

Trajectory Optimization for Continuous Ergodic Exploration

Lauren M. Miller, *Student Member, IEEE*, and Todd D. Murphey, *Member, IEEE*,

Abstract—An algorithm is presented for generating trajectories for efficient exploration that takes into account a probabilistic representation of information density over a sampling region. The problem is cast as a continuous-time trajectory optimization problem, where the objective function directly involves the relationship between the probability density functions representing the spatial distribution and the statistical representation of the time-averaged trajectory. The difference is expressed using ergodicity. It is shown that the trajectory optimization problem can be solved using descent directions that are solutions to linear quadratic optimal control problems. The proposed method generates continuous-time optimal feedback controllers, demonstrated in simulation for a nonlinear sensor model.

I. INTRODUCTION

Dynamic exploration is particularly important in the area of automated tactile sensing. In exploration tasks, there is a tradeoff between completeness of coverage and exploration time or energetic cost. Intuitively, an efficient exploration strategy should spend more time exploring regions of space with a higher information density or likelihood of feature localization. Sensing strategies that are observed to depend on sensory goals and maximize expected information gain have been widely observed in biological systems [1]. For example, observations of human subjects during feature localization tasks imply that sensing strategies reflect the sensing goal [2] and a subject's internal expectation of the feature location [3]. Similar observations have been made during exploration experiments featuring the rat vibrissal system [4].

Motivated by observations of exploratory behavior in biological systems, the objective of this paper is to define a strategy for planning exploratory motion to sample a given region in an efficient manner with respect to the spatial sensory information density over that region. The planning problem is posed as an infinite-dimensional trajectory optimization, where the objective function involves a metric on the difference between the time-averaged behavior of the sensor trajectory and a spatial probability density function (PDF) representing the information density. The metric relies on the measure of ergodicity proposed by Mathew and Mezić [5], discussed in Section III. We show that iterative first order optimization can be used to minimize the ergodic metric, using the projection-based trajectory optimization method derived in [6], which allows the descent direction at each iteration to be calculated using linear quadratic regulator techniques. Using trajectory optimization allows calculation

of an optimally ergodic trajectory over a continuous exploration time horizon, which can be applied to receding horizon or fixed-time control scenarios.

II. RELATED WORK

General research in planning for mobile sensors has received much attention in the context of nonuniform coverage for distributed control. For example, optimal placement of stationary sensors [7], [8] and positioning of mobile sensors [9] have been analyzed in the context of distributed sensor networks. The focus of these strategies are efficient algorithms for distributed control and generally do not directly take into account nonlinear sensor dynamics.

In the context of tactile sensing and object recognition, several strategies have been derived for choosing between sensing strategies or movements based on the sensing goal or likelihood of information gain [10]–[13]. In [10], for example, a method is presented to determine the best straight sensor path that is guaranteed to produce data unique to only one possible object. While these techniques present different methods of selecting sensing strategies or choosing from a set of parameterized movements, the trajectory planning problem for continuous exploration is not addressed.

The optimal control problem we develop depends on a metric defined in [5] that relates the statistical properties of a spatial distribution with those of a time-averaged trajectory. In addition to defining the metric, the authors of [5] use the metric to develop an optimization-based feedback controller for uniform sampling of irregular domains for first- and second-order linear systems. While our objective function involves the same metric derived in [5], the optimal control problem and applications are fundamentally different. The strategy in [5] involves discretizing the exploration time and solving for the optimal control input at each time-step which maximizes the rate of decrease of the ergodic metric at that time. The method we derive calculates the optimal control for minimizing the metric over the entire exploration time horizon. Additionally, the feedback controllers derived in [5] are specific to linear first- or second-order integrator systems, whereas our method is derived for general nonlinear dynamic systems. While the continuous-time optimization is more computationally expensive than the discrete time approach in [5], the linear quadratic formulation has the advantage that each iteration of the optimization involves a calculation with known complexity, and all computational cost happens offline.

In prior work, we derived a method for determining contact transition times [14] which also uses the ergodic metric presented in [5]. Although the metric is the same, the

The authors are with the Department of Mechanical Engineering, Northwestern University, Evanston, IL, 60208 USA (e-mail: LaurenMiller@U.Northwestern.edu, T-Murphey@Northwestern.Edu)

control problem in [14] is a finite dimensional switching-time problem.

The contributions of this paper are threefold. Using the ergodicity metric derived in [5] as the objective function for a receding horizon control problem, an iterative descent algorithm is provided. The algorithm is formulated for general nonlinear dynamics, and provides a continuous-time optimal feedback control solution. Although the ergodic optimization problem is both nonlinear and not a Bolza problem, it is shown here that the descent direction at every step of the optimization can be calculated as the solution to a linear quadratic problem. This approach has the benefit of providing a state feedback law, and each iteration has known computational complexity.

In Section III, the concept of ergodicity is introduced, and the metric derived in [5] is defined. Section IV summarizes trajectory optimization methods and describes the optimal control problem formulation. A quadratic model is presented for a steepest descent algorithm, involving derivation of first-order optimality condition for the objective function in Section IV-E. Section V demonstrates the use of the algorithm to generate optimally ergodic trajectories over two-dimensional PDFs in simulation.

III. ERGODICITY

As mentioned in Section I, the aggregate motion of a sensor should result in more samples from the most information-dense areas of a distribution. This concept is quantified using the measure of ergodicity presented in [5]. In order for a system to be ergodic, the fraction of time spent sampling an area should be equal to some metric quantifying the density of information in that area. For a given task, the spatial distribution used for calculating ergodicity will vary. For example, the distribution might be the variance

of some surface parameter over the surface manifold or the probability that some feature is located at any point in the domain.

This concept is expressed in Fig. 1, where the distribution $\phi(x)$, depicted as level sets over the domain X , is sampled by a sensor following the trajectory $x(t)$ from $t = 0$ to $t = T$. The trajectory $x(t)$ is *ergodic* with respect to the PDF $\phi(x)$ if the percentage of time spent in any subset N of X from $t = 0$ to $t = T$ is equal to the measure of N . The equations in Fig. 1 represent the condition for ergodicity for the two subsets shown; this condition must hold for all possible subsets.

A. Metric for Ergodicity

The metric presented in [5] will be used to calculate how far a trajectory is from being ergodic with respect to a distribution. The distance is quantified by sum of the weighted norm of the coefficients of the Fourier series decomposition of the multidimensional spatial and time-averaged distributions. For a more formal discussion, see [5].

In this work, calculation of the optimal control solution is done over a fixed time horizon $[0, T]$. It is assumed that there is a probability distribution function $\phi(x)$ representing information density over an n -dimensional exploration domain $X \subset \mathbb{R}^n$ defined as $[0, L_1] \times [0, L_2] \dots \times [0, L_n]$.

The ergodicity of a trajectory $x(t)$ with respect to a distribution ϕ can be quantified by the sum of the weighted squared distance between the Fourier coefficients of the spatial distribution Φ_k and the distribution representing the time-averaged trajectory c_k . The ergodic metric will be defined as \mathcal{E} , as follows:

$$\mathcal{E} = \sum_{k=0}^K \Lambda_k |c_k - \phi_k|^2 \quad (1)$$

where K is the number of basis functions used in each dimension. Following [5], $\Lambda_k = \frac{1}{(1+|k|^{2s})}$ and $s = \frac{n+1}{2}$, which places larger weight of on lower frequency information.

The Fourier coefficients ϕ_k of a spatial distribution $\phi(x)$ are computed using an inner product as follows

$$\phi_k = \int_X \phi(x) F_k(x) dx,$$

and the Fourier coefficients of the basis functions along a trajectory $x(t)$, averaged over time, are calculated as

$$c_k = \frac{1}{T} \int_0^T F_k(x(t)) dt,$$

where T is the final time and F_k is a Fourier basis function, as derived in [5].

The Fourier basis functions we will use to approximate a distribution defined over n dimensions are

$$F_k(\bar{x}) = \frac{1}{h_k} \prod_{i=1}^n \cos\left(\frac{k\pi}{L_i} \bar{x}_i\right), \text{ for } k = 0, 1, 2 \dots K$$

where h_k is a normalizing factor [5]. It should be noted, however, that any set of basis functions that is differentiable

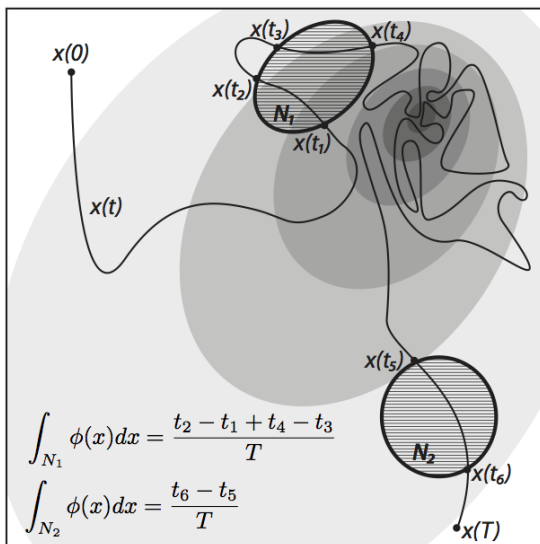


Fig. 1: Conceptual illustration of what it means for the trajectory $x(t)$ to be ergodic with respect to the distribution $\phi(x)$, represented by the level sets shown. Equations representing the condition for ergodicity for the two subsets, N_1 and N_2 are shown [14]

in the state and can be evaluated along the trajectory can be used in both the metric in Eq. (1) and the trajectory optimization algorithm we derive.

The following sections present the use of Eq. (1) as an objective function in a continuous-time trajectory optimization framework, used to solve general, nonlinear optimal control problems.

IV. TRAJECTORY OPTIMIZATION

The problem of generating a continuous-time optimally ergodic trajectory is formulated using the projection-based trajectory optimization method presented in [6]. Using the projection-based method allows us to define a local quadratic model of the ergodic objective function, which can then be used to calculate the steepest descent direction for use in iterative first-order optimization methods, using linear quadratic techniques [6].

The following sections define the equations for the system dynamics and the ergodic objective function which will be minimized, as well as a summary of projection-based trajectory optimization.

A. Dynamics

The dynamics of a general nonlinear dynamic sensor can be expressed as follows:

$$\dot{x}(t) = f(x(t), u(t)) \quad x(t_0) = x_0, \quad (2)$$

where $x(t) \in \mathbb{R}^n$ represents the state and $u(t) \in \mathbb{R}^m$ the control inputs. For notational convenience, define $\eta(t) = (x(t), u(t))$, as in [6], as a feasible curve which lives in the trajectory manifold \mathcal{T} (curves satisfying Eq. (2)).

B. Objective Function

The objective function $J(\cdot)$ will be defined as a function of the metric for ergodicity in Eq. (1) plus the integrated magnitude of the control, which takes as an argument any (feasible or infeasible) curve $\xi(t) = (\alpha(t), \mu(t))$,

$$J(\xi(t)) = q \sum_{k=0}^K \Lambda_k \left(\frac{1}{T} \int_0^T F_k(\alpha(\tau)) d\tau - \phi_k \right)^2 + \int_0^T \frac{1}{2} \mu(\tau) R(\tau) \mu(\tau) d\tau, \quad (3)$$

where $q \in \mathbb{R}$ and $R(\tau) \in \mathbb{R}^{m \times m}$ are arbitrary design parameters which affect the relative importance of minimizing ergodicity vs. control effort in the optimization problem.

The goal is to solve for a feasible continuous time trajectory which minimizes the objective function, i.e.

$$\arg \min_{\xi(t) \in \mathcal{T}} J(\xi(t)). \quad (4)$$

The optimization problem in Eq. (4) is subject to a nonlinear constraint, and is not written in the form of a Bolza problem, both factors which make the optimization problem as written nontrivial. However, we use the projection operator defined in Section IV-C to allow calculations at each iteration of the optimization to occur in the (infinite dimensional) tangent space of the constraint, and define a quadratic model for Eq.

(3) in Section IV-D which is a Bolza problem and linear quadratic in the descent direction.

C. Projection-based Trajectory Optimization

The optimization in Eq. (4) can be reformulated as an unconstrained trajectory optimization problem using a projection operator following [6]. A projection operator takes the form of a stabilizing feedback control law which maps any (feasible or infeasible) trajectory $\xi(t) = (\alpha(t), \mu(t))$ to a feasible trajectory. The projection operator used in this paper is

$$P(\xi(t)) : \begin{cases} u(t) = \mu(t) + K(t)(\alpha(t) - x(t)) \\ \dot{x}(t) = f(x(t), u(t)) \quad x(0) = x_0. \end{cases} \quad (5)$$

The optimal feedback gain $K(t)$ is calculated by solving an additional linear quadratic regulator problem.

Using the projection operator allows the optimization problem to be reformulated as an equivalent, unconstrained optimization problem, where the goal is to minimize $J(P(\xi(t)))$, i.e.

$$\arg \min_{\xi(t)} J(P(\xi(t))). \quad (6)$$

Use of the projection operator has the advantage of removing the nonlinear constraint imposed by the dynamics during the descent direction search at each iteration.

D. Optimization Algorithm

In order to use first-order iterative optimization methods such as steepest descent, a descent direction $\zeta_i(t)$ must be calculated at every iteration i of the optimization algorithm. The steepest descent direction is obtained by minimizing a quadratic model of the form

$$\zeta_i(t) = \arg \min_{\zeta_i(t) \in T_{\xi_i} \mathcal{T}} DJ(\mathcal{P}(\xi_i(t))) \circ \zeta_i(t) + \frac{1}{2} \langle \zeta_i(t), \zeta_i(t) \rangle. \quad (7)$$

where $T_{\xi_i} \mathcal{T}$ is the tangent space of the trajectory manifold. It is shown in [6] that solving Eq. (7), where $\zeta_i(t)$ is constrained to lie on the tangent space of the trajectory manifold, provides a descent direction for the optimal control and that $P(\xi_i + \gamma_i \zeta_i)$ (where γ_i is the step size) is both feasible and satisfies the sufficient decrease condition.

Algorithm 1 Steepest Descent for Ergodic Trajectory Optimization

Calc. ϕ_k
 Init. $\xi_0 \in \mathcal{T}$, tolerance ϵ
while $DJ(\xi_i) \circ \zeta_i > \epsilon$ **do**
 Calculate descent direction:
 $\zeta_i = \arg \min_{\zeta_i \in T_{\xi_i} \mathcal{T}} DJ(\xi_i) \circ \zeta_i + \frac{1}{2} \langle \zeta_i, \zeta_i \rangle$
 Calculate γ_i using Armijo line search [15]
 Project the update:
 $\xi_{i+1} = \mathcal{P}(\xi_i + \gamma_i \zeta_i)$
 $i=i+1$
end while

The basic algorithm is outlined in Algorithm 1. The descent direction search occurs in the tangent space of the trajectory manifold at the current iteration. As shown in the next section, the optimization in Eq. (7) involves only linear and quadratic terms in ζ_i , which can be solved using linear quadratic techniques [6]. An Armijo line search is used to calculate the step size for each descent direction, ensuring sufficient decrease in the objective function [15]. After direction and step size have been calculated, the update step projects the new iterate onto the trajectory manifold.

E. Defining the Linear Quadratic Problem

In this section, we will assume that the objective function can be written as a function of a feasible trajectory $\eta(t)$, as follows

$$J(\eta(t)) = q \sum_{k=0}^K \Lambda_k \left(\frac{1}{T} \int_0^T F_k(x(\tau)) d\tau - \phi_k \right)^2 + \int_0^T \frac{1}{2} u(\tau) R(\tau) u(\tau). \quad (8)$$

This assumption can be made because each update is projected onto the feasible space, therefore each iteration begins on the trajectory manifold.

First, we show that directional derivative in the first term of Eq. (7) is linear in the descent direction $\zeta(t)$. Second, we define the inner product in the second term of Eq. 7, and show that the resulting expression can be rearranged as a Bolza problem. Third, the linearization of the dynamic constraint is provided. These three definitions result in a model that is linear and quadratic in the descent direction, which can be solved using linear quadratic techniques [6].

1) *Linearization of $DJ(\eta(t)) \circ \zeta(t)$* : Equation (8) is linearized by taking the directional derivative of $J(\eta(t))$ in the direction $\zeta(t)$, where $\zeta(t) = (z(t), v(t))$, resulting in

$$DJ(\eta(t)) \circ \zeta(t) = q \sum_{k=0}^K \Lambda_k \left[2 \left(\int_0^T \frac{1}{T} F_k(x(\sigma)) d\sigma - \phi_k \right) \circ \int_0^T \frac{1}{T} DF_k(x(\tau)) \circ z(\tau) d\tau \right] + \int_0^T R(\tau) u(\tau) \circ v(\tau) d\tau.$$

This expression can be rearranged, pulling the expression in parentheses into the second integral over τ , and switching the order of the integral and summation, to the following:

$$DJ(\eta(t)) \circ \zeta(t) = \int_0^T q \sum_{k=0}^K \Lambda_k \left[2 \left(\int_0^T \frac{1}{T} F_k(x(\sigma)) d\sigma - \phi_k \right) \circ \frac{1}{T} DF_k(x(\tau)) \right] \circ z(\tau) + R(\tau) u(\tau) \circ v(\tau) d\tau.$$

$$\text{Defining } a^T(\tau) = q \sum_{k=0}^K \Lambda_k \left[2 \left(\int_0^T \frac{1}{T} F_k(x(\sigma)) d\sigma - \phi_k \right) \circ \frac{1}{T} DF_k(x(\tau)) \right],$$

and $b^T(\tau) = R(\tau)u(\tau)$, $DJ(\eta(t)) \circ \zeta(t)$ can be written

$$DJ(\eta(t)) \circ \zeta(t) = \int_0^T a^T(\tau) \circ z(\tau) + b^T(\tau) \circ v(\tau) d\tau,$$

which is linear in $\zeta(t) = (z(t), v(t))$.

2) *Definition of $\frac{1}{2}\langle \zeta, \zeta \rangle$* : The second term in Eq. (8) can be defined as

$$\frac{1}{2}\langle \zeta, \zeta \rangle = \int_0^T \left[\frac{1}{2} z(\tau)^T Q_n(\tau) z(\tau) + \frac{1}{2} v(\tau) R_n(\tau) v(\tau) \right] d\tau + \frac{1}{2} z(T)^T P_{1n} z(T)$$

Where Q_n and P_{1n} are arbitrary positive semi-definite matrices and R_n is positive definite. See [6] for more information.

3) *Linearization of the Constraint*: Because the descent direction $\zeta_i(t)$ is constrained to lie in the tangent space of the trajectory manifold, it must satisfy the differential equation

$$\dot{z}_i(t) = A(t)z_i(t) + B(t)v_i(t),$$

where $A(t) = D_1 f(x_i(t), u_i(t))$ and $B(t) = D_2 f(x_i(t), u_i(t))$, a linear equation in $\zeta_i(t)$.

The linear quadratic formulation of Eq. (7) is therefore

$$\arg \min_{\xi} \int_0^T a^T z + b^T v + \frac{1}{2} z(\tau)^T Q_n(\tau) z(\tau) + \frac{1}{2} v(\tau) R_n(\tau) v(\tau) d\tau + \frac{1}{2} z(T)^T P_{1n} z(T),$$

subject to $\dot{z}(t) = A(t)z(t) + B(t)v(t)$, $z(0) = z_0$.

The problem reduces to solving a linear quadratic optimal control problem for the descent direction at each iteration. This solution is found using standard Riccati differential equations [6], [16].

V. SIMULATED EXAMPLE

In order to demonstrate the proposed ergodic trajectory optimization for a nonlinear, dynamic system, a simulation was carried out for a sensor traversing two different PDFs in two dimensional space, using a simple kinematic model for the sensor dynamics. The state for this model is $x(t) = (X(t), Y(t), \theta(t))$ where $X(t)$ and $Y(t)$ are cartesian coordinates and $\theta(t)$ is the heading angle. The control is $u(t) = (v(t), \omega(t))$ where $v(t)$ is forward velocity and $\omega(t)$ is the angular velocity. The dynamics of the model are

$$f(x) = \begin{bmatrix} \cos(x_3(t)) & 0 \\ \sin(x_3(t)) & 0 \\ 0 & 1 \end{bmatrix} \cdot u(t). \quad (9)$$

Figure 2 shows an optimally ergodic trajectory for the dynamics in Eq. 9 with respect to a bimodal Gaussian PDF, shown as level sets in Fig. 2. For the trajectory shown in Figure 2, the optimization was initialized to a figure-eight-like pattern, chosen as an initialization which approximates uniform coverage. The trajectory used as the initialization






Fig. 2: Initial and optimal trajectories for a bimodal PDF and the dynamics of Eq. 9. Grayscale contours represent the spatial distribution; the initial trajectory is plotted in black, the optimized trajectory in white.

is shown as a dotted black line. The optimized trajectory is plotted in white, using the termination criteria that $\|DJ(\cdot) \circ \zeta\|$ converges to zero within a tolerance of 10^{-3} .


Qualitatively, whereas the initial trajectory explores the domain approximately uniformly, the optimized trajectory spends more time near the peaks of the PDF. Figure 3 compares the PDF representing the spatial distribution, and the PDFs representing the time-averaged trajectories for the initial trajectory in Fig. 3(b), and the optimized trajectory in



(a) 3D Plot of the spatial PDF shown in Fig. 2



(b) PDF representing time-averaged initial trajectory (dotted black line in Fig. 2)



(c) PDF representing time-averaged optimized trajectory (white line in Fig. 2)

Fig. 3




Fig. 4: Initial (black) and optimal (white) trajectories over the same PDF as in Fig. 2, for two different initial trajectories.




Fig. 5: Initial (black) and optimal (white) trajectories (white) shown over an alternate PDF for two initializations.

Fig. 3(c). After applying the optimization, the PDF in Fig. 3(c) closely matches the spatial PDF in Fig. 3(a).

The optimization was applied over the same distribution for two additional initial trajectories. Figure 4 shows two alternate locally optimally ergodic trajectories (plotted in white) over the same PDF. The optimizations in Fig. 4 were initialized with different initial conditions for the dynamic equations and different initial trajectories for the optimization, once again plotted in black. The objective function derived only requires that the time averaged statistics of the trajectory converge towards the spatial statistics; depending on what initialization is used for the optimization, the trajectory may therefore look very different, although the time-averaged statistics are similar.

The algorithm was also applied over a different spatial distribution using two different initializations, shown in Fig 5. Once again, the optimized trajectories are different for different initializations, although both spend a greater percentage of time in areas of the PDF where the density is highest.

Sensitivity of the optimization method to the initialization is likely the result of several factors. Because a finite number of basis functions were used to represent spatial and time-averaged PDFs, some inherent smoothing of the time-averaged data occurs. The loss of resolution results in similar




Fig. 6: Plot of the decay of the ergodic metric from Eq. 1. The legend corresponds to which figure shows the optimal trajectory the line corresponds to; Fig. 2, Fig. 4(a),(b) or Fig. 5(a),(b).

approximation of different trajectories for a finite number of basis functions. This could be remedied by using larger numbers of basis functions, however increasing the number of basis functions increases computational cost. Analysis of the tradeoff between computational cost and benefits of finer resolution will be left to future work, however Fig. 6 demonstrates that for the number of basis functions used here, the trajectories all provide similarly optimally ergodic solutions.

As mentioned, the optimal trajectories for all five initializations over both PDFs satisfy the necessary conditions of optimality, i.e. $\|DJ(\cdot) \circ \zeta\|$ converges to zero within the chosen tolerance; the trajectories are all local minima of the objective function in neighborhoods around the initial trajectories. The decay of the ergodic metric \mathcal{E} with each iteration of the optimization is shown in Fig. 6 for all five trajectories. The value of the metric decays very quickly, with 99% of the benefit of the continuous time algorithm realized after fewer than 10 iterations of the. Therefore, while this approach is more computationally expensive per iteration compared to greedy approaches, such as the controller designed in [5], a small number of iterations results in significant improvement of the ergodicity.

VI. CONCLUSION

In this paper, a method is presented for determining optimal exploration trajectories with respect to a given information density. The problem is cast as an infinite-dimensional trajectory optimization, where the objective function is defined by the distance from ergodicity between a spatial PDF, representative of the information density over the exploration region and the time-averaged behavior of the continuous trajectory. The ergodic metric presented in [5] is used to formulate an iterative descent algorithm is derived for calculating a receding horizon control strategy. The algorithm is formulated for general nonlinear dynamics, and provides a continuous-time optimal control solution. The method is demonstrated for a kinematic model of sensor dynamics, over two different spatial PDFs from different initial conditions.

Simulations demonstrate that a small number of iterations result in dramatic decrease in the objective function.

Ultimately, this method has potential applications in a multi-scan exploration or detection algorithm. Future work will involve applying this method to experimental sensor models. Additionally, second-order optimality conditions will be derived in order to increase convergence rates. One of the limitations of this approach is that the sampling time-horizon is fixed. Analysis of the effects of the length of the time horizon and the properties of the resulting trajectory is planned for future work.

ACKNOWLEDGMENT

This material is based upon work supported by the National Institute of Health under grant T32 HD007418 and the National Science Foundation under Grant IIS 1018167. Any opinions, findings, and conclusions expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or the NIH.

REFERENCES

- [1] T. Prescott, M. Diamond, and A. Wing, "Active touch sensing," *Philosophical transactions of the royal society biological sciences*, vol. 366, no. 1581, pp. 2989–2995, Nov 2011.
- [2] R. Klatzky, J. Loomis, S. Lederman, H. Wake, and N. Fujita, "Haptic identification of objects and their depictions," *Attention, Perception, and Psychophysics*, vol. 54, pp. 170–178, 1993.
- [3] K. Huynh, C. S. L.W., White, J. Colgate, and Y. Matsuoka, "Finding a feature on a 3d object through single-digit haptic exploration," in *IEEE Haptics Symposium*, March 2010, pp. 83–89.
- [4] R. Grant, B. Mitchinson, C. Fox, and T. Prescott, "Active touch sensing in the rat: Anticipatory and regulatory control of whisker movements during surface exploration," *J Neurophysiol*, vol. 101, pp. 862–874, Feb 2009.
- [5] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D-nonlinear Phenomena*, vol. 240, no. 4-5, pp. 432–442, Feb 2011.
- [6] J. Hauser, "A projection operator approach to the optimization of trajectory functionals," in *IFAC world congress*, vol. 4, 2002, pp. 3428–3434.
- [7] A. Howard, M. Mataric, and G. Sukhatme, "Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem," in *Int. Symposium on Distributed Autonomous Robotics Systems (DARS)*, 2002, pp. 299–308.
- [8] F. Lekien and N. Leonard, "Nonuniform coverage and cartograms," *SIAM J. Control Optim.*, vol. 48, no. 1, pp. 351–372, Feb 2009.
- [9] J. Cortes, S. Martinez, T. K. F., and Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, April 2004.
- [10] J. Schneiter and T. Sheridan, "An automated tactile sensing strategy for planar object recognition and localization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 8, pp. 775–786, Aug 1990.
- [11] K. Roberts, "Robot active touch exploration: constraints and strategies," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, vol. 2, May 1990, pp. 980–985.
- [12] S. Hutchinson and A. Kak, "Planning sensing strategies in a robot work cell with multi-sensor capabilities," *IEEE Transactions on Robotics*, vol. 5, no. 6, pp. 765–783, Dec 1989.
- [13] S. Kristensen, "Sensor planning with bayesian decision theory," in *In Reasoning With Uncertainty In Robotics*, 1995, pp. 273–286.
- [14] L. M. Miller and T. D. Murphey, "Optimal contact decisions for ergodic exploration," in *IEEE Int. Conf. on Decision and Control (CDC)*, 2012.
- [15] L. Armijo, "Minimization of functions having lipschitz continuous first partial derivatives," *Pacific Journal of Mathematics*, vol. 16, no. 1, pp. 1–3, 1966.
- [16] B. Anderson and J. Moore, *Optimal Control: Linear Quadratic Methods*. Prentice-Hall, 1990.