Stable Exploration for Bearings-only SLAM

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

***Abstract*— Recent work on robotic exploration and active sensing has examined a variety of information-theoretic approaches to efficient and convergent map construction. These involve moving an exploring robot to locations in the world where the anticipated information gain is maximized. In this paper we demonstrate that, for map construction using bearings-only information and the Extended Kalman Filter (EKF), driving exploration so as to maximize expected information gain leads to ill-conditioned filter updates and a high probability of divergence between the inferred map and reality. In particular, we present analytical and numerical results demonstrating the effects of blindly applying an information-theoretic approach to bearings-only exploration. Subsequently, we present experimental results demonstrating that an exploration approach that favours the conditioning of the filter update will lead to more accurate maps.**

I. Introduction

This paper considers the problem of designing active exploration policies that maximize the accuracy of a robot's map as it moves through an unknown environment. Autonomous exploration is a problem that has only recently gained attention, largely due to the recent success of the robotics community at formalizing the map construction problem, otherwise known 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 and inferred landmark positions, conditioned on its actions and observations.

Very few authors have considered the problem of how a robot can best navigate through the world so as to maximize the quality of its map. Where results do exist they often assume perfect robot localization (such as the rich literature on bearings-only target tracking [1], [2]). Furthermore, the extant results tend to emphasize an information-optimal approach to exploration, where the sensor is driven to locations that maximize the expected information to be gained from acquiring an observation at that location [3], [4], [5], [6].

This paper makes two important contributions. First, it presents an analysis demonstrating that information-optimal approaches to active exploration can be detrimental to map quality, particularly when coupled with approximating assumptions that are common in SLAM approaches (for example, those applied by the extended Kalman Filter and other Gaussian filters), and sensing modalities that tend towards ill-posed localization estimates. In the extreme, such approaches can be prone to filter divergence, rendering the map useless and the robot effectively lost. Second, this paper presents an alternative approach to exploration that optimizes a map's accuracy by taking a policy that emphasizes the conditioning of the map update step. It will be seen that such an approach results in a highly conservative exploration strategy that greatly reduces the probability of divergence.

We will apply our analysis in the context of performing SLAM with an omnidirectional bearings-only sensor, using the EKF for state estimation. Most SLAM approaches consider the case where the robot is equipped with a sensor that can measure both range and bearings to landmarks in the world. However, there are a wide variety of contexts, such as monocular vision sensing, where the robot has access only to a bearing measurement to each observable landmark. In addition, bearings-only SLAM solutions using the EKF and related filters demonstrate a higher tendency to diverge, due primarily to the ill-posedness of initializing landmark positions from one or two early observations, as well as the underlying non-linear relationship between landmark positions and observations. A limited number of authors have considered this problem, primarily from the perspective of filter design [7].

II. Bearings-only SLAM using the EKF

The extended Kalman Filter has been widely deployed for range-and-bearings SLAM approaches [8], [9]. The state of the world at time t is generally described as a vector $\mathbf{x}_t = [\mathbf{x}_t^r \ \mathbf{l}_1 \ \ldots \ \mathbf{l}_n]^T$, 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 actions $A = \{\mathbf{u}_1, \ldots, \mathbf{u}_t\}$ and observations $Z = \{\mathbf{z}_1, \ldots, \mathbf{z}_t\}$ is approximated as a Gaussian distribution:

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

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 $\mathbf{x}' = f(\mathbf{x})$ describing the uncertain outcome of the robot's actions, and as observations are taken, the map and pose of the robot are updated using a *measurement* model $\mathbf{z} = h(\mathbf{x})$ describing the relationship between poses and observations. For the EKF, the outcomes of these processes are assumed to be normally distributed.

Of particular interest to the exploration problem is how to minimize the uncertainty of the landmark positions over time. This is typically measured as $|P_t|$, the determinant of the state covariance at time t. It has been demonstrated elsewhere ([3], [4]) 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:

$$S = \nabla\mathbf{h} P \nabla\mathbf{h}^T + R \quad (2)$$

where $\nabla\mathbf{h}$ is the gradient of the observation function $h(\cdot)$, evaluated at the current state estimate $\hat{\mathbf{x}}$ and R is the covariance matrix describing the sensor's intrinsic noise characteristics.

Note that S is a function of both the state estimate μ and the map covariance P. Maximizing $|S|$ has the effect of moving 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 $\nabla\mathbf{h}\nabla\mathbf{h}^T$ (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 also parallel results in bearings-only target tracking, which move the sensor so as to maximize the determinant of the Fisher information matrix (FIM) J_f of the set of observations Z [2]:

$$J_f = -E[\frac{\partial}{\partial \mathbf{l}^2} \log p(Z|\mathbf{l})] \quad (3)$$

where $\mathbf{l}$ is the inferred target location.

These results are complicated somewhat by the fact that, in the SLAM problem, the robot's pose is not exactly known. Maximizing $|S|$ would imply maximizing the uncertainty of the robot's pose as well, 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 [10]. The results of one such simulation are depicted in Figure 1. These results indicate that, even in the presense of pose uncertainty, the maximally informative actions move the robot as close as possible to the landmark under observation.

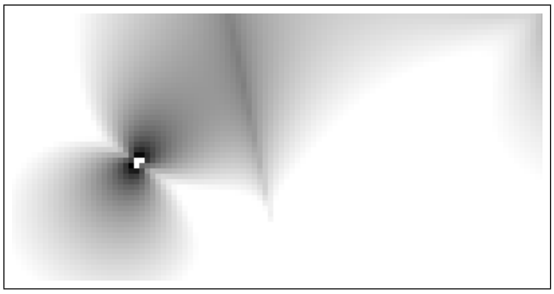

Fig. 1. Expected information gain as a function of position, accounting for noise injected due to robot motion. Darker poses correspond to more informative destinations. The initial robot pose is located at the bottom center of the image. A landmark is located at the center of the 'hole' in the peak (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 are preferred.

The next section will examine the conditioning of the filter update step and demonstrate that the information-optimal approach to computing exploration strategies is prone to divergence.

III. The Conditioning of the Observation Update

In this section we will examine the conditioning of the update step in the EKF. When the robot takes an observation $\mathbf{z}$ the state estimate $\hat{\mathbf{x}}$ is updated by weighting the *innovation* $\mathbf{v} = \mathbf{z} - h(\hat{\mathbf{x}})$ by the Kalman gain W as follows:

$$W = P\nabla\mathbf{h}^T S^{-1} \quad (4)$$

$$\hat{\mathbf{x}}' = \hat{\mathbf{x}} + W\mathbf{v} \quad (5)$$

Here, W transforms the innovation into a state displacement.

The EKF update equations effectively compute the log-likelihood of the observation by solving the linear system

$$S\tilde{\mathbf{x}} = \mathbf{v} \quad (6)$$

through Equations 4 and 5. Here, S is the observation covariance, determined by Equation 2 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 6. 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.

We will compute the conditioning of this system analytically for the two-landmark case. Below, we will demonstrate similar results numerically for three- and four-landmark configurations.

The bearings-only observation function is given by

$$\mathbf{z} = h(\mathbf{x}) = [\tan^{-1}(\frac{\Delta y_1}{\Delta x_1}) - \phi \ \ldots \ \tan^{-1}(\frac{\Delta y_n}{\Delta x_n}) - \phi]^T \quad (7)$$

where Δx_i and Δy_i are the x and y displacements between the robot's pose and the observed landmark i, respectively, and ϕ is the orientation of the robot. This yields a gradient function

$$\nabla h(\mathbf{x}) = [\frac{\partial h}{\partial \mathbf{x}^{r_t}} \frac{\partial h}{\partial \mathbf{l}_1} \dots \frac{\partial h}{\partial \mathbf{l}_n}]^T \quad (8)$$

$$= \begin{bmatrix} \tilde{y}_1 & -\tilde{x}_1 & -1 & -\tilde{y}_1 & \tilde{x}_1 & 0 & 0 \\ \tilde{y}_2 & -\tilde{x}_2 & -1 & 0 & 0 & -\tilde{y}_2 & \tilde{x}_2 \end{bmatrix} \quad (9)$$

where $\tilde{x}_i$ is $\Delta x_i / r_i^2$, $\tilde{y}_i$ is $\Delta y_i / r_i^2$, with $r_i^2 = \Delta x_i^2 + \Delta y_i^2$.

In order to simplify the analysis, we will assume $P = I$. That is, the state covariance is diagonal and all uncertainties are equal. In the general case, of course, this is not the case. However, this simplifying assumption yields results of interest that can extend to the general case. In addition, we will assume that $R = diag(\sigma^2)$, a constant noise model, which is not unreasonable for the bearings-only case. With this assumption in mind,

$$S = \nabla h \nabla h^T + R \quad (10)$$

$$= \begin{bmatrix} 2(\tilde{y}_1^2 + \tilde{x}_1^2) + 1 + \sigma^2 & \tilde{y}_1 \tilde{y}_2 + \tilde{x}_1 \tilde{x}_2 + 1 \\ \tilde{y}_1 \tilde{y}_2 + \tilde{x}_1 \tilde{x}_2 + 1 & 2(\tilde{y}_2^2 + \tilde{x}_2^2) + 1 + \sigma^2 \end{bmatrix} \quad (11)$$

$$= \begin{bmatrix} a & b \\ b & d \end{bmatrix} \quad (12)$$

The condition number of S is equal to the ratio of the larger to the smaller eigenvalue of S, λ_1/λ_2, which evaluates to

$$\frac{\lambda_1}{\lambda_2} = \frac{a + d + \sqrt{(a+d)^2 - 4(ad - b^2)}}{a + d - \sqrt{(a+d)^2 - 4(ad - b^2)}} \quad (13)$$

The stability of the Kalman filter update will be maximized when Equation 13 is minimized, and conversely, the stability will be reduced when the ratio is large. It can be seen that the difference between numerator and denominator is controlled by $|(a+d)^2 - 4(ad - b^2)|$:

$$|(a+d)^2 - 4(ad - b^2)| = 4(((\tilde{x}_1^2 + \tilde{y}_1^2) - (\tilde{x}_2^2 + \tilde{y}_2^2))^2 + (\beta^2 + \beta + 1)) \quad (14)$$

where $\beta = \tilde{y}_1 \tilde{y}_2 + \tilde{x}_1 \tilde{x}_2$.

The system achieves maximal stability when the robot is equidistant from both landmarks (cancelling the first two terms), at a position such that the vectors from the robot to the two landmarks are orthogonal (yielding $\beta = 0$). Furthermore, the system becomes more unstable as the robot moves closer to a landmark. That is, the quantity is maximized by letting any r_i go to zero, since $\tilde{x}_i$ and $\tilde{y}_i$ will become unbounded.

These results indicate that poses yielding maximally informative observations (as per the prediction variance, Equation 2) are also maximally destabilizing for the filter. This is a particular concern for the EKF, since these regions also correspond to regions where the curvature of the observation function is also large. In other words, the neighborhood of the operating point of the local linear approximation at these poses is very small, and that observations may be more likely fall outside this neighborhood. When they do fall outside this neighborhood, the linear extrapolation will be incorrect and the filter will diverge.

Figure 2 illustrates the effect of the relative configuration of three landmarks $\{l_0, l_1, l_2\}$ on the conditioning of S.

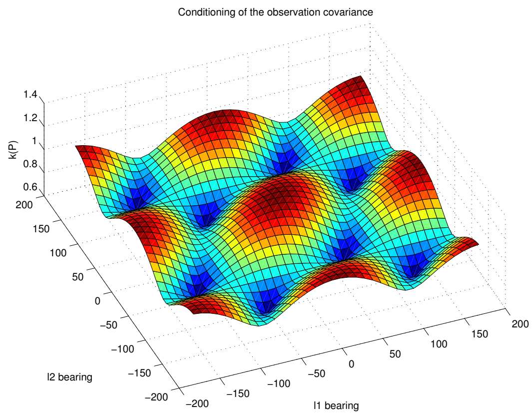

Fig. 2. Conditioning of the update problem based on landmark position. Landmark 0 is assumed to lie at 0 degrees, while the x axis corresponds to varying the position of landmark 1 from -180 to 180 degrees and likewise the y axis varies the position of landmark 2. All three landmarks are assumed to lie on the unit circle, with the robot at the origin.

The state covariance P is assumed to be the identity in this analysis. The $x - y$ plane in this figure corresponds to the bearings to landmarks l_1 and l_2 as they move about the unit circle, while l_0 remains fixed at 0 bearing and one unit distance from the robot. The z axis in the figure corresponds to the log of the condition number of the observation covariance $||S||\,||S^{-1}||$. Peaks in the surface correspond to arrangements where all three landmarks are collinear (an ill-posed localization problem), and valleys generally correspond to arrangements of landmarks that maximize the relative difference in bearings between landmarks.

Figure 3 demonstrates numerical results examining the effect of robot pose on the conditioning of S for three, and four fixed landmarks. The landmarks are located on the unit circle at bearings 0, 120^o, and -120^o in the three-landmark case, and on the intersection of the unit circle with the principal axes in the four-landmark case. Again, the state covariance P is assumed to be the identity matrix. In these plots, the robot pose varies over the $x - y$ plane, with a constant orientation of 0. Peaks in the plot correspond to the pose of the robot falling close to a landmark. The optimal pose (in terms of conditioning the update problem), falls at points equidistant from all the landmarks, with well-defined troughs passing between pairs of landmarks. Note that the the position of the globally optimal pose will move for different landmark configurations. In addition, we assumed that there is no correlation between individual landmarks, or between the robot's pose and individual landmarks. In fact, such correlations play an important role in map inference, and can not be lightly discounted. We will defer this issue for future work.

In the following section, we present experimental results illustrating that exploration that aims for stability, in addition to information gain, yields more accurate maps.

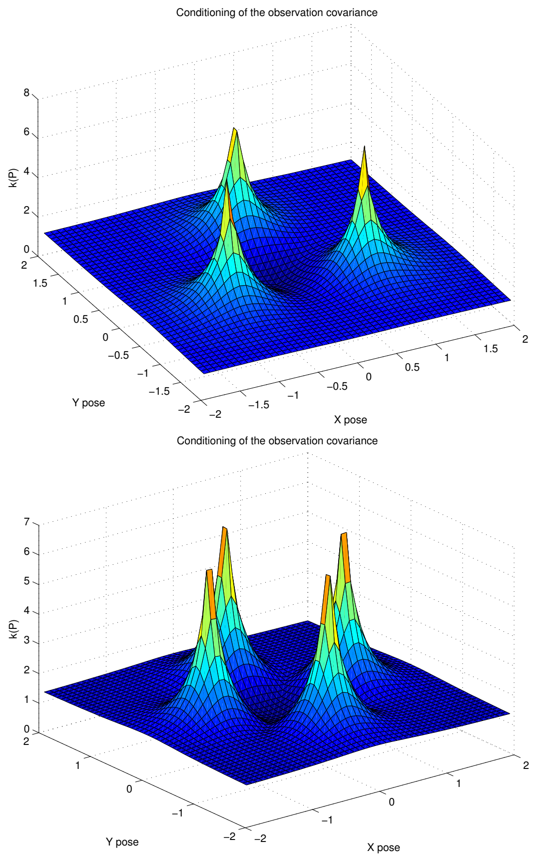

Fig. 3. Conditioning of the update problem based on robot position, for a) three landmarks and b) four landmarks

IV. 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, ϕ) 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 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:

$$\sigma_r = \frac{\texttt{max\_range} - \texttt{min\_range}}{2} \tag{15}$$

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 have designed three exploratory policies to evaluate. They are described below and illustrated in Figure 4.

- `Random Pose`: 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.
- `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.

The choice of the Voronoi graph-based approach is motivated by Figure 3. While the VG doesn't strictly trace the optimal path for filter conditioning, it will steer the robot away from regions that tend towards ill-conditioned updates, and towards regions that are well-conditioned. Furthermore, the VG can be efficiently computed from the landmark estimates and by covering the VG the robot can ensure adequate coverage of the environment. It should be noted that the VG approach does not abandon an information theoretic approach completely, as it ensures that the robot will pass near each landmark, and does so in a way that does not lead to instability. There are other appealing aspects to this approach, since typically landmarks would correspond to obstacles in the world, and the VG maximizes the safety of the robot (see also [11]).

A. Coverage

The latter two exploration policies 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 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 the non-random strategies make strictly local planning decisions in space (such as moving to the nearest unexplored voronoi junction), and in time (by greedily moving to the maximally informative pose without considering future plans). In practice, this approach can lead the robot in large loops, making loop closure difficult. Despite this fact, we have found that in general the robot correctly closes loops, and divergence usually occurs more often where the robot moves to places where the filter

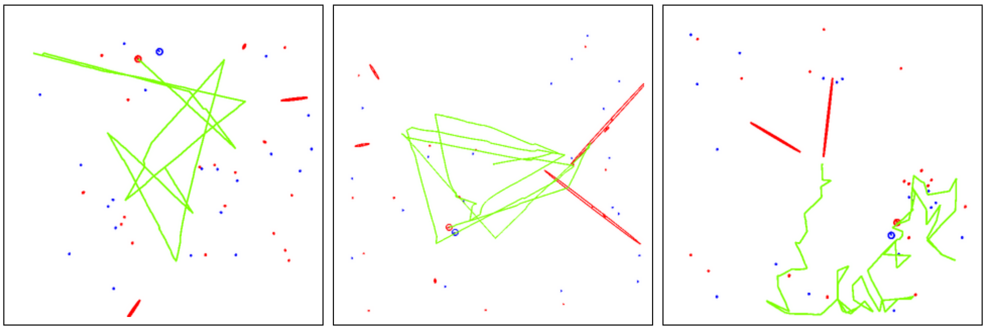

(a) (b) (c)

Fig. 4. Exploration Policies: a) `RandomPose`, b) `InformationGain`, c) `Voronoi`. Refer to the text for details.

update is ill-conditioned or ill-posed. The author notes that the general problem of how to optimally move a robot so as to close loops frequently and accurately remains a rich area of study.

V. Experimental Results

We have run the three 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 distributed a set of dummy landmarks for planning that are removed as real landmarks are discovered, and relocated when the robot explores territory near them. This approach guarantees that under normal conditions all three policies will eventually explore the entire environment.

For individual trials, each run consisted of 2000 time steps in which an action consisted of a rotation to a desired heading 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. 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 4 illustrates a typical map constructed by each 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 5 depicts the results from our experiments. The first plot indicates the mean 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 the locations of unvisited landmarks.

The experimental results support our analysis: tracing the Voronoi graph of the landmarks results in significantly higher accuracy than the other policies. In addition, on average the Voronoi-based approach achieved greater coverage of the environment, further confirming that the approach provides enhanced stability.

VI. Conclusion

We have examined the stability of the bearings-only simultaneous localization and mapping problem. The results presented here suggest that information-driven approaches to SLAM can introduce instabilities that eventually lead to catastrophic divergence. A series of exploratory policies was investigated that support this observation, where the most accurate policy was one that aimed specifically for stability, rather than rapid convergence in the state covariance. It should be noted that in general this policy failed to cover the entire map region in the time alloted, suggesting a trade-off between accuracy and coverage. This trade-off has also been observed in other exploration contexts [12].

While we have examined only an application of the EKF to the bearings-only case, this analysis can be easily extended to range-and-bearings mapping and other representations (e.g. [13], [14]). It should be noted that while several other representations have demonstrated superior efficiency at solving large-scale SLAM problems, the Kalman Filter is generally considered to be more accurate. Given that divergence in the EKF is in large part due to the linearity assumption, other linear, Gaussian models will share a similar susceptibility to divergence. Short of representing the full, non-Gaussian posterior distribution over maps, we anticipate that applying other representations will bear out

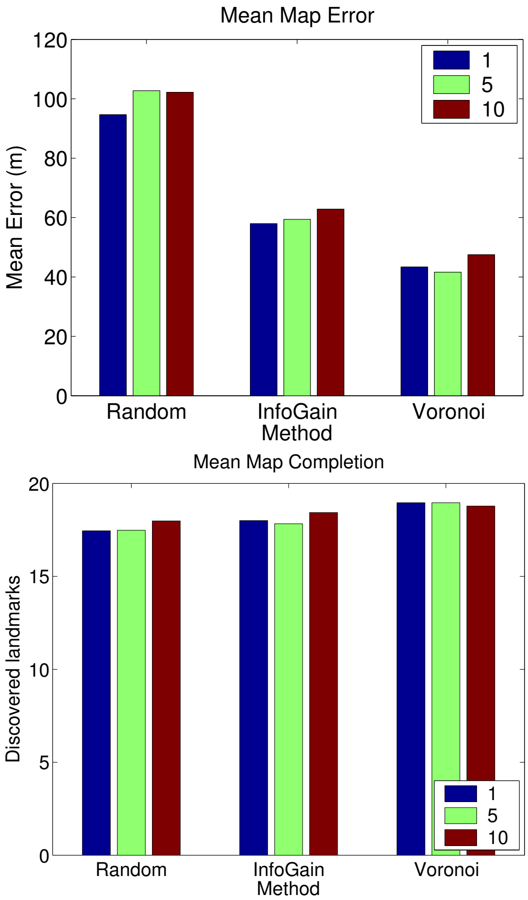

Fig. 5. Mean mapping error, and map coverage, for 100 trials. Bars are reported by observation noise models of 1, 5 and 10 degree standard deviation.

similar results– information driven exploration can come at the cost of accuracy. Formalizing this trade-off (between information gain and stability) is a major topic for future consideration.

It should be noted that the robot's motion model did not figure in our analysis. Instead, we focused on computing optimal observation poses independent of variations (and correlations) in the pose covariance. A full solution will invariably need to take these variations into account. However, we contend that the motion model only indirectly contributes to the stability of the measurement update problem and that the more important quantity is the operating point of the linearization, as determined by the robot pose and landmark estimates.

A second significant assumption is that of perfect data association. While in vision settings data association methods are becoming increasingly reliable, incorrect data associations can occur and cause filter divergence. In the information-driven exploration context, driving a robot to regions where the observation is highly uncertain will surely lead to difficulties in data association. We will defer this question to future work.

Another important direction for future work is to consider a bearings sensor with a limited field of view, such as would be exemplified by a monocular camera. Finally, the applications of these results extend beyond robotic exploration to other state estimation tasks. Our future work will examine related problems and applications in this domain.

REFERENCES

[1] 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.

[2] 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.

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

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] H. Choset and J. Burdick, "Sensor based planning, part 11: Incremental construction of the generalized voroni graph," in *Proc. of IEEE Conf. on Robotics and Automation*. Nagoya, Japan: IEEE Press, May 1995, pp. 1643 – 1648.

[12] R. Sim and G. Dudek, "Effective exploration strategies for the construction of visual maps," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, vol. 3. Las Vegas, NV: IEEE Press, October 2003, pp. 3224–3231.

[13] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," in *Proceedings of the AAAI National Conf. on Artificial Intelligence*. Edmonton, Canada: AAAI, 2002.

[14] M. Paskin, "Thin junction tree filters for simultaneous localization and mapping," in *Proceedings of the Eighteenth Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, G. Gottlob and T. Walsh, Eds. San Francisco, CA: Morgan Kaufmann Publishers, 2003, pp. 1157–1164.

[15] A. Davison, "Mobile robot navigation using active vision," Ph.D. dissertation, Oxford University, 1999.