



10-2016

Mobile Robots as Remote Sensors for Spatial Point Process Models


Paul B. Reverdy

University of Pennsylvania, preverdy@seas.upenn.edu

Daniel E. Koditschek

University of Pennsylvania, kod@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/ese_papers

 Part of the [Electrical and Computer Engineering Commons](#), and the [Systems Engineering Commons](#)

Recommended Citation

Paul B. Reverdy and Daniel E. Koditschek, "Mobile Robots as Remote Sensors for Spatial Point Process Models", *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2847-2852. October 2016.

missing image.gif

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ese_papers/716
For more information, please contact repository@pobox.upenn.edu.

Mobile Robots as Remote Sensors for Spatial Point Process Models

Abstract

Spatial point process models are a commonly-used statistical tool for studying the distribution of objects of interest in a domain. We study the problem of deploying mobile robots as remote sensors to estimate the parameters of such a model, in particular the intensity parameter λ which measures the mean density of points in a Poisson point process. This problem requires covering an appropriately large section of the domain while avoiding the objects, which we treat as obstacles. We develop a control law that covers an expanding section of the domain and an online criterion for determining when to stop sampling, i.e., when the covered area is large enough to achieve a desired level of estimation accuracy, and illustrate the resulting system with numerical simulations.

Keywords

Spatial point processes, remote sensing, mobile robots, coverage, stopping rule

Disciplines

Electrical and Computer Engineering | Engineering | Systems Engineering

Comments

missing image.gif

Mobile robots as remote sensors for spatial point process models

Paul Reverdy and Daniel E. Koditschek

Abstract—Spatial point process models are a commonly-used statistical tool for studying the distribution of objects of interest in a domain. We study the problem of deploying mobile robots as remote sensors to estimate the parameters of such a model, in particular the intensity parameter λ which measures the mean density of points in a Poisson point process. This problem requires covering an appropriately large section of the domain while avoiding the objects, which we treat as obstacles. We develop a control law that covers an expanding section of the domain and an online criterion for determining when to stop sampling, i.e., when the covered area is large enough to achieve a desired level of estimation accuracy, and illustrate the resulting system with numerical simulations.

Mobile robots have been extensively used as platforms for remote sensors in recent years. Much of the statistically-rigorous research on such platforms has focused on sensing and mapping continuous fields using various models. This has led to a series of results that focus on coordinating a set of robotic vehicles [1], [2], [3] in order to cover a domain of interest [4], [5]. These results have shown that the problem of mapping a spatially-distributed continuous field, e.g., by minimizing the uncertainty associated with an estimate of that field, is closely linked to the problem of covering the space with the remote sensors. The fields studied in applications have included seawater temperature and salinity [1], as well as atmospheric variables of interest to meteorology, such as temperature and vorticity [6].

The mobile sensor network literature has placed significantly less emphasis on spatial point process models, which are statistical models that capture the distribution of discrete objects in space. Point process models are the subject of extensive study in the statistical literature [7] and have proven useful in numerous applied fields. In ecology, point process models have been used to study the distribution of plants in a domain [8], [9], [10], while in search and rescue they have been used to study the distribution of distress calls in a service area [11], and in disaster management they have been used for modeling the distribution of forest fires [12].

Consideration of spatial point process models has recently begun to appear in the robotics literature, albeit largely in a theoretical context. For example, the Probability Hypothesis Density (PHD) used in multitarget tracking [13] can be interpreted as approximating the target distribution as a Poisson point process, which is in some sense the simplest spatial point process. The connection between the PHD and Poisson point processes has been utilized by several authors to develop new implementations of the PHD filter [14]

and to propose PHD filter-based methods for performing Simultaneous Localization and Mapping [15]. More recently, [16] used the PHD filter to implement an active strategy to detect and register target locations.

Several works have used spatial point processes to study the problem of robot navigation. A theoretical article [17] studied the problem of high-speed navigation through obstacle fields. Using a Poisson point process model of obstacle fields, the authors of [17] derived fundamental bounds on the speed with which a vehicle could safely traverse a field with a given mean density of obstacles. Spatial point processes have also been applied in [18], [19] to develop guidance laws for navigating through hazard fields, where hazardous regions are distributed at unknown locations throughout a domain. In these works, the spatial point process is used as a model to predict the hazard locations.

Applying spatial point processes to physical problems requires sensing the locations of the objects being modeled, and can be done in a variety of ways. Standard procedures for sensing plant locations in the ecology literature include remote imaging [9] and other, labor-intensive, surveying methods, while in search and rescue survivors needing rescue can often be located by tracking a mobile phone [20]. Accurate measurements of survivor density can assist in planning the allocation of resources. The use of mobile robots could significantly improve such measurements.

In this work, we study the use of mobile robots as remote sensors for spatial point process models. We focus on developing algorithms for deploying vehicles to estimate the parameters of point process models. The most basic such model is the Poisson point process, whose single parameter λ represents the mean density of points. Other, more complicated point processes which model higher-order details of the point structure are built up using the Poisson point process, so we focus on the Poisson model as a place to begin this new work. As we explain below, the problem of estimating the parameter λ of a Poisson point process model can be reduced to a coverage problem. In the context of this reduction, a fundamental question concerns “cost:” how extensive a domain must be covered in order to provide a parameter estimate of specified accuracy. The Poisson point model provides a natural way to answer this question.

The remainder of the paper is structured as follows. In Section I we review the Poisson point process and derive an online criterion for determining the size of the sampling domain as a function of desired accuracy. In Section II we develop a control law to perform the parameter estimation and prove that the law covers the sampling domain while avoiding obstacles. In Section III we show results of applying

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104. Email: {preverdy, kod}@seas.upenn.edu. This work was supported by AFOSR MURI FA9550-10-1-0567.

this control law in simulated experiments. Finally, in Section IV we offer perspectives on future work and conclude.

I. THE POISSON POINT PROCESS

We assume that the objects to be sensed follow a homogeneous Poisson point process in a region $\mathcal{R} \subset \mathbb{R}^2$. In the Poisson point process model, the number of points $N_{\mathcal{D}}$ in a region $\mathcal{D} \subset \mathbb{R}^n$ is a random variable following a Poisson distribution with parameter λA , where $A = |\mathcal{D}|$ is the volume of the region (i.e., area, if \mathcal{D} is two-dimensional as in our present application setting): $N_{\mathcal{D}} \sim \text{Poisson}(\lambda A)$, so the probability mass function of $N_{\mathcal{D}}$ is

$$f(N_{\mathcal{D}}, \lambda A) = \Pr[N_{\mathcal{D}} = k] = \frac{(\lambda A)^k e^{-\lambda A}}{k!}. \quad (1)$$

Each point is assumed to be located according to a uniform random distribution inside the domain \mathcal{D} , independent of all other points, i.e., the model ignores any possible interaction between the $N_{\mathcal{D}}$ points [7]. Samples of the Poisson point process on a region \mathcal{D} with parameter λ can be drawn by first drawing the number of points $N_{\mathcal{D}}$ according to the Poisson distribution (1), then independently drawing locations for each of the points from a uniform distribution on \mathcal{D} [7]. The quantity $N_{\mathcal{D}}$ modeled by the point process is a random variable, but for any realization of the process the number of points is a deterministic quantity which we denote by $n_{\mathcal{D}}$.




Fig. 1. Three example realizations of the Poisson point process with parameter $\lambda = 1$ on the domain $\mathcal{D} = [0, 2] \times [0, 2]$. The points of each of the three realizations are marked with blue triangles, red circles, and green squares, respectively. Note that each realization has a different number of points $n_{\mathcal{D}}$ (two, eight, and three, respectively).

The expected value of $N_{\mathcal{D}}$ is λA , so the mean density of points under this model is $\mathbb{E}[N_{\mathcal{D}}] = \lambda A / A = \lambda$. Similarly, the variance of $N_{\mathcal{D}}$ is λA . Figure 1 shows three example realizations of the Poisson point process model, each of which has a different value of $n_{\mathcal{D}}$. It can be shown [7, p. 147] that $\hat{\lambda} = n_{\mathcal{D}} / A$ is an unbiased estimator for λ ; in Figure 1, the average value of $n_{\mathcal{D}}$ is $(2 + 3 + 8) / 3 \approx 4.33$, so the maximum likelihood estimate of λ is $\hat{\lambda} = 13 / 3 / 2^2 \approx 1.08$, compared to its true value, $\lambda = 1$. In general, the parameter λ can vary from location to location in what is known as a *non-homogenous* Poisson process. However, for the present paper, we consider only the homogeneous Poisson process, where λ is constant throughout the region \mathcal{R} .

A. Sensor model

We want our mobile sensor-bearing robots to estimate the density parameter λ , which they do by sampling a domain $\mathcal{D} \subset \mathcal{R}$ and counting the number of objects it finds. We assume that each robot knows its location in the sampling domain \mathcal{D} and has a simple sensor, modeled as follows.

We model the robot as a rigid body with pose state $q = (x, \theta) \in \mathcal{D} \times \mathbb{S}^1 \subset \text{SE}(2)$, where x is the location in the domain and θ is the robot's heading in the global reference frame. The robot is equipped with a sensor of range $\rho > 0$ and angular field of view $\beta \in (0, 2\pi]$ centered on θ , so the sensor sees a region in the shape of a circular sector. When an object is visible to the sensor, the sensor reports range and bearing information to the robot; we assume the range and bearing measurements to have minimal noise. Based on these measurements and the robot's known location, the robot can then register the object's location in the sampling domain.

In practice, the location and sensor measurement information will not be known precisely but will be corrupted with noise. However, since we will primarily be relying on the sensor as a means to count the number of objects in the domain, we can safely ignore sensor noise at this stage of algorithm development.

B. Estimation

Our robots will estimate the density parameter λ by sampling a domain $\mathcal{D} \subset \mathcal{R}$ and counting the number of objects they find; the natural estimate of λ is then the mean observed density of points $\hat{\lambda} = n_{\mathcal{D}} / A$. This leads to two main algorithmic questions: 1) How to sample the domain, and 2) how to decide when one can stop sensing?

We note here that the problem of sampling the domain can be reduced to a sensor coverage problem in the sense of Choset [21]. To register the locations of all the objects in the sensing domain \mathcal{D} with our robot-borne sensors, and thereby count $n_{\mathcal{D}}$, we must plan a path for the robots such that the footprint of their sensors covers the whole domain.

To perform the estimation quickly, the domain should be chosen such that the coverage problem can be solved efficiently and that a simple criterion for deciding when to stop sensing can be derived. The isotropic nature of the Poisson point process model suggests selecting a circular domain, and a stopping criterion can be derived based on the concept of the *representative volume element*, which is the size of the smallest sampling volume that is representative of the macro-scale properties of the region [22].

Kanit *et al.* [22] suggest a relative error criterion for determining the size of the representative volume element for a point process:

$$\varepsilon_r = \frac{\varepsilon_a}{z} < \varepsilon, \quad (2)$$

where ε_r denotes the relative measurement error of the property being measured, ε_a denotes the absolute measurement error, z the true mean value of the property, and $\varepsilon \ll 1$ a threshold value. The criterion can be interpreted as follows: stop sampling once the relative error of the estimate is below the threshold value ε .

For estimating λ , we choose the standard deviation of the estimate as the measure of absolute error ε_a . This leads to the expression for relative error

$$\varepsilon_r = \frac{2\sigma(\hat{\lambda})}{\lambda}, \quad (3)$$

which is empirically estimated as

$$\hat{\varepsilon}_r = 2\frac{\sigma(\hat{\lambda})}{\hat{\lambda}}, \quad (4)$$

where $\hat{\lambda} = N_{\mathcal{D}}/A$ is the empirical mean density of points. Using a Gaussian approximation to the distribution of measurement error, this expression for relative error can be considered as a worst-case bound on the (unknown) true relative error that captures 95% of the possible values. This is a form of Provably Approximately Correct (PAC) learning [23], where one seeks an estimate that has error less than ε with probability at least $1 - \delta$.

Equation (4) reduces the relative error criterion to a function of $\hat{\lambda}$, the density estimate, and A , the area sampled. The Poisson point process model assumption allows us to compute $\sigma(\hat{\lambda})$ in closed form for some sampling domains \mathcal{D} . In particular, if the sampling domain is a circle of radius R , one finds, using Equations (4.109) and (1.58) of [7], that:

$$\sigma(\hat{\lambda}) = \frac{1}{R} \sqrt{\frac{\lambda}{\pi}}. \quad (5)$$

Therefore, the stopping criterion $\hat{\varepsilon}_r < \varepsilon$ implies that one can stop sampling once the sampling domain radius obeys

$$R^* \geq \frac{2}{\varepsilon \sqrt{\pi \hat{\lambda}}}. \quad (6)$$

If one has a method for covering discs of increasing radius, this criterion is simple to implement online on a vehicle.¹

II. CONTROL

Having reduced the sampling problem to a coverage problem with an online stopping criterion, we can appeal to the literature on robotic coverage to provide a framework for control. In many applications, for example if the robot is a ground-based vehicle mapping plant locations, the points being sensed constitute obstacles that the robot must avoid; this creates a combined requirement for both coverage and obstacle avoidance. To cover a growing disc we use a technique based on following a spiral path, which we encode in a vector field. Following [24], we add an obstacle avoidance strategy using a repelling vector field; the resulting control law is simple to implement and maps naturally to control policies for robots with non-trivial dynamics, e.g., the unicycle. Other authors have also developed methodologies for performing coverage and obstacle avoidance using spiral paths; e.g., see Section 5 of [25]. Our vector field strategy is less mathematically sophisticated than the cellular decomposition strategy of [25] but significantly simpler to implement.

¹For a domain \mathcal{D} of area A , (5) generalizes to $\sigma = \lambda/A$. Using this result, one can extend our results to use a lawnmower coverage strategy.

A. Spiral-based coverage

The derivation of the stopping criterion requires the robot to cover a circular domain of increasing radius until the criterion is met. A commonly-employed method for solving such a coverage problem is to have the vehicle follow an Archimedean spiral, which can be written in polar coordinates as $r(\phi) = (\alpha/2\pi)\phi$. See Figure 2 for an example. The trajectory of such a spiral has a constant pitch of α , so the vehicle crosses a ray from the origin at regular intervals spaced α apart, which is assumed to be less than the vehicles' sensor footprint. Therefore, our mobile sensing strategy is to deploy a robot along such a spiral pattern, have it register point locations, and use these locations to estimate $\hat{\lambda}$, stopping when the criterion (6) is met. Note that the idealized spiral coverage law employed in Figure 2 results in the trajectory coming arbitrarily close to several points, motivating the need for obstacle avoidance in applications.




Fig. 2. An example of the idealized point process estimation law. The vehicle follows an Archimedean spiral with $\alpha = 1$ until it has sampled a sufficiently large domain. The dashed circle shows the footprint of a sensor with range $\rho = 0.5$ and field of view $\beta = 2\pi$ when the vehicle is located at the position marked by the empty dot.

1) *Robot kinematics*: In the simulations that follow, we implement the spiral-based coverage policy by modeling the robot as a fully-actuated point particle. Despite this limitation in the simulations, our control algorithm maps naturally to a vehicle modeled as a kinematic or dynamic unicycle using the transformation reported in Section III of [24].

2) *Reference vector field*: The obstacle avoidance strategy developed in [24] avoids obstacles while performing a reference navigation task, in this case following a spiral trajectory. This reference task is assumed to be encoded in a vector field $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, whose flow generates the desired trajectory. We encode the arithmetic spiral in a vector field that, in polar coordinates $x = (r, \phi)$, takes the form

$$\begin{aligned} \dot{r} &= \alpha(R_0 - r), \quad R_0 \gg \alpha \\ \dot{\phi} &= (R_0 - r), \end{aligned}$$

where $R_0 \gg \alpha$ is a reference radius. Note that in the absence of obstacles, the dynamics $\dot{x} = f(x)$ traces an arithmetic spiral until $r = \|x\|$ reaches the reference value R_0 , at which point it stops. Therefore, R_0 should be larger than the largest required sampling radius, which can be estimated

from expected values of the intensity λ and desired relative error ε . In Cartesian coordinates, f takes the form

$$\dot{x} = \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \alpha \frac{R_0 - \|x\|}{\|x\|} x + (R_0 - \|x\|) J_2 x, \quad (7)$$

where $J_2 = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ is a symplectic matrix.

The dynamics $\dot{x} = f(x)$ admit a Lyapunov function $V(x) = \frac{1}{2}(\|x\| - R_0)^2$. The gradient of V is

$$\nabla V = \frac{\|x\| - R_0}{\|x\|} x,$$

so in terms of V , the vector field f is

$$f(x) = -\alpha \nabla V - \|x\| J \nabla V = -A \nabla V,$$

where $A = \alpha I_2 + \|x\| J_2$ and I_2 is the identity matrix. Setting $\dot{x} = f(x)$, we have $\dot{V} = -\nabla V^T A \nabla V = -\alpha \gamma_V^2$, where $\gamma_V = \|\nabla V\| = |R_0 - \|x\||$.

The Hessian of V is

$$H_V = \frac{R_0}{\|x\|} \hat{x} \hat{x}^T + \frac{\|x\| - R_0}{\|x\|} I_2,$$

where $\hat{x} = x/\|x\|$. Direct computation yields

$$\|H_V\| = \max \left\{ \frac{|\|x\| - R_0|}{\|x\|}, 1 \right\}.$$

Let $0 < \rho_0 < R_0/2$ and $R_0/2 < \rho_1 < R_0$. Then for $\|x\| > \rho_0$, $\|H_V\| < R_0/\rho_0 - 1$. Furthermore, for $\|x\| < \rho_1 < R_0$, $\|\nabla V\| > R_0 - \rho_1 > 0$.

B. Obstacle avoidance

In some cases, for example an aerial vehicle sensing mobile phone signals in order to locate survivors of a disaster, there is no harm in the vehicle coming close to the objects it is trying to locate. However, in other cases, such as a ground vehicle tasked with locating plants or land mines, the vehicle must treat the objects as obstacles to be avoided so as to avoid risks to the objects or the vehicle itself. In these cases, we must modify the coverage control algorithm in such a way to incorporate obstacle avoidance without damaging the coverage property of the control.

To each point obstacle i , located at $x_i \in \mathbb{R}^2$, we associate a local obstacle function ψ_i that is zero everywhere except on a disc centered on the obstacle.

$$\psi_i = \begin{cases} \frac{1}{2\rho_M} (\rho_M - \|x - x_i\|)^2, & x \in \mathcal{B}_i \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where $\mathcal{B}_i = \{x \in \mathcal{D} : \|x - x_i\| < \rho_M\}$ is a punctured open disk. In addition,

$$\nabla \psi_i|_{\mathcal{B}_i} = -\frac{\rho_M - \|x - x_i\|}{\rho_M \|x - x_i\|} (x - x_i), \quad (9)$$

$$H_{\psi_i}|_{\mathcal{B}_i} = -\left(\frac{1}{\|x - x_i\|} - \frac{1}{\rho_M} \right) \bar{n}_{\psi_i} \bar{n}_{\psi_i}^T + \frac{n_{\psi_i} n_{\psi_i}^T}{\rho_M}, \quad (10)$$

where $n_{\psi_i} = (1/\gamma_{\psi_i}) \nabla \psi_i$, $\gamma_{\psi_i} = \|\nabla \psi_i\|$, and $\bar{n}_{\psi_i} = J n_{\psi_i}$. Notice that $\gamma_{\psi_i}|_{\mathcal{B}_i} = \frac{\rho_M - \|x - x_i\|}{\rho_M}$.

Consider the case of a single obstacle and the modified Lyapunov function

$$\varphi = V + \nu \psi, \quad (11)$$

where $\nu > 0$ encodes the relative importance of the spiral coverage task and the obstacle avoidance. Then, we set the control law

$$\dot{x} = -A \nabla \varphi = -A(\nabla V + \nu \nabla \psi). \quad (12)$$

With the dynamics (12), we have $\dot{\varphi} = -\alpha \|\nabla \varphi\|^2$. It can be shown that setting $\nu > R_0 - \rho_0$ implies that the obstacle function dominates the dynamics near an obstacle. Therefore, the dynamics accomplishes the coverage task while avoiding obstacles, with the possible exception of critical points where $\nabla \varphi = 0$, that can trap the robot. It suffices to show that any such critical points are not stable. Then the following result allows us to conclude that the system (12) is almost globally asymptotically stable.

Lemma 1 (Proposition 2.1 of [26]): Let φ be a twice differentiable Morse function (i.e., its Hessian, H_φ , is non-singular at every critical point) on a smooth manifold \mathcal{R} , with no boundary. Let f be a smooth vector field with $\dot{\varphi} = \nabla \varphi^T f \leq 0$, and assume the set of critical points, \mathcal{G}_φ , has a compact subset of local minima, $\mathcal{G}_V \subset \mathcal{G}_\varphi$, and a nonempty set of local maxima and saddle points, $\mathcal{G}_V - \mathcal{G}_\varphi \neq \emptyset$. Then, all initial conditions excluding a nowhere dense set, $\mathcal{G}_V - \mathcal{G}_\varphi$, asymptotically converge to \mathcal{G}_V . We conclude that \mathcal{G}_V is almost globally asymptotically stable.

Theorem 2: Consider an unconstrained planar agent and a point obstacle located at x_i with $\|x_i\| < \rho_1$, where $\rho_1 < R_0$, that limits the workspace to the set $\mathcal{R} = \mathcal{D} \setminus \{x_i\}$, and the obstacle potential ψ_i . Let $\mathcal{G}_V \subset \mathcal{R}$ be a compact goal set admitting a smooth Lyapunov function V satisfying $\mathcal{G}_V = \{x \in \mathcal{R} : \nabla V(x) = 0\}$ where $\mathcal{G}_V \cap \mathcal{B} = \emptyset$. For the combined Lyapunov function (11) and the control law (12), there exists a choice of ν and ρ_M guaranteeing almost global asymptotic stability of \mathcal{G}_V .

Proof: Let λ_φ be the smaller eigenvalue of $H_\varphi(x_c)$. Then, by Weyl's inequality,

$$\begin{aligned} \lambda_\varphi &\leq \|H_V(x_c)\| - \nu \left(\frac{1}{\|x_c - x_i\|} - \frac{1}{\rho_M} \right) \\ &\leq \frac{R_0}{\rho_0} - 1 - \nu \frac{\gamma_\psi(x_c)}{\|x_c - x_i\|} \\ &\leq \frac{R_0}{\rho_0} - 1 - \nu \frac{\gamma_V^{\min}}{\rho_M} \end{aligned}$$

where $\rho_M < \nu \gamma_V^{\min} / (R_0/\rho_0 - 1)$ guarantees that λ_φ is negative. Note that for $\|x\| < \rho_1$, $\gamma_V^{\min} > R_0 - \rho_1$, so the condition reduces to $\rho_M < \nu(R_0 - \rho_1) / (R_0/\rho_0 - 1)$. From Lemma 1, we conclude that \mathcal{G}_V is almost globally asymptotically stable. ■

Corollary 3: This result generalizes to the case of more than one obstacle when, for a pair of obstacles $x_i \in \mathcal{B}_i$ and $x_j \in \mathcal{B}_j$, $\mathcal{B}_i \cap \mathcal{B}_j = \emptyset$.

The implication of the theorem is that any spurious critical point introduced by an obstacle is unstable, so the control

law avoids getting stuck at obstacles. Since the task vector field f dominates away from the (small) obstacle regions, the control law accomplishes the coverage task with the possible exception of some small areas near obstacles and within the immediate vicinity of the origin. For practical purposes, it suffices to deploy the robot away from an obstacle.

C. Extension to multiple vehicles

In this initial work we have focused exclusively on the case of a single vehicle, but extensions to multiple vehicle control policies are straightforward. The stopping criterion remains the same, so the extension requires a multi-vehicle version of the control law (12) and a data association method to ensure the vehicles achieve an accurate count of $n_{\mathcal{D}}$. These can be developed by employing the methods of [27] and [28] for control, and [29] for data association.

III. SIMULATION RESULTS

In this section, we provide results of numerical simulations of the control law developed in the previous section. These simulations demonstrate the effectiveness of the control law in performing estimation of the Poisson point process parameter λ and the accuracy of the stopping criterion.

A. Results

Figure 3 shows the result of a simulation of the point process estimation control law (12) developed in Section II. The simulation was conducted with a realization of a Poisson point process with intensity parameter $\lambda = 0.1$ on a region $\mathcal{R} = [-20, 20] \times [-20, 20]$. The parameter values were set as follows: the sensor radius ρ and field of view width β were set to 1.1 and 2π , respectively, while the spiral pitch α was set to 1 and the reference radius R_0 to 40. Finally, the initial location x_0 was set to $(0.1\alpha, 0.1\alpha) = (0.1, 0.1)$, so $\rho_0 = 0.1$.




Fig. 3. An example simulation of the point process estimation law. The process has $\lambda = 0.1$ in the region $[-20, 20] \times [-20, 20]$, and the vehicle uses the control law (12) to register obstacles (in red) in order to estimate λ with $\varepsilon = 0.33$. Red dots represent the locations of the obstacles that have been observed and registered, while blue dots represent those yet to be observed. The dashed line shows the nominal spiral trajectory.

In the simulation shown in Figure 3, the repulsion strength parameter ν was set to 40, the obstacle radius ρ_M set to 0.3,

and the stopping criterion parameter ε was set to 0.33. The simulated trajectory is initially in reasonable agreement with the desired spiral trajectory; eventually the obstacle interactions cause the trajectories to diverge, but the controller still covers the space, as can be seen by observing that there are no unobserved points inside the simulated trajectory. When the stopping criterion is met, the vehicle's estimate of λ is $\hat{\lambda} = 0.1072$, which corresponds to a true relative error of 0.07195, well below the desired level of 0.33.

Table I demonstrates the effectiveness of the stopping criterion (6) for moderate values of α . The control law (12) was simulated 100 times for different realizations of the obstacle field, first with $\lambda = 0.2$ and $\alpha = 1$. For the stopping criterion, the desired relative error was set to $\varepsilon = 0.2$, so 95% of the simulations should have relative error $\varepsilon_r < 0.2$; in fact 97% did, so the stopping criterion is somewhat conservative. The mean estimate is $\hat{\lambda} = 0.2015$, resulting in a mean relative error of 0.733%.

Repeating the exercise with a sparser obstacle field $\lambda = 0.05$ and a wider sensor footprint $\alpha = 3$, we see that only 69% of simulations achieved the desired relative error and that the estimates $\hat{\lambda}$ exhibit significant systematic bias. This is due to approximating the area covered by the sensor as $A = \pi R^2$, where R is the current distance of the robot to the origin. When α is small relative to R , the approximation is accurate, but when α is large the approximation tends to underestimate the covered area, leading to an overestimate of $\hat{\lambda}$. In this case, a more accurate approximation of the area could correct the bias.

λ	α	# simulations	$\varepsilon_r < \varepsilon$	mean $\hat{\lambda}$	mean $\hat{\lambda}/\lambda - 1$
0.2	1.0	97		0.2015	0.733%
0.05	3.0	69		0.6331	26.62%

TABLE I

RESULTS FROM SIMULATING THE CONTROL LAW (12) FOR 100 DIFFERENT REALIZATIONS OF THE OBSTACLE FIELD.

B. Discussion

The simulation results show that the simple control law is effective at estimating the intensity parameter λ of a Poisson point process with the desired level of accuracy. However, the obstacle avoidance exacts a price by deforming the idealized spiral trajectory. We are able to ensure that the coverage property of the algorithm is robust to these deformations by following a tighter spiral than the sensor footprint strictly requires. This is costly because it requires covering a longer trajectory than otherwise necessary; a different obstacle avoidance method that resulted in less deformation in the reference trajectory could carry out the same task more quickly.

Finally, the stopping criterion (6) is somewhat conservative. As explained in the derivation in Section I-B, (6) requires sampling until the relative error ε_r is less than ε with high probability, approximately equal to 95%. As seen in the simulation examples, the mean relative error attained by the stopping criterion can be significantly less than ε , which means that the size of the sampling domain was larger than required to obtain a relative error of ε . One could be less

conservative by requiring $\varepsilon_r < \varepsilon$ with a more moderate probability. In general, the stopping criterion $R^* \geq \frac{n}{\varepsilon\sqrt{\pi\lambda}}$, where $n > 0$, results in stopping when $\varepsilon_r < \varepsilon$ with probability approximately equal to $\Phi(n) - \Phi(-n) = \text{Erf}(n/\sqrt{2})$, where $\Phi(\cdot)$ is the cumulative distribution function for the normal distribution and $\text{Erf}(\cdot)$ is the error function. The criterion (6) corresponds to the case $n = 2$.

IV. CONCLUSION

In this work we considered spatial point processes as models of the locations of objects of interest in a region. We then specialized to the Poisson point process, whose single parameter λ represents the mean density of these objects in the region, and studied the problem of deploying mobile robots as remote sensors to estimate λ . This problem combines aspects of the coverage problem for mobile robots with the requirement to avoid the objects of interest, which we treat as obstacles. We developed a criterion for deciding online the size of domain to be covered in order to achieve a desired level of precision in the estimate, and a control law that covers the domain while avoiding obstacles. We analyzed the control law and provided conditions under which it achieves the task. Through simulations, we showed that the estimation framework is effective; an implementation on a ground robotic platform is in progress and will be the subject of a future report.

This preliminary work focused on the Poisson point process because of its simplicity and its central role in the field of spatial point processes; its single parameter models the mean object density, i.e., first-order moment, of the object distribution. Future work will consider estimating higher-order moments of the distribution. We will also consider estimation problems for marked point processes, which are point processes where the mark is a real number associated with a point representing, e.g., size, type, or orientation. Such estimation problems arise in numerous applications [7].

Another interesting direction for future work would be to study probabilistically complete coverage in the point process framework. In this setting a control policy could use the real-time estimate of λ to adapt its coverage process such that it is guaranteed, e.g., to detect a fraction $1 - \epsilon$ of the points with probability $1 - \delta$.

REFERENCES

- [1] N. E. Leonard, D. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. E. Davis, "Collective motion, sensor networks, and ocean sampling," *Proc. IEEE*, vol. 95, no. 1, pp. 48–74, 2007.
- [2] N. E. Leonard, D. A. Paley, R. E. Davis, D. M. Fratantoni, F. Lekien, and F. Zhang, "Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in Monterey Bay," *J. of Field Robotics*, vol. 27, no. 6, pp. 718–740, 2010.
- [3] D. A. Paley and C. Peterson, "Stabilization of collective motion in a time-invariant flowfield," *J. of Guidance, Control, and Dynamics*, vol. 32, no. 3, pp. 771–779, 2009.
- [4] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Trans. Robot. Autom.*, vol. 20, no. 2, pp. 243–255, April 2004.
- [5] R. Graham and J. Cortés, "Adaptive information collection by robotic sensor networks for spatial estimation," *IEEE Trans. Autom. Control*, vol. 57, no. 6, pp. 1404–1419, 2012.

- [6] H.-L. Choi and J. How, "Efficient targeting of sensor networks for large-scale systems," *IEEE Trans. Control Syst. Tech.*, vol. 19, no. 6, pp. 1569–1577, Nov 2011.
- [7] S. N. Chiu, D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic Geometry and Its Applications*. John Wiley & Sons, Jun. 2013.
- [8] L. Gough, "The spatial ecology of an endemic desert shrub," Ph.D. eThesis, University of Nottingham, Jul. 2010.
- [9] I. O. McGlynn and G. S. Okin, "Characterization of shrub distribution using high spatial resolution remote sensing: Ecosystem implications for a former Chihuahuan Desert grassland," *Remote Sensing of Environment*, vol. 101, no. 4, pp. 554–566, Apr. 2006.
- [10] G. Okin, D. Gillette, and J. Herrick, "Multi-scale controls on and consequences of aeolian processes in landscape change in arid and semi-arid environments," *J. of Arid Environments*, vol. 65, no. 2, pp. 253–275, Apr. 2006.
- [11] D. Afshartous, Y. Guan, and A. Mehrotra, "US Coast Guard air station location with respect to distress calls: A spatial statistics and optimization based methodology," *European J. of Operational Research*, vol. 196, no. 3, pp. 1086–1096, 2009.
- [12] S. Trilles, P. Juan, L. Diaz, P. Arago, and J. Huerta, "Integration of environmental models in spatial data infrastructures: a use case in wildfire risk prediction," *IEEE J. Sel. Topics Appl. Earth Observ. in Remote Sens.*, vol. 6, no. 1, pp. 128–138, 2013.
- [13] R. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, Oct 2003.
- [14] M. Pace and P. Del Moral, "Mean-field PHD filters based on generalized Feynman-Kac flow," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 3, pp. 484–495, 2013.
- [15] C. S. Lee, D. E. Clark, and J. Salvi, "SLAM with single cluster PHD filters," in *IEEE Int. Conf. Robotics and Automation*, 2012, pp. 2096–2101.
- [16] P. Dames and V. Kumar, "Autonomous localization of an unknown number of targets without data association using teams of mobile sensors," *IEEE Trans. Autom. Sci. Eng.*, vol. 12, no. 3, pp. 850–864, July 2015.
- [17] S. Karaman and E. Frazzoli, "High-speed flight in an ergodic forest," in *IEEE Int. Conf. Robotics and Automation*, 2012, pp. 2899–2906.
- [18] C. E. Priebe, D. E. Fishkind, L. Abrams, and C. D. Piatko, "Random disambiguation paths for traversing a mapped hazard field," *Naval Research Logistics (NRL)*, vol. 52, no. 3, pp. 285–292, 2005.
- [19] D. E. Fishkind, C. E. Priebe, K. E. Giles, L. N. Smith, and V. Aksakalli, "Disambiguation protocols based on risk simulation," *IEEE Trans. Syst. Man, Cybern., A, Syst., Humans*, vol. 37, no. 5, pp. 814–823, 2007.
- [20] M. Roberts, T. Apker, B. Johnson, B. Auslander, B. Wellman, and D. W. Aha, "Coordinating robot teams for disaster relief," in *Proc. 28th Int. Florida Artificial Intelligence Research Society Conf.*, 2015, pp. 366–371.
- [21] H. Choset, "Coverage for robotics—a survey of recent results," *Annals of mathematics and artificial intelligence*, vol. 31, no. 1-4, pp. 113–126, 2001.
- [22] T. Kanit, S. Forest, I. Galliet, V. Mounoury, and D. Jeulin, "Determination of the size of the representative volume element for random composites: statistical and numerical approach," *International J. of solids and structures*, vol. 40, no. 13, pp. 3647–3679, 2003.
- [23] L. G. Valiant, "A theory of the learnable," *Communications of the ACM*, vol. 27, no. 11, pp. 1134–1142, Nov. 1984.
- [24] B. D. Ilhan, A. M. Johnson, and D. E. Koditschek, "Autonomous legged hill ascent," (*in preparation*).
- [25] E. U. Acar, H. Choset, A. A. Rizzi, P. N. Atkar, and D. Hull, "Morse decompositions for coverage tasks," *The Int. J. of Robotics Research*, vol. 21, no. 4, pp. 331–344, 2002.
- [26] D. E. Koditschek, "The application of total energy as a Lyapunov function for mechanical control systems," *Contemporary Mathematics*, vol. 97, pp. 131–158, February 1989.
- [27] M. Pavone and E. Frazzoli, "Decentralized policies for geometric pattern formation and path coverage," *J. of Dynamic Systems, Measurement, and Control*, vol. 129, no. 5, pp. 633–643, 2007.
- [28] C. S. Karagoz, H. I. Bozma, and D. E. Koditschek, "Coordinated navigation of multiple independent disk-shaped robots," *IEEE Trans. Robot.*, vol. 30, no. 6, pp. 1289–1304, 2014.
- [29] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and data fusion*. Storrs, CT: YBS publishing, 2011.