Simukai W. Utete
Department of Engineering Science
University of Oxford
Oxford OX1 3PJ, UK

Billur Barshan
Birsel Ayrulu
Department of Electrical Engineering
Bilkent University
Bilkent 06533, Ankara, Turkey

Voting as Validation in Robot Programming

Abstract

This paper investigates the use of voting as a conflict-resolution technique for data analysis in robot programming. Voting represents an information-abstraction technique. It is argued that in some cases a voting approach is inherent in the nature of the data being analyzed: where multiple, independent sources of information must be reconciled to give a group decision that reflects a single outcome rather than a consensus average. This study considers an example of target classification using sonar sensors. Physical models of reflections from target primitives that are typical of the indoor environment of a mobile robot are used. Dispersed sensors take decisions on target type, which must then be fused to give the single group classification of the presence or absence and type of a target. Dempster-Shafer evidential reasoning is used to assign a level of belief to each sensor decision. The decisions are then fused by two means. Using Dempster's rule of combination, conflicts are resolved through a group measure expressing dissonance in the sensor views. This evidential approach is contrasted with the resolution of sensor conflict through voting. It is demonstrated that abstraction of the level of belief through voting proves useful in resolving the straightforward conflicts that arise in the classification problem. Conflicts arise where the discriminant data value, an echo amplitude, is most sensitive to noise. Fusion helps to overcome this vulnerability: in Dempster-Shafer reasoning, through the modeling of nonparametric uncertainty and combination of belief values; and in voting, by emphasizing the majority view. The paper gives theoretical and experimental evidence for the use of voting for data abstraction and conflict resolution in areas such as classification, where a strong argument can be made for techniques that emphasize a single outcome rather than an estimated value. Methods for making the vote more strategic are also investigated. The paper addresses the reduction of dimension of sets of decision points or decision makers. Through a consideration of combination order, queuing criteria for more strategic fusion are identified.

The International Journal of Robotics Research
Vol. 18, No. 4, April 1999, pp. 401-413,

KEY WORDS—voting, Dempster-Shafer theory, evidential reasoning, sonar sensing, ultrasonic transducers, multisensor data fusion, spatial placement, robot programming, target classification, learning.

1. Introduction

Sensing systems are fundamental to robot control and programming. In multisensor systems, diverse geometric, spatial, and physical information is acquired from the environment, through the deployment of multiple sensors, for example (Durrant-Whyte 1987; Luo and Lin 1987). Increasingly, systems of multiple-sensor *nodes* are being exploited, as by Chong, Mori, and Chang (1990). In sensor nodes, sensing is augmented by processing and communication capabilities. Therefore, local processing of sensor data, and in some cases local decision making, takes place. Numerous advantages accrue: processing of data before transmission, as opposed to bulk transfer of raw data, means more-compact messages can be delivered, with extraneous material being removed prior to transmission. Local processing can result in increased speed of operation, which is important in real-time robotics applications; operations may be performed more efficiently in parallel at every node rather than in parallel in software at a single, central, processing site. Increased speed of data transfer is possible where there is communication between the sensing nodes. Increased reliability can be achieved by devolving processing and communication activities to sensing sites. These are general advantages; the degree to which they arise is dependent on the design of the multiple-node system. Multisensor systems fall into a taxonomy based on processing and communication complexity. In the simplest case, sensors may relay data directly to a central processing node. Increasing the complexity of communication, nodes in hierarchical systems may process some data locally prior to communicating to a central site (Chong, Mori, and Chang

1990). Eliminating the central processor altogether, individual sensor nodes can operate autonomously, on the basis of communicated information, as in some decentralized systems (Durrant-Whyte, Rao, and Hu 1990; Grime, Durrant-Whyte, and Ho 1991). This paper is concerned with a problem that is common to systems of multiple sensor nodes—obtaining a group decision on the basis of the combination of individual views.

Individual sensor views must be combined so that a consistent and coherent system outcome is determined. This outcome may be an estimate, say of a robot's position or of a target's velocity. A consequence of having multiple decision makers is that conflicts can arise. These conflicts may be data dependent, predicated on the diversity of views in the sensing system, or in more-complex sensing scenarios, owing to processing diversity—tactical differences in the decision-making strategies of the sensor nodes. Fusing diverse sensor data is a fundamental requirement of multisensor systems. Numerous fusion methods exist, but these can broadly be divided into parametric and nonparametric approaches. In parametric methods, a model of the process and of the observations can be adduced and/or an underlying probability distribution is assumed. The models derived are employed to estimate the state of the system or environment under observation. For example, sensor estimates may be combined using Bayesian methods. In nonparametric methods, no assumptions about distribution are embedded in the fusion process. This can yield greater robustness in certain situations; for example, where noise is nonadditive or non-Gaussian, or where the system model does not adequately represent the process. Nonparametric methods based on voting have been applied widely in reliability problems. An example is their use to achieve fault tolerance in satellite systems (Takano et al. 1996). Other applications include voting-based algorithms such as the Hough transform (see, for example, Duda and Hart 1973). A majority-voting scheme for fusing features in 3-D object recognition is presented by Mao, Flynn, and Jain (1995). Some application of voting has been made to the analysis of sonar data; for example, for underwater target recognition (Foresti et al. 1997) and for data analysis using neural networks (Tan and Teow 1997). Examples of applications in robotics exist (Sukthankar 1997; Rosenblatt 1997; Toye and Leifer 1994). However, although potentially a very powerful approach, voting methods have not been widely exploited in robotics.

Klein (1993a, 1993b) has applied voting fusion to target detection and discusses Dempster-Shafer evidential reasoning in contrast to this. Evidential reasoning (Shafer 1976) requires no a priori information. This method has been applied to robotics by several researchers recently (Murphy 1996; Pagac, Nebot, and Durrant-Whyte 1996; Reece 1997). A concise description of the Dempster-Shafer approach appears in Krause and Clark's (1993) book.

This paper is concerned with the application of nonparametric methods based on voting to the problem of multisensor fusion for target classification, and draws a comparison with fusion based on evidential reasoning. In the work presented, data abstraction through voting is employed as a means of combining diverse sensor information in robot localization or navigation problems. In particular, the problems of resolving conflicts where multiple, dissenting decision makers are present is considered. Thus, the paper is concerned with the highest level in the taxonomy of sensor systems—the case where individual nodes hold decision-making power. The scope of this work is wider in the sense that abstraction is often useful in determining single outcomes from conflicting respondents. In this case, the respondents are multiple sensor nodes, but the essence of the problem is that of resolving conflict between data points that support different outcomes. Comparison is made between the voting approach and Dempster-Shafer evidential reasoning.

The paper looks also at ways in which the fusion techniques can be made more strategic through the extraction of informative decision points. Order is used to determine sensor nodes that improve the rate at which a preset belief value is attained. The nodes are queued according to selected criteria, including similarity in level of belief and physical separation. These groups act as logical sensors with preference orders on target type. Comparisons are made with a randomly generated fusion order. Results using Dempster fusion are again contrasted with the results of the combination of sensor beliefs through voting. An analogy between this approach—data ordering until breakdown of the fused result—and learning is explored in Section 4, which discusses the incorporation of voting strategies in robot programming. All methods are verified by experiments with a practical sonar system.

In the experimental work on target classification, multiple sensor agents, located in a rectangular room containing target primitives, combine their views to determine target type. Sonar sensors that make use of time-of-flight and amplitude data are employed. Conflict resolution is a particular concern in this type of problem. A unique group decision is required; therefore, the fusion strategy must be able to incorporate or exclude variant views. Sonar time-of-flight data is often sensitive to physical conditions, and conflicting responses by the sensor respondents are common in experimental systems. In addition, although the strategies of the sensor nodes are the same at each site, differences in views arise owing to geographical diversity. The next section describes the Dempster-Shafer belief assignment used by the sensors, and the underlying mass function that is used for target differentiation. Conflict resolution through voting is then implemented for the sonar decisions on target type in the experimental room.

2. Conflict Resolution through Voting

Multisensor systems exploit sensor diversity to acquire a wider view of a scene or target under observation. This diversity can give rise to conflicts, which must be resolved when the

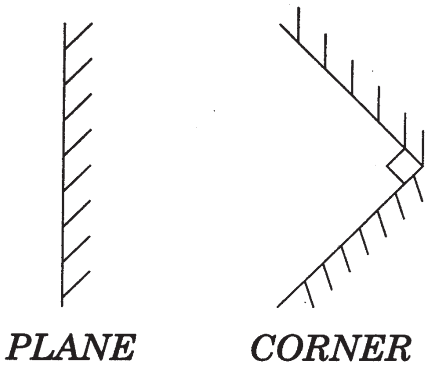

Fig. 1. Target primitives modeled and differentiated in this study.

system information is combined to take a group decision or to form a group value or estimate. The way in which conflict is resolved is encoded in the fusion method. In Dempster-Shafer evidential reasoning, each sensor's view or belief is tied to a belief measure or basic probability assignment. A priori information is not required, and belief is assigned only to those values for which the sensor maintains a view. Therefore, ignorance in view can be represented explicitly. Conflict between views is represented by a conflict measure that is used to normalize the sensor belief assignments.

An abstract means of analyzing sensor data is to combine the views of sensors in a vote, for example by making the group decision on the basis of simple majority. The level of belief (as in Dempster-Shafer methods) or degree of uncertainty (Bayesian methods) is completely abstracted to give a binary output. A wealth of voting strategies exists (see, for example, Kim and Roush 1980), and it is possible to encapsulate information such as strength of belief in the voting process. Abstracting information is a useful tool for conflict resolution. Voting, in its simplest form, has the advantage of being computationally inexpensive and, to a degree, fault tolerant. In cases where the decisions to be made by the sensing system themselves abstract the data, it may be more efficient to employ the instrument of a vote in preference to the fine tuning of parametric information. This is useful in the example of target classification considered here.

Drawbacks of voting include problems of consistency (Arrow 1951). The discrimination algorithm employed by the sensors here is mutually exclusive over target type. For more-complicated decision-making scenarios, preference orders could exist over the possible target types, based on strength of belief. Such preferences admit the possibility of order-dependent outcomes in the vote. A problem of order that can arise in voting is that of cyclical majorities (see, for example, Kim and Roush 1980; and Harary, Norman, and Cartwright 1965). A cycle can arise in the preference order of the sensor group, such that the preference relationship is not transitive. In such a case, the order of voting is important. By delaying the point at which it enters the vote, a sensor could influence the outcome in its favor. Alternatively, the vote could generate an outcome that no voter prefers. Harary, Norman and Cartwright treat the problem, which is related to Arrow's impossibility theorem (Arrow 1951), in terms of tournaments on digraphs. Modifying an example of Harary, Norman, and Cartwright (1965) to reflect a sensing scenario, three sensors might make observations of a target that can be of type plane (P), corner (C), edge (E), or cylinder (Cyl). Let each have a different preference ordering for the four outcomes, determined by the level of belief apportioned to each target type. Where a cycle exists, the final outcome of pairwise votes on the target types may favor sensors that enter their first choice later in the voting process. For example, for the following sensor preferences over the possible outcomes:

Group 1 P, C, E, Cyl;
Group 2 C, E, Cyl, P; and
Group 3 E, Cyl, P, C,

a cycle exists:

$$\begin{array}{ccc} \mathrm{P} & \rightarrow & \mathrm{C} \\ \uparrow & & \downarrow \\ \mathrm{Cyl} & \leftarrow & \mathrm{E} \end{array}$$

and pairwise votes taken in the following order will favor choice P:

E versus C → C wins,
C versus Cyl → C wins,
C versus P → P wins.

For the sensing problem considered here, a two-way decision is required. Planes and corners, as illustrated in Figure 1, form the set of target primitives encountered by the mobile robot. Sonar sensors scan the room, making decisions about target types at each viewing angle. While these decisions may hold information about level of belief, and the final decision may include some marker of strength of belief, the eventual, required outcome is a single decision for each target encountered: in effect, a Boolean value as opposed to an estimate. This work hypothesizes that in such cases, abstraction is a useful tool for fast combination of the sensor data. Even where this is not the case, the gains in employing parametric methods may be marginal by comparison, and in some cases outcomes may be degraded because of the emphasis given to sensors with strong, but misplaced, levels of belief. Abstraction through voting is a leveling system, and offers the advantage of rewarding replicated views. This may be significant for certain problems. It is argued here that in many instances, the classification of environmental features is one such problem.

2.1. Algorithm for Plane/Corner Differentiation

The classification algorithm employed by the individual sonar sensor nodes is the plane/corner differentiation algorithm de-

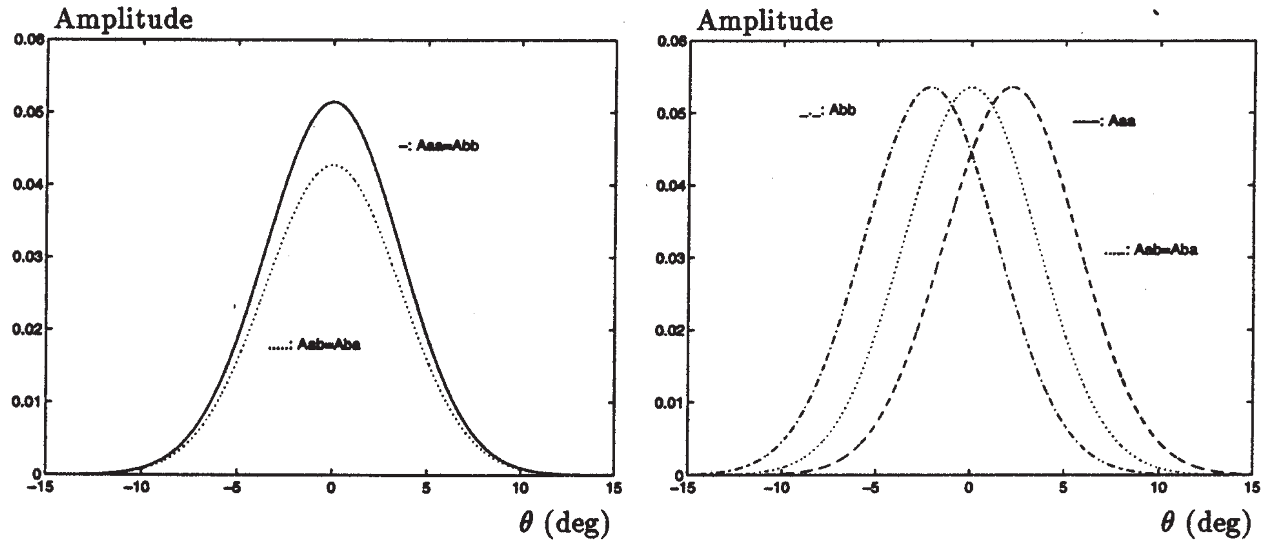

Fig. 2. Plane and corner amplitude characteristics.

rived by Barshan and Kuc (1990). The basis of the algorithm is a realization that, in resolving target types, amplitude differentials are significant. The plane and corner amplitude characteristics are provided in Figure 2 as an example. The algorithm is detailed in the work by Ayrulu and Barshan (1998), and an extension to other target primitives, based on time-of-flight and amplitude differentials, is given by Ayrulu (1996).

To emphasize the decision procedure, a summary is presented here in the form of rules. Each sensing unit consists of two horizontally spaced transducer pairs, labeled a and b. Amplitudes are denoted A, with the subscript representing the transmitter and receiver, in that order. For example, A_{ab} denotes the signal amplitude detected at b due to a pulse from a. Ideally, A_{ab} and A_{ba} should be identical. Characteristic amplitudes are shown in Figure 2. The algorithm discriminates on the basis of amplitude differences. To enhance robustness in decision making, amplitude differences are considered significant only if they exceed a lower bound $k_A\sigma_A$, where the factor k_A is a multiple of the amplitude noise standard deviation, σ_A:

plane-differentiation algorithm
if $[A_{aa}(\theta)-A_{ab}(\theta)] > k_A\sigma_A$ **and** $[A_{bb}(\theta)-A_{ab}(\theta)] > k_A\sigma_A$
then **plane**
else **corner** or **unknown**;
corner-differentiation algorithm
if $[A_{ab}(\theta) - A_{aa}(\theta)] > k_A\sigma_A$ **or** $[A_{ab}(\theta) - A_{bb}(\theta)] > k_A\sigma_A$
then **corner**
else **unknown**.

From these rules, it can be seen that each decision involves only a two-way choice. Differentiation is achieved when the discriminant data value, the difference in amplitudes, exceeds the $k_A\sigma_A$ bound; all other cases are classified as unknown. Therefore, the differentiation algorithm is mutually exclusive on target type, a feature which makes it a particularly strong candidate for abstraction through voting.

2.2. Experimental Setup

The two fusion methods were tested on amplitude data acquired in experiments using scanning sonar sensors. The sensors acquire data from scans of a room, making unilateral decisions on target type at each of several viewing angles. These decisions are then fused to yield the group decision. The data was collected at Bilkent University Robotics Research Laboratory, in a small (1.0-m × 1.4-m) rectangular "room" created by partitioning off a section of a laboratory. The test area was calibrated by lining the floor space with metric paper, to allow the sensors and targets to be positioned accurately. The room offers an uncluttered environment, with specularly reflecting surfaces. Sensor nodes occupy the 15 sites shown in Figure 3. The transducers used are Panasonic sensors (Panasonic 1989), with an aperture radius of $a = 0.65$ cm and a resonant frequency of $f_\circ = 40$ kHz; the resulting beamwidth is $\theta_\circ \cong 54°$. The transducer's transmission and reception characteristics are distinct. In the experimental setup, a transmitter and a receiver are vertically closely spaced. Two such pairs form a single logical sensing unit. A typical unit is shown in Figure 4. The horizontal center-to-center separation of the transducers is $d = 24.0$ cm. The sensing unit is mounted on a small 6-V stepper motor with step size 0.9°. The stepping action is controlled through the parallel port of an IBM-PC 486, with the aid of a microswitch. The sensor data is acquired using a DAS-50 A/D card with 12-bit resolution and 1-MHz sampling frequency. The echo signals are

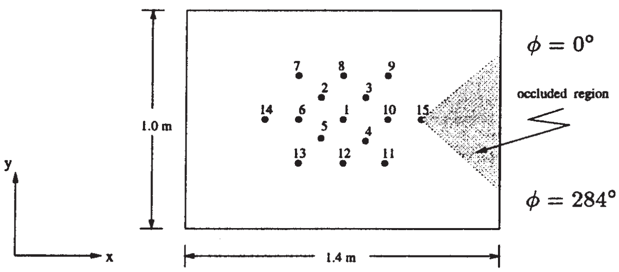

Fig. 3. The 15 sensing sites in the rectangular room.

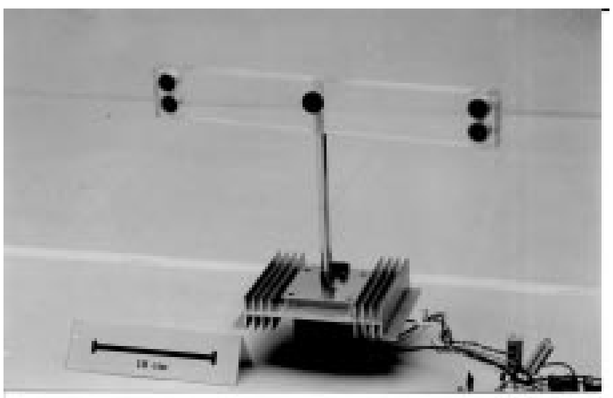

Fig. 4. Configuration of the Panasonic transducers in the real sonar system. The two transducers on the left collectively constitute one transmitter/receiver. Similarly, those on the right constitute another.

processed on an IBM-PC 486 using a C language program. From the time of transmission, 10,000 samples of each echo signal are collected and thresholded. The amplitude information is extracted by finding the maximum value of the signal after the threshold value is exceeded.

As an example, the range readings of the sensor node located at (−0.1,0.1) m (sensor node 2) are given in Figure 5. The physical limitations of the hardware prevent any sensor from covering the entire angular range of ϕ. Instead, rotation is over the range $0° \leq \phi \leq 284°$. A typical map of the plane and corner locations, obtained by one of the sensors, is shown in Figure 6.

2.3. Dempster-Shafer Evidential Reasoning

The sensors are assigned beliefs using Dempster-Shafer evidential reasoning, and their opinions are combined using Dempster's fusion rule (Shafer 1976). Dempster-Shafer theory is based on the use of belief functions. These are set functions that assign numerical degrees of support on the basis of evidence, but allow for the expression of ignorance: belief can be committed to a set or proposition without commitment to its complement.

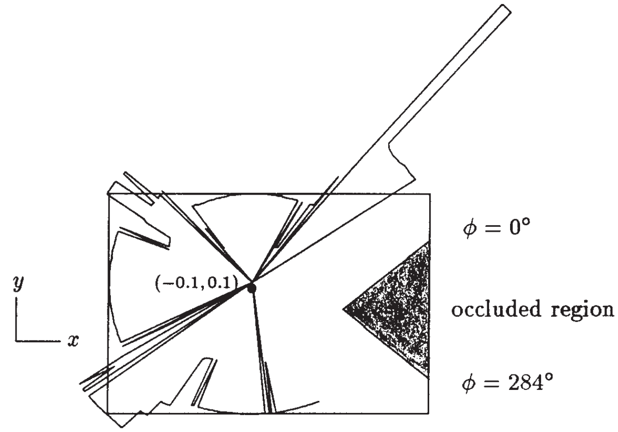

Fig. 5. Range readings of the sensor located at (−0.1, 0.1) m in the rectangular room.

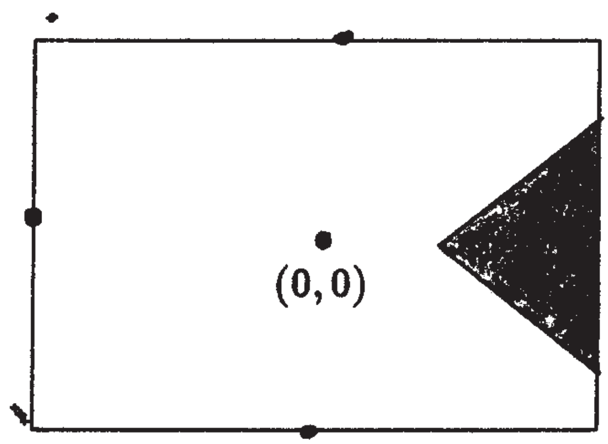

Fig. 6. A typical map of planes and corners in the room, obtained from the sensor at the center of the room. Planes are shown with circles and corners are shown with crosses throughout the scan. When a decision cannot be made (unknown), no mark is made on the map.

In Dempster-Shafer theory, a frame of discernment, Ω, represents a finite universe of propositions, and a basic probability assignment, m, maps the power set of Ω to the interval [0, 1]. The basic probability assignment satisfies the following conditions:

$$m(\emptyset) = 0, \tag{1}$$

$$\sum_{A \subseteq \Omega} m(A) = 1. \tag{2}$$

A set that has a nonzero basic probability assignment is termed a focal element.

The belief or total support that is assigned to a set or proposition A is obtained by summing the basic probability assignments over all subsets of A:

$$Bel(A) = \sum_{B \subseteq A} m(B). \tag{3}$$

Evidence that does not support A directly does not necessarily support its complement. The plausibility of A, denoted $Pl(A)$, expresses the amount of evidence for propositions that do not support the complement of A:

$$Pl(A) = 1 - Bel(\overline{A}). \tag{4}$$

Each proposition A is therefore associated with a belief $Bel(A)$, which represents the evidence directly supporting it, and a plausibility $Pl(A)$, representing evidence that fails to support the negation of A. Evidential reasoning is defined by Shafer (1976). A concise description, with applications, can be found in the work of Krause and Clark (1993).

The assignments for the target-classification problem are made as follows. The uncertainty in the measurements of each sonar pair (sensor node) is represented by a belief function having target type or *feature* as a focal element with the basic probability assignment (or mass assignment) $m(.)$ associated with this feature:

$$BF = \{\mathit{feature}; m(\mathit{feature})\}. \tag{5}$$

The mass function is the underlying function for decision making using the Dempster-Shafer belief assignment. It is defined here according to the algorithm outlined above, and is thus dependent on a difference in signal amplitudes; the greater the amplitude difference, the higher the degree of belief. The structure of mass assignment is chosen to be consistent with the **and/or** rules of the differentiation algorithm. The planar **and** rule is represented by a product; the corner **or** rule is given by a summation. The mass assignment levels are scaled to fall in the interval [0,1]. The basic probability assignment is described below, where $m(p)$ and $m(c)$ correspond to the plane and corner assignments, respectively:

$$m(p) = I_1 \frac{[A_{aa}(\theta) - A_{ab}(\theta)][A_{bb}(\theta) - A_{ab}(\theta)]}{\max[A_{aa}(\theta) - A_{ab}(\theta)] \max[A_{bb}(\theta) - A_{ab}(\theta)]}, \tag{6}$$

$$m(c) = \begin{cases} \frac{I_2[A_{ab}(\theta)-A_{aa}(\theta)]+I_3[A_{ab}(\theta)-A_{bb}(\theta)]}{I_2 \max[A_{ab}(\theta)-A_{aa}(\theta)]+I_3 \max[A_{ab}(\theta)-A_{bb}(\theta)]} \\ \text{if } I_2 \neq 0 \text{ or } I_3 \neq 0, \\ \text{else } 0 \end{cases}. \tag{7}$$

where $A_{ab}(\theta)$ denotes the maximum value over time of $A_{ab}(r, \theta, d, t)$, which is the signal transmitted by a and received by b. Definitions of $A_{aa}(\theta)$ and $A_{bb}(\theta)$ are similar. I_1, I_2, and I_3 are indicators of the conditions given below:

$$\begin{aligned} I_1 &= \begin{cases} 1 & \text{if } [A_{aa}(\theta) - A_{ab}(\theta)] > k_A\sigma_A \\ & \text{and } [A_{bb}(\theta) - A_{ab}(\theta)] > k_A\sigma_A, \\ 0 & \text{otherwise;} \end{cases} \\ I_2 &= \begin{cases} 1 & \text{if } [A_{ab}(\theta) - A_{aa}(\theta)] > k_A\sigma_A, \\ 0 & \text{otherwise;} \end{cases} \\ I_3 &= \begin{cases} 1 & \text{if } [A_{ab}(\theta) - A_{bb}(\theta)] > k_A\sigma_A, \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{8}$$

Mass that is distributed neither to p nor to c is here assigned to a type termed "unknown." This remaining belief represents ignorance, or undistributed probability mass:

$$m(u) = 1 - [m(p) + m(c)]. \tag{9}$$

Dempster's fusion rule applies where independent opinions are to be combined. Given two sources with belief functions

$$\begin{aligned} BF_1 &= \{f_i, m_1(f_i)\}_{i=1}^3 = \{p, c, u; m_1(p), m_1(c), m_1(u)\}, \\ BF_2 &= \left\{g_j, m_2(g_j)\right\}_{j=1}^3 = \{p, c, u; m_2(p), m_2(c), m_2(u)\}, \end{aligned} \tag{10}$$

consensus is obtained as the orthogonal sum

$$\begin{aligned} BF &= BF_1 \oplus BF_2 \\ &= \{h_k, m_c(h_k)\}_{k=1}^3 = \{p, c, u; m_c(p), m_c(c), m_c(u)\}, \end{aligned} \tag{11}$$

which is both associative and commutative. The sequential combination of multiple bodies of evidence can be obtained for n sensor pairs as

$$BF = (((BF_1 \oplus BF_2) \oplus BF_3) \ldots \oplus BF_n). \tag{12}$$

Using Dempster's rule of combination:

$$m(h_k) = \frac{\sum\sum_{h_k = f_i \cap g_j} m_1(f_i) m_2(g_j)}{1 - \sum\sum_{h_k = f_i \cap g_j = \emptyset} m_1(f_i) m_2(g_j)}, \tag{13}$$

where $\sum\sum_{h_k = f_i \cap g_j = \emptyset} m_1(f_i) m_2(g_j)$ is a measure of conflict. The consensus belief function representing the feature fusion process has the measure

$$\begin{aligned} m(p) &= \frac{m_1(p)m_2(p) + m_1(p)m_2(u) + m_1(u)m_2(p)}{1 - \text{conflict}}, \\ m(c) &= \frac{m_1(c)m_2(c) + m_1(c)m_2(u) + m_1(u)m_2(c)}{1 - \text{conflict}}, \\ m(u) &= \frac{m_1(u)m_2(u)}{1 - \text{conflict}}. \end{aligned} \tag{14}$$

Disagreement in the consensus of two logical sensing units is represented by the "conflict" term in the equations above. Thus, it represents the degree of mismatch in the fusion of

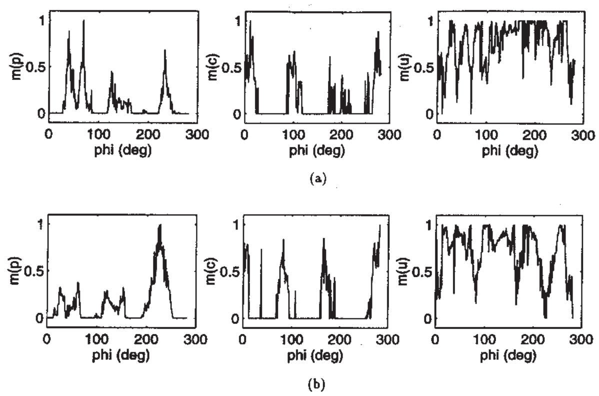

Fig. 7. Belief assignments by the sensors located at (0,0) m (a) and (−0.1,0.1) m (b).

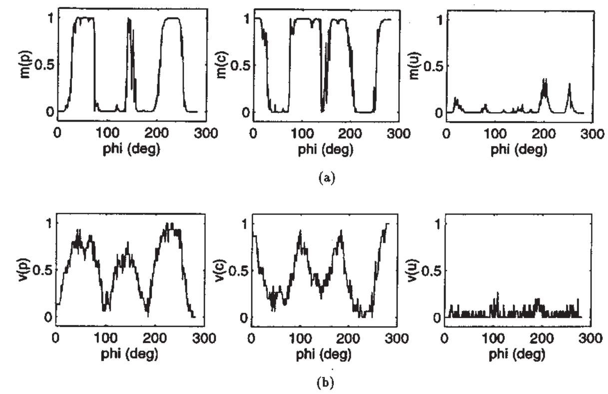

Fig. 8. Results for Dempster's rule (a) and simple majority voting (b) using an arbitrary order of fusion.

features perceived at two different sonar sites. The measure evaluating conflict is expressed as

$$\text{conflict} = m_1(p)m_2(c) + m_1(c)m_2(p). \tag{15}$$

After discounting this conflict, the beliefs can be rescaled and used in further data fusion processes.

Examples of basic probability assignments by individual sensors are given in Figure 7. During a scan, each sensor estimates the range and angle of the target under observation. The values for a target are weighted by the beliefs assigned to the estimates. When only a single logical sensor is employed, a high degree of uncertainty is observed. The aim of fusion is to reduce this uncertainty.

Belief levels with a single sensor and after fusion over 15 sensors are given in Table 1 for the 6 targets when the targets are along the line of sight.

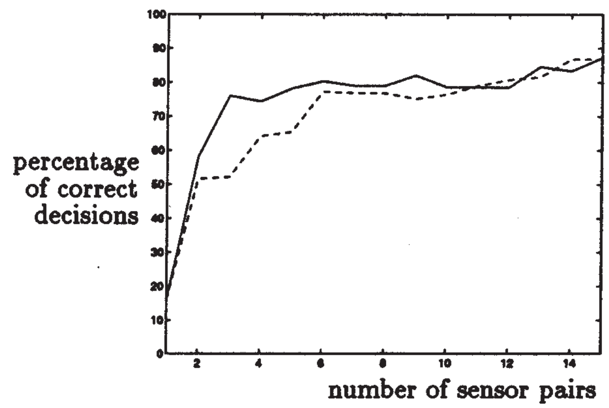

Fig. 9. Correct decision percentage with Dempster's rule (dashed line) and the simple voting algorithm (solid line), versus the number of sensors employed in the fusion process when an arbitrary order of fusion is used.

2.4. Simple Majority Voting

Independent sensor opinions can also be combined in a vote. In this case, conflict is resolved through some form of majority decision. For the sake of comparison, the sensors' beliefs, which are later abstracted, are the probability assignments of the Dempster-Shafer approach. Beliefs could have been assigned in other ways; for example by learning assignments using decision trees (Utete, Barshan, and Srinivasan, in preparation). However, belief abstraction rather than belief assignment is the focus of this work.

The sensors' beliefs about target type are counted as votes, and the majority vote is taken as the outcome. The results of fusing beliefs by simple majority vote and using Dempster's rule of combination are compared. The opinions of all 15 sensors are initially combined, from site 1 to site 15, without regard to intermediate fused results. Results after voting over 15 sensors are given in Table 2 for the 6 targets when they are along the line of sight. Results after fusion are provided in Figure 8. To illustrate the accumulation of evidence, the percentage of correct decisions is plotted as a function of the number of sensor pairs used; plots are given in Figure 9. When the decisions of all of the sensors are fused, the correct decision on target type is made for all targets. However, the maximum percentage of correct decisions achievable is below 100%, because at certain viewpoints during a scan the targets may not be visible. When the decisions of 15 pairs are fused using Dempster's fusion rule, the correct decision percentage improves to 86.75%. With simple majority voting, using the same ordering as for the Dempster's fusion rule case, the number rises to 87.50%. It can be noted that after simple voting fusion from about five pairs, the correct decision percentage remains approximately constant around 81%, indicating redundancy in the number of sensors employed. When a single sensor is used, only about 15% of its decisions are correct. The outstanding 85%, the incorrect and unknown decisions, can be attributed to noise, choice of k_A, and complete uncertainty that occurs when the target is not visible to the sensor.

Table 1. Single Sensor Beliefs versus Fusion over 15 Sensors Using Dempster's Rule

	$\phi = 51.3°$	$\phi = 104.4°$	$\phi = 140.4°$	$\phi = 175.5°$	$\phi = 230.4°$	$\phi = 283.5°$
Sensor 1	$m(p) = 0.260$	$m(p) = 0$	$m(p) = 0.167$	$m(p) = 0$	$m(p) = 0.684$	$m(p) = 0$
	$m(c) = 0$	$m(c) = 0.509$	$m(c) = 0$	$m(c) = 0.618$	$m(c) = 0$	$m(c) = 0.429$
	$m(u) = 0.740$	$m(u) = 0.491$	$m(u) = 0.833$	$m(u) = 0.382$	$m(u) = 0.316$	$m(u) = 0.571$
Dempster's rule	$m(p) = 0.9988$	$m(p) = 0.0002$	$m(p) = 1.0$	$m(p) = 0$	$m(p) = 0.9999$	$m(p) = 0$
(all 15 sensors)	$m(c) = 0.0007$	$m(c) = 0.9992$	$m(c) = 0$	$m(c) = 0.993$	$m(c) = 0$	$m(c) = 1.0$
	$m(u) = 0.0005$	$m(u) = 0.0005$	$m(u) = 0$	$m(u) = 0.007$	$m(u) = 0.00001$	$m(u)$=0

Table 2. Simple Majority Voting Results over 15 Sensors

	$\phi = 51.3°$	$\phi = 104.4°$	$\phi = 140.4°$	$\phi = 175.5°$	$\phi = 230.4°$	$\phi = 283.5°$
Majority voting	$v(p) = 12/15$	$v(p) = 1/15$	$v(p) = 10/15$	$v(p) = 3/15$	$v(p) = 14/15$	$v(p) = 2/15$
(all 15 sensors)	$v(c) = 2/15$	$v(c) = 10/15$	$v(c) = 3/15$	$v(c) = 11/15$	$v(c) = 0/15$	$v(c) = 12/15$
	$v(u) = 1/15$	$v(u) = 4/15$	$v(u) = 2/15$	$v(u) = 1/15$	$v(u) = 1/15$	$v(u)$=1/15

3. Strategic Voting

In this section, the sensors are ordered—essentially, placed in a queue—on the basis of a selected criterion. Fusion occurs in the determined order. Dempster's fusion rule is independent of order (Krause and Clark 1993). For a fixed group of nodes, the resulting belief will be the same, independent of the order in which the beliefs are combined. However, by varying order, it is possible to achieve a preset belief level in a shorter time. The aim is to determine the more informative nodes in the fusion process. Order can also be varied to eliminate redundant or less informative sensors, thus allowing the preset targeted belief level to be reached using fewer nodes. The problem of order is not negated by parallelism, as it can arise through network structure. In networks that do not allow direct communication between all nodes, propagation delays affect the transfer of information. This means, in effect, that through the network's structure, a decision is being made to serve potentially more informative nodes first (Utete 1995). Looking at the problem another way, these nodes become more informative because they are served first; the situation can be viewed as a problem of synthesis as well as of analysis. The other aspect of order is its use to extract a parsimonious set of informative nodes. This process may be repeated for different sensing scenarios, and thus different criteria. Heuristics based on ordering are important in ordinal optimization methods (Ho 1994), and the extraction of parsimonious decision sets is related to the problem of determining criteria for clustering data (Pearl 1988).

3.1. Ordering Strategies

The order of combination of beliefs is varied in a number of ways. The fusion order is first generated by taking the level of belief as the criterion of node informativeness. The nodes are placed in order, based on increasing belief level, irrespective of target type. In making group decisions, the sensor nodes evaluate their decisions relative to those of the group. Starting with the sensor with the highest belief, nodes are added to the fusion list in the order of smallest distance in belief (highest belief). The objective is to determine whether strength of belief forms a natural selection for sensor nodes. This is analogous to dimensionality reduction in pattern recognition (Duda and Hart 1973), where, among a large number of features, more informative ones, or those with large variances, are selected to improve the efficiency of the classification process. In a similar fashion, in this study, those sensors with larger beliefs are fused first. The objective is to select, from a group of decision makers, a parsimonious set of accurate experts that can achieve a given bound on the correct decision rate. A threshold can be set on the belief level so that the fusion process is limited to sensors that exceed this level.

The results of maximum-toward-minimum belief fusion are illustrated in Figure 10 for the two methods. Here, the fusion process begins with the sensor that has the highest belief in a target type, and continues in the direction of decreasing belief. The performances of the two fusion methods are comparable, and the average correct-decision percentage is around 85%. The zigzag pattern in the voting results arises with the switch between odd and even numbers of decision makers.

For comparison purposes, at each viewing angle, fusion was performed only with those five sensors that possess the highest belief levels. Clearly, these need not be the same five sensors at each step throughout the scan. Using the Dempster-Shafer approach, the results were 84.62% correct on average, while voting gave 79.91%. Similarly, both fusion methods were again applied to compare the performance of the five sensors with the lowest belief levels. In this case, Dempster's rule yields only 3.42% correct decisions, whereas simple majority voting gives 65.81%. This significant difference in performance indicates that voting is insensitive to belief levels and can be more robust when high uncertainty prevails (Turney 1994a, 1994b). Since voting emphasizes *numbers* of voters supporting an outcome, as opposed to the *strength of belief* of voters, which is significant in Dempster fusion, this result is expected.

A metric is also defined based on the physical distance relative to an arbitrary origin. Starting with a randomly selected

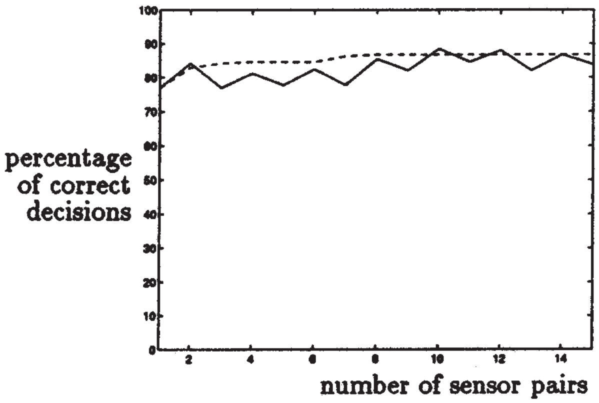

Fig. 10. Decision fusion from maximum toward minimum belief with Dempster's rule (dashed line) and the simple majority voting (solid line), versus the number of sensors employed in the fusion process.

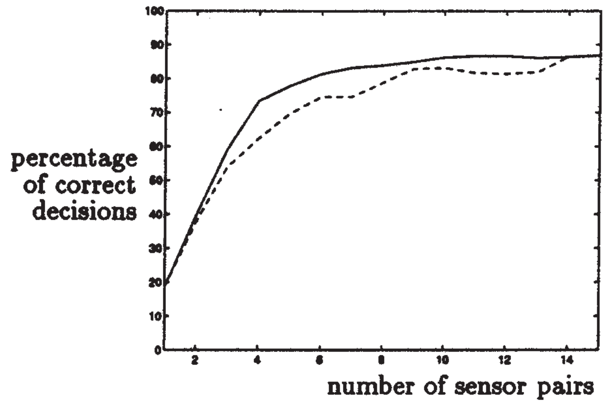

Fig. 11. Average percentage of correct decisions versus the starting sensor in fusion with Dempster's rule when the decisions of sensors are fused according to maximum distance (solid line) and minimum distance (dashed line).

node, the beliefs are fused in the order of greatest physical separation. The starting node selects for fusion the node at greatest separation. The next node selected is the one whose distance is greatest from the two nodes that have already combined their beliefs. In this method, the objective is to acquire a comprehensive view of the room more quickly. Distance measures other than physical separation, for example correlation, could be used. Similarly, a minimum-distance criterion can be established.

Distance calculations are made as follows. Suppose after fusion over n sensors, the average x- and y-positions of the group are

$$x_{\mathrm{av}}(n) = \frac{1}{n}\sum_{i=1}^{n} x_i, \tag{16}$$

$$y_{\mathrm{av}}(n) = \frac{1}{n}\sum_{i=1}^{n} y_i. \tag{17}$$

The $n+1^{\mathrm{th}}$ sensor is chosen such that the distance

$$[x_{n+1} - x_{\mathrm{av}}(n)]^2 + [y_{n+1} - y_{\mathrm{av}}(n)]^2 \tag{18}$$

is maximized (or minimized) over the remaining sensors. In the next step, the new average x- and y-positions can be found recursively:

$$x_{\mathrm{av}}(n+1) = \frac{n}{n+1}\, x_{\mathrm{av}}(n) + \frac{1}{n+1}\, x_{n+1}, \tag{19}$$

$$y_{\mathrm{av}}(n+1) = \frac{n}{n+1}\, y_{\mathrm{av}}(n) + \frac{1}{n+1}\, y_{n+1}. \tag{20}$$

The results of distance-based fusion are illustrated in Figures 11 and 12 for the two fusion methods. In both figures, the results reflect averaging over the 15 possibilities for the starting sensor. In Figure 11, results for Dempster's fusion rule using maximum- and minimum-distance criteria are compared. Note that, as should be expected, the results for maximum- and minimum-distance fusion are identical for a single sensor and also when all 15 sensors are included in the fusion process. For the case of 15 sensors, since Dempster's fusion rule is commutative and associative, the end result is the same when the same 15 sensors are used but sorted differently based on minimum- and maximum-distance criteria. The concern is with intermediate results. In the intermediate fused stages, the maximum performance difference is about 10% between minimum- and maximum-distance fusion. In Figure 12, a similar plot is given for simple majority voting. In this case, note that the average percentage of correct decisions is much larger than the Dempster-Shafer result for up to five or six sensing nodes. The maximum difference is 25%, which occurs when two sensors are used. The effect of choosing the maximum or minimum distance appears to be insignificant for voting. After fusion over six sensors, the performances of the two methods become comparable.

In Figure 13, sensors are eliminated one at a time from the group, and the performances of the two methods are compared after fusion over 14 sensors. The horizontal axis indicates which sensor node is eliminated in the fusion process. From the results, sensors 1, 6, 13, and 14 appear to be most informative. This elimination method can be generalized from individual sensors to groups, to enable the effect of elimination of certain groups of sensors to be studied.

3.2. Grouping

Strategies for grouping sensors during the fusion process are investigated. The sensor nodes are grouped on the basis of the selected criterion. Fusion occurs first within the clusters. In

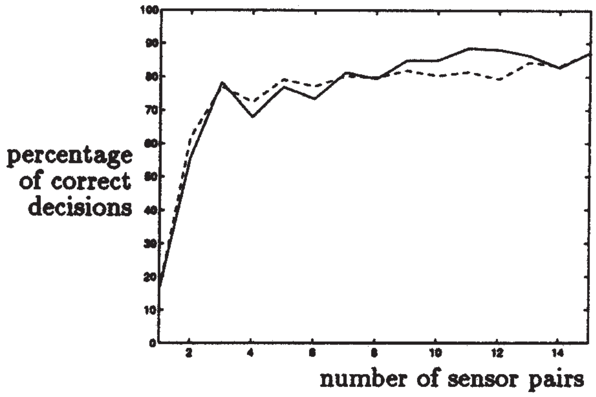

Fig. 12. Average percentage of correct decisions versus the initial sensor in simple majority voting in which the decisions of the sensors are fused according to maximum distance (solid line) and minimum distance (dashed line).

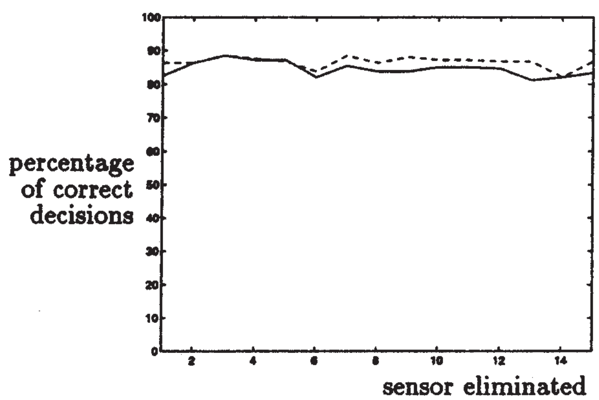

Fig. 13. Fusion with Dempster's rule (dashed line) and simple majority voting (solid line), versus number of the sensor that is eliminated in the fusion process.

this way, beliefs that support the same criterion are enhanced prior to their fusion with the beliefs of dissenting sensors.

The strategy investigated is to organize adjacent nodes in groups of three. The decisions of the nodes in each group form a sequence of votes, whose order is determined by physical distance. Since the differentiation algorithm is mutually exclusive on target type, a sensor node's decision at each step allocates its total belief between a single target type or unknown. By organizing the sensors, a group acts as a single logical sensor with a preference order on target type. For example, a sensor that observes a planar target within its field of view should correctly identify the target based on amplitude difference. When the target is at a large angular deviation from the line of sight, the difference in amplitudes will be small and the target may be incorrectly identified as a corner. Therefore, a triad of sensors may generate a sequence <P,C,C>, the two corner decisions being taken by sensors on either side of the line of sight.

A comparison is made between the simple majority vote outcome and the decision reached when sensors group themselves according to minimum physical distance, fusing only within groups of three. Following this, fusion takes place based on the results of each group. This comparative approach investigates the importance of the *numbers* of voters supporting an outcome, which voting emphasizes, as opposed to the *strength of belief* of voters, which is significant in Dempster's fusion rule.

Initially, five groups of three sensors are formed, based on minimum physical distance: (1–2–3), (4–9–10), (5–6–13), (7–8–11), and (12–14–15). The percentage of correct classification within each group using Dempster's rule of combination was 52.14%, 49.57%, 48.29%, 52.56%, and 70.51%, respectively. The total correct percentage after fusion over the five groups was 86.75%. Using simple majority voting, the same groups yielded 77.78%, 72.65%, 84.19%, 69.66%, and 91.03%, respectively. Taking the majority vote in each group as a vote, the total correct percentage of decisions after voting over the five groups was 89.32%. Note that with voting, success rates of individual groups were much larger than those achieved when Dempster's rule was applied. The overall average was also slightly higher.

Further tests were performed, this time using three groups of five sensing nodes. The groups were selected as: a group of four nodes at minimum distance and one at the furthest distance (2–3–4–5–15); a line configuration (14–6–1–10–15); and a star configuration (1–14–15–8–12); (1–2–3–8–15); (7–2–1–4–11); (13–5–1–3–9); (7–1–9–13–11); and (1–2–6–5–15).

The percentages for correct classification within each group using Dempster's fusion rule and using voting are shown below:

Sensor Grouping	Dempster's Fusion Rule	Voting
(2–3–4–5–15)	63.25%	76.07%
(14–6–1–10–15)	81.20%	89.74%
(1–14–15–8–12)	87.18%	91.88%
(1–2–3–8–15)	58.97%	82.48%
(7–2–1–4–11)	71.37%	76.92%
(13–5–1–3–9)	63.68%	84.62%
(7–1–9–13–11)	69.66%	80.34%
(1–2–6–5–15)	76.07%	85.04%

In all of these groupings, voting gives a higher total correct-decision percentage. The results also show the importance of grouping: the group (1–14–15–8–12) gives a higher correct-decision percentage than all 15 sensors.

The superior performance of voting is partly explained by its relative insensitivity to outliers. Further fusion tests were performed using a group of four sensors that are in agreement. In this case, fusion with Dempster's rule yields a slightly higher correct-decision percentage: 84.62% compared with 83.33% for voting. At this point, a dissenting sensor is introduced and fusion is performed over the five node values. The voting percentage is stable at 83.33%, but fusion using Dempster's rule shows a marked decline, from 84.62% to 58.97%. This demonstrates the relative benefit of strength of belief where sensors are in agreement, as opposed to numbers supporting an outcome. For small sets of sensors, unanimity is favored by feature fusion using Dempster-Shafer methods, but the introduction of dissent motivates a more robust approach. Significant improvement in decision accuracy can be achieved using simple majority voting.

4. Discussion

The resolution of conflicting data from multiple sources is fundamental to a host of problems in robotics and sensing. This is particularly the case in problems such as classification, where a fused outcome must reflect the group's consensus rather than a compromise value. The voting strategies implemented provide means of resolving such issues, for example, when determining features in map building by a mobile robot or identifying targets in localization problems. However, the main extension of this work is to a system for programming robots to learn about environmental features and system behavior.

In the experimental work, planes and corners are identified by mutually exclusive rules. The nature of the discrimination algorithm (Barshan and Kuc 1990) makes it a good candidate for voting fusion, since it partitions the data into two sets based on target type. In situations involving a greater number of choices, determining a fair voting strategy can be a problem (Arrow 1951). This can make voting a less attractive strategy. However, a voting approach may be the best solution, at least initially, in certain problems. Work by Dawes (1979) shows that counting (or voting) strategies can be the best starting point in problems where reliable covariance estimates are unavailable. This suggests a role for voting in sensor fusion, as an indicator of the reliability of underlying assignments as learning is incorporated in a fusion process. The incorporation of learning can improve the performance of fusion methods. Voting can be used to gauge such improvement. Comparison of the performance of the fusion strategy relative to that achieved through voting can be a guide to the extent to which further learning must be incorporated.

Inherent in data analysis is some form of data compression, abstraction of the raw data so that inferences may be drawn. In doing this, a system is essentially determining a set of surrogates for the data or decision points. This principle lies behind the selection of principal components (Maxwell 1977) or dimensionality reduction (Duda and Hart 1973). The process of surrogate selection is made explicit in the organization by grouping as detailed above. By selecting a criterion for grouping and testing the fused results, the sensor surrogates for the particular problem are determined. Alternatively, by eliminating successive members of the decision group, the system can be reduced to a set of informative decision makers, whose members can stand in for the group as a whole. Organization of the sensor system in this way, by analysis, is a means of extracting information from the observations such that a coherent and consistent description of the environment is formed with the minimal number of nodes or decision points. This is of importance where decisions must be made in real time, or with little room for conflict, or where there is a possibility of large outliers (for example, through sensor failure). Even where all sensors are used in the fusion process, the surrogates form a natural means of validating the sensor decisions.

Validation through surrogates is a useful technique for sensor systems of the dimension typically found in mobile robots. This is indicated by drawing an analogy to cross-validation in learning (Chmielewski and Grzymalabusse 1996; Turney 1994a, 1994b). Stability of voting in learning is investigated by Turney (1994a, 1994b). In learning, using small data sets, one or more elements can be left out for testing, the bulk of the data being used for training (Napolitano et al. 1996; Opper and Winther 1996). The process of cross-validation can be performed leaving out one element or many, as in k-fold cross-validation (Turney 1994a, 1994b). By varying the partitions of the data, in other words, leaving out a different element each time, the limitation in the size of the test data set can be overcome. In the sensor decision problem, a criterion can be selected for the combination of the group views and a distance metric defined to relate the decision of a single sensor to that of the group. The decision criterion essentially orders or queues the sensors. Initially, the views of all of the sensors can be combined. Sensors can then be eliminated from the decision process in the order determined by the decision criterion. If a sensor's decision has little impact on the group decision, this is reflected in insignificant change (or increase) in the level of accuracy of the group determination of target type when that sensor is eliminated from fusion. The point where exclusion leads to a decrease in accuracy can be used to determine when the exclusion of sensors can be terminated. It is arguable that in some cases, dissenting sensors should be excluded from the group decision. (For example, in the previous section, a higher decision percentage was achieved for the sensor group (1-14–15–8–12) than for the total population of 15 sensors.) The elimination strategy described above is a form of data validation, as it allows the group decision to be taken as the hypothesis (about target type, in the experimental example given here). The surrogate decisions form a *strategic* partition of the sensor data, one which is learned through the analysis process and which can be exploited to validate other sensors' decisions.

The selection of a group of surrogates by criterion involves learning about the informativeness of the decision makers. This information is used in subsequent fusion steps to determine outliers and resolve conflicts. Learning can be applied to identify better grouping criteria for sensor nodes in fusion. The elimination procedure detailed above presents an effective strategy for fault-tolerant operation by allowing all of the sensor decision makers to be tested against the group hypothesis with their exclusion. For the classification problem, results demonstrated that the form of this test could profitably be a vote on outcome drawn from decisions of the individual sensor nodes.

5. Conclusion

This work presents a novel and comparative application of the theory of evidence and voting to target classification in robotics. Many problems in robot programming and control rely on multisensor solutions. An inevitable consequence of exploiting sensor diversity is that conflicts can arise. These must be resolved through the fusion process. In some cases, the group decision is profitably informed by appealing to the opinions of a number of surrogates—sensors selected for their informativeness in decision making. In this paper, voting is suggested as a means of resolving sensor conflicts in multinode decisions for certain classes of problems in robot programming. In the experimental example of this work, target features are generated as being evidentially tied to degrees of belief. The target decisions of multiple sonar sensors at distinct geographical sites are fused. Using both time-of-flight and amplitude data in the feature fusion process allows more-robust differentiation. The evidential reasoning approach is contrasted with combination of sensor beliefs by voting. A difference between the two fusion approaches exploited is the effect on eventual outcome of the numbers as opposed to the strength of belief of the sensors supporting an outcome. In voting, the belief level is abstracted, and emphasis is placed on the number of voters supporting an outcome. For the simple targets in our room, in most cases, resolution of conflicts by majority decision achieved more-robust target decision making. Initially, the voting process takes the view of the majority as a surrogate for the view of the whole network of sensor nodes. In subsequent analysis, other criteria are determined for formation of the surrogate group. The aim is to improve the classification process through the generation of sets of informative decision makers whose views can stand in for the views of the system as a whole. This parsimony in network exploitation is useful in resolving conflicts between sensors and in validating sensor information. As demonstrated in the experimental work, both fusion methods are suitable for real-time applications in which multiple sensing sites are used. The results give ground for the use of strategic voting as a technique for the resolution of problems in robotics where the views of multiple sensing agents or robots must be reconciled; for example, in surveying an unknown environment composed of primitive target types. The proposed fusion method could be extended to include physically different sensors such as infrared and laser-ranging systems for map building, target identification, localization, and tracking applications. The fusion method can also be enhanced by the incorporation of learning. The use of learning for classification of features from sonar data is being investigated in a work in progress (Utete, Barshan, and Srinivasan, in preparation). Learning could assist also in the definition of more effective grouping criteria for the sensor nodes. Future work will investigate learning strategies for more robust decision making.

Acknowledgments

This collaboration was supported by NATO collaborative research grant CRG.951262. The experiments were performed at Bilkent University Robotics and Sensing Research Laboratory, which is supported in part by TÜBİTAK under grant EEEAG-92. The first author was supported by the Lady Wolfson Junior Research Fellowship in Engineering Science at St. Hugh's College, Oxford. The third author is partially supported by a TÜBİTAK-BDP grant.

The authors would like to thank the anonymous reviewers for their comments.

References

Arrow, K. J. 1951. *Social Choice and Individual Values*. New York: John Wiley.

Ayrulu, B. 1996. Classification of target primitives with sonar using two nonparametric data-fusion methods. Master's thesis, Department of Electrical Engineering, Bilkent University, Ankara, Turkey.

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.

Chmielewski, M. R., and Grzymalabusse, J. W. 1996. Global discretization of continuous attributes as preprocessing for machine learning. *Intl. J. Approx. Reasoning* 15(4):319–331.

Chong, C., Mori, S., and Chang, K. 1990. Distributed multitarget-multisensor tracking. In Bar-Shalom, Y. (ed.) *Multitarget-Multisensor Tracking: Advanced Applications*. Norwood, MA: Artech House, pp. 247–295.

Dawes, R. 1979. The robust beauty of improper linear models in decision making. *Am. Psychologist* 34(7):571–582.

Duda, R. O., and Hart, P. E. 1973. *Pattern Classification and Scene Analysis*. New York: John Wiley.

Durrant-Whyte, H. F. 1987. Consistent integration and propagation of disparate sensor observations. *Intl. J. Robot. Res.* 6(3):3–24.

Durrant-Whyte, H. F., Rao, B. S. Y., and Hu, H. 1990. (Cincinnati, OH, May 13–18). Toward a fully decentralized architecture for multisensor data fusion. *Proc. of the IEEE Intl. Conf. on Robot. and Automat.* New Jersey: IEEE, pp. 1331–1335.

Foresti, G. L., Murino, V., Regazzoni, C. S., and Trucco, A. 1997. A voting-based approach for fast object recognition in underwater acoustic images. *IEEE J. Ocean. Eng.* 22(1):57–65.

Grime, S., Durrant-Whyte, H. F., and Ho, P. 1991. Communication in decentralized data-fusion systems. *Proc. of the Am. Control Conf.* New Jersey: IEEE, pp. 3299–3305.

Harary, F., Norman, R., and Cartwright, G. 1965. *Structural Models: An Introduction to the Theory of Directed Graphs*. New York: John Wiley.

Ho, Y.-C. 1994. Heuristics, rules of thumb and the 80/20 proposition. *IEEE Trans. Automatic Control*. 39(5):1025–1027.

Kim, K. H., and Roush, F. W. 1980. *Introduction to Mathematical Consensus Theory*, vol. 59 of *Lecture Notes in Pure and Applied Mathematics*. New York: Marcel Dekker.

Klein, L. A. 1993a. A Boolean algebra approach to multiple sensor voting fusion. *IEEE Trans. Aerospace Electron. Sys.* 29(2):317–327.

Klein, L. A. 1993b. Voting fusion, chapter 5. In O'Shea, D. C. (series ed.) *Sensor and Data Fusion Concepts and Applications*, vol. TT 14 of *Tutorial Texts in Optical Engineering*. Bellingham, WA: SPIE Optical Engineering Press, pp.73–90.

Krause, P., and Clark, D. 1993. *Representing Uncertain Knowledge: An Artificial Intelligence Approach*. Oxford: Intellect Books.

Luo, R. C., and Lin, M.-H. (1987, St. Charles, IL, October 5–7). Multisensor integrated intelligent robot for automated assembly. *Proc. of the 1987 Workshop on Spatial Reasoning and Multisensor Fusion*. Los Altos, CA: Morgan-Kaufmann, pp. 351–360.

Mao, J. C., Flynn, P. J., and Jain, A. K. 1995. Integration of multiple feature groups and multiple views into a 3-D object-recognition system. *Comp. Vision Image Understanding*. 62(3):309–325.

Maxwell, A. E. 1977. *Multivariate Analysis in Behavioural Research*. London: Chapman and Hall.

Murphy, R. R. (1996, Washington DC, December 8–11). Adaptive rule of combination for observations over time. *Proc. IEEE/SICE/RSJ Intl. Conf. on Multisensor Fusion and Integration for Intelligent Systems*. New Jersey: IEEE, pp. 125–131.

Napolitano, M. R., Casdorph, V., Neppach, C., Naylor, S., Innocenti, M., and Silvestri, G. 1996. Online learning neural architectures and cross-correlation analysis for actuator failure-detection and identification. *Intl. J. Contr.* 63(3):433–455.

Opper, M., and Winther, O. 1996. Mean-field approach to Bayes learning in feedforward neural networks. *Phys. Rev. Lett.* 76(11):1964–1967.

Pagac, D., Nebot, E. M., and Durrant-Whyte, H. F. 1996. (Minneapolis, MN, April 22–28). An evidential approach to probabilistic map building. *Proc. of the 1996 IEEE Intl. Conf. on Robot. and Automat.* New Jersey: IEEE, pp. 745–750.

Panasonic Corp. 1989. Ultrasonic ceramic microphones. 12 Blanchard Road, Burlington, MA 01803, USA.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Francisco: Morgan-Kaufmann.

Reece, S. 1997 (Orlando, FL, April 20–25). Qualitative model-based multisensor data fusion and parameter estimation using ∞-norm Dempster-Shafer evidential reasoning. *SPIE Proc.*, vol. 3068. SPIE.

Rosenblatt, J. K. 1997. DAMN: A distributed architecture for mobile navigation. *J. Exp. Theoret. Art. Intell.* 9(2–3):339–360.

Shafer, G. 1976. *A Mathematical Theory of Evidence*. Princeton: Princeton University Press.

Sukthankar, R. 1997. Situation awareness in tactical driving: The role of simulation tools. *Trans. Soc. Comp. Simulation* 14(4):181–192.

Takano, T., Yamada, T., Shutoh, K., and Kanekawa, N. 1996. In-orbit experiment on the fault-tolerant space computer aboard the satellite Hiten. *IEEE Trans. Reliability* 45(4):624–631.

Tan, A. H., and Teow, L. N. 1997. Inductive neural logic network and the SCM algorithm. *Neurocomputing* 14(2):157–176.

Toye, G., and Leifer, L. J. 1994. Hellenic fault-tolerance for robots. *Comp. Electrical Eng.* 20(6):479–497.

Turney, P. 1994a. Theoretical analysis of cross-validation error and voting in instance-based learning. *J. Exp. Theoret. Art. Intell.* 6(4):331–360.

Turney, P. 1994b. A theory of cross-validation error. *J. Exp. Theoret. Art. Intell.* 6(4):361–391.

Utete, S. W. 1995. Network management in decentralised sensing systems. PhD thesis, Department of Engineering Science, University of Oxford.

Utete, S. W., Barshan, B., and Srinivasan, A. (In preparation.) Classification by learning for sonar sensors.