

Environment Exploration Using “Just-in-Time” Sensor Fusion

Ioannis M. Rekleitis¹ Gregory Dudek¹
Paul Freedman²

¹Centre for Intelligent Machines, McGill University,
3480 University St., Montreal, Québec, Canada H3A 2A7

² Centre de recherche informatique de Montréal
1801 McGill College Avenue, Suite 800, Montréal Québec H3N 2N4

Abstract

This paper describes an approach to combining range data from both a set of sonar sensors as well as from a directional laser range finder to efficiently take advantage of the characteristics of both types of devices when exploring and mapping unknown worlds. We call our approach “just in time sensing” because it uses the more accurate but constrained laser range sensor only as needed, based upon a preliminary interpretation of sonar data. In this respect, it resembles “just in time” inventory control which attempts to judiciously obtain materials for industrial manufacturing only when and as needed. Experiments with a mobile robot equipped with sonar and a laser rangefinder demonstrate that by judiciously using the more accurate but more complex laser rangefinder to deal with the well-known ambiguity which arises in sonar data, we are able to obtain a much better map of an interior space at little additional cost (in terms of time and computational expense).¹

1 Introduction

We are interested in sensor-based environment exploration with a mobile robot. Most approaches to sensing involve tradeoffs between speed, time and accuracy. We propose an approach we call “just in time sensing” to combine data from both two sensor modalities: sonar and active stereo (BIRIS) to efficiently take advantage of the characteristics of both types of devices when exploring and mapping unknown worlds. We call our approach “just in time

sensing” because it uses sonar data for a coarse preliminary mapping phase followed by a judicious application of laser range sensing only where needed. In this respect, it resembles “just in time” inventory control which attempts to judiciously obtain materials for industrial manufacturing only when and as needed.

The key to the approach is that one sensor provide a large-scale but (in some sense) low resolution depiction of the environment while a second sensor provides a more costly but higher resolution view. In principle, the approach is applicable to various combinations of sensing methodologies. In distinction to standard approaches to fusion or multi-scale analysis, the high-resolution data is *acquired* only where it is deemed appropriate. As such, this can also be portrayed as a form of active sensing.

Sonar sensing is ubiquitous on mobile robots due to its low cost, the simplicity of the required processing, and the rapidity with which it can return results reflecting range measurements over a large region of space. Despite these advantages however, it suffers from several shortcomings as a source of range data. Most importantly, the measurements tend to have low spatial resolution and the observed data (with most interpretation strategies) are typically confounded by the effects of multi-bounce specular reflections. For this reason, the use of sonar is often confined to collision avoidance rather than mapping.

Laser range sensors, on the other hand, are typically able to obtain comparatively accurate data with fewer artifacts when compared to simple sonar sensors. The laser rangefinder we have developed at McGill University’s Centre for Intelligent Machines

¹Appeared in “Vision Interface”, pages 175–82, Toronto, May 1996.

is based on BIRIS [4] technology: a special lens with two pinholes near the nodal point is used. This produces a double image of objects in the scene with a disparity that depends directly on the distance of the object from the focal plane of the lens. In short, we obtain a stereo image using only a single CCD array. By projecting a laser stripe onto the scene, a target is made available which can be used for unambiguous stereo correspondence, allowing true depth to be readily computed. Two of these BIRIS sensors mounted on pan tilt units in a specific configuration comprise the McGill QUADRIS sensor platform.

Like all sensing technologies, the BIRIS sensor also has disadvantages. Its primary shortcoming is that although the accuracy of range measurements is reasonable over short distances (up to one or two meters), accuracy degrades rapidly with longer distances (this is one of the design parameters). Furthermore, obtaining either a dense range image or range data over a wider field of view than the 25 degrees or so covered by a typical camera² lens implies physically sweeping the camera and laser across the scene, and hence involves a time delay.

Early work to validate our “just-in-time” sensing strategy using real sonar data and simulated BIRIS data was reported earlier [10]. In this paper, we describe results using real sonar data and real BIRIS data obtained from the McGill QUADRIS platform.

1.1 Background

Various approaches have been considered for the exploration of unknown or partially-known environments. Most work dealing with real noisy sensor data must cope with the management of sensor uncertainty and exploration strategies, i.e. selecting successive unknown locations or objects to visit (investigate) e.g. [22, 20, 23, 2, 1, 7, 3]. The best way to manage exploration in the face of real sensor noise remains an open problem. The question is further complicated by the fact that the choice of an optimal mapping strategy is sensitive to the specific task at hand. A somewhat distinct research stream deals with the complexity issues in autonomous robot exploration of an unknown environment [5, 16, 21, 17].

In general, work on sensor fusion has tended to focus on issues of how best to combine measurements from different sensors e.g. [24, 6], or how best to extract data with a single sensor and fuse the measurements over time e.g. [11, 12, 25], rather than how to selectively extract measurements from different *types* of sensor. This later problem of combining measurements over time has, in fact, two variants: that prob-

²Using cameras with a very wide field of view leads to distortions and a loss of resolution and thus is not an acceptable solution.

lem of fusing a set of measurements obtained somehow over some time period e.g. [19, 9], and the problem of efficiently selecting where or when to obtain additional measurements e.g. [27].

The approach described here is based on the construction of a geometric map using sonar data and BIRIS data. Several approaches to the construction of maps from such range data collected by a mobile robot [8, 13, 14, 18, 15]. Most methods are plagued by the noisy or otherwise limited natures of sonar range sensing, or involve very costly or time consuming data acquisition. Our approach uses a combination of robust fitting and a split-and-merge segmentation of accomplish line segment generation and is documented in detail elsewhere [18].

An apparently general finding is that the difficulty of the exploration task is rather sensitive to the level of sensor noise and the fidelity of the geometric inferences made about the objects in the environment. Thus, it seems appropriate to focus some effort on efficiently obtaining good geometric models since using all of the sensors on the robot all of the time can lead to serious inefficiencies in exploration/mapping time and in using the on-board computational resources.

2 Just in time sensing

2.1 The Exploration Context: An Art Gallery

As an experimental testbed for our sensing methodology, we are examining the performance of our approach in an indoor environment resembling an art gallery or museum; such an environment has certain practical advantages as specified below. The environment consists of a large enclosed room containing several stationary objects. We make the assumption that a collision-free paths exists between any two points within the free space of the room. Although, the objects inside a room are assumed to be static, the configuration of the room is subject to change between successive visits by the robot and must therefore be mapped again every time it is visited. The floors are flat (except for some well defined places where stairs exist) and smooth without any anomalies.

Our robot, an RWI B-12, is equipped with a ring of sonar sensors which provide coarse range measurements of distances up to 8 meters omnidirectionally, and a BIRIS sensor system that can give accurate estimates over at short range and only in one direction; see Figure 1.

2.2 Processing sonar data

As previously mentioned, the main robot sensor is the sonar range finder. Our RWI robot is equipped




Figure 1: RWI mobile robot equipped with McGill QUADRIS sensor platform.

with a ring of 12 sonars positioned equidistantly around it, returning 12 measurements at the same time. In order to obtain a denser range map which is less susceptible to artifacts, the robot is commanded to rotate in place and more measurements are made. More precisely, we collect a total of 180 sonar measurements at each robot (x,y) position but not all of them are used.

But as described above, multiple bounces or echos associated with sonar data may give rise to the following errors. Firstly, a phantom “third” wall may appear where two walls meet at a corner. Secondly, false measurements may be obtained suggesting the presence of objects far away.

In order to deal with such errors in the sonar data, certain measurements are suppressed from further processing. Firstly, measurements (range estimates) beyond a certain “logical” distance are discarded, under the assumption that they represent multiple bounces. Although this could result in the loss of legitimate information, we assume that such information is sufficiently distant from the robot as to be of little importance in mapping that part of the environment which is nearby.

In practice, the “discard” threshold must be established so as to maintain measurements which are close enough to the robot and therefore sufficiently accurate (almost no data from multiple bounces) since otherwise, the range of the sensor is wasted and in larger environments, this could result into some exploration/mapping inefficiency.

A second criterion for discarding erroneous sonar data is related to knowledge about the immediate surroundings of the robot. Our robot is approximately 30cm in diameter and we shall therefore assume that after a complete sonar scan, if no object “appears” in a distance less than half a meter from

the robot, then the circle around the robot with a diameter of about one meter is empty. Consequently, any range measurement within this circle can be assumed to be in error, due to multiple bounces, and therefore is discarded.

Once those measurements deemed to be in error are discarded, the sonar data measurements are then clustered together using the “Sphere of Influence Graph” [26]. assuming a minimum distance in order to avoid clustering points at the same position as separate small clusters. Line segments are then fitted to the data clusters by splitting and merging them using a certain confidence measure to obtain new line segments with a certain confidence [MD94]. From that point on, the construction of the environmental map is performed in terms of lines with a confidence measure attached to them.

2.3 Processing laser range data

In order to improve the results we obtain from just sonar data, the BIRIS laser range finder is judiciously called into play, according to our “just-in-time” sensing strategy. As previously described, the useful range of BIRIS is about three to four meters but in order to have very good accuracy, most of the scans are performed at one and a half meters range. Another limitation is its field of view of 25 degrees, which is, in practice, even smaller due to artifacts appearing at the edges of the scan lines. The BIRIS sensor may be used to obtain up to 512 range measurements per scan whenever something is present (an object or a series of objects) in the whole field of view.

Once again, the “Sphere of Influence Graph” is used to cluster measurement data into clusters but this time, no minimum distance threshold is used, since the data points are already closer together. When line segments are fitted to the clusters, we have observed almost ideal confidence levels.

Note too that when BIRIS is used in addition to sonar, a third criterion for discarding sonar data in error may be defined as follows. Whenever a line segment is observed by BIRIS, we assume that no object is present inside the triangle formed by the BIRIS line segment and the position of the robot at the apex. Moreover, as we shall see, BIRIS is only used in places where sonar is ambiguous, i.e. where much sonar data is error.

2.4 Just-in-time BIRIS sensing

We now describe the way in which the mobile robot explores and maps its unknown environment with just-in-time BIRIS sensing to complement sonar sensing. Simply, BIRIS data is only acquired to ac-

curately pinpoint the corners and the borders of objects, where the sonar data is ambiguous.

The exploration strategy developed as a testbed for the fusion of sonar and laser data is a kind of wall following, or more accurately, “closest object following”. The world is modelled in terms of line segments with marked endpoints; an endpoint is marked as either “terminal” (confirmed) or “non terminal” (to be confirmed). The algorithm of the exploration strategy is outlined next.

The robot explores its unknown environment by proceeding to the nearest object (or wall). (When multiple objects are present in the same sonar scan, the robot applies a modified breadth first search algorithm in order to structure the way in which it will explore one object after another.) The robot then navigates around the object while increasing the length of the current line segment until reaching a non-terminal endpoint of the current line segment in a trajectory parallel to it at a distance of approximately one meter from it (making, if necessary, the proper adjustments). It is at this stage that BIRIS is used.

Non-terminal endpoints will arise when the robot is approaching the corner of two walls whereby the current object (the wall that the robot was following) comes to an end. The robot calculates the potential corner using sonar data, aims BIRIS at the corner, and then maps the corner accurately, marking the endpoints of the two lines as terminals, i.e. the end of the wall that the robot was following, and the beginning of the other wall.

The other situation arises when the robot reaches the physical extremity of a line/wall. Here BIRIS is needed to accurately map the endpoint position since the sonar cone identifies any object inside it as a point without precision. Once again, the robot aims BIRIS at the endpoint of the object currently marked as “non-terminal”, obtains data, transforms the data into line segments, and then calculates the exact position of the endpoint while marking it as “terminal”. When both endpoints of the closest object line segment are marked as terminal, then the robot moves away from until a new object, not fully explored (with at least one non-terminal endpoint) becomes the new closest object.

2.5 Constructing the map from sonar and BIRIS data

The incremental construction of the map of the unknown world takes into account the position and orientation of the line segments, their data source (sonar, BIRIS), and the confidence measure attached to them.

In particular, the map consists of three types of lines: *Sonar*, for the line segments created using only sonar points; *BIRIS*, for the ones that use only BIRIS measurements; *Complex*, for the ones that appear after merging Sonar and BIRIS line segments.

If two lines are almost parallel and their separation is less than a critical threshold d (in our experiments, $d = 5cm$), then they are merged. There are two different cases for merging.

The first case is when there is one *Sonar* line segment and one *BIRIS* or *Complex* line segment. Since BIRIS is much more accurate than sonar, we project the *Sonar* line segment onto the *BIRIS/Complex* line segment. In all other cases, we merge two line segments by taking into account the confidence measure (f_i) associated with each one, and its length (l_i). If the two line segments intersect, then a new line segment passing through the point of intersection is created with slope m as given by the equation 1 where m_i is the slope of the line segment with the largest confidence measure (length multiplied by confidence). If the two line segments don't intersect, then we identify their two nearest endpoints and calculate a weighted midpoint using the same weights used for the slope, thereby creating a new line segment passing through the weighted midpoint with a slope m as described above.

$$m = \alpha * m_i + (1 - \alpha)m_j \tag{1}$$

$$\alpha = \max\left(\frac{l_1}{l_1+l_2} * f_1, \frac{l_2}{l_1+l_2} * f_2\right)$$

One difficulty arises when merging changes the position of line segment endpoints. Consider the case of merging two line segments with one having an endpoint marked terminal. It is possible that when we map a *BIRIS* or *Complex* line having a terminal endpoint onto a *sonar* line segment, the length of the sonar line segment will change, in order to preserve the position of the terminal endpoint. In this way, BIRIS data helps “clean up” the *sonar* line segments by discarding erroneous points *and* by adjusting the lengths of the *sonar* line segments.

Each time the robot uses its sonar sensors or BIRIS rangefinder, it updates its map with new *sonar* or *BIRIS* line segments and re-evaluates its agenda regarding how to continue the exploration of the environment. As the partial map evolves it can, of course, be used to navigate more efficiently or to refine the robots position estimate. The entire just-in-time exploration algorithm is summarized below:

- Obtain a sonar scan
- Transform Sonar Points into Lines
- Cluster with Sphere of Influence Graph

```

-Fit Lines
-Measure confidence
-Update Map with New Lines
  -Merge Parallel Lines
-Find Closest Line
-If Closest_Line == Line_Followed AND
  Not after the EndPoint          Then
  -Check EndPoints
  -If Both terminals then
    -Move away from Closest_Line;
  -Else
    -Find Closest Non-Terminal EndPoint
    -Move towards it, in a trajectory
      parallel to the Line keeping one
      meter distance
-Else If Closest_Line == Line_Followed
  AND After the EndPoint          Then
  -Point BIRIS at the EndPoint
  -Obtain BIRIS data,
    Transform them into Lines
  -Decide if EndPoint is Terminal
  -Mark the EndPoint as Terminal

-Else /*Potential Intersection */
  -Calculate if Closest_Line and
    Line_Followed could intersect
  -Point BIRIS at the Intersection
  -Obtain BIRIS data,
    Transform them into Lines
  -If an intersection exist
    -Mark the EndPoints as Terminals
  -Update Map for BIRIS
  -Follow the New Closest Line
-Goto to the beginning

```

3 Experimental results

A series of experiments were carried out to evaluate the performance of the algorithm both in simulation and using real data. The simulation results use an accurate sonar simulator that replicates multipath echos, diffraction and other characteristics of sonar sensing. The results, depicted in Figure 2. The path of the robot is shown in green, the walls are shown in red, and BIRIS scans are illustrated by dark-red lines (which emanate from points on the green path). The results illustrate the effectiveness of the algorithm in generating a map free from traditional anomalies associated with sonar mapping. The BIRIS scanner is used only at a few positions, leading to a comparatively rapid traversal.

To verify these results in practice, a series of experiments were conducted in the Mobile Robotics Lab of McGill's Centre for Intelligent Machines using an RWI B-12 robot equipped with the QUADRIS

sensing platform (which contains a BIRIS sensor). Computations were carried out on a Silicon Graphics INDY workstation. Mapping was carried out using sonar data alone to establish a baseline for comparison; see Figure 3. Note that the robot does not complete a full tour of the indoor space; just three walls are followed. Note the substantial number of artifacts, spurious "walls" and incorrect corner geometries.




Figure 3: Map of our interior space generated just using sonar sensors.

The same environment, when mapped using sonar and BIRIS in the context of our "just-in-time" algorithm is mapped much more effectively. The following four figures are a sequence of "snapshots" during different stages of the map construction of the same world, using "just-in-time sensing" with BIRIS (Figures 3 through fig:comp8). In Figure 3, we observe that the robot started its exploration at a slightly different position and has now followed two walls; the portions in the figure shaded in grey indicate where BIRIS data was used to more accurately map corners.

In Figure 3, partial information about the interior object and about a third wall has been added to the map. Note too that the artifact present in the lower left of the previous figure has been "cleaned-up" as part of the accurate mapping of the third wall.

In Figure 3, we observe that BIRIS is used once again to more accurately map the corner associated with the end of the third wall and the beginning of the fourth wall.

Finally, in Figure 3, the robot has completed the same partial tour as before (using just sonar), by following three walls. When compared to the map shown in Figure 3, we observe that the modelling is

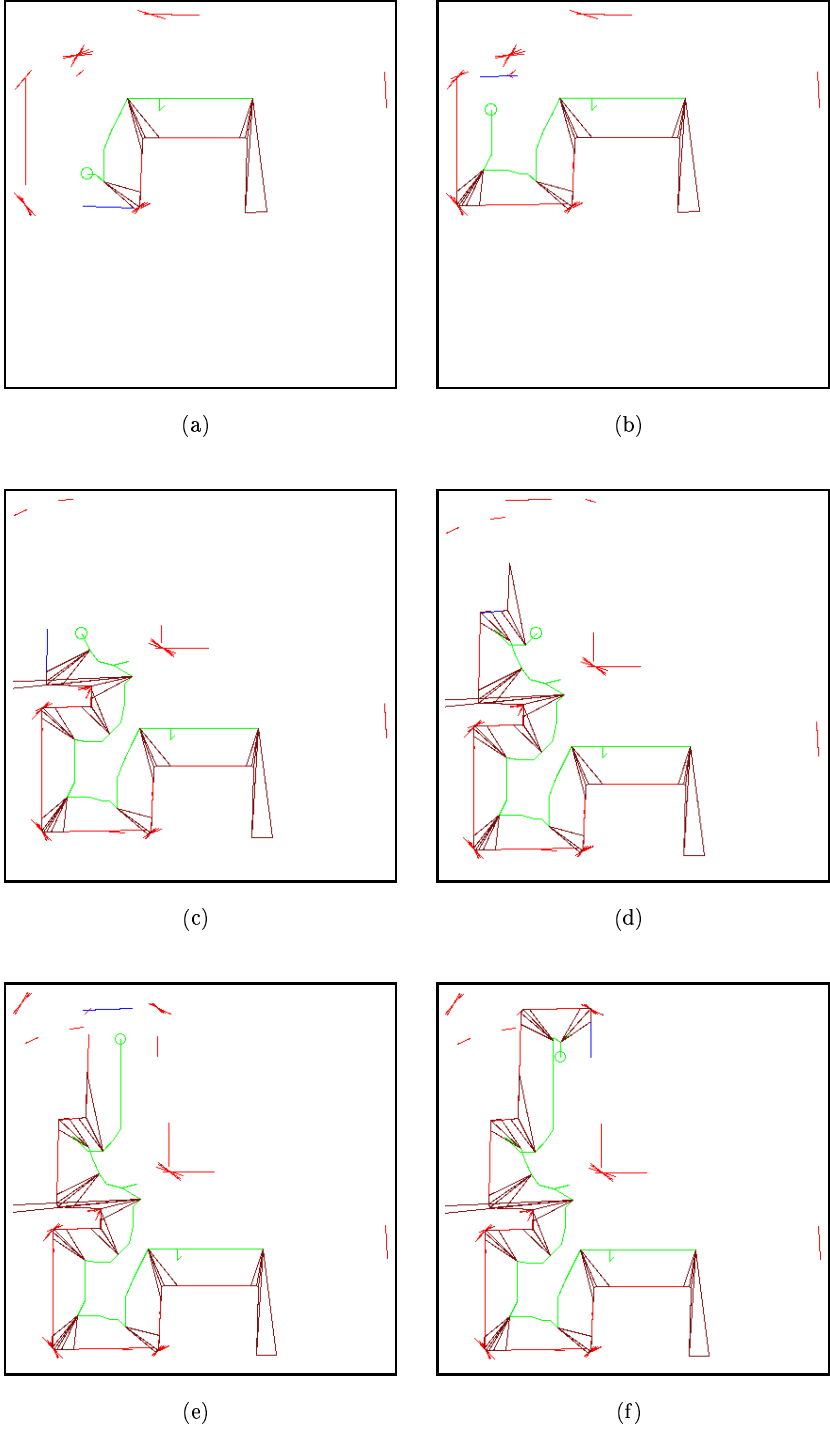


Figure 2: A series of six snap shots of the map constructed by the Robot.




Figure 4: Progressive map construction using just-in-time sensing; Step A




Figure 6: Progressive map construction using just-in-time sensing; Step C




Figure 5: Progressive map construction using just-in-time sensing; Step B




Figure 7: Progressive map construction using just-in-time sensing; Step D

much cleaner, especially with respect to the interior (navigable) space, i.e. artifacts still remain “behind” the walls of the interior space but they have no bearing on the way in which the robot would use its map in the future to navigate within the world now explored. In principle, these spurious walls located in invisible locations could be readily deleted, but such an inference would constitute a step beyond the simple “just-in-time” method we are examining.

4 Conclusions

This paper describes an approach to combining range data from both a set of sonar sensors as well as from a directional laser range finder to efficiently

take advantage of the characteristics of both types of devices when exploring and mapping unknown worlds. We call our approach “just in time sensing” because it uses the more accurate but constrained laser range sensor only as needed, based upon a preliminary interpretation of sonar data. It has the advantage of allowing rapid sensing of the environment with comparatively high accuracy.

Simulations and experiments demonstrate that by judiciously using the more accurate but more complex laser rangefinder to deal with the well-known ambiguity which arises in sonar data, we are able to obtain a much better map of an interior space at little additional cost (in terms of time and computational expense). The algorithm is based on knowledge of how sensor errors manifest themselves as well as how the environment is typically structured. It is

this knowledge that allows the informed selection of locations to probe with the more accurate sensor. In interesting question is how to relate even more abstract domain knowledge to active data acquisition.

References

- [1] R. De Almeida and C. Melin. Exploration of unknown environments by a mobile robot. *Intelligent Autonomous Systems 2*, pages 715–725, 1989.
- [2] Minoru Asada. Map building for a mobile robot from sensory data. *IEEE Transactions on Systems, Man and Cybernetics*, 37(6):1326–1336, November/December 1990.
- [3] Kenneth Basye and Thomas Dean. Map learning with indistinguishable locations. In M. Henrion L. N. Kanal J. F. Lemmer, editor, *Uncertainty in Artificial Intelligence 5*, pages 331–340. Elsevier Science Publishers, 1990.
- [4] F. Blais, M. Rioux, and J. A. Beraldin. Laser range finder development for 3-d vision. In *Proc. Vision Interface '89*, London, Ont., May 1989. CIPRS.
- [5] Avrim Blum, Prabhakar Raghavan, and Baruch Schieber. Navigating in unfamiliar geometric terrain. In *Proc. 23rd ACM Symposium on the Theory of Computing*, pages 494–504. ACM Press, 1991.
- [6] Paul B. Chou and Christopher M. Brown. Probabilistic information fusion for multi-modal image segmentation. In *Proc. International Joint Conference of Artificial Intelligence*, pages 779–782, Milan, Italy, August 1987.
- [7] Gregory Dudek, Paul Freedman, and Souad Hadjes. Using local information in a non-local way for mapping graph-like worlds. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence*, pages 1639–1645, Chambéry, France, August 1993. International Joint Conf. on Artificial Intelligence Inc.
- [8] Gregory Dudek and Paul MacKenzie. Model-based map construction for robot localization. In *Proceedings of Vision Interface 1993*, Toronto, ON, May 1993.
- [9] A. Elfes. Using occupancy grids for mobile robot perception and navigation. *IEEE Computer*, 22(6):46–58, June 1989.
- [10] Ioannis Rekleitis Gregory Dudek Paul Freedman. Mobile robot exploration by using fused data from two sensors. In *PRECARN/IRIS Conference* Vancouver, B.C., June 14–16, 1995.
- [11] Benjamin J. Kuipers and Y. T. Byun. A qualitative approach to robot exploration and map-learning. In *Proceedings of the IEEE workshop on spatial reasoning and multi-sensor fusion*, pages 390–404, Los Altos, CA, 1987. IEEE.
- [12] Simon Lacroix, Philippe Fillatreau, Fawzi Nashashibi, Raja Chatila, and Michael Devy. Perception for autonomous navigation in a natural environment. *Workshop on Comp. Vis. for space applications*, 1993.
- [13] S. Lange, L. Korba, and A. Wong. Characterizing and modelling a sonar ring. In *Proc. SPIE Mobile Robotics IV*, pages 291–304, Philadelphia, 1989.
- [14] J. J. Leonard and H. F. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation*, 7(3):376–382, 1991.
- [15] John J. Leonard, Bradley A. Moran, Ingemar J. Cox, and Matthew L. Miller. Underwater sonar data fusion using an efficient multiple hypothesis algorithm. In *Proc. of IEEE Conference on Robotics and Automation*, pages 2995 – 3002, Nagoya, Japan, May 1995. IEEE Press.
- [16] T. Levitt and D. Lawton. Qualitative navigation for mobile robots. *Artificial Intelligence*, 44(3):305–360, August 1990.
- [17] Vladimir Lumelsky, Snehasis Mukopadhyay, and Kang Sun. Dynamic path planning in sensor-based terrain acquisition. *IEEE Trans. on Robotics and Automation*, 6(4):462–472, 1990.
- [18] Paul MacKenzie and Gregory Dudek. Precise positioning using model-based maps. In *Proceedings of the International Conference on Robotics and Automation*, San Diego, CA, 1994. IEEE Press.
- [19] H. Moravec. Sensor fusion in certainty grids for mobile robots. In A. Casals, editor, *Sensor Devices and Systems for Robotics, NATO ASI Series, Vol. F52*, pages 253–276. Springer-Verlag, Berlin, 1989.
- [20] P. Moutarlier and R. Chatila. Incremental environment modelling by a mobile robot from noisy data. In *Experimental Robotics 1*, pages 327–346. Springer-Verlag, 1990.
- [21] B. Oommen, S. Iyengar, S. Nageswara, S. Rao, and R. Kashyap. Robot navigation in unknown terrains using learned visibility graphs, part i: The disjoint convex obstacle case. *IEEE J. of Robotics and Automation*, 3(6):672–681, 1987.
- [22] M. Penna and J. Wu. Models for map building and navigation. *IEEE Trans. Systems, Man and Cybernetics*, 23(5):1276–1301, Sept 1993.
- [23] Y. Roth-Tabak and T. Weymouth. Environment model for mobile robotics indoor navigation. Technical Report CSE-TR-53-90, The University of Michigan, 1990.
- [24] S. Shafer, A. Stentz, and C. Thorpe. An architecture for sensor fusion in a mobile robot. In *IEEE Int. Conf. on Robotics and Automation*, pages 2002–2011, 1986.
- [25] Jun Tani. Model-based learning for mobile robot navigation from the dynamical systems perspective. *IEEE Trans. System, Man and Cybernetics*, 26(6), 1996.

- [26] Godfried Toussaint. A graph-theoretical primal sketch. *Computational Morphology*, 1988.
- [27] Peter Whaite and Frank P. Ferrie. Uncertain views. *CVPR*, 1992.