



ELSEVIER

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Reliable and efficient landmark-based localization for mobile robots

I. Loevsky^a, I. Shimshoni^{b,*}^a Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel^b Department of Management Information Systems, University of Haifa, Haifa 31905, Israel

ARTICLE INFO

Article history:

Received 7 April 2009

Received in revised form

7 January 2010

Accepted 15 January 2010

Available online xxx

Keywords:

Angle measurements

Landmarks

Robot localization

Mobile robots

AGV

ABSTRACT

This paper describes an efficient and robust localization system for indoor mobile robots and AGVs. The system utilizes a sensor that measures bearings to artificial landmarks, and an efficient triangulation method. We present a calibration method for the system components and overcome typical problems for sensors of the mentioned type, which are localization in motion and incorrect identification of landmarks. The resulting localization system was tested on a mobile robot. It consumes less than 4% of a Pentium4 3.2 GHz processing power while providing an accurate and reliable localization result every 0.5 s. The system was successfully incorporated within a real mobile robot system which performs many other computational tasks in parallel.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Robot localization is an essential component for path planning to desired locations and for preventing the robot from reaching undesired locations. While GPS-based systems provide comfortable solutions for outdoor localization, indoors other methods have to be considered. Many industrial AGVs localize using stripes or wires attached to the floor. The drawback of this approach is that the possible paths the robot can traverse are limited, and therefore its motion is not flexible.

The simultaneous localization and map building (SLAM) approach [1,2] opened a new branch of promising research on the localization problem. A special variant of this problem applicable to our case deals with the calibration of a 1D projective camera and the recovery of the landmark positions [3,4]. Under this setting an omnidirectional camera can be used to detect the positions of vertical lines in the scene [5]. The problem however is that even though these methods are able to build a map of the locations of the landmarks, there exists an unknown similarity transformation between the coordinate system of the recovered map and the coordinate system of the environment specified by the user.

One of the methods to perform indoor localization is to put landmarks at known locations and using a sensor mounted on the robot to measure their bearings. The hardware for implementing

this method can vary from a laser sensor with retro-reflective stripes to an omnidirectional camera with landmarks that create certain image patterns. This hardware is manufactured by several companies and used widely in AGV systems. The traditional geometric method for localization from landmark bearings is based on the idea that two bearing measurements constrain the position of the robot to lie on a specific circle. Thus, when three measurements are available the robot location can be computed [6–12].

When more than three landmarks are observed, these methods yield several dependent localization estimates. The naive idea of averaging these estimates is computationally inefficient and statistically sub-optimal.

In [13] an algebraic technique was suggested, where each landmark defines a linear constraint on the position and the orientation of the robot and an efficient least-squares solution is obtained. When more than three measurements are available, an estimate of the mean square error (MSE) is also obtained. It was shown experimentally that this algorithm yields close to optimal localization results at a fraction of the cost of the non-linear maximum likelihood method.

Mobile robot systems which use landmark bearings for localization are widely used [14–17], and therefore it is important to develop accurate and robust localization algorithms for them.

The method presented in [13] deals correctly with unbiased noise yielding close to optimal results efficiently. However, in real world situations other sources of error are present and a localization system must be able to handle such cases. Dealing with these cases is the main focus of this paper.

* Corresponding author. Tel.: +972 4 8288510; fax: +972 4 8288522.

E-mail addresses: igal@tx.technion.ac.il (I. Loevsky), ishimshoni@mis.haifa.ac.il (I. Shimshoni).

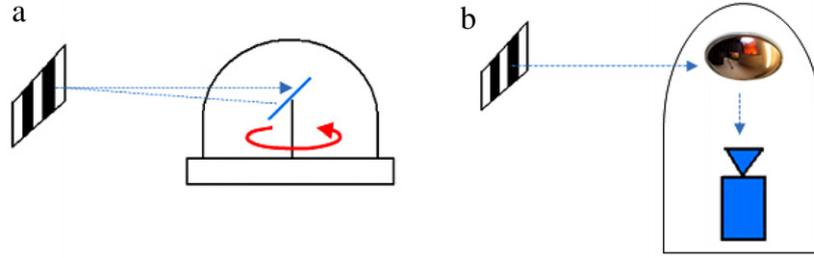


Fig. 1. (a) A laser beam is sent towards an artificial landmark and received, using a rotating mirror. (b) The landmark is detected on the omnidirectional camera image.

There are four reasons for possible failure of the triangulation method presented in [13] that we investigated and solved. The first problem is that the sensor mechanics cause biased measurement errors and therefore has to be calibrated. The second cause of error is that the positions of the landmarks in the environment, which were manually measured and which are given to the algorithm, might be inaccurate. In Section 4 we present an analysis and a combined calibration technique for these two problems. In certain sensors, localization in motion can distort the results as well. In Section 5 we modeled the effects of translation and rotation on the measurements and using it were able to achieve accurate localizations even while the robot was moving. This correction is especially significant when the robot rotates. Finally, there are cases in which even if the landmark bearing is measured accurately its identity is not detected correctly. This causes gross errors in the localization. In Section 6 a RANSAC-based [18] method is proposed for efficient detection of misidentified landmarks.

Our method performs all its operations without requiring ground truth localizations, relying only on MSE calculations. This is in contrast to approaches that deal with such problems by requiring the user to supply ground truth localizations for many robot locations, an extremely tedious and time-consuming process.

All the components of our system were implemented and tested on a Denning MRV4 mobile robot platform with a LaserNav landmark bearings localization laser sensor. In the final experiment described in Section 7 we compared the results of our localization to those achieved by an independent camera-based measurement system and achieved nearly identical results.

2. Localization system hardware

Typical sensors for performing artificial landmark bearings localization are the laser scanner illustrated in Fig. 1a, and an omnidirectional camera sensor illustrated in Fig. 1b.

The laser scanner consists of an emitter, a receiver, and a rotating mirror that reflects a laser beam towards the environment. The landmarks are usually made of retro-reflective material so they can be easily detected by the photodetector inside the scanner. They can be in the form of a single strip [11], or in the form of a barcode. Hence, the returning beam is analyzed and the landmarks are detected and identified.

In the omnidirectional camera setup, the landmarks appearance has to be distinguishable from other objects in the scene. An image processing algorithm detects the observed landmarks.

A set of n artificial landmarks is placed at known locations $P_i \in \mathbb{R}^2$ of which a subset is visible at each robot location. The sensor returns the identity and the bearing measured to each landmark β_i with respect to the sensor orientation, from which the location and the orientation of the sensor are calculated by triangulation. The frame origins of the sensor and the robot are aligned, and therefore the localization is valid for the robot too. See Fig. 2 for illustration of the robot, the sensor and the landmarks.

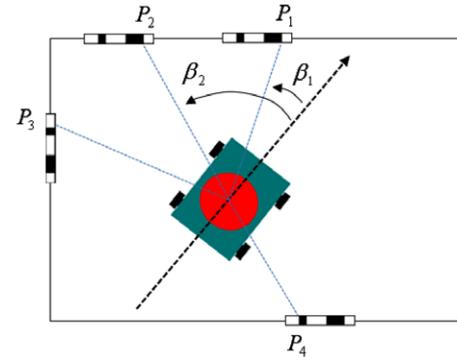


Fig. 2. The robot, the localization sensor and the artificial landmarks. The black dashed line denotes the robot and the sensor orientation.

3. Estimation of localization result accuracy

According to the method described in [13], the relation between the unknown position P and the orientation θ of the robot, and the known position P_i of the landmark is derived from a coordinate transformation from the global coordinates to the robot reference frame. Transformation of a single bearing yields the following set of two equations, each equation is for a row of $R_{2 \times 2}$:

$$M_i = l_i(\cos \beta_i, \sin \beta_i) = R(P_i - P) = RP_i + T, \quad (1)$$

where M_i are the coordinates of the landmark in the robot reference frame, R is a planar rotation matrix by θ , l_i is the unknown distance from P to P_i , and $T = -RP$.

Eliminating l_i leads to a single linear homogeneous equation with four unknowns: $\cos(\theta)$, $\sin(\theta)$, T_x , and T_y . When three measurements are available together with the natural constraint on $\cos(\theta)$ and $\sin(\theta)$, the problem is solved and the position and orientation of the robot are computed. If more than three landmarks are available, the linear least-squares solution is obtained and the error provides a quality measure—however, in this paper we use a more accurate quality measure. It is the mean of squared differences between the observed bearings to the landmarks, and the angles that are calculated from the localization result and the known locations of the landmarks. This measure is statistically correct since it is based on the assumption that measurement errors have a zero mean Gaussian distribution. Let (P_x, P_y) be the position of the robot, i.e the localization result and let (P_{ix}, P_{iy}) be the position of the i th landmark. The estimate of the bearing is therefore

$$\beta_{ci} = \arctan2(P_{iy} - P_y, P_{ix} - P_x). \quad (2)$$

The angular error of a single landmark is

$$\Delta\beta_i = (\beta_{ci} - \theta) - \beta_i. \quad (3)$$

The error value of a localization result is the MSE of all the landmarks of the scan:

$$S = \sum_{i=1}^n \frac{\Delta\beta_i^2}{n}. \quad (4)$$

Under the assumption mentioned above that

$$\Delta\beta_i \sim N(0, \sigma_\beta^2),$$

it follows that

$$\frac{nS}{\sigma_\beta^2} \sim \chi^2(n),$$

and therefore given a certain probability value very close to one, a threshold on S can be derived for each value of n . σ_β^2 is found experimentally for the sensor. Thus, when S is greater than that threshold, the zero mean Gaussian assumption we made about the measurement errors does not hold (e.g., when a landmark is misidentified) and we therefore deduce that the localization result is incorrect. This is the basic measure that we use in all our algorithms to estimate the quality of the localization result. Therefore, at least four landmarks have to be detected in each observation to enable localization with an error estimate. Obviously, as more landmarks are detected in a single observation, the localization accuracy increases.

4. Calibration of the localization system components

Calibration is a necessary step in complex measuring systems in order to assure accurate results. We have to calibrate two components: the hardware which is the sensor and the robot, and the landmarks set. The aim of the hardware calibration is to yield unbiased bearing measurements. This was not the case initially. The reasons for that can be either insufficient leveling of the sensor or distorted readings of the sensor. The landmarks set has to be calibrated in order to overcome the inaccuracies in measurements of the landmark positions that were performed by the user. These inaccuracies can significantly distort the localization results.

This section presents the methods to calibrate each component separately and provides an algorithm to calibrate the components together in a joint procedure.

4.1. Hardware calibration given locations of the landmarks

The uncalibrated hardware problem can be detected when the robot rotates around its axis at a fixed location. The location component of the localization result (P_x, P_y) is supposed to be normally distributed around a certain position with a small variance and dispersion, but instead it is grossly dispersed as can be seen in Fig. 3. This experiment demonstrates a problem which always causes localization errors. We expect that after the hardware calibration step, under these conditions, the localization system should produce a set of locations with a small variance caused by the unbiased measurement errors. A fixed set of landmarks was used, thus isolating the nature of the distortion function. As can be seen in Fig. 3, the distorted localization results are continuous and periodic. This implies that the distortion function that offsets the observed angles is also continuous and periodic.

Thus, the observed value of the bearing can be modeled as $\beta_i = \beta_{corr_i} - f(\beta_i)$, where β_{corr_i} is the correct angle and f is the distortion function. Our goal is to find an approximation for f and use it to compensate for each observed bearing, such that $\beta_{corr_i} = \beta_i + \hat{f}(\beta_i)$, where \hat{f} is the approximation for f . A convenient way to approximate such continuous periodic functions is by their Fourier coefficients. Consider the following second-order approximation:

$$\hat{f}(\beta_i) = A \cos(\beta_i) + B \sin(\beta_i) + C \cos(2\beta_i) + D \sin(2\beta_i) + E. \quad (5)$$

We wish to estimate the coefficients A, B, C, D , and E . The average value of \hat{f} must be zero, because otherwise the correcting function contributes a constant rotation component to the robot. Therefore,

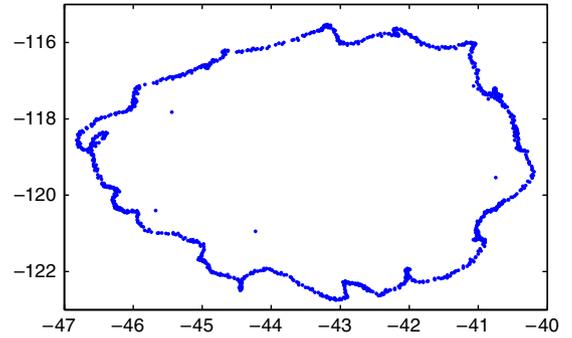


Fig. 3. Distorted localization results due to uncalibrated hardware. The axes: (P_x, P_y) , in cm. The points are the localization system output, which disperses over 7 cm instead of being constant.

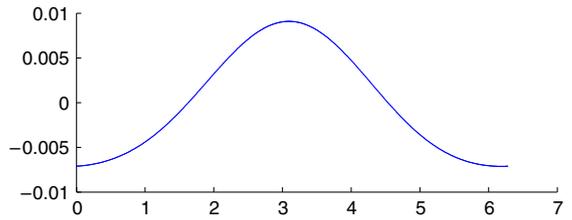


Fig. 4. Illustration of the correction function f for each angle in the range of $[0, 2\pi]$ for $(A, B, C, D) = (-0.0081, 0.0005, 0.00098, -4.12e - 005)$. The horizontal axis represents the bearing in radians, the vertical axis represents the correction value in radians.

we set $E = 0$. Since the distortion transforms the bearing of each landmark differently, the bearings become inconsistent with the real locations of the landmarks, and the MSE of the localization result increases. Therefore, we want to find the coefficients that minimize the sum of the MSE over the set of scans:

$$\{A, B, C, D\} = \operatorname{argmin} \sum_{j=1}^M S_j, \quad (6)$$

where M is the number of scans in the set, and S_j is the MSE of the localization result of scan j given in Eq. (4). In this case however the estimated angular error is calculated as $\Delta\beta_i = \beta_{c_i} - \theta - (\beta_i + \hat{f}(\beta_i))$.

An example of a correction function f obtained by our method is shown in Fig. 4. Our implementation uses the Nelder–Mead simplex method, which is an unconstrained nonlinear minimization technique [19].

To test the hardware calibration method, one can place the robot at different locations, without the restriction of using a fixed set of landmarks. In our experiment we placed the robot at three different locations. Approximately 2000 scans were made at each point. As shown in Fig. 5, the same correction parameters $(A, B, C, D) = (-0.0081, 0.0005, 0.00098, -4.12e - 005)$ transformed the three twisted circles on the top row into normal distributions of points in the second row. The third row of Fig. 5 shows the density distribution of S_j before and after the correction was applied, demonstrating that the error values have been reduced considerably.

The resulting robot positions are scattered within circles of diameter 2 cm, while 95% of the points lie within a diameter of 1.3 cm. These results show that a single correction function with the recovered coefficients corrects the bias accurately invariantly to the robot's position. It is important to stress here that the dispersion of the localization results was not a parameter which the algorithm tried to minimize. Therefore, the fact that the resulting dispersion is small demonstrates that the correction achieved by the algorithm improves the localization quality. The fact that the orientation of the system remains correct is shown in the final test of the system which is discussed in Section 7.

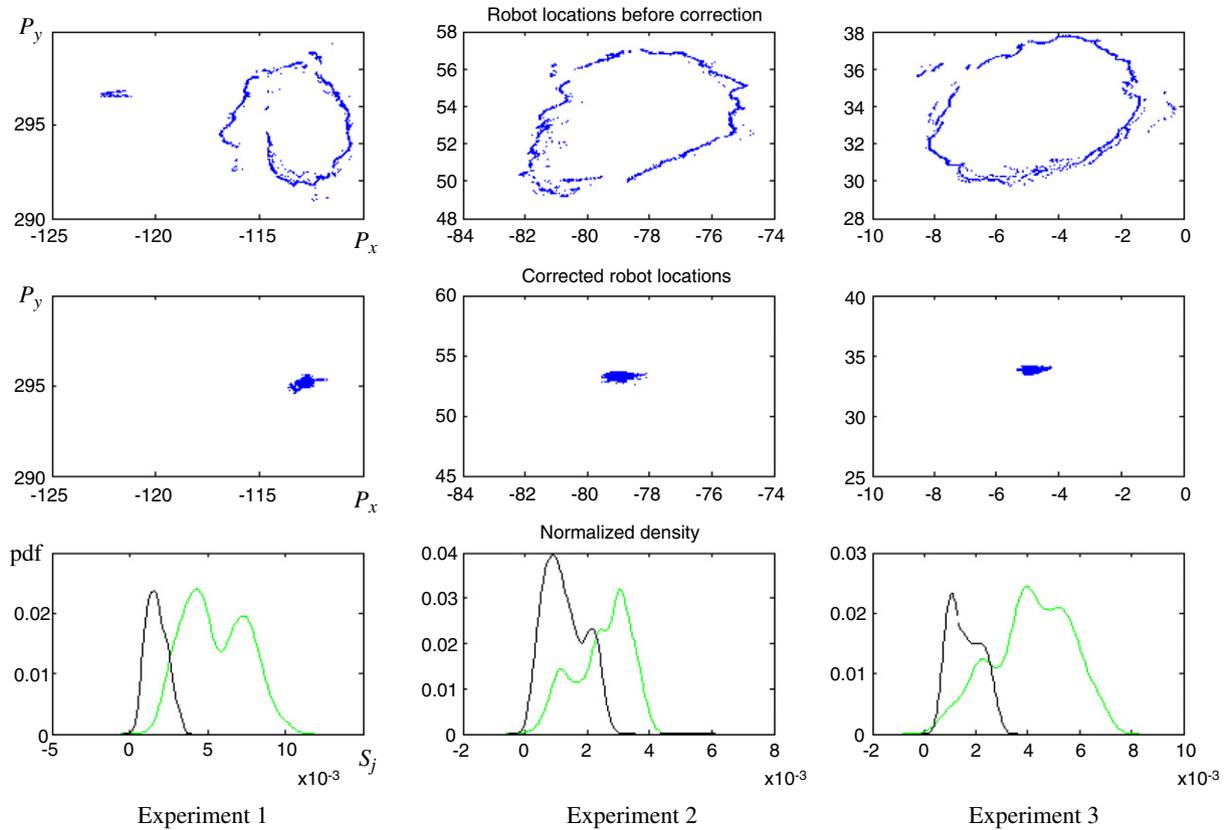


Fig. 5. Experiments for correction of distortion due to uncalibrated hardware. The blue points are the localization system outputs at fixed location but different orientations of the robot. The third row shows the density distribution of S_j before (green) and after (black) the correction was applied. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4.2. Calibration of the landmarks set given calibrated hardware

Inaccurate manual measurement of landmark locations is an additional distortion factor. Its effect on the angular error of a single landmark bearing measurement increases grossly as the robot approaches the landmark. Errors of this type cause bias in the localization results, create regions of bad localization near groups of inaccurately measured landmarks, and therefore must be handled.

The landmarks set calibration consists of two stages. In the first stage, positions of the landmarks that minimize the MSE (Eq. (4)) over a set of measurements at different locations and orientations of the robot are calculated. The result however is invariant to planar similarity transformations. Therefore in the second stage, a transformation which yields the closest positions to those initially measured is found. Thus, the solution found is as close as possible to the scene metric as measured initially by the user.

4.2.1. Calculation of internally consistent locations of the landmarks

At the beginning of the process, we have a set of landmark coordinates measured by the user denoted by $V_0 = \{^0P_i\}_{i=1..n}$. They are inaccurate, but close to the true locations. Our goal is to find the landmark locations that minimize the MSE over a group of M scans. We denote this set by V_1 ,

$$V_1 = \operatorname{argmin}_{\{^1P_i\}_{i=1..n}} \sum_{j=1}^M S_j. \quad (7)$$

To find this set, we use the Nelder–Mead algorithm with a starting point V_0 . During the optimization process, when the landmark location estimates change, so do the estimated robot positions and the MSEs of the scans. The result of this stage is a set of internally consistent locations of the landmarks.

In our experiments, performed in a room of size 10×10 m, each measurement of a landmark location had an error bounded by 5 cm. In order to find more accurate locations of the landmarks, bearing data were collected from several positions of the robot distributed evenly over the whole scene. In our experiment, for $n = 11$ landmarks, $M = 50$ measurements were taken.

4.2.2. Correct the locations of the landmarks to fit the original metrics

Landmark locations found in the first stage of the landmarks set calibration correspond to the metrics of the scene up to a similarity transformation, because when such a transformation is applied to the locations of the robot and the landmarks—the bearings do not change. This is also true if the landmark locations were obtained by a SLAM algorithm such as the one described in [3]. This phenomenon causes the localization results not to be in sync with the objects in the scene and the zones that the robot needs to reach or avoid, all measured in the scene metrics. The manual measurements of initial locations of the landmarks are inaccurate, but they are the only source of information on the metrics of the scene that we have. Therefore they are used to perform a global correction. We wish to find the similarity transformation consisting of a rotation matrix R by angle γ , a translation vector t and a scale parameter s . Applying this transformation to V_1 will yield V_2 whose distance from the measured landmark locations V_0 is minimal. To solve this problem we formulate a new minimization problem, where the target function is

$$\{s, \gamma, t\} = \operatorname{argmin} \|V_2 - V_0\|, \quad (8)$$

where $V_{2_i} = s \cdot R \cdot V_{1_i} + t$. The optimization process starts with the values $s = 1$, $\gamma = 0$, $t = (0, 0)$.

An example of a landmarks set calibration is presented in Table 1. Note that in the first two rows of the table, the $V_0 - V_1$

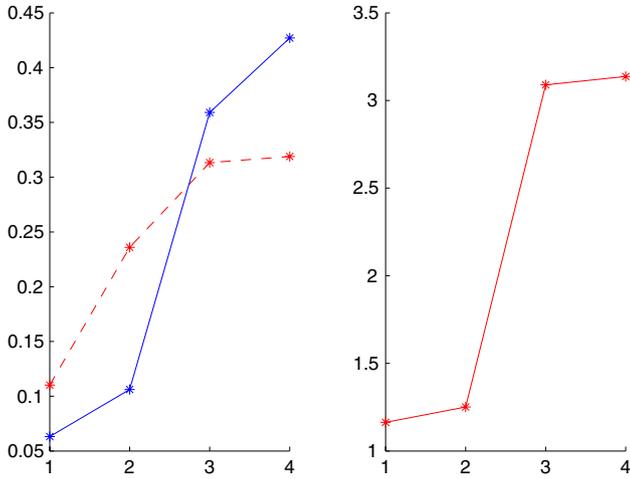


Fig. 6. Correction of the localization results during robot translation. Left: the robot speed in m/s, as estimated by odometry (dashed) compared to the optimization method (solid) for the four experiments. Right: the MSE ratio before and after the correction.

at time $t_i = t_0 + \Delta t_i$. Then, for the i th landmark, $m_i = \mathbf{v}_{rob} \cdot \Delta t_i$ is the distance that the robot moved along its motion axis.

Suppose that the bearing to the first landmark in the scan is β_0 . Then, it takes the mirror rotating at angular velocity of av_{mirr} time interval of Δt_i to cover $\beta_i - \beta_0$ radians. Therefore,

$$\Delta t_i = \frac{\beta_i - \beta_0}{2\pi} / av_{mirr}.$$

The robot has moved in that time period

$$m_i = \Delta t_i \mathbf{v}_{rob}.$$

To compensate for the robot motion we estimate the corrected landmark position C_i at time t_0 as

$$C_i = P_i - m_i.$$

We consider only the case when the robot moves only in the forward or the backward direction. Hence,

$$\mathbf{v}_{rob} = V(\cos(\theta), \sin(\theta)),$$

where θ is the rough estimate of the robot orientation obtained from the localization estimation performed on the biased measurements and V is the robot's speed. V is estimated by minimizing

$$\{P, \theta, V\} = \arg \min \sum \frac{\Delta \beta_i^2}{n},$$

where $\Delta \beta_i = (\beta_{c_i} - \theta) - \beta_i$, and in this case

$$\beta_{c_i} = \arctan2(C_{iy} - P_y, C_{ix} - P_x).$$

To test the correction process, four experiments were performed in which the robot was given a command to move in a straight line. Fig. 6(left) compares the speed of the robot measured by the odometry to the speed measured by the algorithm. Fig. 6(right) presents the ratio of the sum of the MSE of the experiments before and after the correction. As the speed increases so does the ratio indicating that the increased error can be explained by the model. Finally, the average estimated position of the robot changed only slightly as a result of the corrected model, from 0.24 cm in the lowest speed to 1.47 cm in the highest speed (Fig. 7).

5.2. Correction for rotation

We now turn to deal with robot rotation. In this case the robot rotates at an angular speed of av_{rob} . Suppose that we are observing a landmark at bearing β_i , the scan starts at time t_0 and the landmark is detected at time t_i . Then,

$$\Delta t_i = t_i - t_0 = \beta_i / av_{mirr}.$$

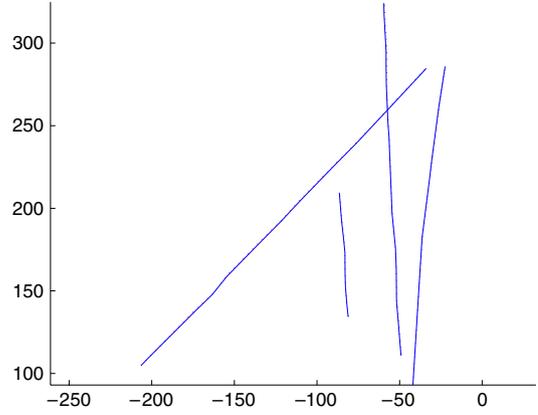


Fig. 7. Correction of the localization results during robot translation obtained from four experiments. The localization results without correction (dashed) compared to the results after correction (solid). The correction is between 0.24 cm for slow speed to 1.47 cm for high speed.

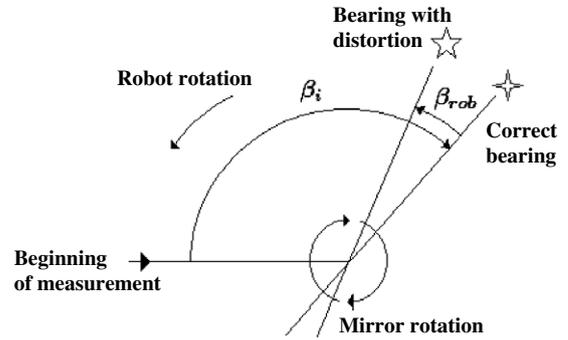


Fig. 8. Illustration of angle measurement distortion due to robot rotation.

At this time the mirror of the sensor has rotated by $\beta_{mirr_i} = av_{mirr} \cdot \Delta t_i$ and the robot rotated by $\beta_{rob_i} = av_{rob} \cdot \Delta t_i$. Both of the angles contribute to the measured bearing, $\beta_i = \beta_{rob_i} + \beta_{mirr_i}$. We are interested to measure only the mirror rotation, so β_{rob} must be subtracted: $\beta_{mirr_i} = \beta_i - \beta_{rob_i}$. An illustration of this model is presented in Fig. 8.

Assuming that within a single scan the robot's angular velocity is close to constant, the velocity of the robot that minimizes the MSE of the result for all the observed landmarks must be found. The setting is similar to Eq. (4), but this time

$$\Delta \beta_i = (\beta_{c_i} - \theta) - \beta_{mirr_i}.$$

In other words, we have to solve the following minimization problem:

$$av_{rob} = \arg \min \sum_{i=1}^n \frac{\Delta \beta_i^2}{n}, \tag{9}$$

where $\Delta \beta_i = (\beta_{c_i} - \theta) - \beta_i(1 - av_{rob}/av_{mirr})$.

The function of the MSE versus the approximated angular velocity is close to parabolic and smooth, as can be seen in Fig. 9. This is because $(\beta_{c_i} - \theta)$ changes only slightly yielding a parabolic function in av_{rob} .

In such cases, when a single minimum exists, Brent's method of parabolic interpolation in one dimension [20] can be applied to search efficiently for the minimum of the error function. Fig. 10 shows the results of the correction of the localization during the robot's rotation. In this experiment the robot was rotating at speeds of up to 45° per second around its axis at a fixed location. 170 scans were taken during the experiment and for each scan the angular velocity was estimated. The corrected results lie within a circle of diameter 2 cm, while the raw results are scattered within

Table 3
An example of a landmarks set with a misidentified landmark (28).

| Index | 2 | 4 | 7 | 8 | 15 | 21 | 26 | 28 |
|-------------|--------|-------|-------|--------|--------|--------|--------|-------|
| X | 524.0 | -78.1 | 523.8 | -429.3 | 71.1 | -339.1 | -434.4 | 523.4 |
| Y | 192.7 | 458.3 | -2.1 | -25.4 | -546.9 | -512.6 | 106.3 | 68.4 |
| Bearing | 2.88 | 1.33 | 3.28 | -0.13 | -1.59 | -0.94 | 0.15 | 2.13 |
| Error value | Active | | | | | | | |
| 0.68 | + | + | + | + | + | + | + | + |
| 1.46 | (+) | + | (+) | - | - | - | - | (+) |
| 1.26e-4 | - | - | - | (+) | (+) | + | - | (+) |
| 2.32e-7 | + | (+) | + | + | + | (+) | (+) | - |

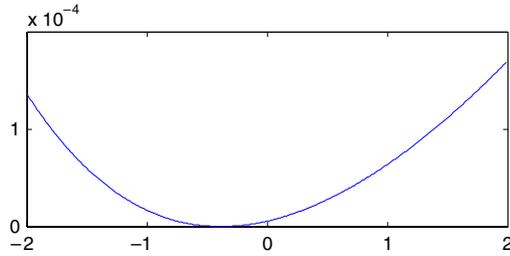


Fig. 9. MSE of a scan versus assumed robot angular velocity. The minimum point corresponds to the correct angular velocity of the robot.

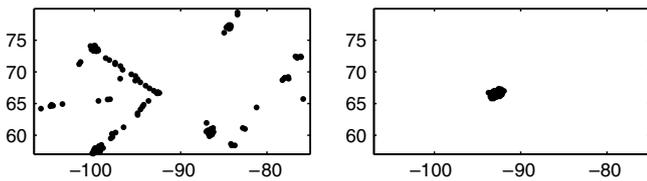


Fig. 10. Correction of the localization results during robot rotation. The axes are the coordinates of the scene in cm. The points are the localization system outputs. The distorted results are shown on the left and the corrected results are shown on the right.

a 30 cm diameter. The sum of the MSE over the set is $2.8e-5$ for the corrected results as compared to $1.7e-2$ for the initial results. As can be seen from the results, without the correction localization is simply not possible during robot rotation.

As a result of our experiments we conclude that the correction obtained from modeling robot rotation far exceeds the correction due to modeling robot translation. We decided that for robots that move straight slower than 4 m/s while the average distance to the observed landmarks is greater than 2 m, the small increase in accuracy is not worth the increased complexity and the additional recognized landmarks required to deal with errors resulting from straight motion.

6. Misidentification of landmarks

In landmark-based localization systems, the sensor has to identify each landmark and measure its bearing. Occasionally a small number of landmarks will be misidentified. The meaning of such an error for the localization system is that it detects a landmark at a completely incorrect angle. In order for a localization system to be robust, this problem has to be addressed.

In order to overcome the problem of index misidentification, we must perform the following steps.

1. Detect the existence of index misidentification.
2. Eliminate the misidentified landmarks from the scan.
3. Calculate a result based on the landmarks that were identified correctly, without losing the scan.

The error values that appear due to index misidentification are greater by an order of magnitude than the error values that

remain after the previous corrections. Thus, it is easy to detect index misidentifications. The problem that remains is to find the misidentified landmarks. We treat it as a set of measurements that contain a majority of true noisy measurements and several grossly erroneous measurements. RANSAC is a generic technique for treating such cases [18]. Based on it we formulated an algorithm for detection of misidentified landmarks, presented in Algorithm 2.

The algorithm treats a set of n bearings I for which a high MSE was computed. d is an experimentally determined threshold value, chosen to identify exceptionally high angular errors. The advantage of the RANSAC-based approach is that it performs the localization at most K times. Its disadvantage is its probabilistic nature which leads to incompleteness. The probability to miss the correct solution p_{err} is

$$p_{err} = (1 - p_{corr}^s)^K,$$

where p_{corr} is the probability that a landmark is correctly identified, s the number of landmarks selected at random, and K the number of iterations performed. Assume for example that $p_{corr} = 0.9$ and using the value of $s = 5$. We chose $K = 20$ for which the value obtained for p_{err} is $2e - 8$. Under these conditions the probability of the RANSAC-based procedure to fail is negligible.

Algorithm 2 Detection of misidentified landmarks

- 1: **for** $i = 1$ to K **do**
- 2: Choose at random 3 landmarks (5 landmarks if the robot moves) from I and denote this set A .
- 3: Calculate the localization result for the set A and the MSE. If the robot moves, estimate also the angular velocity.
- 4: $S = \emptyset$.
- 5: **for** Each landmark $b \in I \setminus A$ **do**
- 6: Compute the angle β_{cb} according to Eq. 2.
- 7: Compute $\Delta\beta_b$ according to Eq. 3.
- 8: **if** $\Delta\beta_b < d$ **then**
- 9: consider that landmark b “agrees” with the localization result of set A , and therefore $S = S \cup \{b\}$.
- 10: **end if**
- 11: **end for**
- 12: **If** $|S| > |I \setminus A|/2$ accept the result of the set S and terminate.
- 13: **end for**
- 14: Report failure.

A typical run of the RANSAC-based method is presented in Table 3. The upper part of the table contains the landmark indices, locations and bearings. The top row of the bottom part shows a result based on all the observed landmarks, that invoked the RANSAC-based method because of its high error value. In the other rows of the bottom part “(+)” denotes landmarks belonging to the set A from which the localization was estimated, “+” denotes landmarks belonging to the set S of landmarks that agree with the localization result, and “-” denotes landmarks which do not. In the first two tests the size of S is only one. In the third test, it was four. The remaining landmark 28 was the misidentified landmark. It appeared in the first two trials and therefore yielded incorrect localizations which were discarded by the algorithm.

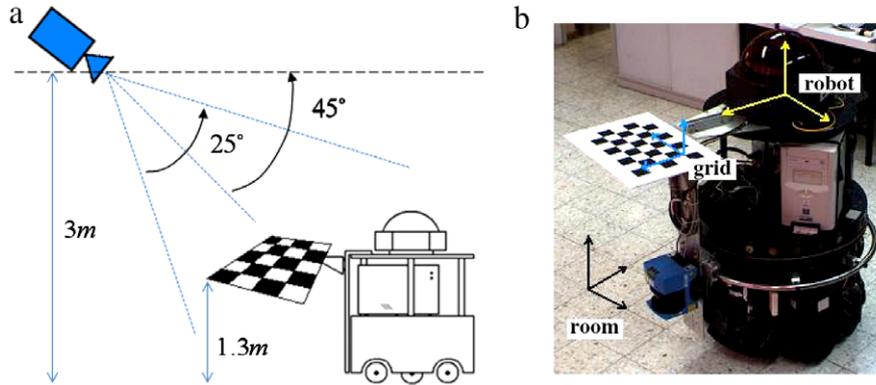


Fig. 11. (a) Placement of the camera relative to the robot; (b) MRV4 mobile robot with a calibration grid mounted on it, and the relevant coordinate frames.

7. Testing the system

To test the localization system, we compared its results to the results of another method that measures position and orientation of the robot using a camera and a checkered board. We used the camera calibration toolbox [21] to perform an alternative localization. Its function “Comp. Extrinsic” computes the transformation matrix from the reference frame of the calibration grid to the camera reference frame, c_gT . A high-resolution and sensitivity camera was used, in order to ensure accurate results. The camera was placed as shown on Fig. 11a. The grid was mounted on the robot, and for some of the measurements of the localization system, corresponding grid location and orientation measurements were taken. To avoid smearing, images were taken only when the robot was stationary. These corresponding measurements test the global correctness of the system, while the experiment in Section 5 demonstrates the local correctness while the robot is in motion. An image of the robot with the mounted grid and the relevant coordinate systems is shown in Fig. 11b.

An image of the grid at the room’s coordinates origin point was taken in order to extract the transformation matrix from the camera system to the room origin, c_rT . From each image, c_gT was extracted. Then, it is possible to calculate the transformation matrix from the room reference frame to the grid reference frame: ${}^r_gT = {}^r_cT \cdot {}^c_gT = {}^c_rT^{-1} \cdot {}^c_gT$. There is a constant transformation g_rT between the robot reference frame and the grid reference frame. This transformation is estimated as the average transformation computed from all the images between these coordinate frames. We are interested in the X and Y translation components of ${}^r_rT = {}^r_gT \cdot {}^g_rT$, and its rotation θ in the ${}^rX, {}^rY$ plane.

Experimental results

In the experiments we used a Denning Robotics MRV4 mobile robot with a LaserNav laser sensor. A new setup of 10 landmarks was used for this experiment, and the system was fully recalibrated. The robot moved on curves between 15 stops, while measuring its position continuously. At each stop an image of the grid was also taken. The standard deviation of robot localization using the Camera Calibration Toolbox for the rX and rY axes respectively was (3.8, 2.5) mm. This STD was calculated internally by the camera calibration toolbox. The difference between the localizations was on average (1.7, 2.5) mm, with the STD of (2.1, 1.7) mm. This shows that the localization system STD is bounded by (5.9, 4.2) mm. The angular error STD derived in the same way was 0.07° . These results confirm the accuracy found in the previous experiments, ± 1.5 cm.

Fig. 12 plots part of the results. The results show that during the motion between the stationary points the localization results

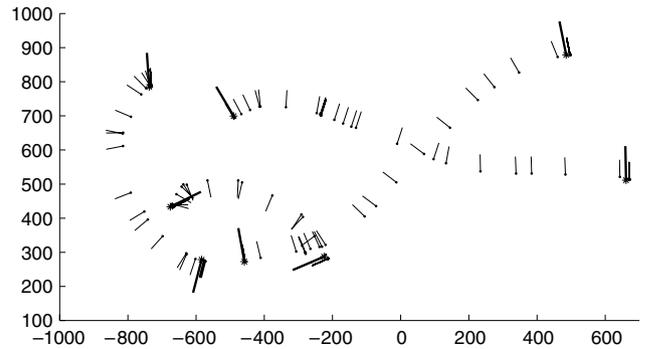


Fig. 12. Robot positions and orientations rotated clockwise by 90° measured by the camera when the robot was stationary (black asterisks and thick lines), and by the localization system (black dots and thin lines). The axes are the scene coordinates in mm.

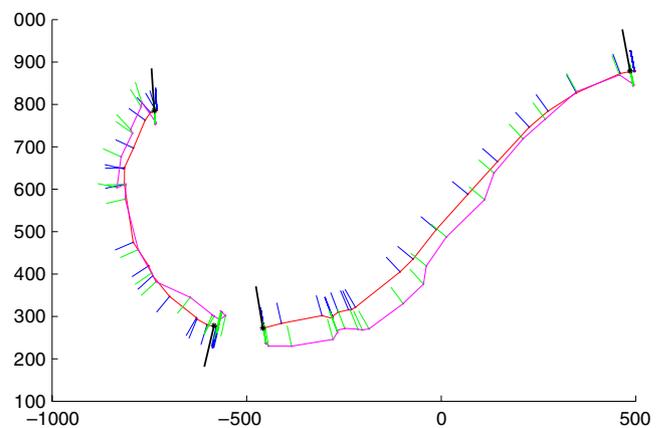


Fig. 13. Comparison between the localization results without (cyan and green) and with (red & blue) the corrections, and the camera based measurements (black). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

are smooth and frequent. Their local accuracy has been shown in Section 5.

Fig. 13 compares the results obtained with and without the various corrections described in this paper. Two of the motion sequences from Fig. 12 are shown with and without the corrections. Note that the corrected motion is much smoother and the localization results are much closer to the ones obtained by the vision system. In the experiment several occurrences of misidentified landmarks occurred. The algorithm dealt with these cases correctly. When no corrections were performed the localization results yielded gross errors yielding errors of up to 600 cm. We therefore do not show these results in the figure.

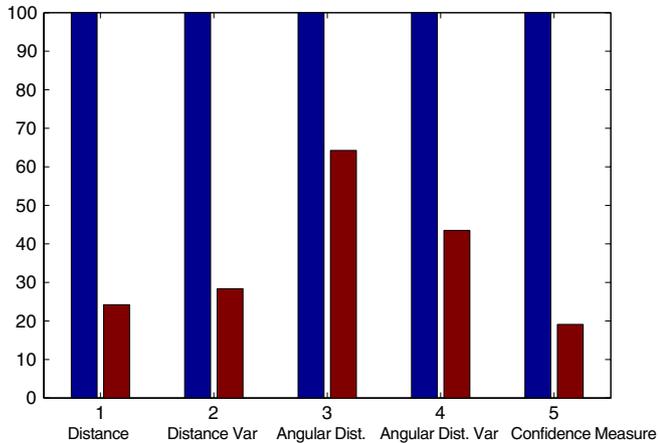


Fig. 14. Comparison between the localization errors without and with the corrections (except landmark index misidentifications) with respect to the camera-based measurements. The confidence measure values are also shown. The values before the correction are plotted in blue as 100% whereas the measures with the correction (in red) are scaled respectively.

Finally, in Fig. 14 a comparison of the location and orientation results, error values and their variances are shown together with the confidence measure values. The measures are only compared when the robot is stationary so the vision-based localization can be performed accurately. Still, the reduction in the positional error and confidence measure are by factors of 4–5. The correction in the orientation is more modest.

8. Discussion and conclusions

In this paper we described the steps that have to be taken to transform an efficient triangulation method and an off-the-shelf landmark bearings sensor to a complete localization system which yields accurate and robust performance. With these enhancements the localization system has been incorporated into a mobile robot application. The robot works successfully in an environment with obstacles and areas that it should avoid. The localization system enables the robot to plan paths, to avoid areas out of its workspace, and to reach important locations like the docking station or the automated storage device.

A simple extension to our algorithm enables automatic addition of new landmarks to the system. The robot has simply to traverse the workspace, record the bearings of the new landmark, calculate its rough location and then run the optimization procedure from Section 4.2.

At this stage our results can be used within the framework of an extended Kalman filter as was done in [17,11]. However, making the maximal effort to generate the most accurate and robust results from a single bearings measurement of a landmarks set, as was presented here, is a prerequisite for a good EKF-based localization system.

To increase the availability of data for localization when landmarks are occluded, it is possible in the future to deal with cases in which less than the minimal number of landmark bearing measurements are available at a single observation, within the framework of the EKF or a particle filter. This will enable mobile robots and AGVs to perform better in dynamic environments where landmarks are frequently occluded by moving objects.

Acknowledgment

The authors would like to thank the Grinblatt Fund for the Promotion of Research and Teaching in Production Engineering.

References

- [1] J.J. Leonard, H.F. Durrant-Whyte, Simultaneous map building and localization for an autonomous mobile robot in: IEEE International Conference on Intelligent Robots and Systems, 1991, pp. 1442–1447.

- [2] G. Dissanayake, P. Newman, S. Clark, H. Durrant-Whyte, M. Csorba, A solution to the simultaneous localization and map building SLAM problem, IEEE Transactions on Robotics and Automation 17 (3) (2001) 229–241.
- [3] O. Faugeras, L. Quan, P. Strum, Self-Calibration of a 1D Projective Camera and Its Application to the Self-Calibration of a 2D Projective Camera, IEEE Transactions on Pattern Analysis and Machine Intelligence 22 (10) (2000) 1179–1185.
- [4] S. Thirthala, M. Pollefeys, Multi-view geometry of 1D radial cameras and its application to omnidirectional camera calibration in: Proc. ICCV, 2005, pp. 1539–1546.
- [5] C. Sagues, A.C. Murillo, J.J. Guerrero, T. Goedeme, T. Tuytelaars, L. Van Gool, Localization with omnidirectional images using the radial trifocal tensor. in: Proc. of the IEEE Int. Conf. on Robotics and Automation, 2006.
- [6] M. Betke, L. Gurvits, Mobile robot Localization using landmarks, IEEE Transactions on Robotics and Automation 13 (2) (1997) 251–263.
- [7] D. Sinriech, S. Shoval, Landmark configuration for absolute positioning of autonomous vehicles, IIE Transactions 32 (7) (2000) 613–624.
- [8] K.T. Sutherland, W.B. Thompson, Localizing in unstructured environments: Dealing with the errors, IEEE Transactions on Robotics and Automation 10 (6) (1994) 740–754.
- [9] U. Wiklund, U. Anderson, K. Hyyppa, AGV navigation by angle measurements in: Proc. of 6th Inter. Conf. on AGV Systems, 1988, pp. 199–212.
- [10] U.D. Hanebeck, G. Schmidt, Set theoretic localization of fast mobile robots using angle measurements in: IEEE Inter. Conf. on Robotics and Automation, 1996, pp. 1387–1394.
- [11] H. Hu, D. Gu, Landmark-based Navigation of Industrial Mobile Robots, International Journal of Industry Robot 27 (6) (2000) 458–467.
- [12] C.B. Madsen, C.S. Andersen, Optimal landmark selection for triangulation of robot position, International Journal of Robotics and Autonomous Systems 23 (4) (1998) 277–292.
- [13] I. Shimshoni, On mobile robot localization from landmark bearings, IEEE Transactions on Robotics and Automation 18 (3) (2002) 971–976.
- [14] K.E. Bekris, A.A. Argyros, L.E. Kavraki, Angle-based methods for mobile robot navigation: Reaching the entire plane in: Proceedings of the International Conference on Robotics and Automation, New Orleans, USA, 2004, pp. 2373–2378.
- [15] J.A. Batlle, J.M. Font, J. Escoda, Dynamic positioning of a mobile robot using a laser-based goniometer. in: Proceedings of the 5th IFAC Symposium on Intelligent Autonomous Vehicles, Lisboa, Portugal, 2004.
- [16] A.W. Stroupe, K. Sikorski, T. Balch, Constraint-based landmark localization, in: G. Kaminka, P. Lima, R. Rojas (Eds.), RoboCup 2002: Robot Soccer World Cup IV., in: LNCS, vol. 2752, Springer-Verlag, 2002, pp. 8–24.
- [17] F. Matia, D. Rodriguez-Losada, R. Galan, A. Jimenez, Experiments at trade fairs with blacky the robot, in: IEEE International Conference on Intelligent Robots and Systems, Lausanne, Switzerland. Proceedings of Robots in Exhibitions Workshop, 2002, pp. 61–66.
- [18] M.A. Fischler, R.C. Bolles, Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography, Communications of the ACM 24 (6) (1981) 381–395.
- [19] J.A. Nelder, R. Mead, A simplex method for function minimisation, Journal of Computer 7 (1965) 308–313.
- [20] W.H. Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling, Numerical Recipes in C: The Art of Scientific Computing, Cambridge University Press, 1992.
- [21] J.Y. Bouguet, Camera calibration toolbox http://www.vision.caltech.edu/bouguetj/calib_doc/ August 2006.



I. Loevsky is a graduate student in the faculty of Industrial Engineering and Management of the Technion in Haifa. Igal conducts research on mobile robots for indoor applications in the laboratory of Computer Integrated Manufacturing of the faculty. He worked several years in computer vision research projects in the Intelligent Systems Laboratory of the Computer Science faculty of the Technion, and in MobilEye, a leading automobile vision systems manufacturer.



I. Shimshoni received his B.Sc. in mathematics from the Hebrew University in Jerusalem, his M.Sc. in computer science from the Weizmann Institute of Science, and his Ph.D. in computer science from the University of Illinois at Urbana Champaign (UIUC). Ilan was a post-doctorate fellow at the faculty of computer science at the technion, from 1995–1998, and was a member of the faculty of industrial engineering and management from 1998–2005. He joined the department of Management Information Systems (MIS) at Haifa University in October 2005. His research areas are: computer vision, robotics and

computer graphics.