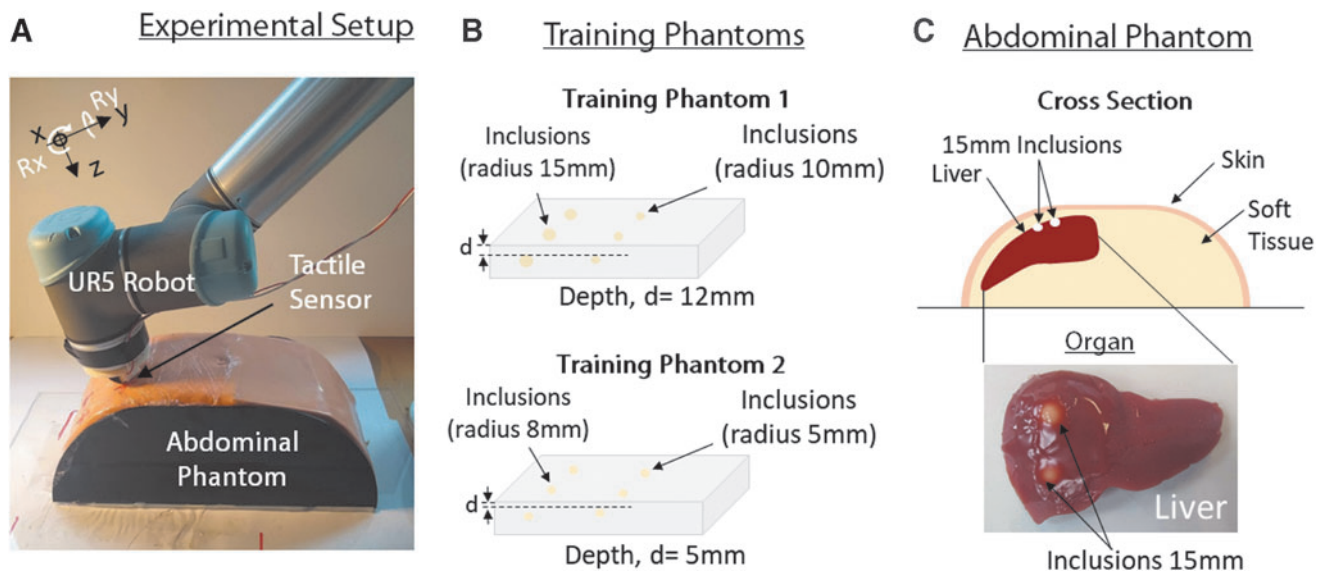


FIG. 1. Robotic medical palpation, including (a) the experimental setup, (b) the training phantoms, and (c) the abdominal phantom developed.

Using this setup, we validate the need for complex motion strategies, and we find those strategies that can improve soft tactile perception for the classification of hard inclusions. The experimental framework we propose to achieve this has three key phases: training, inference, and evaluation (Fig. 2).

Training phase

During the training phase, the robot generates sensory data by repeatedly palpating the different inclusions with different palpation trajectories. Each robot experiment involves a palpation trajectory, or action (A_m), being performed on a specific class of inclusion (C_k) in a phantom. These data are then represented probabilistically as probability density functions (PDFs).

Data sampling. Let \mathbf{X} be an $N \times D$ dimensional vector, where each unique temporal tactile image for a probed location is a D -dimensional row in the matrix. A temporal tactile image is a sequence of tactile images sampled at constant time intervals. Each tactile image corresponds to the normalized raw capacitance values of each taxel in the tactile sensor. By limiting the palpation to three seconds, we gain 35 pressure points over time limiting the dimensionality of the data ($D=35$). The value of N is not constant, but it instead increases with the number of palpation experiments performed. In each experiment, the value of N is initially 0 and for each ‘‘palpation iteration’’ $N=N+K$ where K is the number of discriminative classes, or types of inclusions in the phantom to palpate (examples of rows of \mathbf{X} can be visualized as heatmaps in Fig. 3a–d). The final total number of experiments can be computed by the product of three variables, that is, (number of actions) \times (number of inclusion classes) \times (number of samples). Each of these three variables changes depending on the experimental conditions, and it is reported in Table 1. We use

principal component analysis (PCA) to project the original tactile data \mathbf{X} onto its first principal component p_1 , obtaining a matrix \mathbf{W} , where each row w_i is a one-dimensional (1D) projection of the original 35-dimensional tactile sensor data. A detailed explanation of the dimensionality reduction processes is provided in the Supplementary Data.

Bayesian PDFs update. Using the sensor data, Bayesian approaches can be applied, including Bayesian Exploration, an approach first proposed for tactile discrimination of textures,⁴⁷ and other tactile discrimination tasks.^{48,49} In this article, we additionally derive a measure of confidence for each robot palpation trajectory. The mathematical details behind the representation proposed in this section can be found in Fishel and Loeb,⁴⁷ as well as in the Bayesian Treatment of Sensor Evidence section.

To represent the sensor data probabilistically, we use the 1D tactile evidence computed through PCA to generate a probability density function (PDF) for each class of inclusion C_k , and palpation trajectory A_m via:

$$p(w_i | C_k, A_m) = \frac{1}{\sqrt{(2\pi)^2 |\Sigma_{k,m}|}} e^{-\frac{1}{2}(w_i - \mu_{k,m})^T \Sigma_{k,m}^{-1} (w_i - \mu_{k,m})} \quad (1)$$

where $(\mu_{k,m})$ and $(\Sigma_{k,m})$ denote the mean and standard deviation of the 1D sensor data from a series of palpation iterations. By representing the tactile sensor data probabilistically, the width of the PDF captures the variation of the sensor data for a given palpation action and inclusion type (Supplementary Movie S2).

From these training data, we can generate two key metrics that help assess the quality of a palpation trajectory. The first is an ‘‘Unbiased Benefit Estimator’’ B_m^u , which provides a measure of how useful the sensor data from a are for

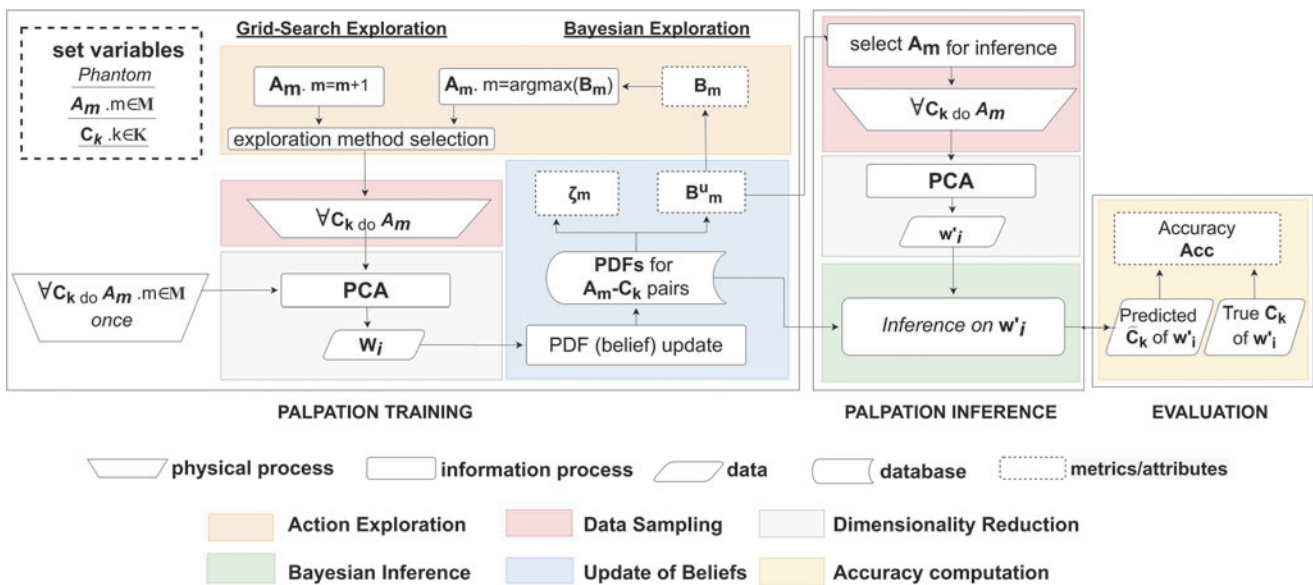


FIG. 2. Flowchart of experimental procedure. During the palpation training phase, the robot performs palpations A_m on different types of inclusions C_k to form PDFs. After an initial set of palpations to generate PDFs, the robot performs additional experiments to improve its classification capabilities based on the biased B_m score. In the palpation inference phase, the PDFs are used to perform inference on new samples. Moreover, an unbiased benefit and a confidence level for each palpation trajectory A_m can be estimated. In the evaluation phase, the performance of the robot can be evaluated if the ground truth classification of the palpated area is known. PDF, probability density function.

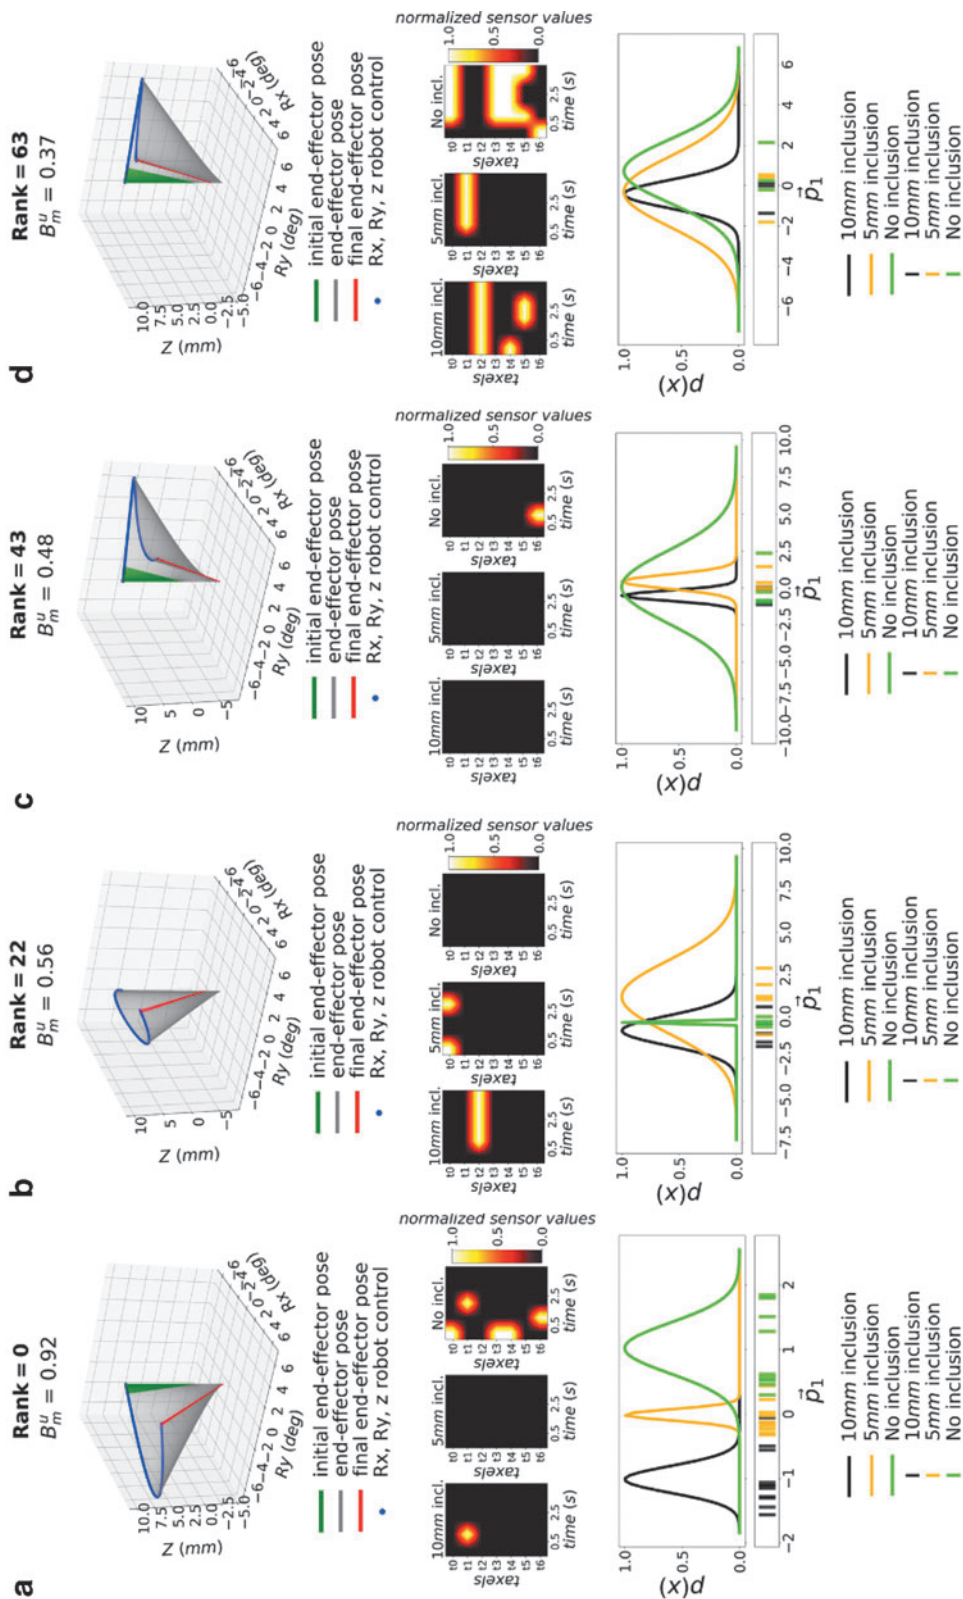


FIG. 3. Influence of the palpation trajectory to the PDFs in *Training Phantom 2*. The 64 robot actions are ranked by B_m^u and Rank 0-22-43-63 (best to worst) are shown, respectively, in (a-d). The *top* plots show the robot trajectory generated during palpation, where R_x , R_y are rotations about the x - y axis and Z is the probing depth. The *middle* plots show the normalized spatiotemporal tactile images generated during palpation. The brightness corresponds to the normalized tactile values at specific time intervals (proportional to pressure). The *bottom* plots show the PDFs generated from 60% of the data detailed in the “Exploration Experiments” of Table 2. Here, p_1 is the first principal component onto which the original sensor data were projected.

classification, and can be quantitatively measured by considering the overlap between the PDFs for the different classes of inclusions. This is based on the Bhattacharyya coefficient,⁵⁰ which for two probability density functions p and q is defined as:

$$\text{BCoeff} = \int \sqrt{p(x)q(x)}dx \quad (2)$$

Based on this coefficient, we can define the unbiased benefit estimator (\mathbf{B}_m^u), for a specific action A_m as:

$$\mathbf{B}_m^u = \frac{\sum_k P(C_k)^2}{\sum_s \Psi_{ks,m} P(C_s)} \quad (3)$$

where $\Psi_{ks,m}$ is calculated from the Bhattacharyya coefficient, and contains the mutual confusion between two classes and C_k and C_s are under action A_m . $P(C_k)$ and $P(C_s)$ are the prior probabilities of inclusions C_k and C_s , respectively. This \mathbf{B}_m^u score can be used to rank palpation trajectories. Using the properties of Gaussian distributions, a lower degree of overlap among distributions implies a higher likelihood for any new sample to fall into an area of sensor space belonging to either class unambiguously. As such, a higher overlap should imply a higher prediction accuracy.

One of the unique advantages of computing an unbiased benefit estimator is that we can also obtain a measure of confidence of the tactile sensor data for a specific palpation trajectory A_m . This measure of confidence, ζ_m , is defined as:

$$\zeta_m = 1 - \frac{1}{2 + e^{-\mathbf{B}_m^u}} \quad (4)$$

This metric increases monotonically when the discriminatory confusion reduces, and it signifies classification confidence for a specific trajectory.

Exploratory action identification. As the number of parameters that describe a palpation trajectory increases, there is an exponential increase in the number of actions to be searched. If each of the six action parameter can take on “ n ” possible values, there are up to n^6 trajectories to search. Efficient search i approaches are required, as it is neither practical nor feasible to perform searches of this scale for each new phantom or patient. Figure 2 illustrates how Bayesian Exploration can be implemented to find optimal palpation actions by iteratively selecting, and exploring, the most “promising” action.⁴⁷ This search requires a metric to guide the selection of actions. For this we can use a “biased” benefit score based on \mathbf{B}_m^u , that is:

$$\mathbf{B}_m = 1 - (1 - \mathbf{B}_m^u)^{\frac{1}{n_m}} \quad (5)$$

where n_m is the number of times action m was performed iteratively during the palpation experiments. Thus, the biased benefits are discounted by the number of times the action has already been performed during action exploration to discourage excessive exploitation and eventually encourage the explorative update of belief states under less exploited actions.

Initially, the robot palpates each class of inclusion under every action once, to gather initial experimental evidence. After this, the action is selected by using \mathbf{B}_m , each class is palpated by using this action, and the PDFs are then updated accordingly.

Bayesian inference phase

In the second phase of the framework, Bayesian inference, the robot performs the classification of abnormal inclusions, identifying the class of unseen sensor data obtained through additional robotic palpations. This classification is made via Bayesian Inference, using the PDFs generated in the palpation training phase through several palpation iterations. To perform inference on a new tactile sample w_i^0 , we evaluate the sample at $p(w_i^0|C_k, A_m)$, under every C_k for a chosen action A_m . The C_k of the PDF yielding the highest value will be inferred as a class for w_i^0 .

$$\tilde{C}_k = \operatorname{argmax}\{p(w_i^0|C_k, A_m) : k \in \mathcal{C}\} \quad (6)$$

where C_{ek} is the class estimated for C_k . This inference process is used throughout the results section to test the abilities of different palpation trajectories, and it will be referred to as “Bayesian inference classification.”

Evaluation phase

In this final phase, we evaluate the performance of the classification by comparing the “true” class of inclusion C_k against the class that was inferred. Over several iterations, we can count the number of correctly classified abnormal inclusions as True Positives, and the number of correctly classified inclusion-free areas of the phantom as True Negatives. For a total of N_C classifications, or palpation inferences, the accuracy can be formally computed as:

$$\text{Acc} = \frac{\text{TP} + \text{TN}}{N_C} \quad (7)$$

Results

Exploring action complexity in robot medical palpation

The first set of experiments investigates the influence of the palpation trajectory on soft tactile sensing capabilities and the ability to distinguish different classes of inclusion.

In these experiments, we examine 64 different palpation trajectories, and we analyze how they influence the separation of PDFs. The 64 palpation trajectories are generated through the combination of six parameters that describe the trajectory ($2^6=64$). We conducted these experiments on *Training Phantom 1*, performing all the palpation actions on all the different inclusion types. For each type of inclusion and palpation trajectory, we perform the palpation 20 times. This brings the total number of experiments to $64 \times 20 \times 3$, where 3 is the number of classes of inclusions present in the phantom. More details can be found in the “Exploration Experiments” column of Tables 1 and 2.

Figure 3 shows the PDFs for four exemplar palpation trajectories, ordered with respect to the \mathbf{B}_m^u scores. The PDFs are created via Equation (1), and the unbiased benefit score \mathbf{B}_m^u is computed for each palpation action by using Equation (3). The different motion parameters result in

TABLE 1. EXPERIMENTAL BREAKDOWN OF ROBOTIC PALPATIONS AND PALPATED PHANTOMS

<i>Phantom</i>	<i>Class of inclusions (mm)</i>	<i>Depth of inclusions (mm)</i>	<i>No. of inclusions</i>	<i>Complexity demonstration</i>	<i>Exploration experiments</i>	<i>Validation experiments</i>	<i>Total number</i>
<i>Training Phantom 1</i>	10	12	3	1920 (10 iterations per trajectory)	3840 (20 iterations per trajectory)	—	5760
	15	12	3				
	No incl.	N/A	4				
<i>Training Phantom 2</i>	5	5	3	—	3840 (20 iterations per trajectory)	—	3840
	8	5	3				
	Healthy	N/A	4				
<i>Abdominal test phantom</i>	15	12	2	—	800 (20 iterations per trajectory)		800
	No incl.	N/A	2				

The values in parenthesis represent the number of samples gathered by the robot for any trajectory (A_m) and class of inclusion (C_k) pair. N/A, not applicable.

different palpation trajectories with very diverse PDFs. In Figure 3, an example of reference raw tactile sensor data of the inclusions is shown in the middle figures, where the y-axis represents a layout of all taxels (Supplementary Fig. S2), and the x-axis the experiment time. The corresponding PCA projected points are shown in the x-plot of the lower figures, together with their corresponding PDFs. The raw tactile sensor data for each class are influenced by the palpation strategy itself. Ideally, PDFs for different inclusion classes should have minimal overlap for discrimination purposes. The results show how it is possible to have motion parameters, and hence trajectories, that give rise to PDFs that are fully separated across the PCA principal component p_1 for all classes of inclusion (Fig. 3a).

The figure also shows that the degree of these overlaps can be represented by using the \mathbf{B}''_m scores. The trajectories with less overlap (Fig. 3a, b) result in higher \mathbf{B}''_m scores, whereas those with more overlap (Fig. 3c, d) have a far lower score. As such, the score represents the discriminative performance of the palpation trajectory and can be used to compute a ranking for the different trajectories. \mathbf{B}''_m also indicates the degrees to which each PDF is separated from the others, in addition to measuring the amount of overlap. This allows similarly overlapped PDFs to be ranked; for example, although the PDFs of the 43rd and 63rd ranked actions are similarly overlapped, the former is ranked higher because of the lower overlap between the PDFs of the 10 and 5 mm inclusion classes.

In the next experiment, we compare the separation of the PDFs for the same trajectories but across different phantoms, that is, *Phantom 1*, *Phantom 2*, and the *Abdominal Phantom*. We perform this experiment to assess whether a palpation trajectory optimized for one phantom can perform well on other phantoms.

To achieve this, we identified the best trajectories for *Phantom 1*, *Phantom 2* and the *Abdominal Phantom*, by finding the trajectory with the highest \mathbf{B}''_m score for each phantom. These three top-ranking trajectories, together with the resulting PDFs, are compared in Figure 4a. The experimental data used are detailed in the ‘‘Exploration Experiments’’ of Tables 1 and 2. The first observation is that the optimum trajectories are significantly different for the different phantoms. The best palpation trajectory for *Phantom 1* is a counter clockwise rotation in an almost horizontal plane, whereas that for *Phantom 2* is a clockwise trajectory with a similar amplitude. The optimum trajectory for *Abdominal Phantom* is significantly different, with a clockwise rotation occurring with smaller amplitude, and a higher palpation depth. The second observation that can be made considers the PDF overlaps. Figure 4a shows that the highest ranked actions do not show high separation of the PDFs on the other phantoms. The best trajectory for *Phantom 1*, for example, does not perform well in *Phantom 2*, with the action resulting in high overlaps of PDFs belonging to different classes of inclusions. The *Abdominal Phantom* is a relatively easier task, in comparison to *Phantom 1* and *Phantom 2*, with all of three palpation strategies achieving high separation of the PDFs. However, the trajectory ranking higher for *Abdominal Phantom* still achieves higher separation of PDFs in the same phantom, whereas it does not perform well in *Phantom 1* and *Phantom 2*.

In the next set of experiments, we examine the need for more complex trajectories for more accurate palpation of soft bodies. This is achieved by comparing palpation trajectories that are described by a different number of control parameters. As each axis of motion is controlled by a specific pair of parameters (i.e., $A_{rx}-\omega_{rx}$, $A_{ry}-\omega_{ry}$, and $A_z\omega_z$ to control R_x , R_y , and Z , respectively); reducing the number of parameters decreases the complexity of the trajectory.

TABLE 2. EXPERIMENTAL BREAKDOWN OF ROBOTIC PALPATION TRAJECTORIES AND PARAMETERS OVER EXPERIMENTS

	<i>Complexity demonstration</i>	<i>Exploration experiments</i>	<i>Validation experiments</i>
No. of trajectories attempted	64	64	20
Parameter combinations	$A_{rx} \in [0, \frac{\pi}{18}]$ $A_{ry} \in [0, \frac{\pi}{18}]$ $A_z \in [0, 0.01]$ $\omega_{rx} \in [0, 1]$ $\omega_{ry} \in [0, 1]$ $\omega_z \in [0, 0.5]$	$A_{rx} \in [-\frac{\pi}{18}, \frac{\pi}{18}]$ $A_{ry} \in [-\frac{\pi}{18}, \frac{\pi}{18}]$ $A_z \in [0.002, 0.01]$ $\omega_{rx} \in [1, 3]$ $\omega_{ry} \in [1, 3]$ $\omega_z \in [0.5, 2]$	20 Highest score trajectories from ‘‘exploration experiments’’

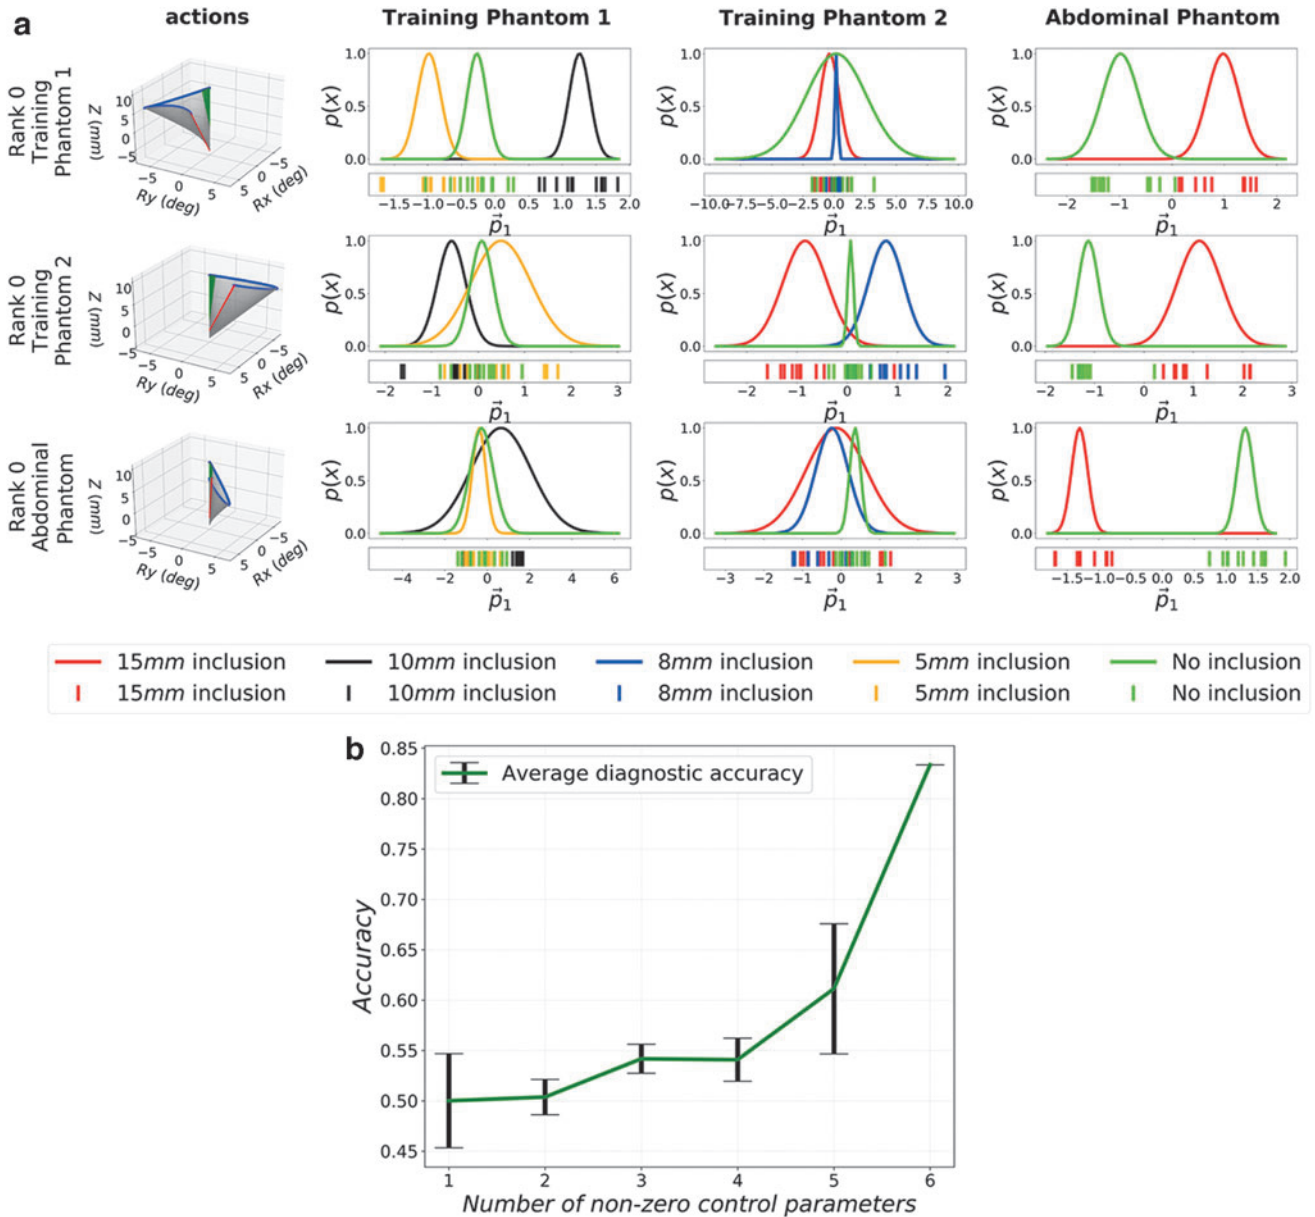


FIG. 4. The figure shows the complexity of robotics palpation. The diagonal plots in (a) show the PDFs of the best performing palpation trajectories for each phantom, whereas the off-diagonal plots show PDFs of the same trajectories in all other phantoms. (b) Shows the accuracy of a Bayesian Inference classifier trained on sensor data generated via palpation trajectories with a varying number of parameters.

To systematically vary and reduce the complexity, every possible combination of the six parameters is set to zero in turn. The 64 palpation trajectories defined by these parameters are performed 10 times on all types of inclusions in *Training Phantom 1*, and the corresponding tactile data are stored. These data correspond to tactile information from 2160 palpations; these are depicted in the columns of Tables 1 and 2.

To evaluate the performance of each set of motion parameters, Bayesian Inference classification is performed on the computed PDFs. The classification inference is performed on each palpation trajectory separately, with 60% of the sampled palpations used for training, and the remaining 40% used for testing. Figure 4b shows the average performance of the classifier across all palpation trajectories with different

numbers of active parameters. As illustrated in Figure 4b, trajectories described by one or two parameters achieve accuracy rates of 50% on average, thus a little above random selection (33%). With the full employment of the six descriptive parameters, the generated trajectories can achieve accuracies above 60%. As shown in Figure 4b, when the dimensionality of the actions, and hence number of motion parameters, is increased, there is up to 35% improvement in the average classification accuracy of the robot. This justifies and demonstrates the need for complex trajectories when performing palpation.

From this first set of experiments, we can make several conclusions. First, the palpation trajectory influences the tactile sensor data significantly, where, slight changes in the

palpation trajectory can significantly affect the discriminatory abilities of the robot. Second, the optimum trajectories vary from phantom to phantom. There is not one “optimum” motion for all phantoms. Third, introducing more complex palpation trajectories allows for better action profiles to emerge, demonstrating that increasingly complex actions increase the ability to make more accurate classification of abnormalities in soft bodies.

Bayesian approaches for confident abnormality detection

The next set of experiments examines the levels of confidence (ζ_m) and the experimental accuracy (Acc) when computing the PDFs based on a different number of training samples.

In these experiments, the same dataset from the previous experiments was used, where palpation training was

performed on each class of inclusion, using each action 20 times (see “Exploration Experiment” columns of Tables 1 and 2).

Out of the 20 palpation samples for each class-action pair, 40% of the data (corresponds to 8 samples) is held out for testing. For every trajectory, then, we consider the remaining 12 samples and compute the PDFs with a varying number of samples, from 1 to 12. Every time the PDFs are computed, we also compute the benefit and the confidence as previously described. The resulting PDFs are also used to compute the accuracy, as described in Materials and Methods Section.

We show that it is possible to achieve high classification accuracy if appropriate actions are selected. Figure 5a shows the highest accuracy of all palpation trajectories, as a function of the number of training samples used to compute the PDFs. As the number of training samples increases, the evidence used to build the PDFs increases, leading to the best classifiers performing more accurate classification. In this set of

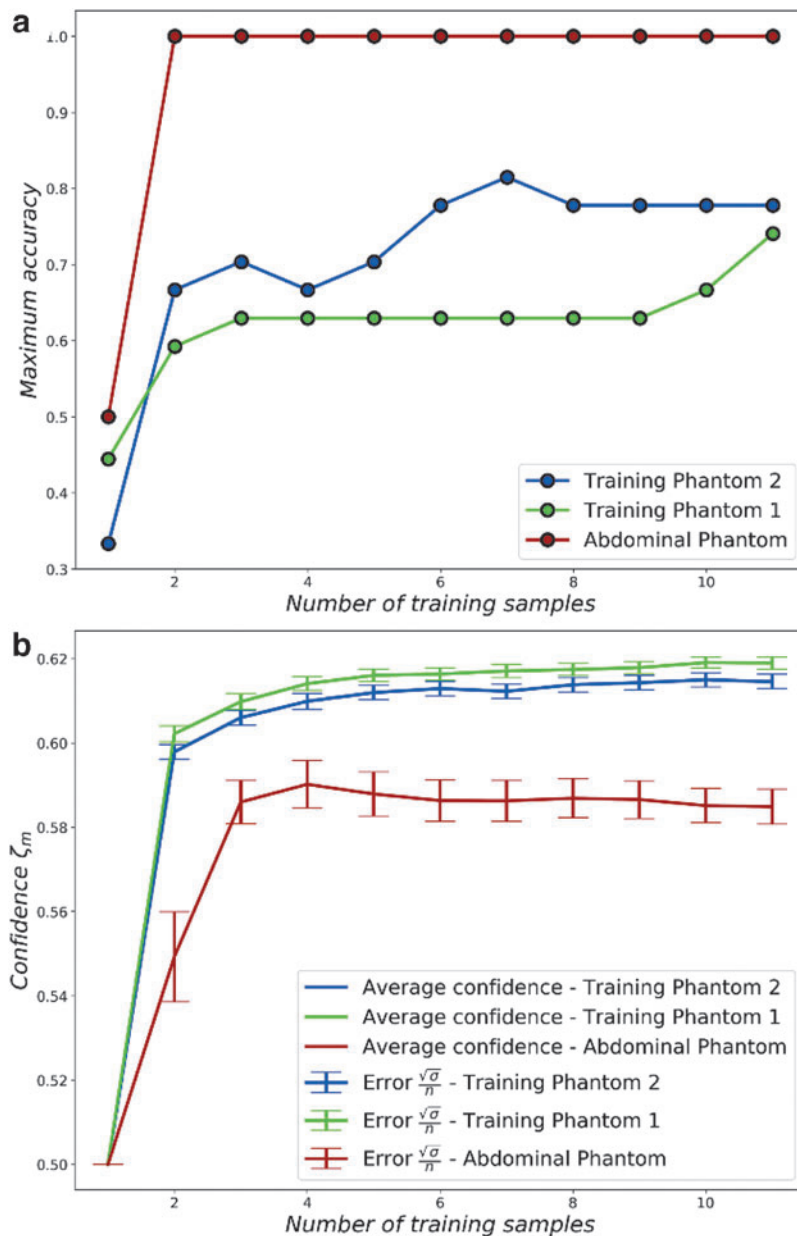


FIG. 5. The figure shows the performance of a Bayesian inference classifier within the framework developed. (a) Shows the relationship between the maximum classifier accuracy and the number of samples gathered for each palpation trajectory-class pair. (b) Shows the relationship between the developed confidence level and the number of samples gathered for each palpation trajectory-class pair. The vertical bars in the plot illustrate the errors of the confidence at that point.

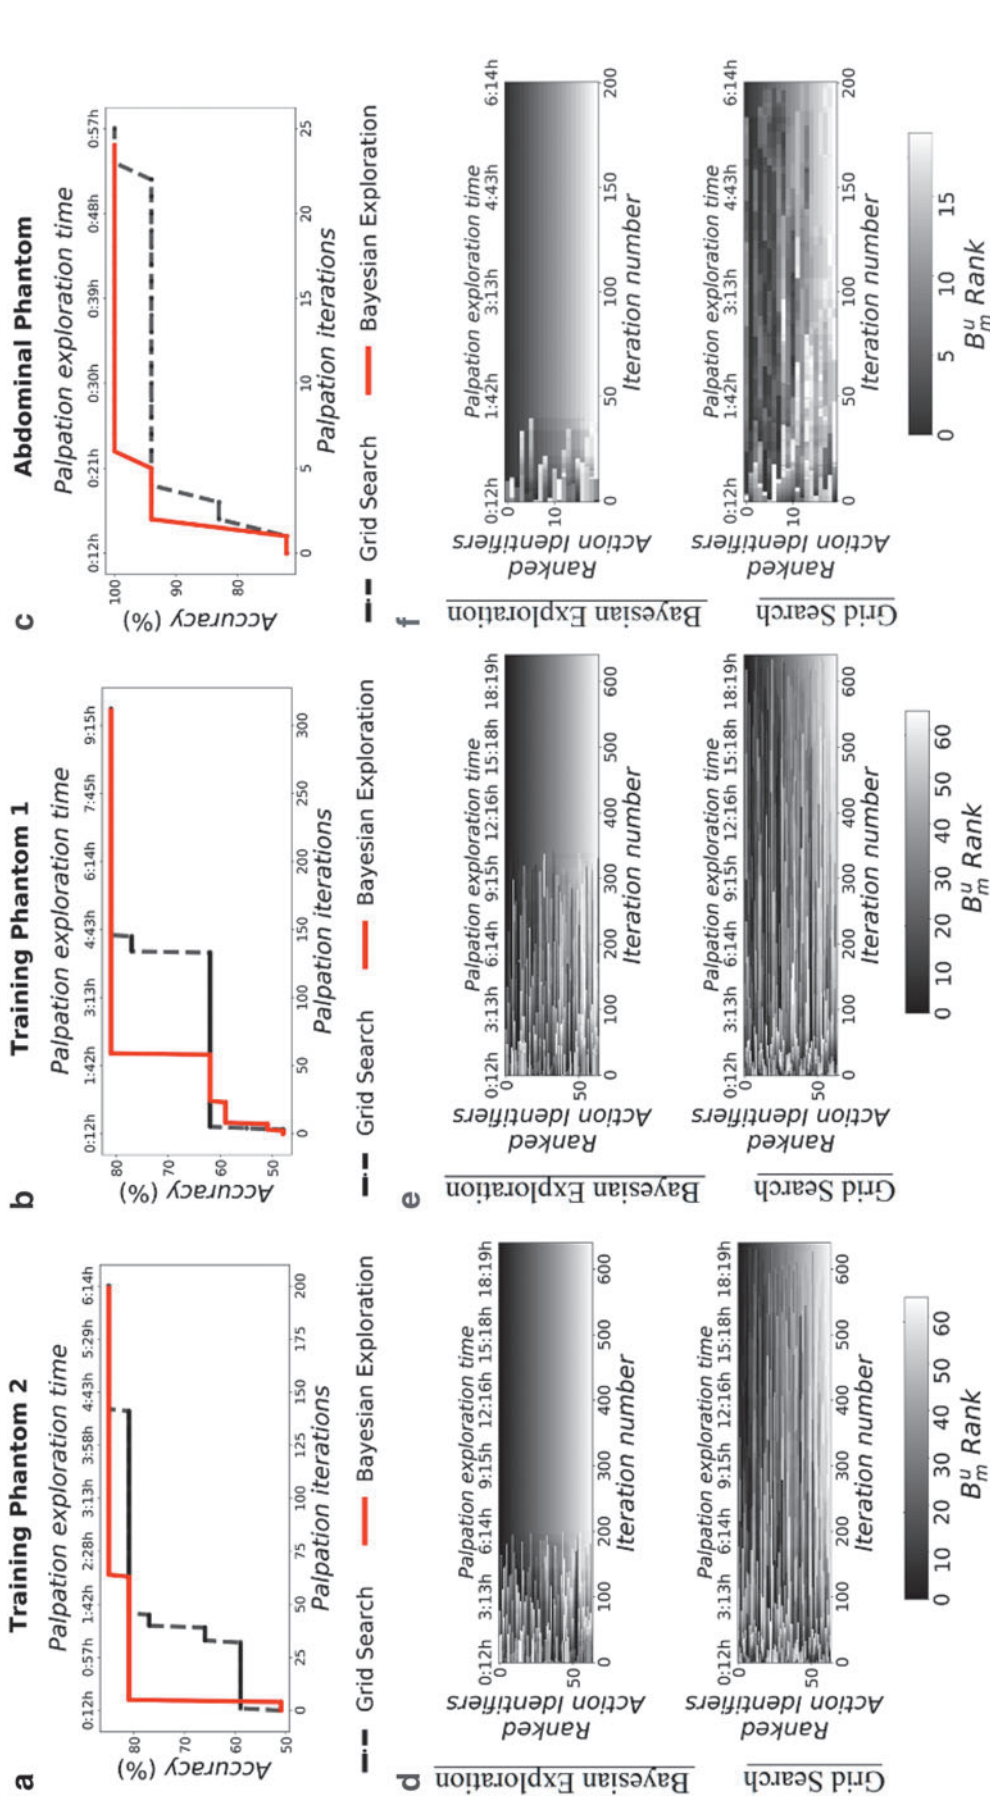


FIG. 6. Training comparison between Bayesian online exploration and Grid Search. (a–c) Show the maximal accuracy achieved at each training iteration by both Grid Search and Bayesian Exploration, on a left-out validation set of eight palpation samples for every class of inclusion. (d–f) Show the trajectories on the y-axis, ordered based on their final B_m^u rank during training for both Bayesian Exploration and Grid Search. The color intensity for each action indicates its rank at each iteration. Bayesian Exploration achieves a final ranking faster, bringing the robot’s trajectories to a ranked solution in at least half the time than when training through a systematic action space search.

TABLE 3. HIGHEST AND AVERAGE CLASSIFICATION ACCURACIES ACHIEVED BY THE PALPATION SYSTEM WHEN TRAINING THE BAYESIAN CLASSIFIER ON 14 SAMPLES OF EACH CLASS OF INCLUSION AND TESTING ON SIX UNSEEN SAMPLES

Accuracy % highest (average)	15 mm vs. NA	10 mm vs. NA	8 mm vs. NA	5 mm vs. NA	15 mm vs. 8 mm vs. NA	10 mm vs. 5 mm vs. NA
Training Phantom 1	0.944 (0.462)		0.778 (0.394)		0.740 (0.412)	
Training Phantom 2		0.889 (0.438)		0.833 (0.387)		0.778 (0.423)
Abdominal phantom	1.0 (0.85)					

NA, no inclusion.

experiments, the robot is also observed to reach maximum classification accuracy on the *Abdominal Phantom* by employing actions with more than two parameters.

In Figure 5b, the confidence metric is also plotted as a function of the number of samples used for training. As expected, we can see that the confidence metric increases with the number of training samples. The confidence, however, saturates at different values for each phantom. These values indicate how “reliable” the classification of the robot is under a specific trajectory. This measure will first and foremost depend on the overlap of the PDFs, which will, in turn, depend on the similarity of the tactile sensor data for different classes of inclusions. In Figure 5b, the robot achieves highest confidence for *Phantom 1*, followed by *Phantom 2* and the inclusion-action pair. This initial gathering of evidence allows the Bayesian Exploration process to then start (Fig. 2).

In this exploration process, all palpation trajectories are ranked by using the biased B_m score. The action that has the highest B_m score is then used to palpate each class of inclusion once, and the PDFs are updated with the new tactile information. This corresponds to one iteration of the Bayesian Exploration framework. The B_m score is then computed again and used to select the next palpation trajectory to test, with the steps then iteratively repeated. To evaluate each iteration of the exploration process, we take the top scoring action at that time, as defined by the unbiased benefit score, and use this action to perform Bayesian Inference. The inference is performed on 40% of unseen data from “Exploration Experiments,” and it provides the robot with the “best accuracy” for every iteration of the exploration process. Importantly, the top scoring action is selected by the unbiased benefit score, as we want to find the top performing action that is purely based on the ability to separate the PDFs in sensor space. As a benchmark, the results from the grid-search method are also presented. During grid search, contrarily to Bayesian Exploration, the action is selected based on a breadth-first parametric search, with the rest of the experiments performed in the same manner.

We compare the performance of these methods by considering the number of “palpation iterations” necessary to train the robot. As previously described, a “palpation iteration” involves the palpation of all classes of inclusions C_k under a specific action A_m . The action A_m is here iteratively selected through Bayesian Exploration or Grid-Search. As shown in Figure 6a–c, Bayesian Exploration achieved its highest performance after around only 60 iterations in both training phantoms. On the *Abdominal Phantom* this took ~ 7 iterations. Conversely, a grid-based systematic search performed poorly, finding equally good palpation strategies after

150 palpation iterations on the training phantoms, and 23 iterations on the *Abdominal Phantom*.

In Figure 6d–f, the intensity of the color shows the final ranking of the actions. The figure shows how this ranking is “unstable” for grid search, that is, the ranking keeps changing throughout the experiments, before reaching the final rank. Bayesian Exploration, however, induces a stable ranking much sooner, where the final ranking of trajectories is found much earlier on in the experiments.

By applying Bayesian Exploration, and leveraging the ranking provided by the score, the actions that best separate the PDFs across different classes of inclusions are preferentially explored. By using this exploration technique, the robot can efficiently search a high-dimensional parameter space. This complex high-dimensional action space has previously been demonstrated to be necessary for accurate classification of abnormal inclusions in soft tissues. From these results, we can observe that by using Bayesian Exploration, the time taken to find the optimal strategy is halved in comparison to a systematic grid search.

Finally, after performing Bayesian Exploration, we can report the final accuracy of the entire framework across all palpated phantoms. As previously explained, this is computed as the accuracy achieved on 40% of unseen palpation samples from each phantom. Table 3 reports the final highest test accuracy observed after training. Since the hypothesis in this article hinges on the postulate that appropriate palpation trajectories can aid in abnormality detection via palpation, we also report the average accuracy across all attempted palpation trajectories in Table 3. These results show how on average the palpation trajectories perform quite poorly, and appropriate optimization procedures are necessary to find the highest performing palpations. This highlights the importance of Bayesian Exploration in this context.

Notably, the system is capable of achieving more than 80% accuracy when discriminating between 5 mm inclusions and no inclusions. On the *Abdominal Phantom*, the robot achieves 100% accuracy when discriminating between 15 mm inclusions and no inclusions. Moreover, the highest performing motion strategies outperform the average performance of any one action by approximately a factor of two in almost all scenarios, confirming and emphasizing the need for appropriate palpation trajectories during abnormality detection.

Discussion and Conclusion

Medical palpation is an impactful preliminary diagnosis tool that is used widely by primary care physicians, yet it is extremely challenging for a robot to perform due to the

complexity of the interactions. The interactions between the palpation device and the soft human body are nonlinear; the complexity of the action space and the interactions is significant; and the solutions are different for every “patient.” Thus, to gain a more insightful understanding of this problem, we need to go beyond typical robotic approaches, including modeling and optimization. In this work, we perform large-scale physical experiments to understand whether and how multi-axis palpation trajectories can influence a robot’s soft tactile perception to make accurate classification of abnormal inclusions in soft bodies. The framework presented in this work (Fig. 2) allows for the fast exploration of a high-dimensional action space, which arises from the palpation of soft bodies. The framework identifies palpation strategies that allow for a confident classification of the presence, or absence, of abnormal inclusions. The identified palpation strategies have been shown to enable the confident detection of abnormal inclusions that are as small as 5 mm in diameter (Table 3).

In this experimental approach to palpation, we have identified that increasing the complexity of the palpation trajectory can be beneficial for soft tactile perception in the context of palpation. In addition, we have shown that slight changes in the trajectory, or the patient, significantly affect the performance. This demonstrates that the optimum palpation trajectory must be found or identified for each patient through physical experimentation, and mirrors the method in which human practitioners find the best palpation motion for each patient. To make intelligent decisions in this soft, nonlinear, and highly complex space, we have demonstrated how a probabilistic Bayesian approach allows for accurate and efficient search and decision making. However, the parameterization of the trajectory is still based on human design and intuition, and as such, they are limited. In future scenarios, the parameterization and trajectories would ideally emerge from the haptic interaction with the soft tissue itself.

Going forward, this knowledge is important in several ways. In the long term, we can use the methods to develop “robot doctors” who can perform accurate and confident diagnosis. The framework development provides a starting point for the experiment procedure for such a robot. However, to achieve this, it is necessary to find appropriate ways to perform knowledge transfer across patients or phantoms. In the short term, we can use this understanding to improve robot tactile sensing in soft environments/settings. We can also apply the methods and approaches to other similar problems, where the Bayesian treatment and large-scale physical experiments would further our understanding of the problem at hand.

Finally, it would be interesting to explore the relationship between the investigated, point-based, palpation trajectories and sliding trajectories explored for localization.

Author Disclosure Statement

No competing financial interests exist.

Funding Information

This work was funded by the U.K. Agriculture and Horticulture Development Board (CP 172), Physical Sciences Research Council (EPSRC) MOTION grant [EP/N03211X/2] and RoboPatient grant [EP/T00519X/1 and EP/T00603X/1].

Supplementary Material

Supplementary Data
 Supplementary Figure S1
 Supplementary Figure S2
 Supplementary Figure S3
 Supplementary Movie S1
 Supplementary Movie S2

References

1. Bosch AM, Kessels AG, Beets GL, *et al.* Preoperative estimation of the pathological breast tumour size by physical examination, mammography and ultrasound: a prospective study on 105 invasive tumours. *Eur J Radiol* 2003;48:285–292.
2. Thomson H, Francis D. Abdominal-wall tenderness: a useful sign in the acute abdomen. *Lancet* 1977;310:1053–1054.
3. Bickley LS, Szilagy PG, Bates B. *Bates’ Guide to Physical Examination and History Taking*. Philadelphia, PA: Lippincott Williams & Wilkins, 2009.
4. Schiessel R, Wunderlich M, Herbst F. Local recurrence of colorectal cancer: effect of early detection and aggressive surgery. *Br J Surg* 1986;73:342–344.
5. Lederle FA, Simel DL. Does this patient have abdominal aortic aneurysm? *JAMA* 1999;281:77–82.
6. Hardin DM, Jr. Acute appendicitis: review and update. *Am Fam Physician* 1999;60:2027.
7. Ferguson CM. Inspection, auscultation, palpation, and percussion of the abdomen. In: Walker HK, Hall WD, Hurst JW (eds) *Clinical Methods: The History, Physical, and Laboratory Examinations*. 3rd ed. Boston: Butterworths, 1990.
8. Li CG, Mahon C, Sweeney N, *et al.* PPAR γ interaction with UBR5/ATMIN promotes DNA repair to maintain endothelial homeostasis. *Cell Rep* 2019;26:1333–1343.e7.
9. McDonald S, Saslow D, Alciati MH. Performance and reporting of clinical breast examination: a review of the literature. *CA Cancer J Clin* 2004;54:345–361.
10. Sarvazyan A, Egorov V, Sarvazyan N. Tactile sensing and tactile imaging in detection of cancer. In: *Biosensors and Molecular Technologies for Cancer Diagnostics*. Abingdon, United Kingdom: Taylor & Francis, 2012:337–352.
11. Scimeca L, Iida F, Maiolino P, *et al.* Human-robot medical interaction. In: *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI ’20)*. New York, NY: Association for Computing Machinery. 2020:660–661.
12. Scimeca L, Maiolino P, Bray E, *et al.* Structuring of tactile sensory information for category formation in robotics palpation. *Auton Robot* 2020;44:1377–1393.
13. Palacio-Torralba J, Reuben RL, Chen Y. A novel palpation-based method for tumor nodule quantification in soft tissue—computational framework and experimental validation. *Med Biol Eng Comput* 2020;58:1369–1381.
14. Kato I, Koganezawa K, Takanishi A. Automatic breast cancer palpation robot: Wapro4. *Adv Robot* 1988;3:251–261.
15. Herzig N, Maiolino P, Iida F, *et al.* A variable stiffness robotic probe for soft tissue palpation. *IEEE Trans Robot* 2018;3:1168–1175.
16. Cleary K, Nguyen C. State of the art in surgical robotics: clinical applications and technology challenges. *Computer Aided Surg* 2001;6:312–328.

17. Taylor RH, Menciassi A, Fichtinger G, *et al.* Medical robotics and computer-integrated surgery. In: Siciliano B, Khatib O (Eds). Springer Handbook of Robotics. Cham: Springer, 2016:1657–1684.
18. Shen Z, Chen F, Zhu X, *et al.* Stimuli-responsive functional materials for soft robotics. *J Mater Chem B* 2020;8:8972–8991.
19. Shih B, Shah D, Li J, *et al.* Electronic skins and machine learning for intelligent soft robots. 2020;5:eaz9239.
20. Davaria S, Najafi F, Mahjoob M, *et al.* Design and fabrication of a robotic tactile device for abdominal palpation. In: 2014 Second RSI/ISM International Conference on Robotics and Mechatronics (ICRoM). IEEE 2014:339–344.
21. Talasaz A, Patel RV. Telerobotic palpation for tumor localization with depth estimation. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE 2013:463–468.
22. Daniulaitis V, Alhalabi MO, Kawasaki H, *et al.* Medical palpation of deformable tissue using physics-based model for haptic interface robot (HIRO). In: 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No. 04CH37566), vol. 4. IEEE 2004: 3907–3911.
23. Sornkarn N, Nanayakkara T. The efficacy of interaction behavior and internal stiffness control for embodied information gain in haptic perception. In: IEEE International Conference on Robotics and Automation (ICRA) 2016. IEEE, 2016:2657–2662.
24. Hui JC, Block AE, Taylor CJ, *et al.* Robust tactile perception of artificial tumors using pairwise comparisons of sensor array readings. In: 2016 IEEE Haptics Symposium (HAPTICS). IEEE, 2016:305–312.
25. Gwilliam JC, Pezzementi Z, Jantho E, *et al.* Human vs. robotic tactile sensing: detecting lumps in soft tissue. In: 2010 IEEE Haptics Symposium. IEEE, 2010:21–28.
26. Li B, Shi Y, Fontecchio A, *et al.* Mechanical imaging of soft tissues with a highly compliant tactile sensing array. *IEEE Trans Biomed Eng* 2017;65:687–697.
27. Go E, Del Pozo F, Quiles J, *et al.* A telemedicine system for remote cooperative medical imaging diagnosis. *Comput Methods Programs Biomed* 1996;49:37–48.
28. Mort M, May CR, Williams T. Remote doctors and absent patients: acting at a distance in telemedicine? *Sci Technol Hum Values* 2003;28:274–295.
29. Saracino A, Vrieling TO, Menciassi A, *et al.* Haptic intracorporeal palpation using a cable-driven parallel robot: a user study. *IEEE Trans Biomed Eng* 2020;67:3452–3463.
30. Patlan-Rosales P, Krupa A. Robotic assistance for ultrasound elastography providing autonomous palpation with teleoperation and haptic feedback capabilities. In: 8th IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics, BioRob. 2020: 1018–1023.
31. Kim S-Y, Kyung K-U, Park J, *et al.* Real-time area-based haptic rendering and the augmented tactile display device for a palpation simulator. *Adv Robotics* 2007;21:961–981.
32. Ullrich S, Kuhlen T. Haptic palpation for medical simulation in virtual environments. *IEEE Trans Visual Comput Graphics* 2012;18:617–625.
33. Kim S-Y, Park J, Kwon D-S. Palpation simulator for laparoscopic surgery with haptic feedback. In: Proceedings of the Second International Conference on Biomedical Engineering, Innsbruck, Austria, 2004:478–482.
34. Trejos AL, Jayender J, Perri M, *et al.* Robot-assisted tactile sensing for minimally invasive tumor localization. *Int J Robot Res* 2009;28:1118–1133.
35. Yakubo S, Kinoshita Y, Aki T, *et al.* Improvement of a simulator production project for abdominal palpation in kampo medical training. *Kampo Med* 2008;59:595–600.
36. Hamza-Lup FG, Bogdan CM, Seitan A. Haptic simulator for liver diagnostics through palpation. *MMVR* 2012;19: 156–160.
37. Nichols KA, Okamura AM. Autonomous robotic palpation: machine learning techniques to identify hard inclusions in soft tissues. In: 2013 IEEE International Conference on Robotics and Automation. IEEE, 2013: 4384–4389.
38. Nichols KA, Okamura AM. Methods to segment hard inclusions in soft tissue during autonomous robotic palpation. *IEEE Trans Robot* 2015;31:344–354.
39. Ayvali E, Srivatsan RA, Wang L, *et al.* Using bayesian optimization to guide probing of a flexible environment for simultaneous registration and stiffness mapping. In: 2016 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2016:931–936.
40. Garg A, Sen S, Kapadia R, *et al.* Tumor localization using automated palpation with gaussian process adaptive sampling. In 2016 IEEE International Conference on Automation Science and Engineering (CASE). IEEE, 2016: 194–200.
41. Xiao B, Xu W, Guo J, *et al.* Depth estimation of hard inclusions in soft tissue by autonomous robotic palpation using deep recurrent neural network. *IEEE Trans Automation Sci Eng* 2020;17:1791–1799.
42. Dario P, Bergamasco M. An advanced robot system for automated diagnostic tasks through palpation. *IEEE Trans Biomed Eng* 1988;35:118–126.
43. Hui JC, Kuchenbecker KJ. Evaluating the biotac’s ability to detect and characterize lumps in simulated tissue. In: International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. Berlin, Heidelberg: Springer, 2014:295–302.
44. Hughes J, Maiolino P, Nanayakkara T, *et al.* Sensorized phantom for characterizing large area deformation of soft bodies for medical applications. In: 2020 3rd IEEE International Conference on Soft Robotics (RoboSoft). IEEE, 2020:278–284.
45. Konstantinova J, Li M, Mehra G, *et al.* Behavioral characteristics of manual palpation to localize hard nodules in soft tissues. *IEEE Trans Biomed Eng* 2014;61:1651–1659.
46. Maiolino P, Maggiali M, Cannata G, *et al.* A flexible and robust large scale capacitive tactile system for robots. *IEEE Sensors J* 2013;13:3910–3917.
47. Fishel JA, Loeb GE. Bayesian exploration for intelligent identification of textures. *Front Neurobotics* 2012;6:4.
48. Xu D, Loeb GE, Fishel JA. Tactile identification of objects using Bayesian exploration. In: 2013 IEEE International Conference on Robotics and Automation. IEEE, 2013: 3056–3061.
49. Scimeca L, Maiolino P, Iida F. Efficient Bayesian exploration for soft morphology action co-optimization. In: 2020 IEEE International Conference on Soft Robotics (RoboSoft). IEEE, 2020:639–644.
50. Ray S. On a theoretical property of the bhattacharyya coefficient as a feature evaluation criterion. *Pattern Recognit Lett* 1989;9:315–319.

51. Jiang A, Xynogalas G, Dasgupta P, *et al.* Design of a variable stiffness flexible manipulator with composite granular jamming and membrane coupling. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, 2012:2922–2927.
52. Scimeca L, Maiolino P, Cardin-Catalan D, *et al.* Non-destructive robotic assessment of mango ripeness via multi-point soft haptics. In: 2019 International Conference on Robotics and Automation (ICRA). IEEE, 2019:1821–1826.
53. Scimeca L, Hughes J, Maiolino P, *et al.* Model-free soft-structure reconstruction for proprioception using tactile arrays. IEEE Robot Automation Lett 2019;4:2479–2484.
54. Scimeca L, Maiolino P, Iida F. Soft morphological processing of tactile stimuli for autonomous category formation. In: 2018 IEEE International Conference on Soft Robotics (RoboSoft). IEEE, 2018;639–644.
55. Tipping ME, Bishop CM. Probabilistic principal component analysis. J R Stat Soc B Stat Methodol 1999;61:611–622.
56. Armstrong K, Handorf EA, Chen J, *et al.* Breast cancer risk prediction and mammography biopsy decisions: a model-based study. Am J Prev Med 2013;44:15–22.

Address correspondence to:

Fumiya Iida
Bio-Inspired Robotics Laboratory
Department of Engineering
University of Cambridge
Trumpington Street
Cambridge CB2 1PZ
United Kingdom

E-mail: fi224@cam.ac.uk