Stabilizing information-driven exploration for bearings-only SLAM using range gating

Robert Sim
Department of Computer Science
University of British Columbia
2366 Main Mall,
Vancouver, BC V6T 1Z4
simra@cs.ubc.ca

***Abstract*— This paper examines the problem of information-driven exploration for the purposes of simultaneous localization and mapping (SLAM) with a bearings-only sensor. In recent work, we have demonstrated that employing an information-driven approach to exploration with an extended Kalman filter (EKF) can drive drive the robot to locations in the world where filter updates are ill-conditioned and linearization constraints are violated, potentially destabilizing the filter, and increasing the probability of divergence from the true state estimate. In this paper, we demonstrate an information-driven approach to exploration that preserves the stability of the EKF and produces maps that are significantly more accurate than a conventional information-driven approach. Our method is based on range-gating observations so as to avoid potentially destabilizing updates. We provide simulated experimental results demonstrating the superior performance of our approach over simple outlier gating and over heuristic-driven exploration.**

***Index Terms*— Bearings-only SLAM, Exploration, Extended Kalman Filter**

I. Introduction

This paper considers the problem of optimally controlling a mobile robot as it explores an environment and constructs a map. Until only recently, the control problem has been largely ignored in favor of solving the underlying map construction problem, usually referred to as simultaneous localization and mapping (SLAM). The key problem in SLAM is that a robot's actions and observations are noisy, and as such the robot can never know its precise position, or the precise positions of features in the environment. Solutions to SLAM generally amount to simultaneously maximizing the probability of the robot's trajectory estimate and inferred landmark positions, conditioned on its actions and observations.

With recent successes in formalizing the SLAM problem, and the development of convergent solutions, it now becomes feasible to inquire as to how a robot should behave as it explores. In particular, we are interested in optimizing the quality of the map that the robot produces. Very few authors have considered this problem to date. The widely-accepted method is to apply an information-optimal approach to data acquisition, driving the robot and its sensors to locations that maximize the expected information to be gained from acquiring an observation at that location [1]–[4].

In recent work, we demonstrated that for a robot with a bearings-only sensor (that is, a sensor which can observe only the direction to features of interest, and not the distance, one such sensor being a monocular camera), the performance of an information-optimal approach to exploration suffers compared to other (theoretically sub-optimal) approaches. Furthermore, we identified the reason for this poor performance as being the fact that information-optimal viewpoints also tend to be locations where the EKF update is ill-conditioned and linearization constraints are more likely to be violated. Intuitively, the robot is driven to locations where the uncertainty of its observation is high (such as moving close to uncertain landmarks), but these regions of high uncertainty are also regions where small changes in observation will result in large changes to the robot's state (both map and pose) estimate. This poses a significant problem for the EKF, as incorporating observations that lie outside the (potentially narrow) operating point of the filter's linearization can result in divergence.

In our previous work, we presented simulation results indicating the poor performance of an information-optimal exploration policy compared to a heuristic policy that simply traced the Voronoi graph of the estimated map [5]. Despite the fact that tracing the Voronoi graph is sub-optimal in the information theoretic sense, this heuristic policy succeeded by maintained a 'safe' distance from all of the landmarks in the map, thereby preventing catastrophic divergence.

There are two weaknesses to our previous work. First, despite the fact that the Voronoi-based approach achieves good results, it is based on a sub-optimal heuristic, and as the robot incorporates new observations into its map, the Voronoi graph can change significantly, resulting in an erratic exploration trajectory. Second, one might argue that an information-driven approach to exploration would be more successful by applying simple outlier gating to the observations.

This paper has two main contributions. First, we demonstrate that simple outlier removal does not significantly improve the performance of an information-driven exploration policy. Second, we present an alternative approach to observation filtering, which we refer to as *range gating*, that, when used in conjunction with an information-optimal exploration policy, out-performs a Voronoi-based approach to exploration. These results are important for devising autonomous systems that require accurate maps, and in particular when the best map possible is required within certain time constraints. We will present experimental results demonstrating the performance of our range-gated approach in a simulated environment.

The remainder of this paper proceeds as follows. First, we will examine the bearings-only SLAM problem and present the conventional information-driven approach to exploration. This will be followed by a presentation of our range gating method. We will conclude with a series of experiments illustrating the approach and a discussion of the results.

II. Related Work

The bearings-only SLAM problem was considered by Deans and Hebert from the context of filter design and selection [6]. Recent work, such as that by Kwok and Dissanayake and Solà *et al.* has focused on the landmark initialization problem, as incorrect initializations can cause filter divergence [7], [8].

A related problem to the exploration problem is that of bearings-only target tracking, in which an optimal control mechanism must be determined to determine the position of a fixed target using a bearings-only sensor [9], [10]. These approaches usually assume that the position of the robot can be determined exactly.

Several authors have examined the exploration problem using a variety of different sensing modalities [1]–[4], [11], [12]. These all generally aim to maximize information gain over time. In other related work, Kwok and Fox used a reinforcement learning strategy for controlling a robot in a way that optimized its state estimate for the purposes of conducting a specific task, in this case, kicking a ball towards a goal [13].

III. Information-driven Exploration using the EKF

The extended Kalman Filter has been widely deployed for SLAM [14], [15], although most implementations assume a range-and-bearings sensor. The state of the system at time t is generally described as a vector $\mathbf{x}_t = [\mathbf{x}_t^r \ \mathbf{l}_1 \ \dots \ \mathbf{l}_n]^T$, where $\mathbf{x}_t^r = [x \ y \ \phi]$ describes the pose of the robot in a planar environment, and $\mathbf{l}_i$ describes the positions of n landmarks in the world, all expressed in a global coordinate frame. Because sensor observations $\mathbf{z}_t$ and robot actions $\mathbf{u}_t$ are generally noisy, a probabilistic framework is applied to the state estimation problem. In the Kalman Filter, the probability of the state $\mathbf{x}_t$, conditioned on the sequence of robot actions $A = \{\mathbf{u}_1, \dots, \mathbf{u}_t\}$ and observations $Z = \{\mathbf{z}_1, \dots, \mathbf{z}_t\}$ is approximated as a Gaussian distribution:

$$p(\mathbf{x}|A, Z) \approx k \exp\{(\mathbf{x} - \hat{\mathbf{x}})^T P^{-1} (\mathbf{x} - \hat{\mathbf{x}})\} \quad (1)$$

with the mean $\hat{\mathbf{x}}$ representing a maximum-likelihood state estimate with covariance P, and k is a normalizing constant.

As the robot performs actions (that is, moves through the environment), the pose distribution is propagated according to a *plant*, or *motion* model:

$$\mathbf{x}' = f(\mathbf{x}, \mathbf{u}) + \nu(\mathbf{u}) \quad (2)$$

which describes the noisy outcome of the robot's actions $\mathbf{u}$, where $\mathbf{v}$ is normally distributed noise.

As observations are taken, the map and pose of the robot are updated using a *measurement* model

$$\mathbf{z} = h(\mathbf{x}) + w(\mathbf{x}) \quad (3)$$

describing the relationship between poses and landmark observations, where $w(\cdot)$ is normally distributed noise. For the EKF, $f(\cdot)$ and $h(\cdot)$ might be non-linear functions. In the bearings-only case, $h(\cdot)$ describes the direction to each of the landmarks in range of the sensor:

$$h_l(\mathbf{x}_t^r) = \tan^{-1} \frac{y - l_y}{x - l_x} - \phi \quad (4)$$

where $\mathbf{x}_t^r = [x \ y \ \phi]$ is the pose of the robot at time t and $l = [l_x \ l_y]$ is the position of a landmark.

State updates for a given observation $\mathbf{z}$ are determined by defining the innovation $\mathbf{v}$ and computing the associated measurement covariance S:

$$\begin{aligned} \mathbf{v} &= \mathbf{z} - h(\hat{\mathbf{x}}) && (5) \\ S &= E\left[\mathbf{v}(t+1)\mathbf{v}^T(t+1)\right] && (6) \\ &= \nabla h P \nabla h^T + R && (7) \end{aligned}$$

where R is a covariance matrix describing the sensor noise model and $E[x]$ indicates the expectation of the random variable x. ∇h is the Jacobian of $h(\cdot)$. The innovation describes the extent to which the current observation differs from what the robot expects to see from its current pose. Given the innovation $\mathbf{v}$, and covariance S the state estimate is updated by first computing the *Kalman gain* W:

$$W = P \nabla h^T S^{-1} \quad (8)$$

and applying W to transform the innovation into a state displacement:

$$\hat{\mathbf{x}} := \hat{\mathbf{x}} + W\mathbf{v} \quad (9)$$

and covariance

$$P := P - WSW^T. \quad (10)$$

Of particular interest for the exploration problem is how to minimize the uncertainty of the landmark positions over time. There are a variety of measures for quantifying the map uncertainty. For this paper, we compute the sum of the determinants of the block-2 matrices describing the individual landmark covariances:

$$Err(P) = \sum_{i=1}^{n} |P_i| \quad (11)$$

where each P_i is a 2x2 sub-matrix of P_t corresponding to landmark i. Alternative measures include computing the trace or determinant of P_t [11].

It has been demonstrated elsewhere that, if the robot has perfect localization, a locally optimal strategy for data collection is to drive the robot to positions that maximize the *prediction variance* $|S|$ of the observation (Equation 7) [1], [2].

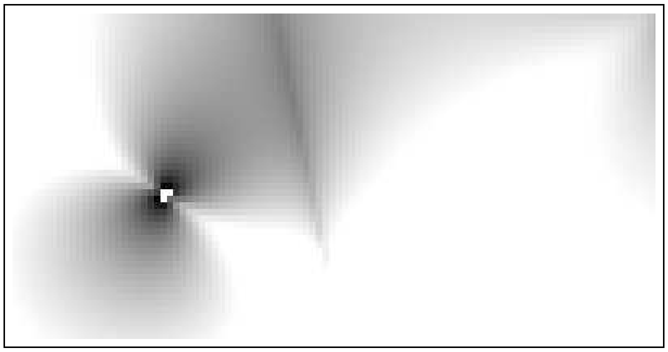

Fig. 1. Expected information gain as a function of position, accounting for noise injected due to robot motion. Darker poses correspond to more informative destinations. The initial robot pose is located at the bottom center of the image. A landmark is located at the center of the 'hole' in the peak (the hole being due to the minimum range of the sensor). The diagonal ridge reflects the current heading of the robot, where motions that don't induce a rotation (and hence don't inject rotation-dependent noise) are preferred.

Note that S is a function of both the state estimate $\hat{\mathbf{x}}$ and the map covariance P. Maximizing $|S|$ moves the robot to locations in the world where the least information is known about the observation. For example, in the bearings-only case, it is advisable to move the robot to take an observation from a direction orthogonal to the principal direction of a landmark's uncertainty covariance. The second point to note is that, all other things being equal, the maximally informative pose will be one that maximizes the determinant of the gradient covariance $\nabla h \nabla h^T$ (by setting P to identity and keeping R constant). Put simply, the robot should move to locations where the observation changes rapidly as a function of pose. In the bearings-only case, this amounts to moving as close as possible to the landmark.

These results are complicated by the fact that for the SLAM problem the robot's pose is not exactly known. $|S|$ can always be increased by simply increasing the uncertainty of the robot's position, which is clearly not desirable. While computing the optimal trajectory analytically seems difficult, if not intractable in this case, one can compute an "information surface" numerically by simulating robot actions and observations and examining their effects on the posterior covariance P [11].

Formally, the informativeness of an action $\mathbf{u}$ is described by:

$$H(\mathbf{u}) = -(Err(P') - Err(P)) \tag{12}$$

where P' is the state covariance obtained by simulating action $\mathbf{u}$ and then simulating an observation and state update. The results of one such simulation are depicted in Figure 1. In this figure, darker poses correspond to more informative destinations. These results indicate that, even in the presence of pose uncertainty, the maximally informative actions move the robot as close as possible to the landmark under observation.

IV. Range Gating

In our prior work, we demonstrated that poses corresponding to maximal information gain simultaneously maximize the condition number of the EKF state update system. Specifically, when the robot takes an observation $\mathbf{z}$ the state estimate $\hat{\mathbf{x}}$ is updated by by solving the linear system

$$S\tilde{\mathbf{x}} = \mathbf{v} \tag{13}$$

through Equations 8 and 9. Here, S is the observation covariance, determined by Equation 7 and the solution vector $\tilde{\mathbf{x}}$ is a displacement that will be subsequently projected into the state space through $P\nabla h^T \tilde{\mathbf{x}}$. Hence, the stability of the EKF update is dependent on the conditioning of the linear system defined in Equation 13. Two quantities contribute to the stability of this system. First, the configuration of visible landmarks plays a key role, and second, the pose of the robot relative to this configuration is also important.

In our previous work, we computed the conditioning of this system analytically for the two-landmark case (i.e. two landmarks are visible) and provided numerical results for special instances of the three- and four-landmark cases. In each case, the condition number of Equation 13 is maximized as the robot approaches any given landmark. It should also be noted that these configurations also correspond to regions where the linearization constraints introduced in the EKF approximation are more likely to be violated.

One possible option for stabilizing the filter update is to perform outlier gating, such as that employed in [16]. Specifically, when a landmark observation is obtained, we compute the log likelihood of the observation based on the current state estimate:

$$\log p(\mathbf{z}_i|\hat{\mathbf{x}}) = \mathbf{v}_i^T S_i^{-1} \mathbf{v}_i \tag{14}$$

Estimates whose log-likelihood exceeds a user-defined threshold of g^2 can be removed:

$$\mathbf{v}_i^T S_i^{-1} \mathbf{v}_i > g^2 \tag{15}$$

The difficulty with taking this approach is two-fold. First, when the robot closes a large loop, it may be the case that the observations of landmarks that come into view will deviate significantly from the expected value. If these observations are not incorporated into the filter update, the robot cannot successfully close the loop. Secondly, as a robot moves close to a landmark, the measurement covariance S will naturally grow, so that even observations that deviate significantly from their expected values will be more probable.

It has been observed that destabilizing observations correspond to those obtained from nearby landmarks. Our approach to outlier removal is to ignore these nearby observations. To accomplish this we employ a virtual sensor that has a user-defined minimum range. When observations arrive, the virtual sensor first uses the observed landmark's current state estimate to determine whether that landmark might be inside the sensor's minimum range (Figure 2). Observations derived from these nearby landmarks are ignored. We refer to this method as *range gating*.

Our gating approach will have an impact on the robot's map estimates and the information gained from observing from

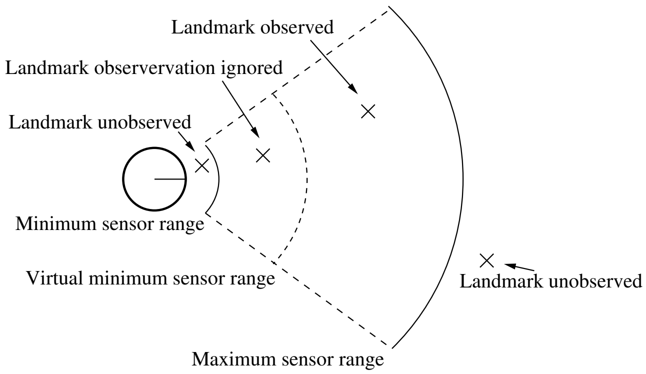

Fig. 2. The virtual sensor ignores landmark observations that derive from landmarks that are too close to the robot. Since the sensor is a bearings-only sensor, this decision is based on the current state estimate.

new locations. Therefore, when we estimate the information gain for any particular action, we assume that the robot is taking observations with the virtual sensor, rather than its real sensor. This ensures that the actual outcome of a robot's exploration is consistent with its predictions.

In the remaining sections, we will experimentally compare the performance of our range gating method using information-driven exploration with a variety of other approaches.

V. Exploration Policies

Our exploration approach will operate under the following assumptions. First, the world is populated with a set of n landmarks, whose positions are initially unknown (but the value of n is assumed to be known). The environment is free of obstacles. The EKF will maintain the pose $[x\ y\ \phi]$ of the robot, and the landmark positions $[x_1\ y_1 \ldots x_n\ y_n]$.

At each time step, the robot executes an action u, followed by an observation z. The observation z returns (noisy) bearing measurements to those landmarks in the environment that are within the range of the robot's (real) sensor $[\texttt{min\_range}, \texttt{max\_range}]$. Data association is assumed to be perfect, and the sensor is assumed to have a full 360 degree field of view. When a landmark is observed for the first time, its position is initialized in the filter to be located at the mean of the sensor range, and its covariance initialized to have a standard deviation of half the sensor range:

$$\sigma_r = \frac{\texttt{max\_range} - \texttt{min\_range}}{2} \tag{16}$$

In real-world applications with obstacles, these sensor ranges may not be so easily estimated, and other initialization schemes may be required. For the experiments presented here, one might consider initializing the landmark position at the maximum sensor range, since the robot will typically see landmarks for the first time when the robot moves within range of them. However, our experiments suggest that the current initialization scheme is reasonable.

We will compare the performance of five exploration policies, as described below:

- *Random*: The robot drives to successive random poses in the environment.
- *InformationGain*: The robot drives directly to the globally optimal position for maximizing information gain. The global maximum is found by hill-climbing from each landmark estimate, as well as from the robot's current pose. Observations are not filtered.
- *OutlierGated*: Like *InformationGain* above, the robot maximizes information gain but filters outlier observations according to Equation 15, where $g = 1.5$.
- *RangeGated*: Like *InformationGain* above, the robot maximizes information gain but filters observations using a virtual sensor with a user defined minimum range (defined below).
- *Voronoi*: The robot traces out the Voronoi graph (VG) defined by the landmarks. Specifically, the robot attempts to visit each junction of the VG at least once by following routes that pass between nearby landmarks.

In our previous work, the *Voronoi* approach out-performed *InformationGain* and *Random* by a significant margin.

A. Coverage

Our exploration policies (with the exception of *Random*) assume that the robot has a list of landmarks to consider, either for computing the VG or for locally optimizing information gain. However, some landmarks may not yet be discovered. In order to ensure coverage of the environment so that each landmark is discovered, the world is initially populated with a set of dummy landmarks, as described in [5]. As each landmark is observed for the first time, a dummy landmark is removed. Planning uses the set of known and dummy landmarks. When the robot moves to a pose where it expects to observe a dummy, and fails to do so, the dummy is relocated to an unexplored region of space. While this approach can have drawbacks, it drives exploration and in the limit will guarantee that all landmarks are observed.

It should be noted that none of the exploration strategies considered here employ any special mechanisms for successful loop closing. In our view, successful loop closing should be an emergent property of the exploration policy and should not require handling as a special case. It is also worth noting that the policies considered make strictly local decisions in time– no attempt is made to consider plans that optimize over the long term. This latter problem remains a rich area of study (see, for example [11], [12]).

VI. Experimental Results

We have run our exploration policies in a simulated environment with a variety of settings. For all experiments, the map region was confined to a 200m by 200m plane. For each trial, a map containing 20 landmarks was generated by randomly sampling uniformly from the map region. To ensure coverage, we initially employed the dummy landmark mechanism described above.

For individual trials, each run consisted of 2000 time steps in which an action consisted of a rotation to a desired heading

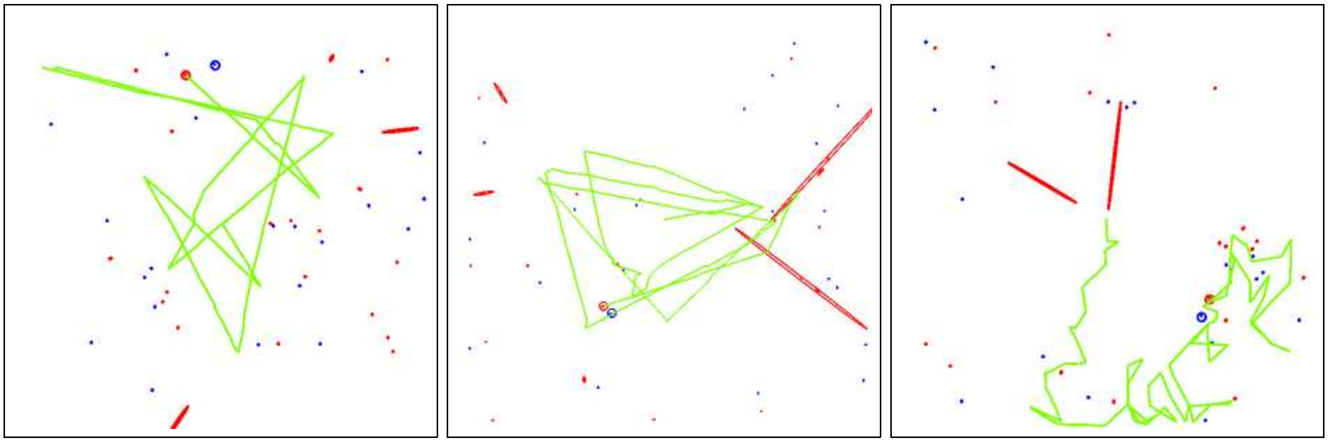

(a) (b) (c)

Fig. 3. Selected Exploration Policies: a) *RandomPose*, b) *InformationGain*, c) *Voronoi*. Refer to the text for details.

followed by a translation of at most 1m. Note that this motion model is non-linear. The robot's maximum sensor range was 40m and the minimum sensor range was 2m. The minimum range of the virtual sensor employed for range gating was 10m. Finally, three observation noise models were employed. All three models assumed that bearing observations were zero-mean normally distributed with a variance that is constant with respect to landmark distance. The standard deviations of the noise were set to 1 degree, 5 degrees and 10 degrees respectively.

For each exploration policy, and for each observation noise model, a total of 100 trials was conducted. Figure 3 illustrates a typical map constructed by the main types of exploration strategy. The estimated trajectory of the robot is marked and the landmark estimates are plotted, along with (sometimes elongated) ellipses indicating the landmark covariances.

Figure 4 depicts the results from our experiments. The first plot indicates the mean total error in the map (ground truth landmarks versus landmark estimates). Since the map is invariant to rigid transformations, the estimated map is first corrected by a global rotation about the robot's starting pose to bring it to closest correspondence to the actual map. The second graph plots the mean number of landmarks discovered by each exploration policy over the 100 trials. There are two reasons why all 20 landmarks might not be discovered in any particular trial: the policy might spend too much time in explored regions to cover all the unexplored space within 2000 time steps, or, the policy might be susceptible to divergence, so that the robot is incapable of successfully navigating to unexplored regions of space.

It is interesting to note that in most cases the accuracy of the map improves as sensor noise increases. This is due to the fact that the filter is more brittle when the sensor noise is low–behaviors that violate the underlying Gaussian and linearity assumptions become more problematic for the filter. In our experimentation, we assumed that the sensor noise model is known, but even in such situations our results demonstrate that there is an argument to be made for over-estimating the sensor noise.

These experimental results clearly indicate the success of our range gating method over both the Voronoi-based approach and traditional outlier gating. It should be noted, however, that the Voronoi-based approach results in improved coverage of the pose space, due to its systematic traversal of the map (the information-driven approaches will spend more time in regions where the map is uncertain).

VII. Conclusion

We have presented an approach to information-driven exploration for bearings-only SLAM that successfully overcomes stability issues that are inherent to the EKF update. Specifically, we employ a virtual sensor that performs range gating, removing observations that are expected to pose problems for the filter update. In experimental results simulating an exploring robot, we demonstrated that this approach produces maps that are significantly more accurate than maps produced using simple outlier gating or a heuristic exploration approach based on tracing the Voronoi graph of the observed landmarks.

Despite these successes, there are several directions for future work. First, while our simulated results illustrate the numerical properties of our exploration approach, it remains to be tested on a real robot. In addressing a real-world scenario, we will also need to consider issues in data association, and unknown noise models in the sensor and robot's actuators. Second, our approach employs a user-defined threshold for gating observations. This threshold was set according to the user's intuitive sense of what observations should be excluded, and depends on both the sensor's characteristics and the robot's odometric noise model.

An alternative approach to gating is to employ an outlier detector where the threshold is based on raw deviation in the

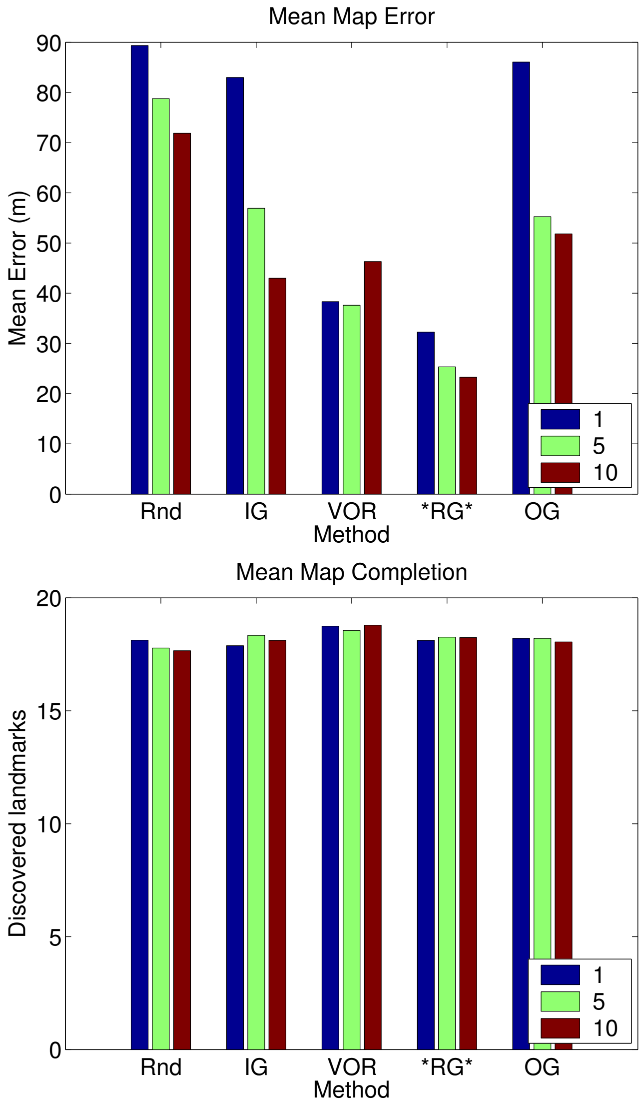

Fig. 4. Mean mapping error, and map coverage, for 100 trials. Bars are reported by observation noise models of 1, 5 and 10 degree standard deviation. The horizontal axis is sorted by method: Rnd: *Random*, IG: *InformationGain*, VOR: *Voronoi*, RG: *RangeGated*, and OG: *OutlierGated*

bearings observation (that is, a threshold set at k degrees, rather than k standard deviations from the mean). It may be possible to analytically derive an optimal threshold setting based on the deviation of the linearized sensor model from the underlying non-linear process. Future work will continue to address these questions.

REFERENCES

[1] D. MacKay, "Information-based objective functions for active data selection," *Neural Computation*, vol. 4, no. 4, pp. 590–604, 1992.

[2] P. Whaite and F. P. Ferrie, "Autonomous exploration: Driven by uncertainty," in *Proceedings of the IEEE Computer Society Conf. on Computer Vision and Pattern Recognition*. Los Alamitos, CA, USA: IEEE Computer Society Press, June 1994, pp. 339–346.

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

[4] C. Stachniss and W. Burgard, "Exploring unknown environments with mobile robots using coverage maps," in *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, Acapulco, Mexico, 2003.

[5] R. Sim, "Stable exploration for bearings-only SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, IEEE. Barcelona, Spain: IEEE Press, April 2005.

[6] M. Deans and M. Hebert, "Experimental comparison of techniques for localization and mapping using a bearings only sensor," in *Proc. of the ISER '00 Seventh Int. Symposium on Experimental Robotics*, December 2000.

[7] N. Kwok and G. Dissanayake, "An efficient multiple hypothesis filter for bearing-only slam," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*. Sendai, Japan: IEEE Press, October 2004, pp. 736–741.

[8] J. Solà, A. Monin, M. Devy, and T. Lemaire, "Undelayed initialization in bearing only slam," in *Proc. Int. Conf on Intelligent Robots and Systems (IROS)*, IEEE/RSJ. Edmonton, AB: IEEE Press, August 2005.

[9] J. A. Fawcett, "Effect of course maneuvers on bearings-only range estimation," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 36, no. 8, pp. 1193–1199, August 1988.

[10] Y. Oshman and P. Davidson, "Optimization of observer trajectories for bearings-only target localization," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 35, no. 3, pp. 892–902, July 1999.

[11] R. Sim and N. Roy, "Global a-optimal robot exploration in SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, IEEE. Barcelona, Spain: IEEE Press, April 2005.

[12] S. Huang, N. M. Kwok, G. Dissanayake, Q. P. Ha, and G. Fang, "Multi-step look-ahead trajectory planning in SLAM: Possibility and necessity," in *Proc. Int. Conf. on Robotics and Automation*, IEEE. Barcelona, Spain: IEEE Press, April 2005, pp. 1103–1108.

[13] C. Kwok and D. Fox, "Reinforcement learning for sensing strategies," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*. Sendai, Japan: IEEE Press, October 2004.

[14] J. J. Leonard and H. J. S. Feder, "A computationally efficient method for large-scale concurrent mapping and localization," in *Robotics Research: The Ninth Int. Symposium*, J. Hollerbach and D. Koditschek, Eds. London: Springer-Verlag, 2000, pp. 169–176.

[15] J. Guivant, E. Nebot, and H. Durrant-Whyte, "Simultaneous localization and map building using natural features in outdoor environments," in *Sixth Int. Conf. on Intelligent Autonomous Systems*, vol. 1. Italy: IOS Press, 2000, pp. 581–588.

[16] J. J. Leonard and H. F. Durrant-Whyte, "Mobile robot localization by tracking geometric beacons," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 3, pp. 376–382, 1991.