

The International Journal of Robotics Research

<http://ijr.sagepub.com>

Directional Processing of Ultrasonic Arc Maps and its Comparison with Existing Techniques

Billur Barshan

The International Journal of Robotics Research 2007; 26; 797

DOI: 10.1177/0278364907078888

The online version of this article can be found at:
<http://ijr.sagepub.com/cgi/content/abstract/26/8/797>

Published by:

 SAGE Publications

<http://www.sagepublications.com>

On behalf of:



Multimedia Archives

Additional services and information for *The International Journal of Robotics Research* can be found at:

Email Alerts: <http://ijr.sagepub.com/cgi/alerts>

Subscriptions: <http://ijr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations (this article cites 36 articles hosted on the
SAGE Journals Online and HighWire Press platforms):
<http://ijr.sagepub.com/cgi/content/abstract/26/8/797#BIBL>

Billur Barshan

Department of Electrical Engineering
Bilkent University
TR-06800 Bilkent, Ankara, Turkey
billur@ee.bilkent.edu.tr

Directional Processing of Ultrasonic Arc Maps and its Comparison with Existing Techniques

Abstract

A new technique for processing ultrasonic arc maps is proposed and compared to six existing techniques for map-building purposes. These techniques are simple point marking along the line-of-sight, voting and thresholding, morphological processing, Bayesian update scheme for occupancy grids, arc-transversal median algorithm, and triangulation-based fusion. The directional maximum technique, newly proposed in this paper, employs directional processing to extract the map of the environment from ultrasonic arc maps. It aims at overcoming the intrinsic angular uncertainty of ultrasonic sensors in map building, as well as eliminating noise and cross-talk related misreadings. The compared techniques are implemented with a wall-following motion-planning scheme for ground coverage. The comparison is based on experimental data and three complementary error criteria: mean absolute error, correct detection rate for full and empty regions, and computational cost in terms of CPU time. The directional maximum technique offers a very good compromise between mean absolute error and correct detection rate, with a processing time less than one-tenth of a second. Compared to the existing techniques, the directional maximum method is also superior in range accuracy and in eliminating artifacts, resulting in the best overall performance. The results indicate several trade-offs in the choice of ultrasonic arc-map processing techniques.

KEY WORDS—sonars, range sensing, wheeled robots, sensing and perception

List of Abbreviations

TOF:	time-of-flight
LOS:	line-of-sight
DOI:	direction of interest

UAM:	ultrasonic arc map
PM:	point marking
VT:	voting and thresholding
DM:	directional maximum
MP:	morphological processing
BU:	Bayesian update scheme for occupancy grids
ATM-org:	arc transversal median—original version
ATM-mod:	arc transversal median—modified version
TBF:	triangulation-based fusion
MAE:	mean absolute error
CDR:	correct detection rate
CDR _O :	overall correct detection rate
CDR _F :	correct detection rate for full regions
CDR _E :	correct detection rate for empty regions

1. Introduction

Sensing and becoming aware of their environment is an essential feature of intelligent mobile robots. This awareness can be accomplished by utilizing simple sensors and processing the acquired sensory data according to the perceptive needs of the robot, such as path-planning, navigation, obstacle avoidance, map building, and localization. Due to the limited resources of autonomous systems, the available resources need to be exploited as much as possible. For this reason, it makes sense to first exploit the potential of simple and inexpensive sensor systems to extract information about the environment as much as possible before more expensive sensing modalities with higher resolution and higher resource requirements (such as computing power) are considered for a given task. Therefore, one of the aims of this study is to explore the limits of simple and low-cost ultrasonic sensors in map building through improvements in processing the raw ultrasonic data. Ultrasonic sensors

have been widely used in robotic applications due to their accurate range measurements, robustness, low cost, and simple hardware interface. We believe that ultrasonic sensing, when coupled with intelligent processing, is a useful alternative to more complex laser and camera systems. Furthermore, it may not be possible to use laser and camera systems in some environments due to surface characteristics or insufficient ambient light. Despite their advantages, the frequency range at which air-borne ultrasonic transducers operate is associated with a large beamwidth that results in low angular resolution and uncertainty in the location of the echo-producing object. Thus, having an intrinsic uncertainty of the actual angular direction of the range measurement and being prone to various phenomena such as multiple and higher-order reflections and crosstalk between transducers, a considerable amount of modeling, processing, and interpretation of ultrasonic data is necessary.

For autonomous operation in static but unknown environments, where a prior map of the workspace is not available, or in dynamically changing environments, the robot needs to build an accurate map of its surroundings by using the sensory data it acquires. Therefore, the selection of an appropriate map-building scheme is an important issue. The different geometric approaches in map building primarily fall into two categories: *feature-based* and *grid-based*. Feature-based approaches are based on extracting the geometry of the environment from sensor data as the first step in data interpretation (e.g. edge detection, straight-line or curve fitting to obstacle boundaries) (Crowley 1985; Drumheller 1987; Grimson and Lozano-Perez 1984; Cox 1991). Most commonly, the environment is modeled in terms of straight lines. 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.

Many researchers have reported the extraction of line segments from ultrasonic data as being difficult and brittle (Leonard and Durrant-Whyte 1992). Straight lines obtained from time-of-flight measurements do not necessarily match or align with the world model, and may yield many erroneous line segments. Improving the algorithms for detecting line segments and including heuristics does not really solve the problem. Leonard and Durrant-Whyte (1992) advocate another feature-based representation using *regions of constant depth* (RCDs) as features, extracted directly from raw ultrasonic range readings. RCDs are circular arcs which are natural features of ultrasonic data from specularly reflecting surfaces, first reported in Kuc and Siegel (1987). Approaches based on physical model-based reasoning, where the environment is described in terms of its natural features such as planes, corners, edges, and cylinders, are also considered as feature-based methods (Leonard and Durrant-Whyte 1992; Kuc and Siegel 1987; Barshan and Kuc 1990; Peremans et al. 1993; Kleeman and Kuc 1995; Hong and Kleeman 1997a, 1997b; Wijk and Christensen 2000a, 2000b).

An alternative representation is the *certainty* or *occupancy grid* where the environment is discretized into an array of cells. The difficulty of line segment-based feature extraction was an important factor in the development of the grid concept proposed by Elfes (1987, 1989). 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. 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 ultrasonic data will be observed from a given position and to give account of individual ultrasonic 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 (1991a), Gilbreath and Everett (1988), Zelinsky (1988), and Beckerman and Obrow (1990) have all used variations of the grid-based approach to construct ultrasonic maps for path planning. In Pagac et al. (1998), the problem of constructing and maintaining a 2-D occupancy-grid map of a robot environment is examined using evidential reasoning, where new ultrasonic readings are incorporated in the map using Dempster's rule of combination. Fuzzy logic has also been employed to represent uncertainty in map-building (Gasos and Martin 1996; Oriolo et al. 1997). In Kurz (1996), free-space is partitioned into regions in which a specific situation can be recognized based on ultrasonic ranging. These regions are then attached to graph nodes generating a map of the environment in the form of a graph representation. The use of grids has been extended to 3-D ultrasonic range sensing underwater in Auran and Silven (1996).

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

The different approaches to motion planning for ground coverage in map building may be categorized into three: in the first class, the environment is covered by taking pseudo-random steps while avoiding collisions (Bruce and Veloso 2002; Hsu et al. 2002). The second class includes the more systematic wall-following type coverage which can be further divided as simple rule-based (Yata 1998) and potential field-




Fig. 1. (a) Single transducer configuration (monostatic mode) resulting in a *circular* arc, (b) dual transceiver configuration (bistatic mode) resulting in an *elliptical* arc.

based (Yun and Tan 1997) wall following. In the second case, each obstacle/wall creates a potential field and the mobile platform aims at staying on constant potential lines. There also exist different approaches to the wall-following problem that cannot be included under the above two categories (Bemporad 1997; Lee and Recce 1994), for example, those which employ fuzzy-logic based control methods (Tunstel and Jamshidi 1994). The third and most sophisticated coverage scheme involved in map building is based on Voronoi diagrams that can be constructed iteratively (Choset and Burdick 2000; Choset et al. 2000).

The main contribution of this paper is that it provides a valuable comparison between the performances of one newly proposed and six existing techniques for processing ultrasonic arc maps (UAMs) for map-building purposes. We have also modified one of the existing techniques and included it in the comparison. The comparison is based on experimental data and three complementary error criteria. The newly proposed directional maximum technique introduces a sense of direction in data processing. Such directional awareness is proven to be cost-effective and beneficial. The directional maximum technique also offers a very good compromise between two of the error criteria, higher range accuracy, and effective elimination of spurious arcs, resulting in the best overall performance. Some of the existing techniques may also have certain advantages that make them suitable for map building under different constraints and requirements.

This paper is organized as follows: in Section 2, representation of angular uncertainty by ultrasonic arc maps is described. Descriptions of the newly proposed and existing UAM processing techniques are provided in Section 3 together with results from a curved surface example and from acute corners. Three complementary error criteria are defined. In Section 4, a comparison and discussion of the results is provided based on experimental data from indoor environments. Final notes, con-

clusions, and future research directions are provided in Section 5.

2. Representing Angular Uncertainty by UAMs

In this study, simple ultrasonic range sensors are employed that measure the time-of-flight (TOF), which is the round-trip travel time between the transducer and the object. The received echo is usually contaminated by noise and the time at which the reflection is received can be estimated by means of *simple thresholding*, using a constant threshold level. Alternatives to simple thresholding such as adaptive, variable or double thresholding (McMullan et al. 1996) and curve-fitting (Barshan and Kuc 1992) techniques have been proposed. Once the TOF value is estimated, the range is calculated from $r = \frac{ct_0}{2}$, where t_0 represents the TOF and c is the speed of sound.¹

Since most air-borne ultrasonic sensing systems operate below a resonance frequency of 200 kHz, frequencies involved correspond to wavelengths well above several millimeters, and the reflections from typical surfaces in indoor environments are specular, not diffused. Due to this mirror-like reflection, when the transducer is operated in the monostatic mode (Figure 1(a)), an echo can be detected only if the incoming ray is perpendicular to the surface at some point. In the bistatic mode (Figure 1(b)), there should be a path between the transmitter and receiver such that at the point of reflection, the angle of incidence and the angle of reflection made with the surface normal are equal.

Although such devices return accurate range data, typically they cannot provide direct information on the angular position of the point on the surface from which the reflection was

1. $c = 331.4\sqrt{T/273}$ m/s = 343.3 m/s at room temperature ($T = 293$ K).

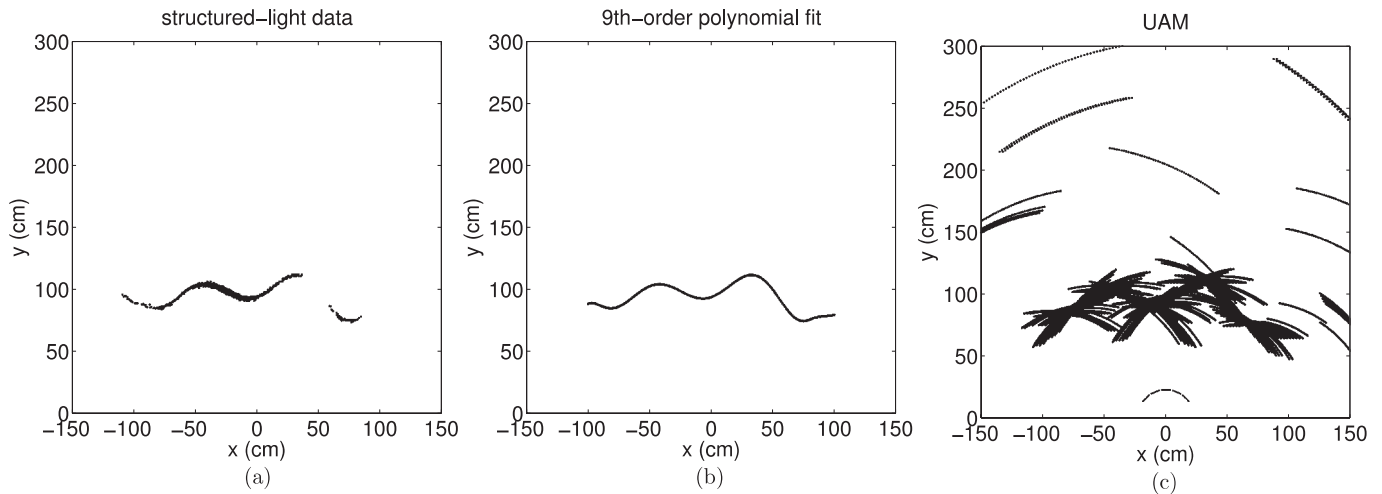


Fig. 2. (a) Structured-light data collected from a curved surface, (b) 9th-order polynomial fit to part (a), (c) the UAM and the sensor positions when the robot's center is at (0,0) cm.

obtained. Most commonly, the large beamwidth of the transducer is accepted as a device limitation that determines the angular resolving power of the system, and the reflection point is assumed to be along the line-of-sight (LOS). According to this naive approach, a simple mark is placed along the LOS of the transducer at the measured range, resulting in inaccurate maps with large angular errors and artifacts. Alternatively, the angular uncertainty in the range measurements has been represented by UAMs (Başkent and Barshan 1999; Barshan and Başkent 2001a, 2001b) that preserve more information. This is done by drawing arcs spanning the beamwidth of the sensor at the measured range, representing the angular uncertainty of the object location and indicating that the echo-producing object can lie anywhere on the arc. Thus, when the same transducer transmits and receives, all that is known is that the reflection point lies on a circular arc of radius r , as illustrated in Figure 1(a). More generally, when one transducer 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 elements (Figure 1(b)). The arcs are tangent to the reflecting surface at the actual point(s) of reflection (please refer to Figures 2(c) and 9(b) and (d) for sample UAMs). Arc segments near the actual reflection points tend to reinforce each other. Arc segments not actually corresponding to any reflections and simply representing the angular uncertainty of the transducers remain more sparse and isolated. Arcs generated by spurious readings, cross-talk, higher-order reflections, and noise also remain sparse and lack reinforcement. These are not enhanced as much as the arcs resulting from the actual surface profile. The proposed directional maximum technique is capable of effectively suppressing these effects, and, although not implemented here, it has the intrinsic ability to process echoes returning from surface features further away than the

nearest (i.e. multiple reflections) informatively. By combining the information inherent in a large number of such arcs, much improved angular resolution is obtained.

The device that is used in this study is the Polaroid 6500 series transducer (Polaroid Corporation 1997) operating at a resonance frequency of $f_o = 49.4$ kHz, corresponding to a wavelength of $\lambda = c/f_o = 6.9$ mm. The half-beamwidth angle of the transducer is $\theta_o = 12.5^\circ$, the transducer aperture radius is $a = 2.0$ cm and $r_{\min} = 5.7$ cm. In the experiments, we use three or five of these transducers, located on an arc of radius of curvature 23 cm, with a center-to-center separation of approximately 9 cm (this corresponds to the configuration of the ultrasonic sensors on the Nomad 200 mobile robot). Each transducer is fired in sequence and all transducers detect the resulting echoes. After a pulse is transmitted, if echoes are detected at all transducers, this corresponds to one circular and two (or four) elliptical arcs. When multiple echoes are detected after a single transmission, only the very first echo is processed, usually coming from the obstacle that is closest to the transducer that was fired.

3. UAM Processing Techniques

In this section, we introduce the directional maximum (DM) technique developed in this work for processing UAMs, and give summarizing descriptions of six existing techniques which are:

- point marking along the LOS (Kuc and Siegel 1987; Moravec and Elfes 1986; Borenstein and Koren 1991b),
- voting and thresholding (Barshan 1999),

- morphological processing (Başkent and Barshan 1998; Barshan and Başkent 2001a, 2001b),
- Bayesian update scheme for occupancy grids (Elfes 1987),
- arc-transversal median algorithm (Choset et al. 2003) and its modified version,
- triangulation-based fusion (Wijk and Christensen 2000a, 2000b; Wijk 2001).

The processing techniques will be demonstrated on examples based on experimental data where the same UAM will be processed by each technique in order to provide a uniform comparison. Each technique is summarized below:

3.1 Point Marking (PM)

This is the simplest approach, mentioned in Section 2, where a mark is placed along the LOS at the measured range. This method produces reasonable estimates for the locations of objects if the arc of the cone is small. This can be the case at higher frequencies of operation where the corresponding sensor beamwidth is small or at nearby ranges. The PM technique is demonstrated on an example first, where ultrasonic data from the curved surface shown in Figure 2(a) are collected. The details of this experiment are given at the beginning of Section 4.2. The resulting UAM is illustrated in Figure 2(c). The result of applying the PM approach to this UAM is given in Figure 3(a). Since every arc is reduced to a single point, this technique cannot eliminate any of the outlying TOF readings. The resulting map is inaccurate with large angular errors and artifacts.

3.2 Voting and Thresholding (VT)

Another simple technique for processing UAMs is a voting scheme where each pixel stores the number of arcs crossing that pixel, resulting in a 2-D array of occupancy counts for the pixels (Barshan 1999). By simply thresholding this array, i.e. zeroing the pixels that have a value lower than a suitably chosen threshold level, artifacts can be eliminated and the map of the surface is extracted.

The result of applying the VT approach to the curved surface example of Figure 2 with a threshold value of 4 is given in Figure 3(b). Note that most of the erroneous readings visible in Figure 2(c) are cleared out. While the profile obtained by using VT has some arc branches and point artifacts that cannot be removed without a higher threshold, the overall performance observed in Figure 3(b) is acceptable. Using a larger threshold value would create gaps in the curved profile.

3.3 Directional Maximum (DM)

This technique is based on the idea that in processing the acquired range data, there is a direction-of-interest (DOI) associated with each detected echo. Ideally, the DOI corresponds to the direction of a perpendicular line drawn from the sensor to the nearest surface from which an echo is detected. However, in practice, due to the angular uncertainty of the object position, the DOI can be approximated as the LOS of the sensor when an echo is detected. Since prior information on the environment is usually unavailable, the DOI needs to be updated while sensory data are being collected and processed on-line. It may also be possible to determine the DOI by post-processing, based on the distribution of the acquired data by choosing it perpendicular to the direction along which the spread of the data collected over a given region is maximum. It should be noted that ideally, the DOI is different for each detected echo and sensor position.

Using the UAM, the number of arcs crossing each pixel is counted and stored, and a suitable threshold value is chosen. Up to this point, the DM method is exactly the same as the VT method. The novelty of the DM method is the processing done along the DOI. Once the DOI for a measurement is determined using a suitable procedure, the UAM is processed along this DOI as follows: The array of pixels along the DOI is inspected and the pixel(s) exceeding the threshold with the maximum count is kept, while the remaining pixels along the DOI are zeroed out. If there exist more than one maxima, the algorithm takes their median (If the number of maxima is odd, the maxima in the middle is taken; if the number is even, one of the two middle maxima is randomly selected.) This way, most of the artifacts of the UAM can be removed.

The result of applying the DM technique to the curved surface example of Figure 2 with a threshold value of 4 is presented in Figure 3(c). The highly populated UAM given in Figure 2(c) is successfully cleaned to give a sufficiently accurate surface profile. Note that the branches and dark regions in Figure 2(c) have been successfully removed. Only a few minor point artifacts remain which can easily be eliminated by polynomial fitting. The DOI in this example is taken to be the positive y direction, perpendicular to the path of the robot. Simulation examples based on this technique can be found in Kurt (2005).

3.4 Morphological Processing (MP)

Morphological processing techniques have been used in pattern recognition applications since they were first introduced (Serra 1982). Opening, closing, dilation, and thinning are the fundamental binary morphological operations. Erosion and pruning are special cases of thinning. Being easy-to-use yet powerful tools in image processing, morphological operators

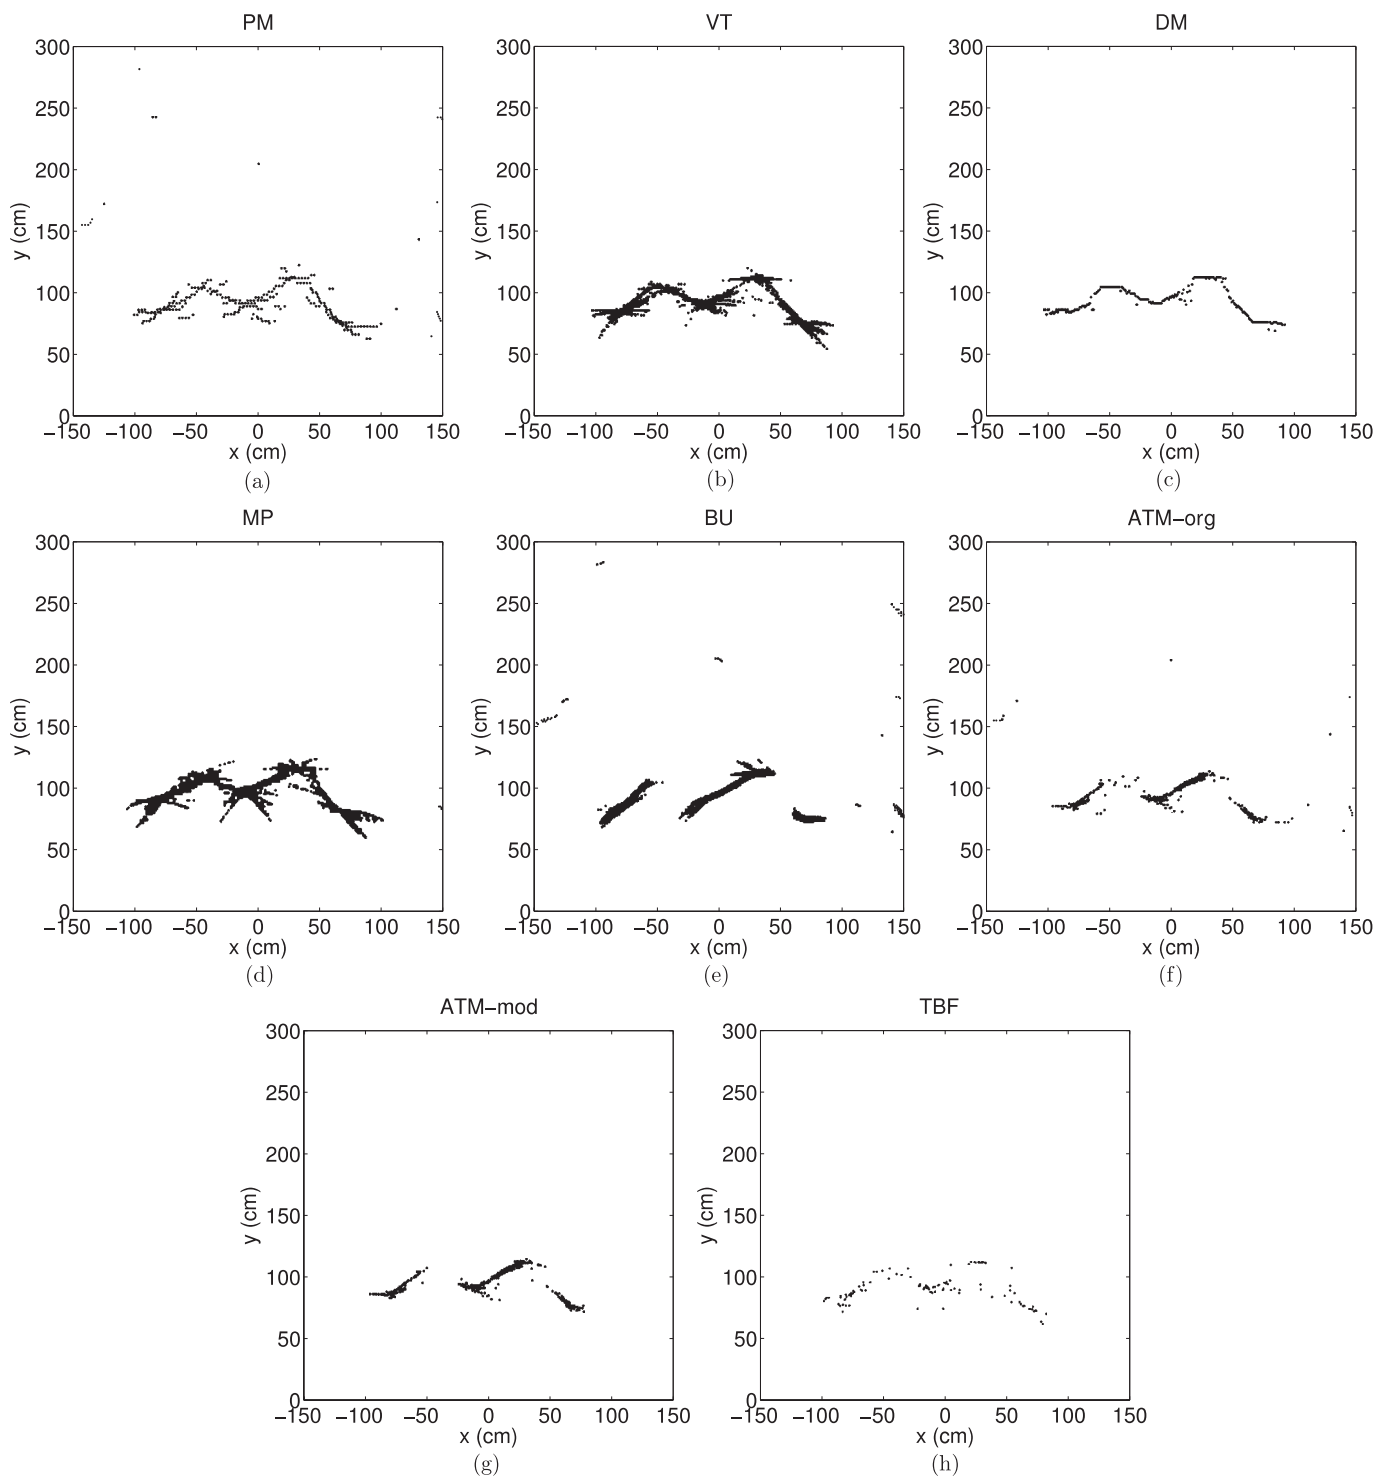


Fig. 3. Results of (a) PM, (b) VT with threshold 4, (c) DM with threshold 4, (d) MP with $m = 8$, (e) BU with threshold 0.999, (f) original ATM, (g) modified ATM, and (h) TBF with $n_t \geq 4$.

have been widely employed in a number of areas including biomedical engineering, machine vision, and remote sens-

ing. The use of these techniques was also extended to gray-scale images (Chen and Dougherty 1994). In addition to con-

ventional images, range data (Verly and Delanoy 1993) and medical ultrasound data (Mojsilovic et al. 1997) have also been processed by morphological operators. A comparison of MP with VT is provided in our earlier work in Barshan and Bařkent (2000) for surfaces of varying roughness.

The processing of UAMs using morphological processing techniques was first proposed by Bařkent and Barshan (1999). This approach exploits neighboring relationships and provides an easy to implement yet effective solution to ultrasonic map building. By applying binary morphological processing operations, one can eliminate the artifacts of the UAM and extract the surface profile. In Bařkent and Barshan (1999), simulation results based on a large number of transducers were presented, where the transducers were configured linearly, circularly or randomly.

In MP, we chose to apply the thinning operation to a given UAM. First, the thinning parameter m is set to a suitable value between 0 and 8 ($m = 1$ corresponds to pruning, $m = 8$ to erosion, and $m = 0$ does not modify the UAM). Only the pixels with value 1 are considered, indicating that there exists at least one arc crossing that pixel. The number of nonzero neighbors of each pixel is counted. If the number of nonzero neighbors of a pixel is less than m , this means that the pixel does not have sufficient support and its value is equated to zero; otherwise, its value remains as 1. The result of applying thinning with parameter $m = 8$ (erosion) to the UAM of the curved surface can be seen in Figure 3(d). Although the outlying artifacts of the raw arc map of Figure 2(c) are cleared, there is a considerable amount of dark regions remaining around the surface profile. Note that the remaining pixels are those with complete support, that is, all 8 of their neighbors have the value 1. Since 8 is the maximum possible value of the thinning parameter m , further cleaning is not possible by increasing the order of thinning.

3.5 Bayesian Update Scheme for Occupancy Grids (BU)

Occupancy grids were first introduced by Elfes, and a Bayesian scheme for updating their probabilities of occupancy and emptiness was proposed in Elfes (1987) and verified by ultrasonic data. Besides occupancy grids, Bayesian methods can be applied to many other algorithms. This technique is included in this paper as an earlier example of related work. Starting with a blank or completely uncertain occupancy grid, each range measurement updates the grid formation in a Bayesian manner. For a certain range measurement, the following sensor characteristics and acquired data are listed for each pixel $P(x, y)$ of the map to be updated:

- r range measurement returned by the ultrasonic sensor
- r_{\min} lower range threshold (near-field limit)
- r_{ϵ} maximum ultrasonic range measurement error

- θ_{\circ} sensor half-beamwidth angle
- $S = (x_s, y_s)$ position of the ultrasonic sensor
- σ distance from $P(x, y)$ to $S = (x_s, y_s)$.

Occupancy probability of the scanned pixels are defined over two distinct probability measures: p_E , probability of emptiness and p_O , probability of occupancy. These probability density functions are defined as follows:

$$p_E(x, y) = p[\text{point } (x, y) \text{ is empty}] = E_r(\sigma) \cdot E_a(\theta) \quad (1)$$

where $E_r(\sigma)$ and $E_a(\theta)$ are respectively the range and angle dependence of the probability density function for emptiness, given by:

$$E_r(\sigma) = \begin{cases} 1 - [(\sigma - r_{\min}) / (r - r_{\epsilon} - r_{\min})]^2 & \text{for } \sigma \in [r_{\min}, r - r_{\epsilon}] \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

and

$$E_a(\theta) = 1 - (\theta / \theta_{\circ})^2, \quad \text{for } \theta \in [-\theta_{\circ}, \theta_{\circ}]. \quad (3)$$

Likewise, p_O is defined as:

$$p_O(x, y) = p[\text{position } (x, y) \text{ is occupied}] = O_r(\sigma) \cdot O_a(\theta) \quad (4)$$

where $O_r(\sigma)$ and $O_a(\theta)$ are respectively the range and angle dependence of the probability density function for occupancy, and defined as:

$$O_r(\sigma) = \begin{cases} 1 - [(\sigma - r) / r_{\epsilon}]^2, & \text{for } \sigma \in [r - r_{\epsilon}, r + r_{\epsilon}] \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

and

$$O_a(\theta) = 1 - (\theta / \theta_{\circ})^2, \quad \text{for } \theta \in [-\theta_{\circ}, \theta_{\circ}]. \quad (6)$$

The p_E and p_O described above are illustrated in Figure 6. The value of r_{ϵ} was taken as 2 cm.

Initially, the map-building process starts with a maximally uncertain map where all the pixel values are set to zero, corresponding to the mid-point of the interval $[-1, 1]$ for pixel values. For each range reading r , p_E values are updated for those pixels within the sensitivity region of the transducer that fall in the range interval $[r_{\min}, r - r_{\epsilon}]$ using (1)–(3). Similarly, p_O values are calculated for those pixels within the sensitivity region that fall in the range interval $[r - r_{\epsilon}, r + r_{\epsilon}]$, using

(4)–(6). The following BU rules are employed to update the existing values in the pixel array:

$$\text{updated value of } p_E(\text{pixel}) \quad (7)$$

$$= p_E(\text{pixel}) + p_E(\text{reading}) - p_E(\text{pixel}) \times p_E(\text{reading})$$

$$\text{updated value of } p_O(\text{pixel}) \quad (8)$$

$$= p_O(\text{pixel}) + p_O(\text{reading}) - p_O(\text{pixel}) \times p_O(\text{reading})$$

The map of the environment is constructed by iteratively updating the contents of each pixel using all the available range measurements comprising the UAM. In the end, p_E and p_O arrays contain modified probability distribution functions whose values vary between -1 and 1 . These are then transformed to the interval $[0, 1]$ and thresholded. If we consider pixels with occupancy probabilities above 0.5 as full, the resulting map very much resembles the UAM for the environment. Therefore, in this study, we used larger threshold values to consider a pixel as full. Values of occupancy probabilities exceeding 0.997 worked very well in many experiments.

By using the technique described above on the example given in Figure 2, the surface profile in Figure 3(e) is obtained. The resultant map contains several large gaps along the profile and a considerable amount of artifacts. Considering the fact that the robot moved from left to right, the gaps are located at positions on the curve that are “shadowed” towards the end of the robot’s path. The probability of occupancy of the corresponding pixels were lowered by the updates in (7) and (8) due to this shadowing effect.

3.6 Arc-Transversal Median (ATM)

The Arc-Transversal Median Algorithm was developed by Choset and co-workers, and requires both extensive bookkeeping and a considerable amount of processing (Choset et al. 2003). The algorithm can be summarized as follows: for each arc in the UAM, the positions of the intersection(s) with other arcs, if they exist, are recorded. For arcs without any intersections, the mid-point of the arc is taken to represent the actual point of reflection (as in the PM method) which corresponds to the intersection point of the arc with the LOS. If the arc has a single intersection, the algorithm uses the intersection point as the location of the reflecting object. For arcs with more intersections, the median of the positions of the intersection points with other arcs is chosen to represent the actual point of reflection. In the work reported in [45], the median operation is applied when an arc has *three or more* intersection points. If

there is an even number of intersections, the algorithm uses the mean of the two middle values (except that arcs with two intersections are ignored). It can be considered as a much improved version of the PM approach, mentioned in Section 2, where a single mark is placed along the LOS at the measured range. When this original ATM algorithm (ATM-org) is applied to the curved surface example, the profile in Figure 3(f) is obtained. Note that apart from a few gaps along the surface profile, since the mid-points of the arcs with no intersections are kept, the algorithm is not very good at eliminating single, outlying arcs due to erroneous readings or higher-order reflections.

To improve the performance of this technique in eliminating the artifacts, we have also implemented a modified version of the algorithm (ATM-mod) where we ignored arcs with no intersections. Furthermore, since we could not see any reason why arcs with two intersections should not be considered for finding the median, we took the mean of the two intersection points. The resulting profile is seen in Figure 3(g). Note that the single point artifacts in part (e) of the same figure are now cleared.

3.7 Triangulation-Based Fusion (TBF)

Another interesting earlier work on ultrasonic sensing is the work of Wijk and Christensen (2000a, 2000b) and Wijk (2001) where the triangulation-based fusion (TBF) method is introduced. The TBF method is primarily developed for accurately detecting and extracting the edge-like features in the environment (vertical edges such as table legs, door posts etc.) as natural landmarks. These landmarks may be used later as reference maps for purposes such as robot pose tracking or localization.

The TBF method assumes that features are of a certain geometrical form (e.g. point-like target or edge) and performs the triangulation accordingly. The TBF method is not suitable for accurately detecting positions of planar walls which are widely found in indoor environments. This is because unlike an edge, the intersection point of arcs from a planar wall, taken at different robot positions, do not necessarily correspond to the range readings obtained at the two positions (Bozma 1992; Leonard and Durrant-Whyte 1992). Therefore, the triangulation equations in Wijk and Christensen (2000a) are not suitable for localizing planar features.

Basically, TBF and the other methods described above aim at finding the intersection points of the arcs in the UAM (i.e. extracting the darker features in the UAM). The main difference is that TBF realizes this by using a geometric model suitable for edge-like features, whereas the previously described techniques do this by dividing the environment into grids and processing the information in each grid.

Another difference between TBF and the other techniques we compare is that TBF considers a sliding window of sonar scans where the number of rows of the sliding window corresponds to the number of sonars fired, and the number of

columns corresponds to the number of most recent sonar scans to be processed by the algorithm. TBF is focused on detection of edge-like features at nearby ranges of less than 5 m. The other methods included in our comparison consider all of the arcs in the UAM corresponding to all ranges, and are suitable for detecting all types of features.

We have implemented the TBF method as described in Wijk and Christensen (2000a). We have used exactly the same parameters as those given in Wijk and Christensen (2000a) except that we used sliding windows of dimension 3×20 or 5×20 instead of 16×10 as reported in Wijk and Christensen (2000a). This is because our sonar data were collected every 2.5 cm (which is half of that stated in Wijk and Christensen (2000a)) using the front three or front five sensors of our mobile robot.

An important parameter of the TBF algorithm is the number n_t of successful triangulations performed for a given range reading. If the maximum deviation between successful triangulations associated with a range reading is larger than a preset threshold, the triangulation point is classified as belonging to an object which is not well-represented by an edge, for instance a wall. At the end of the algorithm, features which are not edge-like may or may not be shown in the map. In our comparison, if we extract and show only the edge-like features, the resulting map would contain many gaps and result in a very low CDR_F , defined below. To improve the performance of this technique, we have also considered including all the successful triangulation points. In addition, we have varied the threshold value for n_t from 1 to 8 (see Tables 5 and 6), considered sliding window column sizes of 10 and 20, and presented the best results possible.

The result of applying the TBF algorithm to the UAM of the curved surface can be seen in Figure 3(h). In this figure, we have shown all successful triangulation points where at least four triangulations were performed. Note that since the TBF is suitable for point-like features (which have very high curvature), it does not perform very well for surfaces with much lower curvature as in this example.

As another example, experimental results for corners of wedge angles 90° and 60° are shown in Figures 4 and 5 for the techniques described above. In this example, data were collected by rotating five ultrasonic sensors located on an arc from -30° to 30° in 1° steps. Due to the rotational motion in acquiring the data, there were many overlapping arcs in this example. A visual comparison indicates that the results for the 60° corner are worse than the right angle corner in general. The MP has the worse performance since it produces a gap at the corner point in both cases. We also experimented with a 30° acute corner whose results were worse than the 60° acute corner. The acute corner is a difficult feature to identify, and more so as its wedge angle decreases.

3.8 Error Criteria

We defined and used three complementary error criteria to evaluate and compare the techniques described above. The first and maybe the more important one is the mean absolute error (MAE) between the true and extracted features of the environment. To calculate the MAE, the distance of each nonzero point in the resulting map to the nearest point in the structured-light profile (taken as reference) is found and these distances are averaged over all the nonzero points in the constructed map. The second error criterion is the correct detection rate (CDR). In evaluating the CDR, first, the full and empty directions of the actual map are considered separately. Then, combining these, an overall CDR is calculated, given by the equations:

$$CDR_F = \frac{\text{correctly detected full}}{\text{total number of full}} \quad (9)$$

$$CDR_E = \frac{\text{correctly detected empty}}{\text{total number of empty}} \quad (10)$$

$$CDR_O = \frac{\text{correctly detected full and empty}}{\text{total number of full and empty}} \quad (11)$$

We have evaluated the CDR using the concept of DOI described above, instead of evaluating it pixel by pixel. For a given DOI, if the actual profile revealed by the structured-light data indicates that that direction is full, and the resultant map agrees with this, correct detection of a full DOI occurs. Similarly, if the structured-light data indicates that that direction is empty, and the resultant map agrees so, correct detection of an empty DOI occurs.

The final criterion stands for the computational cost of the techniques in terms of the CPU time t_{CPU} . These error criteria, one standing for quality, the second for quantity, and the last for implementability are more meaningful when considered together.

The results obtained from the curved surface example (Figure 2) according to these criteria are summarized in Table 1. According to the results, ATM-mod produces the smallest error but also has a low CDR, indicating that this technique creates sparsely filled yet quite accurate maps. Although DM and ATM-mod are comparable in error (MAE), the overall CDR for DM is about 21% higher, making DM more advantageous. In terms of MAE, DM is followed by VT, TBF, original ATM (ATM-org), MP, BU, and PM. The MP technique with $m = 8$ produces the highest overall CDR, very closely followed by VT and DM, and then by MP with $m = 7$, PM, ATM, BU, and TBF. However, as seen from the CDR_F and CDR_E in the fourth and fifth columns of the table, MP has a tendency for over-detecting a given direction as full, and under-detecting as empty. Variable parameters of each technique are selected such that CDR_F and CDR_E are as close to each other as possible (see Tables 2–6). In terms of CPU time, PM is the

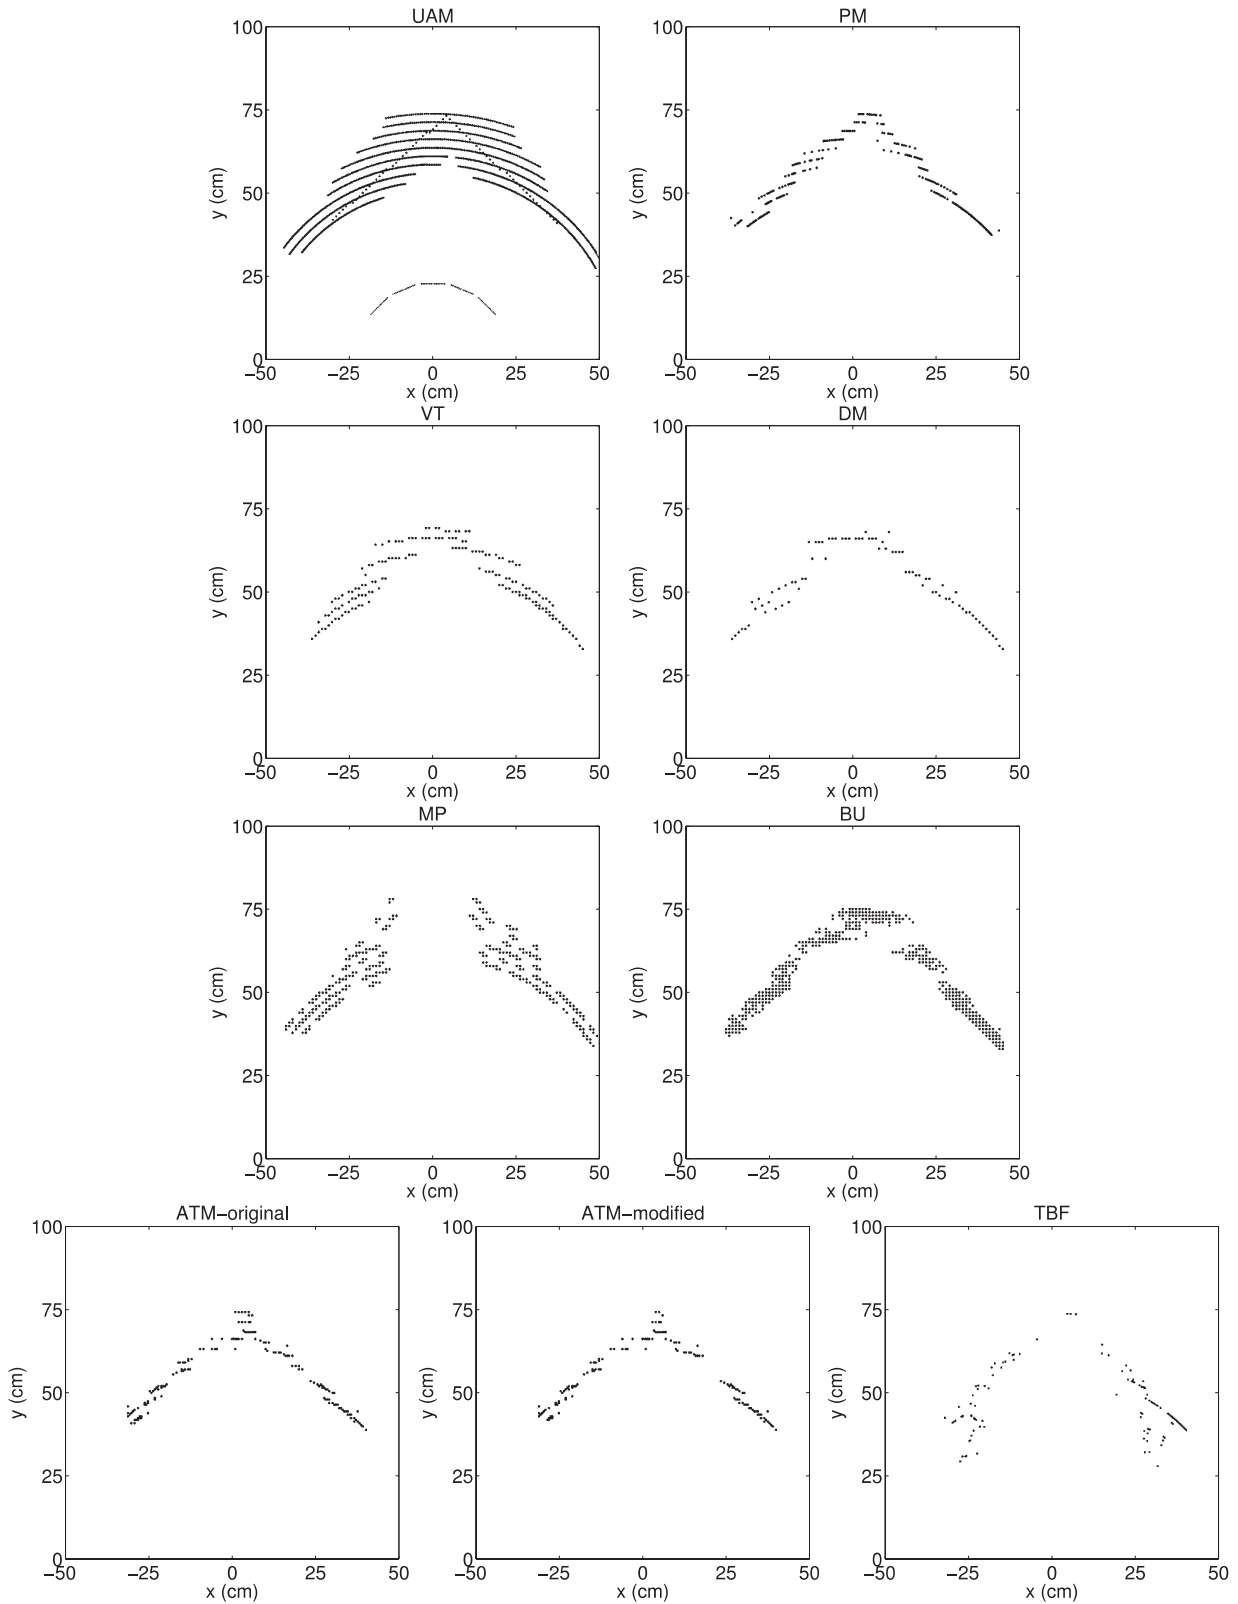


Fig. 4. Top left: 90° corner, its UAM, and the positions of the 5 sensors. The following plots give the results of PM, VT and DM with threshold 25, MP with $m = 4$, BU with threshold 0.999999, ATM-org, ATM-mod, and TBF for $n_t \geq 3$.

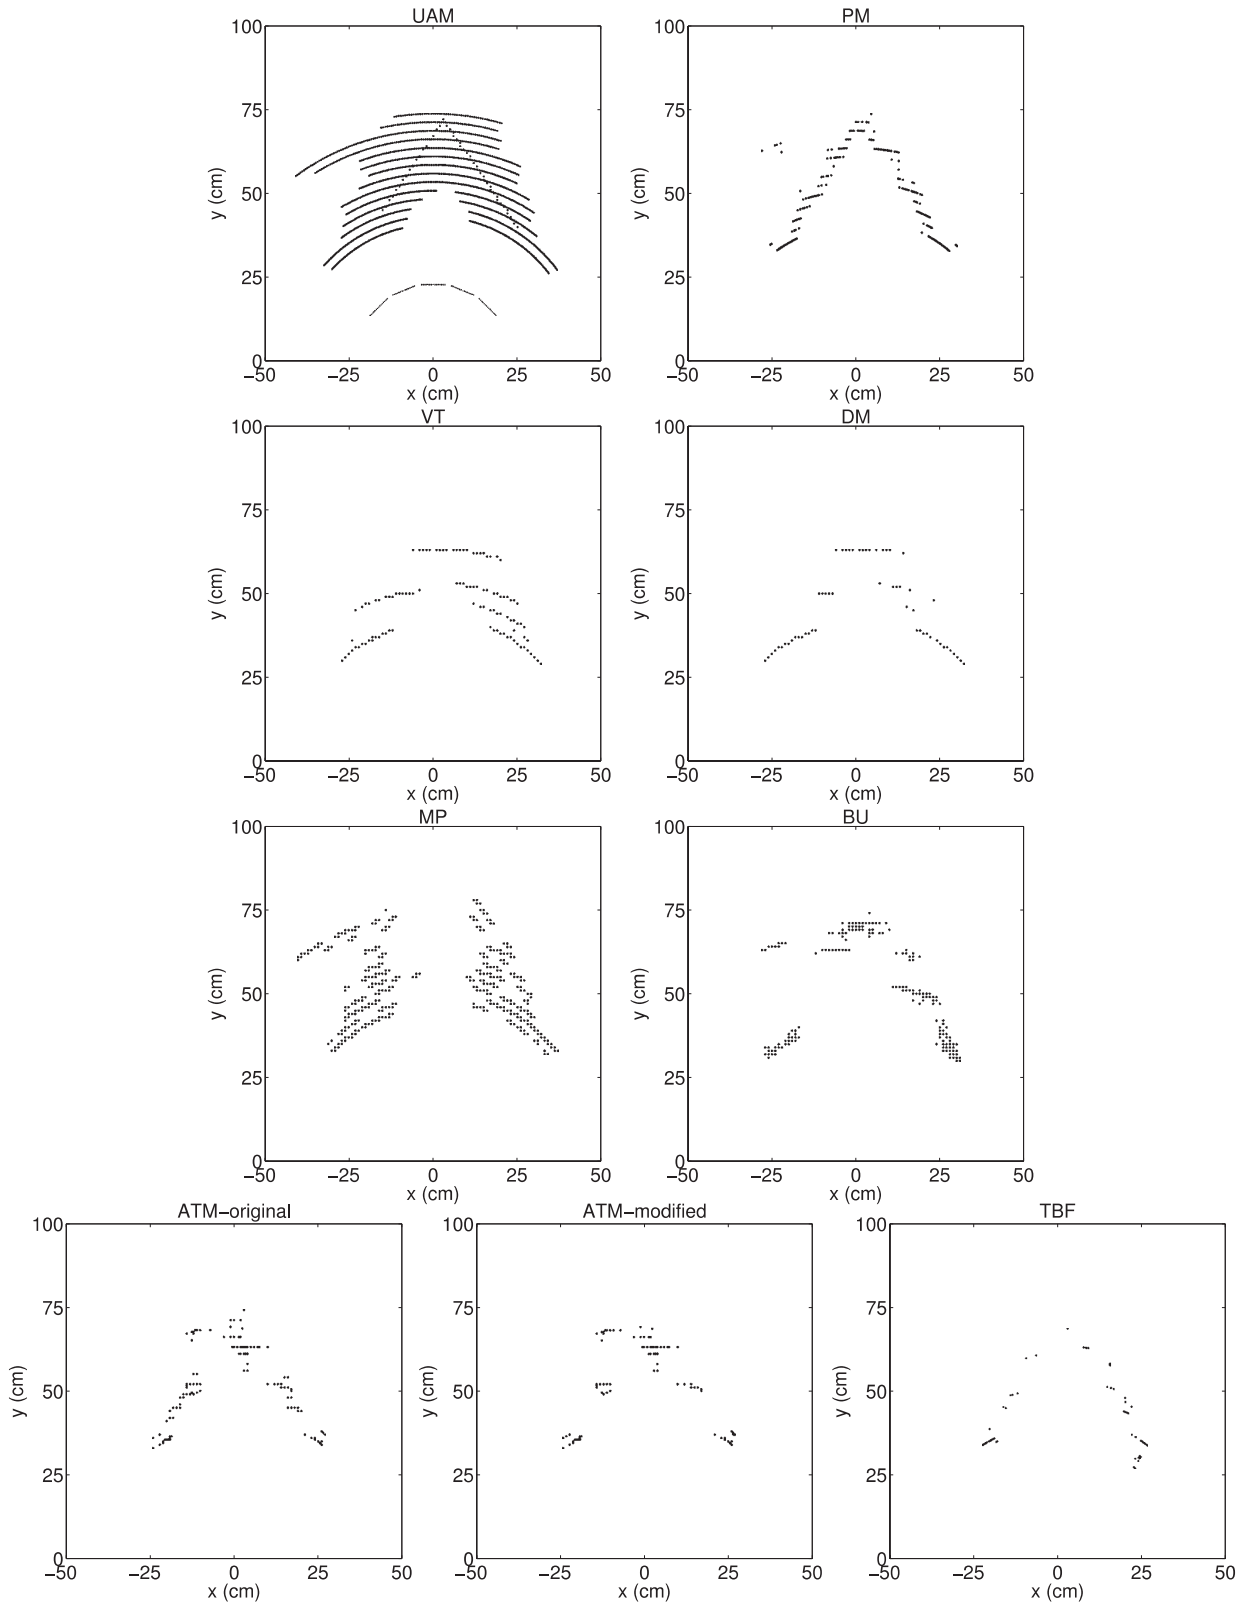


Fig. 5. Top left: 60° corner, its UAM, and the positions of the 5 sensors. The following plots give the results of PM, VT and DM with threshold 25, MP with $m = 4$, BU with threshold 0.999999, ATM-org, ATM-mod, and TBF for $n_t \geq 3$.

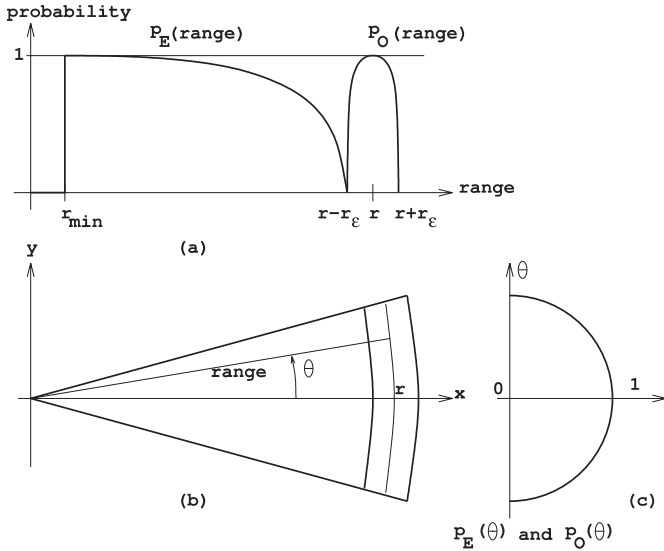


Fig. 6. (a) Range dependence, (b) x and y dependence, and (c) θ dependence of p_O and p_E (adopted from [14]).

Table 1. Experimental Results for the Curved Surface given in Figure 2(a) for the Different Techniques

Method	MAE (cm)	CDR _O	CDR _F	CDR _E	t_{CPU} (s)
PM	11.10	0.817	0.796	0.860	0.005
VT (thld=4)	3.46	0.960	0.955	0.970	0.035
DM (thld=4)	2.23	0.960	0.955	0.970	0.031
MP ($m = 7$)	5.87	0.897	1.000	0.690	0.030
MP ($m = 8$)	4.62	0.967	0.995	0.910	0.031
BU (thld=0.999)	10.23	0.734	0.761	0.680	0.352
ATM-org	4.33	0.748	0.682	0.880	0.310
ATM-mod	1.81	0.754	0.632	1.000	0.304
TBF ($n_t \geq 4$)	4.28	0.595	0.393	1.000	0.010

simplest technique, and therefore has the lowest CPU time, followed by TBF, and then by VT, DM, and MP about equally. The CPU times of BU and ATM are considerably larger, but still implementable in real time. The latter two techniques that take longer for processing also produce lower overall CDRs. However, the lowest overall CDR is produced by TBF, even if we include all the successful triangulation points (which is the case given in Table 1).

To summarize, we have introduced the DM technique and demonstrated it through some examples based on experimental data. We have also reviewed six existing techniques for UAM processing, modified one of them, and compared them with DM on the same example, using three complementary error criteria.

Table 2. Experimental Results for the Indoor Environment given in Figure 9(a) for the VT and DM Techniques for Varying Threshold Values

Thld	VT		DM			VT		DM	
	MAE (cm)	MAE (cm)	CDR _O	CDR _F	CDR _E	t_{CPU} (s)	t_{CPU} (s)	t_{CPU} (s)	t_{CPU} (s)
1	14.76	5.59	0.906	1.000	0.348	0.082	0.082		
2	7.02	3.10	0.946	1.000	0.627	0.083	0.083		
3	4.58	2.75	0.946	0.977	0.764	0.078	0.081		
4	2.57	2.36	0.938	0.951	0.863	0.080	0.085		
5	1.87	1.96	0.858	0.847	0.926	0.074	0.083		
6	1.80	1.87	0.834	0.813	0.957	0.078	0.081		
7	1.54	1.63	0.778	0.746	0.969	0.075	0.083		
8	1.53	1.55	0.753	0.717	0.969	0.077	0.083		
9	1.48	1.47	0.704	0.659	0.969	0.074	0.083		
10	1.44	1.45	0.657	0.604	0.969	0.075	0.082		

Table 3. Experimental Results for the MP Technique for the Curved Surface in Figure 2(a) for Different Thinning Parameter Values

m	MAE (cm)	CDR _O	CDR _F	CDR _E	t_{CPU} (s)
0	20.18	0.671	1.000	0.010	0.031
1	17.85	0.671	1.000	0.010	0.031
2	12.69	0.671	1.000	0.010	0.030
3	11.11	0.698	1.000	0.090	0.030
4	9.91	0.748	1.000	0.240	0.030
5	8.63	0.771	1.000	0.310	0.031
6	7.19	0.834	1.000	0.500	0.031
7	5.87	0.897	1.000	0.690	0.031
8	4.62	0.967	0.995	0.910	0.031

4. Experimental Verification

In this section, the techniques briefly described and demonstrated above are experimentally verified using the sensor systems on the Nomad 200 mobile robot in our laboratory.

4.1 System Description

The Nomad 200 mobile robot, shown in Figure 7, is used in the experiments. It is an integrated mobile robot including tactile, infrared, ultrasonic, and structured-light sensing

Table 4. Experimental Results for the BU Technique for the Curved Surface in Figure 2(a) for Different Threshold Values for Occupancy Probabilities

Thld	MAE (cm)	CDR _O	CDR _F	CDR _E	$t_{CPU}(s)$
0.990	13.58	0.694	0.846	0.390	0.353
0.991	13.69	0.691	0.841	0.390	0.350
0.992	13.51	0.701	0.841	0.420	0.352
0.993	12.61	0.714	0.841	0.460	0.350
0.994	12.65	0.718	0.836	0.480	0.350
0.995	12.13	0.721	0.831	0.500	0.350
0.996	12.24	0.718	0.821	0.510	0.351
0.997	11.34	0.721	0.796	0.570	0.351
0.998	11.02	0.718	0.776	0.600	0.350
0.999	10.23	0.734	0.761	0.680	0.350
1.000	16.54	0.508	0.313	0.900	0.350

Table 5. Experimental Results for the TBF Technique for the Indoor Environment in Figure 9(a) for Different Threshold Values for n_t . All Successful Triangulation Points where $n_t \geq \text{thld}$ are Included

Thld	MAE (cm)	CDR _O	CDR _F	CDR _E	$t_{CPU}(s)$
1	6.63	0.517	0.443	0.957	0.01
2	2.82	0.427	0.334	0.981	0.01
3	2.44	0.395	0.296	0.981	0.01
4	2.34	0.368	0.262	0.994	0.01
5	2.29	0.329	0.216	0.994	0.01
6	2.36	0.297	0.179	0.994	0.01
7	2.46	0.272	0.149	1.000	0.01
8	1.77	0.244	0.117	1.000	0.01

systems, with dimensions 76.2 cm (height) and 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 Nomad 200 are 60 cm/s and 60°/s respectively. The Nomad 200 is powered by a 840 Wh removable battery package (Nomadic Technologies 1997).

Nomad 200 has an on-board 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.

Table 6. Experimental Results for the TBF Technique for the Indoor Environment in Figure 9(a) for Different Threshold values for n_t . Only those Triangulation Points that Correspond to Edge-like Features where $n_t \geq \text{thld}$ are Included

Thld	MAE (cm)	CDR _O	CDR _F	CDR _E	$t_{CPU}(s)$
1	9.68	0.382	0.283	0.963	0.01
2	1.59	0.261	0.139	0.988	0.01
3	1.02	0.238	0.111	0.988	0.01
4	1.09	0.219	0.088	0.994	0.01
5	1.04	0.191	0.056	0.994	0.01
6	1.06	0.176	0.038	0.994	0.01
7	1.00	0.161	0.019	1.000	0.01
8	0.73	0.155	0.013	1.000	0.01




Fig. 7. Nomad 200 mobile robot. The ring of ultrasonic sensors can be seen close to the top rim of the turret, and the structured-light system is seen pointing rightwards on top.




Fig. 8. An example curved surface and a meter stick.

We next give a brief description of the two sensor modules used in the experiments:

The Sensus 200 Ultrasonic Ranging System consists of 16 ultrasonic transducers in a circular configuration, which can yield range information from 15 cm to 10.7 m with $\pm 1\%$ accuracy. The sensors are Polaroid 6500 series transducers (Polaroid Corporation 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 in order 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, ultrasonic and structured-light data are collected from three different environments constructed in our laboratory. Both systems are rigidly fixed to the turret of the mobile robot so that the correspondence between them is never altered. The structured-light system is much more expensive and complex, requiring higher-power and sufficient ambient light for operation. Furthermore, this mode of sensing does not work in all environments, such as those with dark-colored upholstery and glass. Since it reveals a very accurate surface

profile, the profile detected by this system is taken as an absolute reference in the experimental calculation of the MAE and CDR using ultrasonic data.

To prevent any cross-talk between consecutive pulses, the ultrasonic transducers should be fired with at least 62 ms intervals since the maximum range of operation of Polaroid transducers is 10.7 m. In the experiments, the ultrasonic transducers are fired at 40 ms intervals. This prevents much of the cross-talk, and in the few cases where erroneous readings are obtained due to cross-talk, these false readings can be eliminated by several of the processing techniques. This is another aspect in which these techniques (in particular, DM, VT, and modified ATM) exhibit their robust character.

4.2 Experimental Results

First, curved surfaces have been constructed in our laboratory with varying curvature and dimensions, using thin cardboard of height 1.05 m and length 3.65 m (see Figure 8 for an example). In these experiments, only the front five ultrasonic sensors of the robot have been fired.

The structured-light data obtained from one of the cardboard surfaces constructed are presented in Figure 2(a). This is the example used in introducing and demonstrating the different techniques in Section 3 and will be described in more detail here. As expected, the structured-light data provide a very accurate surface profile. However, since the structured-light data have a large gap, we fitted a 9th-order polynomial to be used in calculating the errors. If the fitted curve in Figure 2(b) is compared with the structured-light data (Figure 2(a)), it can be observed that a close fit to the original surface is obtained.

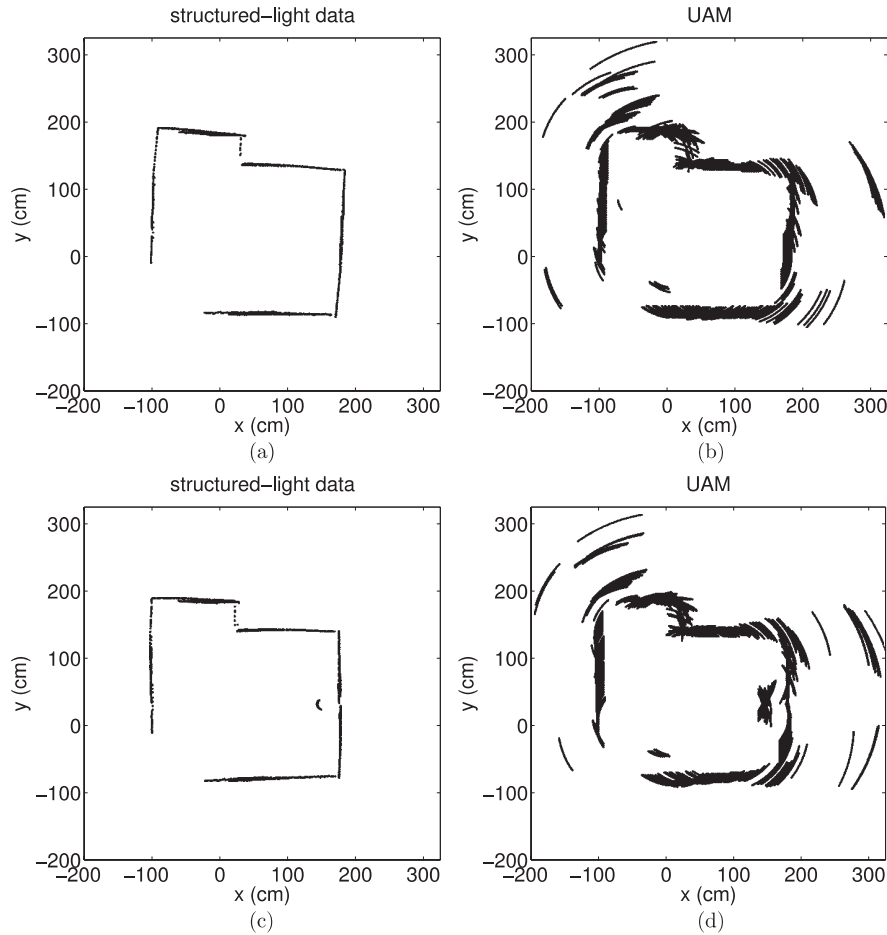


Fig. 9. (a) Structured-light data collected from an environment comprised of planar walls, corners, an edge, and a doorway; (b) the UAM for part (a); (c) a cylinder has been added to the environment in part (a); (d) the UAM for part (c).

In this experiment, the mobile robot simply translated along a straight path from $(-75, 0)$ cm to $(75, 0)$ cm alongside the surface at an average distance of 1 m and collected data by firing the ultrasonic transducers at every 2.5 cm. The turret was oriented such that both the structured-light and the front five ultrasonic sensors were directed towards the surface and it did not rotate throughout the translational movement. Therefore, the DOI for this example was taken to be the positive y direction, perpendicular to the path of the robot. A total of 350 arcs were collected, 284 of which fall in the region shown in Figure 2(c). The environment was divided into 300×300 square elements or pixels with size 1 cm. In the resulting UAM shown in Figure 2(c), there are some arcs which are not tangent to the actual surface at any point (e.g. the isolated arcs in the upper part of Figure 2(c)). These correspond to spurious data due to higher-order reflections, readings from other objects in the environment, or totally erroneous readings. These points are readily eliminated by several of the processing techniques (see Figure 3).

We also present experimental results obtained by using the front three ultrasonic sensors of the Nomad 200 robot, following the walls of the indoor environment in Figure 9(a). The environment comprises of smooth wooden (top and left) and painted (right) walls, and a window shade with vertical slats of 15 cm width (bottom). Some of the corners of the room were not perfect (e.g. where the shade and the right wall make a corner). The resulting UAM is given in Figure 9(b). A total of 738 arcs are obtained from this environment, 697 of which fall in the 525×525 region shown in the figures. Again, the UAM includes many artifacts, especially exterior to the surrounded region where the robot can move. In Figure 9(c), a cylindrical object of radius 15 cm has been added to the environment at a distance of 30 cm from the center of the right wall. Views of the environment showing the wooden and painted walls, the window shade, and the cylinder are given in Figure 10. A total of 702 arcs are obtained, 659 of which fall in the 525×525 region (Figure 9(d)). The results of the different processing techniques for these two experiments are shown in Figure 11



(a)



(b)

Fig. 10. Views of the environment in Figure 9(c): (a) looking towards the right, showing the top, right, and bottom walls; (b) looking towards the lower right corner, showing the right and bottom walls in Figure 9(c).

and 12. In these two experiments, we used a simple rule-based wall-following scheme for motion planning and took the DOI as the direction of the currently followed wall. In other words, the DOI is perpendicular to the path of the robot and is in the direction of the nearest wall from which a TOF reading is acquired.

For VT and DM, the variable parameter is the pixel threshold value, for MP, the thinning parameter m , for BU, the threshold used for probability of occupancy, and for TBF, the number n_t of successful triangulations associated with a range reading. In Tables 2–6, sample results of varying the param-

eters of VT, DM, MP, BU, and TBF techniques are tabulated for some of the experiments. It should be noted that processing with a threshold value of 1 in VT, and thinning with $m = 0$ in MP do not modify the original UAMs. Thinning with $m = 1$ (pruning) removes only the isolated points. For each technique (VT, DM, MP, BU, TBF), a suitable parameter value is selected such that the corresponding CDR_F and CDR_E are as close to each other as possible. This is needed to eliminate any tendency to over-detect pixels as full or empty and to ensure that full and empty regions are detected with approximately the same correct detection rate. In Tables 5 and 6, two different cases for TBF are considered: in the case presented in Table 5, we included all the successful triangulation points where n_t was greater than or equal to a threshold value varied between 1 and 8. In the second case, given in Table 6, only those triangulation points that correspond to edge-like features were included where n_t was again greater than or equal to a threshold between 1 and 8. We note that in this case, the errors are reduced but CDR is also lower, resulting in a very accurate but very sparse map.

Using the most suitable parameter values, a comparison of the seven UAM processing techniques is provided in Tables 7 and 8, for the two indoor environments. In some cases, the results corresponding to more than one competing parameter value are given. Again, ATM-mod produces the smallest error but has the second lowest CDR. The maps extracted with this technique are quite accurate but sparse. The errors of DM and VT are comparable to that of ATM-mod, followed by TBF, ATM-org, MP, BU, and PM. The larger MAEs usually result from artifacts that a certain technique was unable to remove. The DM and ATM-mod techniques are the best in eliminating artifacts, followed by VT, MP, and TBF. The processing results of PM, BU, and ATM-org still contain artifacts. The highest overall CDR is obtained by DM and VT. The CDR of MP is also quite high, followed by BU and PM. The CDRs of PM and ATM-org are comparable and those of ATM-mod and TBF are the lowest among all seven techniques.

The results of MP given in parts (d) of Figures 11 and 12 correspond to applying a single thinning operation of order $m = 7$. Artifacts have been removed to some extent and planar surfaces are satisfactorily represented. The MP results have larger range uncertainty indicated by thick solid line features at planar surfaces, while corners are too much eroded. Furthermore, a substantial amount of arc branches remain and applying further thinning operations not only removes these unwanted branches but also reduces the CDR considerably, while increasing CPU time. From the given resulting maps, DM and ATM techniques are superior to the others in terms of range uncertainty, indicated by thinner borders and more accurately placed features, as well as their lower MAE. Artifacts are also still present in the results of BU given in parts (e) of Figures 11 and 12. Thresholding the probability of occupancy with smaller values is found to result in more artifacts, while larger thresholds reduce the CDRs to too low values.

Table 7. Experimental Results for the Indoor Environment given in Figure 9(a) for the Different Techniques

Method	MAE (cm)	CDR _O	CDR _F	CDR _E	<i>t</i> _{CPU} (s)
PM	6.17	0.627	0.575	0.932	0.004
VT (thld=4)	2.57	0.938	0.951	0.863	0.077
VT (thld=5)	1.87	0.858	0.847	0.926	0.077
DM (thld=4)	2.36	0.938	0.951	0.863	0.085
DM (thld=5)	1.96	0.858	0.847	0.926	0.086
MP (<i>m</i> = 7)	5.85	0.835	0.826	0.888	0.083
BU (thld=0.997)	8.13	0.788	0.787	0.795	0.587
BU (thld=0.998)	7.56	0.775	0.771	0.795	0.584
BU (thld=0.999)	6.93	0.762	0.752	0.820	0.584
ATM-org	4.47	0.624	0.570	0.944	1.126
ATM-mod	1.70	0.535	0.463	0.963	1.128
TBF (<i>n_t</i> ≥ 2)	2.82	0.427	0.334	0.981	0.010

Table 8. Experimental Results for the Indoor Environment given in Figure 9(c) for the Different Techniques

Method	MAE (cm)	CDR _O	CDR _F	CDR _E	<i>t</i> _{CPU} (s)
PM	6.06	0.631	0.579	0.955	0.001
VT (thld=5)	2.63	0.883	0.881	0.896	0.074
DM (thld=5)	2.37	0.883	0.881	0.896	0.078
MP (<i>m</i> = 7)	4.89	0.820	0.824	0.799	0.082
BU (thld=0.997)	5.29	0.785	0.783	0.799	0.563
BU (thld=0.998)	5.20	0.773	0.768	0.799	0.563
BU (thld=0.999)	4.85	0.754	0.741	0.831	0.566
ATM-org	4.78	0.606	0.549	0.955	1.054
ATM-mod	1.68	0.522	0.449	0.947	0.871
TBF (<i>n_t</i> ≥ 2)	2.70	0.414	0.320	1.000	0.010

4.3 Discussion

In general, if a processing technique cannot eliminate artifacts well, the resulting MAE is larger. DM, VT, and ATM-mod are superior to the other techniques in eliminating artifacts, therefore they result in smaller errors. Note that after thresholding but before applying any directional processing in DM, the VT and DM pixel values are exactly equal if the same threshold value is used in both algorithms. If there are multiple maxima along a given DOI, DM directionally processes and selects one of them, whereas VT keeps all. For this reason, DM usually performs better than VT in terms of MAE, especially at smaller values of the threshold between 1 and 4. At larger val-

ues of the threshold, they perform comparably. The CDRs of VT and DM are always equal because if a given DOI is empty, it will remain empty after directional processing; if it is full, it will still be full after the directional maximum is taken along that direction.

The modified ATM is also quite accurate and eliminates the artifacts better than the original ATM. However, comparing the results in parts (f) and (g) of Figures 11 and 12, there are more gaps in parts (g), especially around corners and edges. This is because in those regions of the UAM, there were arcs with no intersections that were removed with the modified ATM but not removed by the original ATM. Generally speaking, the ATM technique creates accurate yet sparsely filled maps with a tendency towards under-filling. Since the same number of arcs are processed by each technique, ATM is found to require a higher number of arcs in order to produce a map with similar CDR to the other techniques.

The PM technique reduces each arc to a single point mark in the middle of the arc. ATM-org places a more accurately positioned point mark on arcs with transversal intersections (except those with 2), reducing many of the arcs to single points also. It treats arcs with no intersections in the same way as PM. For this reason, the CDR of this technique is quite close to that of PM, but as expected, a little lower.

Edge locations obtained with TBF are very accurate as expected. The accuracy of TBF usually falls between ATM-org and ATM-mod. However, due to the large number of gaps in the resulting map, the CDR obtained with TBF is the lowest among all the techniques compared. This is expected because apart from the fact that a smaller number of arcs is used at a given time to begin with (due to the sliding window), TBF eliminates those arcs without any meaningful and accurate correspondence. In addition, planar wall locations are not very accurate. This is also observed in Wijk and Christensen (2000b) in Figures 15 and 16 as many outlying points extracted by the algorithm. For the same reason, the performance of TBF on the curved surface in Figure 2 is not very good. On the other hand, a major advantage of TBF is that it is very fast for the given window sizes (5×20 and 3×20) in our implementation and takes about the same time as the simplest PM method. This is because it does not divide the environment into grids and processes the information geometrically, instead of grid by grid.

Among the seven approaches considered, DM produces one of the lowest MAE and CPU time. Its performance according to the CDR criterion is also very good, where VT and DM usually rank the best in different examples. Considering that DM also has very high range accuracy and is superior in eliminating artifacts and outliers of the UAM, it can be concluded that it results in the best overall performance.

In Barshan and Bařkent (2001a, 2001b) and Barshan and Bařkent (2000), VT and MP were investigated in detail based on simulations and experimental studies for different transducer configurations (linear, circular, random), different

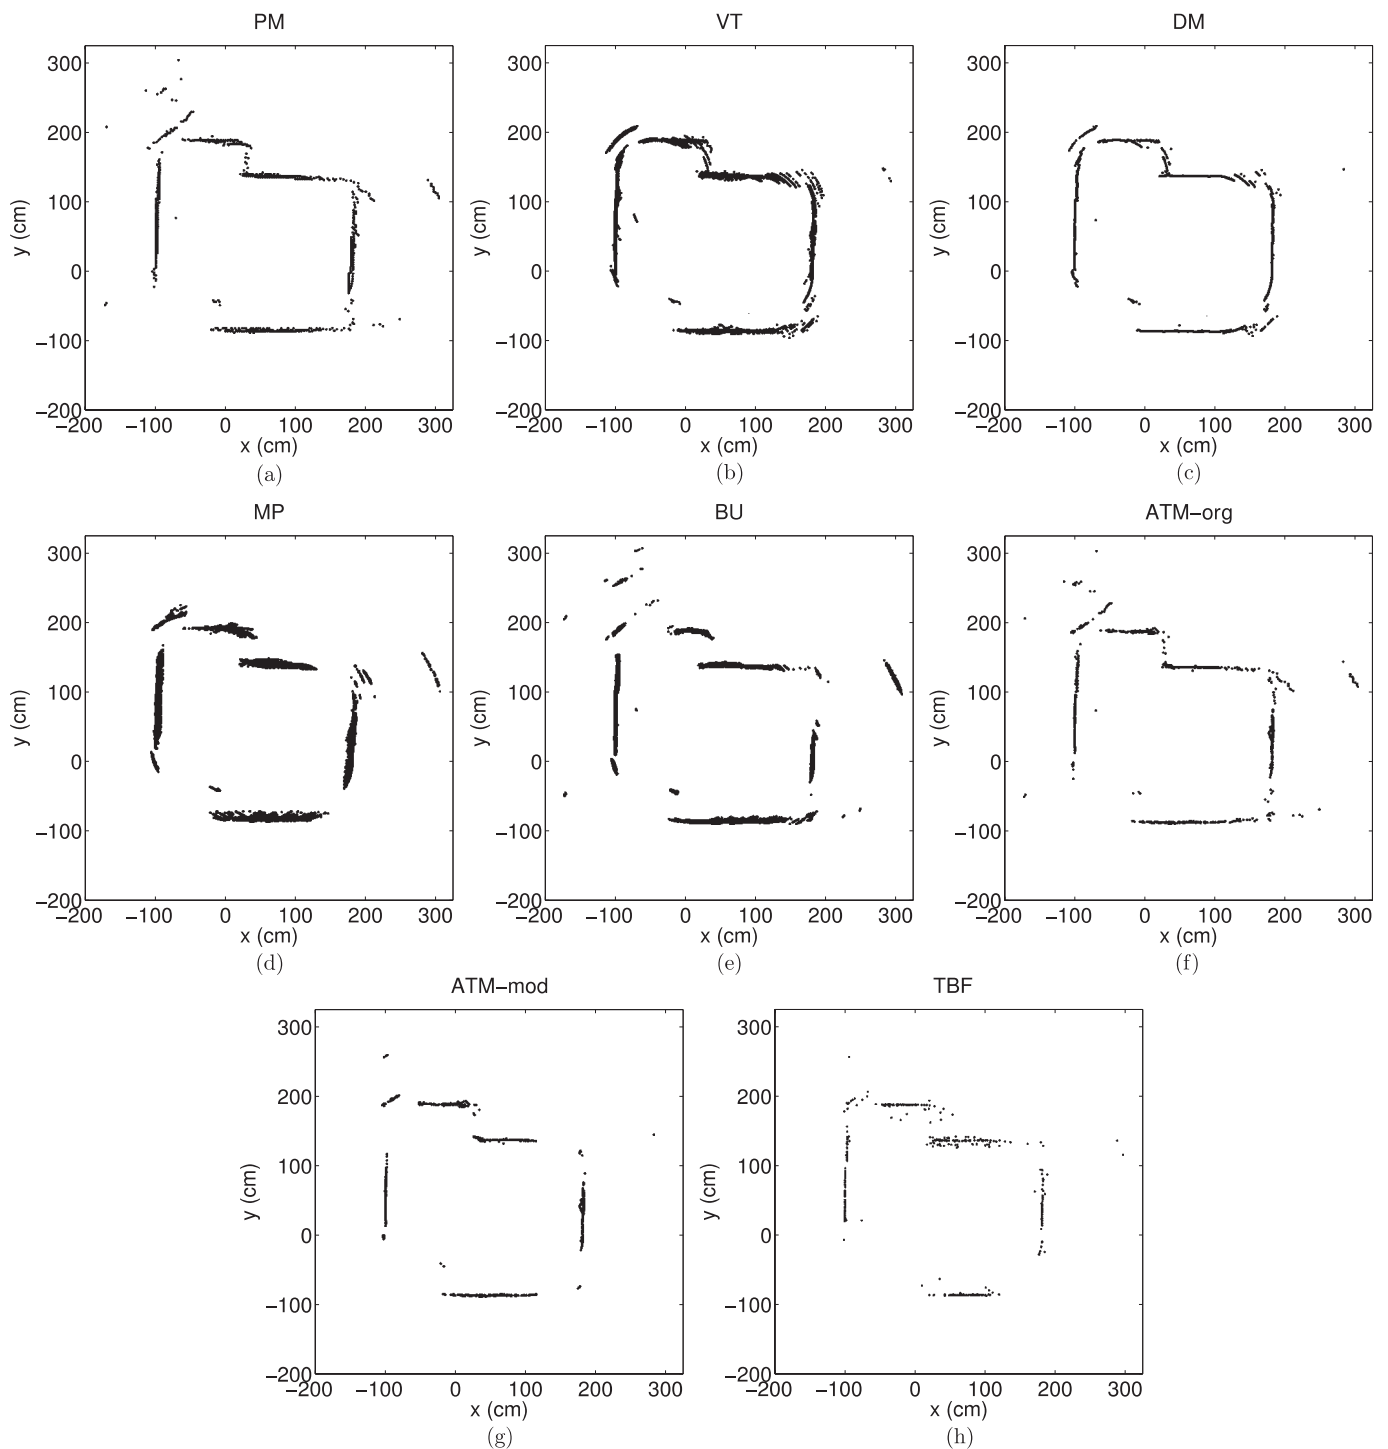


Fig. 11. Results of (a) PM, (b) VT with threshold 4, (c) DM with threshold 4, (d) MP with $m = 7$, (e) BU with threshold 0.998, (f) ATM-org, (g) ATM-mod, and (h) TBF with $n_t \geq 2$.

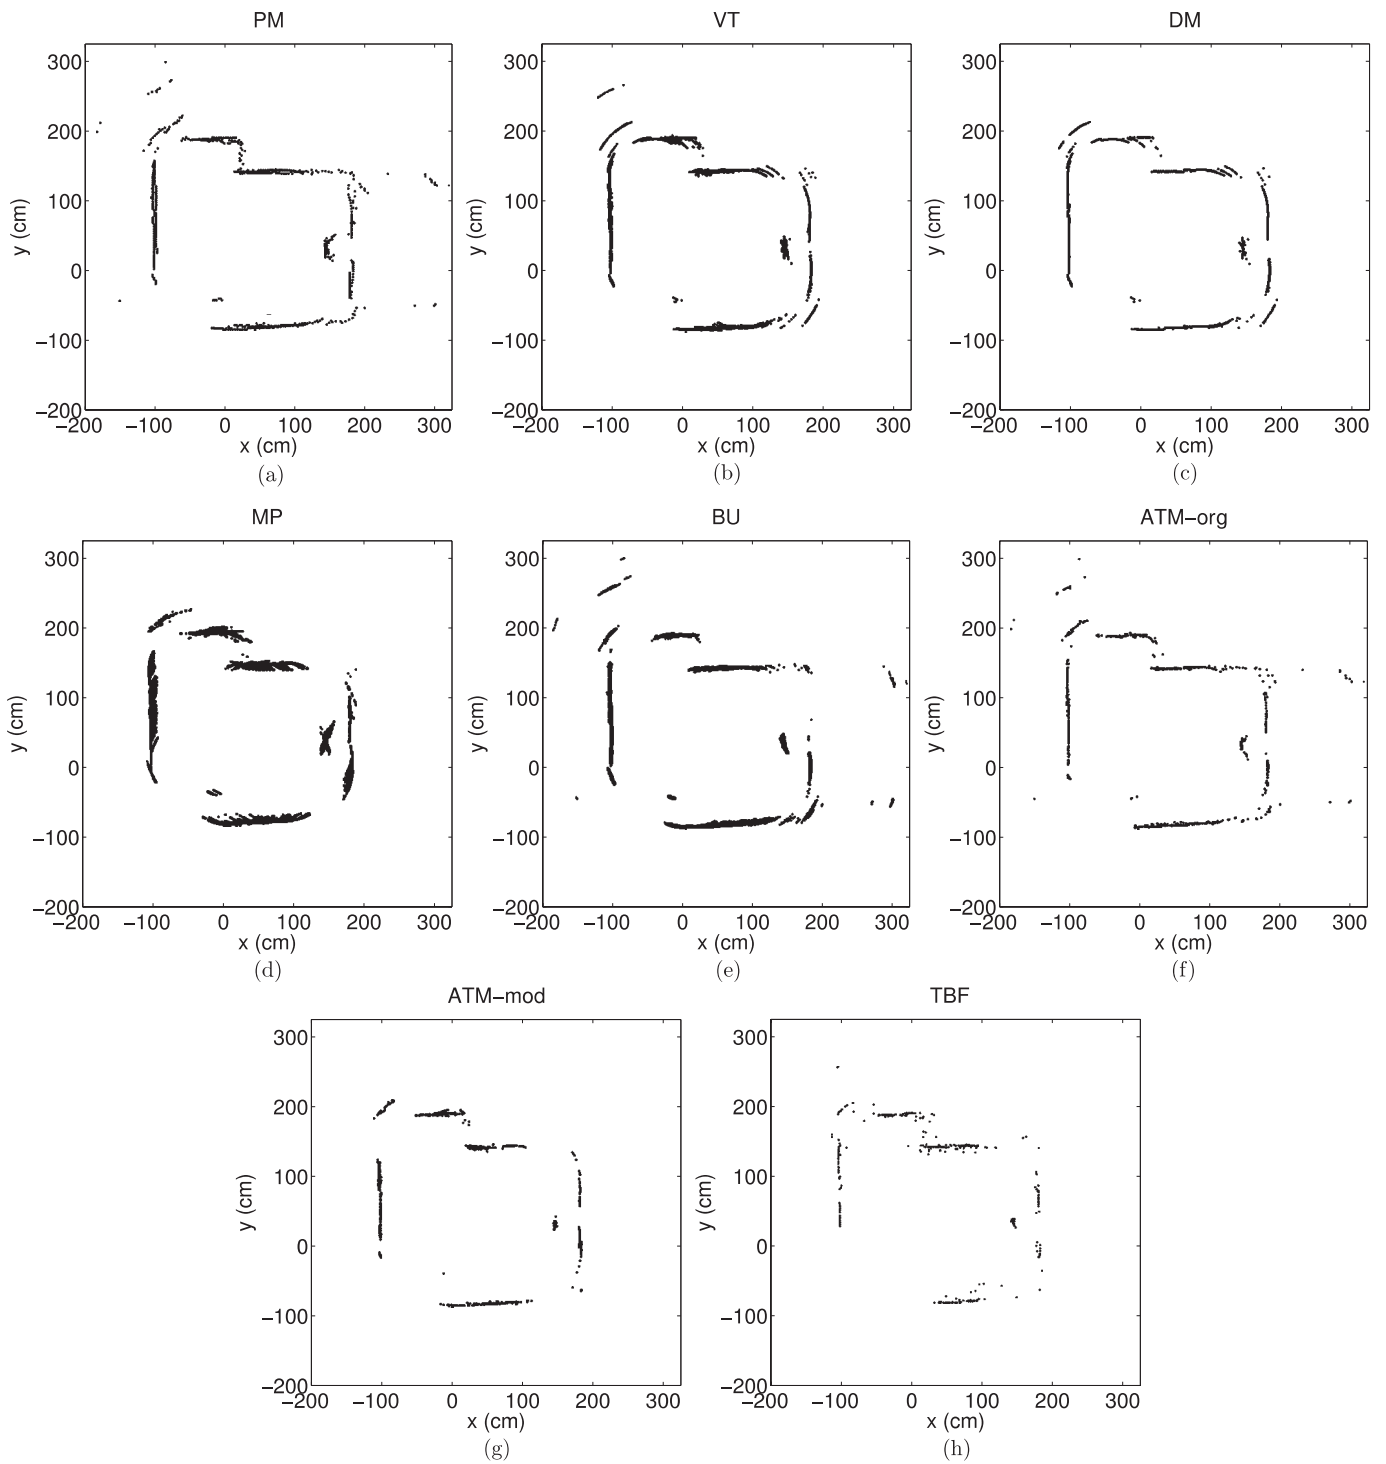


Fig. 12. Results of (a) PM, (b) VT with threshold 5, (c) DM with threshold 5, (d) MP with $m = 7$, (e) BU with threshold 0.998, (f) ATM-org, (g) ATM-mod, and (h) TBF with $n_t \geq 2$.

beamwidths (5° to 105°), different surface curvature, roughness, distance, and different noise levels on the time-of-flight measurements. The best results were obtained with a random configuration of transducers, followed by circular and linear ones. For both methods, the errors were shown to increase with increasing beamwidth, increasing surface distance, curvature, and roughness. Although such detailed studies for the other methods have not been performed, we expect similar results for the remaining techniques. This is because the impact of varying these parameters is primarily to affect the quality of the information inherent in the ultrasonic arc map. This also leads us to expect that for a given choice of these parameters, the comparison of the methods will not be altered significantly.

In Choset et al. (2003), the transversal angle θ_t , which is the intersection angle of the arcs was defined and only the arcs intersecting at angles larger than θ_t were processed. Too small values of θ_t do not provide much improvement in angular resolution, whereas too large values result in too few intersections and a very sparse map. A value of 30° was selected in Choset et al. (2003), arguing that this brings about 10.5-fold accuracy improvement in angular resolution. In processing with ATM, we varied the value of θ_t to be able to choose a suitable value for it experimentally. The results for the curved surface example are presented in Figure 13. In the curved surface example, there are 284 arcs intersecting over a wide range of transversal angles. Referring to the results in Figure 13(a), the MAE of ATM-org increases drastically after about 13° . For ATM-mod, the error appears to decrease beyond $\theta_t = 25^\circ$, however, too few points are left in the map and there is not sufficient coverage after this value of θ_t . Processing with ATM-mod, there were no points left in the map after $\theta_t = 83^\circ$. Referring to Figure 13(b), The CDR of ATM-mod decreases drastically starting around $\theta_t = 20^\circ$. The CDR of ATM-org remains more or less flat at a level close to that obtained by the PM method. Therefore, we decided that $\theta_t = 13^\circ$ is a good compromise between MAE and CDR criteria. Similar plots for the other two experiments are given in Figures 14 and 15. In these experiments, the transversal angles were distributed over a narrower range and we used a smaller value of 5° for θ_t . In fact, we believe that the most suitable value of θ_t should be a parameter of the ATM algorithms and chosen according to the distribution of the arc intersection angles in the map, which depends on the features of the environment, the motion planning scheme used in data acquisition, and the configuration of the sensors. In practice, since the true map of the environment will not be known beforehand, θ_t can be selected based on the histogram of the arc intersection angles in the UAM. A central characteristic value of the histogram such as the mean, mode or the median would be a suitable choice.

4.4 Computational Cost of the Techniques

The average CPU times are of the order of fractions of a second, indicating that processing methods are viable for real-

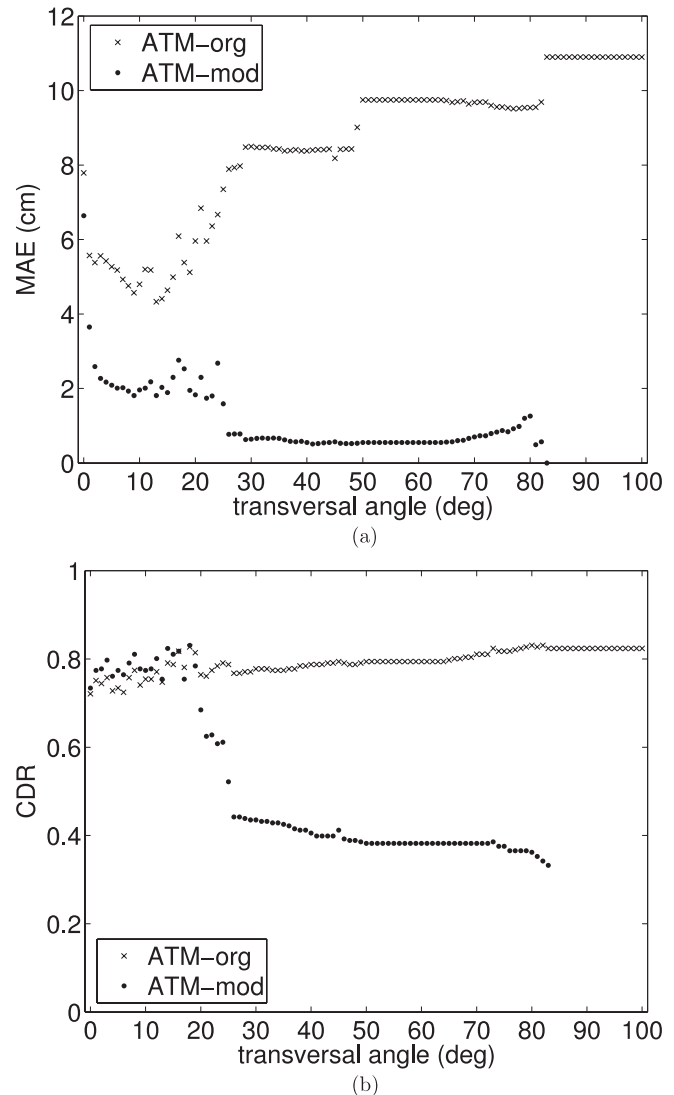


Fig. 13. (a) MAE and (b) CDR_O for ATM-org and ATM-mod techniques for the curved surface in Figure 2(a).

time applications. These represent the total time the computer takes to realize the processing techniques, starting with the raw TOF data (i.e. the UAM). (The processing techniques have been implemented in the C language and run on an Intel Pentium 4 PC with 3.00 GHz Hypertreading processor and 1 GB memory. Internal C commands are used for time keeping.)

Since PM is the simplest technique, it has the smallest processing time as expected (less than one-hundredth of a second). It is closely followed by TBF. The CPU times of VT, DM, and MP are comparable to each other and slightly less than one-tenth of a second in the last two experiments. CPU times of BU and ATM are larger but still suitable for real-time applications. Since the total CPU time depends on the total number of arcs in the UAM, dimensions of the environ-


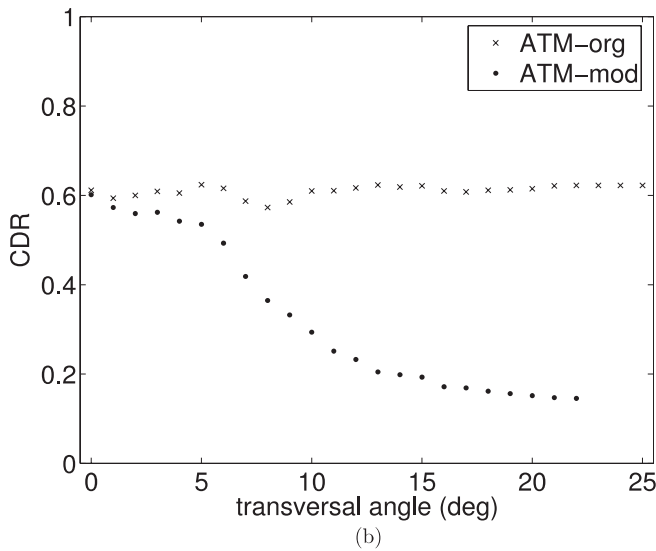




Fig. 14. (a) MAE and (b) CDR_O for ATM-org and ATM-mod techniques for the indoor environment given in Figure 9(a).

Fig. 15. (a) MAE and (b) CDR_O for ATM-org and ATM-mod techniques for the indoor environment given in Figure 9(c).

ment, and the grid size, CPU time per arc is a better indicator of the computational load of the processing techniques. For the BU method, the CPU time per arc is less than 1 ms, whereas for the ATM method, the CPU time per arc is about 1.5 ms. The CPU times of original and modified ATM are comparable.

The computational complexity of TBF depends strongly on the calculation of the triangulation points and their optional refinement using a local grid map. It is also dependent on the dimensions of the sliding window. Since we omitted the optional refinement step and used a relatively small window size, computational complexity was not high in our implementation.

Inclusion of the refinement step would considerably increase the cost but further improve the accuracy of the triangulation points, consequently reducing the MAE.

For comparison, the time it takes for an array of 16 ultrasonic sensors to collect TOF data is $16 \times 40 \text{ ms} = 0.64 \text{ s}$ which is of the same order of magnitude as the processing time for BU. It should be noted that the actual algorithmic processing time is a small fraction of the CPU time, as most of the CPU 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 or extract the map of the environment even faster.

Another important factor is memory usage. Memory usage depends on the number of arcs processed, size of the surface or the environment, grid size, and the nature of the processing technique. These were given above for the different experiments. The data files typically consume about 5–70 kB. Among the techniques considered, ATM is the one that requires the largest amount of bookkeeping and storage.

Among all the methods considered, PT and BU can be implemented on-line while data are still being acquired by the mobile robot. The thresholding part of VT, thresholding and directional processing of DM, MP, and ATM work better if employed off-line, after the complete UAM becomes available. However, it may be possible to develop regional or on-line versions of these techniques.

5. Conclusions

The main contribution of this paper is introducing the directional maximum technique and providing a valuable comparison between this technique and existing techniques for processing ultrasonic arc maps for map building in indoor environments. The results indicate that the newly proposed directional maximum technique has some advantages over existing techniques. Having a small mean absolute error, one of the largest correct detection rates and a CPU time less than one-tenth of a second, the directional maximum technique should be considered a good compromise between the mean absolute error and correct detection rate criteria in many cases. It also has low range uncertainty and is superior in eliminating artifacts. Some of the existing techniques also have characteristics that may be suitable for certain situations or conditions. For example, if a large number of ultrasonic arcs is conveniently available and processing time and memory is not a major issue, the arc transversal median algorithm would be a good choice. On the other hand, if simple and fast on-line processing is desirable, and high accuracy is not crucial, one could use the point marking method, as is done in many cases. Triangulation-based fusion is more suitable for environments in which many edge-like features are present and extracts the positions of these features quite accurately. It is also very fast for moderate sliding window sizes.

The directional maximum technique in itself has added an important aspect to the map-building techniques, which is the direction of interest. The sense of direction readily available in most motion-planning schemes is shown to be an effective addition when incorporated into the map-building process. The directional awareness of the mobile platform for surface profile extraction tasks was proved to deliver satisfactory maps.

As for future extensions of this work, although wall following provided satisfactory results, motion-planning schemes such as Voronoi diagram tracing may be employed since they

offer a more systematic and complete method for ground coverage. However, it may be computationally more demanding and should be used when the computational cost is of secondary importance. Two of the error criteria (mean absolute error and correct detection rate) used in this paper may be combined and reduced to a single criterion by fitting polar polynomials or active, snake-like contour models to the resulting maps. Polar directions of interest may be considered and the maps acquired at certain vantage points may be fused, employing different data fusion techniques. On-line and regional versions of the techniques may be developed. A surface-following algorithm for mapping arbitrarily curved surfaces can be developed. Even though the results demonstrated in this paper are based on ultrasonic range sensing for map-building applications, techniques presented here can be conveniently extended to other sensing modalities such as radar and infrared, as well as other mobile robot applications.

Acknowledgment

This work was supported in part by The Scientific and Technological Research Council of Turkey (TÜBİTAK) under grant number EEEAG-105E065. The author would like to thank Professor O. Aytür for taking some of the photographs.

References

- Auran, P. G. and Silven, O. (1996). Underwater sonar range sensing and 3-D image formation. *Control Engineering Practice*, **4**(3): 393–400.
- Barshan, B. (1999). Ultrasonic surface profile determination by spatial voting. *Electronics Letters*, **35**(25): 2232–2234.
- Barshan, B. and Başkent, D. (2000). Comparison of two methods of surface profile extraction from multiple ultrasonic range measurements. *Measurement Science and Technology*, **11**(6): 833–844.
- Barshan, B. and Başkent, D. (2001a). Morphological surface profile extraction with multiple range sensors. *Pattern Recognition*, **34**(7): 1459–1467.
- Barshan, B. and Başkent, D. (2001b). Map building from range data using mathematical morphology. In *Active Sensors for Local Planning in Mobile Robotics* (ed. P. Probert Smith), Chapter 7, pp. 111–135. Vol. 26 of *World Scientific Series in Robotics and Intelligent Systems*, World Scientific, NJ.
- Barshan, B. and Kuc, R. (1990). Differentiating sonar reflections from corners and planes by employing an intelligent sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **12**(6): 560–569.
- Barshan, B. and Kuc, R. (1992). A bat-like sonar system for obstacle localization. *IEEE Transactions on Systems Man and Cybernetics*, **22**(4): 636–646.

- Başkent, D. and Barshan, B. (1999). Surface profile determination from multiple sonar data using morphological processing. *International Journal of Robotics Research*, **18**(8): 788–808.
- Beckerman, M. and Oblow, E. (1990). Treatment of systematic errors in the processing of wide-angle sonar sensor data for robotic navigation. *IEEE Transactions on Robotics and Automation*, **6**(2): 137–145.
- Bemporad, A., Di Marco, M., and Tesi, A. (1997). Wall-following controllers for sonar-based mobile robots. *Proceedings of the IEEE International Conference on Decision and Control*, San Diego, CA, December, Vol. 8, pp. 3063–3068.
- Borenstein, J. and Koren, Y. (1991a). The vector field histogram: fast obstacle avoidance for mobile robots. *IEEE Transactions on Robotics and Automation*, **7**(3): 278–288.
- Borenstein, J. and Koren, Y. (1991b). Histogramic in-motion mapping for mobile robot obstacle avoidance. *IEEE Transactions on Robotics and Automation*, **7**(4): 535–539.
- Borenstein, J., Everett, H. R., and Feng, L. (1996). *Navigating Mobile Robots*. A K Peters Ltd., Wellesley, MA.
- Bozma, Ö. I. (1992). *A Physical Model-Based Approach to Analysis of Environments using Sonar*. PhD thesis, Yale University, Dept. of Electrical Engineering, New Haven, CT.
- Bruce, J. and Veloso, M. (2002). Real-time randomized path planning for robot navigation. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, September/October, Vol. 3, pp. 2383–2388.
- Chen, Y. D. and Dougherty, E. R. (1994). Gray-scale morphological granulometric texture classification. *Optical Engineering*, **33**(8): 2713–2722.
- Choset, H. and Burdick, J., (2000). Sensor-based exploration: The hierarchical generalized Voronoi graph. *International Journal of Robotics Research*, **19**(2): 96–125.
- Choset, H., Walker, S., Eiamsa-Ard, K., and Burdick, J. (2000). Sensor-based exploration: Incremental construction of the hierarchical generalized Voronoi graph. *International Journal of Robotics Research*, **19**(2): 126–148.
- Choset, H., Nagatani, K., and Lazar, N. (2003). The arc-transversal median algorithm: a geometric approach to increasing ultrasonic sensor azimuth accuracy. *IEEE Transactions on Robotics and Automation*, **19**(3): 513–522.
- Cox, I. J. (1991). Blanche: an experiment in guidance and navigation of an autonomous robot vehicle. *IEEE Transactions on Robotics and Automation*, **7**(2): 193–204.
- Crowley, J. L. (1985). Navigation for an intelligent mobile robot. *IEEE Transactions on Robotics and Automation*, **RA-1**(1): 31–41.
- Drumheller, M. (1987). Mobile robot localization using sonar. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAMI-9**(2): 325–332.
- Elfes, A. (1987). Sonar based real-world mapping and navigation. *IEEE Transactions on Robotics and Automation*, **RA-3**(3): 249–265.
- Elfes, A. (1989). Using occupancy grids for mobile robot perception and navigation. *IEEE Computer*, **22**(6): 46–57.
- Everett, H. R. (1995). *Sensors for Mobile Robots, Theory and Application*. A K Peters Ltd., Wellesley, MA.
- Gasos, J. and Martin, A. (1996). A fuzzy approach to build sonar maps for mobile robots. *Computers in Industry*, **32**(2): 151–167.
- Gilbreath, G. and Everett, H. (1988) Path planning and collision avoidance for an indoor security robot. In *Proceedings SPIE: Mobile Robots III*, Vol. 1007 (ed. W. J. Wolfe), pp. 19–27, Boston, MA.
- Grimson, W. E. and Lozano-Perez, T. (1984). Model-based recognition and localization from sparse range or tactile data. *International Journal of Robotics Research*, **3**(3): 3–35.
- 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.
- Hsu, D., Kindel, R., Latombe, C., and Rock, S. (2002). Randomized kinodynamic motion planning with moving obstacles. *International Journal of Robotics Research*, **21**(3): 233–255.
- Kleeman, L. and Kuc, R. (1995). Mobile robot sonar for target localization and classification. *International Journal of Robotics Research*, **14**(4): 295–318.
- Kuc, R. and Siegel, M. W. (1987). Physically-based simulation model for acoustic sensor robot navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAMI-9**(6): 766–778.
- Kurt, A. (2005). *A Comparative Study of Five Algorithms for Processing Ultrasonic Arc Maps*. Master's thesis, Bilkent University, Department of Electrical Engineering, Ankara, Turkey.
- Kurz, A. (1996). Constructing maps for mobile robot navigation based on ultrasonic range data. *IEEE Transactions on Systems Man and Cybernetics*, **26**(2): 233–242.
- Lee, D. and Recce, M. (1994). Quantitative evaluation of the exploration strategies of an intelligent vehicle. *International Journal of Robotics Research*, **16**(4): 413–447.
- Leonard, J. J. and Durrant-Whyte, H. F. (1992). *Directed Sonar Sensing for Mobile Robot Navigation*. Kluwer Academic Publishers, Boston, MA.
- McMullan, W. G., Delanghe, B. A., and Bird, J. S. (1996). A simple rising-edge detector for time-of-arrival estimation. *IEEE Transactions on Instrumentation and Measurement*, **45**(4): 823–827.
- Mojsilovic, A., Popovic, M., Amodaj, N., Babic, R., and Ostojic, M. (1997). Automatic segmentation of intravascular ul-

- trasound images: A texture-based approach. *Annals of Biomedical Engineering*, **25**(6): 1059–1071.
- Moravec, H. P. and Elfes, A. (1986). High resolution maps from wide angle sonar. *Proceedings of IEEE International Conference on Robotics and Automation*, San Francisco, CA, April, pp. 116–121.
- Nomadic Technologies Inc. (1997) *Nomad 200 Manual*, Mountain View, CA.
- Oriolo, G., Ulivi, G., and Vendittelli, M. (1997). Fuzzy maps: a new tool for mobile perception and planning. *Journal of Robotic Systems*, **14**(3): 179–197.
- Pagac, D., Nebot, E. M., and Durrant-Whyte, H. F. (1998). An evidential approach to map-building for autonomous vehicles. *IEEE Transactions on Robotics and Automation*, **14**(4): 623–629.
- Peremans, H., Audenaert, K., and Van Campenhout, J. M. (1993). A high-resolution sensor based on tri-aural perception. *IEEE Transactions on Robotics and Automation*, **9**(1): 36–48.
- Polaroid Corporation (1997). *Polaroid Manual*, Ultrasonic Components Group, Cambridge, MA.
- Serra, J. (1982). *Image Analysis and Mathematical Morphology*. Academic Press, New York.
- Tunstel, E. and Jamshidi, M. (1994). Embedded fuzzy logic-based wall-following behavior for mobile robot navigation. *Proceedings of 1st International Joint Conference of NAFIPS/IFIS/NASA*, San Antonio, TX, December, pp. 329–330.
- Verly, J. G. and Delanoy, R. L. (1993). Some principles and applications of adaptive mathematical morphology for range imagery. *Optical Engineering*, **32**(12): 3295–3306.
- Wallner, F. and Dillman, R. (1995). Real-time map refinement by use of sonar and active stereo vision. *Robotics and Autonomous Systems*, **16**(1): 47–56.
- Wijk, O. (2001). *Triangulation Based Fusion of Sonar Data with Application in Mobile Robot Mapping and Localization*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden.
- Wijk, O. and Christensen, H. I. (2000a). Triangulation-based fusion of sonar data with application in robot pose tracking. *IEEE Transactions on Robotics and Automation*, **16**(6): 740–752.
- Wijk, O. and Christensen, H. I. (2000b). Localization and navigation of a mobile robot using natural point landmarks extracted from sonar data. *Robotics and Autonomous Systems*, **31**: 31–42.
- Yata, T., Kleeman, L., and Yuta, S. (1998). Wall following using angle information measured by a single ultrasonic transducer. *Proceedings of IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May, Vol. 5, pp. 1590–1596.
- Yun, X. and Tan, K.-C. (1997). A wall-following method for escaping local minima in potential field based motion planning. *Proceedings of the IEEE International Conference on Advanced Robotics*, Monterey, CA, July, pp. 421–426.
- Zelinsky, A. (1988). Environment mapping with a mobile robot using sonar. In *Proceedings of the Australian Joint Artificial Intelligence Conference-AI'88*, November, pp. 373–388.