

Subpixel Localization and Uncertainty Estimation Using Occupancy Grids

Clark F. Olson

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Mail Stop 125-209, Pasadena, CA 91109-8099
<http://robotics.jpl.nasa.gov/people/olson/homepage.html>

Abstract

We describe techniques for performing mobile robot localization using occupancy grids that allow subpixel localization and uncertainty estimation in the pixelized pose space. The techniques are based on a localization method where matching is performed between the visible landmarks at the current robot position and a previously generated map of the environment. A likelihood function over the space of possible robot positions is formulated as a function of the probability distribution for the map matching error. Subpixel localization and uncertainty estimation are performed by fitting the likelihood function with a parameterized surface. The performance of the method is analyzed using synthetic experiments and an example is given using the Rocky 7 Mars rover prototype.

1 Introduction

Localization is a critical issue in mobile robotics. If the robot does not know where it is, it cannot effectively plan movements, locate objects, or reach goals. It is important to not only perform accurate localization when possible, but also to know when the localization estimate has a large uncertainty and when it is possible that a qualitative failure in localization has occurred. For example, in environments where the recognizable landmarks are sparse, many locations may appear to be very similar to the robot. It is crucial to know when this is the case, so that additional data can be collected to improve the robot localization.

This paper describes techniques to perform accurate subpixel localization (in the pixelized pose space) when sufficient information is available and to generate uncertainty estimates in the localization result, both in the standard deviation of the localization and in the probability of qualitative failure, regardless of the quality of the data. The techniques that we describe are general in nature and can be applied to most map representations. However, we concentrate on the

application of these techniques to three-dimensional occupancy grids, building upon previously reported results that perform localization by matching terrain maps using a maximum-likelihood comparison measure [7].

We first review the basic localization method that is used. This technique computes a map similarity measure using the probability distribution function (PDF) of the distance from each occupied cell in the local terrain map that is computed at the current robot position to the closest occupied cell in a previously computed map of the environment. An accurate approximation for this probability distribution function is given using the weighted sum of a normal distribution (for cases where the cell is an inlier) and a constant distribution (for cases where the cell is an outlier). This PDF forms the core of the likelihood function for each robot pose.

Subpixel localization and uncertainty estimation are performed by fitting the tallest peak in the likelihood function with a parameterized surface. We approximate the peak in the likelihood function as a normal distribution. Operating in the log-likelihood domain allows us to fit the peak with a second-order polynomial. The location of the summit of this surface yields the subpixel localization estimate discretized pose space. Furthermore, the rate at which the surface falls off from the peak provides an estimate on the uncertainty in the localization. The probability of a qualitative failure is estimated by comparing the likelihood scores under the tallest peak and the likelihood scores in the rest of the pose space.

These techniques have been applied to Rocky 7[2], which is a six-wheeled mobile robot (of approximately the same size and shape as the Sojourner rover) that has been built at the Jet Propulsion Laboratory in order to develop new technologies for future missions to Mars (Fig. 1). The techniques were tested using Rocky 7 in the JPL Mars Yard to simulate mission-like conditions.



Figure 1: Rocky 7 generating a range map in the JPL Mars Yard using its mast cameras.

2 Maximum-likelihood localization

The basic localization technique that we use is to compare a map generated at the current robot position (the *local map*) to a previously generated map of the environment (the *global map*) [7]. This technique is reviewed here. See [1, 4] for related techniques.

2.1 Terrain maps

We generate both the local map and the global map using stereo vision on-board the robot. (The global map may consist of the combined result of the previous local maps or it may be generated using panoramic imagery from the robot's starting location.) A dense range image is first generated using the stereo cameras [5]. The range image is then converted into an occupancy grid representation at some canonical orientation using a binning operation. (It is assumed that the robot orientation is known through other sensors.) The average height of the range points that fall into each bin is taken to be the height of the grid at that location. Finally, 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.

Figure 2 shows an example of a terrain map that was generated using data from the Mars Pathfinder mission.

2.2 Map similarity measure

In order to formulate the matching problem in terms of maximum-likelihood estimation, we use a set of measurements that are a function of the robot position. A convenient set of measurements 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. For convenience, we work in the $\ln L(X)$ domain:

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

The map similarity measure that is used is defined entirely by the probability distribution function (PDF) of the distances, $p(D_i^X)$. This probability distribution function is discussed in detail in Section 3.

2.3 Search strategy

A multi-resolution search strategy is used to determine the most likely robot position [3, 6, 7]. This method is guaranteed to locate the optimal position in the discretized search space. The pose space is first discretized at the same resolution as the occupancy grids so that neighboring positions in the pose space move the relative positions of the grids by one grid cell. We then test the nominal position of the robot given by dead-reckoning so that we have an initial position and likelihood to compare against. Next, the pose space is divided into rectilinear cells. Each cell is tested to determine whether it could contain a position that is better than the best position found so far. Cells that cannot be pruned are divided into smaller cells, which are examined recursively. When a cell is reached that contains a single position in the discretized pose space, then this position is tested explicitly.

To determine whether a cell C could contain a pose superior to the best found so far, we examine the pose c at the center of the cell. A bound is computed on the maximum distance between the location to which a cell in the local map is transformed by c and by any other pose in the cell. We call this distance Δ_C . For the space of translations, Δ_C is simply the distance between c and any corner of the cell. To place a bound on the quality of any position within the cell, we bound each of the distances that can be achieved by features in the local map over the cell. This is done by subtracting the maximum change of the cell, Δ_C , from the distance achieved at the center of the cell, D_i^c :

$$D_i^C = \max(D_i^c - \Delta_C, 0)$$

The values obtained are then propagated through the likelihood function to bound the score that can be achieved by any position in the cell.

$$P_i^C = \ln p(D_i^C)$$

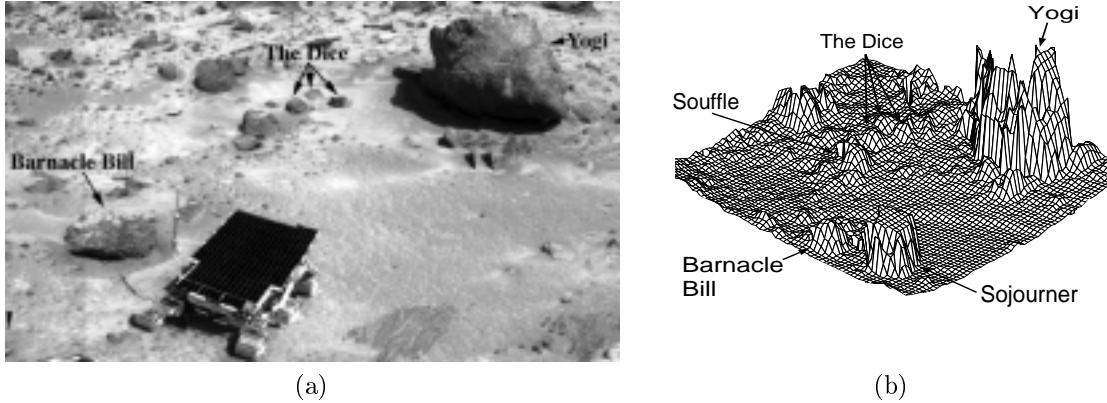


Figure 2: Terrain map generated from Pathfinder imagery. (a) Image mosaic of Sojourner and rocks on Mars. (b) Terrain map generated from stereo imagery.

P_i^C is now the maximum score that the i th feature of the local map can contribute to the likelihood for any position in the cell¹.

A bound on the best overall likelihood that can be found at a position in the cell is given by:

$$\max_{X \in C} \ln L(X) \leq \sum_{i=1}^n P_i^C$$

If this bound does not surpass the best that we have found so far, then the entire cell is pruned from the search. Otherwise, the cell is divided into two cells by slicing it along the longest axis and the process is repeated recursively on the subcells.

3 Estimating the PDF

For the uncertainty estimation to be accurate, it is important that we use a probability distribution function (PDF) that closely models the sensor uncertainty. This can be accomplished using a PDF that is the weighted sum of two terms:

$$p(D_i^X) = \alpha p_1(D_i^X) + (1 - \alpha)p_2(D_i^X)$$

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_i^X is a combination of the errors in the local and global maps at this position. In the absence of additional information with

¹This assumes that the PDF is monotonically non-increasing, which is true for any reasonable PDF, since we desire closer matches to yield higher scores.

respect to the sensor error, we approximate $p_1(D_i^X)$ as a normal distribution:

$$p_1(D_i^X) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(D_i^X)^2/2\sigma^2}$$

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. This may be due to range shadows that were present when the global map was constructed or outliers that are present in the range data when the local map is constructed. In theory, this term should also decrease as D_i^X increases, since even true outliers are likely to be near some occupied cell in the global map. However, this allows pathological cases to have an undue effect on the likelihood for a particular robot position. In practice, we have found that modeling this term as a constant is both convenient and effective:

$$p_2(D_i^X) = K$$

Let us now consider the constants in this probability distribution function. First, α is the probability that any particular cell in the local map is an inlier. For our occupancy grids, we shall assume that this value is relatively large ($\alpha = 0.95$). In practice, the localization is insensitive to the precise value of this variable. Next, 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. Finally, K is the expected probability density for the measurement generated for a random outlier point:

$$K = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} p(D_i^X)^2 dx dy$$

This value can be estimated quickly through examination of the Euclidean distance transform of the map [8].

4 Subpixel localization

Using this probabilistic formulation of the localization problem, we can estimate the uncertainty in the localization in terms of both the variance of the estimated positions and the probability that a qualitative failure has occurred. 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. In addition, we can perform subpixel localization in the discretized pose space by fitting a surface to the peak that occurs at the most likely robot position.

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.

Now, since we actually perform our computations in the domain of the natural logarithm of the likelihood function, we must fit these values with a polynomial of order 2. If we assume independence in x and y , we have:

$$\begin{aligned} \ln L(x, y) &= \ln \frac{1}{2\pi\sigma_x\sigma_y} e^{-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2}} \\ &= -\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2} + \ln \frac{1}{2\pi\sigma_x\sigma_y} \end{aligned}$$

In order to estimate the parameters in which we are interested (x_0 , y_0 , σ_x , and σ_y), we project this polynomial onto the lines $x = x_0$ and $y = y_0$, yielding:

$$\ln L(x, y_0) = -\frac{(x-x_0)^2}{2\sigma_x^2} + \ln \frac{1}{2\pi\sigma_x\sigma_y}$$

$$\ln L(x_0, y) = -\frac{(y-y_0)^2}{2\sigma_y^2} + \ln \frac{1}{2\pi\sigma_x\sigma_y}$$

We now fit these equations to the x and y cross-sections of the likelihood function at the location of the peak. If the peak in the discretized search space occurs at position (x_p, y_p) , then we fit $L(x, y_0)$ to the values at the surrounding 5 positions along $y = y_p$, which we label $\{l_{-2}, l_{-1}, l_0, l_1, l_2\}$.

The least-squares fit to a parabola ($y = ax^2 + bx + c$) with $x = \{-2, -1, 0, 1, 2\}$ yields:

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} \frac{1}{7} & -\frac{1}{14} & -\frac{1}{7} & -\frac{1}{14} & \frac{1}{7} \\ -\frac{1}{5} & -\frac{1}{10} & 0 & \frac{1}{10} & \frac{1}{5} \\ -\frac{3}{35} & \frac{12}{35} & \frac{17}{35} & \frac{12}{35} & -\frac{3}{35} \end{bmatrix} \begin{bmatrix} l_{-2} \\ l_{-1} \\ l_0 \\ l_1 \\ l_2 \end{bmatrix}$$

We can now solve for x_0 and σ_x using:

$$x_0 = x_p - \frac{b}{a} \quad \sigma_x = \frac{1}{\sqrt{-2a}}$$

The derivation for y_0 and σ_y is the same, except that we project onto the line $x = x_p$. The values of x_0 and y_0 yield the subpixel localization result, since this is the estimated location of the peak in the likelihood function. In addition, σ_x and σ_y now yield direct estimates for the uncertainty in the localization result.

5 Probability of failure

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. This is particularly useful when the terrain yields few landmarks or other references for localization and thus many positions appear similar to the robot.

We estimate the probability of failure by summing the likelihood scores under the peak selected as the most likely robot position and comparing to the sum of the likelihood scores that are not part of this peak. In practice, we can usually estimate the sum under the peak by examining a small number of values around the peak, since they fall off very quickly.

The values for the remainder of the pose space can be estimated efficiently with some additional computation during the search. Whenever a cell in the search space is considered, we compute not only a bound on the maximum score that can be achieved, but also an estimate on the average score that is achieved by determining the score for the center of the cell. If the cell is

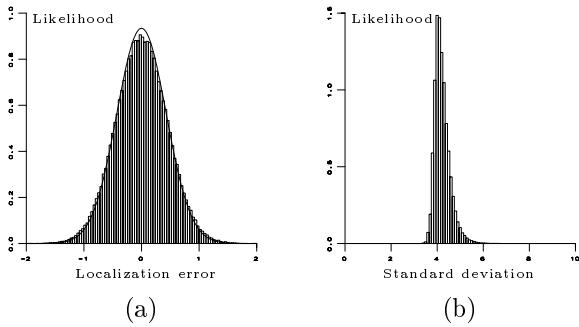


Figure 3: Distribution of errors and estimated standard deviations in synthetic landmark localization experiment. (a) Comparison of estimated distribution of localization errors (solid line) to observed distribution of localization errors (bar graph). (b) Distribution of estimated standard deviations in the localization estimate.

pruned, then the sum is incremented by the estimated score multiplied by the size of the cell. In practice, this yields a very good estimate, since regions with large scores cannot be pruned until the cells become small.

Let S_p be the sum obtained for the largest peak in the pose space and S_t be the sum for the entire pose space as described above. We can estimate the probability of correctness for the largest peak as:

$$P_c = \frac{S_p}{S_t}$$

6 Results

These techniques have been tested on synthetic data for which we can compare the performance of the techniques with precise ground-truth and in real experiments on the Rocky 7 rover prototype.

6.1 Synthetic data

We first applied these techniques to localization using landmarks in synthetic experiments. In these experiments, we randomly generated a synthetic environment containing 160 landmarks on a 256×256 square. Let us say that each unit is 10 cm (though the entire problem scales to an arbitrary size). In each trial, seven of the ten landmarks closest to some random robot location were considered to be observed by the robot (with Gaussian error in both x and y with standard deviation $\sigma = 1$ unit) along with 3 spurious landmarks not included in the map. Localization was

then performed using these 10 observed landmarks in a discrete occupancy map, with no knowledge of the position of the robot in this environment. Over 100000 trials, the robot was correctly localized in 99.8% of the cases, with an average error in the correct trials of 0.356 units in each dimension. The average estimated standard deviation in the localization using the techniques from the previous section was 0.427 units.

Figure 3(a) shows the distribution of actual errors observed versus the distribution that we expect from the average standard deviation estimated in the trials. The close similarity of the plots indicates that the average estimated standard deviation is a very good estimate of the actual value. It appears that this estimate is slightly smaller than the true value since the frequency of the observed errors is slightly above the curve at the tails and lower at the peak. However, the overall similarity is quite high. Figure 3(b) shows the distribution of the estimated standard deviations in this experiment. It can be observed that the estimate is very consistent between trials, since the plot is very strongly peaked near the location of the average estimate. Taken together, these plots indicate that the standard deviation estimates are very likely to be accurate for each individual trial.

We have also tested the probability of correctness measure in these trials. For trials that resulted in correct localization, the average computed probability of correctness was .993, while this value was .642 for trials that resulted in failures. The probability of correctness measure thus yields information that can be used to evaluate whether the localization result is reliable.

A comparison of these techniques with a version that does not perform subpixel localization indicates that the subpixel localization reduces the error in the localization by 16.2%.

6.2 Real example

Additional experiments have been performed using the Rocky 7 Mars rover prototype in the JPL Mars Yard. Figure 4 shows one example. In this case, the rover generated a map of the terrain at a starting position using 4 stereo pairs of images covering the area around a rock that was designated as the localization target. After moving, the rover again captured a stereo pair of images directed at the localization target. Localization was performed by matching the range map generated after the move to the terrain map that was generated before the move. In this case, the rover determined that it had moved 4.14 meters

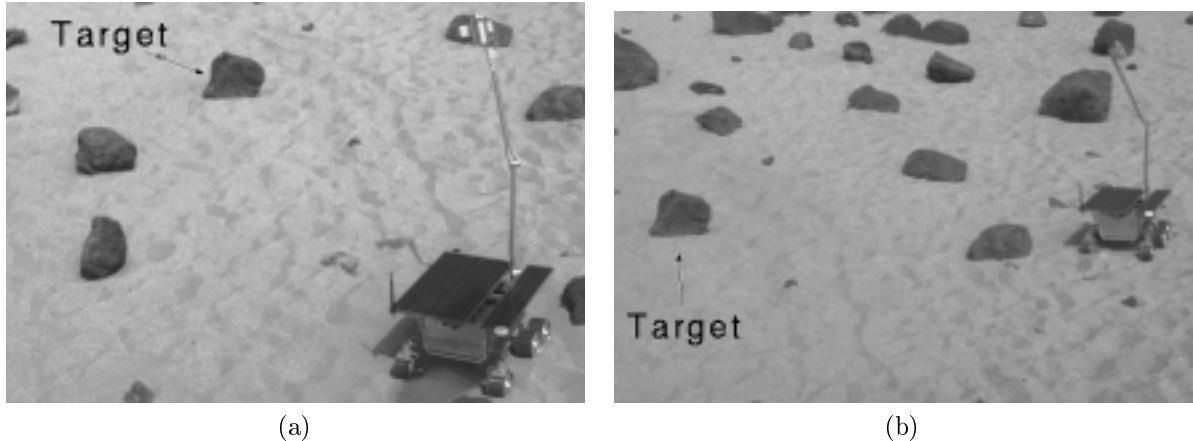


Figure 4: Rocky 7 performing localization. (a) Before moving the terrain is mapped. (b) After moving the same area is imaged to perform localization.

from the original location, which agrees closely with the measured result of 4.10 meters.

7 Summary

This paper has described techniques for performing accurate localization with uncertainty estimation using discrete occupancy grids. The method is based upon a maximum-likelihood method to register occupancy grids representing the robot's visible environment and a previously generated map. A maximum-likelihood map registration measure is first formulated using the probability distribution function of the distances from the cells in the local map to the closest cells in the global map. In order to perform sub-pixel localization and uncertainty estimation, the likelihood function is fit with a parameterized surface in the neighborhood of the highest peak. In addition, the probability of a qualitative failure is estimated by examining the scores over the entire pose space. Experiments on synthetic data have demonstrated that this approach yields superior results to cases where these techniques are not used. Furthermore, the uncertainty estimates that are generated can be used to integrate multiple localization steps or other techniques in a Kalman filtering framework. These techniques were applied to localization of the Rocky 7 Mars rover prototype.

Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the Na-

tional Aeronautics and Space Administration. This work is an element of the Long Range Science Rover project, which is developing technology for future Mars missions. This project is funded as part of the NASA Space Telerobotics Program, through the JPL Robotics and Mars Exploration Technology Office.

References

- [1] A. Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, 3(3):249–265, June 1987.
- [2] S. Hayati et al. The Rocky 7 rover: A Mars sciencecraft prototype. In *Proceedings of the IEEE Conference on Robotics and Automation*, volume 3, pages 2458–2464, 1997.
- [3] D. P. Huttenlocher and W. J. Rucklidge. A multi-resolution technique for comparing images using the Hausdorff distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 705–706, 1993.
- [4] I. S. Kweon and T. Kanade. High-resolution terrain map from multiple sensor data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):278–292, February 1992.
- [5] L. Matthies. Stereo vision for planetary rovers: Stochastic modeling to near real-time implementation. *International Journal of Computer Vision*, 8(1):71–91, July 1992.
- [6] C. F. Olson. A probabilistic formulation for Hausdorff matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 150–156, 1998.
- [7] 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.
- [8] S. Pavel and S. G. Akl. Efficient algorithms for the Euclidean distance transform. *Parallel Processing Letters*, 5(2):205–212, 1995.