Optimal Planning for Target Localization and Coverage Using Range Sensing

Lauren M. Miller and Todd D. Murphey

Abstract— **This paper presents an algorithm for autonomously calculating active sensing strategies applied to range sensing. The receding-horizon algorithm we use, called Ergodic Exploration of Distributed Information (EEDI), involves two major components: a) calculation of an expected information density map over the search space based on prior information and a model of the sensor, and b) ergodic trajectory optimization over the sensor configuration space with respect to that information map. The ergodic control algorithm does not rely on discretization of the search or action spaces, and is well posed for coverage with respect to the expected information density whether the information is diffuse or localized. We simulate successful localization and discrimination of targets in a two dimensional workspace using a fixed-location range sensor under various noise levels, and compare performance to an information maximizing strategy.**

I. Introduction

Active sensing, i.e. control of sensor parameters to acquire information or reduce uncertainty, is of critical importance in many robotic applications including search and rescue [1], [2], [3], mine detection [4], and object recognition [5]–[7]. The ability to actively steer or focus sensory attention in each region of the sensor configuration space can significantly improve performance for a broad class of sensing tasks. In this paper we apply a receding-horizon algorithm, called *Ergodic Exploration of Distributed Information* (EEDI) to active steering of a range sensor. Active steering of range sensors is likely to be beneficial in scenarios where sensor steering may be slow or expensive, for low frequency or for long distance measurements [8]. In addition to improved performance in search and tracking applications, active sensing for range sensors is also likely to improve performance of robot localization and map building [9]–[11].

In Section IV, we present a series of simulated experiments using the EEDI algorithm. The sensing objective is to estimate the planar location of a circular target using only a single range sensor that can be steered around a fixed axis. We compare the EEDI algorithm to two standard control alternatives: a uniform sweep strategy (the sensor repeatedly covers the whole search space at a constant velocity) and an information maximizing controller (the sensor is steered to the configuration that maximizes the expected information gain). We first demonstrate localization when the sensor position is fixed (only the beam angle is controlled). We then perform the same localization task, but consider a larger search space and allow the sensor to translate along 1D as well as rotate around the (translating) axis.

The authors are with the Department of Mechanical Engineering at R.R. McCormick School of Engineering and Applied Science at Northwestern University, Evanston, Illinois, USA

This work was supported by Army Research Office grant W911NF-14-1-0461

A. Background and related work

A number of active sensing strategies have been proposed for controlling mobile sensors based on expected information, including gradient-based methods [12], [13] and information maximizing strategies [14]–[17]. While such myopic (selecting only the optimal next configuration or control input) approaches have an advantage in terms of computational tractability, they are likely to fail in the presence of local information maxima or high estimate uncertainty and are typically applied to situations where the sensor dynamics are not considered, and are likely to suffer when uncertainty is high and information diffuse (as argued in [18]–[20]). To avoid sensitivity of single-step optimization methods to local optima, methods of planning control actions over longer time horizons—nonmyopic approaches—are often used. Various heuristics have been developed to approximate the general (otherwise computationally intensive) dynamic programming solution to the nonmyopic optimal control problem [19], [21], [22]. Alternatively, variants of commonly used sampling-based motion planners for maximizing the expected information over a path for a mobile robot have been applied to sensor path planning problems [4], [23], [24].

In our approach, we solve for continuous trajectories based on a map of the expected information density (EID), which is then used in a receding horizon framework as the EID is updated. The aspect of the EEDI algorithm that is distinct from other strategies is the way in which the control is calculated based on the expected information. Ergodic trajectory optimization, originally presented in [25] (for open-loop control), uses iterative trajectory optimization to calculate dynamically feasible, fixed time-horizon trajectories that minimize an objective function based on the *ergodicity* of the trajectory with respect to the EID and a norm on the control effort.

A maximally ergodic trajectory is one for which the percentage of time spent (equivalent to the distribution of measurements for fixed sampling rate) in regions of the state space is proportional to the expected information gain in those regions (characterized by the EID). As opposed to many commonly used information maximization or entropy minimization approaches, the control is designed explicitly to distribute the resulting measurements proportionally to the EID [25] (not the maximum or the mean). We therefore expect increased robustness in the presence of uncertainty or modeling errors. Additionally, ergodic trajectory optimization produces uniform coverage-like trajectories when the expected information is uniformly distributed and localized, and focused search when the information is minimally diffuse. Using ergodicity as an objective therefore results in an algorithm that is suitable for both exploration-prioritizing coverage sampling or exploitation-prioritizing localized sampling, without modification (policy switching or user-defined weighted objectives [20], [22], [26]–[28]).

A preliminary version of the EEDI algorithm was presented in [29]. In [29], only estimation of a single parameter (1D location) of a target is considered, and the sensor movement is constrained to a single dimension. In this paper, we present an extended EEDI algorithm for estimation of multiple parameters, applied to target localization in 2D. Additionally, we demonstrate EEDI when both sensor orientation (bearing angle) and location are considered control design variables. In [29], EEDI was used to estimate underwater target location using a bioinspired robot. That scenario involved a sensor model that was very nonlinear, and therefore is expected to require careful control of sensor motion for estimation. Here—in addition to broadening the scope of the algorithm to multiple parameter estimation in higher dimensional search space—we verify results from [29] using a simpler, more commonly used sensor model. In addition to demonstrating successful localization using EEDI, we demonstrate improved performance over more traditional information maximizing approaches.

II. Problem description

The sensing objective in the following examples is to estimate the location in $\mathbb{R}^2$ of a stationary *target object* in the presence of several similar, unmodeled *distractor objects*. All target and distractor objects are circular objects with varying radii, placed at different locations in a 2D workspace. We use a beam-based sensor model. Assuming that the sensor position and bearing angle are known, and that the shape of the target is known, the expected measurement can be calculated using ray casting operations on the expected target location. The measurement model used for estimation only includes a geometric model of the target object; the model used to simulate measurements includes additive Gaussian noise and multiple unmodeled distractor objects. An illustration of the configuration is shown in Fig. 1.

A sensing trajectory—a "scan"—is produced by rotating the sensor angle for a fixed time period. An example of a simulated measurement scan with the distractor objects and additive noise is shown in Fig. 2, as well as the expected information density (EID) (explained in the next section) for the target object. The objective in using ergodic trajectory optimization is to calculate the optimal scan based on the current EID.

Ergodic trajectory optimization requires a map of the expected information density (EID), and a way of updating the EID iteratively based on measurements. While ergodic trajectory optimization does not rely on the specific choices and assumptions that follow (appropriate for this particular problem), we describe the calculations involved in updating

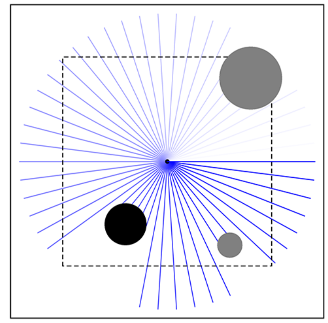

Fig. 1: The objective is to estimate the location of a circular object with known radius (black) within a square workspace. The workspace is indicated by the dashed box. The target object is shown in black, unmodeled distractor objects are shown in gray. The sensor is located in the center of the workspace. Blue rays illustrate the ray lengths produced by the simulation model as a function of the sensor steering angle from zero degrees (light blue) to 2π (dark blue).

and representing the EID for this sensor model and sensing task As we will see in Section IV, calculation of the EID produces a map over the sensor configuration space (bearing angle and position), although the belief is defined over the workspace (potential target locations in $\mathbb{R}^2$).

III. EEDI for stationary target localization using active range sensing

The objective is to estimate the unknown 2D location, the parameters $\boldsymbol{\alpha} \in \mathbb{R}^2$ describing a single, stationary target. A measurement z is made according to a known measurement model $z = \Upsilon(\boldsymbol{\alpha}, \boldsymbol{x}) + \delta$, where $\Upsilon(\cdot)$ is a function of sensor location and parameters (calculated using ray-casting, assuming circular target geometry), and δ represents zero mean Gaussian noise with variance σ^2. Pseudocode for the closed-loop EEDI algorithm is shown in Algorithm 1. The algorithm is initialized with the sensor state at the initial time $\boldsymbol{x}(0)$ and an initial probability distribution $p(\boldsymbol{\alpha})$ for the parameters $\boldsymbol{\alpha}$. The initial distribution can be chosen based on prior information, or in the case of no prior knowledge, assumed to be uniform on bounded domains. The EID is calculated over the sensor state space, given a measurement model that maps sensor state and an estimate of the target parameters (e.g. location, size) to an expected measurement. The normalized EID is then used to calculate an optimally ergodic search trajectory $\boldsymbol{x}_k(t)$ for a finite time horizon $[0, T]$. The trajectory is then executed, collecting measurements $z_k(t)$ along $\boldsymbol{x}_k(t)$. These measurements are used to update $p(\boldsymbol{\alpha})$, which is then used to calculate the EID in the next EEDI iteration.

A. Bayesian Probabilistic Update

We update the parameter estimate—the joint distribution $p(\boldsymbol{\alpha})$—at every iteration k of the EEDI algorithm using a Bayesian filter based on collected measurement $z_k(t)$, the measurement model, and the sensor trajectory $x_k(t)$, $p(\boldsymbol{\alpha}|z_k(t)) = \eta\, p(z_k(t)|\boldsymbol{\alpha})p(\boldsymbol{\alpha})$, where $p(\boldsymbol{\alpha})$ is the belief calculated at the previous iteration, η is a normalization factor, and $p(z_k(t)|\boldsymbol{\alpha})$ is the likelihood function for $\boldsymbol{\alpha}$ given $z_k(t)$. The set of measurements $z_k(t)$ (e.g. the measurements collected during a single scan) is actually a set of discrete measurements (size dependent on trajectory time horizon and sensor sampling rate). Similar to standard beam-based models [30], we assume each measurement is independent and normally distributed. Therefore the likelihood function for all measurements taken along $x_k(t)$ is the product of the likelihood of taking a single measurement $z_k(t_j)$ at time t_j, for all times $t_j \in [0, T]$. We use a standard Gaussian likelihood function,

$$p(z_k(t)|\boldsymbol{\alpha}) = \prod_{j=1}^{T} \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(z(t_j) - \Upsilon(\boldsymbol{\alpha}, \boldsymbol{x}_k(t_j)))^2}{2\sigma^2}\right], \tag{1}$$

although alternative likelihood functions for range sensors have been proposed [31], [32]. Measurement independence is commonly assumed, for example in occupancy grid problems [19], however likelihood functions that do not rely on this assumption of independence, [33], could be used without significantly changing the structure of the algorithm.

B. Calculating Expected Information Density using Fisher Information

Using tools from information theory, a measurement model can be used to predict which sensor parameter values (e.g. sensor state) will provide the most informative measurement for an estimation task. At every iteration of the EEDI algorithm, the expected information density (EID) is calculated by taking the expected value of the Fisher information matrix with respect to $p(\boldsymbol{\alpha})$. Fisher information quantifies the ability of a random variable, in our case a measurement, to estimate an unknown parameter [11], [17], [34], [35]. The Fisher information is represented as an $m \times m$ matrix, where m is the number of parameters being estimated (in the examples that follow, $m = 2$). Each element of the Fisher information matrix (FIM), assuming independent parameters and a Gaussian noise model, can be reduced to

$$\mathcal{I}_{i,j}(\boldsymbol{x}, \boldsymbol{\alpha}) = \frac{1}{\sigma^2} \frac{\partial^2 \Upsilon(\boldsymbol{\alpha}, \boldsymbol{x})}{\partial \alpha_i \partial \alpha_j}. \tag{2}$$

Algorithm 1 Ergodic Exploration of Distributed Information (EEDI), Single Target

Define $\boldsymbol{x}(t_0), \Upsilon(\boldsymbol{\alpha}, \boldsymbol{x}), \epsilon$,
Init. $p(\boldsymbol{\alpha})$ (e.g. to a uniform distribution)
Calculate the Fisher Information $I(\boldsymbol{\alpha}, \boldsymbol{x})$ (Section III-B)
k = 0
while $|SD(\boldsymbol{\alpha})| > \epsilon$ **do**
  a) Calculate EID_k (Section III-B)
  b) Calculate ergodic control $\boldsymbol{u}_k(t)$ (Section III-C)
  c) Execute trajectory $\boldsymbol{x}_k(t)$, measuring $z_k(t)$
  d) Update $p_{k+1}(\boldsymbol{\alpha})$ given $z_k(t)$ (Section III-A)
  k++
end while

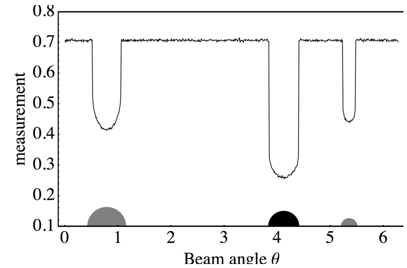

(a) Simulated measurements as a function of sensor steering angle.

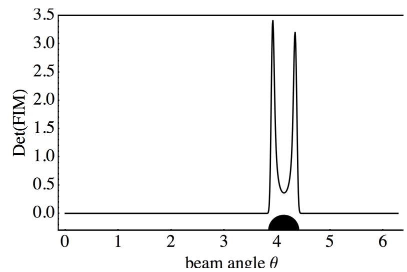

(b) The expected Fisher information

Fig. 2: A range scan is simulated using the beam model with additive Gaussian noise with a target object (black) and two unmodeled distractor objects (gray). In (b), the expected Fisher information is shown for a circular target object located as indicated by the back circle.

The Fisher information can be pre-calculated for a discrete set of α, x to reduce online computation.[1]

Since the estimate of $\boldsymbol{\alpha}$ is represented as a probability distribution function, we take the expected value of each element of $\mathcal{I}(\boldsymbol{x}, \boldsymbol{\alpha})$ with respect to the joint distribution $p(\boldsymbol{\alpha})$ to calculate the expected information matrix, $\Phi(\boldsymbol{x})$. This is an $m \times m$ matrix, where the i, j^{th} element is

$$\Phi_{i,j}(\boldsymbol{x}) = \frac{1}{\sigma^2} \int_{\alpha_i} \int_{\alpha_j} \frac{\partial^2 \Upsilon(\boldsymbol{\alpha}, \boldsymbol{x})}{\partial \alpha_i \partial \alpha_j} p(\alpha_i, \alpha_j)\, d\alpha_j d\alpha_i. \tag{3}$$

This expression can be approximated as a discrete sum as required for computational efficiency.

The often-used D-optimality metric on the expected information matrix, equivalent to maximizing the determinant of the expected information [11], [35]–[37], is used to define a scalar metric on the information matrix over the sensor state space. Therefore, the expected information distribution (EID) that is passed to the ergodic trajectory optimization is

$$EID(\boldsymbol{x}) = \det \Phi(\boldsymbol{x}). \tag{4}$$

Note that different choices of optimality criteria may result in different performance for different problems based on, for example, the conditioning of the information matrix. D-optimality is commonly used for similar applications and we found it to work well experimentally, however the rest of the EEDI algorithm is not in any way dependent on this choice of optimality criterion.

C. Ergodic Trajectory Optimization

Ergodicity compares the difference between the time-averaged statistics of a trajectory to the spatial statistics of

[1]the measurement model generated using ray-casting operations are not continuously differentiable; we use a piecewise approximation.

the EID. The time-averaged statistics of a trajectory $x(t)$ are expressed as a distribution over the spatial domain by calculating the percentage of time the trajectory spends in a neighborhood of each point x. The distance from ergodicity, $\mathcal{E}(x(t))$, can be quantified by defining a norm on the Fourier coefficients of both distributions [38]. This norm is the sum of the weighted squared distance between the Fourier coefficients of the spatial distribution (the EID), ϕ_k, and those of the distribution representing the time-averaged trajectory, $c_k(\boldsymbol{x}(t))$. The ergodic metric $\mathcal{E}$ is therefore

$$\mathcal{E}(x(t)) = \sum_{\boldsymbol{k}=0\in\mathbb{Z}^n}^{\boldsymbol{K}\in\mathbb{Z}^n} \Lambda_k \left[c_{\boldsymbol{k}}(\boldsymbol{x}(t)) - \phi_{\boldsymbol{k}}\right]^2, \qquad (5)$$

where K is the number of coefficients calculated along each of the n dimensions, $\boldsymbol{k}$ is a multi-index $(k_1, k_2, ..., k_n)$, and Λ_k is a weighting factor [38]. When $\mathcal{E}(\boldsymbol{x}(t)) = 0$, the statistics of the trajectory perfectly match the statistics of the distribution.

Ergodic trajectory optimization assumes a general dynamic model for a mobile sensor $\dot{\boldsymbol{x}}(t) = f(\boldsymbol{x}(t), \boldsymbol{u}(t))$ where $\boldsymbol{x} \in \mathbb{R}^N$ is the state and $\boldsymbol{u} \in \mathbb{R}^n$ the control. The motion model can be nonlinear and dynamic, and is assumed to be deterministic. For this paper we consider calculating search trajectories in 1D and 2D, however the trajectory optimization of the objective in Eq. (6) can be extended to search in higher dimensional search spaces such as $\mathbb{R}^3$ and $SE(2)$, so long as a Fourier transform exists for the manifold [39], [40].

We can solve for a continuous trajectory that minimizes an objective function based on both the measure of the ergodicity of the trajectory with respect to the EID and the control effort, defined as

$$J(x(t), u(t)) = Q\mathcal{E}(x(t)) + \int_0^T \tfrac{1}{2} u(\tau)^T R u(\tau) d\tau, \qquad (6)$$

where $\mathcal{E}(x(t))$ is a norm on the ergodicity of the trajectory $x(t)$. Q (scalar) and R ($n \times n$ matrix) determine the relative importance of minimizing ergodicity vs. control effort in the optimization. Minimization of the objective function in Eq. (6) is accomplished using an extension of the trajectory optimization method presented in [41]. The infinite dimensional optimization we use provides a tractable method for optimizing continuous, dynamically constrained trajectories and does not require discretization of search space or control actions in space or time.[2] For details, see [25].

IV. Examples

The results we present were performed in simulation, as this allows us to control noise levels and characteristics and to evaluate the qualitative behavior of the controllers in a systematic way. In Section IV-A, we compare the performance of the EEDI algorithm to the following control strategies.

1) **Uniform Sweep** Using the uniform sweep controller, the sensor is steered from $\theta = 0$ to $\theta = 2\pi$ at a constant velocity.
2) **Information Maximization Controller (IM)** The IM controller performs an uniform sweep initially to initialize the belief, then uses proportional gain feedback control to drive the sensor to the the EID maximum. [17].

In both Sections IV-A and IV-B we introduce uncertainty by including unmodeled distractor objects, as mentioned previously (see Figs. 1 and 2). In Section IV-A, we consider a second form of uncertainty—unmodeled error in the $\mathbb{R}^2$ location of the sensor. While the probabilistic update and the EID calculation do take the noise parameters of the measurement model into account, the uncertainty in the sensor position is not explicitly taken into account in the PDF update or EID calculation.

In all examples, the ergodic optimal control calculation is performed using single integrator dynamics, although this could easily be replaced with a more complex, nonlinear model. We assume that no information about the object location is initially available, i.e. the initial estimate is a uniform probability over the workspace, and as mentioned above, that there is a model of the expected measurement for a given target object, in this case a solid, two-dimensional circle with known radius.

We compare the performance of the different controllers using two metrics; 1) the number of scans required before the norm of the covariance of the estimate PDF drops below a threshold ($Tr(\sigma^2) < 0.03$), and 2) the norm on the error between the estimated object location and the true object location when the termination criterion is reached.

A. 2D Target localization: bearing only control

For the results presented in this section, only the bearing angle of the sensor is controlled (the position of the sensor is held constant). The search space is limited to a 1×1 unit square region, where all points in the region are within range of the sensor. Each iteration of the EEDI algorithm calculates a trajectory for a fixed scan time T of 5 seconds. Measurements are simulated for each trajectory at 100 Hz.

The number of scans required for convergence at different levels of added measurement noise is plotted in Fig. 3 for the EEDI, IM, and US controllers. Each point on the plot represents the mean of five simulated trials. All three controllers exhibit similar convergence rates at low noise levels, and show reduced convergence rates as the noise level is increased. The IM controller results in the fastest convergence of the estimate for low noise levels, but is requires the most scans as the noise level increases. Using the IM controller, there were also instances of the belief converging to the wrong location (finding a distractor instead of the target) at higher noise levels. This was not observed using the ergodic controller.

Examples of the trajectories produced using the IM and EEDI algorithms, plotted with the evolving EID, are shown for two different noise levels in Figs. 4 and 5. The US

[2] The EID map and ergodic objective function could, however, also be utilized within an alternative trajectory optimization framework.

controller is not shown. In Fig. 4, the added noise is small, and both trajectories quickly focus on the correct target. Figure 5 is an example of a trial with higher noise levels, and demonstrates an instance where, although the IM controller produces the fastest convergence early on, it converges to the incorrect value. We also see that even for the low noise scenario depicted in Fig. 4, the EID after the same number of scans is less diffuse (although it should be noted that reducing the variance on the EID is not the same as reducing the variance on the PDF). The evolving PDF and corresponding EID for both EEDI and IM controllers for a different trial (similar to those plotted in Fig. 5, but at a lower noise level), are shown in Fig. 9. The corresponding EID for each PDF is projected onto a circle as a function of the sensor angle.

A summary of the performance of the three algorithms for unmodelled sensor location error is shown in Fig. 6. The measurement model assumes that the sensor position in the workspace is known exactly; simulations were performed with a sensor position that was perturbed by various amounts. For all trials summarized in Figs. 6, the noise level is the same (added noise variance= 0.0025). All three algorithms still perform reasonably well when the modeled sensor location is only slightly perturbed, and all three controllers fail when the true sensor position is significantly different from the modeled position. In this scenario, the performance in terms of estimate error drops off most quickly for the IM controller as the error in the sensor position increases.

Examples of the IM and EEDI Trajectories, plotted with the evolving EID, are shown for two different error levels in Figs. 7 and 8. In Fig. 7, the sensor location is only slightly perturbed and we see both algorithms converge to the correct estimate (given the perturbed sensor position). In Fig. 8, the IM controller converges to the incorrect value as the error increases. The evolving PDF estimate for the same

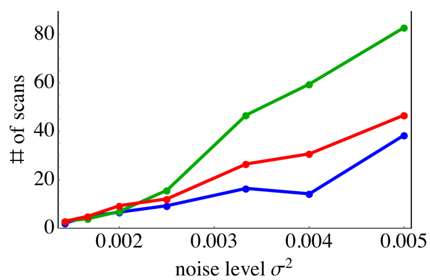

Fig. 3: The number of scans required for the variance of the PDF to meet the termination criterion is plotted as a function of measurement noise level.

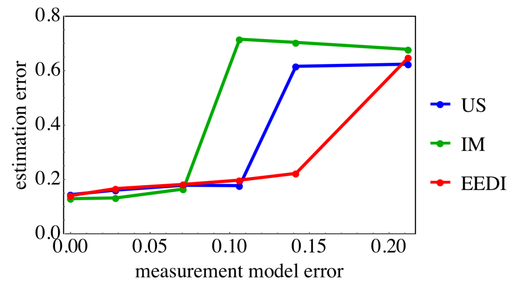

Fig. 6: The norm of the estimation error is plotted at different levels of error between the modeled sensor position and the true (simulated) sensor position. Measurement noise is held constant at 0.0025.

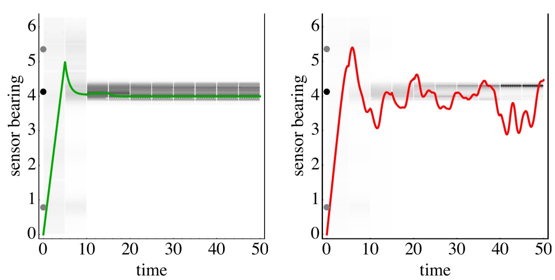

(a) The IM sensor trajectory **(b)** The EEDI sensor trajectory

Fig. 4: The sensor trajectories calculated using both the IM and EEDI controllers are plotted with the evolving EID (grayscale). For both trajectories, the variance of the added measurement noise was 0.0016.

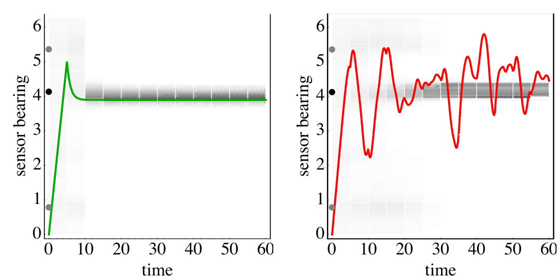

(a) IM sensor trajectory **(b)** EEDI sensor trajectory

Fig. 7: The sensor trajectories calculated using both the IM and EEDI controllers are plotted with the evolving EID (grayscale). The variance of the measurement noise was 0.0025 and the norm on the sensor position error was 0.028.

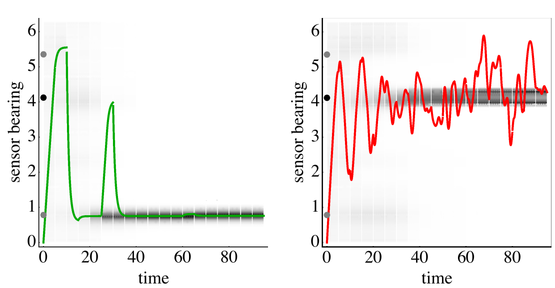

(a) IM sensor trajectory **(b)** EEDI sensor trajectory

Fig. 5: The sensor trajectories calculated using both the IM and EEDI controllers are plotted with the evolving EID (grayscale). The variance of the measurement noise was 0.0033. This example illustrates the IM controller failing in the presence of local maxima in the EID.

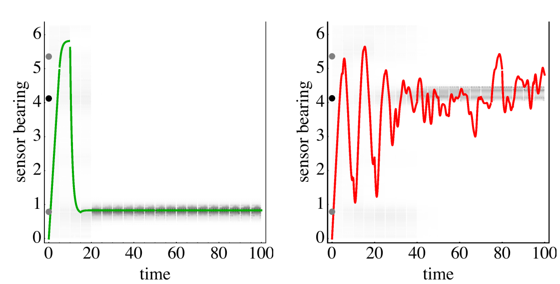

(a) IM sensor trajectory **(b)** EEDI sensor trajectory

Fig. 8: The sensor trajectories calculated using both the IM and EEDI controllers are plotted with the evolving EID (grayscale). The variance of the measurement noise was 0.0025 and the norm on the sensor position error was 0.14

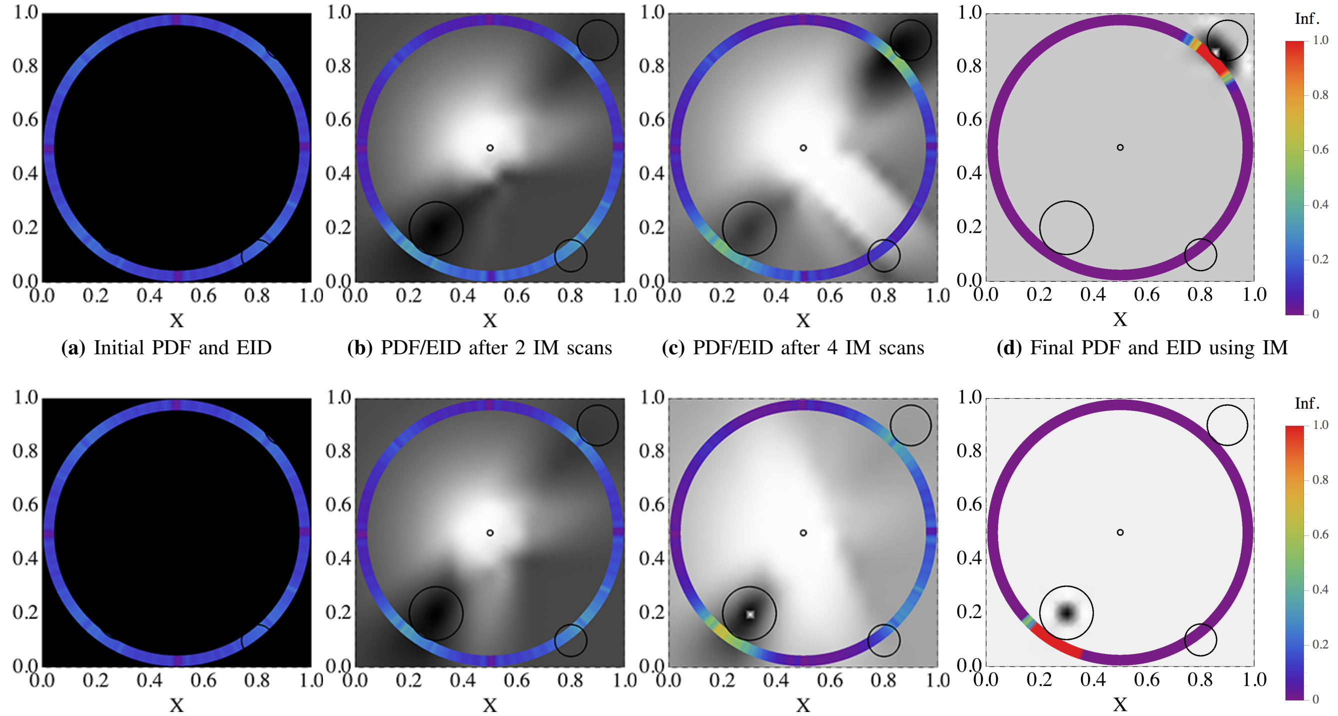

(a) Initial PDF and EID (b) PDF/EID after 2 IM scans (c) PDF/EID after 4 IM scans (d) Final PDF and EID using IM

(e) Initial PDF and EID (f) PDF/EID after 2 EEDI scans (g) PDF/EID after 4 EEDI scans (h) Final PDF and EID using EEDI

Fig. 9: The evolution of the PDF (grayscale) of the target object in $\mathbb{R}^2$ is plotted over the workspace with the EID (colored), at different iterations of the EEDI algorithm. Dark regions represent a high probability of target location, white a low probability. The EID, a function of the sensor bearing angle, is also shown projected onto a circle. The locations of the target and distractor objects are indicated by the unfilled black circles. Figures (a)-(d) illustrate the progression of the PDF and EID using the EEDI algorithm Figures (e)-(h) show the progression using the IM controller, ultimately converging to the wrong object The additive noise level for this trial was 0.0025.

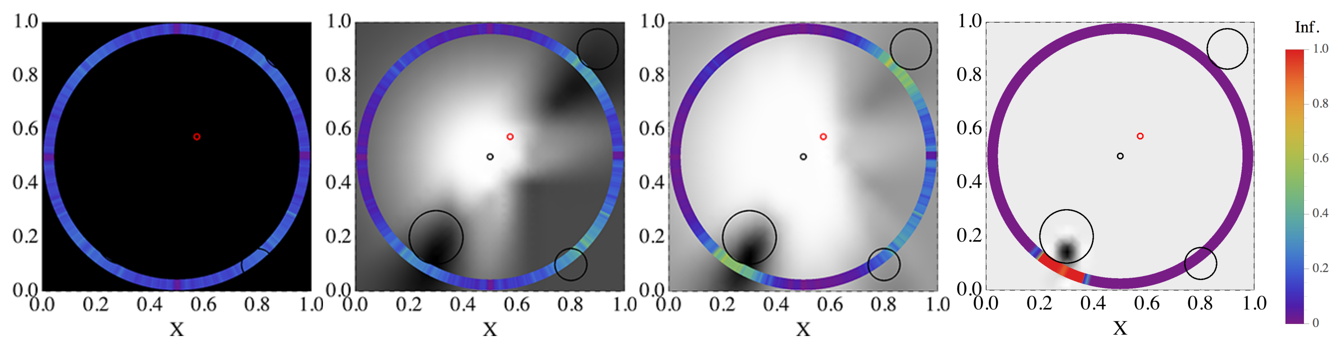

(a) Initial PDF and EID (b) PDF/EID after 2 EEDI scans (c) PDF/EID after 4 EEDI scans (d) Final PDF and EID using EEDI

Fig. 10: In this example, the simulated measurement level is 0.0025, and the norm on the error between modeled and true sensor position is 0.14 (10% of the workspace). The sensor location used in the EID/PDF measurement model is as indicated by the central black circle, the perturbed sensor location indicated by the small red circle. These plots correspond to the trajectory in Fig. 8b.

trial shown in Fig. 8b, using the EEDI controller, is shown in Fig. 10. We see that the ergodic controller produces a PDF that converges to a value corresponding to the correct target estimate, offset by the bias in the sensor model. The IM controller on the other hand (shown in Fig. 8a only) simply fails to "explore" enough to be robust with levels of increasing model uncertainty.

B. 2D Target localization: bearing and translation control

We also demonstrate successful target localization, in the presence of distractors, when controlling both the bearing angle of the sensor and the location (along a single dimension). In this case, he search space is extended to a 6×1 rectangular region, where not all points in the region are within range of the sensor. Three distractor objects were included in the search space. Each iteration of the EEDI algorithm calculates a trajectory for a fixed scan time T of 20 seconds. Two examples of the evolution of the PDF, EID, and the search trajectories as a function of algorithm iterations are shown in Fig. 11, for two different noise levels (0.0014 and 0.0017). Measurements are simulated for each trajectory at 100 Hz, and the locations of the target and distractor objects were randomly selected.

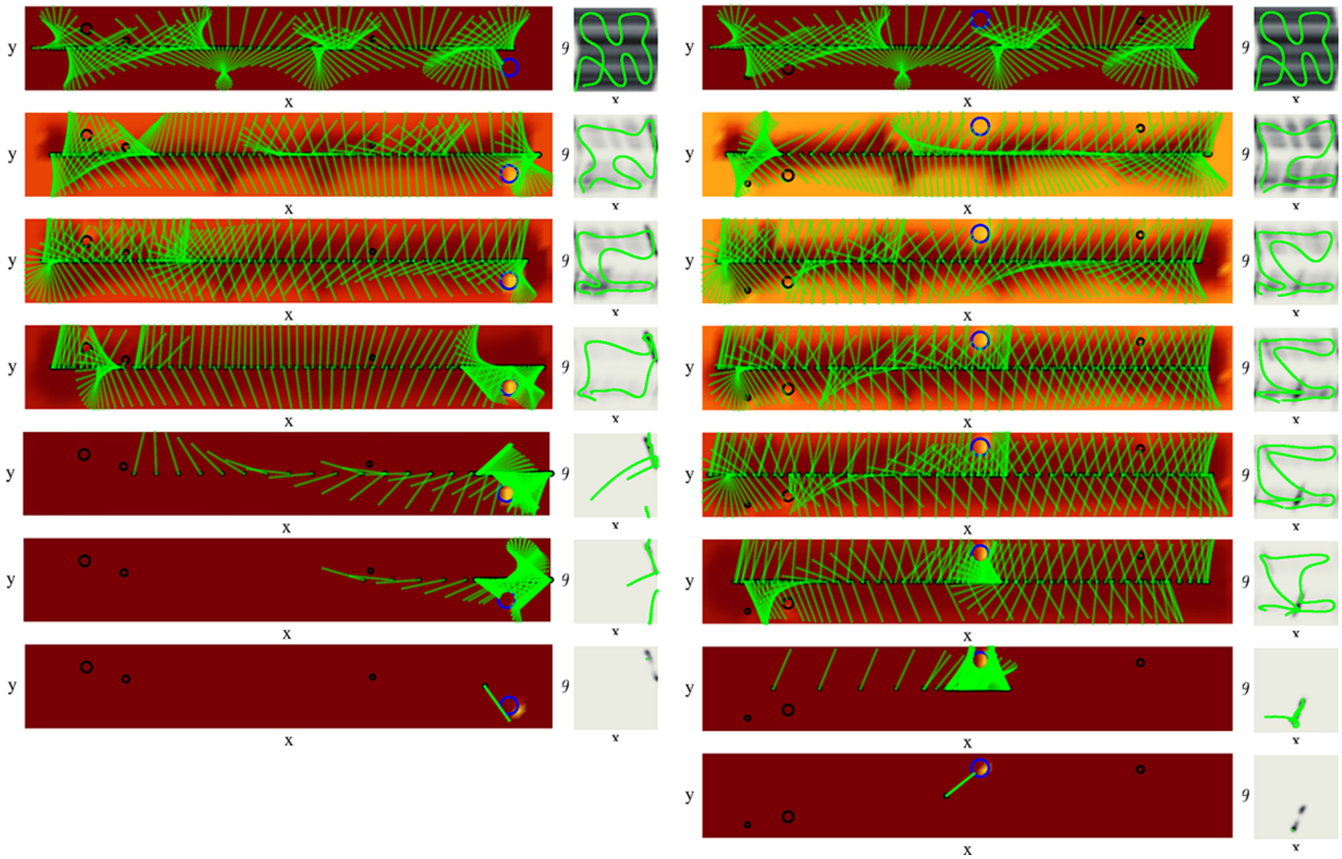

(a) PDF, search trajectory, and EID evolution using EEDI, $\sigma^2 = 0.0014$ **(b)** PDF, search trajectory, and EID evolution using EEDI, $\sigma^2 = 0.0017$

Fig. 11: The evolution of the PDF (heatmap), is plotted over the the 6×1 workspace. The position of the sensor is indicated with filled black circles, with the simulated range scans in green, both potted at 0.1 second intervals. The target object is indicated by the unfilled blue circle, three distractors are shown as unfilled black circles. The EID evolution, and the ergodic trajectory, are plotted to the right of the PDF at each iteration, over the sensor configuration space—position x and orientation θ.

V. Conclusion

We present a receding horizon control algorithm for active estimation using mobile sensors. Information theory, the measurement model, and the belief on the estimates are used to create a spatial map of expected information gain. Because the EEDI algorithm performs an optimization based on the distribution of measurements over a particular scan, the EEDI is more robust in situations where information maximization strategies (e.g. high variance or noise, multimodal distributions) are likely to fail.

The EEDI strategy, because it calculates trajectories based on a distribution, also seamlessly transitions between sensing trajectories similar to a uniform sweep strategy (when the distribution being sampled is uniform) and an information maximizing strategy (as the distribution approaches a delta function). This is observed in the series of trajectories, plotted over the evolving EID in Fig. 11. The information maximizing controller, on the other hand, requires the belief to be initialized using a uniform sweep in order to find the EID maximum. We expect that when the search problem is extended to higher dimensional sensing spaces—for example if we are allowed to control the angular position of the sensor as well as the position in $\mathbb{R}^2$—differences in performance using the different controllers will be even more apparent. The uniform scan approach becomes increasing inefficient in higher dimensions, and we expect the higher dimensional space to result in more local minima in the information space, resulting in poorer performance of information maximizing controllers.

In this paper, we make several assumptions for implementation. We assume a kinematic model for dynamics, although a major advantage of ergodic trajectory optimization is that the formulation is suitable for systems with nonlinear motion constraints. The simulations presented deal exclusively with localizing a single, stationary target, and the formulation of ergodic exploration provided here assumes deterministic dynamics. However, ergodic search generalizes to both time-varying systems as well as estimation of a continuum of targets (e.g., fields [21], [42]) in a reasonably straightforward fashion. Field exploration can be achieved by using an appropriate choice of measurement model and belief update in the EID calculation, and estimation of time-varying parameters could be achieved by extending the state in Section III-C to use time as a state. The determinism restriction primarily makes calculations and exposition simpler; adding stochastic process noise to the model can be achieved by replacing the deterministic, finite-dimensional equations of motion with the Fokker-Planck equations [43] for the nonlinear stochastic flow, without changing the mathematical formulation of ergodic control. These extensions will direct future work.

REFERENCES

1. J. Toh and S. Sukkarieh, “A Bayesian formulation for the prioritized search of moving objects,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2006, pp. 219–224.
2. J. Cooper and M. Goodrich, “Towards combining UAV and sensor operator roles in UAV-enabled visual search,” in *IEEE Int. Conf. on Human Robot Interaction (HRI)*, 2008, pp. 351–358.
3. G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, “Active planning for underwater inspection and the benefit of adaptivity,” *International Journal of Robotics Research*, vol. 32, no. 1, pp. 3–18, 2013.
4. C. Cai and S. Ferrari, “Information-driven sensor path planning by approximate cell decomposition,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 39, no. 3, pp. 672–689, 2009.
5. J. Denzler and C. Brown, “Information theoretic sensor data selection for active object recognition and state estimation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 2, pp. 145–157, 2002.
6. T. Arbel and F. Ferrie, “Viewpoint selection by navigation through entropy maps,” in *IEEE Int. Conf. on Computer Vision*, vol. 1, 1999, pp. 248–254 vol.1.
7. Y. Ye and J. K. Tsotsos, “Sensor planning for 3D object search,” *Computer Vision and Image Understanding*, vol. 73, no. 2, pp. 145 – 168, 1999.
8. J. J. Leonard and H. F. Durrant-Whyte, *Directed sonar sensing for mobile robot navigation*. Kluwer Academic Publishers Dordrecht, 1992, vol. 448.
9. F. Lu and E. Milios, “Robot pose estimation in unknown environments by matching 2d range scans,” *Journal of Intelligent and Robotic Systems*, vol. 18, no. 3, pp. 249–275, 1997.
10. M. Ruhnke, R. Kummerle, G. Grisetti, and W. Burgard, “Highly accurate maximum likelihood laser mapping by jointly optimizing laser points and robot poses,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, May 2011, pp. 2812–2817.
11. S. Martínez and F. Bullo, “Optimal sensor placement and motion coordination for target tracking,” *Automatica*, vol. 42, no. 4, pp. 661 – 668, 2006.
12. B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, “Cooperative air and ground surveillance,” *IEEE Robotics and Automation Magazine*, vol. 13, no. 3, pp. 16–25, 2006.
13. C. Kreucher, J. Wegrzyn, M. Beauvais, and R. Conti, “Multiplatform information-based sensor management: an inverted UAV demonstration,” in *SPIE Defense Transformation and Network-Centric Systems*, vol. 6578, 2007, pp. 65 780Y–1–65 780Y–11.
14. C. Kreucher, K. Kastella, and A. O. Hero, “Sensor management using an active sensing approach,” *Signal Processing*, vol. 85, no. 3, pp. 607–624, 2005.
15. P.-P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich, “Viewpoint selection using viewpoint entropy.” in *Vision Modeling and Visualization Conference*, vol. 1, 2001, pp. 273–280.
16. 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.
17. X. Liao and L. Carin, “Application of the theory of optimal experiments to adaptive electromagnetic-induction sensing of buried targets,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 8, pp. 961–972, 2004.
18. M. Rahimi, M. Hansen, W. Kaiser, G. Sukhatme, and D. Estrin, “Adaptive sampling for environmental field estimation using robotic sensors,” in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, Aug 2005, pp. 3692–3698.
19. C. Stachniss and W. Burgard, “Exploring unknown environments with mobile robots using coverage maps,” in *IJCAI*, 2003, pp. 1127–1134.
20. K. H. Low, J. M. Dolan, and P. Khosla, “Adaptive multi-robot wide-area exploration and mapping,” in *Conference on Autonomous agents and multiagent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 23–30.
21. N. Cao, K. H. Low, and J. M. Dolan, “Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms,” in *International Conference on Autonomous Agents and Multi-agent Systems*, 2013, pp. 7–14.
22. T. N. Hoang, K. H. Low, P. Jaillet, and M. Kankanhalli, “Nonmyopic -bayes-optimal active learning of gaussian processes,” in *International Conference on Machine Learning*, 2014, pp. 739–747.
23. G. Zhang, S. Ferrari, and M. Qian, “An information roadmap method for robotic sensor path planning,” *Journal of Intelligent and Robotic Systems*, vol. 56, no. 1-2, pp. 69–98, 2009.
24. G. A. Hollinger and G. S. Sukhatme, “Sampling-based robotic information gathering algorithms,” *International Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.
25. L.M. Miller and T.D. Murphey, “Trajectory optimization for continuous ergodic exploration,” in *American Controls Conf. (ACC)*, 2013, pp. 4196–4201.
26. A. Krause and C. Guestrin, “Nonmyopic active learning of gaussian processes: an exploration-exploitation approach,” in *International Conference on Machine learning*. ACM, 2007, pp. 449–456.
27. R. Martinez-Cantin, N. de Freitas, E. Brochu, J. Castellanos, and A. Doucet, “A bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot,” *Autonomous Robots*, vol. 27, no. 2, pp. 93–103, 2009.
28. R. Marchant and F. Ramos, “Bayesian optimisation for informative continuous path planning,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 6136–6143.
29. 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. 5974- 5980, 2013.
30. S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, Mass: The MIT Press, 2005.
31. P. Pfaff, C. Plagemann, and W. Burgard, “Improved likelihood models for probabilistic localization based on range scans,” in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, Oct 2007, pp. 2192–2197.
32. C. Plagemann, K. Kersting, P. Pfaff, and W. Burgard, “Gaussian beam processes: A nonparametric bayesian measurement model for range finders,” in *Robotics: Science and Systems (RSS)*, Atlanta, Georgia, USA, June 2007.
33. S. Thrun, “Learning occupancy grid maps with forward sensor models,” *Autonomous Robots*, vol. 15, no. 2, pp. 111–127, 2003.
34. R. B. Frieden, *Science from Fisher Information: A Unification*. Cambridge University Press, 2004.
35. A. Emery and A. V. Nenarokomov, “Optimal experiment design,” *Measurement Science and Technology*, vol. 9, no. 6, p. 864, 1998.
36. Y. Oshman and P. Davidson, “Optimization of observer trajectories for bearings-only target localization,” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 35, no. 3, pp. 892–902, Jul 1999.
37. T. Chung, J. Burdick, and R. Murray, “A decentralized motion coordination strategy for dynamic target tracking,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, May 2006, pp. 2416–2422.
38. 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.
39. G. S. Chirikjian and A. B. Kyatkin, *Engineering Applications of Noncommutative Harmonic Analysis*. CRC Press, 2000.
40. L. M. Miller and T. D. Murphey, “Trajectory optimization for continuous ergodic exploration on the motion group SE(2),” in *IEEE Int. Conf. on Decision and Control (CDC)*, 2013, p. In Press.
41. J. Hauser, “A projection operator approach to the optimization of trajectory functionals,” in *IFAC world congress*, vol. 4, 2002, pp. 3428–3434.
42. A. Bender, S. B. Williams, and O. Pizarro, “Autonomous exploration of large-scale benthic environments,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 390–396.
43. G. S. Chirikjian, *Stochastic Models, Information Theory, and Lie Groups, Volume 1: Classical Results and Geometric Methods*. Springer Science & Business Media, 2009.