

- [14] L. Stark *et al.*, "Telerobotics: Display, control, and communication problems," *IEEE J. Robotics Automat.*, vol. RA-3, pp. 67-75, Feb. 1987.
- [15] D. Regan, K. Beverly, and M. Cynader, "The visual perception of motion in depth," *Sci. Amer.*, vol. 241, p. 136-151, July 1979.
- [16] C. W. Tyler, "Stereoscopic depth movement: Two eyes less sensitive than one," *Science*, vol. 174, pp. 958-961, Nov. 1971.
- [17] W. J. Zinn and H. Solomon, "A comparison of static and dynamic stereoacuity," *J. Amer. Optom. Assoc.*, vol. 56, pp. 712-715, Sept. 1985.
- [18] G. Westheimer and S. P. McKee, "Stereoscopic acuity for moving retinal images," *J. Opt. Soc. Amer.*, vol. 68, pp. 450-455, Apr. 1978.
- [19] G. Perlman, *ANOVA(ICSL)-UNIX Programmer's Manual*, 1980.

## Mobile Robot Localization by Tracking Geometric Beacons

John J. Leonard and Hugh F. Durrant-Whyte

**Abstract**—This short paper presents the application of the extended Kalman filter (EKF) to the problem of mobile robot navigation in a known environment. We have developed an algorithm for model-based localization that relies on the concept of a *geometric beacon*—a naturally occurring environment feature that can be reliably observed in successive sensor measurements and can be accurately described in terms of a concise geometric parameterization. The algorithm is based on an EKF that utilizes matches between observed geometric beacons and an *a priori* map of beacon locations. We describe two implementations of this navigation algorithm, both of which use sonar. The first implementation uses a simple vehicle with point kinematics equipped with a single rotating sonar. The second implementation uses a "Robuter" mobile robot and employs six static sonar transducers to provide localization information while the vehicle moves at typical speeds of 30 cm/s.

### I. THE NAVIGATION PROBLEM

Stated most simply, the problem of navigation can be summarized into answering the following three questions: "where am I?", "where am I going?" and "how should I get there?". The first question is one of localization; how can I work out where I am in a given environment, based on what I can see and what I have previously been told? The second and third questions are essentially those of specifying a goal and being able to plan a path that results in achieving this goal. We are principally concerned with the first, localization, question, and maintain that finding a robust and reliable solution to this problem is an essential precursor to answering the remaining two questions.

The problem of position determination has been of considerable interest over the last 4000 years. The basic process of distance measurement, correlation, and triangulation was known to the Phoenicians,<sup>1</sup> who successfully managed to build and maintain quite accurate maps of the Mediterranean area. Today, navigation is a well-understood quantitative science, used routinely in maritime and

aviation applications [22]. Given this, the question must be asked as to why robust and reliable autonomous mobile robot navigation remains such a difficult problem. In our view, the reason for this is clear; it is not the navigation process *per se* that is a problem, it is the reliable acquisition or extraction of information about navigation beacons, from sensor information, and the automatic correlation or correspondence of these with some navigation map that makes the *autonomous* navigation problem so difficult.

Implementing a navigation system that uses artificial beacons together with sensors that provide accurate and reliable measurements of beacon location is a straight forward procedure used by many commercial mobile robots today. For example, the GEC-Caterpillar AGV [3] uses a rotating laser to locate itself with respect to a set of bar codes that are fixed at known locations through the AGV's environment. More recently, the TRC Corporation has developed a system for localization that uses retro-reflective strips and ceiling lights as beacons that are observed by vision and active infrared sensors. Our goal for a *competence of localization* is to use the naturally occurring structure of typical indoor environments to achieve comparable performance to artificial beacon systems without modifying the environment.

We have developed a system in which the basic localization algorithm is formalized as a vehicle-tracking problem, employing an extended Kalman filter (EKF) to match beacon observations to a navigation map to maintain an estimate of mobile robot location. Kalman filtering techniques have been used extensively in location estimation problems such as missile tracking and ship navigation [21]. There have been many notable applications of the EKF in mobile robot systems. For example, Dickmanns uses an EKF in a real-time vision system that achieves autonomous road-following at speeds over 80 km/h [7]. Ayache and Faugeras [1], Matthies and Shafer [19], and Kriegman *et al.* [14] have used the EKF for visual map building and motion estimation. These systems address a much more complex task than that considered here, as they start without an *a priori* model. However, the *motion* estimation formulation does not by itself meet our requirements for long-term autonomous *position* estimation, for despite the high accuracy with which the relative motion between frames can be estimated, uncertainty in the globally referenced vehicle position estimate must accumulate with time. Hallam has developed an undersea navigation system that maintains an absolutely referenced estimate of vehicle position in an environment comprised of moving targets and clutter [13]. Our formulation to the problem deals with the much simpler case of a static environment, but has been demonstrated to be successful with real sonar data on several different robots.

### II. THE LOCALIZATION ALGORITHM

In man-made indoor environments, we model the world in terms of geometry and consider each feature of the environment to be a geometric target. A *geometric beacon* is a special type of target that can be reliably observed in successive sensor measurements and that can be accurately described in terms of a concise geometric parameterization. Hence, geometric beacons are stable, naturally occurring environment features that are useful for navigation. The idea of a generalized geometric beacon arises as the result of our earlier work in describing sensors and processing algorithms as "geometry extractors" [9], [11], allowing many different types of information to be integrated easily in a common geometric framework.

With reference to Fig. 1, we denote the position and orientation of the vehicle at time step  $k$  by the state vector  $x(k) =$

Manuscript received August 30, 1989; revised August 30, 1990. This work was supported in part by ESPRIT 1560 (SKIDS) and by the SERC-ACME under Grant GRE/42419. A portion of this work was presented at the IEEE/RSJ International Conference on Intelligent Robot Systems, 1989.

The authors are with the Department of Engineering Science, University of Oxford, Oxford, United Kingdom.  
IEEE Log Number 9042230.

<sup>1</sup>The Phoenicians were pre-Greek seafarers from Syria whose main claim to fame was their hypothesis that the world was flat.

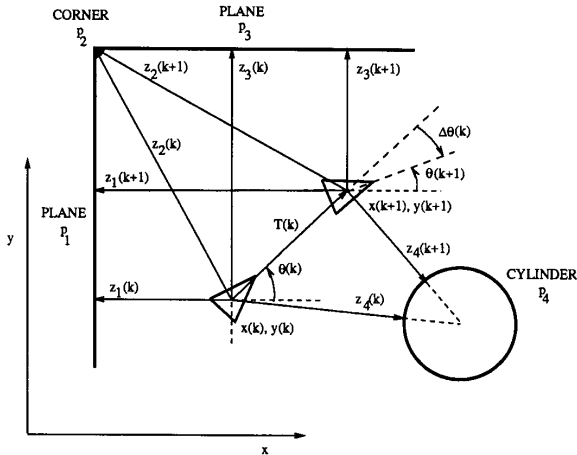


Fig. 1. Localization by concurrent observation of several beacons. Four geometric beacons are in view to an ultrasonic sensor at time  $k$  and time  $k + 1$ : plane  $p_1$ , corner  $p_2$ , plane  $p_3$ , and cylinder  $p_4$ . The sonar measurements  $z_1(k)$  and  $z_3(k)$  are the shortest distance from the vehicle to planes  $p_1$  and  $p_3$  at time  $k$ . The measurement  $z_2(k)$  is the distance from the vehicle to corner  $p_2$  at time  $k$ . Measurement  $z_4(k)$  is the distance to the central axis of cylinder  $p_4$  less the radius of the cylinder.

$[x(k), y(k), \theta(k)]^T$  comprising a Cartesian location and a heading defined with respect to a global coordinate frame. At initialization the robot starts at a known location, and the robot has an *a priori* map of the locations of geometric beacons  $p_i$ . We stress that the map is just a set of beacon locations, not an exhaustively detailed world model. Each beacon is assumed to be precisely known. At each time step, observations  $z_j(k)$  of these beacons are taken. Our goal in the cyclic process is to associate measurements  $z_j(k)$  with the correct beacon  $p_i$  to compute an updated estimate of vehicle position.

The Kalman filter relies on two models: a *plant* model and a *measurement model*. The plant model describes how the vehicle's position  $x(k)$  changes with time in response to a control input  $u(k)$  and a noise disturbance  $v(k)$

$$x(k+1) = F(x(k), u(k)) + v(k), \quad v(k) \sim N(0, Q(k)) \quad (1)$$

where  $F(x(k), u(k))$  is the nonlinear state transition function. We use the notation  $v(k) \sim N(0, Q(k))$  to indicate that this noise source is assumed to be zero mean Gaussian with variance  $Q(k)$  [12].

The measurement model expresses a sensor observation in terms of the vehicle position and the geometry of the beacon being observed and has the form

$$z_j(k) = h_i(p_i, x(k)) + w_j(k), \quad w_j(k) \sim N(0, R_j(k)). \quad (2)$$

The observation function  $h_i(p_i, x(k))$  expresses an observed measurement  $z_j(k)$  as a function of the vehicle location  $x(k)$  and beacon location  $p_i$ . This observation is assumed corrupted by a zero-mean Gaussian noise disturbance  $w_j(k)$  with variance  $R_j(k)$ . The form of the observation function  $h_i(\cdot, \cdot)$  depends on the sensor employed and the type of beacon being observed.

The goal of the cyclic computation is to produce an estimate of the location of the robot<sup>2</sup>  $\hat{x}(k+1|k+1)$  at time step  $k+1$

<sup>2</sup>The term  $\hat{x}(i|j)$  should be read as "the estimate of the vector  $x$  at time step  $i$  given all observations up to time step  $j$ ."

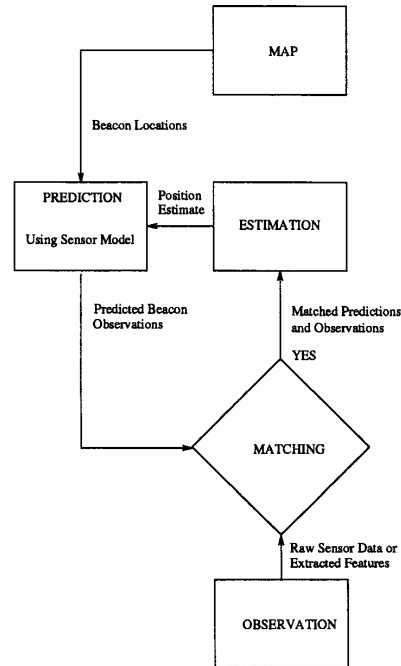


Fig. 2. The localization algorithm.

based on the estimate of the location  $\hat{x}(k|k)$  at time step  $k$ , the control input  $u(k)$  and the new beacon observations  $z_j(k+1)$ . The algorithm employs the following steps: prediction, observation, matching, and estimation. Fig. 2 presents an overview of this cyclic process. We state the Kalman filter equations without derivation and refer the reader to Bar-Shalom and Fortmann [2] and Smith *et al.* [20] for more details.

#### A. Prediction

First, using the plant model and a knowledge of the control input  $u(k)$ , we predict the robot's new location at time step  $k+1$ :

$$\hat{x}(k+1|k) = F(\hat{x}(k|k), u(k)). \quad (3)$$

We next compute  $P(k+1|k)$ , the variance associated with this prediction

$$P(k+1|k) = \nabla F P(k|k) \nabla F^T + Q(k) \quad (4)$$

where  $\nabla F$  is the Jacobian of  $F(\cdot, \cdot)$  obtained by linearizing about the updated state estimate  $\hat{x}(k|k)$ . Next, we use this predicted robot location to generate predicted observations of each geometric beacon  $p_i$

$$\hat{z}_i(k+1) = h_i(p_i, \hat{x}(k+1|k)), \quad i = 1, \dots, N_k. \quad (5)$$

#### B. Observation

The next step is to actually take a number of observations  $z_j(k+1)$  of these different beacons and compare these with our predicted observations. The difference between a prediction  $\hat{z}_i(k+1)$  and an observation  $z_j(k+1)$  is termed the innovation, and it is written as

$$v_{ij}(k+1) = [z_j(k+1) - \hat{z}_i(k+1)] \\ = [z_j(k+1) - h_i(p_i, \hat{x}(k+1|k))]. \quad (6)$$

The innovation covariance can be found by linearizing (2) about the

prediction, squaring, and taking expectations

$$\begin{aligned} S_{ij}(k+1) &\equiv E[v_{ij}(k+1)v_{ij}^T(k+1)] \\ &= \nabla h_i P(k+1|k) \nabla h_i^T + R_j(k+1) \end{aligned} \quad (7)$$

where the Jacobian  $\nabla h_i$  is evaluated at  $\hat{x}(k+1|k)$  and  $p_i$ .

### C. Matching

Around each predicted measurement, we set up a validation gate [2] in which we are prepared to accept beacon observations

$$v_{ij}(k+1)S_{ij}^{-1}(k+1)v_{ij}^T(k+1) \leq g^2. \quad (8)$$

This equation is used to test each sensor observation  $z_j(k+1)$  for membership in the validation gate for each predicted measurement. When a single observation falls in a validation gate, we get a successful match. Measurements that do not fall in any validation gate are simply ignored for localization. More complex data association scenarios can arise when a measurement falls in two validation regions or when two or more measurements fall in a single validation region. At this stage, such measurements are simply ignored by the algorithm as outlier rejection is vital for successful localization. This has proven acceptable thus far, but we are investigating the use of probabilistic data association techniques [2] to resolve these ambiguities.

### D. Estimation

The final step is to use successfully matched predictions and observations to compute  $\hat{x}(k+1|k+1)$ , the updated vehicle location estimate. To do so, we use a parallel update procedure [23]. We first stack the validated measurements  $z_j(k+1)$  into a single vector to form  $z(k+1)$ , the composite measurement vector for time  $k+1$ , and designate the composite innovation  $v(k+1)$ . Next, we stack the measurement jacobians  $\nabla h_i$  for each validated measurement together to form the composite measurement jacobian  $\nabla h$ . Using a composite noise matrix  $R(k+1) = \text{diag}[R_j(k+1)]$ , we then compute the composite innovation covariance  $S(k+1)$  as in (7). We then utilize the standard result that the Kalman gain can be written as

$$W(k+1) = P(k+1|k) \nabla h^T S^{-1}(k+1) \quad (9)$$

to compute the updated vehicle position estimate

$$\hat{x}(k+1|k+1) = \hat{x}(k+1|k) + W(k+1)v(k+1) \quad (10)$$

with associated variance

$$\begin{aligned} P(k+1|k+1) &= P(k+1|k) \\ &\quad - W(k+1)S(k+1)W^T(k+1). \end{aligned} \quad (11)$$

Fig. 3 shows a variety of simulated sonar runs that illustrate the localization process. Fig. 3(a) is a run with no visible beacons, showing how position uncertainty grows in accordance with the system plant model [see (1)] if no sensor observations are available. Fig. 3(b) shows the use of a wall as a beacon. The only part of a specular planar target that is visible to a sonar sensor is the portion of the wall that is perpendicular to the incident sonar beam. Thus, uncertainty grows as before for the first two cycles of the filter. After the wall comes into view, range measurements to the wall provide a position update perpendicular to the wall, while uncertainty continues to grow in the parallel direction. Fig. 3(c) and (d) shows position estimation by using range updates from multiple walls. As different walls come in and out of the sonar's view, uncertainty is reduced in different degrees of freedom.

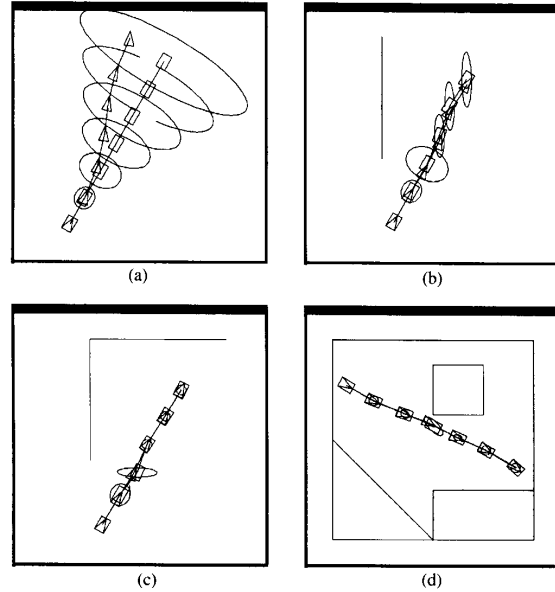


Fig. 3. Simulations illustrating the use of walls as geometric beacons.

## III. LOCALIZATION USING A SINGLE ROTATING SONAR

In this section we present an implementation of the navigation algorithm based on a simple mobile robot with point kinematics, equipped with a single servo-mounted sonar. (The Appendix summarizes our approach to extracting and predicting beacon information for the standard Polaroid sonar system that is used in both of these implementations.) The results we present here are taken from an experiment where the vehicle's position was accurately measured by hand at a sequence of positions from which real sonar scans were taken. The Kalman filter was then run off-line by introducing artificial process noise to the position estimates.

Fig. 4 illustrates one cycle of the localization algorithm. The system has a 2-D global representation of the environment consisting of a list of line segments and the corners they define, assumed known with absolute confidence. The vehicle starts from a known location. This starting point is the left-most triangle in Fig. 4(a). The initial state covariance matrix  $P(0|0)$  is set to zero. We show a run in which observation is suppressed for the first two time steps.

### A. Prediction and Observation

Fig. 4(a) shows a sonar scan of 612 measurements taken at time step 3. Based on the predicted vehicle location, Fig. 4(b) shows predicted beacon observations generated using (5), with corresponding validation regions generated using (7). The validation regions take the form of circular arcs blurred in the perpendicular direction by range measurement noise and uncertainty in the *a priori* vehicle position estimate. In conjunction with this, Fig. 4(c) shows regions of constant depth (RCD's) extracted from the original scan using a simple thresholding algorithm.

### B. Matching

Fig. 4(d) shows the result of matching the predicted RCD's in Fig. 4(b) with the observed RCD's in Fig. 4(c) using (8). The figure shows the validation regions and observed RCD's for the seven matches found.




Fig. 4. Localization with a single rotating sonar. (a) Original scan displayed with reference to the *a priori* location estimate. (b) Predicted RCD's, with validation gates for each prediction. (c) Extracted RCD's. (d) Matched predictions and observations. (e) Observed RCD's displayed with reference to the updated position. (f) Original scan displayed with reference to the updated position. The triangle and rectangle show the estimated and true vehicle positions, respectively.

### C. Estimation

Using these seven matches, the vehicle's *a posteriori* position estimate  $\hat{x}(k+1|k+1)$  and associated variance  $P(k+1|k+1)$  are computed using (10) and (11). Fig. 4(e) and (f) shows the extracted RCD's and the original scan displayed with respect to the updated position. This approach to localization with sonar scans, based on matching observed RCD's to RCD's predicted from the map, should be contrasted to the approaches of Crowley [5] and Drumheller [8], which extract line segments from sonar scans for matching with the model. We have found that matching line segments is unreliable because the features inherent to sonar data are RCD's—circular arcs in Cartesian coordinates—not straight-line segments.

## IV. LOCALIZATION USING SIX FIXED SONARS

In this section we present an implementation of model-based navigation using the "Robuter" mobile vehicle. This vehicle uses six fixed sonars in a "ring" formation to provide range information. Fig. 5 shows a complete run with localization being performed repeatedly as the vehicle moves at an approximate speed of 30 cm/s. Fig. 6 shows various stages in this run.

### A. Initialization

The vehicle starts from a known location in the corner of the room as shown in Fig. 6(a). In the figure, the triangle represents *a*

*priori* location and the rectangle represents the *a posteriori* estimate. The initial covariance  $P(0|0)$  was set to zero. The environment is approximately 12 m in length. The line segment model was obtained by measuring the room by hand.

### B. Motion Prediction

The vehicle moves continuously, following a sequence of straight-line motions and on-the-spot rotations supplied by a trajectory planner. The plant noise matrix  $Q(k)$  was chosen to be diagonal, with the values  $\sigma_x = \sigma_y = 1.5$  cm and  $\sigma_\theta = 1^\circ$ .<sup>3</sup>

### C. Data Acquisition and Measurement Prediction

Each time a set of six range readings is acquired, a reading from the vehicle's odometry system is made. The odometry measurement is used to provide a prediction [using (3)] of the vehicle's location at which these sonar range readings were acquired. At this predicted location, predicted range measurements are generated to the nearest wall to each transducer using the planar target model described in the Appendix. We assume a maximum value of  $\beta$  for all planar targets of  $26^\circ$ . If the vehicle orientation estimate is not within  $13^\circ$  of being perpendicular to the target, a NULL prediction that cannot be matched is generated.

### D. Matching

These predictions are matched to the observed range values using (8). The innovation variance  $S(k+1)$  is computed using  $\sigma_r = 1$  cm for the standard deviation of range measurement error. A value of 1 was used for  $g$ , the validation gate "number of sigmas," in (8). Fig. 6(b) shows the vehicle at a location in which only two of these six range readings have been matched to their predictions. In this case only these two matches are passed on to the vehicle position update step. Likewise, Fig. 6(c) and (d) shows five and three matched range readings, respectively. Approximately 30% of the 798 range measurements taken during the run shown were matched for an average of just under two matches out of six readings per time step.

### E. Vehicle Position Update

Matched range readings are then used to update the predicted vehicle location using (10). The resulting location estimate and its confidence depends on which range readings have been matched. For example, Fig. 6(b) shows that two matched range readings have reduced the size of the error ellipse in a direction perpendicular to the readings themselves, while Fig. 6(c) shows that the five matched range readings have reduced uncertainty in all degrees of freedom.

The localization system was written in C and run on a Sun-3 workstation, communicating to a vehicle controller running in LISP on a Sun-4, which in turn supplied updated position estimates to the vehicle's on-board 68000 microprocessor only at the completion of each straight-line motion segment. The filter achieved a cycle time of approximately 1 s. The most computationally expensive part of the algorithm as implemented is the time spent generating predicted observations. This time can be significantly reduced by presorting the map of beacon locations, but this facility was not in use when the run shown here was carried out.

Sonar measurements can only be matched if the sonar transducer was estimated to be nearly perpendicular to a planar target when the measurement was obtained (within  $13^\circ$ ). Note that, in Fig. 5, during the second leg of the vehicle's clockwise journey around the room, the error ellipse is quite large in comparison with the rest of the run.

<sup>3</sup>These values were chosen by an empirical evaluation of the localization capabilities of the vehicle's built-in odometry-based controller.




Fig. 5. A complete run of the SKIDS navigation system. The vehicle moves at a speed of approximately 30 cm per second for each straight line component of this motion sequence, matching 30 percent of all sonar measurements taken during the run.




Fig. 6. A localization sequence for the skids vehicle. The triangle is the *a priori* position estimate, the rectangle is the *a posteriori* position estimate, and the ellipse shows the confidence in the estimates of  $x(k)$  and  $y(k)$ . (a) Initial position, with range measurements from all six sensors displayed. (b), (c), and (d) Various stages in the run, only displaying the current *validated* sonar range values that have been used to update position.

During this part of the run, no wall beacons are visible because the vehicle's estimated orientation is about  $110^\circ$ . As a result, no predicted measurements are generated, no matches obtained and no motion update performed. The vehicle recovers when it subsequently rotates to be nearly perpendicular to the walls of the room, as shown in Fig. 6(c).

Hence, with the configuration of the six fixed sensors on this vehicle (one sonar facing forward and backward, two facing left and right), the system is restricted to follow paths that are nearly perpendicular to the walls of the room. To overcome this limitation, for our current research the vehicle is being fitted with servo-mounted sonars of the type used in the first implementation to endow the vehicle with the ability to track corner, cylindrical, and planar beacons in its environment for precise localization while executing an arbitrary trajectory. Despite this restriction, the system demonstrates that accurate localization is achievable with sonar, provided one has a good sensor model. For walls and corners, sonar provides precise distance measurements, but other times misleading

multiple reflections; the key to using sonar is knowing when it is telling the truth. This system achieves this by only attempting to use updates from planar targets and using tight validation gates to reject outliers.

## V. DISCUSSION AND SUMMARY

We have developed a mobile robot localization system that integrates a variety of beacon observations for input to a Kalman filter to maintain a robust vehicle location estimate. We have presented the application of this system on two different mobile robots that employ sonar as the principle navigation sensor.

The two implementations reveal the tradeoffs between the use of servo-mounted sonars versus a ring of fixed sonars. Stopping a vehicle to obtain a densely sampled scan from a single location takes a great deal of time. A fixed ring provides a spread of measurements over a range of orientations very quickly, permitting dynamic "on-the-fly" position estimation. However, a single unsupported

range measurement obtained by an isolated sonar in a ring has no *local support* to aid in its interpretation. Single unsupported range measurements have a large angular uncertainty of about  $25^\circ$ , but this varies widely depending on the target being observed. This angular uncertainty is *uniformly distributed*. For this reason, we do not use the orientation of the sensor directly in the EKF equations—updates of the vehicle's orientation come implicitly through range updates to different sensors around the vehicle's perimeter. For strong targets the range errors are small enough (less than 1 cm) to be adequately modeled as Gaussian.

In contrast, densely sampled sonar scans obtained from a stationary location provide the potential for bottom-up data interpretation without an *a priori* model [17]. Angularly adjacent range measurements of nearly the same range can be combined to constrain the ambiguity in the true bearing to the target. Rotating the sensor to acquire a densely sampled scan takes time, but the *local support* inherent in a densely sampled scan makes data interpretation easier.

There are a number of limitations to the current version of this algorithm. The most obvious is the restriction that the environment needs to be known *a priori*; we are investigating the problem of navigating in the absence of an initial model in our current research [17], [18]. However, our extensions to handle unknown and changing environments firmly rest on the competence of localization provided by the algorithm presented here. Although these implementations have made exclusive use of sonar, the principles apply equally well to other sensors, provided that a suitable sensor model has been developed. In summary, we reiterate our initial comments that navigation for autonomous mobile vehicles remains an important unsolved problem for which a solution must be found before many potential applications can be realized.

#### APPENDIX

##### EXTRACTING BEACON INFORMATION FROM SONAR DATA

The key to interpreting sensor information is to have a good model of sensor behavior [10]. Indeed, a perfect sensor model would be able to exactly *predict* what sensor data can be observed from any given configuration of objects. The sensor model is used by our algorithm in two crucial ways: 1) *predicting* beacon observations based on the *a priori* vehicle position estimate and 2) *extracting* beacon observations from sensor data. Our model of sonar follows from the work of Kuc and Siegel [16]. They describe a physically based model of sonar that considers the responses of corners, walls, and edges in a specular environment. One key conclusion from their work is that corners and walls produce responses that can not be distinguished from a single scan. The responses from these specular targets take the form of a sequence of headings over which the range value measured is *very accurate* (typically within 1 cm.) Fig. 7 shows a typical, densely sampled (612 range readings), sonar scan obtained in an uncluttered office scene. With this high sampling density, one can see that the scan is composed of sequences of headings at which the range value measured is essentially constant. We refer to such sequences as regions of constant depth (RCD's).

The applicability of this model might seem limited by the fact that most environments present a complex mixture of both diffuse and specular targets. However, the vast amount of sonar data that we have taken in our research has led us to conclude that almost all measurements obtained with the off-the-shelf unmodified Polaroid ranging system in typical scenes (e.g., offices and corridors) are in fact the result of specular reflections. Fig. 7 shows that, in a typical indoor scene, many "false" range readings are produced by the system when the beam is oriented at high angles of incidence to




Fig. 7. A typical sonar scan.

planar targets. At these high angles of incidence, the sound energy emitted from the side-lobes of the beam that strikes the wall perpendicularly is not of sufficient strength to exceed the threshold of the receiving circuit. As a result, the first echo detected by the system is a *multiple reflection* by some other part of the beam reflecting specularly off the wall to some other target and then back. These multiple reflections have been observed by many other researchers, and have commonly been referred to in the literature as "specularities" [6]. We feel this term is misleading because Kuc's model shows that most accurate range measurements produced by planes and corners are in fact due to specular reflections.

To prevent this confusion, we define the *order* of a range measurement as the number of surfaces from which the sound has reflected before returning to the transducer. Orienting the transducer perpendicular to a planar surface such as a wall produces a first-order range reading. Corners produce second-order range readings because the sound has reflected specularly off two surfaces before returning back to the transducer. Multiple reflections will produce third-order and higher range readings. A crucial task in interpretation is to eliminate these higher order reflections which, if naively taken to be the distance to the nearest object, yield false range readings.

Another conclusion Kuc and Siegel reach is that (convex) edges give rise to diffuse echoes that will be weaker in intensity than reflections from walls or corners. In a recent paper Kuc [15] shows that the primary task for obstacle avoidance is to "look" for edges since these are the most difficult to observe. This implies that, for the purpose of localization, edges are less useful as beacons. To incorporate diffuse edges in our terminology, we refer to diffuse reflections from edges as zeroth-order range readings.

We define  $\beta$  as the *visibility angle* of a given target, corresponding to the angles over which the RCD is observed. Fig. 8 shows the result of extracting RCD's of width  $\beta \geq 10^\circ$  from the scan in Fig. 7 using a simple thresholding algorithm, superimposed on a line segment model of the room. Here we can see RCD's corresponding to walls, corners, edges, and higher order targets. To use first- and second-order RCD's for navigation, we need to define which geometric targets in the environment produce them. First-order RCD's principally arise from *planes* and *cylinders*. Most second-order RCD's arise from *corners*. Kuc and Siegel show that corners and walls will appear the same in a scan from a given location [16]. A cylinder is a first-order target that appears similar to a plane or corner from a single location. Higher order multiple reflections also




Fig. 8. Regions of constant depth (RCD's) of width  $\beta \geq 10^\circ$  extracted from this sonar scan, superimposed on a hand-measured model of the room. We can see four first-order RCD's from the four walls of the room, three second-order RCD's from corners, and a single fourth-order RCD resulting from a multiple reflection off the top wall into the lower right-hand corner of the room. There is a single zeroth-order RCD resulting from a weaker reflection from the edge in the lower right-hand region of the room.

produce RCD's that if considered in isolation from a single location could be interpreted as any of these targets.

Qualitatively, the shape of the sonar beam pattern dictates that the width of an RCD is determined by a target's ability to reflect acoustic energy—the stronger the target, the wider the RCD. However, computing in general the precise visibility angle  $\beta_i$  of target  $p_i$  is a complicated process involving a multitude of factors, including the complex pattern of the beam, limitations in the ranging system hardware, and occlusion by other targets. We are addressing these issues in our current research, but for our purposes here, the important conclusion is that planes, cylinders, and corners are the strongest targets, and hence we can obtain precise range measurements to these targets over a wide range of viewing angles using the standard Polaroid ranging system.

#### ACKNOWLEDGMENT

The implementations presented here used software developed in cooperation with C. Brown [4]. The assistance of J. M. Valade of Matra MS2I is gratefully acknowledged.

#### REFERENCES

- [1] N. Ayache and O. Faugeras, "Maintaining representations of the environment of a mobile robot," *IEEE Trans. Robotics Automat.*, vol. 5, no. 6, pp. 804–819, Dec. 1989.
- [2] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. New York: Academic, 1988.
- [3] M. Brady *et al.*, "Progress towards a system that can acquire pallets and clean warehouses," in *4th International Symposium on Robotics Research*. Cambridge, MA: MIT Press, 1987.
- [4] C. Brown, H. Durrant-Whyte, J. Leonard, B. Rao, and B. Steer, "Kalman filter algorithms, applications, and utilities," Oxford U. Robotics Res. Group, Tech. Rep. OUEL-1765/89, 1989.
- [5] J. L. Crowley, "Navigation for an intelligent mobile robot," *IEEE J. Robotics Automat.*, vol. RA-1, no. 1, pp. 31–41, Mar. 1985.
- [6] —, "World modeling and position estimation for a mobile robot using ultrasonic ranging," in *Proc. IEEE Int. Conf. Robotics Automat.*, 1989, pp. 674–681.
- [7] D. D. Dickmanns, "4D-dynamic scene analysis with integral spatio-

- temporal models," in *4th International Symposium on Robotics Research*. Cambridge, MA: MIT Press, 1987, pp. 311–318.
- [8] M. Drumheller, "Mobile robot localization using sonar," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-9, no. 2, pp. 325–332, Mar. 1987.
- [9] H. F. Durrant-Whyte, *Integration, Coordination and Control of Multi-sensor Robot Systems*. New York: Kluwer 1987.
- [10] —, "Sensor models and multi-sensor integration," *Int. J. Robotics Res.*, vol. 7, no. 6, pp. 97–113, 1988.
- [11] —, "Uncertain geometry in robotics," *IEEE J. Robotics Automat.*, vol. RA-4, no. 1, pp. 23–31, Feb. 1988.
- [12] A. C. Gelb, *Applied Optimal Estimation*. Cambridge, MA: MIT Press, 1973.
- [13] J. Hallam, "Intelligent automatic interpretation of active marine sonar," Ph.D. dissertation, Univ. of Edinburgh, Edinburgh, U.K., 1984.
- [14] D. Kriegman, E. Triendl, and T. Binford, "Stereo vision and navigation in buildings for mobile robots," *IEEE Trans. Robotics Automat.*, vol. 5, no. 6, Dec. 1989.
- [15] R. Kuc and B. Barshan, "Navigating vehicles through an unstructured environment with sonar," in *Proc. IEEE Int. Conf. Robotics Automat.*, May 1989, pp. 1422–1426.
- [16] R. Kuc and M. W. Siegel, "Physically based simulation model for acoustic sensor robot navigation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-9, no. 6, pp. 766–778, Nov. 1987.
- [17] J. J. Leonard, I. J. Cox, and H. F. Durrant-Whyte, "Dynamic map building for an autonomous mobile robot," in *Proc. IEEE Int. Workshop Intelligent Robots Syst.*, 1990.
- [18] J. J. Leonard and H. F. Durrant-Whyte, "Application of multi-target tracking to sonar based mobile robot navigation," in *Proc. 29th IEEE Int. Conf. Decision Control*, 1990.
- [19] L. Matthies and S. Shafer, "Error modeling in stereo navigation," *IEEE J. Robotics Automat.*, vol. RA-3, no. 3, pp. 239–248, June 1987.
- [20] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Autonomous Robot Vehicles*, I. Cox and G. Wilfong, Eds. New York: Springer-Verlag, 1990.
- [21] H. W. Sorenson, Ed., *IEEE Trans. Automatic Control (Special Issue on Applications of Kalman Filtering)*, vol. AC-28, no. 3, 1983.
- [22] D. J. Torrieri, "Statistical theory of passive location systems," in *Autonomous Robot Vehicles*, I. Cox and G. Wilfong, Eds. New York: Springer-Verlag, 1990.
- [23] D. Willner, C. B. Chang, and K. P. Dunn, "Kalman filter algorithms for a multi-sensor system," in *Proc. IEEE Int. Conf. Decision and Control*, 1976, pp. 570–574.

### Learning Impedance Parameters for Robot Control Using an Associative Search Network

Moshe Cohen and Tamar Flash

**Abstract**—This work presents an evaluation of the associative search network (ASN) learning scheme when used for learning control parameters for robot motion. The control method used is impedance control in which the controlled variables are the dynamic relations between the motion variables of the robot manipulator's tip and the forces exerted by the tip. The main task used is that of wiping a surface whose

Manuscript received January 11, 1989; revised December 1, 1989. Portions of this work were presented at the 11th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Seattle, WA, November 9–12, 1989.

The authors are with the Department of Applied Mathematics and Computer Science, The Weizmann Institute of Science, Rehovot 76100, Israel. IEEE Log Number 9042229.