

Ergodic Exploration using Binary Sensing for Non-Parametric Shape Estimation

Ian Abraham¹, Ahalya Prabhakar¹, Mitra J.Z. Hartmann^{1,2}, and Todd D. Murphey¹

Abstract—Current methods to estimate object shape—using either vision or touch—generally depend on high-resolution sensing. Here, we exploit ergodic exploration to demonstrate successful shape estimation when using a low-resolution binary contact sensor. The measurement model is posed as a collision-based tactile measurement, and classification methods are used to discriminate between shape boundary regions in the search space. Posterior likelihood estimates of the measurement model help the system actively seek out regions where the binary sensor is most likely to return informative measurements. Results show successful shape estimation of various objects as well as the ability to identify multiple objects in an environment. Interestingly, it is shown that ergodic exploration utilizes non-contact motion to gather significant information about shape. The algorithm is extended in three dimensions in simulation and we present two dimensional experimental results using the Rethink Baxter robot.

Index Terms—Sensor-Based Control; Force and Tactile Sensing

I. INTRODUCTION

TACTILE sensing is often associated with shape estimation [1]–[4] and mapping problems [5] in conditions where visual sensing may be limited. In some instances, tactile sensing is used to supplement vision-based sensing to improve shape estimates [6], [7]. The richness of touch as a sensing modality is underscored by the development of novel tactile sensors [8]–[11] for use in a myriad of applications ranging from robot-assisted tumor detection [7], to texture recognition, and feature localization [12], [13]. These advances in tactile sensor technology require corresponding advances in active exploration algorithms and the interpretation of tactile-based sensor data.

Recent approaches for active exploration use random sampling-based search algorithms for sensor motion planning that require a separate controller for path following [14]. Other methods use task-specific probabilistic spatial methods for shape estimation that are updated based on sensor poses that

minimize measurement uncertainty [2], [15]–[18]. A common approach in the literature is motion planning for sensing followed by feedback regulation of the generated plan [17]–[20]. Moreover, most methods focus on one object at a time. In contrast, the presented work integrates planning and control into a feedback law. As a result, the method uses sensor motion to actively sense for time-varying spatial information. We take the framework in [21] and use the feedback law developed here to enable real-time execution during shape estimation. Lastly, the feedback in the active sensing algorithm compensates for low resolution sensors and the specification of the algorithm is independent of the number of objects.

We show that ergodic exploration with Sequential Action Control (eSAC) can be used for active exploration with respect to time-varying tactile-based distributed information [22], [23]. Previous applications of ergodic theory utilize parametric measurement models for localization tasks [21], [24]. The current work demonstrates localization and estimation of non-parametric shapes using a binary sensor model with classification methods. In contrast to other methods of active sensing for shape estimation [16], [18], [19], [25], [26], the proposed algorithm automatically encodes dynamical constraints without any overhead spatial discretization or motion planning. In addition, the algorithm incorporates sensor measurement information to actively adjust shape estimates and synthesize tactile-information based control actions. As a result, the algorithm automatically adjusts the control synthesis for multiple objects in an environment. Notably, ergodic exploration uses non-contact motion data (sensor motion not in contact with an object) [27], [28] to improve the shape estimate. The idea of utilizing free space is often found in other related works of pose estimation and tracking [27]–[30] and is emphasized in our work. As a final contribution, we show the algorithm is modular with respect to the choice of shape representation and tactile information distribution.

The paper outline is as follows: Section II motivates the use of a binary contact sensor and outlines the implications for tactile sensing. Section III describes the robot control algorithm for active exploration and motivates ergodicity as a measure for exploration. Section IV explains how shape estimation is achieved and how measurements update the shape estimate and the control policy. Experimental and simulated results and conclusion are shown in sections V and VI respectively.

II. TACTILE INFORMATION

In biological systems, tactile sensing is generally an active process that incorporates feedback from multiple environmental cues [31], [32]. Humans, for instance, use fingers to grasp

Manuscript received: September, 10, 2016; Revised December, 10, 2016; Accepted December, 26, 2016.

This paper was recommended for publication by Editor Francois Chaumette upon evaluation of the Associate Editor and Reviewers' Comments. This work was supported by the National Science Foundation award IIS-1426961 to TDM and National Institute of Health award R01-NS093585 to MJZH and TDM.

Authors are with the Neuroscience and Robotics Laboratory (NxR) at the Department of Mechanical Engineering¹ and Biomedical Engineering², Northwestern University, 2145 Sheridan Road Evanston, IL 60208 USA.

Email: ianabraham2015@u.northwestern.edu,
AhalyaPrabhakar2013@u.northwestern.edu,
m-hartmann@northwestern.edu,
t-murphey@northwestern.edu

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

objects for manipulation and feature detection. Many rodents use hair-like appendages (whiskers) for tactile sensing [33]–[35]. These biological sensors have a large set of actuators and sensor channels, thus the process control for active sensing is complex. We investigate a lower resolution version of tactile sensing that is simple to control for active sensing. Specifically, we show that a binary form of tactile sensing (i.e., collision detection) [8], [10], has enough information for shape estimation, when combined with an active exploration algorithm that automatically takes into account regions of shape information.

A binary tactile measurement model consists of a transition state from “no collision” to “collision” and vice versa. We denote this measurement model as

$$\Upsilon(x) = \begin{cases} 1, & \phi(x) \leq 0 \\ 0, & \phi(x) > 0 \end{cases} \quad (1)$$

where x is the sensor state and $\phi(x)$ is a boundary function that determines a transition state if $\phi(x) \leq 0$ (output of 1). The goal of shape estimation is to determine $\phi(x)$.

III. ACTIVE EXPLORATION USING ERGODIC CONTROL

A. Motivation

Typical algorithms used in active exploration and information maximization cast the problem of information acquisition as “exploratory” (wide spread search for diffuse information such as localization) or “exploitative” (direct search for highly structured information such as shape contours) [36], [37]. Ergodic exploration [21]–[24], [38] is responsive to both diffuse information densities and highly focused information. Thus, ergodic control seamlessly encodes wide spread coverage for diffuse information and localized search for focused information for both needs.

B. Ergodic Metric

A trajectory $x(t)$ of a sensor is ergodic with respect to a probability (“target”) distribution $\Phi(x)$ when the fraction of time the trajectory spends in an area within the search space is equal to the spatial statistics across the search space [22], [39]. The ergodic metric that measures this characteristic is given as [22], [23],

$$\mathcal{E}(x(t)) = \sum_{k \in \mathbb{Z}^n} \Lambda_k [c_k(x(t)) - \phi_k]^2 \quad (2)$$

where

$$c_k(x(t)) = \frac{1}{t_f - t_0} \int_{t_0}^{t_f} F_k(x(\tau)) d\tau, \quad (3)$$

$$\phi_k = \int_{\mathcal{X}} \Phi(x) F_k(x) dx. \quad (4)$$

Here, $F_k(x)$ is a Fourier basis function, and Λ_k are weights described in [22]. The target distribution $\Phi(x)$ is what drives the control synthesis for active-exploration.

Comparison between robot trajectory and distribution is done with equation (2) which takes the Fourier power series decomposition using (3) and (4) and directly compares the statistics of a trajectory with that of the spatial statistics

$\Phi(x)$. As the ergodic metric is convex [22], convex objectives subject to linear affine constraints are still convex. As a consequence, for linear affine dynamical systems, there are no local minimizers and must eventually search the whole space. A path is “ergodic” with respect to a distribution if the fraction of time the trajectory spends in a region in space is equivalent to the measure associated with that region.

C. Ergodic Sequential Action Control

Ergodic control determines a discrete set of control inputs that aims to reduce the ergodic metric over time. Ergodic Sequential Action Control (eSAC) is a control scheme that generates ergodic control actions. Although trajectory optimization has been used to generate control synthesis for ergodic exploration [21], [23], sequential action control provides a closed-form control that can immediately respond to changes in the ergodic metric in an infinite horizon setting [40]. In comparison with sample-based planners, SAC incorporates dynamical constraints. In particular, non-linear feedback control is desired when dealing with collisions and contact information that must be incorporated at the time of measurement acquisition. Thus, the motivating factor for using eSAC is a combination of compactness in its formulation and the ability to quickly react to measurements.

Assuming control-affine dynamics of the form

$$\dot{x} = f(x(t), u) = g(x) + h(x)u \quad (5)$$

where $x \in \mathbb{R}^n$ and $u \in \mathbb{R}^m$, the trajectory cost is

$$J(x(t)) = q\mathcal{E}(x(t)) + \int_{t_0}^{t_f} \frac{1}{2} u(\tau)^T R u(\tau) d\tau \quad (6)$$

where $q \in \mathbb{R}$ weights the ergodic metric and $R \in \mathbb{R}^{m \times m}$ is a positive definite matrix. During each sampling instance, eSAC computes the optimal action to reduce the cost function $J(x(t))$ in the following steps:

1) *Prediction*: eSAC forward simulates the system dynamics starting from the sampled sensor state x_0 at time t_0 to a finite time horizon $t_0 + T$ assuming a default control u_0 similar to NMPC methods [41], [42]. Cost sensitivity is computed by backwards simulating an adjoint variable, $\rho \in \mathbb{R}^n$ that satisfies the following differential equation

$$\begin{aligned} \dot{\rho} &= -\frac{q}{T} \sum_{k \in \mathbb{Z}^n} \Lambda_k (c_k(x(t)) - \phi_k) \frac{\partial}{\partial x} F_k(x(t)) \\ &\quad - \frac{\partial}{\partial x} f(x, u_0)^T \rho \\ &\text{subject to } \rho(t_0 + T) = \vec{0}. \end{aligned} \quad (7)$$

The pair $x(t)$ and $\rho(t)$ are used to compute the optimal control action.

2) *Calculate Optimal Control Actions*: eSAC produces a schedule of control actions $u^*(t) \in [t_0, t_0 + T]$ that optimizes

$$J_u = \frac{1}{2} \int_{t_0}^{t_0 + T} \left[\frac{dJ_{\text{track}}}{d\lambda} - \alpha_d \right]^2 + \|u(t)\|_R^2 dt \quad (8)$$

$$\text{where } \frac{dJ_{\text{track}}}{d\lambda} = \rho(t)^T (f(x(t), u) - f(x(t), u_0)). \quad (9)$$




Fig. 1. (A) Simulated time sequence of tactile-based eSAC estimating a square (green) in $\mathcal{X} \subset \mathbb{R}^2 \subset [0, 1] \times [0, 1]$. Robot trajectory (magenta) translates in time from the previous final state (shown as a circle) to the current initial condition (shown as a magenta square). The boundary of the surface estimate is shown as a black line and the posterior likelihood is shown in the density plot. High likelihood of acquiring a collision is shown in red and low likelihood is shown in blue. **Bottom Row:** Γ metric is defined as $\Gamma = \int_{\mathcal{X}} (\phi_{\text{emp}}(x) - \phi_{\text{actual}}(x))^2 dx$ the integrated difference of the shape and the estimated shape squared. The Γ measure drops quickly after first contact and then remains at an equilibrium as the shape estimate is updated. (B) Note the posterior at the end of the simulation (top subplot) resembles the Expected Information Density (EID) (see [21] for derivation of EID) of the parametric model of the square shape (bottom subplot).

The term $\frac{dJ_{\text{track}}}{d\lambda}$ computes the rate of change of the cost with respect to a switch of infinitesimal duration λ [40], [43]. The value $\alpha_d \in \mathbb{R}^-$ dictates the aggressiveness of the control (more negative values tend to saturate the control). In this work a value for $\alpha_d = -555$ was used. The closed form solution of u^* is then given as

$$u^* = (\Lambda + R^T)^{-1}[\gamma u_0 + h(x)^T \rho \alpha_d] \quad (10)$$

where $\gamma \triangleq h(x)^T \rho \rho^T h(x)$. Additional information on the derivation of SAC can be found in [40]. In this work, we used values of $q = 30$, $R = \text{diag}([0.01, 0.01])$, $T = 0.8$ for eSAC.

Algorithm 1 eSAC

- 1: **given** $x_0, c_{k,0}, \phi_{k,0}, t_{\text{curr}}, T, t_s, i = 0$
- 2: $(x(t), \rho(t)) \leftarrow \text{prediction}(x_0, t_{\text{curr}}, T)$
- 3: $u^*(t) \leftarrow \text{calcControl}(x(t), \rho(t))$
- 4: **return** $u^*(t) \in [t_{\text{curr}}, t_{\text{curr}} + t_s]$

IV. SHAPE ESTIMATION

The process for shape estimation is described in Fig. 2. Shape estimates are obtained through repeated samples of the search space using directed motion of the sensor, as calculated by eSAC. The samples are processed using kernel basis functions in order to create decision boundaries. The decision boundaries are then processed into a posterior likelihood estimate using a common method in machine learning to generate statistics on a classification fit known as Platt Scaling [44]–[47]. The returned spatial statistics on the fit is used for active sampling of the search space, thus updating the shape estimates. The process iterates at a user-specified sampling frequency. We further explain the process in the following sections.




Fig. 2. Block diagram for shape estimation. The measurements are processed by a Gaussian Kernel and made into a shape estimate. With Platt Scaling, the shape estimate is transformed into a posterior estimate of the likelihood of acquiring a specific binary sensor measurement, which then is converted into control using eSAC.

A. Binary Measurements for Shape Estimation

Given a set of measurements $y_k \in [0, 1]$ at time indexed by k sampled at the corresponding set of sensor state x_k , shape estimation is accomplished by probabilistic classification methods [16], [18], [25], [26]. Using the set of measurements, a decision boundary for the object is approximated as

$$\phi(x) \approx \sum_k \alpha_k y_k K(x_k, x) + b \quad (11)$$

where $K(x_k, x)$ is the kernel basis function that determines the basis shape of the decision boundary and the parameters α_k and b are optimized parameters based on the set of y_k and x_k . The choice of kernel basis function determines the shape of the decision function [48]. Arbitrary shape estimation is desired so a Gaussian kernel basis (also known as a radial basis function) is chosen [49]. However, any kernel basis function can be used as a design choice. The Gaussian Kernel is given

as

$$K(x_k, x) = e^{-\frac{(x_k - x)^2}{\sigma^2}}. \quad (12)$$

A kernel of this form maps data into infinite dimensional feature space which provides flexibility for decision boundaries [50]. (In this work, the python package *Scikit-Learn* [51] is used for optimizing the boundary fit.)

B. Posterior Likelihood Estimate (Platt Scaling)

As the measurement model is unknown, we design the controller to be ergodic with respect to the likelihood distribution of acquiring a collision measurement. The distribution is updated for each measurement via Platt Scaling, defined as

$$P(y_k = 1|x) = \frac{1}{1 + e^{A\phi(x)+B}}, \quad (13)$$

where A and B are solved through a regression fit and $\phi(x)$ is the current shape estimate.

V. RESULTS

A. Simulation Results

Active shape estimation in $\mathbb{R}^2$ is done with double integrator dynamics given as

$$\dot{x} = f(x, u) = [\dot{x}, \dot{y}, u_1, u_2]^T. \quad (14)$$

The binary measurements are collisions detected during the crossing of the boundary function $\phi(x) = 0$. Estimation in $\mathbb{R}^3$ is done with a similar double integrator model with dynamics

$$\dot{x} = f(x, u) = [\dot{x}, \dot{y}, \dot{z}, u_1, u_2, u_3]^T. \quad (15)$$

No assumptions are needed on the location of the object and we initialize with a uniform distribution as the target distribution.

First, we demonstrate that the proposed algorithm's shape estimate converges to the actual shape. Figure 1 shows time evolution of the posterior as well as the resulting shape estimate. The estimate depicted by the black line in Fig. 1 eventually converges to the shape, although much of the shape restructuring is accomplished from non-contact motion up until 6s. Moreover, by comparing the posterior and the windowed sensor path, it is shown in Fig. 1 that the sensor path is drawn to high likelihood densities. This ensures that the posterior estimate is verified and updated as new sensor information is acquired. Note that the sensor is unable to access the interior of the shape, resulting in high likelihood probabilities within the shape. The posterior likelihood of the square shape estimate shows high likelihood probabilities near the corners of the square. As an aside, we take note of the similarities that the resultant likelihood has with the expected information density (EID) of the known shape measurement model (see [21] for EID derivation). Specifically, high likelihood near the corners correspond to the similar large expectation of information for a square shape. If we define the measurement model of the known square as $\Upsilon(x, x)$ where x is search space and $\mathbf{x}$ is the robot's state space, then a measure of information is the Fisher Information Matrix [52] defined by

$$I(\mathbf{x}, x) = \frac{\partial \Upsilon^T}{\partial x} \Sigma^{-1} \frac{\partial \Upsilon}{\partial x}, \quad (16)$$

where Σ is the measurement covariance. Assuming the measurement model is of the same form as in equation (1) then the region with the largest information is at the corners where the slopes of the edges collide. Thus, large likelihood estimates should exist near corners if they have not been previously searched.

The non-contact motion shown in Fig. 4 is beneficial for tactile-based exploration. The sensor trajectory from time 4.5s to 6s updates the shape estimation of the sine wave boundary without measuring a collision for $t \in [4.5, 6s]$. Knowing that no contact has occurred indicates that the shape is not in the current sensor state, but likely in other unexplored regions of sensor state. (Interestingly, tactile sensor arrays appear often in biological systems such as the rat whisker system [53]). The algorithm is shown in Fig. 3 to estimate both a clover shape and a diamond shape. Between times 2s to 6s in Fig. 3, non-contact motion is shown to drive the sensor trajectory towards an enclosing path that estimates the location of the center of the shape. After a few seconds, the sensor regularly is in contact with the shape as the estimate converges. The final estimates of the diamond and clover are shown by black lines in the rectangular box in Fig. 3.

The final time posterior shown in Fig. 1 and Fig. 3 show a resemblance in the posterior likelihood near the edges of the shape. This similarity in posterior estimates shows the algorithm approximates regions where there is high expected tactile information.

Shown in Fig. 6, multiple objects are allocated randomly in the sensor state. The sensor has no prior knowledge of where the shapes are, nor the number of shapes. The sensor contacts the first shape around 1s (Fig. 6), then between 2 and 6s the sensor comes into contact with the remaining two shapes. Following the trajectory (magenta) traversed from time 11 – 30s shown in Fig. 6, the sensor distributes the time spent amongst the three shapes. Because the ergodic metric is convex with respect to the information distribution and thus convex with respect to the shape, unexplored regions of sensor state must be explored [22].

We further expand upon multiple shape detection with a comparison with a version of Greedy expected entropy reduction (gEER) exploration algorithm used in [21], [54]–[57]. The algorithm samples nearby states centered at the robot probe's current location for regions with the largest probability of acquiring a collision measurement. The algorithm then moves the probe to that location, sampling along the way to adjust the shape estimate. In Fig. 9, we run 20 trials of uniformly sampled initial conditions for both eSAC and gEER. Figure 9(A) shows sample shape estimates for 80s simulations of both algorithms. eSAC is shown to detect all the shapes and begin shape refinement while gEER algorithm at most detects two shapes, but refines only one. We can see in (C) that eSAC detects all shapes within 40s of simulation time whereas the gEER detects two shapes at most. The time it takes to detect objects does depend on the distance of the shape and the control saturation of the robot probe which we maintained constant with both algorithms. In (B), we see the difference in the algorithm's area coverage. In particular, eSAC covers most of the search space coarsely with densely collected




Fig. 3. Time series shape estimation of diamond and clover shape. The posterior likelihood estimate is shown at 0.1, 1, 2, 6, 11, and 30 seconds. Final shape estimates at 30s is shown in the enclosed box. Likelihood of measuring a collision measurement is denoted by the contours. The sensor trajectory (shown in magenta) traverses from the previous time window onto the current posterior. Shape estimates are depicted in the black line. Actual continuous contact motion does not begin until around 6 seconds into the simulation. Nonetheless, through non-contact motion, the algorithm is able to estimate the shapes. After continuous contact occurs, the algorithm refines the shape estimate. The final shape estimates are highlighted in the enclosed rectangular box. Note that the shape estimate (denoted as the black line) for the square is underneath the actual shape. With the clover, part of the shape estimate is outside the shape boundary.




Fig. 4. Time series shape estimation of a spatial sine wave. High likelihood probabilities of a collision are shown in red and low probabilities in blue. Start point of sensor shown as the magenta square. By following the trajectory from 4.5s to 5.5s, the shape estimate depicted by the black line is updated by the sensor even without a collision. At 6s the middle peak of the sine wave has been estimated and the sensor trajectory has not collided with the sine wave boundary.

collision measurements around the shape.

Modularity with respect to shape representation is demonstrated by comparing with the shapes used in [19] with the Gaussian Process (GP) for shape estimation (see Fig. 5). Here, eSAC is shown to work with both a GP and SVM for shape estimates. In particular, in Fig. 5(A), eSAC uses the touch point selection metric used in [19] for control synthesis which is based on the covariance of the GP fit. In comparison with eSAC, the work done in [19] uses an Markov Decision Process (MDP) in order to control a robot manipulator towards acquiring touch data for shape estimation. In [19], a grid is




Fig. 5. Modularity in the eSAC algorithm is shown with the example shapes used in [19] and their approach for generating shape estimate and a target distribution. (A) A Gaussian Process (GP) is used for shape representation and the uncertainty in the GP fit as the target distribution (this implies the robot will search near regions of high uncertainty in the fit). (B) Support Vector Machine (SVM) approach for shape representation using the posterior likelihood as the target distribution. Both methods using eSAC provide similar results in final shape estimates.

defined and a path planner is used to traverse the grid. The large overhead state discretization and motion planning code that is necessary for the MDP formulation is not required for eSAC. Moreover, eSAC automatically encodes robot dynamic constraints. Notably, eSAC can be used with other visual-based estimators that feed initial shape estimates such that tactile data is used to refine the estimate.

This algorithm can be trivially extended to $\mathbb{R}^3$ as shown in simulation. Figure 8 shows the sensor trajectory in the top plot (magenta) actively sensing the torus shape. Shape estimation of the torus is shown to be successful in the lower sub-plot of Fig. 8 resulting in scalability to objects in $\mathbb{R}^3$. Performance for estimation in $\mathbb{R}^3$ depends on the robot dynamical constraints as well as the sensor area coverage.




Fig. 6. Time series of eSAC exploring multiple objects. The posterior is viewed at times 0.1, 1, 2, 6, 11, 30. Trajectories start at the magenta square. In each subplot, the sensor trajectory (magenta) starts from the endpoint of the previous time instant. Although the robot has no prior knowledge of the number of objects in the sensor state, the algorithm is still able to correctly estimate all the objects. Therefore, regardless of the number of shapes, the algorithm is still able to estimate all the shapes in the environment given a large enough time horizon.




Fig. 7. Baxter Robot used for experiments. An object is pinned down to the table for shape estimation using the end-effector probe. The probe itself has no sensors. Thus, the measured joint torques are used to detect collisions based on rapidly changing measurements at specific joint locations near the end-effector.

B. Experimental Results

Experiments are done using the Baxter robot (Fig. 7). Binary measurements are taken by the end-effector probe during a large change of joint torques, indicating a collision has occurred. During collision, controller weight R prevents the robot from dragging along the object. In the case that the control weight R is not significant, heuristics are used to prevent overexertion. Baxter's end-effector is used as a probe for the $\mathbb{R}^2$ search space. No assumptions are needed on the location of the object or the height of the object.

Experimental results show that shape estimation can be accomplished with a robotic system with multiple sensors (Fig. 10). Although there is odometry error within the robot's joint states, shape estimations are still visually seen to match. Within the posterior regime seen in Fig. 10, the probability densities for the shapes with corners are shown to have unique features. In particular, the posterior for the triangle in Fig. 10 extends outwards along the three corners indicating high probability of obtaining a collision measurement. Further mea-




Fig. 8. eSAC trajectory for a Torus $\mathbb{S}^2$ in $\mathcal{X} \subset \mathbb{R}^3 \subset [0, 1] \times [0, 1] \times [0, 1]$ search space. **Top:** In magenta, the robot trajectory on a Torus. **Bottom:** Density plot of the shape estimate. Although the simulation was run for 40s, the use of eSAC is shown to capture the doughnut shape inner circle feature which defines the torus shape. Utilizing an array of tactile sensors would reduce the error on the outer ring of the torus as larger coverage is needed for $\mathbb{R}^3$.

surements surrounding the edge regions reduce the likelihood of acquiring estimate which can be seen in with results from estimating the square shape.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, an algorithm is presented for shape estimation using a binary sensor combined with ergodic exploration. It is shown that a binary measurement sensor can be treated as a lower resolution tactile sensor for shape estimation with active




Fig. 9. Comparison of gEER and eSAC for detecting and refining multiple shape estimates. We randomly selected 20 initial conditions with uniform distribution in the search space for both gEER and eSAC methods. (A) Trial samples shown from both eSAC and gEER. eSAC consistently detects all objects in the search space. (B) Refined estimates of the shape after detection. eSAC is shown to coarsely explore the whole region while spending a larger fraction of time around the shape estimates. gEER is driven by larger values which reduce uncertainty which is susceptible to local minima. (C) eSAC is shown to detect all shapes in the environment within 40s of simulation. In all 20 trials, gEER does not successfully detect all shapes in the environment.




Fig. 10. Baxter experimental results for shape estimation. (A) shows the shapes used for estimation. (B) shows the location of Baxter's probe at the end of the estimation as well as the estimate of the shape as the dotted black line. The posterior at the end of the experiment is shown in the contour map. Note the prominent posterior likelihood in the triangle estimation around the edges are colored with high likelihood estimate, indicating that the estimated measurement model retains information about corners that is useful for localization. Error in shape area [19] is within 2.5%, 7.5%, and 10.0% for the circle, triangle, and square respectively. (Coarse discretization of the distribution was created to generate the figure to prevent run-time degradation on Baxter which results in an off-centered visual error of 1cm in the distribution.)

sensing. In addition, using ergodic exploration as the principle for active sensing uses sensor motion to capture shape features during estimation. The algorithm is shown to estimate shapes in $\mathbb{R}^2$ and $\mathbb{R}^3$ as well as an any number of shapes in the sensor domain. The algorithm is shown to be modular with respect to choice of shape representation. The algorithm is shown to have potential applications for active sensing with low resolution sensors with respect to shape estimation as well as applications in localization.

Future research directions include the use of multiple sensors for estimation. Another direction is to augment the current algorithm to include localization using the estimated mea-

surement model. Thus, simultaneous localization and shape estimation should be possible.

REFERENCES

- [1] M. Moll and M. A. Erdmann, "Reconstructing the shape and motion of unknown objects with active tactile sensors," in *Algorithmic Foundations of Robotics V*. Springer, 2004, pp. 293–309.
- [2] M. Meier, M. Schopfer, R. Haschke, and H. Ritter, "A probabilistic approach to tactile shape reconstruction," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 630–635, 2011.
- [3] S. J. Lederman and R. L. Klatzky, "Hand movements: A window into haptic object recognition," *Cognitive psychology*, vol. 19, no. 3, pp. 342–368, 1987.
- [4] K. Huynh, C. E. Stepp, L. W. White, J. E. Colgate, and Y. Matsuoka, "Finding a feature on a 3d object through single-digit haptic exploration," in *2010 IEEE Haptics Symposium*. IEEE, 2010, pp. 83–89.
- [5] G. Cannata, S. Denei, and F. Mastrogianni, "Tactile sensing: Steps to artificial somatosensory maps," in *19th International Symposium in Robot and Human Interactive Communication*. IEEE, 2010, pp. 576–581.
- [6] J. Ilonen, J. Bohg, and V. Kyriki, "Three-dimensional object reconstruction of symmetric objects by fusing visual and tactile sensing," *The International Journal of Robotics Research*, vol. 33, no. 2, pp. 321–341, 2014.
- [7] A. L. Trejos, J. Jayender, M. Perri, M. D. Naish, R. V. Patel, and R. Malthaner, "Robot-assisted tactile sensing for minimally invasive tumor localization," *The International Journal of Robotics Research*, vol. 9, pp. 1118–1133, 2009.
- [8] R. S. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile sensing :from humans to humanoid," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 1–20, Feb 2010.
- [9] P. Mitterdorfer and G. Cheng, "Humanoid multimodal tactile-sensing modules," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 401–410, June 2011.
- [10] M. Lee and H. Nicholls, "Review article tactile sensing for mechatronics state of the art survey," *Mechatronics*, vol. 9, no. 1, pp. 1 – 31, 1999.
- [11] M. I. Tiwana, S. J. Redmond, and N. H. Lovell, "A review of tactile sensing technologies with applications in biomedical engineering," *Sensors and Actuators A: Physical*, vol. 179, pp. 17 – 31, 2012.
- [12] D. Hughes and N. Correll, "A soft, amorphous skin that can sense and localize textures," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1844–1851.
- [13] D. Hughes and N. Correll, "Texture recognition and localization in amorphous robotic skin," *Bioinspiration & biomimetics*, vol. 10, no. 5, p. 055002, 2015.

- [14] Z. Pezzementi, E. Plaku, C. Reyda, and G. D. Hager, "Tactile-object recognition from appearance information," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 473–487, 2011.
- [15] J. Bimbo, S. Luo, K. Althoefer, and H. Liu, "In-hand object pose estimation using covariance-based tactile to geometry matching," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 570–577, Jan 2016.
- [16] M. Ye and R. Yang, "Real-time simultaneous pose and shape estimation for articulated objects using a single depth camera," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 2345–2352.
- [17] C. Dune, E. Marchand, C. Collin, and C. Leroux, "Active rough shape estimation of unknown objects," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2008, pp. 3622–3627.
- [18] N. F. Lepora, U. Martinez-Hernandez, and T. J. Prescott, "Active bayesian perception for simultaneous object localization and identification," in *Robotics: Science and Systems*, 2013.
- [19] T. Matsubara, K. Shibata, and K. Sugimoto, "Active touch point selection with travel cost in tactile exploration for fast shape estimation of unknown objects," in *Advanced Intelligent Mechatronics (AIM), 2016 IEEE International Conference on*. IEEE, 2016, pp. 1115–1120.
- [20] R. Mur-Artal and J. D. Tardós, "Probabilistic semi-dense mapping from highly accurate feature-based monocular slam," *Proceedings of Robotics: Science and Systems, Rome, Italy*, vol. 1, 2015.
- [21] L. M. Miller, Y. Silverman, M. A. MacIver, and T. D. Murphey, "Ergodic exploration of distributed information," *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 36–52, 2016.
- [22] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4, pp. 432–442, 2011.
- [23] L. M. Miller and T. D. Murphey, "Trajectory optimization for continuous ergodic exploration," in *2013 American Control Conference*. IEEE, 2013, pp. 4196–4201.
- [24] L. M. Miller and T. D. Murphey, "Optimal planning for target localization and coverage using range sensing," in *2015 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2015, pp. 501–508.
- [25] T. Coates and C. Taylor, "A mixture model for representing shape variation," *Image and Vision Computing*, vol. 17, no. 8, pp. 567 – 573, 1999.
- [26] M. Rousson and D. Cremers, "Efficient kernel density estimation of shape and intensity priors for level set segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2005, pp. 757–764.
- [27] L. L. Wong, L. P. Kaelbling, and T. Lozano-Pérez, "Not seeing is also believing: Combining object and metric spatial information," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1253–1260.
- [28] M. C. Koval, M. R. Dogar, N. S. Pollard, and S. S. Srinivasa, "Pose estimation for contact manipulation with manifold particle filters," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4541–4548.
- [29] V. Ganapathi, C. Plagemann, D. Koller, and S. Thrun, "Real-time human pose tracking from range data," in *European conference on computer vision*. Springer, 2012, pp. 738–751.
- [30] R. A. Newcombe, D. Fox, and S. M. Seitz, "Dynamic fusion: Reconstruction and tracking of non-rigid scenes in real-time," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 343–352.
- [31] J. J. Gibson, "Observations on active touch," *Psychological review*, vol. 69, no. 6, pp. 477–491, 1962.
- [32] H. J. Chiel, L. H. Ting, Ö. Ekeberg, and M. J. Hartmann, "The brain in its body: motor control and sensing in a biomechanical context," *The Journal of Neuroscience*, vol. 29, no. 41, pp. 12807–12814, 2009.
- [33] E. Guicé-Robles, C. Valdivieso, and G. Guajardo, "Rats can learn a roughness discrimination using only their vibrissal system," *Behavioural brain research*, vol. 31, no. 3, pp. 285–289, 1989.
- [34] G. E. Carvell and D. Simons, "Biometric analyses of vibrissal tactile discrimination in the rat," *The Journal of Neuroscience*, vol. 10, no. 8, pp. 2638–2648, 1990.
- [35] J. A. Hobbs, R. B. Towal, and M. J. Hartmann, "Spatiotemporal patterns of contact across the rat vibrissal array during exploratory behavior," *Frontiers in behavioral neuroscience*, vol. 9, p. 356, 2015.
- [36] K. H. Low, J. M. Dolan, and P. Khosla, "Adaptive multi-robot wide-area exploration and mapping," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 23–30.
- [37] A. Krause and C. Guestrin, "Nonmyopic active learning of gaussian processes: an exploration-exploitation approach," in *Proceedings of the 24th international conference on Machine learning*. ACM, 2007, pp. 449–456.
- [38] L. M. Miller and T. D. Murphey, "Optimal contact decisions for ergodic exploration," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*. IEEE, 2012, pp. 5091–5097.
- [39] D. A. Shell, C. V. Jones, and M. J. Matarić, "Ergodic dynamics by design: A route to predictable multi-robot systems," in *Multi-Robot Systems. From Swarms to Intelligent Automata Volume III*. Springer, 2005, pp. 291–297.
- [40] A. R. Ansari and T. D. Murphey, "Sequential action control: Closed-form optimal control for nonlinear and nonsmooth systems," *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1196–1214, Oct 2016.
- [41] L. Grüne and J. Pannek, "Nonlinear model predictive control," in *Nonlinear Model Predictive Control*. Springer, 2011, pp. 43–66.
- [42] H. Chen and F. Allgöwer, "A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability," in *Control Conference (ECC), 1997 European*. IEEE, 1997, pp. 1421–1426.
- [43] E. Tzorakoleftherakis, A. Ansari, A. Wilson, J. Schultz, and T. D. Murphey, "Model-based reactive control for hybrid and high-dimensional robotic systems," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 431–438, Jan 2016.
- [44] V. N. Vapnik and V. Vapnik, *Statistical learning theory*. Wiley New York, 1998, vol. 1.
- [45] T. Chateau, V. Gay-Belille, F. Chausse, and J.-T. Lapresté, "Real-time tracking with classifiers," in *Dynamical Vision*. Springer, 2007, pp. 218–231.
- [46] A. Niculescu-Mizil and R. Caruana, "Predicting good probabilities with supervised learning," in *Proceedings of the 22nd international conference on Machine learning*. ACM, 2005, pp. 625–632.
- [47] T.-F. Wu, C.-J. Lin, and R. C. Weng, "Probability estimates for multi-class classification by pairwise coupling," *Journal of Machine Learning Research*, vol. 5, no. Aug, pp. 975–1005, 2004.
- [48] B. Scholkopf and A. J. Smola, *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [49] N. Cristianini and J. Shawe-Taylor, *An introduction to support vector machines and other kernel-based learning methods*. Cambridge university press, 2000.
- [50] M. D. Buhmann, "Radial basis functions," *Acta Numerica 2000*, vol. 9, pp. 1–38, 2000.
- [51] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [52] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [53] M. J. Hartmann, "Active sensing capabilities of the rat whisker system," *Autonomous Robots*, vol. 11, no. 3, pp. 249–254, 2001.
- [54] C. Kreucher, K. Kastella, and A. O. Hero Iii, "Sensor management using an active sensing approach," *Signal Processing*, vol. 85, no. 3, pp. 607–624, 2005.
- [55] D. Fox, W. Burgard, and S. Thrun, "Active markov localization for mobile robots," *Robotics and Autonomous Systems*, vol. 25, no. 3, pp. 195–207, 1998.
- [56] H. J. S. Feder, J. J. Leonard, and C. M. Smith, "Adaptive mobile robot navigation and mapping," *The International Journal of Robotics Research*, vol. 18, no. 7, pp. 650–668, 1999.
- [57] J. R. Souza, R. Marchant, L. Ott, D. F. Wolf, and F. Ramos, "Bayesian optimisation for active perception and smooth navigation," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 4081–4087.