
Deniz Başkent
Billur Barshan

Department of Electrical Engineering
Bilkent University
Bilkent, 06533 Ankara, Turkey
baskent@usc.edu
billur@ee.bilkent.edu.tr

Surface Profile Determination from Multiple Sonar Data Using Morphological Processing

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 approach taken is extremely flexible and robust, in addition to being simple and straightforward. It can deal with arbitrary numbers and configurations of sensors as well as synthetic arrays. 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.

KEY WORDS—sonar sensing, range sensing, morphological processing, mathematical morphology, surface-profile extraction, map building, pattern recognition, mobile robots

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.

Sonar sensors, robust and inexpensive devices that are capable of providing accurate range data, have been widely used in robotics applications. However, because of their large beamwidth, the angular resolution of sonar sensors is low, resulting in uncertainty in the location of the object encountered. Furthermore, data from these sensors are often considered difficult to interpret due to specular reflections from most

objects, especially when multiple or higher-order reflections are involved.

In robotics, most of the research on sonar has concentrated on surfaces with fixed or piecewise-constant curvature, primarily composed of target primitives such as planes, corners, edges, and cylinders. Brown (1986) described an analytical approach to surface-curvature extraction that employs differential geometry. The echoes received from planes, corners, and edges were first modeled by Kuc and Siegel (1987). Planes and corners were differentiated by using both amplitude and time-of-flight information by Barshan and Kuc (1990). In addition to planes and corners, edges and cylinders were used as natural beacons for mobile-robot localization (Leonard and Durrant-Whyte 1991). Peremans, Aude-naert, and Van Campenhout (1993) and Kleeman and Kuc (1995) processed the full echo waveform using matched filtering for more accurate target differentiation. Hong and Kleeman (1997a, 1997b) treated the classification and localization of 3-D room features using maximum-likelihood estimation. The work reported by Ayrulu and Barshan (1998) employed evidential reasoning to fuse the results of multiple sonar sensors for improved target differentiation. Apart from these works on target-primitive differentiation, an acoustic imaging system that combines holography with neural networks for the recognition of 3-D objects was described by Watanabe and Yoneyama (1992). Kuc (1997) reported differentiation of O-rings and coins using an adaptive sonar configuration mounted on the end of a robot arm.

Sonar sensors have also been extensively used for map building and obstacle avoidance in robotics. The different geometric approaches in map building basically fall into two primary categories: *feature based* and *grid based*. In feature-based approaches, the geometry of the environment is extracted from sensor data as the first step in data interpretation (e.g., edge detection, or straight line or curve fitting to obstacle boundaries). Important issues to consider are the

representation of uncertainty, suitability of the selected feature to the environment and type of data, the reliability of feature extraction, and the speed with which the model can be constructed. Crowley employed a feature-based approach to world modeling with sonar using line segments as features, and matching the local line-segment map to a global map (Crowley 1985). Straight lines are extracted from time-of-flight (TOF) measurements, then matched and incorporated into a *composite model* of the environment. The results are then Kalman filtered to estimate and update the positions and orientations of the line segments (Crowley 1989). A different approach for robot localization in a known environment was presented by Drumheller (1987). In this method, *sonar segments* were generated from range readings that were pairwise matched to the environment model by employing the interpretation tree of Grimson and Lozano-Pérez (1984). Possible interpretations were passed through a *sonar barrier test*. If multiple interpretations existed after the test, their number was reduced by *maximum-contact heuristics*. Cox used a positioning system based on matching a local grid map to a global line-segment map (Cox 1991).

Many researchers have reported the extraction of line segments from sonar data as being difficult and brittle (Leonard 1990). Experimental results have indicated that straight lines obtained from TOF measurements do not necessarily match and/or align with the world model, and may yield many spurious and erroneous line segments. Improving the algorithms for detecting line segments and including heuristics does not really solve the problem. Leonard advocates another feature-based representation using regions of constant depth (RCD) as features, extracted directly from raw sonar data (Leonard and Durrant-Whyte 1991; Leonard, Durrant-Whyte, and Cox 1992). Regions of constant depth are circular arcs that are natural features of sonar readings from specularly reflecting surfaces, first reported by Kuc and Siegel (1987). Approaches based on physical model-based reasoning, including classification of environmental features into target primitives discussed earlier (Barshan and Kuc 1990; Bozma 1992; Kleeman and Kuc 1995; Hong and Kleeman 1997a, 1997b; Ayrulu and Barshan 1998), can also be considered feature-based methods.

An alternative representation is the *certainty* or *occupancy grid*. In certainty-grid methods, the environment is discretized into a 2-D array of cells. A “certainty” measure is associated with each cell by assigning the cell a single value, between 0 and 1, representing the probability of that cell being occupied. For each range reading, the values of the cells within the sensor-beam profile are updated to reflect current occupancy. Grid-based methods are particularly useful for obstacle avoidance, since free and occupied regions of space are explicitly represented. The difficulty of line-segment-based feature extraction was an important factor in the development of the grid concept. This representation was first proposed by Moravec and Elfes (Moravec and Elfes 1986; Moravec 1989), and further developed by Elfes and his coworkers (Elfes 1987,

1990; Elfes and Matthies 1987). Although grid-based methods have their limitations in terms of memory and resolution, they are advantageous in that they do not commit to making difficult geometric decisions early in the data-interpretation process. On the other hand, since different target types are not explicitly represented, it is not as easy to predict what sonar data will be observed from a given position and to give an account of individual sonar readings. Typically, when sufficient sensor data have been collected in the grid cells, the data are matched or correlated with a global model. This process can be computationally intensive and time consuming, depending on the cell resolution of the grid.

Borenstein and Koren (1991) also employed a grid-based approach to represent sonar data in their work on obstacle avoidance. The vector-field histogram method used a simplified grid-update procedure that allowed in-motion sampling of the ultrasonic range sensors and a very rapid sensor-firing rate. Initially, a 2-D Cartesian histogram grid was used as a world model, which was then reduced to a 1-D polar histogram centered at the robot. Instead of updating all the cells in the sensor-beam profile (as in the original occupancy grid of Elfes), only a single cell, which was along the line of sight and at the distance corresponding to the range reading, was updated. Sufficient grid coverage was obtained by the extremely fast firing and update rate, and the most suitable steering direction was selected from the resulting polar map for real-time obstacle avoidance. A measure for judging map accuracy was proposed by Raschke and Borenstein (1990), where improved mapping results were reported with this method.

Gilbreath and Everett (1988), Zelinsky (1988), Beckerman, and Oblow (1990) have all used variations of the grid-based approach to construct sonar maps for path planning. The use of grids has been extended to 3-D sonar-range sensing for under water by Auran and Silven (1996). Pagac, Nebot, and Durrant-Whyte (1998) examined the problem of constructing and maintaining a 2-D occupancy-grid map of a robot environment using evidential reasoning. New sonar readings were incorporated in the map using Dempster’s rule of combination. Recently, fuzzy logic has been introduced to represent uncertainty in map building (Gasos and Martin 1996; Oriolo, Ulivi, and Vendittelli 1997, 1998). Kurz (1996) partitioned free-space into regions in which a specific situation could be recognized based on sonar ranging. These regions were then attached to graph nodes, generating a map of the environment in the form of a graph representation.

In the work of Wallner and Dillman (1995), a hybrid method was presented for updating the local model of the perceivable environment of a mobile robot. Local grids could be used in dynamic environments, which was not possible with the earlier grid-based approaches. The method combined the advantages of feature- and grid-based environment-modeling techniques. More detail on the different approaches to map building can be found in Borenstein, Everett, and Feng’s (1996) work.

The approaches described above are often limited to elementary target types or simple sensor configurations. On the other hand, the method presented in this paper is aimed at the determination of arbitrary surface profiles that are typically encountered in mines, rough terrain, or under water. The approach is completely novel in that morphological processing techniques are applied to sonar data to reconstruct the profile of an *arbitrarily* curved surface. It can also be looked at as a novel method for solving a class of nonlinear reconstruction (inverse) problems. It is important to emphasize that morphological processing is employed here to process the sonar map of the surface being constructed in the robot's memory, rather than conventional camera images. The method has sufficient generality to find application in other ranging systems such as radar, optical sensing and metrology, remote sensing, and geophysical exploration.

From a map-building perspective, this method can also be considered as a hybrid of feature- and grid-based methods: initially, the environment is discretized into rectangular grids. After accumulation of a sufficient amount of data, a curve-fitting procedure is employed to extract the geometry of the surface under investigation. The present approach is advantageous over probability-based grid methods since it allows the use of a much finer physical grid. This is because the approach does not rely on the accumulation of multiple measurements in each cell. From a different perspective, although it is possible to interpret this method as a spatial voting scheme where cells are locally supported by their neighbors, we find it more appropriate to look at it from a nonlinear signal-reconstruction perspective, where morphological processing is used to extract reinforced features in the arc map.

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 mentioned above, a commonly noted disadvantage of sonar sensors is the difficulty associated with handling spurious readings, crosstalk, and higher-order and multiple reflections. The method proposed is capable of effectively suppressing spurious readings, crosstalk, and most 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 one can also be processed informatively.

This paper is organized as follows: 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. Using the speed of ultrasonic waves in air ($c = 343.3$ m/sec at

room temperature), the range r can be easily calculated from $r = ct_o/2$, where t_o denotes the TOF. Many ultrasonic transducers operate in this pulse-echo mode (Hauptman 1993). The transducers can function 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 a resonance frequency of 200 kHz so that the propagating waves have wavelengths well above several millimeters. Details on the objects that are smaller than the wavelength cannot be resolved (Brown 1986). The sonars used in our experimental setup are Polaroid 6500 series transducers (Polaroid 1997) operating at a resonance frequency $f_o = 49.4$ kHz, which corresponds to a wavelength of $\lambda = c/f_o = 6.9$ mm at room temperature.

The major limitation of sonar sensors comes from their large beamwidth. Polaroid transducers have a half-beamwidth angle of $\theta_o = \pm 12.5^\circ$. 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 a circular arc whose radius is determined by $r = ct_o/2$, as illustrated in Figure 1a. 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 (Fig. 1b). Notice that the reflecting surface is tangent to these arcs at the actual point(s) of reflection. The angular extent of these arcs is determined by the sensitivity regions of the transducers.

If multiple echoes are detected for the same transmitting/receiving pair, circular or elliptical arcs are constructed to correspond to each echo. (However, not all systems commonly in use are able to detect echoes beyond the first one.)

Most commonly, the large beamwidth of the transducer is accepted as a device limitation that 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.

3. Processing of the Sonar Data

Structured configurations such as linear and circular arrays as well as irregular sensor configurations have been considered.

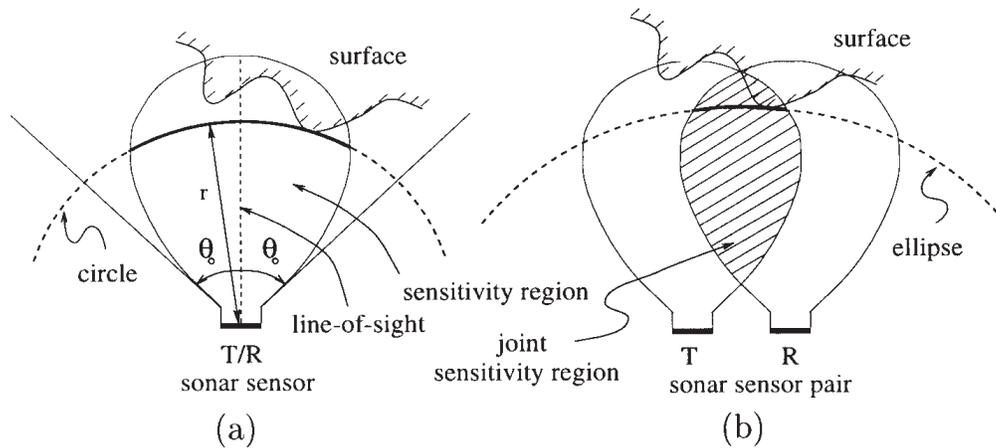


Fig. 1. For the same sonar transmitting and receiving, the reflecting point is known to be on the circular arc shown (a). An elliptical arc is drawn if the wave is transmitted and received by different sensors (b). The intersection of the individual sensitivity regions serves as a reasonable approximation to the joint-sensitivity region.

(The mobile robot used in the experiments has a circular array of sonar sensors.)

Figure 2a shows a surface whose profile is to be determined. Figure 2b shows the circular and elliptical arcs obtained from circular arrays of sensors, which both rotate and translate to increase the number of arcs generated from the available number of sensors. Further sonar maps obtained using a circular configuration are presented in Section 3.3.2.

Notice that although each arc represents considerable uncertainty as to the angular position of the reflection point, nevertheless one can almost extract the actual curve shown in Figure 2a by visually examining the arc map in Figure 2b. Each arc drawn is expected to be tangential 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 2b, 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, remain more sparse and isolated.

Morphological processing is employed to achieve what is natural for the human visual perception system: the extraction of Figure 2a from Figure 2b. Morphological erosion and dilation operators are used to weed out the sparse arc segments, leaving us with the mutually reinforcing segments which will reveal the original surface.

3.1. Morphological Processing

The main application of morphological operations is in modifying regions and shapes of images (Low 1991). Therefore, they are widely used in image processing to accomplish tasks

such as edge detection, enhancement, smoothing, and noise removal (Dougherty 1992; Myler and Weeks 1993).

In this study, morphological processing is used to eliminate the sparse and isolated spikes and segments in the sonar arc map, directly revealing the surface profile.

Erosion and *dilation* (also referred as Minkowski addition and subtraction) are the two fundamental morphological operations used to thin or fatten an image, respectively. These operations are defined according to a structuring element or template. An example 3×3 -square template is shown in Figure 3a.

A simple 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 pixels that overlap with the template (Pitas 1993). If they are all identical, the central pixel with value 1 will preserve its value; otherwise, it is set to 0. For the template shown in Figure 3a, all eight neighbors of the image pixel must be 1 ($n = 8$), and the image is eroded or shrunk accordingly. For better insight on the erosion algorithm, an example is presented in Figures 3b and 3c.

The dilation algorithm is very similar to that for erosion, but it 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 chosen to be the 3×3 -square template shown in Figure 3a with the central pixel encircled. Since the template is symmetric, the image is fattened (dilation) or thinned (erosion) in all directions by 1 pixel.

The direct use of erosion may eliminate too many points and result in the loss of information characterizing the surface. For such cases, the compound operations of *opening*

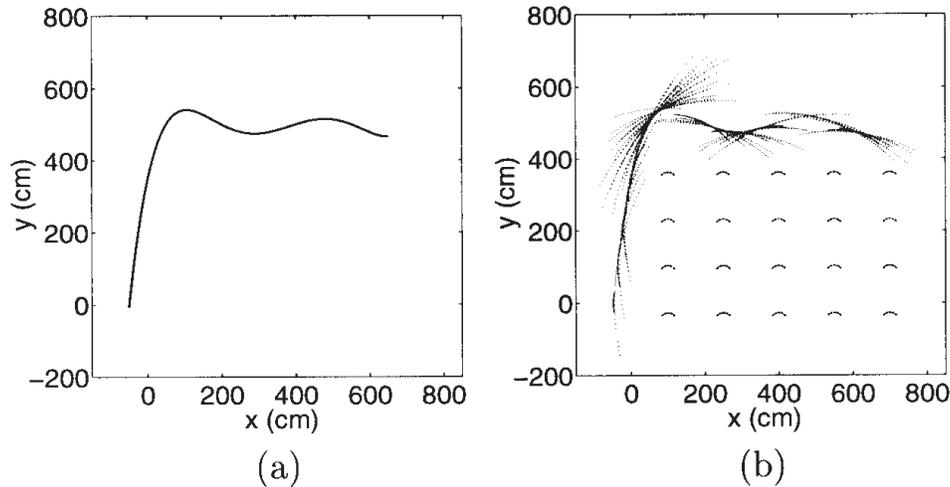


Fig. 2. The original surface (a). The circular sensor array mounted on a mobile robot moves to 20 different locations and collects data by rotating around its center from 45° to 135° in 15° steps (b). The angles are measured with respect to the positive x-axis in the counterclockwise direction. The circular array has been shown at the 20 locations at its 90° position.

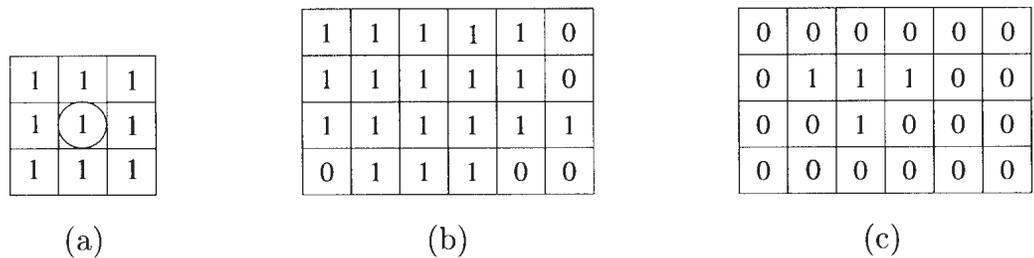


Fig. 3. An erosion example: the template (a), the original image (b), and the image after erosion (c).

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 in the image (Myler and Weeks 1993). Closing is applied prior to thinning (described below) in cases where the points are not closely connected to each other, as 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 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 for the central pixel to remain equal to 1 after the operation. This operation is used to eliminate isolated points (Dougherty 1992). 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:

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

Here, i is the discrete index along the x -direction, and y_i is the discretized function representing the actual surface profile with variance $\sigma_y^2 = \frac{1}{N} \sum_{i=1}^N (y_i - \frac{1}{N} \sum_i y_i)^2$. The total number of columns is N , whereas N_k represents those columns left with at least one point as a result of some morphological operation. The vertical position of the median (centermost) point along the i th column of the map matrix is m_i (e.g., Fig. 4a). 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.

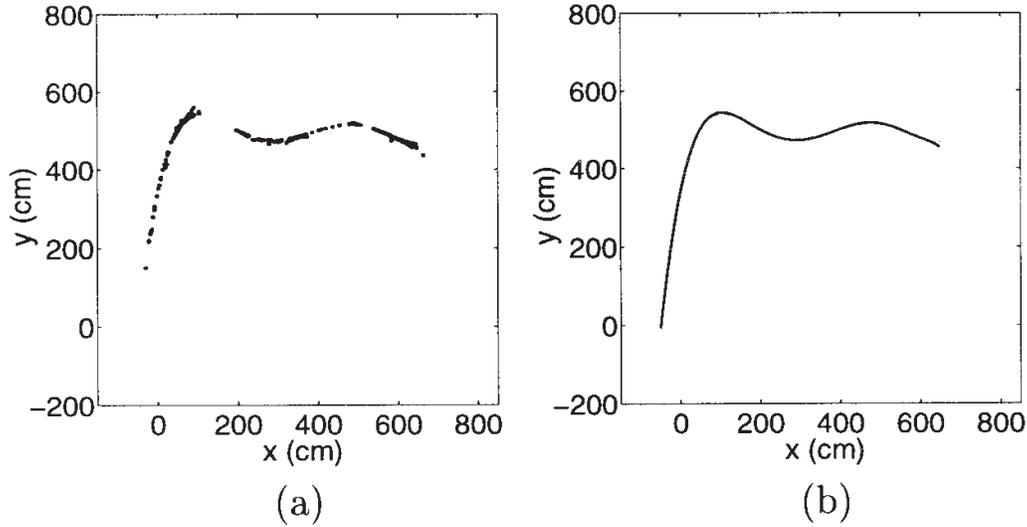


Fig. 4. The result of $n = 3$ thinning: $e = 0.052$, $f_c = 0.341$, $t_{\text{CPU}} = 1.11$ sec (a); the original surface (dashed line) and the polynomial fit of order $m = 9$ (solid line), with $E_1 = 3.18$ cm and $E_2 = 0.036$ (b).

We 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 which method provides a better result.

Additionally, the 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 Pro PC.

The result of applying $n = 3$ thinning to the sonar arc map shown in Figure 2b is presented in Figure 4a, and 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, are defined and are discussed in Section 3.2. 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 .

3.2. Curve Fitting

As a last step, curve fitting is applied 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. Using LSO, we find 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 index j runs through the M_i points along column i , and N is the number of columns. If $M_i = 0$ for a certain column, the inner summation is not evaluated, and is taken as zero for that column. The coefficients of the polynomial are determined by solving the linear equations obtained by setting the partial derivatives equal to zero (Lancaster and Salkauskas 1986). Once an acceptable polynomial approximation is found, the surface can then be represented compactly by storing only the coefficients of the polynomial. Although polynomial fitting has been found to be satisfactory in all of the cases considered, other curve-representation approaches such as the use of splines might be considered as alternatives to polynomial fitting.

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

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

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

Table 1. Results of Various Morphological Operations and Curve Fitting of Order $m = 9^a$

Morphological Operation	e	f_c	E_1 (cm)	E_2	t_{CPU} (sec)
Thinning ($n = 1$: pruning)	0.168	0.924	7.83	0.089	1.15
Thinning ($n = 2$)	0.074	0.637	4.42	0.050	1.12
Thinning ($n = 3$)	0.052	0.341	3.18	0.036	1.11
Thinning ($n = 4$)	0.045	0.160	5.37	0.061	1.10
Thinning ($n = 5$)	0.012	0.063	16.21	0.183	1.09
Thinning ($n = 6$)	0.014	0.021	464.30	5.246	1.10
Thinning ($n = 7$)	0.007	0.003	1915.45	21.641	1.10
Thinning ($n = 8$: erosion)	—	0.000	—	—	1.10
Closing and thinning ($n = 3$)	0.121	0.760	6.72	0.076	7.10
Closing and thinning ($n = 4$)	0.122	0.641	6.88	0.078	7.09

a. Since $f_c = 0$ for $n = 8$ thinning, reflecting the fact that all points are eliminated, the values of e , E_1 , and E_2 are undefined for this case

The first is a root-mean-square absolute-error measure, with dimensions of length, which should be interpreted either with respect to the pixel size or with respect to the wavelength λ , which serves as a natural reference for the intrinsic 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 the thinning shown in Figure 4a is presented in Figure 4b. 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 $n = 3$ 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.

In the simulations, higher-order reflections¹ 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 curve are more sparse. Those spurious arcs caused by higher-order reflections and crosstalk also remain sparse, and lack reinforcement most of the time. The thinning algorithms eliminate these spurious arcs together with the sparse arc segments resulting from the angular uncertainty of the sensors.

3.3. Simulation Results

The approach taken in this paper is very general in that arbitrary configurations and orientations of sensors can be handled equally effectively. We begin by considering the special cases of linear and circular arrays, which might be typical of deliberately designed array structures. Following these, we also consider arrays where the sensors are situated and oriented

randomly. This not only exemplifies the general case where the array structure is irregular, but might also have several applications: for instance, it may correspond to the case where arcs are accumulated by any number of randomly moving and rotating sensors, perhaps mounted on mobile robots as in swarm robot applications. Another potential area of application of the random configuration is in array signal processing, where the individual sensor positions of a regular array are perturbed by wind or waves.

3.3.1. Linear Arrays

First, we consider arrays of sensors configured linearly. The first linear array considered has a horizontal extension of 5 m with 10-cm spacing between the sensors (51 sensors, as shown in Fig. 5). Given the beamwidth of the sensors (25°), the number of arcs collected is relatively small. The situation gets even worse as the array is brought closer to the surface, since interaction between an even smaller number of sensor pairs becomes possible. This is a consequence of the fact that the angular beamwidth subtends a smaller arc on the surface. It is interesting to note that whereas narrow beamwidths are esteemed for their high resolving power in conventional usage of such sonar sensors, here it would have been desirable to have sensors with larger beamwidths. This would enable a greater number of sensor pairs to produce more elliptical arcs, thereby revealing the surface better and faster. However, since we restrict ourselves to the parameters of the most widely available transducer (i.e., the Polaroid), rotating the sensors is considered instead to make up for their limited beamwidth, as shown in Figures 6a and 6b.

A further consideration is that in practice the number of sensors may be limited. One way to overcome this limitation is to move a smaller array much in the same spirit as synthetic aperture radar (SAR) techniques (Skolnik 1985). However, this is not completely equivalent to the full array, since those elliptical arcs corresponding to pairs of sensors not contained within the actually existing array are missing. A further

1. A higher-order reflection refers to an echo detected after bouncing off of object surfaces more than once.

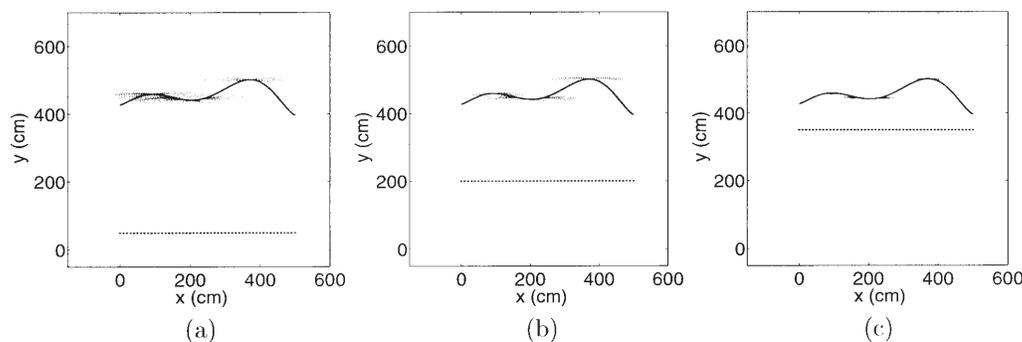


Fig. 5. The actual surface, linear sensor array with 10-cm spacing, and sonar arc map are shown for three locations of the array: the sensor array at $y = 50$ cm (a); the sensor array at $y = 200$ cm (b); and the sensor array at $y = 350$ cm (c).

extension is the combination of such movement and rotation as was the case in our first example in Figure 2.

We now return to Figures 6a and 6b, where the sensors are not translated but are only rotated. The results of morphological processing are presented in Figures 7a and 7b, and the curves fitted to them are shown in the Figures 7c and 7d. In these and later simulations, the values of n and m chosen are those that give the smallest error E_1 , unless otherwise indicated.

3.3.2. Circular Arrays

The circular configuration corresponds to the arrangement of sonar sensors on the Nomad200 mobile robot in our laboratory. Only the five sensors facing the surface are activated, since the others cannot see the surface. Again, the surface shown in Figure 2a is considered. In addition to the array locations in Figure 2b, two further examples are presented in Figures 8a and 8b. The result of applying morphological processing and curve fitting to the sonar map in Figure 8b is presented in Figures 9a and 9b.

3.3.3. Arbitrarily Distributed Sensors

In this section, the locations and the line-of-sight orientations of the sensors are generated randomly, and do not conform to any special structure. In Figure 10a, a surface is shown whose profile is to be determined with such a configuration. Figures 10b and 10c show the sonar arcs obtained using different numbers of sensors. The sensor positions and orientations shown in this figure may correspond to different physical sensors or to different locations and orientations of a smaller number of moving sensors. In Figure 11a, the surface features obtained after applying $n = 2$ thinning to the sonar arcs in Figure 10b are shown. Similarly, applying $n = 5$ thinning to the arc map in Figure 10c yields the result in Figure 11b. The curves fitted to the surface features extracted are presented in Figures 11c and 11d.

In Figure 12a, a surface with a sinusoidal profile is shown. Thinning the arc map given in Figure 12b with $n = 6$ results

in Figure 12c. The curve-fit result is presented in Figure 12d.

Although structured arrays such as linear or circular ones are often preferred in theoretical work for their simplicity and ease of analysis, the method presented here is capable of equally easily handling arbitrary arrays. In fact, the large number of simulations we have undertaken indicates that arrays consisting of irregularly located and oriented sensors tend to yield better results. This seems to result from the fact that the many different vantage points and orientations of the sensors tend to complement each other better than in the case of a uniform array. Although the question of optimal complementary sensor placement is a subject for future research, the results imply that it is preferable to work with irregular or randomized arrays rather than simple-structured arrays such as linear or circular ones.

3.3.4. Robustness to Noise

To investigate the robustness of the method to noise, zero-mean white Gaussian noise was added to the TOF readings. Figure 13 gives an example of a noisy arc map and the result of morphological processing. The noise standard deviation (σ_n) was varied logarithmically to cover a broad range of noise levels, shown in Figure 14. As expected, for σ_n smaller than 1 pixel (1 cm), the performance is approximately the same as for the noiseless case. This is expected, since the system has a resolution of 1 pixel, so that the effect of smaller perturbations is insignificant. This performance can be further improved by reducing the pixel size until it becomes comparable to the TOF measurement accuracy, at the cost of greater computation time.

The error increases significantly as the noise level increases beyond 1 pixel (Fig. 14). Since the method relies on the mutual reinforcement of several arcs to reveal the surface, larger amounts of noise are expected to have a destructive effect on this process by moving the various arc segments out of their reinforcing positions. Consequently, the arc segments that lack each other's mutual reinforcement tend to be eliminated by the morphological operations (Fig. 13). A larger

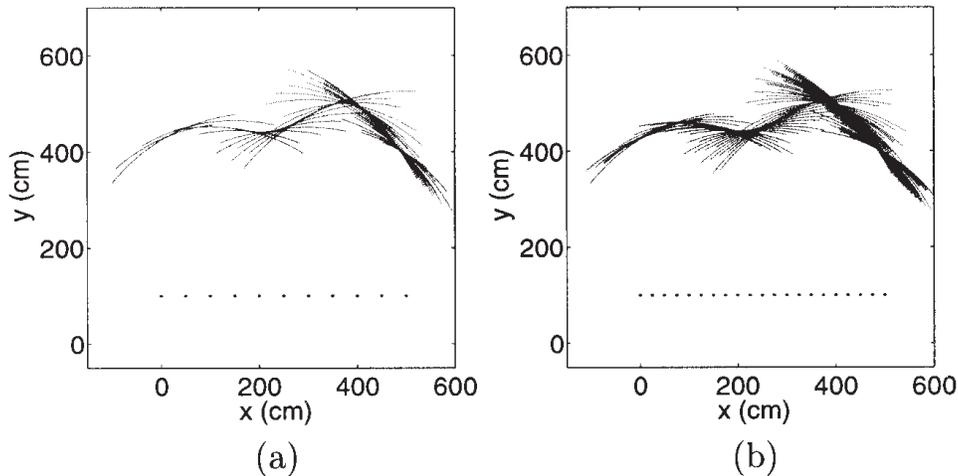


Fig. 6. Two linear arrays where the sensors are individually rotated from 40° to 140° in 10° steps: an array of 11 sensors with 50-cm spacing (a); and an array of 21 sensors with 25-cm spacing (b).

proportion of the arcs is eliminated, resulting in a loss of information characterizing the original curve. Nevertheless, the error-growth rate is not as high as might be suggested by these arguments, and the method seems to be reasonably robust to noise. In Figure 14, the performance is comparable to the noiseless case up to $\sigma_n = 10$ cm. This is partly because the least-squares polynomial fit helps eliminate some of the noise.

4. Experimental Verification

In this section, the method is verified using the sensor systems on the Nomad200 mobile robot in our laboratory.

4.1. System Description

The Nomad200 mobile robot, shown in Figure 15, has been used in the experiments. It is an integrated mobile robot including tactile, infrared, sonar, and structured-light sensing systems, with dimensions of 76.2 cm (height) by 45 cm (diameter). The mechanical system of the robot uses a non-holonomic, three-servo, three-wheel synchronous drive with zero gyro radius. The control of the base translation, base rotation, and turret rotation is performed by three separate motors. The robot can translate only in the forward and backward directions, but not sideways without rotating first. Servo control is achieved by a MC68008/ASIC microprocessor system. The maximum translational and rotational speeds of the Nomad200 are 60 cm/sec and 60° /sec, respectively. The Nomad200 is powered by an 840-Wh removable battery package (Nomadic Technologies 1997).

The Nomad200 has an onboard computer for sensor and motor control and for host-computer communication. The communication is managed with a radio link and a graphics

interface (server). The robot can be run from a C-language program either through the server or directly.

We next give a brief description of the two sensor modules used in the experiments: the Sensus 200 and Sensus 500 systems.

The Sensus 200 Sonar Ranging System consists of 16 sonars that can yield range information from 15 cm to 10.7 m with $\pm 1\%$ accuracy. The sensors are Polaroid 6500 series transducers (Polaroid 1997), which determine the range by measuring the TOF. The transducer beamwidth is 25° . The carrier frequency of the emitted pulses is 49.4 kHz. The system can be interfaced with any type of microcontroller. The power requirements of the system are 100 mA at 5 V or 12 V (Nomadic Technologies 1997).

The Sensus 500 Structured-Light System basically consists of a laser diode (as its light source) and a CCD-array camera. The laser beam is passed through a cylindrical lens to obtain planar light. The operating range of the system is from 30.5 cm to 3.05 m. Within this range, the intersection of the plane of light with the objects in the environment can be detected by the camera. The range is determined by (laser line striping) triangulation, characterized by decreasing accuracy with increasing range (Everett 1995). The power requirement of the system is 2000 mA at 12 V (Nomadic Technologies 1997).

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

To prevent any crosstalk between consecutive pulses, the sonars should be fired at 62-msec intervals, since the

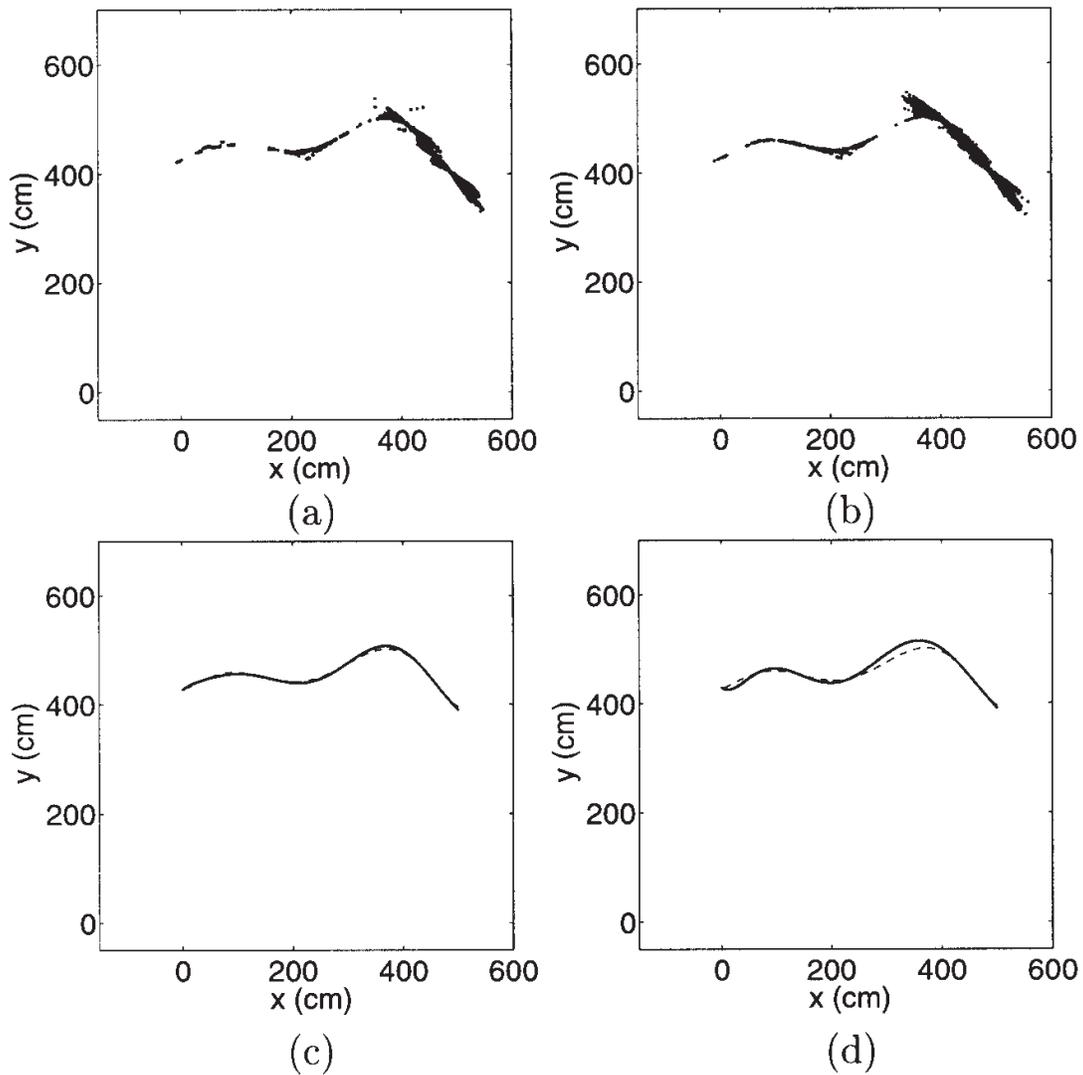


Fig. 7. (a) The result of closing and $n = 5$ thinning applied to Figure 6a: $e = 0.276$; $f_c = 0.772$; and $t_{\text{CPU}} = 4.51$ sec. (b) The result of $n = 3$ thinning applied to Figure 6b: $e = 0.353$; $f_c = 0.864$; and $t_{\text{CPU}} = 1.28$ sec. (c) The solid line indicates the polynomial fit of order $m = 8$ to part a; the dashed line shows the actual surface; $E_1 = 3.03$ cm; and $E_2 = 0.127$. (d) The polynomial fit of order $m = 7$ to part b is shown by solid line; the actual surface is indicated by dashed line; $E_1 = 7.85$ cm; and $E_2 = 0.329$.

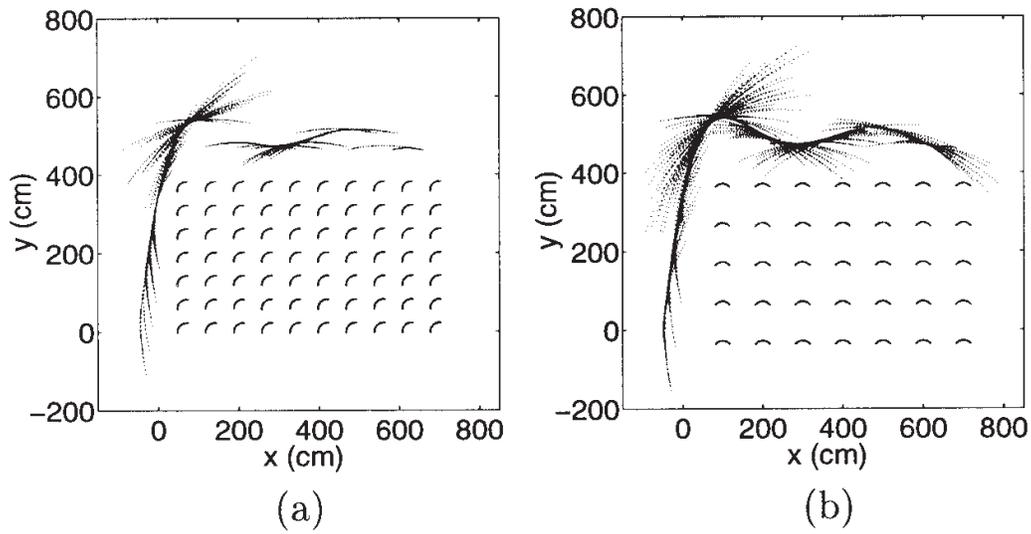


Fig. 8. (a) The robot is located at 70 locations, and the front sonar is oriented at 135° with respect to the positive x -axis; no rotation. (b) The robot rotates at 35 locations, from 45° to 135° in 15° steps.

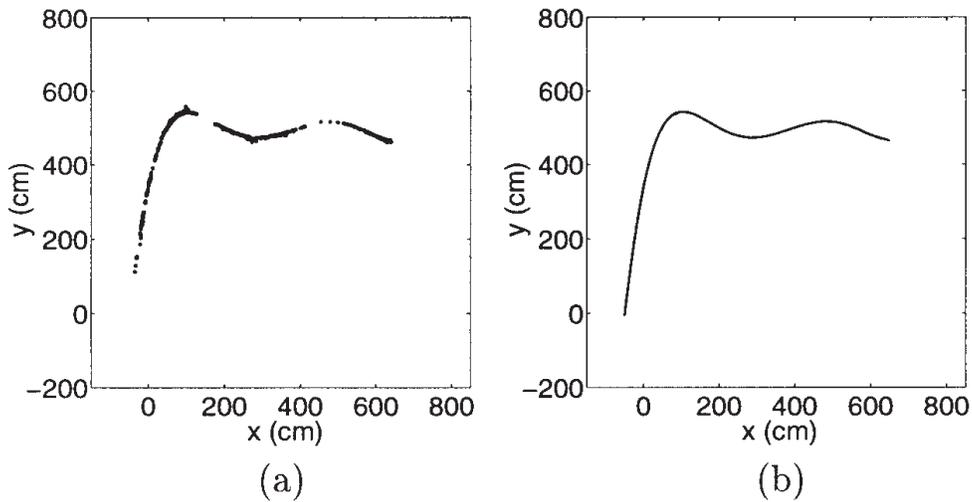


Fig. 9. The result of $n = 5$ thinning applied to Figure 8b: $e = 0.050$; and $f_c = 0.464$ (a). The polynomial fit of order $m = 11$ (solid line) and the actual surface (dashed line); $E_1 = 3.57$ cm, and $E_2 = 0.040$ (b).

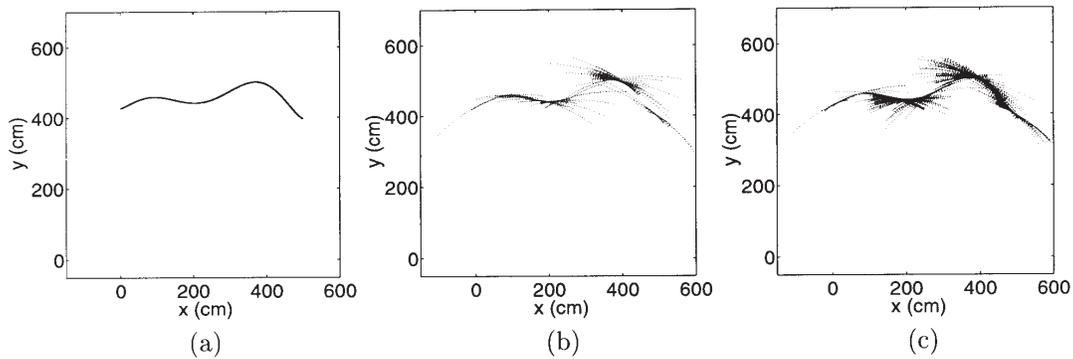


Fig. 10. The original surface (a). The arcs corresponding to the sonar TOF data collected from the surface using 100 sensors scattered and oriented randomly (not shown) (b). The x - and y -coordinates of each sensor are independent and uniformly distributed in the intervals $[0, 500]$ and $[0, 360]$, respectively. The orientation is uniformly distributed in $[40^\circ, 140^\circ]$. 150 sensors are used (not shown) (c).

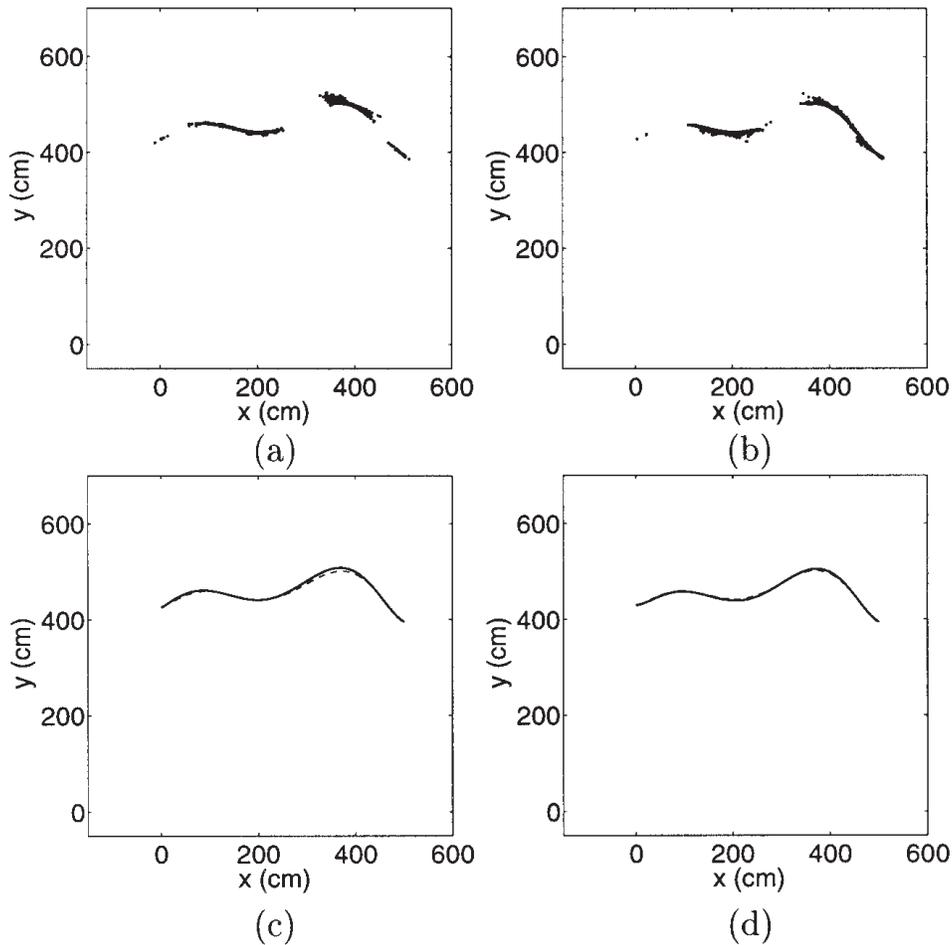


Fig. 11. (a) The result of $n = 2$ thinning applied to Figure 10b: $e = 0.296$, $f_c = 0.566$, and $t_{\text{CPU}} = 0.49$ sec. (b) The result of $n = 5$ thinning applied to Figure 10c: $e = 0.133$, $f_c = 0.610$, and $t_{\text{CPU}} = 0.73$ sec. (c) The polynomial fit of order $m = 6$ to part a (solid line), and the actual surface (dashed line); $E_1 = 3.92$ cm, and $E_2 = 0.164$. (d) The polynomial fit of order $m = 6$ to part b (solid line), and the actual surface (dashed line); $E_1 = 2.30$ cm, and $E_2 = 0.090$.

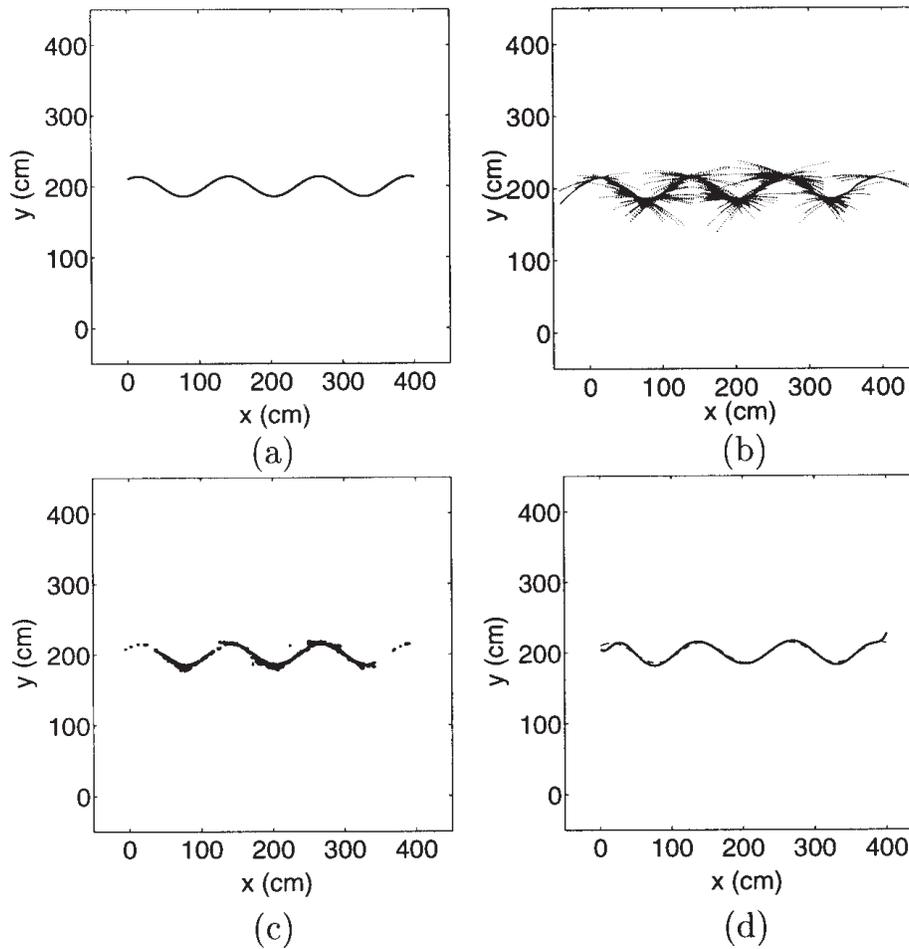


Fig. 12. The sinusoidal surface with a peak-to-peak amplitude of 30 cm and a period of 125 cm (a). The arcs corresponding to the sonar TOF data collected from the surface using 400 sensors scattered and oriented randomly (not shown) (b). The x - and y -coordinates of each sensor are independent, and are uniformly distributed in the intervals $[-100, 500]$ and $[0, 160]$, respectively. The orientation is uniformly distributed in $[40^\circ, 140^\circ]$. The result of $n = 6$ thinning (c): $e = 0.245$, $f_c = 0.777$, and $t_{\text{CPU}} = 0.51$ sec. The polynomial fit of order $m = 10$ to part c (solid line), and the actual surface (dashed line) (d): $E_1 = 2.88$ cm, and $E_2 = 0.283$.

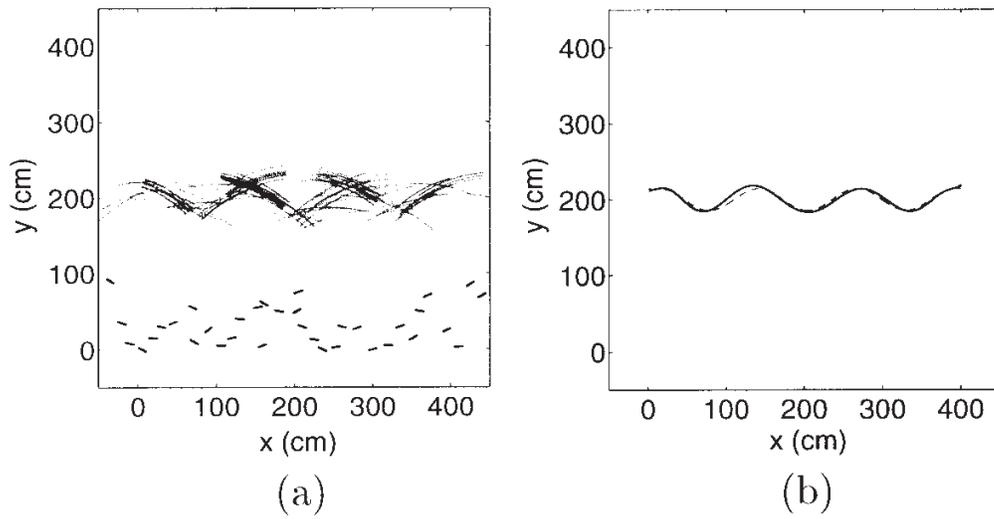


Fig. 13. The 35 sensors with an arbitrary configuration: the arc map obtained from TOF measurements when zero-mean Gaussian noise with $\sigma_n = 5$ cm is added (a); the fitted curve (solid line) after $n = 3$ thinning and the original surface (dashed line) (b); $E_1 = 3.28$ cm, and $E_2 = 0.322$.

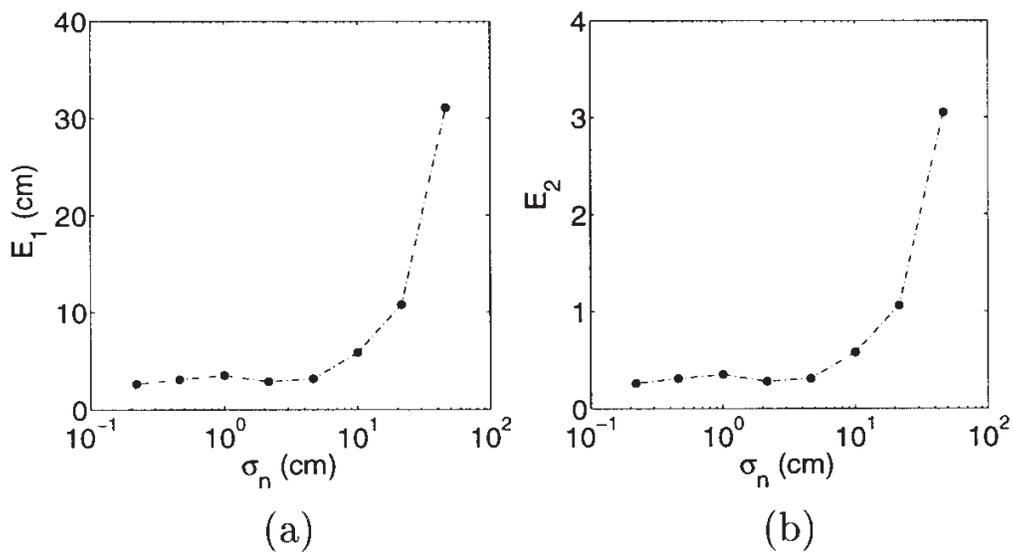


Fig. 14. E_1 (a) and E_2 (b), as the standard deviation of the noise σ_n on the TOF readings is increased.

maximum range of operation of Polaroid transducers is 10.7 m. In the experiments, the sonars were fired at 40-msec intervals. This prevented much of the crosstalk, and in the few cases where erroneous readings were obtained due to crosstalk, these false readings were readily eliminated by the algorithm. This is another aspect in which the algorithm exhibited its robust character.

4.2. Experimental Results

The sonars on the Nomad200 were in a circular configuration. Since the robot had a limited number of sensors that could detect the surface, by moving the robot and rotating its turret, the equivalent of a much larger number of sensors was created synthetically. First, the robot remained stationary and collected data by rotating its turret. However, there were many locations on the surface that could not be detected by the robot if only the turret rotated. On the contrary, pure translation alongside the surface generally provided satisfactory results.

First, several surfaces were constructed in our laboratory with different curvature and dimensions, using thin cardboard of height 1.05 m and length 3.65 m. In these experiments, only the front five sensors were activated.

The structured-light data obtained from one of the cardboard surfaces constructed are presented in Figure 16a. The sonar data presented in Figure 16b were 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 was oriented such that both the structured-light and the front five sonars were directed toward the surface, and it did not rotate throughout the translational movement.

If the sonar data in Figure 16b are examined, it can be observed that there are many points on the surface that were not detected by the sonars. Since the surface reflected specularly, segments of the surface that were relatively perpendicular to the transducer lines of sight were easily detected, whereas those segments that were more-or-less aligned with the lines of sight could not be sensed.² This is less of a problem when the surface is relatively smooth and its curvature is small. The extent to which this becomes a problem depends on the radius of curvature, as is evident when we consider a second example (below) with less curvature.

As expected, the structured-light data provided a very accurate surface profile. In the arc map obtained by sonar, there were some arcs that were not tangent to the actual surface at any point (e.g., the isolated arcs in the upper-right and upper-left parts of Fig. 16b). These correspond to spurious data due to higher-order reflections, readings from other objects in the

2. Since specular reflections involve negligible scattering and mirrorlike reflections, the transmitted waveform could be received back at the transducer only if some of the rays emerging from the transducer were perpendicular to the surface. For the case of a separate transmitter and receiver, the surface must be perpendicular to the normal determined by the transmitting and receiving elements.

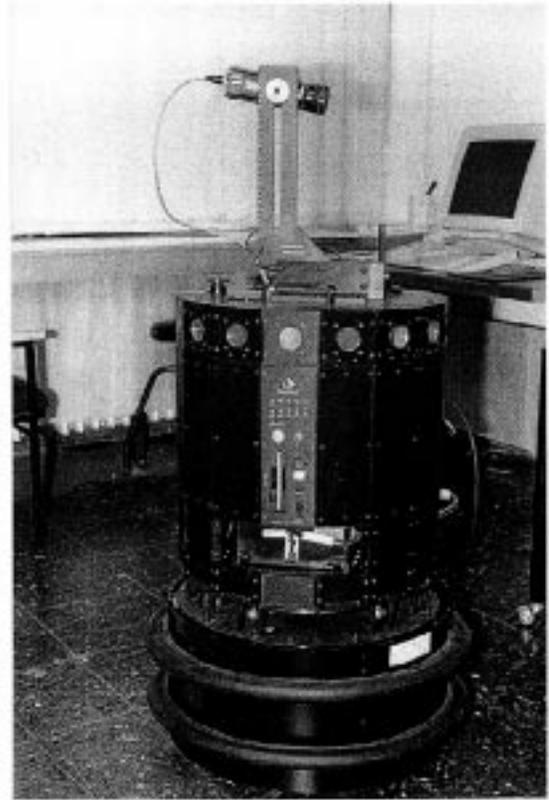


Fig. 15. The Nomad200 mobile robot. The ring of sonars can be seen close to the top rim of the turret, and the structured-light system is seen pointing rightward on top.

environment, or totally erroneous readings. These points are readily eliminated by morphological processing (Fig. 17a). If the final curve in Figure 17b is compared with the structured-light data (Fig. 17c), it can be observed that a close fit to the original surface is obtained. The errors in this case were $E_1 = 3.59$ cm and $E_2 = 0.263$, and the CPU time was $t_{\text{CPU}} = 0.27$ sec.

Next, a surface with smaller maximum curvature (hence with larger minimum radius of curvature), shown in Figure 18a, was considered. The results of morphological processing and curve fitting are shown in Figure 19, resulting in $E_1 = 1.41$ cm, $E_2 = 0.156$, and $t_{\text{CPU}} = 0.39$ sec. It was indeed observed that E_1 and E_2 were reduced significantly with respect to the previous case, as was also evidenced by the much better fit seen in Figure 19c.

Several results obtained for this surface are summarized in Table 2. All polynomials were of degree $m = 10$. The minimum estimation error E_1 was not much larger than the wavelength $\lambda = 6.9$ mm.

It is worth noting that the present approach was aimed toward *estimation* of a curve, as opposed to *detection* of target primitives. When corners or edges are encountered in the environment, the method does not fail, but the results slightly

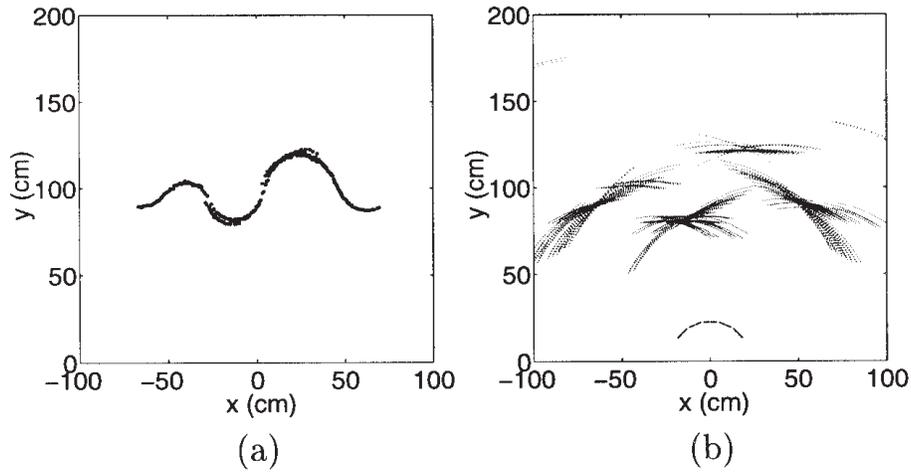


Fig. 16. The surface profile revealed by the structured-light data (a) and the sonar data (b). The data in both parts were collected from the surface at every 2.5 cm by translating the mobile robot from $(-75, 0)$ to $(75, 0)$.

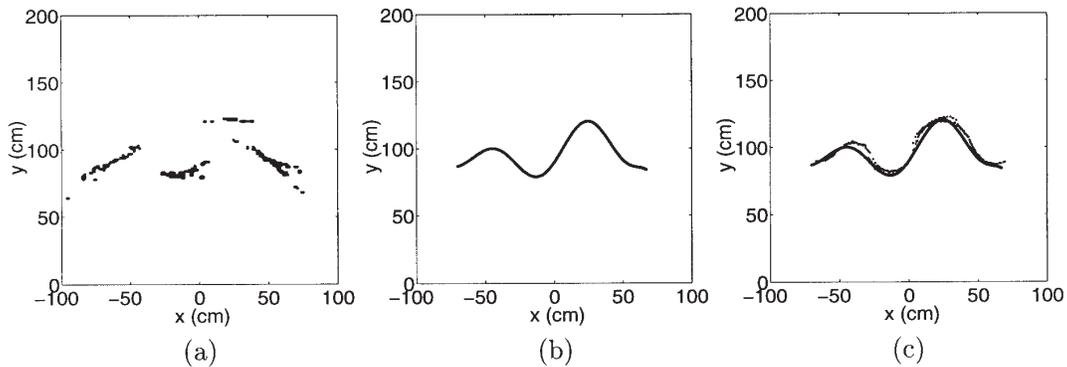


Fig. 17. The result of erosion ($n = 8$) followed by pruning ($n = 1$) (a); the polynomial fit of order $m = 10$ (b); and Figure 17(a) superimposed with part b to demonstrate the fit obtained (c); $E_1 = 3.59$ cm, $E_2 = 0.263$, and $t_{CPU} = 0.27$ sec.

Table 2. Experimental Results for the Surface Given in Figure 19a for Different Morphological Operations

Morphological Operation	E_1 (cm)	E_2	t_{CPU} (sec)
Thinning ($n = 1$: pruning)	4.98	0.554	0.41
Thinning ($n = 2$)	4.84	0.539	0.41
Thinning ($n = 3$)	4.07	0.452	0.40
Thinning ($n = 4$)	3.28	0.364	0.39
Thinning ($n = 5$)	2.58	0.287	0.37
Thinning ($n = 6$)	1.96	0.218	0.36
Thinning ($n = 7$)	1.63	0.182	0.35
Thinning ($n = 8$: erosion)	1.42	0.158	0.34
Erosion and pruning ($n = 1$)	1.41	0.156	0.39
Erosion and thinning ($n = 2$)	1.50	0.167	0.39

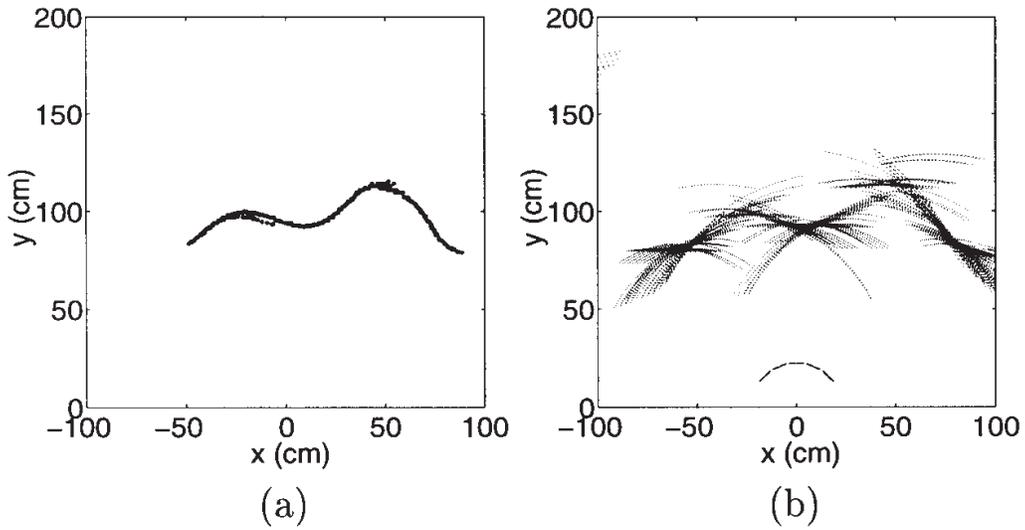


Fig. 18. The surface profile revealed by the structured-light data (a) and the sonar data (b). The data in both parts were collected from the surface at every 2.5 cm by translating the mobile robot from $(-75, 0)$ to $(75, 0)$.

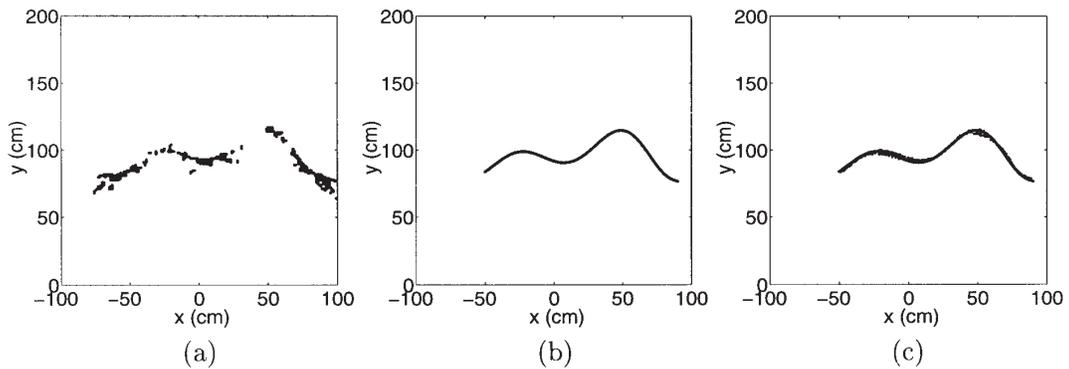


Fig. 19. The result of erosion ($n = 8$) followed by pruning ($n = 1$) (a); the polynomial fit of order $m = 10$ to part a (b); Figure 19a superimposed with part b to demonstrate the fit obtained (c); $E_1 = 1.41$ cm, $E_2 = 0.156$, and $t_{CPU} = 0.39$ sec.

deteriorate, since the method exploits neighboring relationships and local continuity (i.e., smoothness). Experimental results for 60° and 90° corners and a 90° edge are presented in Figure 20. These objects were made of smooth wood and were 1.20 m in height. The results indicate that the method works acceptably in such cases as well, the net effect being that the vertices of the sharp edges are rounded or smoothed (i.e., low-pass filtered) into curved edges. This corresponds to the spatial-frequency resolving power of the system, as determined by the chosen grid spacing.

Finally, we present experimental results obtained by using the front three sonars of the Nomad200 robot, following the walls of the room in Figure 21a. The room comprised smooth wooden (top and left) and plastered (right) walls, and a window shade with vertical flaps of 15-cm width (bottom). Some of the corners of the room were not perfect (e.g., where the shade and the right wall made a corner).

In Figure 21b, the path of the robot and the resulting arc map are given. In Figure 21c, the result of morphological processing is shown. It is clear from this figure that since spurious arc segments were fairly well eliminated, we can expect a good polar polynomial fit (or line-segment matching).

In Figure 22, we also show similar results when a cylindrical object of radius 15 cm was added to the room at a distance of 30 cm from the center of the right wall. It was observed that despite the potential for many higher-order reflections, once again a good polar fit could be expected for Figure 22c.

Even though the method was initially developed and demonstrated for specularly reflecting surfaces, subsequent tests with Lambertian surfaces of varying roughness indicate that the method also works for rough surfaces, with errors slightly increasing with roughness.

Closing operations were not needed in processing the experimental data, because the points were sufficiently dense. If this were not the case, one would first apply closing 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 on the order of several seconds, indicating that the method is viable for real-time applications. These represent the total time the computer took to realize the morphological operations starting with the raw TOF data. (The morphological processing was implemented in C language, and ran on a 200-MHz Pentium Pro PC.) For comparison, the time it took for an array of 16 sonars to collect all the TOF data was $16 \times 40 \text{ msec} = 0.64 \text{ sec}$, which was on the same order of magnitude as the processing time. It should be noted that the actual algorithmic processing time was a small fraction of the CPU time, as most of the CPU time was consumed by file operations, reads and writes to disk, matrix allocations, etc. Thus, it seems possible that a dedicated system could determine the surface profile even faster, bringing the computation time below the data-collection time.

Another important factor is memory usage. In the simulations, the objects were relatively large, and a relatively large number of sensors were employed. This led to memory usage ranging between 100–750 kB. In the experiments, the targets were smaller and relatively close, so that the data files consumed about 50–100 kB of memory. Although memory usage depends on the number of sensors used, the size of the object, and the grid size, these figures are representative of the memory requirements of the method.

5. Discussion and Conclusions

A novel method is described for determining arbitrary surface profiles by applying morphological processing techniques to sonar data. The method is extremely flexible, versatile, and robust, as well as being simple and straightforward. It can deal with arbitrary numbers and configurations of sensors in addition to synthetic arrays obtained by moving a relatively small number of sensors. Accuracy increases with the number of sensors used (actual or synthetic), and was observed to be quite satisfactory. The method is robust in many aspects: it was seen that it has the inherent ability to eliminate most of the 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 and update 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 readily generalized to three-dimensional environments with the arcs being replaced by spherical or elliptical caps and the morphological rules extended to three dimensions. The method was also found successful in determining the profile of surfaces completely surrounding the sensing system. In this case, it may be preferable to reformulate the method in polar or spherical coordinates.

Although the structured-light system was 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 of operation, but are less adept at recognizing the contour details due to their large beamwidth. Despite the possibility of

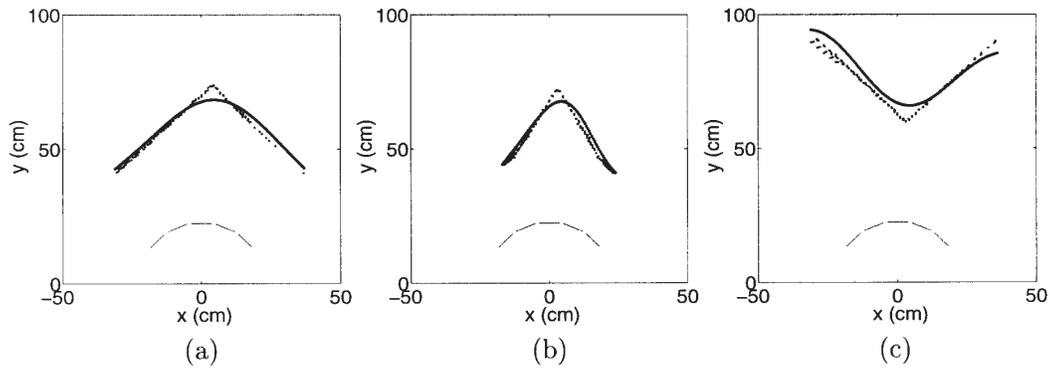


Fig. 20. In all three parts, the result of curve fitting (solid line) was compared to the original surface profile as revealed by the structured-light data. Data were collected by rotating the turret of the robot from -30° to 30° , taking 1° steps. (a) A 90° corner at 80 cm: polynomial fit of order $m = 6$, resulting in $E_1 = 1.90$ cm and $E_2 = 0.217$. (b) A 60° corner at 80 cm: polynomial fit of order $m = 10$, resulting in $E_1 = 1.96$ cm and $E_2 = 0.243$. (c) A 90° edge at 60 cm: polynomial fit of order $m = 8$, resulting in $E_1 = 3.83$ cm and $E_2 = 0.414$.

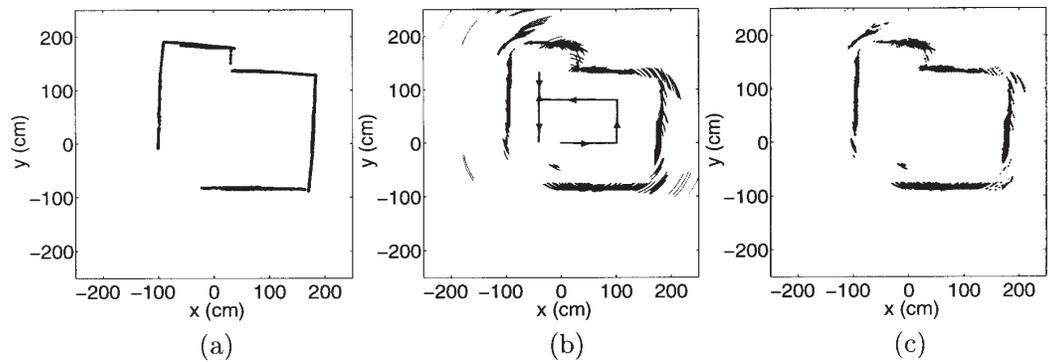


Fig. 21. Structured-light data collected from a room comprised of planar walls, corners, an edge, and a doorway (a); the sonar arc map for the same room and the robot's path (b); and the result of $n = 3$ thinning (c).

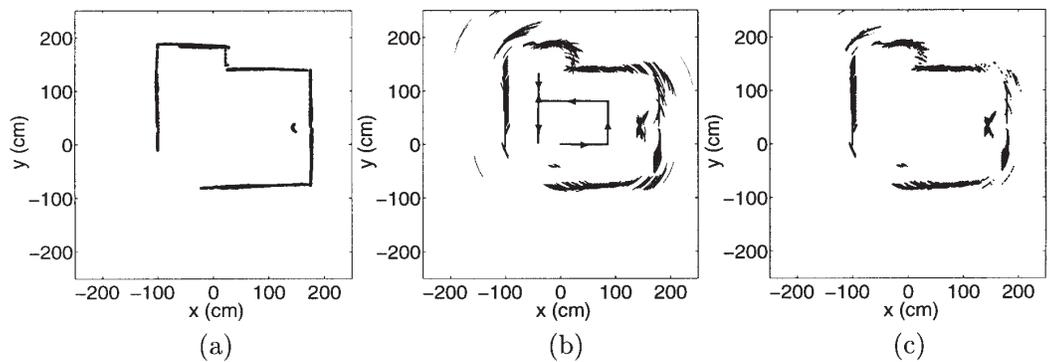


Fig. 22. Structured-light data collected from the room in Figure 21, when a cylinder was added (a); the sonar arc map and the robot's path (b); and the result of $n = 3$ thinning (c).

fusion, the method described in this paper may be preferable in many circumstances, since the structured-light system is much more expensive and complex than sonar sensors.

Although not fully reported here, a detailed quantitative study of the performance of different morphological operations and the dependence of the error on surface curvature, spatial frequency, distance, and sensor beamwidth can be found in the work of Başkent (1998).

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 nonellipsoidal arcs.

Acknowledgments

This work was partially supported by TÜBİTAK under grants EEEAG-92, EEEAG-116, and 197E051. The experiments were performed at Bilkent University Robotics and Sensing Laboratory. The authors would like to thank the anonymous reviewers for their comments and suggestions.

References

- Auran, P. G., and Silven, O. 1996. Underwater sonar range sensing and 3-D image formation. *Control Eng. Practice* 4(3):393–400.
- Ayrulu, B., and Barshan, B. 1998. Identification of target primitives with multiple decision-making sonars using evidential reasoning. *Intl. J. Robot. Res.* 17(6):598–623.
- Barshan, B., and Kuc, R. 1990. Differentiating sonar reflections from corners and planes by employing an intelligent sensor. *IEEE Trans. Pattern Analysis Machine Intell.* 12(6):560–569.
- Başkent, D. 1998. Surface profile determination from multiple sonar data using morphological processing. Master's thesis, Department of Electrical Engineering, Bilkent University, Ankara, Turkey.
- Beckerman, M., and Oblow, E. 1990. Treatment of systematic errors in the processing of wide-angle sonar sensor data for robotic navigation. *IEEE Trans. Robot. Automat.* 6(2):137–145.
- Borenstein, J., Everett, H. R., and Feng, L. 1996. *Navigating Mobile Robots*. Wellesley, MA: A. K. Peters.
- Borenstein, J., and Koren, Y. 1991. The vector-field histogram: Fast obstacle avoidance for mobile robots. *IEEE Trans. Robot. Automat.* 7(3):278–288.
- Bozma, Ö. I. 1992. A physical model-based approach to analysis of environments using sonar. PhD thesis, Department of Electrical Engineering, Yale University, New Haven, CT.
- Brown, M. K. 1986. The extraction of curved-surface features with generic range sensors. *Intl. J. Robot. Res.* 5(1):3–18.
- Cox, I. J. 1991. Blanche: An experiment in guidance and navigation of an autonomous robot vehicle. *IEEE Trans. Robot. Automat.* 7(2):193–204.
- Crowley, J. L. 1985. Navigation for an intelligent mobile robot. *IEEE Trans. Robot. Automat.* 1(1):31–41.
- Crowley, J. L. 1989 (May 14–19, Scottsdale, AZ). World modeling and position estimation for a mobile robot using ultrasonic ranging. *Proc. of the IEEE Intl. Conf. on Robot. and Automat.* Piscataway, NJ: IEEE, pp. 674–681.
- Dougherty, E. R. 1992. *An Introduction to Morphological Image Processing*. Bellingham, WA: SPIE Optical Engineering Press.
- Drumheller, M. 1987. Mobile-robot localization using sonar. *IEEE Trans. Pattern Analysis Machine Intell.* 9(2):325–332.
- Elfes, A. 1987. Sonar based real-world mapping and navigation. *IEEE Trans. Robot. Automat.* 3(3):249–265.
- Elfes, A. 1990 (July 3–6, Tsuchiura, Ibaraki, Japan). Strategies for active robot perception using a probabilistic spatial model. *Proc. of the IEEE/RSJ Intl. Workshop on Intell. Robots and Sys.* Piscataway, NJ: IEEE.
- Elfes, A., and Matthies, L. 1987 (Dec. 9–11, Los Angeles, CA). Sensor integration for robot navigation: Combining sonar and stereo range data in a grid-based representation. *Proc. of the IEEE Intl. Conf. on Decision and Control*. Piscataway, NJ: IEEE, pp. 1802–1807.
- Everett, H. R. 1995. *Sensors for Mobile Robots, Theory and Application*. Wellesley, MA: A. K. Peters.
- Gasos, J., and Martin, A. 1996. A fuzzy approach to build sonar maps for mobile robots. *Comp. Industry* 32(2):151–167.
- Gilbreath, G., and Everett, H. 1988 (Boston, MA). Path planning and collision avoidance for an indoor security robot. *Proc. of the SPIE Conf. on Mobile Robots III*, vol. 1007. Bellingham, WA: SPIE, pp. 19–27.
- Grimson, W. E., and Lozano-Pérez, T. 1984. Model-based recognition and localization from sparse range or tactile data. *Intl. J. Robot. Res.* 3(3):3–35.
- Hauptmann, P. 1993. *Sensors: Principles and Applications*. Englewood Cliffs, NJ: Prentice-Hall.
- Hong, M. L., and Kleeman, L. 1997a. Ultrasonic classification and location of 3-D room features using maximum-likelihood estimation, I. *Robotica* 15(5):483–491.
- Hong, M. L., and Kleeman, L. 1997b. Ultrasonic classification and location of 3-D room features using maximum-likelihood estimation, II. *Robotica* 15(6):645–652.
- Kleeman, L., and Kuc, R. 1995. Mobile robot sonar for target localization and classification. *Intl. J. Robot. Res.* 14(4):295–318.

- Kuc, R. 1997. Biomimetic sonar recognizes objects using binaural information. *J. Acoust. Soc. America* 102(2):689–696.
- Kuc, R., and Siegel, M. W. 1987. Physically-based simulation model for acoustic-sensor robot navigation. *IEEE Trans. Pattern Analysis Machine Intell.* 9(6):766–778.
- Kurz, A. 1996. Constructing maps for mobile-robot navigation based on ultrasonic range data. *IEEE Trans. Sys. Man Cybernet.*, part B. 26(2):233–242.
- Lancaster, P., and Salkauskas, K. 1986. *Curve and Surface Fitting: an Introduction*. London: Academic Press.
- Leonard, J. J. 1990 (Michaelmas Term). Directed sonar sensing for mobile-robot navigation. PhD thesis, Department of Engineering Science, University of Oxford, Oxford, UK.
- Leonard, J. J., and Durrant-Whyte, H. F. 1991. Mobile-robot localization by tracking geometric beacons. *IEEE Trans. Robot. Automat.* 7(3):376–382.
- Leonard, J. J., Durrant-Whyte, H. F., and Cox, I. J. 1992. Dynamic map-building for an autonomous mobile robot. *Intl. J. Robot. Res.* 11(4):286–298.
- Low, A. 1991. *Introductory Computer Vision and Image Processing*. New York: McGraw-Hill.
- Moravec, H. P. 1989. Sensor fusion in certainty grids for mobile robots. In *Sensor Devices and Systems for Robotics*, NATO ASI Series. Berlin: Springer-Verlag, pp. 253–276.
- Moravec, H. P., and Elfes, A. 1986 (April 7–10, San Francisco, CA). High-resolution maps from wide-angle sonar. *Proc. of the IEEE Intl. Conf. on Robot. and Automat.* Piscataway, NJ: IEEE, pp. 116–121.
- Myler, H. R., and Weeks, A. R. 1993. *Computer Imaging Recipes in C*. Englewood Cliffs, NJ: PTR Prentice-Hall.
- Nomadic Technologies, Inc. 1997. *Nomad200 Manual*. Mountain View, CA.
- Oriolo, G., Ulivi, G., and Vendittelli, M. 1997. Fuzzy maps: A new tool for mobile perception and planning. *J. Robot. Sys.* 14(3):179–197.
- Oriolo, G., Ulivi, G., and Vendittelli, M. 1998. Real-time map building and navigation for autonomous robots in unknown environments. *IEEE Trans. Sys. Man Cybernet.* 28(3):316–333.
- Pagac, D., Nebot, E. M., and Durrant-Whyte, H. F. 1998. An evidential approach to map-building for autonomous vehicles. *IEEE Trans. Robot. Automat.* 14(4):623–629.
- Peremans, H., Audenaert, K., and Van Campenhout, J. M. 1993. A high-resolution sensor based on triaural perception. *IEEE Trans. Robot. Automat.* 9(1):36–48.
- Pitas, I. 1993. *Digital Image Processing Algorithms*. Hemel Hempstead, UK: Prentice-Hall International.
- Polaroid Corporation (Ultrasonic Components Group). 1997. *Polaroid Manual*. Cambridge, MA.
- Raschke, U., and Borenstein, J. 1990 (May 13–18, Cincinnati, OH). A comparison of grid-type map-building techniques by index performance. *Proc. of the IEEE Intl. Conf. on Robot. and Automat.* Piscataway, NJ: IEEE, pp. 1828–1832.
- Skolnik, M. I. 1985. *Introduction to Radar Systems*. New York: McGraw-Hill.
- Wallner, F., and Dillman, R. 1995. Real-time map refinement by use of sonar and active stereo vision. *Robot. Autonomous Sys.* 16(1):47–56.
- Watanabe, S., and Yoneyama, M. 1992. An ultrasonic visual sensor for three-dimensional object recognition using neural networks. *IEEE Trans. Robot. and Automat.* 8(2):240–249.
- Zelinsky, A. 1988 (Nov.). Environment mapping with a mobile robot using sonar. *Proc. of the Australian Joint Art. Intell. Conf., AI'88*, pp. 373–388.