

Probably Approximately Correct Coverage for Robots with Uncertainty

Colin Das, Aaron Becker, and Timothy Bretl

Abstract—The classical problem of robot coverage is to plan a path that brings a point on the robot within a fixed distance of every point in the free space. In the presence of significant uncertainty in sensing and actuation, it may no longer be possible to guarantee that the robot covers all of the free space all the time, and so it becomes unclear what problem we are trying to solve. We will restore clarity by adopting a “probably approximately correct” measure of performance that captures the probability $1 - \epsilon$ of covering a fraction $1 - \delta$ of the free space. The problem of coverage for a robot with uncertainty is then to plan a feedback policy that achieves a given value of ϵ and δ . Just as solutions to the classical problem are judged by the resulting path length, solutions to our problem are judged by the required execution time. We will show the practical utility of our performance measure by applying it to several examples in simulation.

I. INTRODUCTION


The classical problem of robot coverage is to plan a path that—if followed exactly—would bring a point on the robot within a fixed distance of every point in the free space. The “fixed distance” in this problem is the size of the robot’s coverage implement, for example the radius of a mower blade. The “free space” (following standard convention) is the region of the workspace that is to be covered, for example a yard excluding trees and sidewalks, and in particular should not be confused with the free part of the robot’s configuration space. This coverage problem has been the focus of considerable research over the past three decades (e.g., see [1]–[25]) and is important for a variety of military, industrial, and domestic applications that include painting, demining, floor cleaning, and lawn mowing.

In many cases, we can assume that the path planned by a coverage algorithm is followed by the robot, if not exactly, then at least with negligible error. Such an assumption would be reasonable for an industrial robot arm or an autonomous tractor with centimeter-level GPS.

However, in the presence of significant uncertainty in sensing and actuation, the path followed by the robot may deviate from the path planned by the coverage algorithm. For example, a robot with consumer-grade GPS in an outdoor environment often has localization error variance that is bigger than the size of its coverage implement, causing parts of the free space to be missed and leading to performance of the sort shown in Fig. 1.

In principle, it is possible to improve performance by using a coverage algorithm that explicitly takes uncertainty into account and that produces a feedback policy rather than

C. Das and T. Bretl are with the Department of Aerospace Engineering and A. Becker is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA {cdas2,abecker5,tbretl}@illinois.edu



(a) coverage with negligible uncertainty in sensing and actuation



(b) coverage with significant uncertainty in sensing and actuation

Fig. 1. Simulation over a fixed time horizon of a robot trying to cover a rectangular free space by following a boustrophedon (square wave) path in the case of negligible uncertainty (top) and of significant uncertainty (bottom). Sensing and actuation uncertainty cause parts of the free space to be missed and, in general, lead to poor performance.

a nominal path. However, it is no longer clear what good performance even means in this context, because it may no longer be possible to guarantee that the robot covers all of the free space in finite time [26]. Prior work has suggested a variety of ways to measure performance, including by the expected fraction covered, the variance of fraction covered, and the probability that the entire free space is covered. Each of these measures is more or less appropriate for any given application, and none have been universally adopted. Our goal in this paper is to introduce a more general measure of performance that unifies this prior work.

In particular, we suggest a “probably approximately correct” (PAC) measure of performance that captures the probability $1 - \epsilon$ of covering a fraction $1 - \delta$ of the free space:

$$P(C \geq 1 - \delta) \geq 1 - \epsilon \quad \epsilon, \delta \in [0, 1].$$

The problem of coverage for a robot with uncertainty is then to plan a feedback policy that achieves a given value of ϵ and δ . Just as solutions to the classical problem are judged by the resulting path length, solutions to our problem are judged by the required execution time. We will consider a specific example system and will illustrate the practical utility of our performance measure by applying it to rank policies with respect to the probability of achieving a desired fraction covered, to compute the minimum time required to meet a given performance objective, and to choose the optimal global sensor for a robot.

The remainder of this paper proceeds as follows. We begin with a brief review of related work, focusing on existing

measures of performance for coverage under uncertainty (Section II). Then, we present our alternative PAC measure of performance (Section III), and consider its application to several examples in simulation (Section IV). We discuss broader implications in our concluding remarks (Section V).

II. RELATED WORK

We give a brief overview of robotic coverage and highlight research on coverage under forms of uncertainty. We also survey coverage statistics relevant to robotics.

A. Robotic Coverage

Robotic coverage algorithms typically use a form of *cellular decomposition* which breaks down the free space into cells whose union fills (or approximately fills) the free space. Planners with an exact cellular decomposition achieve a provable guarantee of complete coverage if the robot visits each cell. Exact cellular decomposition can be achieved using on-board sensing via the *boustrophedon* (a square wave path) decomposition [11] or Morse functions [15], both of which rely on the identification of *critical points* [12], points where the connectivity of a slice moving across the workspace changes. In these works the goal is to construct a minimum-length path that visits and covers each cell. Choset et al. [21] showed that the minimum path length is bounded linearly by the area of the free space, the number of critical points, and the perimeter lengths of the obstacles and outer boundary. These planners assume the presence of an accurate global localization sensor or accurate odometry. We are interested in coverage when these sensors are uncertain. In this work, we modify the boustrophedon path so that the reference trajectory extends beyond the boundary and has path spacing less than the width of the coverage implement. These are heuristic methods of dealing with uncertainty.

While Acar et al. [12], [15] achieve complete coverage under sensing uncertainty, their focus is on rejecting bad sensor readings. They assume that most readings are good and that bad readings are exceptional. This paper considers systems where all sensor readings and actuations are corrupted by Gaussian noise. Algorithms such as in [27] and [28] deal with robotic uncertainty, but consider the navigation problem—moving from one configuration to another—not the coverage problem.

For random-reflection, Acar et al. [18] analyzes policy performance in detecting 80% of mines found versus the total mean search time. The data represents performance over a discrete number of points, i.e, mines, in the free space and is expressed by mean and standard deviation values. If the mines are uniformly distributed, this problem is a special case of the coverage problem we consider.

Finally, this work focuses on systems without a *completion detector*—that is, systems that cannot sense when the coverage task has been completed. Systems with a completion detector, such as a limited number of mines to find [18], perfect localization [14], “ant”-type robots that leave a detectable trail over covered cells [19], or an overhead monitoring system [23] can terminate with guaranteed coverage.

B. Coverage Statistics

Prior art has investigated coverage for many end purposes, resulting in a variety of performance objectives for coverage policies. The ideal performance measure is the probability density function (PDF) of coverage for a given policy because any other statistic can be constructed from the PDF. This ideal has so far proved elusive. World War II led to a surge in research funding for coverage related to search and rescue, reconnoissance, bombing, weapon salvos, and mine-sweeping. This research generated the expected coverage probabilities for both random and boustrophedon search paths. Later work led to additional statistics about the probability distributions for coverage tasks.

As Stone showed in [29], the expected fraction covered $\mathbb{E}[C(t)]$ of a region with area A by a random strategy in time t is

$$\mathbb{E}[C(t)] = 1 - e^{-\frac{2rvt}{A}} \quad (1)$$

for velocity v and a coverage implement radius r . They also considered search via boustrophedon paths under three simplifying assumptions: each pass is long enough that the paths may be considered independent, the paths are corrupted by a zero-mean Gaussian noise with standard deviation σ , and the number of passes is large compared to the width of the area to be covered. With these assumptions, the function $\mathbb{E}[C(r, S, \sigma)]$ is the expected fraction covered after following a boustrophedon path with path centers spaced S units apart. It is computed by first calculating $g(i, x)$, the probability of a point at x being covered by the i -th pass; next finding $h(x)$, the total probability of coverage at x after parallel passes $i = -\infty$ through $i = \infty$; and finally computing the average coverage by integrating $h(x)$ over one path width.

$$\begin{aligned} g(i, x) &= \frac{1}{\sqrt{2\pi}} \int_{(x-iS-r)/\sigma}^{(x-iS+r)/\sigma} e^{-\frac{1}{2}z^2} dz \\ h(x) &= 1 - \prod_{i=-\infty}^{\infty} (1 - g(i, x)) \\ \mathbb{E}[C(r, S, \sigma)] &= \frac{1}{S} \int_{-S/2}^{S/2} h(x) dx \end{aligned} \quad (2)$$

Fig. 2 shows the expected fraction covered versus time for a random path and for a boustrophedon path with several values of path spacing and no uncertainty ($\sigma = 0$). Time is expressed in units of *nominal coverage* Ψ , the time required to fully cover an equal area with no overlap. Fig. 3 shows the expected fraction covered versus uncertainty σ for the same path spacings as in Fig. 2. As σ increases, the $\mathbb{E}[C(\cdot)]$ converges to that of a random strategy traveling an equal distance. With (1) or (2) we can use the Markov inequality to find a loose upper bound of the form $P(C \geq x) \leq \mathbb{E}[C]/x$. To obtain tighter bounds, we need additional information.

The probabilities for vacancy and coverage of a space in \mathbb{R}^n by random n -D spheres is a topic of considerable interest to the mathematics community [26], [30], but the full probability distribution function is only known for the 1-D case. We focus on 2-D results because of their relation




Fig. 2. Expected fraction covered by a random strategy (dashed blue) and four boustrophedon paths with different path spacing (S) and zero uncertainty ($\sigma = 0$). The markers indicate the expected fraction covered by a random strategy after executing for equal times as the boustrophedon paths.

to robotic coverage. The mean and variance for coverage of a plane under the *torus convention*¹ by discs whose centers are distributed randomly are

$$\begin{aligned} \mathbb{E}[C] &= 1 - e^{-\frac{n\pi r^2}{A}} \\ \sigma^2(C) &= A\pi e^{-\frac{2n\pi}{A}} \left(8 \int_0^1 \left(e^{\frac{nB(x)}{A}} - 1 \right) x dx - \frac{n\pi}{A} \right) \quad (3) \end{aligned}$$

where

$$B(x) = \begin{cases} 0 & \text{if } x > 1 \\ -x\sqrt{1-x^2} + \arccos(x) & \text{otherwise} \end{cases}$$

As the ratio of disc size to region area decreases, the contribution by the torus convention asymptotes to zero.

Simulation shows that (3) accurately describes the performance of a random policy. Indeed, the expected values given by (1) and (3) are identical. Having the variance allows us to obtain tighter bounds on coverage estimates using Chebyshev's inequality. Unfortunately, such results do not exist for the boustrophedon policy.

The probability of completely covering a region by randomly placing *shapes* is given in [31]. Shapes are placed isotropically with centers inside the region. Complete coverage has a Bernoulli distribution with success probability p , which increases as the number of shapes n increases, the region area A decreases, or the areas a of the shapes increase. The expected number of gaps (uncovered areas) less the number of isles (unattached covered areas) can be computed exactly:

$$\Phi = e^{-\Psi} \left(\Psi \left(\frac{sSa}{2\pi} - \frac{\chi A}{a} \right) + \frac{\Psi^2 As^2}{4\pi a^2} + X \right) \quad (4)$$

Here s and S are the perimeter lengths of a and A , $\Psi = na$ is the nominal coverage if the n shapes did not overlap, and χ and X are the Euler characteristic [32] of the shapes a and A , which for simple shapes without holes is 1.

¹The *torus convention* treats a rectangular area as a torus such that any disc that protrudes out one side of the rectangle enters again from the opposite side.

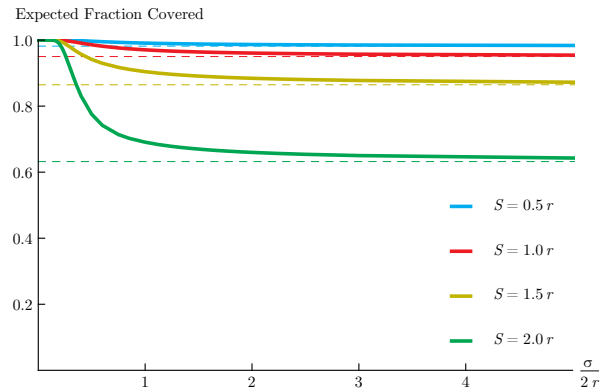


Fig. 3. Expected fraction covered versus the ratio of the uncertainty (σ) to the coverage implement size ($2r$) for boustrophedon paths with different path spacing (S). As the uncertainty increases (x -axis), the expected fraction covered converges to that of a random strategy executing for equal time (the markers in Fig. 2).




Fig. 4. Complete coverage results for simulated random-reflection and (4). Solid lines represent 500 simulations of the random-reflection policy for three region sizes with a coverage implement radius r . Dashed lines show coverage on the same regions by placing $2r \times 1m^2$ rectangles randomly with centers inside the region. Random-reflection outperforms a purely random strategy because random-reflection produces a continuous path.

For large n , the probability for isles becomes vanishingly small, and Φ represents the expected number of gaps. Complete coverage occurs when there are zero gaps. A valid approximation of complete coverage for large n is $1 - e^{-\Phi}$. This function compares reasonably with a random-reflection policy, as shown in Fig. 4. Random-reflection slightly outperforms IID placement of equal-sized shapes because random-reflection produces a continuous path. Sadly, no similar results on the probability of complete coverage exist for a boustrophedon policy.

III. PERFORMANCE MEASURE

In the presence of significant uncertainty in sensing and actuation, it may no longer be possible to guarantee that the robot covers all of the free space all the time, and this leads us to redefine what is meant by coverage. The breadth of the coverage problem with uncertainty is illustrated by two coverage tasks: painting and demining. Painting represents tasks that are considered a success only if the region is completely covered. Demining represents tasks where success is achieved when a robot guarantees that some fraction of the free space has been covered. A ready example of such a criterion is given by the UN mine-clearing standard of 1996,

requiring 99.6% of an area be checked for mines in order to be “cleared” [33]. Our performance measure encompasses the competing demands for complete coverage and fraction covered. Now, painting and demining are points on a 2-D spectrum of coverage tasks.

Previous work on coverage provides several performance measures such as the expected fraction covered for a given distance traveled (1)–(3) and the probability of complete coverage (4). These measures provide only a single point of comparison. A better performance measure would allow us to compare the entire probability distribution.

We define such a measure by:

$$P(C \geq 1 - \delta) \geq 1 - \varepsilon \quad (5)$$

where C is the fraction covered and $\varepsilon, \delta \in [0, 1]$. This is a “probably approximately correct” measure of performance that captures the probability $1 - \varepsilon$ of covering a fraction $1 - \delta$ of the free space [34], [35]. Certain values of ε, δ correspond naturally with common robotic coverage tasks. A painting robot, required to cover an entire space with probability $1 - \varepsilon$, has $\delta = 0$. In contrast, a demining robot that achieves success when it is guaranteed to cover a fraction $1 - \delta$, has $\varepsilon = 0$.

For any policy, the cumulative distribution function (CDF) captures the full range of performance. For a given system and policy, the CDF can either be constructed directly from the generating probability distribution or approximated by Monte Carlo analysis. Because the generating probability distributions are unknown for 2-D, as shown in Section II-B, we construct the CDF from Monte Carlo simulations.

The relationship between the PAC objective and the CDF is shown by the following transformation:

$$\begin{aligned} P(C \geq 1 - \delta) &\geq 1 - \varepsilon \\ \int_{1-\delta}^1 \text{PDF}(\tau) d\tau &\geq 1 - \varepsilon \\ 1 - \int_0^{1-\delta} \text{PDF}(\tau) d\tau &\geq 1 - \varepsilon \\ 1 - \text{CDF}(1 - \delta) &\geq 1 - \varepsilon \end{aligned}$$

Equivalently,

$$\text{CDF}(1 - \delta) \leq \varepsilon. \quad (6)$$

This is the constraint that for a given δ , $\text{CDF}(1 - \delta)$ must be less than or equal to ε . A desired performance, characterized by an ε, δ pair, places a constraint rectangle on the CDF from $(0, 1)$ to $(1 - \delta, \varepsilon)$. To satisfy such a constraint, the CDF must not cross the boundary of this rectangle. Fig. 5 shows an example of this type of constraint.

Multiple ε, δ constraints may be concatenated to fully specify the desired coverage task. For instance, if we must never have less than 50% covered and also require at least 70% covered in 8 out of 10 trials, our constraint set is $\{(\varepsilon_1 = 0, \delta_1 = 0.5) \cup (\varepsilon_2 = 0.2, \delta_2 = 0.3)\}$. This set is illustrated in Fig. 9. The most stringent constraint set ($\delta = 0, \varepsilon = 0$) is achieved by systems with negligible error. The random-reflection and boustrophedon policies achieve this constraint




Fig. 5. Cumulative distribution functions for three policies with an execution time $t_f = 4\Psi$ in a rectangular region ($W = 80\text{m}, H = 5\text{m}$). The grey rectangle depicts the desired performance constraint of ($\varepsilon = 0.5, \delta = 0.02$)—the robot must “cover at least 98% of the free space with probability at least 50%”. The boustrophedon short policy exceeds the requirement, but the random-reflection and boustrophedon long policies fail. All CDFs represent 500 simulations, each with process noise $\sigma_p = 0.1\text{m}$ and measurement noise $\sigma_m = 1.0\text{m}$.

with probability 1 as time goes to infinity. In practice we are interested in tight constraints for finite-horizons.

As (6) shows, the CDF gives an ε -value for every δ , generating a performance curve $\varepsilon(\delta)$. It is now possible to formulate well-posed problems, e.g., finding a policy that maximizes some objective function.

One example is to find a policy π that maximizes the probability of achieving a desired fraction covered given a system and an execution time:

$$\arg \min_{\pi \in \{\pi_1, \pi_2, \dots, \pi_n\}} \varepsilon(\pi) \Big|_{\delta, t_f}. \quad (7)$$

Another example is to find the minimum execution time that meets some desired performance. That is, given a policy, a desired ε, δ pair, and a system

$$\text{minimize } t_f \text{ such that } \text{CDF}(1 - \delta) \leq \varepsilon. \quad (8)$$

The PAC performance measure is a tool for analyzing feedback policies for coverage in systems with sensing and actuation uncertainty.

IV. EXAMPLES

We apply the PAC measure to a simplified system under three feedback policies: random-reflection, boustrophedon long, and boustrophedon short. All three policies are employed in practice [29], [36]. We begin by describing the robotic system and the candidate policies, then describe three applications. The policies are tested on identical free spaces to allow unbiased comparison.

A. System

We consider discrete time, bounded control systems with white Gaussian noise in a bounded 2-D free space. Our process model and prediction is described by

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{x}_k + \mathbf{u}_k \Delta t + \mathbf{w}_k \\ \hat{\mathbf{x}}_{k+1|k} &= \hat{\mathbf{x}}_{k|k} + \mathbf{u}_k \Delta t \\ \sigma_{k+1|k}^2 &= \sigma_{k|k}^2 + \sigma_p^2 \end{aligned}$$

where $\mathbf{x}_k, \mathbf{x}_{k+1} \in \mathbb{R}^n$ are system states at times $k, k + 1$, $\|\mathbf{u}_k\| \leq v$ the control input, $\mathbf{w}_k \sim \mathcal{N}(0, \sigma_p^2)$ zero-mean Gaussian process noise, and $\hat{\mathbf{x}}_{k+1|k}, \sigma_{k+1|k}^2$ the estimated




Fig. 6. A representative path of the random-reflection policy. The robot moves until it detects a boundary, turns to a uniformly random heading, and repeats the process until time t_f .

state and variance. The measurement model at time $k + 1$ is given by

$$\mathbf{z}_{k+1} = \mathbf{x}_{k+1} + \boldsymbol{\nu}_{k+1} \quad (9)$$

where $\boldsymbol{\nu}_{k+1} \sim \mathcal{N}(0, \sigma_m^2)$ is the zero-mean Gaussian measurement noise. The Kalman update is then

$$\begin{aligned} K &= \sigma_{k+1|k}^2 \left(\sigma_{k+1|k}^2 + \sigma_m^2 \right)^{-1} \\ \hat{\mathbf{x}}_{k+1|k+1} &= \hat{\mathbf{x}}_{k+1|k} + K \left(\mathbf{z}_{k+1} - \hat{\mathbf{x}}_{k+1|k} \right) \\ \sigma_{k+1|k+1}^2 &= (1 - K) \sigma_{k+1|k}^2. \end{aligned}$$

For constant σ_p^2 and σ_m^2 , the steady-state localization variance is solved by setting $\sigma_{k+1|k+1}^2 = \sigma_{k|k}^2$, giving

$$\sigma^2 = \frac{\sigma_p^2}{2} \left(-1 + \sqrt{1 + 4(\sigma_m/\sigma_p)^2} \right).$$

We are particularly interested in systems with σ^2 such that a robot attempting to follow a nominal path is not guaranteed to achieve complete coverage. The robot is also equipped with a boundary sensor to prevent it from exiting the free space.

We consider obstacle-free, rectangular workspaces with width W and height H . The free space is a bounded subset of \mathbb{R}^2 with no holes. This class of free space has a simple topology and isolates problems due to uncertainty from problems due to narrow passages as in [37].

B. Policies


A policy

$$\pi : \mathbf{x}_0, \mathbf{z}_0 \cdots \mathbf{z}_k, t_k \rightarrow \mathbf{u}_k$$


assigns the current input \mathbf{u}_k given the initial condition \mathbf{x}_0 and the measurement history $\mathbf{z}_0 \cdots \mathbf{z}_k$. We consider three policies in our experiments.

1) *Random Reflection*: A common industry approach to coverage under uncertainty is random-reflection [18]. It requires no a priori information of the environment nor global localization sensors, but does require a boundary sensor. The robot chooses a uniformly random heading and attempts to move straight with speed v until it detects a boundary at which point the process repeats until time t_f . A representative path is shown in Fig. 6.

2) *Boustrophedon (Long and Short)*: The boustrophedon long and short policies are motivated by [11]. They command the robot to follow a square wave trajectory as shown in Figs. 7(a) and 7(b). The boustrophedon long policy assigns a reference trajectory of square wave motions with passes



(a) Boustrophedon Long Policy



(b) Boustrophedon Short Policy

Fig. 7. The boustrophedon path for policies (2) and (3). The geometry of the path is parameterized by path spacing S and overshoot b , which is the same in both x - and y -directions. The robot computes its state estimate using a KF and tracks its trajectory with an LQR controller.

aligned with the long axis of the boundary. The boustrophedon short policy is similar, but with passes aligned with the short axis of the boundary. The trajectory is parameterized by path spacing S and by path overshoot b . Upon completion of the reference trajectory, the robot retraces it in reverse until time t_f .

The robot relies on measurements from the global localization sensor described in (9) and uses its boundary sensor to remain in the workspace. The robot uses a Linear Quadratic Regulator (LQR) controller as an optimal dynamic regulator about the reference trajectory and a Kalman Filter as a state estimator.

C. Application

We demonstrate the practical utility of our performance measure by applying it to three examples. For each example we run 500 simulations following the policies in Section IV-B. In each simulation the robot executes a policy until complete coverage. We record the actual fraction covered at each time increment, then construct the CDF from this data set. This free space is 400m^2 and the robot has an implement radius $r = 0.15\text{m}$ and a maximum velocity $v = 0.5\text{m/s}$, yielding a nominal coverage $\Psi = 43.74\text{min}$. Ψ is the time required for a robot to cover an equal area perfectly—with no process or measurement noise and no overlap. The boustrophedon reference trajectories use path spacing $S = 0.15\text{m}$ and overshoot $b = 0.6\text{m}$.

1) *Ranking Candidate Policies*: In this example, we use ε, δ curves to compare the three policies from Section IV-B. Policies are ranked by their probability of achieving a desired fraction covered of 97.5% ($\delta = 0.025$). The system is the same for all cases with process noise $\sigma_p = 0.1\text{m}$ and measurement noise $\sigma_m = 1.0\text{m}$ for a region with width $W = 5\text{m}$ and height $H = 80\text{m}$. Applying (7) for $t_f = 4\Psi$, the boustrophedon short policy performs better than random-reflection and boustrophedon long policies, implying that in




Fig. 8. Performance curves (lower is better) for three policies in a system with uncertainty ($\sigma_p = 0.1, \sigma_m = 1.0$) in a rectangular region ($W = 5\text{m}, H = 80\text{m}$). Policies are ranked by their probability of achieving $\delta = 0.025$ (black vertical line) with an execution time $t_f = 4\Psi$ using (7). The boustrophedon short policy has the highest probability ($\varepsilon = 0.01$), followed by boustrophedon long ($\varepsilon = 0.36$), and random-reflection ($\varepsilon = 0.59$).




Fig. 9. Cumulative distribution functions for execution time $t_f = \{0.25, 0.5, 1, 2, 4\}\Psi$ for a random-reflection policy under process noise $\sigma_p = 0.1$. The grey rectangles depict the desired performance constraints of $(\varepsilon_1 = 0.0, \delta_1 = 0.5) \cup (\varepsilon_2 = 0.2, \delta_2 = 0.3)$ —the robot must “cover 50% or more of the free space every time and 70% or more with probability at least 80%.” A search yields that $t_f = 1.33\Psi$ (58min) is the minimum time to achieve the desired performance.

many cases it is better to align passes with the short axis of a boundary than with the long axis. This is clear from Fig. 8 which shows the performance curves of each policy for $t_f = \{1, 2, 4\}\Psi$.

2) *Optimizing Path Length:* To minimize wear-and-tear on the robot, it is often desirable to know the minimum execution-time to achieve some performance. Consider a random-reflection policy with t_f the design parameter and a desired performance of $(\varepsilon_1 = 0.0, \delta_1 = 0.5) \cup (\varepsilon_2 = 0.2, \delta_2 = 0.3)$ —that the robot must “cover 50% of the free space every time and 70% or more with probability at least 80%.” In Fig. 9 we show the performance for $t_f = \{0.25, 0.5, 1, 2, 4\}\Psi$. The first three do not meet the requirement; $t_f = \{2, 4\}\Psi$ exceed the requirement. Applying (8) for the constraint set optimizes this policy. The minimum time that achieves the desired performance is $t_f = 1.33\Psi$ (58min). The same analysis may be extended to optimize other policy parameters.

3) *Robot Sensor Selection:* Sensor selection is a relevant problem in robot design. In this example, we examine the effect of increasing measurement accuracy σ_m for the boustrophedon long policy and compare it to random-reflection, independent of σ_m . Fig. 10 shows results for a region with width $W = 10\text{m}$, height $H = 40\text{m}$, process noise $\sigma_p = 0.1$, and measurement noise $\sigma_m = \{0.1, 0.5, 1.0, 2.5, 10\}$. We compare strategies at execution times $t_f = \{0.5, 1.5, 2.5\}\Psi$. For $t_f = \{0.5, 1.5\}\Psi$, the

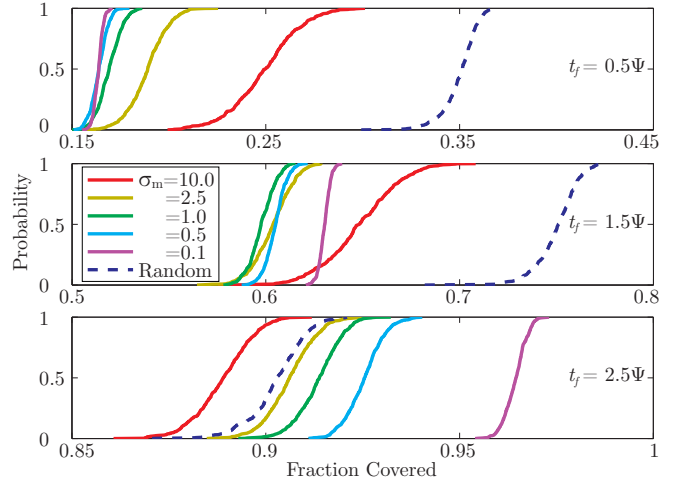


Fig. 10. Cumulative distribution functions for execution times $t_f = \{0.5, 1.5, 2.5\}\Psi$ under process noise $\sigma_p = 0.1$ and varying measurement noise σ_m . These results are compared to a random-reflection strategy, independent of measurements. For short execution times $\{0.5, 1.5\}\Psi$, the random strategy performs best, followed by high measurement noise, and low measurement noise. As t_f increases boustrophedon paths result in better coverage—performance increases as σ_m decreases.

random strategy performs best. For $t_f = 2.5\Psi$, the boustrophedon policy with $\sigma_m = \{0.1, 0.5, 1.0, 2.5\}$ results in better performance than random-reflection (for most values of δ). These results imply that for small t_f a random strategy performs best. This is expected because random policies tend to explore more than boustrophedon policies, giving them an early advantage. This dependence is also shown in Figs. 2 and 8. For larger t_f decreasing values of σ_m yield better performance for the boustrophedon long policy. Modifying (7) to vary over candidate σ_m values—perhaps from a catalog of sensors—is straightforward. Applying (7) for a given execution time and fraction covered determines the optimal sensor selection. This analysis provides a tool for robot design and design trade-offs, e.g., cost and accuracy of a sensor versus time to achieve a desired performance objective.

V. CONCLUSION

Many existing coverage policies are provably complete in that they guarantee a robot visits all points in the free space. However, these policies assume negligible localization error and perform differently in systems with uncertainty in sensing and actuation. In this work we provided a performance objective to evaluate policies under such systems, which enabled us to form well-posed problems. We demonstrated the usefulness of this problem formulation by applying it to an example system and several policies. In particular, we showed that under certain levels of uncertainty boustrophedon policies outperform random policies. We also showed how this performance objective can be used to optimize policy and system design. As an example, we showed that orienting boustrophedon paths in a rectangular free space parallel to the short axis yields better performance than orienting perpendicular to the long axis.

The CDFs for the three policies in this work can be approximated by two-parameter logistic (sigmoid) functions, $(1 + e^{-c_1(x+c_2)})^{-1}$. Future work should classify policies based on this or a related parameterization, enabling ε, δ comparisons in a functional form. In this work we applied the PAC objective to three policies in rectangular free spaces using a global localization sensor. Future lines of inquiry include applying the PAC objective to investigate other policies, to optimize policies (e.g., the S and b boustrophedon parameters), to explore topologically interesting free spaces, and to integrate information from other sensors. In particular, the PAC objective could provide a useful benchmark for multi-robot coverage algorithms.

ACKNOWLEDGMENT

Thanks to S. Hutchinson and M. Johnson for helpful discussion, A. Akce for assistance with simulation, and the anonymous reviewers for valuable comments. This material is based upon work supported by the National Science Foundation under Grant Nos. 0931871 and 0955088.

REFERENCES

- [1] R. A. Brooks, "A robust layered control system for a mobile robot," *IEEE J. Robot. Autom.*, vol. 2, no. 1, pp. 14–23, 1986.
- [2] Z. L. Cao, Y. Huang, and E. L. Hall, "Region filling operations with random obstacle avoidance for mobile robots," *J. Robotic Systems*, vol. 5, pp. 87–102, 1988.
- [3] A. Zelinsky, R. Jarvis, J. C. Byrne, and S. Yuta, "Planning paths of complete coverage of an unstructured environment by a mobile robot," in *In Proceedings of International Conference on Advanced Robotics*, 1993, pp. 533–538.
- [4] D. W. Gage, "Many-robot MCM search systems," in *Proceedings of Autonomous Vehicles in Mine Countermeasures Symposium*, 1995, pp. 9–55.
- [5] S. Hert, S. Tiwari, and V. Lumelsky, "A terrain-covering algorithm for an auv," *Autonomous Robots*, vol. 3, pp. 91–119, 1996.
- [6] D. Kurabayashi, J. Ota, T. Arai, S. Ichikawa, S. Koga, H. Asama, and I. Endo, "Cooperative sweeping by multiple mobile robots with relocating portable obstacles," in *Proceedings of IROS 1996*, vol. 3, Nov. 1996, pp. 1472–1477.
- [7] E. M. Arkin, S. P. Fekete, and J. S. B. Mitchell, "Approximation algorithms for lawn mowing and milling," *Computational Geometry*, vol. 17, no. 1-2, pp. 25–50, 1997.
- [8] R. De Carvalho, H. Vidal, P. Vieira, and M. Ribeiro, "Complete coverage path planning and guidance for cleaning robots," in *Industrial Electronics, 1997. ISIE '97., Proceedings of the IEEE International Symposium on*, vol. 2, Jul. 1997, pp. 677–682.
- [9] E. Gelenbe and Y. Cao, "Autonomous search for mines," *European Journal of Operational Research*, vol. 108, no. 2, pp. 319–333, 1998.
- [10] T. W. Min and H. K. Yin, "A decentralized approach for cooperative sweeping by multiple mobile robots," in *Proceedings of IROS 1998*, vol. 1, Oct. 1998, pp. 380–385.
- [11] H. Choset, "Coverage of known spaces: The boustrophedon cellular decomposition," *Autonomous Robots*, vol. 9, pp. 247–253, 2000.
- [12] E. U. Acar and H. Choset, "Robust sensor-based coverage of unstructured environments," in *Proceedings of IROS 2001*, Oct. 2001, pp. 61–68.
- [13] E. U. Acar, H. Choset, and P. N. Atkar, "Complete sensor-based coverage with extended-range detectors: a hierarchical decomposition in terms of critical points and Voronoi diagrams," in *Proceedings of IROS 2001*, 2001, pp. 1305–1311.
- [14] H. Choset, "Coverage for robotics - a survey of recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, pp. 113–126, 2001.
- [15] E. U. Acar, H. Choset, A. A. Rizzi, P. N. Atkar, and D. Hull, "Morse decompositions for coverage tasks," *I. J. Robotic Res.*, vol. 21, no. 4, pp. 331–344, 2002.
- [16] I. Latimer, D., S. Srinivasa, V. Lee-Shue, S. Sonne, H. Choset, and A. Hurst, "Towards sensor based coverage with robot teams," in *Proceedings of ICRA 2002*, vol. 1, 2002, pp. 961–967.
- [17] C. Luo and S. Yang, "A real-time cooperative sweeping strategy for multiple cleaning robots," in *Proceedings of the 2002 IEEE International Symposium on Intelligent Control*, 2002, pp. 660–665.
- [18] E. U. Acar, H. Choset, Y. Zhang, and M. J. Schervish, "Path planning for robotic demining: Robust sensor-based coverage of unstructured environments and probabilistic methods," *I. J. Robotic Res.*, vol. 22, no. 7-8, pp. 441–466, 2003.
- [19] J. Svennebring and S. Koenig, "Building terrain-covering ant robots: A feasibility study," *Autonomous Robots*, vol. 16, no. 3, pp. 313–332, 2004.
- [20] S. S. Ge and C. Fua, "Complete multi-robot coverage of unknown environments with minimum repeated coverage," in *Proceedings of ICRA 2005*, Apr. 2005, pp. 715–720.
- [21] H. Choset, K. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. Kavraki, and S. Thrun, *Principles of Robot Motion*. The MIT Press, 2005.
- [22] I. Rekleitis, A. New, and H. Choset, "Distributed coverage of unknown/unstructured environments by mobile sensor networks," in *Multi-Robot Systems. From Swarms to Intelligent Automata Volume III*, L. Parker, F. Schneider, and A. Schultz, Eds. Springer Netherlands, 2005, pp. 145–155.
- [23] J. Pearce, B. Powers, C. Hess, P. Rybski, S. Stoeter, and N. Papanikolopoulos, "Using virtual pheromones and cameras for dispersing a team of multiple miniature robots," *Journal of Intelligent & Robotic Systems*, vol. 45, pp. 307–321, 2006.
- [24] Y. Guo and M. Balakrishnan, "Complete coverage control for non-holonomic mobile robots in dynamic environments," in *Proceedings of ICRA 2006*, May 2006, pp. 1704–1709.
- [25] J. Rogge and D. Aeyels, "Sensor coverage with a multi-robot system," in *Intelligent Control, 2007. ISIC 2007. IEEE 22nd International Symposium on*, Oct. 2007, pp. 71–76.
- [26] P. G. Hall, *Introduction to the theory of Coverage Processes*. New York, New York: Wiley Series in Probability and Mathematical Statistics, 1988.
- [27] M. Erdmann, "Using backprojections for fine motion planning with uncertainty," *I. J. Robotic Res.*, vol. 5, no. 1, pp. 19–45, 1986.
- [28] T. Lozano-pérez, M. T. Mason, and R. H. Taylor, "Automatic synthesis of fine-motion strategies for robots," *I. J. Robotic Res.*, pp. 60–99, 1984.
- [29] L. D. Stone, "Search theory: A mathematical theory for finding lost objects," *Mathematics Magazine*, vol. 50, no. 5, pp. 248–256, 1977.
- [30] H. Solomon, *Geometric probability*. Philadelphia : Society for Industrial and Applied Mathematics, 1978.
- [31] A. M. Kellerer, "On the number of clumps resulting from the overlap of randomly placed figures in a plane," *J. Applied Probability*, vol. 20, no. 1, pp. 126–135, Mar. 1983.
- [32] M. A. Armstrong, *Basic Topology*. Springer-Verlag, 1983.
- [33] I. C. on Mine Clearance Technology, "Report of the international conference on mine clearance technology," in *International Conference on Mine Clearance Technology*, Elsinore, Denmark, Jul. 1996.
- [34] L. G. Valiant, "A theory of the learnable," in *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, ser. STOC '84. New York, NY, USA: ACM, 1984, pp. 436–445.
- [35] D. Haussler, "Probably approximately correct learning," in *Proceedings of the Eighth National Conference on Artificial Intelligence*. AAAI Press, 1990, pp. 1101–1108.
- [36] J. L. Jones and P. R. Mass, "Method and system for multi-mode coverage for an autonomous robot," Patent Pending, Sep. 2007, 20070213892.
- [37] D. Hsu, T. Jiang, J. Reif, and Z. Sun, "The bridge test for sampling narrow passages with probabilistic roadmap planners," in *Proceedings of ICRA 2003*, vol. 3, September 2003, pp. 4420 – 4426.