

Landmark Selection for Terrain Matching

Clark F. Olson

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Mail Stop 125-209, Pasadena, CA 91109
<http://robotics.jpl.nasa.gov/people/olson/homepage.html>

Abstract

We describe techniques to optimally select landmarks in order to perform mobile robot localization by matching terrain maps. The method is based upon a maximum-likelihood robot localization algorithm that efficiently searches the space of possible robot positions. We use a sensor error model to estimate the probability distribution of the terrain expected to be seen from the current robot position. The estimated distribution is compared to a previously generated map of the terrain and the optimal landmark is selected by minimizing the predicted uncertainty in the localization. This approach can be used to generate a sensor uncertainty field for use by the robot's planning component. Experiments indicate that landmark selection improves not only the localization uncertainty, but also the likelihood of success.

1 Introduction

In the localization process, a robot must decide what landmarks to use in order to determine where it is. Robots that use sensors with a limited field-of-view (for example, a rover with stereo cameras) must decide how to position the sensor(s) in order to optimize the ability of the robot to perform localization.

Several recent papers have discussed strategies for sensor placement or landmark selection for use in navigation or localization. A common approach is to consider which landmarks, from a pre-determined set of landmarks, will yield the best localization result. Sutherland and Thompson [11] developed one of the earliest methods for landmark selection. They applied heuristic functions to select a landmark triple, from the set of such triples, that is likely to yield a good localization result. Greiner and Isukapalli [2] learn a function to select landmarks that minimize the expected localization error. A related technique is given by Thrun [13], who trains a neural network to learn landmarks that optimize the localization uncertainty.

Yeh and Kriegman [14] select the subset of features from a set of possible features that minimizes

a Bayesian cost of localization. Deng *et al.* [1] select a set of landmarks in order to minimize the cost of sensing over a path segment. Murphy *et al.* [6] first select candidate landmark triples using heuristics. The highest ranked candidate (and others, if necessary) are tested using experimentation with a robot. Sim and Dudek [9] consider image locations with high edge density as possible landmarks, which are represented using an appearance-based method. Landmarks are detected by matching in the image subspace and the resulting estimates are combined in a robust manner. Little *et al.* [4] find stable landmarks by first detecting image corners. The corners that lie on depth discontinuities are eliminated using stereo vision.

Each of these papers considers a problem where landmarks are selected from a pre-determined set of possible landmarks. Research that does not assume a pre-determined set of landmarks includes work by Simhon and Dudek [10]. They choose regions in which good metric maps can be established according to a distinctiveness measure. Grudic and Lawrence [3] learn a mapping between an image and the robot location, but they do not address the problem of where to best place the camera to obtain the image.

Little of the research to date can be successfully applied to localization in unstructured outdoor terrain, which is the problem that we address. We describe a technique that selects the best landmark to view for localization knowing only an elevation map of the terrain and an estimate of the robot's position. We assume that the robot is equipped with a limited field-of-view (FOV) range sensor, such as sonar or stereo cameras. Our method selects the position to aim the range sensor in order to optimally perform localization in the unstructured three-dimensional terrain.

The landmark selection technique that we use is based upon performing uncertainty estimation using a maximum-likelihood localization method [7, 8]. Prior to performing localization, the robot analyzes the terrain in the global map to select a *localization target*, which is the position in the terrain that the robot senses in order to generate a local map to compare

against the global map. We desire a location that has distinctive terrain and thus allows the localization to be performed with a low uncertainty. This assumes that the error in the robot position is not so large that the localization target will be outside of the view of the robot when it attempts to sense this location. Active vision techniques can be used if, after the robot attempts to sense the localization target, no distinctive terrain is seen.

The first step in determining the localization target is estimating the error present in the global map and the error expected from sensing the terrain at the robot's current position. These errors are encoded in a probability map of the terrain expected to be seen by the robot. Each cell in this map contains an estimate of the probability that the cell will be seen as occupied by the robot if the robot performs sensing with the cell in the field-of-view. By treating this probability map as a terrain map, we can apply previously developed uncertainty estimation techniques [7] to predict the uncertainty that will occur for any target in the probability map. The location with the lowest predicted uncertainty is selected as the localization target.

In addition to improving localization, these techniques can be applied to determining a sensory uncertainty field for the robot. The sensory uncertainty field is a concept introduced by Takeda *et al.* [12] that measures the expected distribution of errors in the robot position as the robot moves through some environment, performing sensing at some interval in order to improve localization. Given the uncertainty estimation and target selection methods, we can determine the expected localization uncertainty for any robot position in the environment.

These techniques have been applied to localization for Rocky 7, which is a research prototype for Mars exploration and science operations. In addition to body-mounted stereo cameras on the front and back, Rocky 7 has a stereo pair of cameras on a retractable mast that allows it survey the terrain. See Fig. 1. We thus concentrate on localization using stereo range data. However, most of this discussion applies equally well to other range sensors. Experiments on Rocky 7 and with synthetic data indicate that the landmark selection not only decreases the robot's localization uncertainty, but also increases the probability of achieving a qualitatively correct localization result.

2 Terrain matching

The basic localization technique that we use is to compare a map generated at the current robot position



Figure 1: The Rocky 7 Mars rover prototype in the JPL Mars yard with its mast deployed.

(the *local map*) to a previously generated map of the environment (the *global map*) [8]. This technique is reviewed here.

We generate both the local map and the global map (which may be the combined result of previous local maps) using stereo vision on-board the robot. The range image is converted into a digital elevation map under the assumption that we know the robot orientation through other sensors, although this restriction can be removed, if desired. To further simplify the problem, we use a high-pass filter on the heights so that the search for the robot position needs to be performed only in the x and y directions. The generated representation is then converted into a three-dimensional occupancy grid.

2.1 Map similarity measure

We formulate the map matching problem in terms of maximum-likelihood estimation. A convenient set of measurements that can be used for this problem are the distances from the occupied cells in the local map to their closest occupied cells in the global map. Denote these distances D_1^X, \dots, D_n^X for the robot position X . The likelihood function for the robot position can be formulated as the product of the probability distributions of these distances [8]. For convenience, we work with logarithms:

$$\ln L(X) = \sum_{i=1}^n \ln p(D_i^X) \quad (1)$$

For the uncertainty estimation to be accurate, it is important that we use a probability distribution func-

tion (PDF) that closely models the sensor uncertainty. This can be accomplished using a PDF that is the weighted sum of two terms [7]:

$$p(d) = \alpha p_1(d) + (1 - \alpha)p_2(d) \quad (2)$$

The first term describes the error distribution when the cell is an inlier (in the sense that the terrain position under consideration in the local map also exists in the global map). In this case, d is a combination of the errors in the local and global maps at this position. In the absence of additional information with respect to the sensor error, we approximate $p_1(d)$ as a normal distribution:

$$p_1(d) = \frac{1}{\sigma\sqrt{2\pi}} e^{-d^2/2\sigma^2} \quad (3)$$

The second term describes the error distribution when the cell is an outlier. In this case the position represented by the cell in the local map does not appear in the global map (e.g. due to range shadows or stereo outliers). In practice, we have found that modeling this term as a constant is both convenient and effective [7].

$$p_2(d) = K \quad (4)$$

Although, $p_2(d)$ is not a probability distribution (it does not integrate to one), using the expected probability density for a measurement generated by a random outlier point yields excellent results:

$$K = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} p(d)^2 dx dy \quad (5)$$

This value can be estimated quickly through examination of the Euclidean distance transform of the map.

In Equation (2), α is the probability that any particular cell in the local map is an inlier. For our occupancy grids, we assume that this value is relatively large ($\alpha = 0.95$). In practice, the localization is insensitive to the value of this variable. Finally, σ is the standard deviation of the measurements that are inliers. This value can be determined from the characteristics of the sensor, or it can be estimated empirically by examining real data, which is the method that we have used for localization on Rocky 7.

2.2 Uncertainty estimation

We determine the uncertainty in the localization estimate by fitting a parameterized surface to the likelihood function in the neighborhood of the highest peak [7]. Since the likelihood function measures the probability that each position in the pose space is the actual

robot position, the uncertainty in the localization is measured by the rate at which the likelihood function falls off from the peak.

We assume that the likelihood function can be approximated as a normal distribution in the neighborhood around the peak location. Fitting such a normal distribution to the computed likelihoods yields both an estimated variance in the localization estimate and a subpixel estimate of the peak location. While the approximation of the likelihood function as a normal distribution may not always be ideal, it yields a good fit to the local neighborhood around the peak and our experimental results indicate that very accurate results can be achieved under this assumption.

We fit the peak in the likelihood function with:

$$k e^{-\frac{1}{2(1-\rho^2)} \left[\left(\frac{x-\mu_x}{\sigma_x} \right)^2 - 2\rho \left(\frac{x-\mu_x}{\sigma_x} \right) \left(\frac{y-\mu_y}{\sigma_y} \right) + \left(\frac{y-\mu_y}{\sigma_y} \right)^2 \right]}, \quad (6)$$

where $k = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}}$, μ_x and μ_y represent the subpixel position estimate, σ_x and σ_y are the standard deviations along the axes, and ρ describes the orientation of the axes with respect to the global coordinate frame. The function is fit using the peak value and the eight neighboring values using a least-squares criterion in the log-likelihood domain.

In addition to estimating the uncertainty in the localization estimate, we can use the likelihood scores to estimate the probability of a failure to detect the correct position of the robot [7]. This is particularly useful when the terrain yields few landmarks or other references for localization and thus many positions appear similar to the robot.

3 Probability mapping

In order to predict the uncertainty achievable by sensing some location (or combination of locations), we make a probabilistic prediction of the appearance of the terrain to the sensor. Each cell in this probabilistic map stores a probability estimate that the cell will be seen as occupied in the sensed map. We call this the *probability mapping* of the terrain. This mapping should encompass the errors present both in the generation of the global map and the expected errors in the new local map.

For the case of stereo vision, Matthies [5] has found that the errors are well approximated by a two-dimensional normal distribution with the major axis aligned away from the cameras. We thus convolve the global map with two normal distributions, one representing the error in the global map and one represent-

$$N(x, y; i, j) = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1-\rho^2}} e^{-\frac{1}{2(1-\rho^2)} \left[\left(\frac{x}{\sigma_1}\right)^2 - 2\rho\left(\frac{x}{\sigma_1}\right)\left(\frac{y}{\sigma_2}\right) + \left(\frac{y}{\sigma_2}\right)^2 \right]} \quad (7)$$

$$\sigma_1 = E_{\sigma_1}[x + i, y + j] \quad \sigma_2 = E_{\sigma_2}[x + i, y + j] \quad \rho = \left(\frac{\sigma_1}{\sigma_2}\right) \tan \theta[x + i, y + j]$$

ing the error in the local map. Note, however, that the error in the map is a function of the location being sensed. The expected error grows with the square of the distance to the camera. We thus allow the width of the normal distributions to vary with the position in the environment.

Our position-variant spreading function is given by Equation (7) above. In this equation, $E_{\sigma_1}[x + i, y + j]$ and $E_{\sigma_2}[x + i, y + j]$ are the expected standard deviations at the location $(x + i, y + j)$, and $\theta[x + i, y + j]$ is the orientation of the distribution (i.e. the direction of the sensor position with respect to $(x + i, y + j)$ when the map is created).

To estimate the error in the global map, we use:

$$P_G(i, j) = \sum_{x=-W}^W \sum_{y=-W}^W M(x+i, y+j)N(x, y; i, j), \quad (8)$$

where $M(x, y)$ is the global map, $2W + 1$ is the size of the convolution window, and $N(x, y; i, j)$ is the distribution described above. Incorporating the expected error in the local map, we get:

$$P(i, j) = \sum_{x=-W}^W \sum_{y=-W}^W P_G(x+i, y+j)N(x, y; i, j). \quad (9)$$

Of course, the instances of $N(x, y; i, j)$ in (8) and (9) will be somewhat different since the expected standard deviations and orientations will be different for the points in the global map versus the local map.

4 Landmark selection

Given the probability mapping of the terrain, we can now estimate the uncertainty that will result from pointing the range sensor at some location in the environment and performing localization using the visible terrain. This is performed by treating the corresponding terrain patch in the probability mapping as the local map and using the uncertainty estimation equations described above.

In our implementation, we approximate the general normal distributions used to model the sensor error as rotationally symmetric 2-D normal distributions. While the error due to stereo vision is much greater

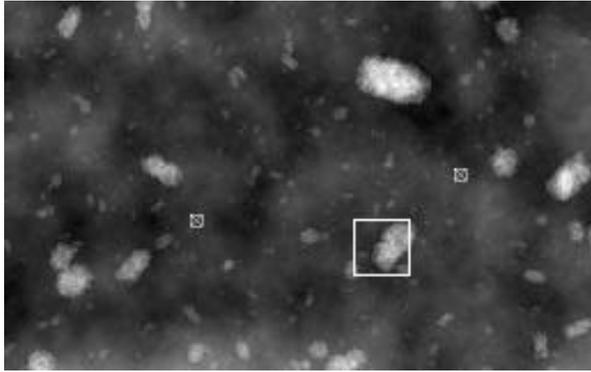
along the direction parallel to the camera axes, error in the robot's knowledge of its orientation will yield additional errors in the perpendicular direction. Furthermore, our experiments indicate that the precise shape of the distribution does not have a large effect on the landmark selected. The use of rotationally symmetric normal distributions makes the function separable and we can thus perform the convolutions efficiently by treating the x and y directions sequentially. On the other hand, it is crucial to use a wider and flatter distribution at locations further from the sensor, in order to model the increase in error with distance. We must thus continue to vary the distribution as a function of the location in the space.

We can make the computation even more efficient by discretizing the space of allowable standard deviations and pre-computing the normal distributions corresponding to them. In our implementation, we select ten standard deviations (related by powers of $\sqrt{2}$). For each position in the space, we select the distribution with the closest standard deviation to the desired value. This approximation allows the probability map to be computed quickly upon demand for a region of the terrain map.

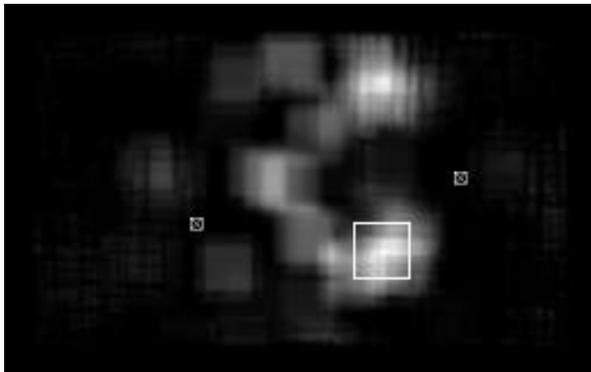
Finally, we use dynamic programming to compute the likelihood function from Section 2.1 for each of the terrain regions considered as a possible landmark for performing localization. This is performed at the optimal localization position and the neighboring locations in the pose space in order to apply the uncertainty estimation techniques from Section 2.2. The terrain landmark yielding the lowest localization uncertainty is selected as the localization target.

5 Results

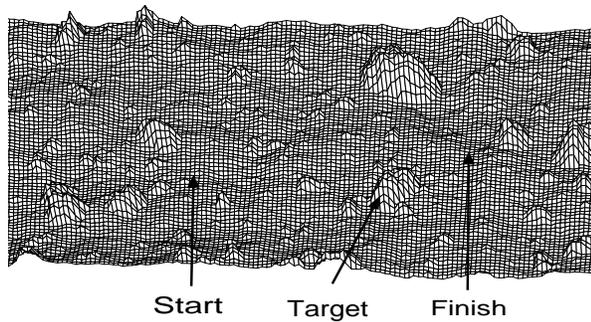
We have tested these techniques in several experiments using real and synthetic data. An example using a synthetically generated elevation map is shown in Fig. 2. This case models a scenario where the robot is moving in a terrain consisting of rocks of various sizes, like the terrain a rover would encounter on the surface of Mars. The positions near large rocks are considered to be good targets, as shown by the uncertainty scores represented by Fig. 2(b). The target



(a)



(b)



(c)

Figure 2: Landmark selection example. The boxed \times 's are the (interchangeable) robot beginning and ending positions. The selected target region is marked by a larger box. (a) Digital elevation map. (b) Estimated uncertainties for landmarks centered at each location. (c) Three-dimensional map.

that is chosen is a position that contains not only a large rock, but also smaller rocks that are also useful in performing the localization.

In order to test the localization performance when using the target selection techniques, we simulated localization problems by sampling local maps from the distribution specified by the probability map of the terrain and then performing localization against the global map. Our experiment selected robot positions at random from the terrain. We next performed target selection and, finally, localization using the selected target. In addition, we tested localization using the target at the position directly between the robot starting and ending positions, and eight other targets on an evenly space grid around this position.

The results of this experiment are dramatic. When target selection was used in 1000 trials using the terrain shown in Fig. 2, localization found the qualitatively correct position in 97.8% of the trials. However, when target selection was not used, the localization succeeded in only 29.5% of the trials, since much of the terrain provides little useful information for localization. In addition, the successful cases were 15.3% more accurate when target selection is used. This experiment thus demonstrates a case where target selection is not only useful in reducing the localization uncertainty, but also critical in obtaining the correct qualitative position.

6 Sensor uncertainty field

Our approach can be used to generate a sensor uncertainty field [12] for a known terrain map. This field is the expected distribution of error in the sensed robot position as a function of the robot location. While, in general, the uncertainty will depend on the path taken to each position, we consider the uncertainty as a function of only the robot position.

We can, of course, compute the sensor uncertainty field using a brute-force method, where the best landmark is selected for each location of the robot and the resulting expected uncertainty is stored for each. Unfortunately, this process requires much computation. Note, however, that the uncertainties change slowly as the robot position that is examined is moved in the pose space. Our strategy is to first sample the pose space at a coarse resolution and then examine locations of interest, such as those that yield low uncertainties, subsequently at a finer resolution.

Figure 3 shows an example where a sensor uncertainty field was generated for the terrain shown in

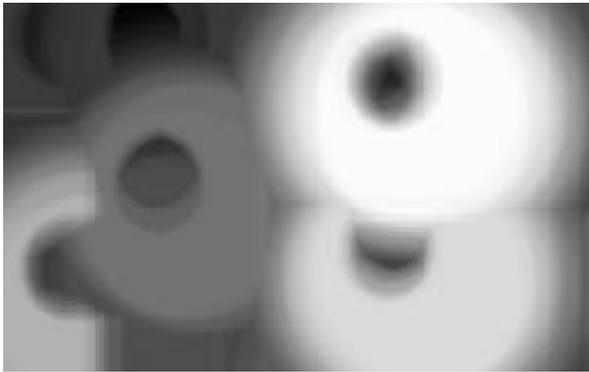


Figure 3: Sensor uncertainty field generated for the terrain in Fig. 2. Bright values correspond to low uncertainties.

Fig. 2. As expected, lower uncertainties occur near large rocks. However, the uncertainty is increased at the location of the rocks, since we use a method where the rock is not useful for localization if the robot is directly on top it.

7 Summary

We have described a method to select the sensing location for performing mobile robot localization through matching terrain maps. The localization method that we use constructs a likelihood function in the space of possible robot positions. The uncertainty is estimated for localization using a local map by fitting a normal distribution to the likelihood function generated. We select the best landmark for localization by minimizing the expected uncertainty in the robot localization. In order to predict the uncertainty obtained by localization using various landmarks, our method constructs a probabilistic representation of the terrain expected to be sensed at any position in the global map. Treating the patches of this “probability mapping” of the terrain as a local map allows the uncertainty expected by sensing the terrain patch to be estimated using the surface fitting techniques. We have applied this technique to robot localization in rocky terrain with excellent results.

Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work is

an element of the Long Range Science Rover project, which is developing technology for future Mars missions. It is funded as part of the NASA Space Tele-robotics Program, through the JPL Robotics and Mars Exploration Technology Office.

References

- [1] X. Deng, E. Miliotis, and A. Mirzaian. Landmark selection strategies for path execution. *Robotics and Autonomous Systems*, 17(3):171–185, May 1996.
- [2] R. Greiner and R. Isukapalli. Learning to select useful landmarks. *IEEE Transactions on Systems, Man, and Cybernetics - Part B : Cybernetics*, 26(3):437–449, June 1996.
- [3] G. Z. Grudic and P. D. Lawrence. A nonparametric learning approach to vision based mobile robot localization. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 724–729, 1998.
- [4] J. J. Little, J. Lu, and D. R. Murray. Selecting stable image features for robot localization using stereo. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1072–1077, 1998.
- [5] L. Matthies and S. A. Shafer. Error modeling in stereo navigation. *IEEE Transactions on Robotics and Automation*, 3(3):239–248, June 1987.
- [6] R. R. Murphy, D. Hershberger, and G. R. Blauvelt. Learning landmark triples by experimentation. *Robotics and Autonomous Systems*, 22(3-4):377–392, 1997.
- [7] C. F. Olson. Subpixel localization and uncertainty estimation using occupancy grids. In *Proceedings of the International Conference on Robotics and Automation*, volume 3, pages 1987–1992, 1999.
- [8] C. F. Olson and L. H. Matthies. Maximum-likelihood rover localization by matching range maps. In *Proceedings of the International Conference on Robotics and Automation*, volume 1, pages 272–277, 1998.
- [9] R. Sim and G. Dudek. Mobile robot localization from learned landmarks. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1060–1065, 1998.
- [10] S. Simhon and G. Dudek. Selecting targets for local reference frames. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 2840–2845, 1998.
- [11] K. T. Sutherland and W. B. Thompson. Localizing in unstructured environments: Dealing with the errors. *IEEE Transactions on Robotics and Automation*, 10(6):740–754, December 1994.
- [12] H. Takeda, C. Facchinetti, and J.-C. Latombe. Planning the motions of a mobile robot in a sensory uncertainty field. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(10):1002–1017, October 1994.
- [13] S. Thrun. Bayesian landmark learning for mobile robot localization. *Machine Learning*, 33(1):41–76, October 1998.
- [14] E. Yeh and D. J. Kriegman. Toward selecting and recognizing natural landmarks. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, pages 47–53, 1995.