Trajectory Optimization for Continuous Ergodic Exploration on the Motion Group SE(2)

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

***Abstract*— Autonomous active sensing presents the need for control of sensor motion in both position and orientation. This paper presents a method of planning continuous search trajectories over the Euclidean motion group $SE(2)$. The method allows one to calculate an optimal search trajectory with respect to a distribution representing probable information gain over the search space. The ergodicity of a trajectory with respect to the information distribution is used to formulate an objective function for a projection-based optimal control strategy. Results from previous work on ergodic trajectory optimization are extended to consider trajectories and information densities defined on $SE(2)$.**

I. Introduction

We consider the problem of efficient planning for active exploration. For both biological and robotic systems, exhaustive exploration of the entire search space is impractical, as there is the tradeoff between coverage and energy or time expenditure. This tradeoff cannot be ignored for realistic applications. Additionally, complete coverage in terms of both sensor position and orientation may not be possible for a physical, dynamically constrained mobile sensor system. We therefore develop a method of generating continuous, dynamically constrained search strategies which take advantage of probabilistic knowledge of the most information-dense regions of the search space. In previous work, we developed a method for trajectory planning with respect to probabilistic search tasks [1]. In this paper, we extend the approach to develop a method of generating optimal control strategies for planar exploration tasks in which sensor position and orientation are considered on the motion group $SE(2)$.

Autonomous active search and exploration relies on development of efficient control strategies for sensor movement. For a wide range of applications, either the available sensing modality or the sensing task make control of both sensor position and orientation important. Successful docking for underwater autonomous vehicles, for example, involves localizing the dock and determining the orientation of the dock [2]. Sensing modalities typically used for docking in low-light conditions, such as electromagnetic homing or optical sensing, are typically short range and orientation dependent; the ship must be close to the dock and correctly oriented to acquire the information needed to dock successfully [3].

As a biological example, the primary sensing modality used by the South American weakly electric fish is similarly position and orientation dependent [4]. The fish must be sufficiently close and oriented correctly with respect to an object's geometry in order to successfully localize or identify it by sensing disturbances in a self-generated electric field. Both position and orientation dependence also factor into vision-based imaging and object recognition, in which some views provide better information than others [5], [6], as well as tactile feature recognition, which is often sensitive to the orientation of a feature with respect to the sensor motion [7]. These examples are a subset of exploration tasks for which the choice of control strategy should directly consider both position and orientation.

The approach we use involves iteratively solving for a continuous trajectory which maximizes the amount of information gained over a finite time horizon. We assume that there is a previously defined measure representing the distribution of information for a given search task over the search space, a bounded domain on $SE(2)$. There are many potential choices of an information measure. We leave the determination of the optimal measure for future work, and consider a general probabilistic representation of information at all points in the search domain, given prior knowledge or integration of previous measurements. It should be noted that the formulation of the control algorithm we present is general to the choice of information measure over the search space.

The trajectory planning method presented here will ultimately be integrated into an iterative closed loop search algorithm, e.g. 1) plan exploration trajectory based on probabilistic distribution of information over the search space, 2) execute trajectory and collect data, and 3) process data, update probabilistic representation of information, repeat.

A. Related Work

The question we consider is how to most efficiently use knowledge of the distribution of information to plan and execute the movement necessary for search. A broad category of search methods employ greedy algorithms, where the optimal sensor movement is incrementally calculated, choosing a motion which locally maximizes a search metric [8]–[10]. These methods, while computationally simple, do not take advantage of sensor dynamics or the global shape of the information density.

Control for active exploration is also closely related to the field of simultaneous localization and mapping (SLAM). While a large percentage of SLAM research is focused on environmental or dynamic uncertainty, a subset also considers action uncertainty [11]–[13]. The technique we present in this paper is related to the idea of action uncertainty, where

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.

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 action is chosen to maximize a measure of predicted information gain. Most closely related to our approach are active SLAM algorithms which optimize information gain over a sequence of control actions, as opposed to a single action, typically with respect to both information gain and the cost of exploration [11], [13]–[15]. Optimization in these contexts is typically combinatorial. Therefore only short time horizons are considered, and the set of control choices is limited [11], [13], [16]. Expanding the search and control action space over $SE(2)$ would increase these effects. We therefore present an alternate method of computing a trajectory which balances control effort and optimal information gain in the continuous space of control. This approach does not require discretization of the search domain or control actions, and the optimization does not suffer from combinatorial complexity. The method generalizes to $SE(2)$, and the optimized trajectories are guaranteed to be feasible.

B. Ergodicity

Our approach involves using projection-based trajectory optimization to solve for the continuous optimal control with respect a PDF representing information density over the search space. In order to formulate the objective for the optimization, we use the concept of ergodicity to develop a metric relating the finite-time horizon sensor trajectory to the spatial information density.

A trajectory that evolves in time is *ergodic* with respect to a spatial PDF if the percentage of time spent by the trajectory in any subset of the spatial domain is equal to the measure of that subset. This is demonstrated in Fig. 1 for the distribution $\phi(x)$, depicted as level sets over the domain X, and 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 subsets [17].

The use of ergodicity in the context of search or exploration is quite intuitive; given that there are regions of the search space which are more likely to contain important information, an efficient search trajectory should spend more time in those areas. We use the *distance from ergodicity* of the time-averaged trajectory from the spatial information PDF as a metric to be minimized.

Note that we are not maximizing the information over the trajectory, which would result in behavior along the lines of travel directly to the state where the PDF is maximized and remaining there (i.e. maximizing the integral of the PDF along the trajectory). This is certainly an alternative approach to solving this type of problem, however we argue that using an ergodic metric, which results in a trajectory that does not maximize the information acquired but rather distributes the information according to the PDF, allows effective search strategies in scenarios where an information-maximizing approach is likely to fail, e.g. when the variance of the PDF is high or when the PDF is multimodal. Our recent experimental work utilizing ergodic control in $\mathbb{R}^1$ during one dimensional search demonstrates performance increase compared to locally information maximizing strategies, and similar comparison in $\mathbb{R}^2$ and $SE(2)$ is planned [18].

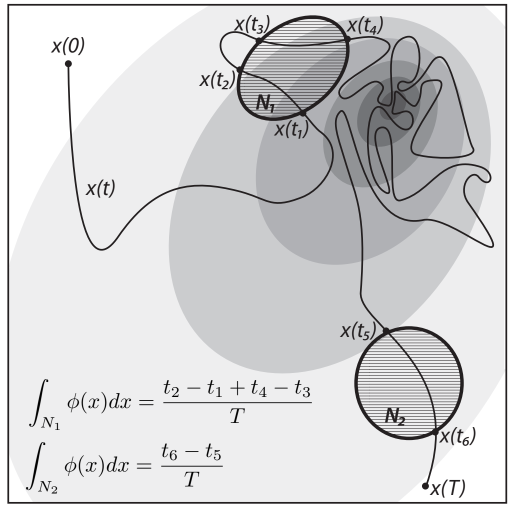

Fig. 1: Conceptual illustration of what it means for the trajectory $x(t)$ to be ergodic with respect to the distribution $\phi(x)$, represented by the level sets shown. Equations representing the condition for ergodicity for the two subsets, N_1 and N_2 are shown [17].

The use of an ergodic metric for determining optimal control strategies was originally presented in [19] for a nonuniform coverage problem. A greedy approach is used to solve for a control input at each point in time that maximizes the decrease of the metric at each timestep. A similar method is employed in [20], using a Mix Norm for coverage on Riemannian manifolds. Our previous work involved using the same metric used in [19] in the context of continuous trajectory optimization for a search problem in $\mathbb{R}^2$. In the current work, we extend this method to calculate optimal control solutions on the motion group $SE(2)$, which is critical in applications where the sensing task or sensor dynamics are both position and orientation dependent.

C. Trajectory Optimization

Optimal control using trajectory optimization methods has several benefits for the type of sensing task we are interested in. While the idea of optimal coverage with respect to a distribution has received a lot of attention in the distributed robotics community for static coverage [21], [22], we are interested in optimization over a continuous trajectory, i.e. dynamic coverage for a single robot. Trajectory optimization simultaneously outputs both the optimal path, with respect to a given objective function, and the control necessary to achieve this path; path planning and control generation are simultaneous. The dynamics are therefore explicitly considered as a constraint in the trajectory optimization, and a weighted term on the control effort is included. Thus, the optimal paths generated are feasible and the method can be adapted to simultaneously minimize control effort and optimize search performance, critical for an active sensing system.

II. Problem Definition

We consider planar dynamic systems which are orientation dependent. In particular, systems where the configuration space is curves in the special Euclidean group $SE(2)$. We call the position component of the configuration $q = (X, Y, \theta)$, and elements of the group $g = (R_\theta, \boldsymbol{x}) \in SE(2)$, where $\boldsymbol{x} \in \mathbb{R}^2 = (X, Y)$ represents translation and $R_\theta \in SO(2)$ is a rotation parameterized by the angle θ.

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 $\Phi(\boldsymbol{q})$ representing information density over the bounded exploration domain $X \subset SE(2)$.

III. Metric for Ergodicity

The ergodicity of a trajectory $q(t)$ with respect to a distribution $\Phi(q)$ can be quantified by the sum of a weighted norm on the difference between the Fourier coefficients, $\phi_{m,n,p}$, of the spatial distribution and those of the distribution representing the time-averaged trajectory $c_{m,n,p}$ where m, n, p are the indices of the coefficients along each dimension.

The ergodic metric will be defined as $\mathcal{E}(q(t))$, as follows:

$$\mathcal{E}(q(t)) = \sum_{m,n,p=0}^{M,N,P} \Lambda_{m,n,p} ||c_{m,n,p}(q(t)) - \phi_{m,n,p}||^2 \quad (1)$$

where M, N, P are the numbers of basis functions used in each dimension and $\Lambda_{m,n,p}$ is a weighting factor which places larger weight of on lower frequency information [19].

Note that the notion of ergodicity here does not strictly necessitate the use of Fourier basis functions in constructing an objective function. The basis functions provide a computationally tractable way of quantifying the spatial statistics of a trajectory, allowing us to ensure the distribution of time spent exploring the search domain is proportional to the PDF.

A. Fourier Coefficients on $SE(2)$

Fourier coefficients are essentially the Fourier transform $\hat{f}(s)$ of a function $f(q)$ sampled on a grid in the frequency domain. We use the definition of the Fourier transform on $SE(2)$ defined in [23].

The transform involves defining a unitary representation of $SE(2)$. This unitary operator can then be expressed in the standard Fourier orthonormal basis functions using an inner product. We will write the transform, parameterized by m, n, p, as

$$\hat{f}(m, n, p) = \int_{SE(2)} f(g) U(g^{-1}, m, n, p) d(g),$$

where

$$U(g^{-1}(r, \phi, \theta), m, n, p) = i^{n-m} \exp^{i[m\phi + (n-m)\theta]} J_{m-n}(pr).$$

$d(g)$ is the volume element on SE(2) and J_{m-n} is the $m - nth$ order Bessel function [23]. Note that this representation is in terms of elements of $SE(2)$ in polar coordinates, i.e.

$$g(r, \phi, \theta) = \begin{pmatrix} \cos\theta & -\sin\theta & r\cos\phi \\ \sin\theta & \cos\theta & r\sin\phi \\ 0 & 0 & 1. \end{pmatrix}$$

Since q can be redefined in terms of polar coordinates or (2) can be modified for Cartesian coordinates, we write the transform moving forward in terms of q for simplicity.

The coefficients for a function $f(q)$ can therefore be calculated by evaluating the transform at (m, n, p). We will use $F_{m,n,p}(q)$, as the notation for basis functions of index (m, n, p), i.e.

$$F_{m,n,p}(q) = U(g^{-1}(q), m, n, p).$$

In general, use of a higher number of basis functions will result in a more accurate representation of the PDF or time-averaged trajectory, however the computational cost of calculating the coefficients also increases.

The Fourier coefficients for a spatial distribution $\Phi(q)$ can be computed using an inner product of the function $\Phi(q)$ and the set of orthonormal basis functions on $SE(2)$ as follows

$$\phi_{m.n.p} = \int_X \Phi(q) F_{m,n,p}(q) dq. \quad (2)$$

The Fourier coefficients of the basis functions along a trajectory $q(t)$, averaged over time, are calculated by evaluating the basis along $q(t)$,

$$c_{m,n,p}(q(t)) = \frac{1}{T} \int_0^T F_{m,n,p}(q(t)) dt. \quad (3)$$

Previous implementations of the ergodic metric [1], [19] use basis functions that return real results. The standard transform over $SE(2)$ may return imaginary results, so some care must be taken when choosing the norm used in Eq. (1), as the norm must be differentiable. We will use the standard Euclidean norm for complex numbers on the difference of the coefficients, i.e.

$$\sqrt{\text{Re}(c_{m,n,p} - \mu_{m,n,p})^2 + \text{Im}(c_{m,n,p} - \mu_{m,n,p})^2}$$

where $Re(\cdot)$ and $Im(\cdot)$ designate the real and imaginary parts of a number, respectively.

IV. Trajectory Optimization

The problem of generating a continuous-time optimally ergodic trajectory is formulated using the projection-based trajectory optimization method presented in [24]. The optimization is initialized with an initial trajectory and termination criteria. Using the projection-based method allows us to define a local quadratic model of the ergodic objective function at each iteration of the optimization, which can then be used to calculate the steepest descent direction for use in iterative first-order optimization methods, using linear quadratic techniques.

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:

$$\dot{q}(t) = f(q(t), u(t)) \qquad q(t_0) = q_0,$$

where $q(t) \in SE(2)$ represents the state and $u(t) \in \mathbb{R}^m$ the control inputs with dimension m.

B. Objective Function

An objective function $J(\cdot)$ is defined in terms of the metric for ergodicity in Eq. (1) plus the integrated magnitude of the control, which takes as an argument the curves $q(t), u(t)$,

$$J(q(t), u(t)) = \gamma \mathcal{E}(q(t)) + \int_0^T \frac{1}{2} u(\tau) R(\tau) u(\tau) d\tau. \quad (4)$$

In this equation, $\gamma \in \mathbb{R}$ 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.

C. Descent Direction

In order to use first-order iterative optimization methods such as steepest descent, a descent direction must be calculated at every iteration.

Lemma 1: The steepest descent direction $\zeta_i(t) = (z_i(t), v_i(t))$ for the objective function at iteration i is the solution to the linear quadratic problem

$$\begin{aligned}\arg\min_{\zeta_i} \int_0^T & a^T z_i + b^T v_i + \tfrac{1}{2} z_i(\tau)^T Q_n(\tau) z_i(\tau) \\ & + \tfrac{1}{2} v_i(\tau) R_n(\tau) v_i(\tau) d\tau,\end{aligned}$$

subject to $\dot{z}_i(t) = A(t) z_i(t) + B(t) v_i(t), \quad z_i(0) = z_0$ where $A(t) = D_1 f(q_i(t), u_i(t))$ and $B(t) = D_2 f(q_i(t), u_i(t))$. In this problem

$$\begin{aligned} a^T(\tau) =& \gamma \sum_{m,n,p=0}^{M,N,P} \Lambda_{m,n,p} \Bigg[2 \Bigg(\int_0^T \frac{1}{T} F_{m,n,p}(q_i(t)) d\sigma - \\ & \phi_{m,n,p} \Bigg) \circ \frac{1}{T} D F_{m,n,p}(q_i(\tau)) \Bigg], \text{ and} \\ b^T(\tau) =& R(\tau) u_i(\tau). \end{aligned}$$

Proof: Projection-based optimization involves projecting each update onto the feasible trajectory manifold, therefore the objective function at every iteration can be written as a function of a feasible trajectory $\eta_i(t) = (q_i(t), u_i(t))$ [24],

$$\begin{aligned} J(\eta_i(t)) =& \\ & \gamma \sum_{m,n,p}^{M,N,P} \Lambda_{m,n,p} \left\| \frac{1}{T} \int_0^T F_{m,n,p}(q_i(\tau)) d\tau - \phi_{m,n,p} \right\|^2 \\ & + \int_0^T \tfrac{1}{2} u_i(\tau) R(\tau) u_i(\tau). \end{aligned} \quad (5)$$

The steepest descent direction $\zeta_i(t)$ is obtained by minimizing a quadratic model of the form

$$\zeta_i(t) = \arg\min_{\zeta_i(t) \in T_{\eta_i} \mathcal{T}} DJ(\eta_i(t)) \circ \zeta_i(t) + \tfrac{1}{2} \langle \zeta_i(t), \zeta_i(t) \rangle, \quad (6)$$

where $T_{\eta_i} \mathcal{T}$ is the tangent space of the trajectory manifold.

The first term in Eq. (6) is obtained by taking the directional derivative of (5) in the direction $\zeta(t)$. Simple manipulation of the derivative results in:

$$\begin{aligned} & DJ(\eta_i(t)) \circ \zeta_i(t) = \\ & \int_0^T \gamma \sum_{m,n,p}^{M,N,P} \Lambda_{m,n,p} \Bigg[2 \Bigg(\int_0^T \frac{1}{T} F_{m,n,p}(q_i(\sigma)) d\sigma - \phi_{m,n,p} \Bigg) \\ & \circ \frac{1}{T} D F_{m,n,p}(q_i(\tau)) \Bigg] \circ z_i(\tau) + R(\tau) u_i(\tau) \circ v_i(\tau) d\tau. \end{aligned}$$

$a(t)$ and $b(t)$ are defined as the quantities operating on $z_i(t)$ and $v_i(t)$, respectively.

Following [24], the second term in Eq. (5) is defined as

$$\tfrac{1}{2} \langle \zeta_i, \zeta_i \rangle = \int_0^T \left[\tfrac{1}{2} z_i(\tau)^T Q_n(\tau) z_i(\tau) + \tfrac{1}{2} v_i(\tau) R_n(\tau) v_i(\tau) \right] d\tau$$

Where Q_n is a arbitrary positive semi-definite matrix and R_n is positive definite. Using the identity for Q_n and R_n results in steepest descent. The constraint is obtained by linearizing the dynamics, as the descent direction is constrained to be in the tangent space of the trajectory manifold. ∎

Given the directional derivative, we can now define the first order optimality condition as consequence of Lemma 1.

Theorem 1: The necessary condition for optimality of an ergodic trajectory on $SE(2)$ is $|DJ(\eta_i(t)) \circ \zeta_i(t)| = 0$

Note that while the optimization problem in Eq. (4) is not written in the form of a Bolza problem, by using the projection operator defined in [24] and constraining calculations at each iteration of the optimization to the tangent space of the constraint, the quadratic model for Eq. (4) is a Bolza problem. The solution is therefore found using standard Riccati differential equations [24], [25]. For a more detailed treatment of the formulation of an ergodic trajectory optimization, see [1].

V. Simulated Example

We present a simulated example of the ergodic trajectory optimization using a kinematic model for the sensor dynamics. The state for this model is $q(t) = (X(t), X(t), \theta(t))$ where $X(t)$ and $Y(t)$ are Cartesian coordinates and $\theta(t)$ is a heading angle, measured from the X axis in the global frame. 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

$$\dot{\boldsymbol{q}}(t) = \begin{bmatrix} \cos(\theta(t)) & 0 \\ \sin(\theta(t)) & 0 \\ 0 & 1 \end{bmatrix} \cdot \boldsymbol{u}(t). \quad (7)$$

We define the distribution of information over the search domain as a PDF on $SE(2)$. The algorithm was run for two different information PDFs, for a time horizon of 10 seconds. We define the information PDF as a unimodal Gaussian on $SE(2)$. PDF A is a Gaussian is defined on $SE(2)$ with mean $X, Y, \theta = 0$ and covariance $0.1I$ where I is the 3×3 identity matrix. PDF B is centered around $X, Y = 0$, $\theta = -\frac{\pi}{2}$ with the same covariance. Plots of both PDFs are shown in Fig. 4 over X, Y, for two different values of θ.

Figure 3 shows optimized trajectories for both PDFs given a particular initialization (plotted in gray). This initialization

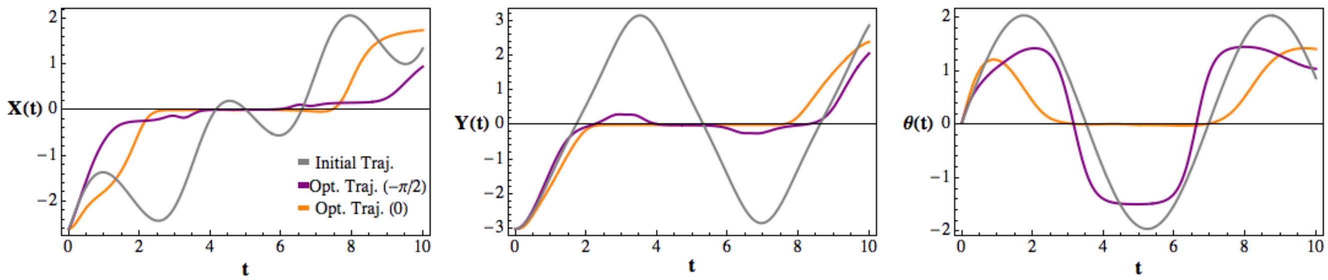

Fig. 2: The trajectories along X, Y, and θ are plotted as functions of time for the trajectories shown in Fig 3. The initial trajectory is plotted in gray, the optimized trajectory for PDF A in orange, and the optimized trajectory for PDF B in purple.

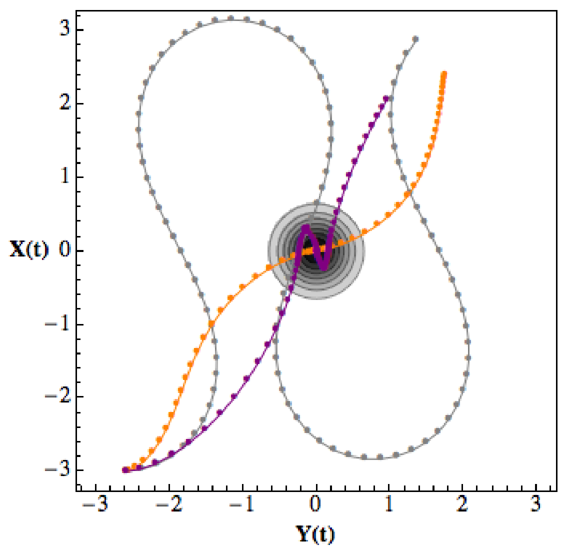

Fig. 3: Plot of the optimized trajectories for PDF A (orange) and PDF B (purple), and initial trajectory for both optimizations (gray). The PDF is depicted as contour lines over X, Y at the mean value of θ.

is similar to a uniform sweep-type search strategy. The optimized trajectory for PDF A, where the Gaussian is centered at $X, Y, \theta = 0$, is plotted in orange. The optimized trajectory for PDF B, where the Gaussian is centered at $X, Y = 0$, $\theta = -\frac{\pi}{2}$, is plotted in purple. Solid lines show the path taken by all three trajectories, dots are plotted at time intervals of 0.1 seconds. The trajectories along $X(t)$, $Y(t)$, and $\theta(t)$ are plotted separately as functions of time for the initial trajectory, optimized trajectory for PDF A and optimized trajectory for PDF B in Fig. 2.

The initial trajectory in Fig. 3 passes very close to the origin, but does not do so at an optimal orientation as defined by either PDF. During the optimized paths for both PDFs, however, the sensor moves through the origin at or close to the optimal angles. The optimization method, while not sensitive to local minima in the information PDF, is subject to local minima with respect to the initial trajectories (plotted in gray) used. We also ran optimizations for the same two PDFs from a different initialization of the trajectory. The trajectories along $X(t)$, $Y(t)$, and $\theta(t)$ and path for a different initialization of the optimization are shown in Figs 5 and 6, respectively.

For both information PDFs, and both initial conditions, in addition to position and orientation tending towards the spatial mean of the distributions, the distribution of time also changes. For both optimized trajectories, the markers are more dense near the mean of the distributions.

Looking at the optimal trajectories, it is also possible to see the effect of simultaneously optimizing over the position and orientation, considering the dynamic constraints. In particular, all trajectories spend more time near the peaks of the Gaussians in the X, Y space than at the optimal orientation. Since the orientation is not independent of the motion of the sensor, there is a balance between the need to move forward, moving towards the origin where the probability of information is highest, and orienting the sensor correctly. All four trajectories however spend a relatively large percentage of time at the correct orientation for the time in which they are close to the origin. Note, however, that in this example the maximum value of this normalized PDF is 0.25, so we expect that an unconstrained ergodic trajectory would spend only 25% of the total search time horizon in the state where the PDF is maximal.

VI. Discussion

With application to sensing modalities or tasks that rely on both orientation and position of a mobile sensor with respect to some target, we present a method of automatically calculating trajectories on $SE(2)$ for sensing tasks. We use a metric on the ergodicity of a sensor trajectory with respect to a distribution representing prior probabilistic information density over the search domain. Infinite-dimensional trajectory optimization is used, minimizing the deviation from

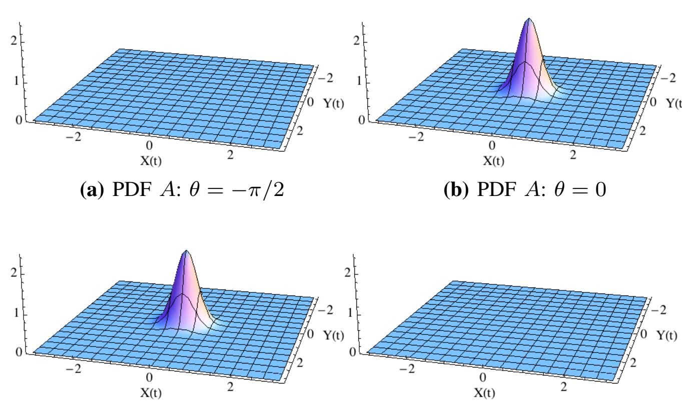

(a) PDF A: $\theta = -\pi/2$ **(b)** PDF A: $\theta = 0$

(c) PDF B: $\theta = -\pi/2$ **(d)** PDF B: $\theta = 0$

Fig. 4: PDF A, with mean $X, Y, \theta = 0$, and PDF B, with mean $X, Y = 0, \theta = -pi/2$, are shown as plots over X, Y for $\theta = (-\pi/2, 0)$.

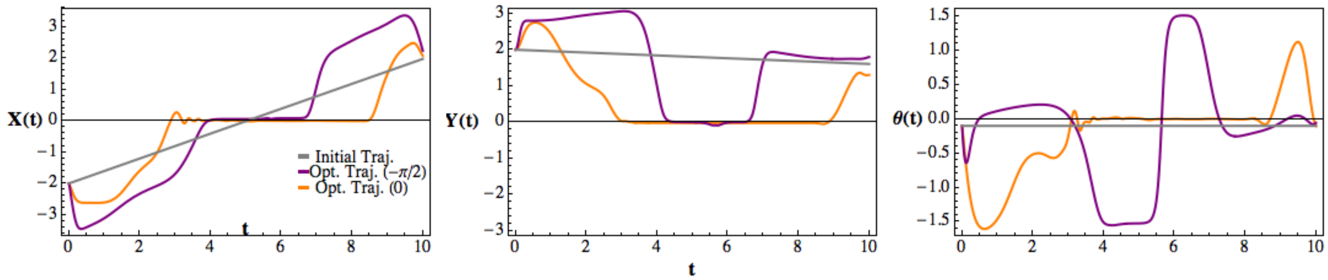

Fig. 5: The trajectories along X, Y, and θ are plotted as functions of time for the trajectories shown in Fig 6.

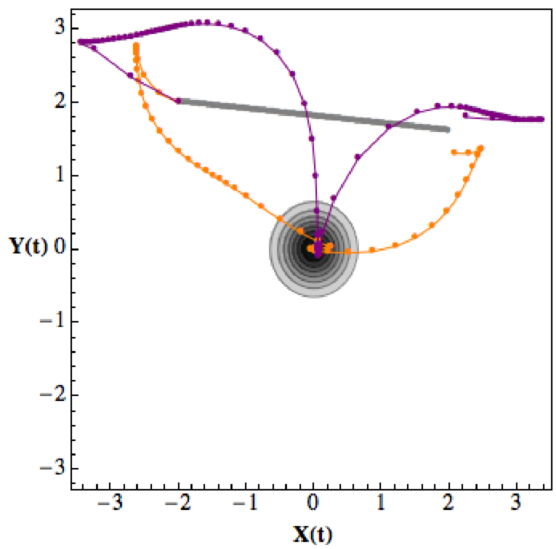

Fig. 6: Plot of the optimized trajectories for PDF A (orange) and PDF B (purple) and alternate initial trajectory for both optimizations (gray).

ergodicity, to generate trajectories which optimally search a given information density.

The method is applied to a Gaussian distribution on $SE(2)$, yielding feasible resulting trajectories that match the input distribution well in both position and orientation. These examples present an early implementation of a control strategy planning step for a closed loop search algorithm. Further work will involve code optimization in order to efficiently calculate the integrations necessary for calculation of a larger number of coefficients over $SE(2)$, and allow extension to $SE(3)$. Additionally, experimental implementation for closed-loop search using a robotic electrosense robot, which has already been implemented along one dimension, independently of orientation [18] is planned for future work.

References

[1] L. Miller and T. Murphey, "Trajectory optimization for continuous ergodic exploration," in *American Controls Conf. (ACC)*, 2013, pp. 432–442.

[2] M. Feezor, P. Blankinship, J. Bellingham, and F. Sorrell, "Autonomous underwater vehicle homing/docking via electromagnetic guidance," in *MTS/IEEE OCEANS Conf.*, vol. 2, 1997, pp. 1137–1142.

[3] S. Cowen, S. Briest, and J. Dombrowski, "Underwater docking of autonomous undersea vehicles using optical terminal guidance," in *MTS/IEEE OCEANS Conf.*, vol. 2, Oct 1997, pp. 1143–1147.

[4] L. Chen, J. H. House, R. Krahe, and M. E. Nelson, "Modeling signal and background components of electrosensory scenes," *J. Comp. Physiol. A*, vol. 191, pp. 331–345, 2005.

[5] S. Roy, "Active recognition through next view planning: a survey," *Pattern Recognition*, vol. 37, no. 3, pp. 429–446, 2004.

[6] W. Burgard, D. Fox, and S. Thrun, "Active mobile robot localization," in *Joint Conf. on Artificial Intelligence (IJCAI)*, 1997.

[7] 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.

[8] Y. F. Li and Z. G. Liu, "Information entropy based viewpoint planning for 3-d object reconstruction," *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 324–327, 2005.

[9] P. Whaite and F. P. Ferrie, "Autonomous exploration: driven by uncertainty," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 3, pp. 193–205, 1997.

[10] H. J. S. Feder, J. J. Leonard, and C. M. Smith, "Adaptive mobile robot navigation and mapping," *International Journal of Robotics Research*, vol. 18, no. 7, pp. 650–668, 1999.

[11] C. Leung, S. Huang, N. Kwok, and G. Dissanayake, "Planning under uncertainty using model predictive control for information gathering," *Robotics and Autonomous Systems*, 2006.

[12] D. Fox, W. Burgard, and S. Thrun, "Active markov localization for mobile robots," *Robotics and Autonomous Systems*, vol. 25, no. 3-4, pp. 195–207, 1998.

[13] R. Sim and N. Roy, "Global a-optimal robot exploration in SLAM," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2005.

[14] F. Bourgault, A. A. Makarenko, S. Williams, B. Grocholsky, and H. Durrant-Whyte, "Information based adaptive robotic exploration," in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.

[15] B. Tovar, L. Munoz-Gomez, R. Murrieta-Cid, M. Alencastre-Miranda, R. Monroy, and S. Hutchinson, "Planning exploration strategies for simultaneous localization and mapping," *Robotics and Autonomous Systems*, vol. 54, no. 4, pp. 314 – 331, 2006.

[16] H. H. González-Baños and J. . Latombe, "Navigation strategies for exploring indoor environments," *International Journal of Robotics Research*, vol. 21, no. 10-11, pp. 829–848, 2002.

[17] L. M. Miller and T. D. Murphey, "Optimal contact decisions for ergodic exploration," in *IEEE Int. Conf. on Decision and Control (CDC)*, 2012.

[18] Y. Silverman, L. M. Miller, M. A. MacIver, and T. D. Murphey, "Optimal planning for information acquisition," in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, vol. In Press, 2013.

[19] G. Mathew and I. Mezic, "Metrics for ergodicity and design of ergodic dynamics for multi-agent system," *Physica D-nonlinear Phenomena*, vol. 240, no. 4-5, pp. 432–442, Feb 2011.

[20] H. Jacobs, S. Nair, and J. Marsden, "Multiscale surveillance of Riemannian manifolds," in *American Controls Conf. (ACC)*, 2010, pp. 5732–5737.

[21] 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.

[22] F. Lekien and N. Leonard, "Nonuniform coverage and cartograms," *SIAM J. Control Optim.*, vol. 48, no. 1, pp. 351–372, Feb 2009.

[23] G. S. Chirikjian and A. B. Kyatkin, *Engineering Applications of Noncommutative Harmonic Analysis*. CRC Press, 2000.

[24] J. Hauser, "A projection operator approach to the optimization of trajectory functionals," in *IFAC world congress*, vol. 4, 2002, pp. 3428–3434.

[25] B. Anderson and J. Moore, *Optimal Control: Linear Quadratic Methods*. Prentice-Hall, 1990.