

Autonomous Robotic Tissue Palpation and Abnormalities Characterisation via Ergodic Exploration

Luca Beber¹, Graduate Student Member, IEEE, Edoardo Lamon¹, Member, IEEE, Matteo Saveriano¹, Senior Member, IEEE, Daniele Fontanelli¹, Fellow, IEEE, and Luigi Palopoli¹, Senior Member, IEEE

Abstract—We propose a novel autonomous robotic palpation framework for real-time elastic mapping during tissue exploration using a viscoelastic tissue model. The method combines force-based parameter estimation using a commercial force/torque sensor with an ergodic control strategy driven by a tailored Expected Information Density, which explicitly biases exploration toward diagnostically relevant regions by jointly considering model uncertainty, stiffness magnitude, and spatial gradients. An Extended Kalman Filter is employed to estimate viscoelastic model parameters online, while Gaussian Process Regression provides spatial modelling of the estimated elasticity, and a Heat Equation Driven Area Coverage controller enables adaptive, continuous trajectory planning. Simulations on synthetic stiffness maps demonstrate that the proposed approach achieves better reconstruction accuracy, enhanced segmentation capability, and improved robustness in detecting stiff inclusions compared to Bayesian Optimisation-based techniques. Experimental validation on a silicone phantom with embedded inclusions emulating pathological tissue regions further corroborates the potential of the method for autonomous tissue characterisation in diagnostic and screening applications.

Index Terms—Medical robots and systems, sensor-based control, force and tactile sensing.

I. INTRODUCTION

PALPATION is a fundamental clinical procedure that allows physicians, by touch, to assess tissue stiffness, texture, and

Received 15 October 2025; accepted 23 February 2026. Date of publication 13 March 2026; date of current version 7 April 2026. This article was recommended for publication by Associate Editor G. Dagnino and Editor J. Burgner-Kahrs upon evaluation of the reviewers' comments. This work was supported in part by PNRR Project FAIR - Future AI Research under Grant PE00000013, in part by European Union Projects INVERSE under GA no. 101136067, in part by MAGICIAN under GA no. 101120731, and in part by the MUR Departments of Excellence 2023-27 Program under Grant L.232/2016. (Corresponding author: Luca Beber.)

Luca Beber is with the Interdepartmental Robotics Labs (IDRA), Università di Trento, 38123 Trento, Italy, also with Department of Information Engineering and Computer Science, Università di Trento, 38123 Trento, Italy, and also with the PhD Programme in Robotics and Intelligent Machines (DRIM), University of Genoa, 16126 Genoa, Italy (e-mail: luca.beber@unitn.it).

Edoardo Lamon and Luigi Palopoli are with the Interdepartmental Robotics Labs (IDRA), Università di Trento, 38123 Trento, Italy, and also with Department of Industrial Engineering, Università di Trento, 38123 Trento, Italy (e-mail: edoardo.lamon@unitn.it).

Matteo Saveriano and Daniele Fontanelli are with the Interdepartmental Robotics Labs (IDRA), Università di Trento, 38123 Trento, Italy, and also with Department of Industrial Engineering, Università di Trento, 38123 Trento, Italy (e-mail: matteo.saveriano@unitn.it).

Digital Object Identifier 10.1109/LRA.2026.3673907




Fig. 1. Experimental setup for the proposed ergodic search: (bottom-left) silicone sample; (top-left) ground-truth elasticity map obtained using a grid-based approach; (right) the robotic manipulator palpating the silicone sample.

the presence of abnormalities such as tumours or swollen lymph nodes [1]. Its speed and non-invasiveness make it crucial for early diagnosis, particularly of conditions such as breast cancer. However, palpation is inherently subjective and sensitive to inter-practitioner variability, which may compromise diagnostic accuracy and repeatability for small or deep-seated lesions. For this reason, physical examinations are often complemented or replaced by imaging techniques such as mammography, the gold standard for early detection [2], and ultrasound [3].

Accurate assessment of tissue mechanics is equally critical in robot-assisted minimally invasive surgery (RMIS), where robots must adapt to variable stiffness for safe navigation and manipulation, especially for autonomous systems operating with minimal human supervision [4]. To support such needs, several technologies have emerged. Imaging-based elastography methods, e.g., ultrasound [5], magnetic resonance [6], or laser-based methods [7], estimate tissue stiffness but typically provide relative measures, are costly or slow, and cannot be easily integrated into RMIS. Tactile sensing, in contrast, offers real-time, quantitative stiffness measurements by capturing force and deformation. Previous works have shown its potential to reduce applied forces, shorten procedures, and improve surgical outcomes [8], [9].

Early stiffness mapping used uniform probing grids [10], [11], while later methods adopted Bayesian Optimisation (BO) to reduce the number of probing actions [12], [13], [14], often coupled with custom tactile sensors [8], [9], [15], guided by anatomical priors [16] or with ad-hoc refined segmentation strategies [15], [17].

In this work, we propose a method for autonomous robotic palpation that enables continuous estimation of elastic tissue stiffness using a viscoelastic contact model, relying solely on a commercial force/torque (F/T) sensor. The use of a viscoelastic model allows more accurate representation of tissue dynamics during palpation compared to purely elastic formulations, while abnormality characterisation in this work focuses on the estimated elastic stiffness. This sensing strategy provides a practical alternative to methods based on embedded or specialised sensors, which typically require customised hardware integration and lack compatibility with standard robotic platforms. In contrast, the use of an off-the-shelf F/T sensor ensures portability, cost-effectiveness, and broader applicability of the proposed framework. We couple a force-based contact model implemented via the Dimensionality Reduction Method (DRM) [18] with an ergodic exploration strategy to efficiently construct stiffness maps, where the robot explores the environment in proportion to the Expected Information Density (EID). To plan trajectories, we employ the Heat Equation Driven Area Coverage (HEDAC) algorithm [19], a real-time ergodic controller that balances exploration and exploitation at 100Hz. This replaces earlier ergodic control methods, like Spectral Multiscale Coverage (SMC), which was limited to 4Hz as reported in [16]. The proposed framework is evaluated in simulation and real-world experiments, including validation on a silicone sample with a stiff inclusion (Fig. 1), demonstrating robust performance and, in some cases, superior detection and boundary delineation compared to BO. To summarise, the main contributions of this paper are:

- The design of an EID tailored to autonomous stiffness mapping, which explicitly combines uncertainty, stiffness magnitude, and spatial gradients to guide ergodic exploration toward diagnostically relevant regions;
- A closed-loop combination between online force-based viscoelastic parameter estimation and ergodic trajectory planning, where stiffness estimates continuously reshape the information objective driving exploration;
- A real-time implementation of ergodic palpation using HEDAC, enabling continuous exploration-exploitation trade-offs at 100 Hz with continuous motion;
- A convergence-based stopping criterion derived from the ergodic metric.

II. RELATED WORK

Typically, search algorithms rely on three key components: (i) a model of the underlying distribution, (ii) a strategy to select the next sampling location, and (iii) a mechanism to store and update the collected data. Although Gaussian Process Regression (GPR) [20] is widely adopted to model spatially varying physical properties (ii), such as tissue stiffness, the choice of sampling strategy (i) and planning (iii) remains an active area of research.

A. Bayesian Optimisation for Palpation

Several studies have employed BO to guide robotic palpation to reduce the number of probing actions while maximising

information gain. For example, Garg et al. [12] investigated how different acquisition functions affect the estimated stiffness map, showing how their choice influences both peak detection and regional estimation. Yan et al. [15] highlighted a key limitation of BO: while it excels at locating stiffness peaks, it struggles to delineate their boundaries accurately. To address this, they proposed a two-step approach combining BO-based search and boundary refinement using a radial strategy originating from the centroid of the estimated region. However, this strategy requires a reliable estimation of the centroid, which may not be guaranteed after exploration, and may be less adaptive to irregular or non-convex boundaries. Expected Improvement has been used as an acquisition function to improve stiffness estimation [13], and later extended with a utility-guided approach [16] that incorporates prior beliefs and penalises excessive motion. Chalasani et al. proposed a method that employs GPR to estimate tissue stiffness, which is then used as a prior to fit a second GPR model for reconstructing the tissue surface geometry [21]. In a follow-up study, they extended this approach to support online stiffness mapping during surgery [22]. A critical limitation of BO is that exploration and exploitation are handled sequentially through the choice of the acquisition function. As a result, BO may focus excessively on promising regions, thereby overlooking smaller or subtler areas that could yield informative measurements.

B. Ergodic Exploration

Ergodic control has emerged as a compelling alternative to BO for tasks that involve spatially distributed information, where exploration must account for system dynamics. Unlike point-wise optimisation strategies, which select discrete sampling locations, ergodic control generates trajectories that allocate time across the environment proportionally to the EID, making it well suited for continuous spatial estimation tasks. Mathew and Mezić [23] introduced the Spectral Multiscale Coverage (SMC) framework, using Fourier-based metrics to quantify how well a trajectory covers a given spatial distribution. Miller et al. [24] demonstrated the use of ergodic control for information-guided exploration, integrating the Linear Quadratic Regulator (LQR) formulation of SMC [25]. However, the high computational cost of optimisation limited real-time adaptation. Ayvali et al. [16] presented the first application of ergodic control to palpation tasks using SMC, concluding that BO achieved superior performance, particularly when prior information was available. Conversely, our results indicate the opposite trend, highlighting the potential of ergodic control for robotic palpation. The main reason is that SMC often prioritises global exploration and can suffer from over-sampling already explored regions, while its reliance on spectral transforms poses challenges for real-time deployment. To address these limitations, Ivic et al. [19] proposed HEDAC, which leverages radial basis functions and a potential field derived from a stationary heat equation to compute smooth local trajectories. This method improves scalability and enables real-time implementation, avoiding the computational burden of Fourier analysis.

III. METHODOLOGY

In this work, we combine online force-based parameter estimation using a viscoelastic tissue model with real-time ergodic control implemented via HEDAC, enabling continuous robotic palpation without tactile sensors and without sacrificing spatial coverage or responsiveness. While ergodic control and online

viscoelastic estimation have been previously investigated, we show that effective ergodic palpation critically depends on the definition of the target distribution guiding exploration. We therefore propose an EID tailored to stiffness mapping and lesion detection, which balances exploration and exploitation by jointly considering uncertainty, stiffness magnitude, and spatial gradients. Although viscoelastic modelling supports robust estimation, lesion search and characterisation in this work are driven solely by elasticity. Our framework is structured as a closed-loop search algorithm comprising four main components:

- (A.) a target distribution over tissue stiffness, which is updated at every iteration;
- (B.) an Expected Information Density (EID) map to prioritise sampling locations;
- (C.) a trajectory planner using ergodic control, guided by the EID;
- (D.) a data acquisition executed during trajectory execution, which enables updates in the target distribution.

These steps are executed cyclically, updating the target distribution (E.) until a predefined stopping criterion is met (F.). Each component is designed to be modular and independent, allowing flexible integration of alternative models or control strategies.

A. Target Distribution

We use GPR to construct a continuous map of the stiffness distribution across the palpated surface. The robot end-effector positions serve as the input domain, and the estimated elasticity values form the output. After each regression update, the mean function $\mu(\mathbf{x})$ estimates the stiffness at unobserved locations, while the variance function $\sigma^2(\mathbf{x})$ quantifies the model uncertainty. As expected with GPR, uncertainty increases with distance from the observed samples. The squared exponential kernel ensures smooth, differentiable estimates, which are critical for downstream gradient computation.

1) *Gaussian Process Regression*: Is a non-parametric Bayesian approach for estimating continuous functions from data, offering both predictions and uncertainty estimates. Unlike parametric models, which assume a fixed functional form, GPR models a distribution over functions that is updated as new observations are incorporated [20]. A Gaussian Process (GP) is a collection of random variables such that any finite subset follows a joint Gaussian distribution. Formally, it is written as:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')),$$

where $m(\mathbf{x})$ is the mean function (often assumed zero) and $k(\mathbf{x}, \mathbf{x}')$ is the covariance function, or kernel, which encodes the similarity between inputs. The kernel defines how closely related the outputs are at different input points. One of the most widely used kernels is the squared exponential (SE):

$$k_{\text{SE}}(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2l^2}\right),$$

where l is the length scale parameter controlling smoothness (in our experiments, $l = 2.5$ cm is equal to the radius of the indenter), and $\|\mathbf{x} - \mathbf{x}'\|$ is the Euclidean distance between inputs. This kernel implies smooth, infinitely differentiable functions, making it ideal for modelling physical quantities such as tissue stiffness.

2) *GPR Training*: Given a set of training observations \mathbf{X} with corresponding outputs \mathbf{y} , GPR computes a posterior predictive distribution at new input locations \mathbf{X}_* . This yields a

predicted mean function μ_* , representing the estimated stiffness at unobserved locations, and an associated covariance Σ_* , which quantifies model uncertainty. As expected, the uncertainty increases with distance from the observed samples, reflecting reduced confidence in sparsely explored regions.

B. Expected Information Density for Elasticity Maps

To prioritise regions for further palpation, we define an EID function $\xi_{\text{EID}}(\mathbf{x})$ that balances two objectives: exploration of uncertain areas, and exploitation of predicted high-stiffness regions and refinement near region boundaries. Formally, we define

$$\tilde{\xi}_{\text{EID}}(\mathbf{x}) = (1 - \alpha) \left(\frac{g(\mathbf{x})}{\int_{\mathbf{X}} g(\mathbf{x})} + \frac{\mu(\mathbf{x})}{\int_{\mathbf{X}} \mu(\mathbf{x})} \right) + \alpha \frac{\sigma(\mathbf{x})}{\int_{\mathbf{X}} \sigma(\mathbf{x})} \quad (1)$$

the un-normalised distribution, where $g(\mathbf{x}) = \|\nabla\mu(\mathbf{x})\|$ is the gradient norm of the predicted elasticity, highlighting boundaries. The distribution (1) is then normalised as:

$$\xi_{\text{EID}}(\mathbf{x}) = \frac{\tilde{\xi}_{\text{EID}}(\mathbf{x})}{\int_{\mathbf{X}} \tilde{\xi}_{\text{EID}}(\mathbf{x}) d\mathbf{x}}. \quad (2)$$

The trade-off between exploration and exploitation is controlled by the scalar $\alpha \in [0, 1]$. Lower values bias the system towards known high-stiffness or high-gradient areas, while higher values encourage sampling in regions of high uncertainty. In our implementation, α varies over time depending on the value of the ergodic metric in the following way:

$$\alpha(t) = \frac{\int_{\mathbf{X}} e(\mathbf{x}, t) d\mathbf{x}}{\int_{\mathbf{X}} e(\mathbf{x}, 0) d\mathbf{x}}, \quad (3)$$

where $e(\mathbf{x}, t)$ is the ergodic metric that will be defined in the following section in (4). According to (3), at the onset of the exploration, the trajectory is not ergodic and $\alpha(0) = 1$. As a consequence, the EID in (1) is dominated by the variance of the GPR, resulting in a purely exploratory behaviour. Regions that have not yet been sampled exhibit higher uncertainty and are therefore prioritised for exploration. As informative measurements are collected, the ergodic metric progressively decreases, leading to a gradual rebalancing between exploration and exploitation. The spatial coverage of the trajectory increasingly matches the target EID distribution, and the exploration naturally shifts towards exploitation of regions associated with higher estimated stiffness and pronounced spatial gradients, such as the boundaries of stiff inclusions. We want to highlight that the EID is computed on-the-fly and does not exploit prior knowledge about the inspected area.

C. Trajectory Planning

The normalised EID ξ_{EID} is used as the target distribution for the ergodic planner, which computes a trajectory that visits each region in proportion to its expected information content. The HEDAC algorithm [19] is used to generate the trajectory online.

1) *Ergodic Control*: Aims to guide a robot trajectory to match a target spatial distribution, balancing exploration and exploitation by spending more time in regions of higher interest. It models the coverage density using radial basis functions (RBFs). Given a trajectory $\mathbf{z} : [0, t] \rightarrow \Omega \subset \mathbb{R}^n$, the normalised

coverage density $c(\mathbf{x}, t)$ is:

$$c(\mathbf{x}, t) = \frac{\tilde{c}(\mathbf{x}, t)}{\int_{\Omega} \tilde{c}(\mathbf{x}, t) d\mathbf{x}}, \quad \tilde{c}(\mathbf{x}, t) = \frac{1}{t} \int_0^t \phi_{\sigma}(\mathbf{x} - \mathbf{z}(\tau)) d\tau,$$

where ϕ_{σ} is a Gaussian kernel. The coverage error with respect to a target distribution $m(\mathbf{x})$ is defined as the spatial diffusion of $m(\mathbf{x})$ through the RBF kernel compared with the coverage density

$$e(\mathbf{x}, t) = (\phi_{\sigma} * m)(\mathbf{x}) - c(\mathbf{x}, t), \quad (4)$$

To avoid local minima and improve spatial coordination, HEDAC smooths this error using a stationary heat equation:

$$\rho \Delta u(\mathbf{x}, t) = \beta u(\mathbf{x}, t) + \gamma a(\mathbf{x}, t) - s(\mathbf{x}, t), \quad \dot{\mathbf{z}} = v_a \cdot \frac{\nabla u}{\|\nabla u\|},$$

where s highlights under-sampled regions, and a ensures collision avoidance between the agents.

2) *Practical Implementation:* Different from [19], this work implemented the non-stationary diffusion equation, as was done in [26], for a more locally consistent exploration behaviour. This means that $\dot{u} \neq 0$ is defined as:

$$\dot{u}(\mathbf{x}, \tau) = \Delta u(\mathbf{x}, \tau), \quad (5)$$

allowing us to control the desired smoothness. In this formulation, an initial condition is required: $u(\mathbf{x}, 0) = s(\mathbf{x}, 0)$. During the experiments, a second-order agent is used to generate the trajectory. The trajectory is updated with a frequency ΔT chosen to ensure smooth motion and adequate control loop resolution (typically 100 Hz in our implementation, 5 times slower than the low-level controller). The planning phase respects constraints on maximum velocity to ensure physical plausibility during data collection imposed by the second-order dynamics inside the ergodic controller.

D. Trajectory Execution and Data Collection

While the trajectory is executed, tissue elasticity is estimated continuously using an Extended Kalman Filter (EKF). Data are sampled at a rate that depends on the robot motion and the indenter geometry. Specifically, higher velocities or smaller indenters require denser sampling to avoid aliasing effects in the stiffness map.

1) *Online Elasticity and Viscosity Estimation:* Accurately estimating the mechanical response of soft tissues in real time is a challenging task, as it requires modelling the contact interaction between the probe and the tissue surface and inferring tissue parameters from force measurements and robot motion. In this work, we adopt a viscoelastic contact model to improve the fidelity of force prediction during dynamic palpation, while subsequent abnormality detection focuses on the estimated elastic stiffness. To achieve this, we adopt DRM [27], which allows mapping the three-dimensional contact mechanics into an equivalent two-dimensional representation. Under this framework, the force exerted by the tissue during indentation is a nonlinear function of penetration depth and velocity:

$$F_{\text{TOT}}(d, \dot{d}) = \frac{4}{3} \frac{E_f}{1 - \nu^2} \sqrt{R} d \dot{d} + \frac{4}{1 - \nu} \eta \sqrt{R} \dot{d}, \quad (6)$$

where d is the penetration depth, \dot{d} its derivative (velocity), E_f is the Young's modulus of the material, η is the viscosity

coefficient, ν is the Poisson's ratio, and R is the radius of the spherical indenter [18]. The first term models elastic deformation according to Hertzian contact mechanics, while the second term accounts for viscoelastic damping.

2) *Filter Design:* The main challenge lies in the fact that d is not directly measurable during robotic interaction. Therefore, we use an EKF to estimate the system state, which includes both d and \dot{d} , as well as the unknown tissue parameters. The system is modelled using a discrete-time nonlinear state-space representation, as detailed in [18]. The state vector is defined as $[d, k, \lambda, \dot{d}, \dot{k}, \dot{\lambda}, \ddot{k}, \ddot{\lambda}]^T$, where κ and λ correspond to the local stiffness and damping parameters, respectively. The filter assumes as input the measured contact force and comprises the effective interaction mass. Process and measurement noise are also included to account for model uncertainties and sensor imperfections. This filtering approach allows real-time estimation of the viscoelastic parameters even during continuous contacts with indentation speeds up to 5 mm/s [28], being suitable for dynamic exploration tasks.

E. Target Distribution Update

Once a sufficient number of new samples are acquired, the GPR model is updated to refine the target stiffness distribution. Since GPR inference has cubic time complexity with respect to the number of samples, we avoid updating the model at every time step. Instead, updates are performed at a fixed frequency (e.g., 1Hz), allowing batch accumulation of new data while keeping computation time bounded.

F. Search Termination Criterion

Different from existing approaches that rely on predefined stopping conditions [15], or fixed time/step limits [16], we employ an information-theoretic stopping criterion on α , which is based on the ergodic metric. The value of α governs the trade-off between exploration and exploitation. Larger values of α correspond to an early exploration phase and favour earlier termination, but may lead to incomplete spatial coverage. Conversely, smaller values of α indicate that the exploration has largely converged toward the target EID distribution, resulting in extended exploitation phases, longer execution times, and only marginal refinement of already identified regions. In practice, the stopping threshold on α is selected within the exploitation-dominated regime, such that further reductions correspond to diminishing returns in reconstruction accuracy. This choice is supported by empirical observations indicating saturation of the RMSE as α approaches a plateau, and reflects a practical balance between coverage completeness, estimation accuracy, and exploration duration. The advantage of using this metric lies in its ability to quantify how thoroughly the area has been explored, without relying on prior knowledge such as fixed time limits or the number and characteristics of stiff regions that need to be found. This criterion is preferable to a fixed-time limit, as it adapts the exploration duration to the complexity of the stiffness distribution and the desired level of map details, avoiding both premature termination and unnecessary probing.




Fig. 2. Synthetic elasticity distributions for simulated scenarios, from 1 to 3 stiffer regions. The colour scale refers to the elastic modulus (in kPa).

IV. SIMULATION RESULTS

A. Setup and Baselines

The proposed search algorithm was evaluated on synthetic stiffness distributions representative of biological soft tissues. They were generated as mixtures of Gaussian components with varying locations, covariances, and amplitudes, modelling localised stiff inclusions (over 100 kPa in pathological areas) in a softer background (few kilopascal in healthy regions) [29], similarly to related research [15], [16], [17]. Three scenarios in a square domain (with side 45 mm) were considered:

- 1) a single stiff region,
- 2) two stiff regions, and
- 3) three distinct stiff areas,

whose distributions are illustrated in Fig. 2. These regions were first densely sampled with a grid-based approach to generate the ground truth, used to compute the estimation and segmentation errors. The simulation environment was implemented in Python, with GPR and BO components built using the PyTorch framework, thus ensuring efficient GPU-based inference and optimisation during online execution. All simulations and experiments were executed on a mini PC equipped with an Intel i7-13700HX CPU (24 cores), 64 GB RAM, and an NVIDIA RTX 4070 GPU with 8 GB of memory. The proposed framework runs in real time, with HEDAC trajectory updates and EKF-based viscoelastic estimation both requiring less than 1 ms per control loop iteration at 100 Hz. GPR updates are performed asynchronously at a lower frequency (1 Hz), with a maximum observed update time of 105 ms. Elasticity samples were collected at a frequency of 4Hz and a maximum velocity of 1cm/s. In this configuration, the maximum distance between two consecutive sampled points equals half of the indenter radius, which is set here to 2.5mm. These simulations were conducted to evaluate the performance of the proposed ergodic search strategy (Section IV-B) across the three distribution scenarios.

We compared ergodic search with two variants of Bayesian Optimisation (BO), which is the standard baseline for robotic palpation (Section IV-C). BO models an expensive-to-evaluate objective function $f(\mathbf{x})$ using a Gaussian Process (GP) and selects sampling locations by maximising an acquisition function that balances exploration and exploitation. Given data $\mathcal{D} = (\mathbf{x}_i, y_i)_{i=1}^n$, the GP provides a posterior predictive distribution $f(\mathbf{x}) \sim \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x}))$. As an acquisition function, the Expected Improvement (EI) [15], [16] is employed, $a_{EI}(\mathbf{x}) = (f_{\min} - \mu(\mathbf{x}))\Phi(z) + \sigma(\mathbf{x})\phi(z)$, where f_{\min} is the minimum observed value and $z = \frac{f_{\min} - \mu(\mathbf{x})}{\sigma(\mathbf{x})}$. This variant is denoted BO-EI. However, the standard point-to-point implementation of BO has limitations for robotic palpation. Due to safety constraints, especially in medical settings, the robot must follow smooth

TABLE I
PERFORMANCE METRICS ACROSS 50 RUNS WITH NOISE-FREE ELASTICITY MEASUREMENTS (w/o) AND WITH WHITE NOISE $\mathcal{N}(0, 6.25 \text{ kPa}^2)$ (w/). NR IS THE NUMBER OF STIFF REGIONS, DR THE DETECTION RATE, L THE LENGTH OF THE TRAJECTORY, AND RMSE IS THE ROOT MEAN SQUARE ERROR.

NR	DR (%) \uparrow	l [mm]	Time [s] \downarrow	RMSE [kPa] \downarrow
w/o Measurement Noise				
1	100%	664	70	0.37
2	100%	660	69	0.97
3	99%	665	70	1.96
w/ Measurement Noise				
1	100%	633	66	1.57
2	99%	630	65	2.17
3	97%	580	60	3.14

and predictable trajectories and cannot move rapidly between distant sampling points as typically generated by the BO. Moreover, each probing action can take up to 5s due to indentation and viscoelastic estimation time. To ensure a fair comparison with the proposed ergodic planner, we adapted BO to perform continuous palpation, which characterises our ergodic search. This prevents BO from penalising non-sampled regions during motion. Specifically, additional samples are collected every 2.5 mm (half the indenter radius) along the path between BO-selected targets, enabling BO to exploit intermediate data, similarly to the ergodic approach (variant termed BO-EIS).

To assess the spatial accuracy of the reconstructed stiffness maps, we applied a segmentation pipeline to identify distinct stiff regions. Based on the elasticity estimates in Section IV-B and IV-C, KMeans clustering (SciPy) was used to separate stiffer regions from the surrounding soft tissue. Unlike fixed thresholding approaches [12], this method does not require manual parameter tuning and adapts to the distribution of estimated stiffness values. The number of detected regions was obtained by counting connected components in the stiff cluster. In Section IV-D, the segmented regions are converted into polygonal boundaries using Shapely and compared against ground truth using standard segmentation metrics.

B. Robustness of Ergodic Exploration

In this section, the performance of the ergodic controller is tested with different numbers of stiffer regions. To assess the robustness of the algorithm, for each map, the simulation is executed 50 times with a random starting position, while three performance metrics were considered:

- 1) the detection rate (DR), i.e., the percentage of trials in which the algorithm correctly estimates the number and position of the stiff regions present in the ground truth (higher is better);
- 2) the elapsed time (lower is better);
- 3) the root mean square error (RMSE, lower is better).

For completeness, we report also the trajectory length l . Each exploration terminates when the ergodic metric α reaches 0.4 (set experimentally). Two experimental conditions were examined: one involved noise-free elasticity measurements and another incorporated noisy measurements $\mathcal{N}(0, 6.25 \text{ kPa}^2)$, where the variance equals the noise of the elasticity estimation in the worst case. The results of the study are reported in Table I. Overall, the algorithm exhibited comparable performance in

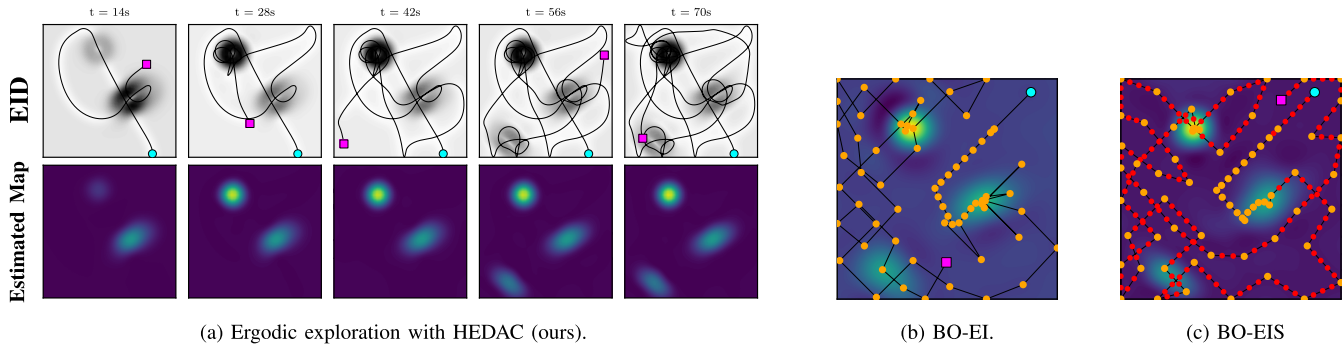


Fig. 3. End-effector tip trajectory during the exploration of an elasticity distribution with three stiff regions: comparison of different planners. The cyan circle indicates the beginning of the trajectory, while the magenta square denotes the end of the trajectory. In plots (b) and (c), the orange markers indicate the locations selected by BO, whereas the red markers represent the sampled points collected along the executed trajectory.

TABLE II

COMPARISON OF PERFORMANCE METRICS ACROSS 100 RUNS OF ERGODIC SEARCH AND BO, WITH (w/) AND WITHOUT (w/o) NOISE. DR IS THE DETECTION RATE, T THE EXPLORATION TIME, AND RMSE THE ROOT MEAN SQUARE ERROR.

Traj. Length Method	500 mm			650 mm			800mm		
	DR (%) \uparrow	t [s] \downarrow	RMSE [kPa] \downarrow	DR (%) \uparrow	t [s] \downarrow	RMSE [kPa] \downarrow	DR (%) \uparrow	t [s] \downarrow	RMSE [kPa] \downarrow
BO-EI (w/o)	39	59	12.09	73	88	7.98	92	124	5.47
BO-EIS (w/o)	93	60	4.61	99	92	2.81	99	124	1.87
Ergodic (w/o)	78	41	4.69	92	62	2.74	99	84	1.42
BO-EI (w/)	38	60	12.1	60	89	7.84	91	118	5.49
BO-EIS (w/)	85	63	5.54	95	95	3.35	99	130	2.59
Ergodic (w/)	79	41	5.87	98	62	3.03	99	84	2.16

terms of DR and execution time across the three scenarios, regardless of the presence of measurement uncertainty. These results demonstrate the robustness of the method with respect to both the number of stiff regions to be detected and the influence of measurement noise. The RMSE, instead, grows under both experimental conditions as the number of stiff regions increases. This behaviour may be attributed to uncertainty at the boundaries of the stiff regions, which are typically under-sampled, and will be examined in greater detail in Section IV-D. As expected, we observe an increased RMSE in the second experimental condition, attributable to the presence of the noise.

Fig. 3(a) shows an example of the evolution of the agent trajectory over time, when exploring the map with 3 stiff regions. All stiffer regions are successfully detected within 40s. Over time, the ergodic metric α decreases from 0.79 at $t = 14$ s to 0.38 at $t = 70$ s, reflecting improved coverage of the target distribution. In parallel, the RMSE shows a substantial reduction, dropping from 13.75 kPa to 2.85 kPa. Intermediate values reinforce this trend: at $t = 28$ s, $\alpha = 0.54$ and RMSE = 11.19 kPa; at $t = 42$ s, $\alpha = 0.51$ and RMSE = 3.7 kPa; and at $t = 56$ s, $\alpha = 0.43$ and RMSE = 2.56 kPa. This consistent evolution confirms that as the exploration becomes more ergodic, the accuracy of the reconstructed stiffness map improves accordingly.

C. Comparison With BO Baselines

In this section, we focus on the comparison of the results of our ergodic approach with the two variants of BO described in Section IV-A. We consider the distribution with 3 stiff regions, which has proven to be the most challenging for the ergodic exploration (see Fig. 3). The simulation time is the

sum of the simulated agent motion with a velocity of 1cm/s and the computation time required to determine the next sampling location. Each simulation is conducted under two conditions: one assuming noise-free elasticity measurements, and another incorporating additive white noise $\mathcal{N}(0, 6.25 \text{ kPa}^2)$, similar to Section IV-B. Fixed length of 400, 600, and 800 mm, were used to evaluate the performance evolution over time, ensuring a fair comparison independently of the method-specific stopping criteria. The average results of 100 simulations are listed in Table II.

This comparison highlights that BO-EI consistently underperforms compared to the other two methods, yielding the lowest DR and the highest RMSE across all tested conditions. Both BO-EIS and the ergodic strategy achieve comparable detection rates once sufficient exploration has been performed, indicating successful identification of the number and location of stiff regions. However, the ergodic approach consistently provides lower reconstruction errors, resulting in more accurate stiffness maps. At shorter travelled distances (500 mm–650 mm), BO-EIS achieves higher detection rates, reflecting faster initial identification of stiff regions. This behaviour is expected given the exploitative nature of BO-based strategies, which prioritise sampling high-uncertainty locations. However, in clinical settings, once suspicious regions are identified, accurate post-detection refinement becomes critical. Therefore, it is more relevant to compare the approaches once detection performance is maximised (800 mm, DR \approx 99–100%), and differences between the methods are primarily reflected in reconstruction accuracy and execution duration. In this regime, the ergodic strategy achieves lower RMSE while requiring substantially less execution time than BO-EIS under the same travelled-distance budget, both in noise-free and noisy conditions. These results indicate that, in




Fig. 4. Clustering and boundary extraction results after the exploration phase. The original shapes (in blue) are shown alongside the estimated ones (in red).

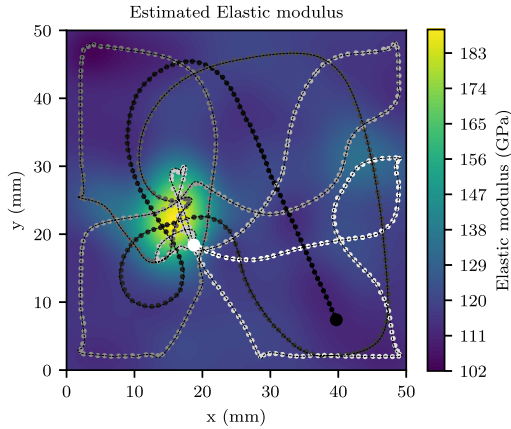


Fig. 5. Estimated stiffness distribution obtained through palpation of the soft phantom. The stiffer region is shown in yellow. The end-effector trajectory during the search motion is visualised using a grayscale gradient, from black (beginning) to white (end). A video of the experiment is provided in the Supplemental Materials and is also available online at <https://youtu.be/a3BUSx7wXY>.

the post-detection phase, ergodic exploration enables more efficient refinement of stiffness maps, yielding improved accuracy per unit motion and reduced execution time under equal motion cost. Compared to BO-based approaches, the ergodic strategy also exhibits greater robustness to noise, maintaining a consistent performance advantage across all evaluated conditions. Overall, these findings confirm the benefit of ergodic exploration in balancing detection reliability, reconstruction accuracy, and execution efficiency once suspicious regions have been identified.

D. Segmentation

The spatial accuracy of the proposed method was evaluated based on its ability to segment stiff regions within the reconstructed stiffness maps and compared against the BO. To this end, we designed two simulated stiffness distributions: the first one with two inclusions, one irregularly shaped (Cluster 1) and one circular (Cluster 2), while the second distribution consisted of a single donut-shaped inclusion (Cluster 3), as shown in Fig. 4. A total of 50 simulations were conducted under identical conditions, incorporating additive white noise $\mathcal{N}(0, 6.25kPa^2)$. Each run was terminated after 70s to ensure consistent comparison across trials. For each run, the estimated stiffness map was segmented using the clustering and boundary extraction pipeline described earlier. The resulting segmented areas were compared with ground truth using two standard metrics [15], sensitivity and specificity, defined as:

- 1) Sensitivity = $TP/(TP+FN)$,
- 2) Specificity = $TN/(TN+FP)$,

TABLE III
AVERAGE SEGMENTATION RESULTS ACROSS THREE CLUSTERS

	Cluster 1			Cluster 2			Cluster 3		
	BO-EI	BO-EIS	Erg.	BO-EI	BO-EIS	Erg.	BO-EI	BO-EIS	Erg.
Sens.	0.561	0.950	0.975	0.657	0.997	0.971	0.746	0.904	0.978
Spec.	0.654	0.993	0.993	0.655	0.998	0.997	0.829	0.953	0.992

where TP, FP, FN, and TN denote the number of pixels of true positives, false positives, false negatives, and true negatives, respectively. The metrics were computed separately for each of the three ground-truth regions, and the final values were averaged across the simulations. The results of the segmentation algorithm (Table III) demonstrate elevated and consistent specificity across all three clusters for the two methods that sample continuously, thereby confirming the model effectiveness in correctly excluding background regions. Conversely, the ergodic search exhibited superior sensitivity compared to both the BO with and without sampling when clustering non-uniform shapes (Clusters 1-3). This finding indicates that the ergodic search is more adept at accurately delineating regions of interest. This is especially relevant in medical contexts, where false negatives can lead to missed detections of pathological regions. Consequently, achieving high sensitivity is imperative to ensure clinically reliable performance.

V. EXPERIMENTAL RESULTS

The validity of the simulation results in Section IV was assessed in a realistic setting by testing the proposed method on a manipulator interacting with a physical sample (Fig. 1). An UR3e, position-controlled at 500 Hz, was used to execute the trajectories generated by the Ergodic Control. The force at the end-effector was measured by a 6-axis F/T sensor, a Bota SensOne, sampling at 1 kHz. A 3D-printed indenter with a spherical tip with 5 mm radius is attached to the end-effector. The communication between the robot, the F/T sensor, and the algorithms was implemented with ROS2 Humble. To simulate a biological tissue with a mass, we fabricated a silicone sample consisting of a soft matrix made from Ecoflex-0030 and a stiffer spherical inclusion (radius 1cm) made from DragonSkin-30NV. This configuration increases local stiffness by up to 90%, consistent with values reported in previous studies [15], [29]. The ground truth is built with point-wise sampling on a grid with a palpation every 2.5 mm in a square of 5 cm \times 5 cm.

We conducted 5 experiments on the physical setup, each starting from a different random initial position. In all cases, the algorithm successfully identified the location of the spherical intrusion, demonstrating consistent target localisation. The reconstruction accuracy (RMSE = 9.76 ± 2.97 kPa) shows good agreement with the reference stiffness distribution obtained from grid-based probing used as ground truth. The method also achieved good classification performance, with an average sensitivity of 0.92, and specificity of 0.98. Compared to simulation results, the experimental metrics confirm the high sensitivity and specificity, which are particularly important for avoiding false negatives in medical applications. Despite the experimental RMSE values being marginally higher than those observed in simulation, the measurement uncertainty is still below the 5% (average absolute error 8.00 ± 2.7 kPa) and the clustering performance appears largely unaffected. This difference can be attributed to several factors, including the non-ideal behaviour of real materials, noise in force measurements, slight

deviations from the assumed contact model, and uncertainties in ground-truth stiffness calibration. Despite these challenges, the observed trends are consistent with the simulation, confirming the robustness of the method in real-world conditions. As shown in Fig. 5, the trajectory of the end-effector adapts to the underlying stiffness distribution: it concentrates more exploration in the stiffer region on the left while still covering the rest of the map.

VI. CONCLUSION

This work presents an autonomous robotic palpation framework for real-time elastic stiffness mapping of soft tissues. By combining force-based viscoelastic parameter estimation with ergodic trajectory planning implemented via a heat-equation-driven controller, the method enables continuous palpation without embedded tactile sensing. Exploration is guided by an EID tailored to stiffness mapping, allowing adaptive focus on diagnostically relevant regions. Simulation results demonstrate improved stiffness reconstruction accuracy, with reliable detection of multiple stiff regions and lower RMSE and segmentation errors compared to BO, particularly under noisy conditions. Experimental validation with a robotic arm and a silicone phantom confirms the robustness of the approach, showing consistent detection of stiff inclusions and good agreement with ground-truth stiffness maps. Overall, the proposed ergodic strategy effectively balances exploration and accurate estimation, making it well suited for autonomous palpation as an assistive diagnostic tool. Future work will extend the method to three-dimensional mapping on non-flat surfaces and evaluate it on anatomically realistic phantoms and ex vivo tissues.

REFERENCES

- [1] R. Tozzi et al., "Laparoscopic treatment of early ovarian cancer: Surgical and survival outcomes," *Gynecologic Oncol.*, vol. 93, no. 1, pp. 199–203, 2004.
- [2] N. Eisemann et al., "Nationwide real-world implementation of AI for cancer detection in population-based mammography screening," *Nature Med.*, vol. 31, no. 3, pp. 917–924, 2025.
- [3] Q. Dan et al., "Diagnostic performance of deep learning in ultrasound diagnosis of breast cancer: A systematic review," *npj Precis. Oncol.*, vol. 8, no. 1, pp. 1–13, 2024.
- [4] A. Attanasio et al., "Autonomy in surgical robotics," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 4, pp. 651–679, 2021.
- [5] L. Zhang et al., "Innovative integration of automated breast volume scan and ultrasound elastography for enhanced differentiation of benign and malignant breast lesions," *Sci. Rep.*, vol. 15, no. 1, pp. 1–8, 2025.
- [6] I. Sack, "Magnetic resonance elastography from fundamental soft-tissue mechanics to diagnostic imaging," *Nature Rev. Phys.*, vol. 5, no. 1, pp. 25–42, 2023.
- [7] N. Pacheco et al., "Virtual palpation: A novel method to identify the physical properties of tissue for laser surgery," in *Proc. Hamlyn Symp. Med. Robot.*, 2025, pp. 163–164.
- [8] A. L. Trejos et al., "Robot-assisted tactile sensing for minimally invasive tumor localization," *Int. J. Robot. Res.*, vol. 28, no. 9, pp. 1118–1133, 2009.
- [9] N. Enayati, E. De Momi, and G. Ferrigno, "Haptics in robot-assisted surgery: Challenges and benefits," *IEEE Rev. Biomed. Eng.*, vol. 9, pp. 49–65, 2016.
- [10] S. McKinley et al., "An interchangeable surgical instrument system with application to supervised automation of multilateral tumor resection," in *Proc. IEEE Int. Conf. Automat. Sci. Eng.*, 2016, pp. 821–826.
- [11] N. Zevallos et al., "A real-time augmented reality surgical system for overlaying stiffness information," in *Proc. Robot. Sci. Syst.*, 2018, pp. 1–10.
- [12] A. Garg et al., "Tumor localization using automated palpation with gaussian process adaptive sampling," in *Proc. IEEE Int. Conf. Automat. Sci. Eng.*, 2016, pp. 194–200.
- [13] E. Ayvali, R. A. Srivatsan, L. Wang, R. Roy, N. Simaan, and H. Choset, "Using Bayesian optimization to guide probing of a flexible environment for simultaneous registration and stiffness mapping," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 931–936.
- [14] H. Salman et al., "Trajectory-optimized sensing for active search of tissue abnormalities in robotic surgery," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 5356–5363.
- [15] Y. Yan and J. Pan, "Fast localization and segmentation of tissue abnormalities by autonomous robotic palpation," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 1707–1714, Apr. 2021.
- [16] E. Ayvali, A. Ansari, L. Wang, N. Simaan, and H. Choset, "Utility-guided palpation for locating tissue abnormalities," *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 864–871, Apr. 2017.
- [17] K. A. Nichols and A. M. Okamura, "Methods to segment hard inclusions in soft tissue during autonomous robotic palpation," *IEEE Trans. Robot.*, vol. 31, no. 2, pp. 344–354, Apr. 2015.
- [18] L. Beber, E. Lamon, G. Moretti, D. Fontanelli, M. Saveriano, and L. Palopoli, "Towards robotised palpation for cancer detection through online tissue viscoelastic characterisation with a collaborative robotic arm," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2024, pp. 2380–2386.
- [19] S. Ivic, B. Crnkovic, and I. Mezić, "Ergodicity-based cooperative multi-agent area coverage via a potential field," *IEEE Trans. Cybern.*, vol. 47, no. 8, pp. 1983–1993, Aug. 2017.
- [20] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [21] P. Chalasani, L. Wang, R. Roy, N. Simaan, R. H. Taylor, and M. Kobilarov, "Concurrent nonparametric estimation of organ geometry and tissue stiffness using continuous adaptive palpation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 4164–4171.
- [22] P. Chalasani, L. Wang, R. Roy, N. Simaan, and R. H. Taylor, "Preliminary evaluation of an online estimation method for organ geometry and tissue stiffness," *IEEE Robot. Autom. Lett.*, vol. 3, no. 3, pp. 1816–1823, Jul. 2018.
- [23] G. Mathew and I. Mezi, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D, Nonlinear Phenomena*, vol. 240, no. 4–5, pp. 432–442, 2011.
- [24] L. M. Miller, Y. Silverman, M. A. MacIver, and T. D. Murphey, "Ergodic exploration of distributed information," *IEEE Trans. Robot.*, vol. 32, no. 1, pp. 36–52, Feb. 2016.
- [25] L. M. Miller and T. D. Murphey, "Trajectory optimization for continuous ergodic exploration," in *Proc. Amer. Control Conf.*, 2013, pp. 4196–4201.
- [26] C. Bilaloglu, T. Löw, and S. Calinon, "Whole-body ergodic exploration with a manipulator using diffusion," *IEEE Robot. Autom. Lett.*, vol. 8, no. 12, pp. 8581–8587, Dec. 2023.
- [27] V. L. Popov and M. Heß, *Method of Dimensionality Reduction in Contact Mechanics and Friction*. Berlin, Germany: Springer, 2015.
- [28] L. Beber et al., "Force-based viscosity and elasticity measurements for material biomechanical characterisation with a collaborative robotic arm," *IEEE Trans. Instrum. Meas.*, vol. 74, 2025, Art. no. 4013314.
- [29] P. N. Wells and H.-D. Liang, "Medical ultrasound: Imaging of soft tissue strain and elasticity," *J. Roy. Soc. Interface*, vol. 8, no. 64, pp. 1521–1549, 2011.