

Data-Driven Measurement Models for Active Localization in Sparse Environments

Ian Abraham, Anastasia Mavrommati, Todd D. Murphey
Department of Mechanical Engineering
Northwestern University, Evanston, IL 60208

Abstract—We develop an algorithm to explore an environment to generate a measurement model for use in future localization tasks. Ergodic exploration with respect to the likelihood of a particular class of measurement (e.g., a contact detection measurement in tactile sensing) enables construction of the measurement model. Exploration with respect to the information density based on the data-driven measurement model enables localization. We test the two-stage approach in simulations of tactile sensing, illustrating that the algorithm is capable of identifying and localizing objects based on sparsely distributed binary contacts. Comparisons with our method show that visiting low probability regions lead to acquisition of new information rather than increasing the likelihood of known information. Experiments with the Sphero SPRK robot validate the efficacy of this method for collision-based estimation and localization of the environment.

I. INTRODUCTION

Sensors such as cameras [15, 17, 36] and LiDAR [40, 41] provide high quality data about an environment with remarkably low sensitivity to the location of the sensor. As a result, the resulting sensing and perception are quite rich and are typically independent of the control. However, some sensing modalities, such as contact/impact sensors and range sensors, only provide useful information when they are in the correct configuration or sequence of configurations—for these sensors, high quality data is sparse in the environment and must therefore be sought out by utilizing control authority. In this paper, we investigate the use of active sensing, through means of ergodic control [29, 32, 33], to provide control authority that seeks out useful information from sensors for which informative measurements are sparse in the environment.

In biology, we see this combination of movement and sensing in rodents that use their whiskers for tactile sensing [7, 16, 18, 34]. These biological sensors are active sensors (i.e., sensing is done through movement) and provide a rich set of contact and force information. Notably, no two whiskers are the same mechanically—more so across rodents of the same species—and yet each rodent is able to use these biological sensors to sense and explore their environment. Similarly, humans use their hands to grasp and to feel objects by actively moving their fingers across objects. New objects are rapidly incorporated into memory without needing to understand and model the complex physical interactions between the hand and the object. We are interested in how one automates the generation of these models, and what principles extend to robotic settings.

The combination of environmentally sparse, unmodeled sensor-environment interactions with motion leads us to the topic of this work: equip a robot with the means of active sensing when the environment is unknown and the sensor interaction with the environment is spatially sparse. In doing so, we seek to improve the utility of sparse sensors (e.g. contact sensors [18, 30, 34]), and their interaction with the environment by synthesizing movement, both for model generation and later identification and localization. We approach this problem of active sensing as a coverage problem where a robot’s motion (and the time spent in regions of space) should coincide with an *expectation of informative sensory input*. Ergodic control policies [27, 31, 33, 42] are used to compute active sensing strategies where uncertainty about the environment, the robot’s sensors, and the interaction is reasoned about as a dynamically evolving area coverage problem. Thus, the contribution of the work is the construction and use of online data-driven distributions that, when used to specify the ergodic control, enables a robot to compensate, through motion, for spatially sparse sensor information and unknown environments to estimate and localize the environment. Simulations illustrate the phases of the algorithm and experiments validate our approach for collision-based robot sensing.

The paper outline is as follows: Section II discusses literature related to this work. Section III introduces the problem statement that is addressed in the paper. Section IV introduces ergodicity and the ergodic control policy. Section V then formulates a method for constructing data-driven measurement models of the environment. Section VI then describes how to use these empirically determined measurement models for localization of the environment. Simulated and experimental results are in Sections VII and VIII with conclusion in Section IX.

II. RELATED WORK

Existing work in robotics that deals with unknown, environmentally sparse sensors typically focus on tactile sensing [9, 13, 25, 25, 26, 28, 30, 35, 37, 48]. Uncertainty distributions are generated that guide the robot towards regions where the robot is likely to acquire a positive contact measurement. Sample-based methods [9, 13, 25, 28, 35, 37] or maximum likelihood methods [9, 25, 26, 30, 48] are used to determine the robot’s next area to sample. An issue with these methods is the omission of low probability regions. By omitting low probability regions, only information near prior

high information areas is obtained to make decisions about where to move. Our approach uses ergodic exploration to enable a robot to systematically explore all nonzero probability regions, avoiding overly focusing on already explored regions.

Ergodic exploration has been used previously to enable search for spatially distributed information [29, 32, 33], similar to sample-based methods. In sample-based methods, the procedure is often to sample the search space and then control the robot to the desired states calculated from the samples. The primary benefit of using ergodic control is the ability to avoid controllers that get stuck in local minima and cyclical patterns as we show in Section VII-C. By using ergodic exploration, our method is able to continuously explore an environment subject to the physical constraints imposed by the robot, the sensor, and the environment. In particular, these physical constraints are used to localize unknown environments directly from empirically determined models. Prior work from the authors have illustrated the use of ergodic exploration for shape estimation [1]. Here, we extend this work for a more general class of sparse non-null measurements and show that the constructed models can be used as a method for data-based environment estimation and localization.

III. PROBLEM STATEMENT

The goal of this work is to enable a robot to autonomously localize an environment when sensory information is sparsely distributed in the environment. Consider the state of an environment $\sigma \in \mathcal{O}$ where \mathcal{O} is the set of points that defines the surface of elements in the environment that a sensor can interact with.¹ A robot's measurements, given its state $x(t) \in \mathbb{R}^n$, is defined as

$$y = \Gamma(\sigma, x_v) + \delta \in \mathbb{Y}^P \quad (1)$$

where $\delta \sim \mathcal{N}(0, \Sigma)$ is zero mean Gaussian noise with variance Σ , $\Gamma(\sigma, s) : \mathcal{O} \times \mathbb{R}^v \rightarrow \mathbb{Y}^P$ is the measurement model, $x_v \in \mathbb{R}^v \subset \mathbb{R}^n$ is the state of the robot that interacts with the environment, and $s \in \mathbb{R}^v$ is a point in the search space. The measurements in (1) are then sparse in the environment and discrete, yielding non-null measurements in particular regions of the search space (e.g., collision measurement in the environment with a finite number of objects).

Given that $\Gamma(\sigma, s)$ is initially unknown for a sparse environment defined by σ , how should a robotic agent empirically construct $\Gamma(\sigma, s)$? Moreover, provided that $\Gamma(\sigma, s)$ was constructed in some search space \mathcal{P} , how can the robotic agent localize within the environment? Thus, the problem statement that this paper addresses is the following: Given some search space \mathcal{W} that has the environment σ within it, estimate the transformation $g(\theta)$ given that $\Gamma(\sigma, s)$ is initially unknown and cannot be estimated by local exploration of \mathcal{P} .

In this work we approach this problem by dividing it into two stages. The first stage is to explore the environment and generate a data-driven measurement likelihood distribution

¹For generality, we define the set of points in \mathbb{R}^v , but one can consider the set of points in \mathbb{R}^3 as the outer surface of objects that interacts with a contact sensor.




Fig. 1. Objects in the environment are defined by the blue shapes. (a) A robot with an exploratory trajectory $x(t)$ is shown as the gray line estimating the measurement model of an environment in the fixed frame \mathcal{P} . Contact measurements are used to construct the measurement model. (b) The same environment originally shown in \mathcal{P} is transformed into search space \mathcal{W} under transform $g(\theta)$. The robot exploratory trajectory used to generate an estimate of the transform $g(\theta)$ is created using the data-driven measurement model in (a).

from a set of N measurements and a resulting measurement model for the environment. This measurement likelihood distribution is necessary because the resulting distribution is not solely a sensor model, but rather a model of how the robot, environment, and the sensors interact with one another. The result is the construction of (1), which is the measurement model for how the robot, the sensor, and the environment interact based on data.

The second stage is to use this constructed measurement model to localize and estimate the transform $g(\theta)$ from $\mathcal{P} \rightarrow \mathcal{W}$. We illustrate the problem in Fig. 1 for $\mathcal{P}, \mathcal{W} \in \mathbb{R}^2$. Sections V and VI illustrate the environment measurement model construction and localization.

The following section introduces ergodicity and the ergodic metric.

IV. ERGODICITY AND ERGODIC CONTROL FOR ACTIVE SENSING

In this section, we review ergodicity and the ergodic metric for active sensing and control for robotic systems [27, 42]. Ergodicity and ergodic control enables us to specify the motion of the robot based on requiring that the time it spends in regions of state-space be proportional to desired spatial statistics in that region, where the desired spatial statistics are determined by an estimate of the information density. It is later shown in Section VII-C through a comparison why using ergodic control is beneficial in the application of sparse sensing.

A. Ergodicity

Consider an agent whose state at time $t \in \mathbb{R}^+$ is $x(t) : \mathbb{R}^+ \rightarrow \mathbb{R}^n$ and the control input to the robot as $u(t) : \mathbb{R}^+ \rightarrow \mathbb{R}^m$. The time evolution of the robot is assumed to be governed by a control-affine dynamical system of the form

$$\dot{x} = f(x(t), u(t)) = g(x(t)) + h(x(t))u(t) \quad (2)$$

where $g(x) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the free dynamic response of the robot, and $h(x) : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ is the driving dynamic re-

sponse subject to input u . Next, define a bounded search space domain $\mathcal{X}_v \subset \mathbb{R}^v$ whose limits are $[0, L_1] \times [0, L_2] \times \dots \times [0, L_v]$ with $v \leq n$. The time-averaged statistics $c(s, x(t))$ of the robot's trajectory $x(t)$ (i.e., where the robot spends most of its time) for some time interval $t \in [t_i, t_i + T]$ is given by

$$c(s, x(t)) = \frac{1}{T} \int_{t_i}^{t_i+T} \delta(s - x_v(t)) dt \quad (3)$$

where δ is a Dirac delta function, $T \in \mathbb{R}^+$ is the time horizon, $t_i \in \mathbb{R}^+$ is the i^{th} sampling time, $s \in \mathbb{R}^v$ is a point in the search space, and $x_v(t) : \mathbb{R}^+ \rightarrow \mathbb{R}^v$ is the state of the robot that exists in the search space \mathcal{X}_v . We then define a "target" distribution $\phi(s) : \mathcal{X}_v \rightarrow \mathbb{R}^+$ with respect to which the robot agent is to be ergodic (i.e., time spent during a trajectory $x(t)$ is proportionate to the spatial statistics of that region). An ergodic metric (based on a Sobolev space norm) [27] which relates the two distributions $c(s, x(t))$ and $\phi(s)$ is:

$$\begin{aligned} \mathcal{E}(x(t)) &= q \sum_{k \in \mathbb{N}^v} \Lambda_k (c_k - \phi_k)^2 \quad (4) \\ &= q \sum_{k \in \mathbb{N}^v} \Lambda_k \left(\frac{1}{T} \int_{t_i}^{t_i+T} F_k(x_v(t)) dt - \phi_k \right)^2 \end{aligned}$$

where

$$\phi_k = \int_{\mathcal{X}_v} \phi(s) F_k(s) ds \quad (5)$$

and c_k are the Fourier decompositions² of $c(s, x(t))$ and $\phi(s)$ with

$$F_k(x) = \frac{1}{h_k} \prod_{i=1}^v \cos\left(\frac{k_i \pi x_i}{L_i}\right)$$

being the cosine basis function for a given coefficient $k \in \mathbb{N}^v$, h_k is the normalization factor defined in [27], $\Lambda_k = (1 + \|k\|^2)^{-\frac{v+1}{2}}$ is a weight on the frequency coefficients, and $q \in \mathbb{R}^+$ is a weight on the ergodic metric. Equation (4) now relates the time spent by the robot to the spatial statistics. A robot whose control inputs result in a trajectory $x(t)$ that minimizes (4) is then said to be optimally ergodic with respect to the target distribution.

B. Ergodic Control

Control inputs that result in a maximally ergodic trajectory with respect to a target distribution are synthesized following the work in [29]. Instead of directly minimizing (4) with respect to $x(t)$ and $u(t)$, we consider the sensitivity of (4) with respect to an infinitesimal application duration time $\lambda \rightarrow 0 \in \mathbb{R}^+$ of the best possible control $u_*(t) : \mathbb{R}^+ \rightarrow \mathbb{R}^m$ that sufficiently reduces (4) at time $\tau \in \mathbb{R}^+$ from some default control $u_{\text{def}}(t) : \mathbb{R}^+ \rightarrow \mathbb{R}^m$. This sensitivity (known as the mode insertion gradient [4, 6, 10, 46]) is given by

$$\left. \frac{\partial \mathcal{E}}{\partial \lambda} \right|_{\tau} = \rho(\tau)^T (f_2(\tau, \tau) - f_1(\tau))$$

²The cosine basis function is used here, however, any choice of basis function F_k can be used.

where $f_2(t, \tau) = f(x(t), u_*(\tau))$, $f_1(t) = f(x(t), u_{\text{def}}(t))$, and $\rho \in \mathbb{R}^n$ is the adjoint variable which is the solution to

$$\dot{\rho} = -2 \frac{q}{T} \sum_{k \in \mathbb{N}^v} \Lambda_k (c_k - \phi_k) \frac{\partial F_k}{\partial x} - \frac{\partial f}{\partial x} \rho$$

where $\rho(t_i + T) = \mathbf{0} \in \mathbb{R}^n$. The control is then

$$u_* = (\Omega + R^T)^{-1} [\Omega u_{\text{def}} + h(x)^T \rho \alpha_d] \quad (6)$$

where $\Omega \triangleq h(x)^T \rho \rho^T h(x)$, $\alpha_d \in \mathbb{R}^-$ parametrizes the aggressiveness of the control, and $R \in \mathbb{R}^{m \times m}$ is a positive definite matrix that weighs the control u_* . Note that as in [29], saturation is taken into account and duration time λ is found using a line search.

Further information about the algorithm provided in Algorithm 1 can be found in [29].

Algorithm 1 Ergodic Control

- 1: **given:** $x_i, \phi_{k,i}, t_i, T, \Delta t_{\mathcal{E}}$
 - 2: simulate $(x(t), \rho(t))$ from initial condition (x_i, ρ_i, t_i, T) and default control u_{def}
 - 3: compute $c_{k,i}$ using past history $\Delta t_{\mathcal{E}}$ and simulated trajectory $x(t)$
 - 4: calculate the control $u_*(t)$ that reduces the ergodic metric with $(x(t), \rho(t), c_{k,i}, \phi_{k,i})$
 - 5: $\tau \leftarrow \text{argmin } \partial \mathcal{E} / \partial \lambda$
 - 6: saturate $u_*(\tau)$
 - 7: compute λ from line search
 - 8: **return** $u_*(\tau)$ for duration λ
-

The following two sections use the ergodic control policy for active sensing to first empirically determine the measurement likelihood of the environment in response to the sensor and then use the measurement model for localization.

V. CONTROL FOR MEASUREMENT MODEL CONSTRUCTION

In this section, we focus on empirically constructing the measurement model of an environment from the movement of a robot and its sensor interaction with the environment. We show how we can leverage the construction of the measurement model into the exploration procedure using ergodic control for active sensing.

A. Construction of Measurement Likelihood with Active Sensing

The approach here is similar to [1] for shape estimation. The controller is initialized with a uniform target distribution over a finite domain (e.g., a volume in which an object resides); however, the controller can be initialized with a prior if information is known. The ergodic control policy then provides control input for the robot based on an initial measurement likelihood distribution (often uniform to start). At each $t_i = t_{i-1} + t_s$ with sampling interval t_s , the robot collects measurements x_i, y_i .³

³Measurements can be collected asynchronously in order to prevent measurements with duration time $t_c < t_s$.

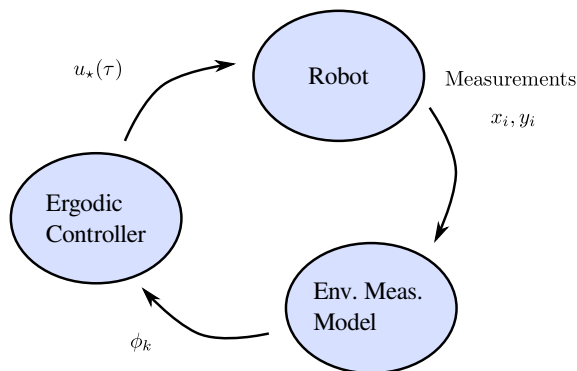


Fig. 2. Block diagram for measurement model construction. Measurements are processed in the measurement likelihood probability. This information is then passed onto the ergodic controller which returns an active sensing policy for the robot.

In this work, we use a Support Vector Machine [38, 45, 47]⁴ to calculate $p(y_i | x_i, \sigma)$; however, any method—such as Gaussian processes [22, 28] or density estimators [43]—can also work to approximate the measurement likelihood so long as one can generate a distribution over the measurements. In the case of collision-based sensing, the coefficients ϕ_k in Eq.(5) are calculated using $p(y = 1 | x, \sigma)$ for $y = 1$ yielding a collision measurement. Since $p(y = 1 | x, \sigma)$ is defined over the search space s , we numerically integrate (5) for the current estimate of $p(y = 1 | x, \sigma)$. This allows the algorithm to leverage a particular desired measurement as an active learning procedure. The process for the construction of the measurement model is illustrated in Fig. 2 and shown in Algorithm 2.

Algorithm 2 Ergodic Control for Measurement Model Construction

- 1: **define:** $x_0, \phi_{k,0}, i = 0, t_0, t_f, t_s, T, \Delta t_{\mathcal{E}}$
 - 2: **while** $t_i < t_f$ **do**
 - 3: compute $u_*(\tau)$ from Algorithm 1
 - 4: sample x_i, y_i from robot (assume asynchronous measurements)
 - 5: determine $p(y_i | x_i, \sigma)$ from data x_i, y_i using SVM (or Gaussian process)
 - 6: calculate target distribution $\phi_{k,i}$ from $p(y | x, \sigma)$
 - 7: $i \leftarrow i + 1$
 - 8: **end while**
-

VI. CONTROL FOR LOCALIZATION USING ENVIRONMENTAL INFORMATION FROM DATA-DRIVEN MEASUREMENT MODELS

In this section, we define the information of a environment measurement model and how one would actively sense with respect to the information of the model for localization. Assume that we have a measurement model for an environment.

⁴A Gaussian kernel is used with variance parameter $\sigma^2 = 0.01$ with input as x_i and output y_i .

Given a new search space $\mathcal{W} \subseteq \mathbb{R}^v$ and that the model was acquired in some search space $\mathcal{P} \subseteq \mathbb{R}^v$, we can define a transform $g(\theta) \in \text{SE}(2)^5$ parametrized by $\theta \in \mathbb{R}^3$ that transforms elements in $\mathcal{P} \rightarrow \mathcal{W}$ (see Fig. 1 for illustration). Localizing based on the data-driven measurement model is then reduced to identifying the parametrization θ that describes the transformation $g(\theta)$.

A. Information about $g \in \text{SE}(n)$

In order to use ergodic control for active localization of the environment, we first need to define a distribution which captures where the robot should explore next based on the environment sensitivities with respect to its dependences. In order to obtain this relationship, the Fisher information [2, 8, 12] of the measurement model of environment σ with respect to changes in the search space is used. While other measures of information such as entropy can be used, the Fisher information directly encodes parameter sensitivities that, as we will show in Lemma 1, includes the Fisher information of how the robot explored the search space.

Lemma 1: The expected information of an environment measurement model with respect to a transformation $g(\theta)$ is given by

$$\phi(s) = \det \left[\int_{\theta} \frac{\partial (g(\theta)^{-1}s)^T}{\partial \theta} I(s) \frac{\partial (g(\theta)^{-1}s)}{\partial \theta} p(\theta) d\theta \right] \quad (7)$$

where

$$I(s) = \frac{\partial \Gamma^T}{\partial s} \Sigma^{-1} \frac{\partial \Gamma}{\partial s}$$

is the Fisher information with respect to the search space of the robot, $p(\theta)$ is the belief of the parameter θ , and D-optimality [12, 19] (determinant) is chosen a measure of information.

Proof: We first calculate the Fisher information matrix of the empirically determined measurement model Γ (Eq. 1) with respect to the transformation:

$$I(s, \theta) = \frac{\partial \Gamma^T}{\partial \theta} \Sigma^{-1} \frac{\partial \Gamma}{\partial \theta}. \quad (8)$$

Applying chain rule to the derivative terms in (8) gives us the expanded derivative terms

$$\frac{\partial \Gamma}{\partial \theta} = \frac{\partial \Gamma(\sigma, \overbrace{g(\theta)^{-1}s}^{\bar{s}})}{\partial \bar{s}} \frac{\partial (g(\theta)^{-1}s)}{\partial \theta}. \quad (9)$$

Replacing (9) into (8) gives

$$\begin{aligned} I(s, \theta) &= \left(\frac{\partial \Gamma}{\partial \bar{s}} \frac{\partial g(\theta)^{-1}s}{\partial \theta} \right)^T \Sigma^{-1} \left(\frac{\partial \Gamma}{\partial \bar{s}} \frac{\partial g(\theta)^{-1}s}{\partial \theta} \right) \quad (10) \\ &= \frac{\partial (g(\theta)^{-1}s)^T}{\partial \theta} \underbrace{\left(\frac{\partial \Gamma^T}{\partial \bar{s}} \Sigma^{-1} \frac{\partial \Gamma}{\partial \bar{s}} \right)}_{I(s)} \frac{\partial (g(\theta)^{-1}s)}{\partial \theta}. \end{aligned}$$

⁵For simplicity we define the problem in the planar space, but more abstract and high dimensional transforms can be utilized.

Taking the expectation of (10) and the D-optimality measure [12, 19] gives us the expected information density of an environment measurement model

$$\phi(s) = \det \left[\int_{\theta} \frac{\partial (g(\theta)^{-1}s)^T}{\partial \theta} I(s) \frac{\partial (g(\theta)^{-1}s)}{\partial \theta} p(\theta) d\theta \right].$$

We can see that the Fisher information with respect to θ includes spatial information about the measurement model with respect to the search space. Thus, by specifying the expected information density as the target distribution for ergodic control, the robot driven towards regions where there is information about the transformation $g(\theta)$ and spatial features that are present in the measurement model.

B. Estimating $p(\theta)$ using Environment Measurement Likelihood

We can estimate $p(\theta)$ from the measurements using a Bayesian update [21, 45] of the form

$$p_{i+1}(\theta) = \eta p(y_i | \theta, x_i, \sigma) p_i(\theta) \quad (11)$$

starting from a uniform prior $p_0(\theta)$. Note that $p(y_i | \theta, x_i, \sigma)$ is simply a likelihood function as a function of θ . Therefore, we update the belief of the parameters θ using only the likelihood function which contains the family of all possible measurement models subject to transformation $g(\theta)$. Because non-null measurements are sparse, and contact measurements are discontinuous, we use a manifold particle filter [3, 21, 44] in order to approximate (11). Other methods for updating $p(\theta)$ such as a Kalman filter [20] can be used; however, note that the sparsity in measurements can make the filter unstable due to the discontinuous measurement model [21].

The algorithm for active localization of the environment using Algorithm 1 is provided in Algorithm 3. An initial prior $p_0(\theta)$ and measurement likelihood function is defined. The expected information of the environment information is calculated using finite difference. Fourier coefficients ϕ_k of the expected information are computed numerically using the particles of the particle filter as a Monte Carlo integration. Given some initial condition $x(t_0)$, the controller computes a control action subject to ϕ_k . Measurements are collected and used to update $p(\theta)$ using $p(y_i | \theta, x_i, \sigma)$.

Algorithm 3 Active Sensing for Localization using Data-driven Measurement Models

- 1: **initialize:** Algorithm 1, $x_0, p_0(\theta), \phi_k$
 - 2: **while** $t_{curr} < t_f$ **do**
 - 3: calculate $u_*(\tau)$ from Algorithm 1
 - 4: apply $u_*(t_i)$ to robot
 - 5: sample x_i, y_i from robot (assume asynchronous measurements)
 - 6: update $p_i(\theta)$ from $p(y_i | \theta_i, x_i, \sigma)$
 - 7: compute ϕ_k from expected information (7)
 - 8: $i \leftarrow i + 1$
 - 9: **end while**
-

In the following sections, we present simulated and experimental examples of constructing environment measurement model and using the models for sensing and localization of the transformed environment.

In this section, we present simulated examples of environment localization using empirically determined environment measurement models. We first focus on examples for measurement model construction for objects in an environment using contact in \mathbb{R}^2 and \mathbb{R}^3 and then show examples of localization using a known measurement model in \mathbb{R}^2 . We then provide a comparison with an entropy-based method to illustrate the benefits of using ergodic control for sensors with spatially sparse information. The experimental section then demonstrates the end-to-end algorithm where we use a robot with a collision-based sensor for localizing a transformed environment from an empirically determined environment measurement model.

VII. SIMULATED EXAMPLES OF ENVIRONMENTAL MEASUREMENT MODEL CONSTRUCTION AND LOCALIZATION

A. Measurement Model Estimation

Examples of environment measurement models in \mathbb{R}^2 and \mathbb{R}^3 are shown for multiple objects in Fig. 3 and 4. The dynamics of the robot are given as a double integrator system with dynamics $\dot{x} = [\dot{x}, \dot{y}, u_1, u_2]^T$ in \mathbb{R}^2 and $\dot{x} = [\dot{x}, \dot{y}, \dot{z}, u_1, u_2, u_3]^T$ in \mathbb{R}^3 . Here, we use $p(y = 1 | x, \sigma)$ as the target distribution for active sensing.

The robot is initialized with a uniform distribution and the measurement likelihood is constructed as measurements are taken. This is shown in sliced regions in \mathbb{R}^3 where high values of the measurement likelihood guide the motion of the robot (the same method is used when there are multiple objects in \mathbb{R}^2). Note that while measurements with higher probability are frequently acquired, the ergodic policy requires sampling from low probability regions. This enables the robot to acquire measurements from the environment where the measurement model is well established, refining the overall model, while incorporating measurements where the environment is uncertain. Thus, ergodic exploration promotes effective exploration of the environment for construction of the measurement model.

Using this measurement model that is constructed from active exploration, we perform active localization of the environment using the resulting information of the measurement model as the target for ergodic control. We illustrate this in the following subsection.

B. Measurement Model-based Localization

Simulated results for active sensing for localizing the environment from an environment measurement model are done with the double integrator model in \mathbb{R}^2 . In this example, three objects are in the environment: a circle, a square, and an oval as shown in Fig. 5. The environment measurement model is known and is used for active sensing using Algorithm 3. A uniform distribution over θ is defined as the prior.

As shown in Fig. 5, at first contact, the belief in the environment's location starts to approach the ground truth

Environment Measurement Model Generation




Fig. 3. Estimating the measurement model of the environment with multiple objects in $\mathbb{R}^2 \subset [0, 1] \times [0, 1]$ search space using collision measurements. The objects are shown as with the purple outline. The simulation is initialized using a uniform distribution over the \mathbb{R}^2 search space. The measurement likelihood from the data-driven measurement model is shown as the dark regions beneath the objects. This distribution drives the ergodic controller as it is being created from data. The resulting robot trajectory is shown as the red line. Note that the robot still visits low probability region in search for new information. Thus, the robot is able to circle around objects and better explored the environment.

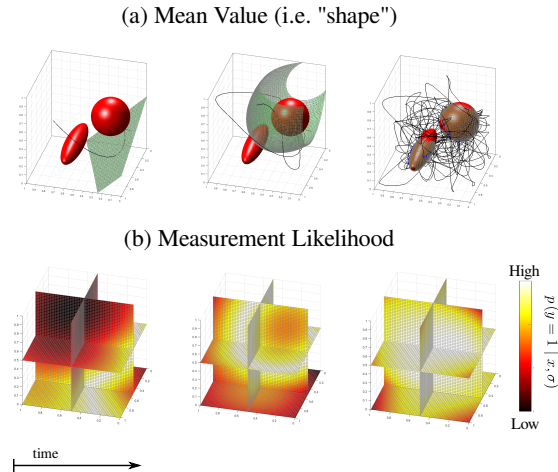


Fig. 4. Estimating the measurement model of the environment with multiple objects in $\mathbb{R}^3 \subset [0, 1] \times [0, 1] \times [0, 1]$ search space using collision measurements. The simulation is initialized using a uniform distribution over the \mathbb{R}^3 search space. (a) Mean values of the contact measurements $p(y | x, \sigma)$ are depicted as the green surfaces. Robot trajectory is shown as the black line. (b) Slices of the measurement likelihood distribution are shown being constructed as the robot collects measurements. Under ergodic control, the robot is able to explore the whole environment to generate the measurement model.

shown as the black shapes. The expected information density shown in the background of Fig. 5 begins to collapse from the family of possible measurement models to a probabilistically certain measurement model. As more measurements are aggregated, the prior on θ starts to converge. Ergodic control then facilitates correction of the belief $p(\theta)$ via active sensing with respect to the environment information calculated from the measurement model.

C. Comparison with Entropy Reduction

We compare our method using ergodic control with another method that is similar in nature known as expected entropy reduction (EER) [1, 32, 33]. The algorithm is similar to that used in [11, 14, 23, 24] where a location in the search space is chosen based on where the expected change in entropy is maximized. The location is chosen from a set of 200 samples uniformly distributed in the search space. The trajectory that maximizes the entropy reduction is then chosen and an LQR controller is used to control the robot with the double integrator dynamics mentioned in the previous subsection.

Here, the reduction in entropy is calculated as

$$H(\theta) - \mathbb{E} [H(\theta) | y(t)^+]$$

where

$$H(\theta) = - \int p(\theta) \log p(\theta) d\theta$$

is the entropy for a random variable θ , and $y(t)^+$ is the expected measurement given the current model along the sampled trajectory. In the case of measurement model construction with impact, θ is replaced with $y = 1$ whereas in environment localization, the random variable is the transform parameters θ for $g(\theta) \in \text{SE}(2)$.

We compare a 20 second simulation of environment construction and localization (with the known environment measurement model) in Fig. 6. Because EER prioritizes immediate reduction in entropy, only the most recent and known measurements are used. In situations where measurements are sparse, this causes the EER algorithm to become stuck as shown in Fig. 6(d). Similarly, in Fig. 6(b), EER is only able to acquire measurements with two out of the three objects in the environment, and only one side of each object. In contrast, using ergodic control as shown in Fig. 6 (a) and (c) illustrates the benefit for sensing sparse impact signals in unknown environments. Low probability regions are not omitted and are important to both stages of our method. As time goes to infinity, the ergodic controller would drive the robot to explore the whole domain such that the time spent is proportional to the statistics in the region [1, 29, 32, 33].

In the following section, we realize our method end-to-end on a colliding robot experiment.

VIII. EXPERIMENTAL EXAMPLES OF ENVIRONMENTAL MEASUREMENT MODEL CONSTRUCTION AND LOCALIZATION USING SPHERO SPRK

In this section, we present experimental examples of the end-to-end algorithm for localizing an environment using an empirically determined environment measurement model. A rolling robot driven by an internal differential drive known as the Sphero SPRK is used for collision-based sensing. The SPRK robot has an internal contact sensor which uses its inertial measurement unit to detect a collision with an external object. Position of the robot is acquired using OpenCV [5] and a Kinect camera. Communication between the position tracker, controller, and robot is done through ROS [39]. Once




Fig. 5. Localization of objects using a given measurement model is illustrated for an environment containing three objects. The expected information of the measurement model is shown as the distribution in the background of each time frame t_i space 1.25 seconds apart. The transparent green objects are the depiction of the current belief of their location $p(\theta)$. The ground truth object location is shown in black and the robot trajectory in shown as the blue line. The robot estimates the transformation using the information of the environment measurement model as a target for active sensing. We refer the reader to the multimedia video for an animation of this result using a manifold particle filter [21].

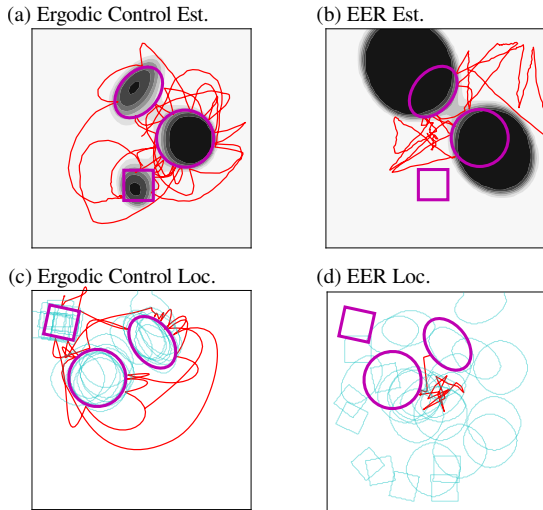


Fig. 6. Comparison of Ergodic control versus entropy reduction for environment measurement estimation (a-b) and localization using the environment measurement model that is known (c-d). Contact likelihood shown as the dark regions in (a-b). Particle filter beliefs are shown as the cyan colored estimates of the shapes in (c-d). Using ergodic control enables a robot to consider low probability regions which are beneficial in acquiring impact measurements that are sparse in the environment. Entropy reduction results in search patterns that immediately reduce the entropy which results in a lack of exploration in sparse environments.

a measurement model is generated in the first stage, we use this model to localize the transformed environment. Note that the only assumption that is made is that the robot moves subject to the double integrator dynamics in \mathbb{R}^2 .

A. Measurement Model Construction

Experimental results for constructing the environment measurement model using collisions with the SPRK robot are shown in Fig. 7. Within the first few measurements, the environment measurements are not useful as the robot has not explored the environment domain sufficiently to find impacts. However, due to the low probability value of unexplored regions, the robot under ergodic control must still visit these regions. In doing so, the robot encounters more impact mea-

surements by exploring the environment. By the termination time ($t_f = 120$ seconds), the robot has constructed an environment measurement model that is sufficient for subsequent localization.

It is worth noting that the collision sensor of the robot is sensitive to quick movements which will result in noisy measurements. While in an engineered or physics-based approach, the noise of a robot's position and sensors are derived, our method automatically incorporates the noise effects into the measurement model. Moreover, the underlying physical constraints in the robot and the environment are also captured in the environment measurement model. This allows us to model the environment empirically without the need to model or estimate the environment or robot parameters. In the following section, we demonstrate localization of the environment with respect to the information of the constructed measurement model.

B. Localization

Using the measurement model from the previous subsection, we localize the transformed environment. In Fig. 8, the objects, as presented in Fig. 7, are rotated and translated in this localization experiment. The initial uniform prior belief $p_0(\theta)$ is defined as a uniform distribution over the domain $\theta \in [0, 1.2] \times [0, 1] \times [-2\pi, 2\pi]$ for the positional and rotational transformation. Here, the experiment is run for $t_f = 130$ seconds. Once the robot begins to collect measurements, we can see that the localization of the environment starts to approach the ground truth shown as the red outlined shapes.

While collisions are important in this experiment, non-collisions are equally if not more important for localizing the environment based on the empirically determined measurement model. Consider that collision measurements are rare occurrences in the environment that is sparse with respect to the number of objects that occupy the space. As a consequence, the environment measurement model will be based on not having collision measurements in particular regions of space which is used in this example to further estimate the location of the transformed environment. We can see this in the final




Fig. 7. Left: Photo of the experiment setup of the environment state σ . The robot is initialized with a uniform target distribution. Top: Time series of the environment measurement model construction process from left to right. The red boundaries indicate the mean value of the measurement likelihood (i.e., shape). Sphero robot is shown as a green circle. The robot trajectory up to the current time is also shown. Bottom: $p(y = 1 | x, \sigma)$ is shown being built as the robot collides with the objects in the environment. We refer the reader to the multimedia video for the experimental videos.

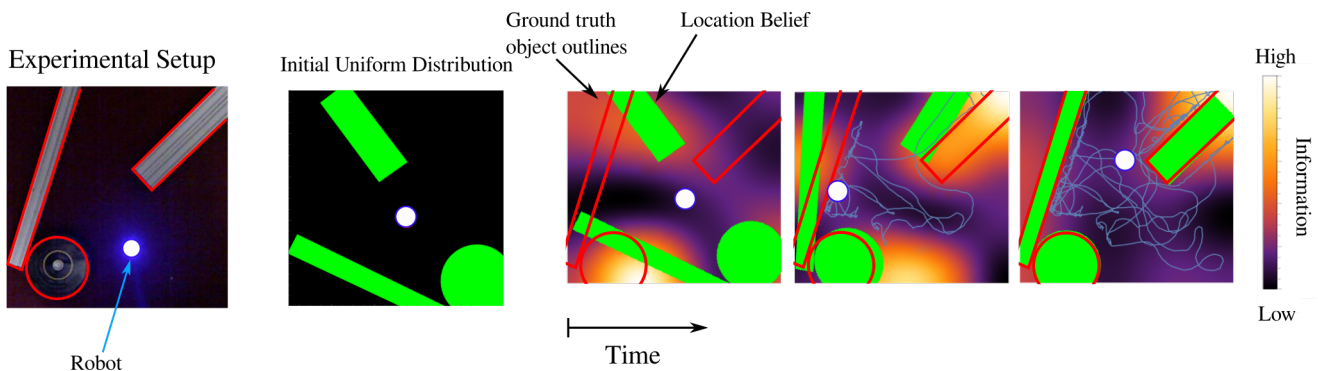


Fig. 8. The robot estimates the transformation of the environment state using the measurement likelihood extracted from Fig. 7 using contact measurements. The global position of the objects is shown as the red outlines. The robot ergodic control policy result in a sequence of contact measurements that localizes the environment based on the environment information. We refer the reader to the multimedia video for the experimental videos of the end-to-end experiment.

frame of Fig. 8 where the robot explored the regions where collision is less likely to occur. This drives the robot to search in regions where there are no collision measurements. In the end of the experiment, the robot is successful in estimating the current configuration ($\theta_{\text{actual}} = [0.5, 0.6, -1.1]$, $\theta_{\text{estimated}} = [0.521, 0.609, -1.103]$) of the environment through the use of the data-driven measurement model of the environment.

IX. CONCLUSION

We present a method for estimating measurement models characterized by spatially sparse information through active sensing and localizing with respect to these data-driven measurement models. The ergodic control policy enables a robot to explore in sparse environments in search for new information. We are able to show that active sensing based on data-driven environmental measurement models can be used to localize within the environment taking into account the physical constraints of the robot and the sensor physics.

Through simulated and experimental examples, it is shown that collision-based sensing, coupled with our method, is sufficient for measurement model generation and later localization. In addition, using an ergodic control policy as the principle for active sensing is shown to direct the robot towards regions that exploit features in an environment discovered from data. Experimental validation of the method shows that this approach can run in real-time on a robotic system. Lastly, this method can be used for any sensor in an environment characterized by spatially sparse information by relying on motion through active sensing.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under awards IIS-1426961. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] I. Abraham, A. Prabhakar, M. J. Z. Hartmann, and T. D. Murphey. Ergodic exploration using binary sensing for non-parametric shape estimation. *IEEE Robotics and Automation Letters*, 2(2):827–834, 2017.
- [2] Hirotugu Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected Papers of Hirotugu Akaike*, pages 199–213. Springer, 1998.
- [3] M Sanjeev Arulampalam, Simon Maskell, Neil Gordon, and Tim Clapp. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.
- [4] Henrik Axelsson, Y Wardi, Magnus Egerstedt, and EI Verriest. Gradient descent approach to optimal mode scheduling in hybrid dynamical systems. *Journal of Optimization Theory and Applications*, 136(2):167–186, 2008.
- [5] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.
- [6] TM Caldwell and TD Murphey. Projection-based iterative mode scheduling for switched systems. *Nonlinear Analysis: Hybrid Systems*, 21:59–83, 2016.
- [7] George E Carvell and DJ Simons. Biometric analyses of vibrissal tactile discrimination in the rat. *The Journal of Neuroscience*, 10(8):2638–2648, 1990.
- [8] Gregory S Chirikjian and Alexander B Kyatkin. *Engineering Applications of Noncommutative Harmonic Analysis: With Emphasis on Rotation and Motion Groups*. CRC press, 2000.
- [9] Claire Dune, Eric Marchand, Christophe Colloutet, and Christophe Leroux. Active rough shape estimation of unknown objects. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3622–3627, 2008.
- [10] Magnus Egerstedt, Yorai Wardi, and Henrik Axelsson. Transition-time optimization for switched-mode dynamical systems. *IEEE Transactions on Automatic Control*, 51(1):110–115, 2006.
- [11] Hans Jacob S Feder, John J Leonard, and Christopher M Smith. Adaptive mobile robot navigation and mapping. *The International Journal of Robotics Research*, 18(7):650–668, 1999.
- [12] Valerii Fedorov. Optimal experimental design. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(5):581–589, 2010.
- [13] Charles Fox, Mat Evans, Martin Pearson, and Tony Prescott. Tactile SLAM with a biomimetic whiskered robot. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4925–4930, 2012.
- [14] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. Active markov localization for mobile robots. *Robotics and Autonomous Systems*, 25(3):195–207, 1998.
- [15] Georgios Georgakis, Arsalan Mousavian, Alexander Berg, and Jana Kosecka. Synthesizing training data for object detection in indoor scenes. In *Proceedings of Robotics: Science and Systems*, 2017. doi: 10.15607/RSS.2017.XIII.043.
- [16] E Guić-Robles, C Valdivieso, and G Guajardo. Rats can learn a roughness discrimination using only their vibrissal system. *Behavioural Brain Research*, 31(3):285–289, 1989.
- [17] Lionel Heng, Gim Hee Lee, and Marc Pollefeys. Self-calibration and visual SLAM with a multi-camera system on a micro aerial vehicle. *Autonomous Robots*, 39(3):259–277, 2015.
- [18] Jennifer A Hobbs, R Blythe Towal, and Mitra JZ Hartmann. Spatiotemporal patterns of contact across the rat vibrissal array during exploratory behavior. *Frontiers in Behavioral Neuroscience*, 9:356, 2015.
- [19] RC St John and Norman R Draper. D-optimality for regression designs: a review. *Technometrics*, 17(1):15–23, 1975.
- [20] Simon J Julier and Jeffrey K Uhlmann. A new extension of the kalman filter to nonlinear systems. In *International Symposium on Aerospace/Defense Sensing, Simulation and Controls*, volume 3, pages 182–193, 1997.
- [21] Michael C Koval, Nancy S Pollard, and Siddhartha S Srinivasa. Pose estimation for planar contact manipulation with manifold particle filters. *The International Journal of Robotics Research*, 34(7):922–945, 2015.
- [22] Andreas Krause and Carlos Guestrin. Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach. In *International Conference on Machine Learning*, pages 449–456, 2007.
- [23] Chris Kreucher, Keith Kastella, and Alfred O Hero Iii. Sensor management using an active sensing approach. *Signal Processing*, 85(3):607–624, 2005.
- [24] Chris Kreucher, John Wegrzyn, Michel Beauvais, and Ralph Conti. Multiplatform information-based sensor management: an inverted uav demonstration. In *Defense Transformation and Net-Centric Systems 2007*, volume 6578, page 65780Y, 2007.
- [25] Nathan Lepora, Uriel Martinez-Hernandez, and Tony Prescott. Active bayesian perception for simultaneous object localization and identification. In *Proceedings of Robotics: Science and Systems*, 2013. doi: 10.15607/RSS.2013.IX.019.
- [26] Uriel Martinez-Hernandez, Tony J Dodd, Mathew H Evans, Tony J Prescott, and Nathan F Lepora. Active sensorimotor control for tactile exploration. *Robotics and Autonomous Systems*, 87:15–27, 2017.
- [27] George Mathew and Igor Mezić. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena*, 240(4):432–442, 2011.
- [28] Takamitsu Matsubara, Kotaro Shibata, and Kenji Sugimoto. Active touch point selection with travel cost in tactile exploration for fast shape estimation of unknown objects. In *IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 1115–1120, 2016.
- [29] Anastasia Mavrommati, Emmanouil Tzorakoleftherakis, Ian Abraham, and Todd D Murphey. Real-time area coverage and target localization using receding-horizon ergodic exploration. *IEEE Transactions on Robotics*, 2017.
- [30] Martin Meier, Matthias Schopfer, Robert Haschke, and Helge Ritter. A probabilistic approach to tactile shape reconstruction. *IEEE Transactions on Robotics*, 27(3):630–635, 2011.
- [31] Lauren M Miller and Todd D Murphey. Trajectory optimization for continuous ergodic exploration. In *IEEE American Control Conference (ACC)*, pages 4196–4201, 2013.
- [32] Lauren M Miller and Todd D Murphey. Optimal planning for target localization and coverage using range sensing. In *IEEE International Conference on Automation Science and Engineering (CASE)*, pages 501–508, 2015.
- [33] Lauren M Miller, Yonatan Silverman, Malcolm A MacIver, and Todd D Murphey. Ergodic exploration of distributed information. *IEEE Transactions on Robotics*, 32(1):36–52, 2016.
- [34] Ben Mitchinson, Chris J Martin, Robyn A Grant, and Tony J Prescott. Feedback control in active sensing: rat exploratory whisking is modulated by environmental contact. *Proceedings of the Royal Society of London B: Biological Sciences*, 274(1613):1035–1041, 2007.
- [35] Raúl Mur-Artal and Juan D Tardós. Probabilistic semi-dense mapping from highly accurate feature-based monocular SLAM. *Proceedings of Robotics: Science and Systems*, 2015.
- [36] Evgeny Nuger and Beno Benhabib. Multicamera fusion for shape estimation and visibility analysis of unknown deforming objects. *Journal of Electronic Imaging*, 25(4):041009–041009, 2016.
- [37] Zachary Pezzementi, Erion Plaku, Caitlin Reyda, and Gregory D Hager. Tactile-object recognition from appearance information. *IEEE Transactions on Robotics*, 27(3):473–487, 2011.
- [38] John Platt. Probabilistic outputs for support vector machines

- and comparisons to regularized likelihood methods. *Advances in Large Margin Classifiers*, 10(3):61–74, 1999.
- [39] Morgan Quigley, Ken Conley, Brian P. Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y. Ng. Ros: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, 2009.
- [40] RH Rasshofer and K Gresser. Automotive radar and lidar systems for next generation driver assistance functions. *Advances in Radio Science*, 3(B. 4):205–209, 2005.
- [41] Brent Schwarz. Lidar: Mapping the world in 3d. *Nature Photonics*, 4(7):429, 2010.
- [42] Dylan A Shell, Chris V Jones, and Maja J Matarić. Ergodic dynamics by design: A route to predictable multi-robot systems. In *Multi-Robot Systems. From Swarms to Intelligent Automata Volume III*, pages 291–297. Springer, 2005.
- [43] Bernard W Silverman. *Density estimation for statistics and data analysis*. Routledge, 2018.
- [44] Rudolph Van Der Merwe, Arnaud Doucet, Nando De Freitas, and Eric A Wan. The unscented particle filter. In *Advances in Neural Information Processing Systems*, pages 584–590, 2001.
- [45] Vladimir Naumovich Vapnik and Vlamimir Vapnik. *Statistical Learning Theory*, volume 1. Wiley New York, 1998.
- [46] Ramanarayan Vasudevan, Humberto Gonzalez, Ruzena Bajcsy, and S Shankar Sastry. Consistent approximations for the optimal control of constrained switched systems—part 1: A conceptual algorithm. *SIAM Journal on Control and Optimization*, 51(6): 4463–4483, 2013.
- [47] Ting-Fan Wu, Chih-Jen Lin, and Ruby C Weng. Probability estimates for multi-class classification by pairwise coupling. *Journal of Machine Learning Research*, 5:975–1005, 2004.
- [48] Zhengkun Yi, Roberto Calandra, Filipe Veiga, Herke van Hoof, Tucker Hermans, Yilei Zhang, and Jan Peters. Active tactile object exploration with gaussian processes. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4925–4930, 2016.