

Morphological Surface Profile Extraction from Multiple Sonars

Deniz Başkent and Billur Barshan
Department of Electrical Engineering
Bilkent University
Bilkent, 06533 Ankara, Turkey

Abstract

This paper presents a novel method for surface profile determination using multiple sensors. Our approach is based on morphological processing techniques to fuse the range data from multiple sensor returns in a manner that directly reveals the target surface profile. The method has the intrinsic ability of suppressing spurious readings due to noise, crosstalk, and higher-order reflections, as well as processing multiple reflections informatively. The algorithm is verified both by simulations and experiments in the laboratory by processing real sonar data obtained from a mobile robot. The results are compared to those obtained from a more accurate structured-light system, which is however more complex and expensive.

1 Introduction

Perception of its surroundings is a distinguishing feature of an intelligent mobile robot. An inexpensive, yet efficient and reliable approach to perception is to employ multiple simple sensors coupled with appropriate data processing.

Since sonar sensors are inexpensive devices which yield accurate range data up to several meters, they have been widely used in robotics applications. However, due to their wide beamwidth, the angular resolution of sonar sensors is low, resulting in an uncertainty about the location of the object encountered.

Most of the approaches for map-building with sonar have concentrated on surfaces with fixed or piecewise-constant curvature, mostly composed of target primitives such as planes, corners, edges, and cylinders [1-6]. In [7], an analytical approach to surface curvature extraction is described which employs differential geometry. The method proposed in this paper is completely novel in that morphological processing techniques are applied to sonar data to reconstruct the profile of an arbitrarily curved surface. It is important to emphasize that morphological processing is employed here to process the sonar map being constructed in the robot's memory, rather than conventional camera images.

Approaches based on geometrical or analytical modeling are often limited to elementary target types or simple sensor configurations. On the other hand, our approach is aimed at the determination of arbitrary surface profiles. The method is extremely flexible in that it can equally easily handle arbitrary sensor configurations and orientations as well as synthetic

arrays obtained by moving a small number of sensors. As already mentioned above, a commonly noted disadvantage of sonar sensors is the difficulty associated with interpreting spurious readings, crosstalk, higher-order, and multiple reflections. The method proposed is capable of effectively suppressing spurious readings, crosstalk, and higher-order reflections. Furthermore, it has the intrinsic ability to make use of echo returns beyond the first one (i.e. multiple reflections) so that echoes returning from surface features further away than the nearest can also be processed informatively.

In Section 2, basic principles of sonar sensing are reviewed. The morphological processing and curve-fitting algorithms are introduced and applied in Section 3. After describing the system setup, experimental results are presented and discussed in Section 4.

2 Basics of Sonar Sensing

The ultrasonic sensors used in this work measure time-of-flight (TOF), which is the round-trip travel time of the pulse from the sonar to the object and back to the sonar. Since the speed of ultrasonic waves is known ($c = 343.3$ m/s), the range r can be easily calculated from $r = ct_0/2$. Many ultrasonic transducers operate in this pulse-echo mode. The transducers act both as receiver and transmitter.

The objects are assumed to reflect the ultrasonic waves specularly. This is a reasonable assumption, since most systems operate below 200 kHz so that the propagating waves have wavelengths well above several millimeters. Details on the objects which are smaller than the wavelength cannot be detected [7]. The sonars used in our experimental setup are Polaroid transducers operating at a resonant frequency $f_0 = 49.4$ kHz [10], which corresponds to a wavelength of $\lambda = c/f_0 = 6.9$ mm at room temperature.

The major limitation of sonar sensors comes from their wide beamwidth. Polaroid transducers have a half beamwidth angle of $\theta_0 = \pm 12.5^\circ$ [10]. Although these devices return accurate range data, they cannot provide direct information on the angular position of the object from which the reflection was obtained. Thus, all that is known is that the reflection point lies on an arc whose radius is determined by $r = ct_0/2$ (Figure 1(a)). More generally, when one sensor transmits and another receives, it is known that the reflection point lies on the arc of an ellipse whose focal points are the transmitting and receiving transducers

(Figure 1(b)). Notice that the reflecting surface is tangent to these arcs at the actual point of reflection. The angular extent of these arcs is determined by the sensitivity regions of the transducers. Most commonly, the wide beamwidth of the transducer is accepted as a device limitation which determines the angular resolving power of the system. In this naive approach, a range reading of r from a transmitting/receiving transducer is taken to imply that an object lies along the line-of-sight of the transducer at the measured range. Consequently, the angular resolution of the surface profile measurement is limited by the rather large beamwidth, which is a major disadvantage. Our approach, as will be seen, turns this disadvantage into an advantage. Instead of restricting oneself to an angular resolution equal to the beamwidth by representing the reflection point as a coarse sample along the line-of-sight, circular or elliptical arcs representing the uncertainty of the object location are drawn. By combining the information inherent in a large number of such arcs, angular resolution far exceeding the beamwidth of the transducer is obtained.

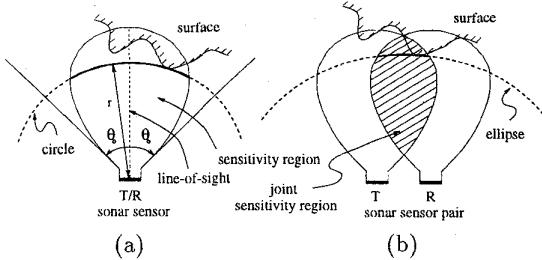


Figure 1: a) For the same sonar transmitting and receiving, the reflecting point is known to be on the circular arc shown. b) The elliptical arc if the wave is transmitted and received by different sensors.

3 Processing of the Sonar Data

Structured sensor configurations such as linear and circular arrays as well as randomly scattered and oriented configurations have been considered.

Figure 2(a) shows a surface, whose profile is to be determined. Figure 2(b) shows the circular and elliptical arcs obtained from a circular array of sensors, which both rotate and translate to increase the number of arcs generated from the available number of sensors.

Notice that although each arc represents considerable uncertainty as to the angular position of the reflection point, nevertheless one can almost visually extract the actual curve shown in Figure 2(a) by examining Figure 2(b). Each arc drawn is expected to be tangent to the surface at least at one point. At these actual reflection point(s), several arcs will intersect with small angles at nearby points on the surface. The many small segments of the arcs superimposed in this manner create the darker features in Figure 2(b), which tend to cover and reveal the actual surface. The remaining parts of the arcs, not actually corresponding

to any reflections and simply representing the angular uncertainty of the sensors, will remain more sparse and isolated.

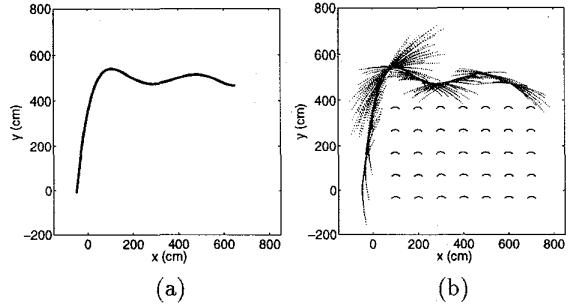


Figure 2: a) The original surface. b) The circular sensor array mounted on a mobile robot moves to 35 different locations and collects data by rotating around its center from 45° to 135° with respect to the positive x axis in 15° steps. The circular array has been shown at the 35 locations at its 90° position.

In the next sections, morphological rules will be employed to achieve what is natural for the human visual perception system: the extraction of Figure 2(a) from 2(b).

3.1 Morphological Processing

In this study, morphological operators are used to eliminate the sparse and isolated segments in the arc map, leaving the mutually reinforcing segments that directly reveal the original surface profile. Erosion, dilation, opening, closing, and thinning are the most widely used morphological operations to accomplish tasks such as edge detection, enhancement, smoothing, and noise removal in image processing [8]. *Erosion* and *dilation* are used to thin or fatten an image respectively. These operations are defined according to a structuring element or template. An example 3×3 template is shown in Figure 3.

1	1	1
1	(1)	1
1	1	1

Figure 3: Template used for both erosion and dilation.

The algorithm for erosion is as follows: The template is shifted over the pixels of the sonar map which take the value 1 one at a time and the template's pixels are compared with those image pixels which overlap with the template [8]. If they are all identical, the central pixel with value 1 will preserve its value; otherwise it is deleted. For the template shown in Figure 3, all eight neighbors of the pixel must be 1 ($n = 8$), and the image is eroded or shrunk accordingly.

The dilation algorithm is very similar to that for erosion, but is used to enlarge the image according to the template. This time, all eight neighbors of those image pixels which originally equal 1 are set equal to 1.

In this study, the structuring element for dilation and erosion is the 3×3 square template, shown in Figure 3 with the central pixel encircled. Since the template is symmetric, the image will be fattened (dilation) or thinned (erosion) in all directions by one pixel.

The direct use of erosion may eliminate too many points and result in loss of information characterizing the surface. For such cases, the compound operations of *opening* and *closing* are considered. Opening consists of erosion followed by dilation and vice versa for closing. Opening helps reduce small extrusions, whereas closing enables one to fill the small holes inside the image. Closing is applied prior to thinning, described below, in cases where the points are not closely connected to each other so that the direct use of thinning may result in the loss of too many points. Filling the gaps using closing first may prevent this from happening.

Thinning is a generalization of erosion with a parameter n varying in the range $1 \leq n \leq 8$. In this case, it is sufficient for any n neighbors of the central image pixel to equal 1 in order for that pixel to preserve its value of 1. The flexibility that comes with this parameter enables one to make more efficient use of the information contained in the arc map.

In *pruning*, which is a special case of thinning, at least one ($n = 1$) of the neighboring pixels must have the value 1 in order for the central pixel to remain equal to 1 after the operation. This operation is used to eliminate isolated points [8]. Thus, pruning and erosion are the two extremes of thinning with $n = 1$ and $n = 8$ respectively.

Since there are many alternatives for morphological processing of sonar data, an error measure is introduced as a success criterion:

$$e = \frac{\sqrt{\frac{1}{N_k} \sum_{i=1}^{N_k} (m_i - y_i)^2}}{\sigma_y} \quad (1)$$

Here, i is the discrete index along the x direction and y_i is the discretized function representing the actual surface with variance $\sigma_y^2 = \frac{1}{N} \sum_{i=1}^N (y_i - \frac{1}{N} \sum_i y_i)^2$. N is the total number of columns whereas N_k represents those columns left with at least one point as a result of some morphological operation. m_i is the vertical position of the median (centermost) point along the i th column of the map matrix (e.g. Figure 4(a)). If there are no points in a particular column, that column is excluded from the summation. If the number of columns thus excluded is large; that is, if the morphological operations have eliminated too many points, the remaining points will not be sufficient to extract the contour reliably, even if e is small. We will denote by $f_c = N_k/N$ the fraction of columns left with at least one point at the end of a morphological operation. This factor must also be taken into account when deciding on which method provides a better result.

Additionally, CPU times of the algorithms (t_{CPU}) are measured. These represent the total time the computer takes to realize the morphological operations

starting with the raw TOF data. Morphological operations are implemented in the C programming language and the programs are run on a 200 MHz Pentium PC.

The result of applying $n = 5$ thinning to the sonar data in Figure 2(b) is presented in Figure 4. The results of various morphological operators applied to the same map are summarized in Table 1. Error measures E_1 and E_2 , given in the same table, will be discussed in the next section. Since simple erosion results in very small values of f_c , we have considered thinning with parameter n . The error e tends to decrease with increasing n . However, larger values of n tend to result in smaller values of f_c so that a compromise is necessary. For the time being, we note that the thinning parameter n allows one to trade off between e and f_c .

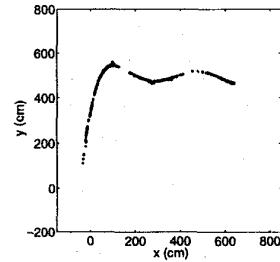


Figure 4: Result of $n = 5$ thinning: $e = 0.0496$, $f_c = 0.464$, $t_{CPU} = 1.07$ s.

3.2 Curve Fitting

As a last step, curve fitting is applied in order to achieve a compact representation of the surface profile in the robot's memory. Since our aim is to fit the best curve to the points, not necessarily passing through all of them, least-squares optimization (LSO) is preferred to interpolation. LSO finds the coefficients of the best-fitting polynomial $p(x)$ of order m (which is predetermined) by minimizing

$$E_p^2 = \sum_{i=1}^N \sum_{j=1}^{M_i} [p(x_i) - f_{ij}]^2 \quad (2)$$

where E_p^2 is the sum of the squared deviations of the polynomial values $p(x_i)$ from the data points f_{ij} . x_i is the horizontal coordinate corresponding to the i th column of the map matrix and f_{ij} is the vertical coordinate of the j th point along the i th column. The polynomial coefficients are obtained by solving linear equations obtained by setting the partial derivatives equal to zero. Once an acceptable polynomial approximation is found, the surface can then be compactly represented by storing only these coefficients.

To assess the overall performance of the method, two final error measures are introduced, both comparing the final polynomial fit with the actual surface:

$$E_1 = \sqrt{\frac{1}{N} \sum_{i=1}^N [p(x_i) - y_i]^2} \quad (3)$$

$$E_2 = \frac{E_1}{\sigma_y} \quad (4)$$

The first is a root-mean-square absolute error measure, with dimensions of length, which should be interpreted with reference to the wavelength λ , that represents the ultimate resolving power of the system. The second is a dimensionless relative error measure which can be interpreted as the error relative to the variation of the actual surface.

The curve fitted to the surface map after thinning shown in Figure 4 is presented in Figure 5. Table 1 shows that increasing n improves e but worsens f_c and that E_1 and E_2 achieve a minimum at some value of n (which in this case happens to occur at $n = 5$ for both E_1 and E_2). In the simulations, where the actual surface is known, it is possible to choose the optimal value of n , minimizing E_1 or E_2 . In real practice, this is not possible so that one must use a value of n judged appropriate for the class of surfaces under investigation.

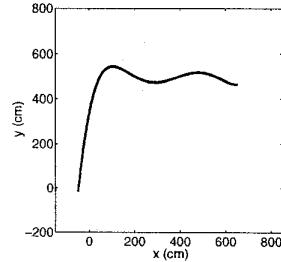


Figure 5: Polynomial fit of order $m = 11$: $E_1 = 3.57$ cm, $E_2 = 0.040$.

morphological operation	e	f_c	t_{CPU} (s)	E_1 (cm)	E_2
thinning ($n = 2$)	0.09	0.97	1.12	6.32	0.068
thinning ($n = 3$)	0.09	0.83	1.11	4.87	0.053
thinning ($n = 4$)	0.05	0.64	1.09	3.74	0.041
thinning ($n = 5$)	0.05	0.46	1.07	3.57	0.039
thinning ($n = 6$)	0.04	0.27	1.07	5.47	0.059
thinning ($n = 7$)	0.03	0.13	1.06	6.93	0.075

Table 1: Results of various morphological operations.

In the simulations, higher-order reflections (i.e. echoes detected after bouncing off from object surfaces more than once) are ignored since they are difficult to model, although they almost always exist in practice. The key idea of the method is that a large number of data points coincide with the actual surface (at least at the tangent points of the arcs) and the data points off the actual surface are more sparse. Those spurious arcs caused by higher-order reflections and crosstalk

also remain sparse and lack reinforcement. The thinning algorithms eliminate these spurious arcs together with the sparse arc segments resulting from the angular uncertainty of the sensors.

4 Experimental Verification

In this section, the method is verified using the sensor systems on the Nomad 200 mobile robot in our laboratory.

4.1 System Description

The Nomad 200 is an integrated mobile robot of height 76.2 cm and diameter 45.7 cm, including tactile, infrared, sonar, and structured-light sensing systems (Figure 6). The robot can translate only in the forward and backward directions but not sideways without rotating first. The maximum translational and rotational speeds of the Nomad 200 are 60 cm/s and 60°/s respectively [10].



Figure 6: Nomad 200 mobile robot.

Nomad 200 has onboard computers for sensor and motor control and for host computer communication. The communication is managed with a graphic interface (server). The robot can also be run from a C language program either through the server or directly [10].

The Sensus 200 Sonar Ranging System on the robot consists of 16 Polaroid transducers which can yield range information from 15 cm to 10.7 m with $\pm 1\%$ accuracy. The Polaroid transducer has beamwidth 25° and resonant frequency 49.4 kHz.

The Sensus 500 Structured-Light System consists of a laser diode (as its light source) and a CCD array camera. The operating range of the system is from 0.305 m to 3.05 m. The range is determined by (laser line striping) triangulation, which causes decreasing accuracy with increasing range and also possible angular measurement errors.

In the experiments, both sonar and structured-light data are collected from various surfaces constructed in our laboratory. The structured-light system is much more expensive and complex, requiring higher-power and sufficient ambient light for operation. Since it

reveals a very accurate surface profile, the surface detected by this system is used as a reference in the experimental calculation of the errors using sonar data.

In order to prevent any crosstalk between consecutive pulses, the sonars should be fired at 62 ms intervals since the maximum range of operation of Polaroid transducers is 10.7 m. In the experiments, the sonars are fired at 40 ms intervals. This prevents much of the crosstalk, and in the few cases where erroneous readings are obtained due to crosstalk, these are readily eliminated by the algorithm. This is another aspect in which the algorithm exhibits its robust character.

4.2 Experimental Results

Several surfaces have been constructed in our laboratory with different curvature and dimensions, using thin cardboard of height 1.05 m and length 3.65 m.

The sonars on the Nomad 200 are in a circular configuration and only the five front sensors are activated. Since the robot has a limited number of sensors which can detect the surface, by moving the robot and rotating its turret, the equivalent of a much larger number of sensors is created synthetically. First, the robot remained stationary and collected data by rotating its turret. However, there were many locations on the surface which could not be seen by the robot if only the turret rotated. On the contrary, pure translation alongside the surface generally provided satisfactory results.

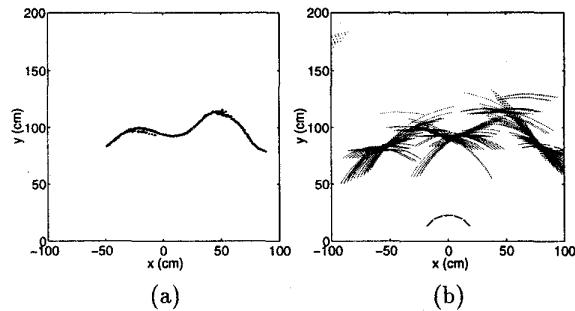


Figure 7: a) The surface profile revealed by the structured-light data, b) sonar data.

The structured-light data obtained from one of the constructed cardboard surfaces is presented in Figure 7(a). The sonar data presented in Figure 7(b) is obtained by translating the mobile robot horizontally over a distance of 1.5 m along the line $y = 0$ and collecting data every 2.5 cm. The turret is oriented such that both the structured-light and the front sonars are directed towards the surface and it does not rotate throughout the translational movement.

As expected, the structured-light data provides a very accurate surface profile. In the arc map obtained by sonar, there are some arcs which are not tangent to the actual surface at any point. These correspond to spurious data due to higher-order reflections, readings from other objects in the room, or totally erroneous

readings. These points are readily eliminated by morphological processing (Figure 8(a)). If the final curve in Figure 8(b) is compared with the structured-light data in Figure 7(a), it can be observed that a close fit to the original surface is obtained. The errors in this case are $E_1 = 1.41$ cm, $E_2 = 0.156$, and $t_{CPU} = 0.15$ s. Generally speaking, the error is larger where curvature is greater.

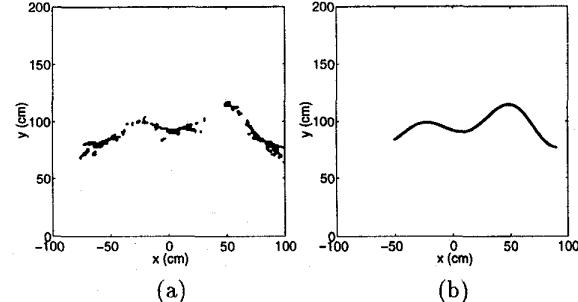


Figure 8: a) Result of erosion ($n = 8$) followed by pruning ($n = 1$) applied after erosion. b) Polynomial fit to part (a): $E_1 = 1.41$ cm, $E_2 = 0.156$, and $t_{CPU} = 0.15$ s.

Several results obtained for this surface are summarized in Table 2. All polynomials are of degree $m = 10$. The minimum estimation error E_1 (obtained from the case in Figure 8) is not much larger than the wavelength $\lambda = 6.9$ mm which represents the fundamental resolution limit of the system.

morphological operation	t_{CPU} (s)	E_1 (cm)	E_2 (cm)
thinning ($n = 5$)	0.12	2.58	0.287
thinning ($n = 6$)	0.12	1.96	0.218
thinning ($n = 7$)	0.12	1.63	0.182
erosion ($n = 8$)	0.11	1.42	0.158
erosion & pruning	0.15	1.41	0.156
erosion & thinning ($n = 2$)	0.14	1.50	0.167

Table 2: Experimental results.

Closing operations were not needed in processing the experimental data because the points were sufficiently dense. If this was not the case, one would first apply closing in order to add extra points to fill the gaps between the points of the original map.

4.3 Computational Cost of the Method

The average CPU times are in general of the order of a few seconds, indicating that the method is viable for real-time applications. For comparison, the time it takes for an array of 16 sonars to collect all the TOF data is 16×40 ms = 0.64 s which is of the same order of magnitude as the processing time. It should be noted that the actual algorithmic processing time is a small fraction of the CPU time, as most of the time is consumed by file operations, reads and writes to disk,

matrix allocations etc. Thus, it seems possible that a dedicated system can determine the surface profile even faster, bringing the computation time below the data collection time.

5 Discussion and Conclusion

A novel method is described for determining arbitrary surface profiles by applying morphological processing to sonar data. The method is both extremely flexible, versatile, and robust, as well as being simple and straightforward. It can deal with arbitrary numbers and configurations of sensors as well as synthetic arrays obtained by moving a relatively small number of sensors. Accuracy increases with the number of sensors used (actual or synthetic) and has been observed to be quite satisfactory, except when the radius of curvature of the surface is very small [9]. The method is robust in many aspects; it has been seen that it has the inherent ability to eliminate undesired TOF readings arising from higher-order reflections as well as the ability to suppress crosstalk when the sensors are fired at shorter intervals than that nominally required to avoid crosstalk. In addition, the method can effectively eliminate spurious TOF measurements due to noise, and process multiple echoes informatively.

The processing time is small enough to make real-time applications feasible. For instance, the system can be used for continual real-time map building purposes on a robot navigating in an environment with vertical walls of arbitrary curvature. Two extensions immediately come to mind: First, it is possible for the robot to continually add to its collection of arcs and reprocess them as it moves, effectively resulting in a synthetic array with more sensors than the robot actually has. Second, the method can be generalized to three-dimensional environments with the arcs being replaced by spherical or elliptical caps and the morphological rules extended to three dimensions. In certain problems, it may be preferable to reformulate the method in polar or spherical coordinates.

Although the structured-light system has been used mainly as a reference in this study, the fact that its strengths and weaknesses are complementary to the sonar system suggests the possibility of fusing the output of the two systems. The structured-light system provides a very accurate surface profile, but introduces errors increasing with range, as a result of the triangulation technique it employs. On the other hand, sonars yield better range information over a wider range but are less adept at recognizing the contour details due to their wide beamwidth. The best properties of these two sources of information can be combined by first calibrating the structured-light range errors using sonar data, and then using the profile found by the structured-light system as a reference for sonar. Despite this possibility, the method described in this paper may be preferable in many circumstances, since the structured-light system is much more expensive and complex compared to sonar sensors.

Although not fully reported here, a detailed quantitative study of the performance of different morpho-

logical operations as well as the effect of TOF noise, and the dependence of the error on surface curvature, spatial frequency, and distance can be found in [9].

The essential idea of this paper — the use of multiple range sensors combined with morphological processing for the extraction of the surface profile — can also be applied to other physical modalities of range finding of vastly different scales and in many different application areas. These may include radar, underwater sonar, optical sensing and metrology, remote sensing, ocean surface exploration, geophysical exploration, and acoustic microscopy. Some of these applications (e.g. geophysical exploration) may involve an inhomogeneous and/or anisotropic medium of propagation. It is envisioned that the method could be generalized to this case by constructing broken or non-ellipsoidal arcs.

References

- [1] R. Kuc and M. W. Siegel, "Physically Based Simulation Model for Acoustic Sensor Robot Navigation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-9, No. 6, pp. 766-778, 1987.
- [2] J. J. Leonard and H. F. Durrant-Whyte, "Mobile Robot Localization by Tracking Geometric Beacons," *IEEE Trans. on Robotics and Automation*, Vol. 7, No. 3, pp. 376-382, 1991.
- [3] B. Barshan and R. Kuc, "Differentiating Sonar Reflections from Corners and Planes by Employing an Intelligent Sensor," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 12, No. 6, pp. 560-569, 1990.
- [4] H. Peremans, K. Audenaert, and J. M. V. Campenhout, "A High-Resolution Sensor Based on Triangular Perception," *IEEE Trans. on Robotics and Automation*, Vol. 9, No. 1, pp. 36-48, 1993.
- [5] L. Kleeman and R. Kuc, "Mobile Robot Sonar for Target Localization and Classification," *Int. J. of Robotics Research*, Vol. 14, No. 4, pp. 295-318, 1995.
- [6] B. Ayrulu and B. Barshan, "Identification of Target Primitives with Multiple Decision-Making Sonars using Evidential Reasoning," *Int. J. of Robotics Research*, Vol. 17, No. 6, 1998.
- [7] M. K. Brown, "The Extraction of Curved Surface Features with Generic Range Sensors," *Int. J. of Robotics Research*, Vol. 5, No. 1, pp. 3-18, 1986.
- [8] E. R. Dougherty, *An Introduction to Morphological Image Processing*. Bellingham, Wash., USA : SPIE Optical Engineering Press, 1992.
- [9] D. Baskent, "Surface Profile Determination from Multiple Sonar Data Using Morphological Processing," M.S. Thesis, Bilkent University, 1998.
- [10] Nomadic Technologies, Inc. Nomad 200 Manual, 1997.