Spatial Cognition and Computation **1:** 205–226, 1999.

Using fuzzy sets to represent uncertain spatial knowledge in autonomous robots *

JORGE GASÓS[1] and ALESSANDRO SAFFIOTTI[2]
[1]*IRIDIA, Université Libre de Bruxelles, 50 av. F. Roosevelt, CP 194/6, B-1050 Brussels, Belgium (e-mail: jgasos@iridia.ulb.ac.be);* [2]*Applied Autonomous Sensor Systems, Dept. of Technology, University of Örebro, S-70182 Örebro, Sweden (e-mail: asaffio@aass.oru.se)*

Abstract. Autonomous mobile robots need the capability to reason from and about spatial knowledge. Due to limitations in the prior information and in the perceptual apparatus, this knowledge is inevitably affected by uncertainty. In this paper, we discuss some techniques employed in the field of autonomous robotics to represent and use uncertain spatial knowledge. We focus on techniques which use fuzzy sets to account for the different facets of uncertainty involved in spatial knowledge. These facets include the false measurements induced by bad observation conditions; the inherent noise in odometric position estimation; and the vagueness introduced by the use of linguistic descriptions. To make the discussion more concrete, we illustrate some of these techniques showing samples from our work on mobile robots.

Key words: environment modeling, fuzzy logic, linguistic descriptions, robot navigation, self localization, spatial maps, uncertainty management

1. Introduction

Current research on autonomous mobile robots aims at building physical systems that can move purposefully and without human intervention in natural environments – that is, real-world environments that have not been specifically engineered for the robot. Despite the recent advances in this field, autonomous robot navigation still requires a better understanding of the processes involved in the perception and representation of space. The use of general purpose sensors that are fixed on the robot, the lack of structure in these environments, and the need to find solutions that do not depend on one specific environment or task, impose strong limitations on the information that will be available to the robot.

Most of these limitations originate in the nature of real-world, natural environments. First, prior knowledge about the environment is, in general,

* A preliminary version of this paper was presented at the Mind III Conference, Dublin, Ireland, August 1998.

incomplete, uncertain, and approximate. For example, maps typically omit some details and temporary features, spatial relations between objects may have changed since the map was built, and the metric information may be imprecise and inaccurate. Second, perceptually acquired information may be unreliable. The limited range of the sensors, combined with the effect of adverse observation conditions (e.g., poor lighting or occlusions), leads to noisy and imprecise data; and errors in the measurement interpretation process may lead to incorrect beliefs about the identity and the position of objects. Finally, real-world environments typically have complex and unpredictable dynamics: objects can move, other agents can modify the environment, and relatively stable features may change with time (e.g., seasonal variations). Despite this pervasive presence of uncertainty, humans and animals are able to deliver excellent performance in natural environments. An important issue in this respect is the ability to build and use partial models of the environment which can accommodate a large amount of uncertainty.

In this paper, we explore this issue by looking at ways in which uncertainty in spatial information can be handled in autonomous mobile robots. We focus in particular on the use of the theory of fuzzy sets. Accordingly, the next Section shortly reviews the use of fuzzy sets for representing uncertain information. We then present different methods to manage spatial information affected by uncertainty in a mobile robot. We consider two separate aspects. First, we focus on the representation of spatial information, and on how to build such a representation, or *map*, from perceptual data. We consider different approaches, classified according to the level of abstraction in the information they use. Second, we discuss how to use these representations in a robot, focusing on two problems: self-localization, and the integration of environment models with other sources of spatial information like linguistic descriptions and image data.

This paper brings together several separate pieces of work by the authors, and connects them to provide a global view of different approaches to manage uncertain spatial knowledge, and of how they may be combined. All the individual approaches have been reported in greater detail in the technical literature (Saffiotti and Wesley, 1996; Gasós and Rosetti, 1999; Gasós and Saffiotti, 1999; Gasós, 2000). The goal of this paper is not to explain the technicalities of any of these approaches, for which we address the reader to the references. Rather, our aim here is to provide a glimpse, accessible to people outside the fields of robotics and artificial intelligence, on the ways in which spatial knowledge is generated, represented, and used in autonomous robots, with a special emphasis on the problem of uncertainty. Our hope is that this glimpse will provide a small contribution to the communication and cross-fertilization between areas that attack the problem of spatial cognition from different viewpoints.

2. Representing uncertainty

Before we embark in our excursion, we need to better clarify the intended meaning of the term “uncertainty”. Strictly speaking, uncertainty is not a property of information, but rather a property of an agent – or, more precisely, of the agent’s mental state. An agent may be uncertain about the existence of an object, the value of a property, the truth of a proposition, or the opportunity to perform an action. Yet, it is common to talk of uncertain information. We maintain that what is actually meant by this is information which is “weak” in some respect, thus inducing a state of uncertainty in the agent. Consider a robot that needs to charge its batteries before continuing with its task; this requires that the position of the battery charger is known with a high degree of precision – how high depends on the robot’s coupling mechanism. The item of information

The battery charger is in room 2 (1)

is too weak for this task, since it does not give us a unique position; we talk in this case of *imprecise* information. The item

The battery charger is close to the left wall (2)

is *vague*, since it does not give us a crisp position. And the item

The battery charger was seen yesterday at coordinates $(5, 3)$ (3)

is *unreliable*, since the charger may no longer be there. All these items may be regarded as uncertain in this context, as they leave the agent in a state of uncertainty about the actual position of the charger. Note that considering an item of information as uncertain depends on the specific task: for instance, item (1) would not induce any uncertainty if the task were to simply move to the same room as the charger.

Fuzzy sets are particularly well suited to represent and reason with weak knowledge. A *fuzzy set* (Zadeh, 1965) extends a classical set in that membership of an element to the set is not a true/false property, but it is a matter of degree, usually measured by a number in the $[0, 1]$ interval. Techniques based on fuzzy sets are able to represent and distinguish different facets of uncertainty (Klir and Folger, 1988; Ruspini, 1991). For example, we may

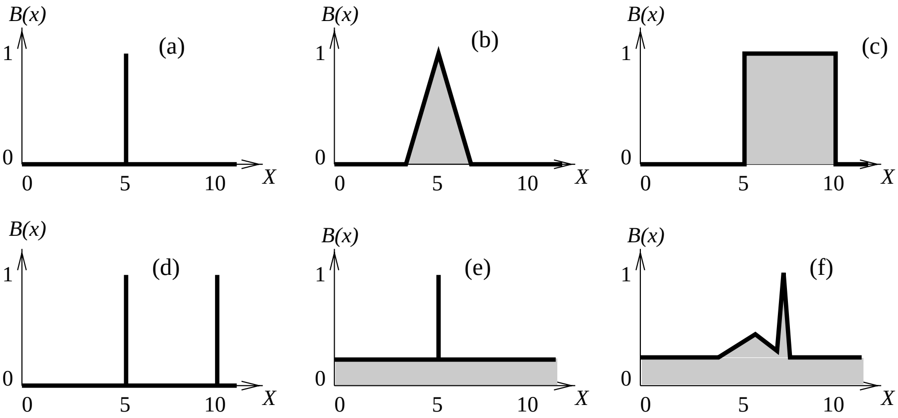

Figure 1. Representing different types of positional uncertainty for an object B by fuzzy sets: (a) crisp; (b) vague; (c) imprecise; (d) ambiguous; (e) unreliable; (f) combined. (Adapted from Saffiotti and Wesley (1996)).

wish to distinguish between the vagueness or inaccuracy in the position of an object, and the uncertainty in its very existence – e.g., the map may be wrong, the object may have been removed from the environment, or its existence may have been inferred from a spurious sensor reading.

Figure 1 shows how we can represent different types of weak information about the location the battery charger by using a fuzzy subset B of the set X of possible positions, taken here in one dimension for graphical clarity. (See, e.g., Bloch (2000) for more complex uses of fuzzy sets to represent spatial information.) In this example, we take a *possibilistic* reading of fuzzy sets (Zadeh, 1978): for each $x \in X$, we read the value of $B(x)$ as the degree of possibility, measured on a $[0, 1]$ scale, that x be the actual location of the charger. In (a) this location is crisp and certain: exactly one location, 5, is possible. Item (b) tells us that the charger is located at approximately 5: this is a situation of vagueness, similar to the one expressed by statement (2) above. In (c), the charger can possibly be located anywhere between 5 and 10: we talk here of imprecision. In (d), it can be either at 5 or at 10 (ambiguity). The situation in (e) is a case of unreliability, similar to the one expressed by statement (3) above: we are told that the charger is at 5, but the information may be invalid, so we put a small "bias" of possibility that it be located just anywhere. Finally, (f) combines vagueness, ambiguity and unreliability. Note that the case of total ignorance is represented by $B(x) = 1$ for all $x \in X$ – as far as we know, the charger can be anywhere.[1]

Fuzzy sets have gained popularity in recent years as a representation for uncertainty which is, in a sense, complementary to probability – see Saffiotti (1997) for an overview of their use in autonomous robotics. Three characteristics of fuzzy sets may make them the favorite choice in some cases: (i) their ability to represent total ignorance; (ii) their appropriateness to represent linguistic expressions (Zadeh, 1975); and (iii) their ease of integration with logical systems, e.g., fuzzy rule-based systems. (The last two properties are partly responsible of the success of the field of fuzzy control.) Needless to say, a probabilistic representation remains the best choice to represent the uncertainty originating by phenomena which are stochastic in nature, and for which we are able to give precise probabilistic assessments (Smith and Cheeseman, 1986). Since most of the uncertainty in the autonomous robotics domain is not of a stochastic nature, and cannot be easily characterized in probabilistic terms, we shall focus on techniques based on fuzzy sets in the rest of this paper.

3. Representing spatial information

Environment information can be represented at several levels of abstraction, ranging from a purely qualitative representation of major objects (e.g., rooms in a building) and of their topological connections, all the way down to a fully detailed geometrical representation of the environment. A common belief in the robotics field is that robots need to represent and reason about information at different levels of abstraction. There are several reasons for this. First is epistemic adequacy: different tasks call for different types of representation. For example, global navigation strategies are more easily planned using a topological map, where we can decide the sequence of rooms and corridors to be traversed; but fine motion control needs geometric information to precisely control navigation among features and obstacles. Second, computational adequacy: geometric information is difficult to collect and expensive to handle, and we cannot pay the price to maintain a detailed geometric representation of the entire environment where the robot can operate. The final reason is ontological adequacy: fine grained information is difficult to obtain *a priori* and is likely to change with time; coarse maps are easier to estimate and more prone to remain valid over time.

Several definitions of the levels of abstraction needed to represent spatial knowledge in robotics have been proposed. One of the earliest and most influential analysis is offered by Kuipers (Kuipers, 1978; Kuipers and Levitt, 1988), whose work was motivated by insights from research on human spatial reasoning. Most of today's work on environment modeling in robotics can be classified into four levels.

Occupancy level At this level, the principal concern is to distinguish between free and occupied areas of the space; representations at this level are typically in a grid form (Moravec and Elfes, 1985; Elfes, 1987; Leonard et al., 1992; Tirumalai et al., 1995; Oriolo et al., 1998), and are mainly used for trajectory planning and local collision avoidance.

Geometric level At this level, we represent geometric features like the contours of objects; these representations are typically in a segmented form (Latombe, 1991; Tunstel, 1995; López-Sánchez et al., 2000), and are mainly used to recognize objects and object configurations, possibly in order to self-localize.

Semantic level At this level, we represent objects which are classified according to their semantic type, e.g., a door or a table; these representations are most adequate for abstract reasoning tasks, like task planning or situation assessment (Saffiotti et al., 1995; Surmann and Peters, 2000).

Topological level At this level, we represent abstract relationships between objects in the environment, like adjacency or containment, without using an absolute reference system; this representation is adequate for generating and communicating navigation strategies in a compact and effective form (Kuipers and Byun, 1991; Kortenkamp and Weymouth, 1994; Thrun, 1999; Gasós and Saffiotti, 1999; Fabrizi and Saffiotti, 2000).

Whatever the level of abstraction, an important aspect of a map representation is the way in which it accounts for the uncertainty in its properties. Historically, most of the techniques developed in the robotic domain have been based on a probabilistic representation of uncertainty. A probability-based representation is adequate when two conditions hold: (i) the underlying uncertainty can be given a probabilistic interpretation; and (ii) the data required by probabilistic techniques is available. Both conditions may be violated in the case of autonomous robots, and several techniques based alternative formalisms have recently been developed. For instance, among the proposals listed above, (Tirumalai et al., 1995) uses Dempster-Shafer's theory of evidence (Shafer, 1976; Klir and Folger, 1988); while (Oriolo et al., 1998; Tunstel, 1995; López-Sánchez et al., 2000; Saffiotti et al., 1995; Surmann and Peters, 2000; Gasós and Saffiotti, 1999; Fabrizi and Saffiotti, 2000) use fuzzy set theory.

In the discussion that follows, we give a few concrete illustrations of techniques at the geometric and the semantic levels which are based on fuzzy sets. These two levels of abstraction possess complementary characteristics: the geometric level is close to the level of abstraction used by the sensory-motor processes, and requires a moderate perceptual effort; while the semantic level is at the level of abstraction used by symbolic reasoning, and it requires a significant perceptual and interpretation effort. The techniques that we consider in this section have been applied to real robots equipped with sonar sensors.

3.1. *Geometric level*

At this level of abstraction we find maps that contain geometric information. Here, sensor readings, that provide information on the existence of objects in certain points of space, are organized into more abstract features such as lines or curves. This is a compact representation that approximates the boundaries of the objects and defines their spatial location, thus, extracting the layout of the environment. It is very efficient in large environments and when the application requires detailed information about position, shape or size of the objects.

Geometric maps can be generated by a careful examination of successive sensor observations. As the robot moves, the consecutive measurements of each sonar sensor provide a discrete sequence of points that forms an outline of the objects located in the proximity of the robot trajectory. In this section, we consider the technique proposed in Gasós and Rosetti (1999), where we build geometric maps from the analysis and aggregation of consecutive measurements according to the following criteria: a sequence of smoothly changing observations indicates that the side of an object is being detected; a discontinuity that starts a new smooth sequence indicates that a new object has been detected; and a single measurement, not related to its neighbors, can be considered noisy and discarded. The approach is based on the following sequence of steps:

– Sensor measurements are preprocessed to eliminate easily detectable misreadings and to group together smoothly changing consecutive measurements from the same sensor.
– Each group of measurements is fitted to a straight line (segment) and the uncertainty in its position and orientation is represented using fuzzy sets.
– Fuzzy segments corresponding to the same side of an object are merged to obtain a single boundary. Uncertainty is propagated from the fuzzy segments to the fuzzy boundaries.

Uncertainty representation is a key aspect of this approach. Boundary extraction from sensor observations is a process that works with inaccurate and noisy data, hence, generating uncertainty on the position, orientation and size of the boundaries. Given a finite set of observations coming from the same side (boundary) of an object, we approximate them by a straight line while representing the uncertainty in the knowledge of its real posi-

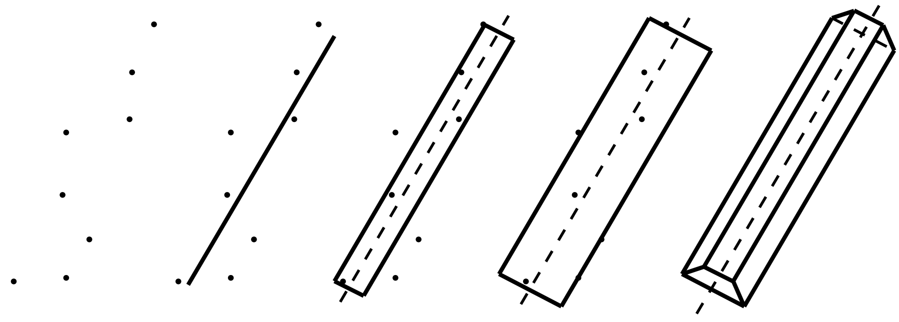

Figure 2. Building a fuzzy segment from a set of scattered points.

tion by a fuzzy set. The degree of membership of a point to the fuzzy set indicates the possibility that this point belongs to the object detected by the sensor. Extending this definition to the case of lines (boundaries), the degrees of membership can also be interpreted as degrees of similarity when two different boundaries are compared. Thus, a low (high) overlap of their membership functions will indicate that they should be considered similar to a low (high) degree. This representation allows us to express the information/uncertainty contained in the sensor data and is closer to the main operation performed on the boundaries in the process of map building: merging boundaries obtained from different sensors or from different positions in the trajectory of the robot.

On the way to build the fuzzy segments, we can observe that different factors influence the level of uncertainty on the location of a segment:

- the relative orientation between sensor and object,
- surface properties and shape of the object,
- where in the sonar beam the ultrasonic signal was reflected,
- the distance between sensor and object,
- ambient temperature and air turbulence,
- dead reckoning estimation of the robot location.

These factors can be classified according to their source of uncertainty. For the first three factors, uncertainty is due to the lack of knowledge on the way the ultrasonic signal reflected on the object, and this implies that their influence on the observations cannot be estimated directly. Since, for example, the texture of the object that reflected the signal is not known, it is not possible to model its effect on the observations. The only way to estimate the influence of these factors is through the scatter they produce when the observations are fitted to form a segment (Figure 2a,b). A good fit to the line indicates that the object surface and the measuring conditions were favorable, while scattered

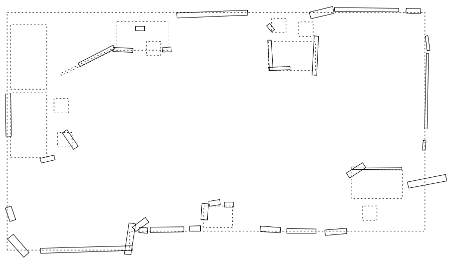

Figure 3. A view of the experimental site, and the corresponding fuzzy model (solid lines) superposed to the real map (dotted lines).

points indicate a high level of uncertainty. The concept of confidence interval, developed in statistics to report uncertainty in the value of a parameter, is used to build fuzzy sets from scatter information. For a given confidence level, the interval provides an estimation of the region within which the segment is likely to lie. In the same way, α-cuts (Klir and Folger, 1988) in the fuzzy set define intervals (regions) within which all segments can be considered similar at least to the degree α (Figure 2c,d). Relating both concepts, we obtain the degrees of membership to the fuzzy sets in terms of the confidence intervals (Figure 2e). A similar analysis is performed to represent the uncertainty due to the three other factors.

Each fuzzy segment provides information on the position and orientation of an object boundary. Since the segments are obtained from different sensors or from different positions of the robot, a key aspect of the process of map building is detecting and combining the segments that may have come from the same boundary. Collinearity is the property that defines the coincidence between two segments but, given the uncertainty in their position, it cannot be defined as a strict overlap between them. Here again, uncertainty representation using fuzzy sets allows a certain flexibility that is related to the quality of the observations and measuring conditions. Its interpretation in terms of degrees of similarity allows one to immediately detect whether two segments correspond to the same object just by computing their degree of overlap. Segments that are found to be collinear are combined and their respective uncertainties are propagated in the integration process.

Figure 3 shows a map built by the robot after navigating in a previously unknown environment. The experimental site, a large rectangular room of 11.6 m $\times$ 6.1 m, contained objects that are normally found in indoor office environments: a printer, three desks, a partition, two tables and six chairs. Figure 3 shows the α-cut of degree 0 of the fuzzy segments, superposed on the real map. The segments in this α-cut provide a good approximation of the environment spatial layout. The outline of the objects with compact and flat surfaces was inferred correctly, and the discontinuities in their shape have been detected. The gaps in the model correspond to areas that were occluded by other objects (a chair or a wall behind a desk cannot be detected), to sides of objects with an orientation perpendicular to the robot trajectory (only the front side of some objects is included in the model), or correspond to non-compact objects, which are difficult to detect by the ultrasonic sensors (tables and chairs generate a small number of segments).

Perhaps the most striking difference between our maps and other geometric representations based on contours (e.g., the polygonal obstacles considered by Latombe (1991)) is the use of fuzzy sets to represent different facets of uncertainty, as discussed above.

The use of fuzzy sets is especially convenient in this case since we need to measure the similarity between segments that have been detected under different measuring conditions: from different sensors or from different positions of the robot. The semantic analysis proposed by Ruspini (1991) shows that fuzzy sets are naturally related to the notion of similarity. This notion of similarity, together with the representation of spatial uncertainty, allows us to immediately detect the degree of collinearity of two segments just by computing their degree of overlap. Fuzzy sets also lend themselves to efficient implementations: computing the degree of matching of n fuzzy segments can be done with very little computation, and the full algorithm for map building, map updating and self-localization (see below) using fuzzy sets can be run continuously inside the 100 millisecond control cycle of our robot, independently on the size of the environment. This contrasts which popular probability-based techniques (e.g., Thrun et al., 1998), that reportedly require computation times of up to one hour in large environments.

3.2. *Semantic level*

At the semantic level, the basic objects that we represent are semantically meaningful objects of the domain: in an indoor environment, these typically include doorways, corridors, halls, and rooms. Each entity is associated to a set of properties which are relevant to the tasks performed by the robot. These often include positional, metric, and topological properties. For example, a doorway is typically associated to the coordinates of its middle point, its width, and an indication of which other objects (rooms and/or corridors) are connected by it. Uncertainty can be attached to any of these properties. In what follows, we focus on the use of fuzzy sets to represent uncertainty about the location property. This discussion follows the proposal presented in Saffiotti and Wesley (1996).

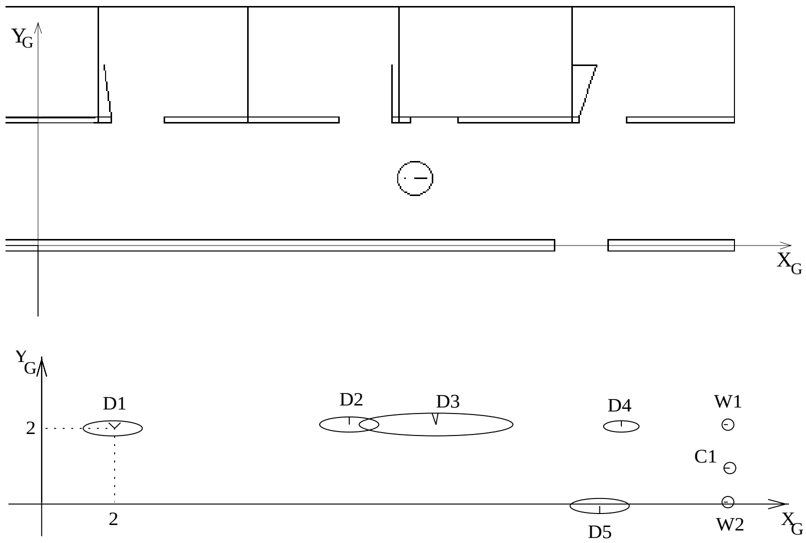

Figure 4. An office environment (top), and an approximate map of it (bottom).

We represent maps by sets of semantic features. Each feature has a position and an orientation, which are given by a triple (x, y, θ) of coordinates in a global Cartesian frame G, where θ is the orientation with respect to the X axis. These coordinates refer to a given point of the feature which is conventionally taken as its representative. For example, the position of a door is given by its center point, and its orientation is given by its axis; for a wall, we consider a point at one extreme and oriented like the wall. In order to account for positional uncertainty, we associate each feature to an approximate position, that is, a fuzzy subset of the three dimensional space $X \times Y \times \Theta$. If P_f is the fuzzy position of some feature f, then, $P_f(x, y, \theta)$ is a number in [0, 1] that measures the degree of possibility that (x, y, θ) be the actual position and orientation of f, in a similar way to the example given in Section 2.

Figure 4 shows (a part of) an office environment, and the approximate map built by our robot to model it. Each ellipsoid in the map represents the fuzzy position of an element. The lengths of the two radii indicate the width of the bases of two triangles representing the fuzzy X and Y coordinates; the fuzzy orientation is indicated by a slice in the ellipsoid, whose width represents the width of the base of a fuzzy triangle on Theta. For instance, door D1 has possibility 1 of being located at $(2, 2, 90°)$; possibility 0 of being at $(2, 2, 60°)$ or at $(1, 2, 90°)$; and intermediate values in between. In this particular map, the positions of the two corridor walls, denoted by W1 and W2, is fairly precise, while there is pretty much uncertainty on the longitudinal position of

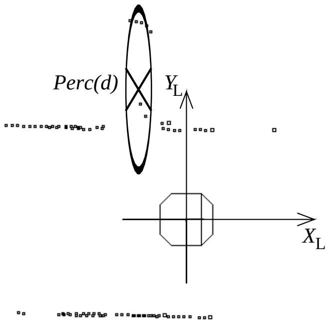

Figure 5. A fuzzy perceptual feature built by "seeing" a doorway.

the doors along the corridor. In particular, the approximate position of door D3 partially overlaps the one of D2; this implies that if the robot perceives a door in the overlapping area, it will not be able to tell whether this door is D2 or D3.

A *semantic approximate map* on a reference frame G is a triple $\langle M, \mathit{type}, \mathit{pos}\rangle$ where: M is a non-empty set of map indexes (the "names" of the objects); *type* associates each index to an element in some given set of object types (e.g., {door, wall, corridor}); and *pos* associates each object i to a fuzzy position P_i on G (its location in the map). The map may also include other properties that are not relevant in this context, e.g., topological properties. In order to reduce the representational and computational burden, we restrict the fuzzy positions P_i's to be pyramidal fuzzy sets, that is, fuzzy sets obtained by the Cartesian product of three triangular fuzzy sets (see Figure 1b) one on each of the X, Y, Θ dimensions.

Representing the position of objects in a global map is one aspect of the spatial representation. Another aspect is the representation of perceived objects in the robot's perceptual space. This space is local to the robot and usually small, being limited by the range and the field of view of the robot's sensors. Semantic feature recognition using the robot's sensors is a difficult task, and the result is inevitably affected by uncertainty. An adequate representation of the percepts, that includes the (different types of) uncertainty, is necessary in order to make intelligent use of them.

We use fuzzy sets to represent the approximate position of perceptual features. Given a perceived feature p, we represent its approximate location in some (usually local) reference frame L by a fuzzy subset of L. Once again, we use triangular fuzzy sets for computational reasons. The precise shape of the sets (i.e., the widths of the three triangles) depends on the quality of the recognition and on the type of feature. Figure 5 shows the fuzzy location $\mathit{Perc}(d)$ of a doorway percept d which was built by the robot's perceptual routines. The figure shows how the robot perceives its surroundings: the robot is drawn in the center in top view and pointing right; the dots indicate the sonar readings cumulated as the robot was running down a corridor. Because of various phenomena of beam reflection, the configuration of the sonar readings around the doorway (its “signature”) is confused and largely unpredictable. This signature only gives a reasonable indication about the position of the doorway along the longitudinal (X_L) axis, leaving its distance and orientation extremely vague; correspondingly, the percept’s fuzzy location has the shape shown in the picture. This fuzzy location may also include a “bias” to reflect the unreliability of perceptual recognition (see Figure 1e).

The most peculiar aspect of our proposal is probably the use of fuzzy sets to capture uncertainty. Although we feel that there is no ultimate conceptual reason to prefer, in this domain, fuzzy sets to other representations of uncertainty, there are a few practical reasons that justify our choice. First, the use of fuzzy locations can be smoothly integrated with fuzzy-logic based controllers. This is an important issue, as locational knowledge has eventually to be used to take decisions. Second, as we noticed already, fuzzy techniques lend themselves to efficient implementations, a critical requirement for robotic applications. Finally, the ability of fuzzy sets to adequately represent ambiguity and total ignorance allows us to postpone the resolution of uncertainty until further information, coming from other observations, is considered.

The last feature suggest that the fuzzy approach may compare favorably to probabilistic techniques (e.g., Smith and Cheeseman, 1986; Moutarlier and Chatila, 1989) in situations of extreme uncertainty: Saffiotti and Wesley (1996) give an example that supports this conjecture. It should be noted, however, that a probabilistic technique, like the Kalman filter, may be a better choice if we have enough control of the uncertainty – that is, if we have a good stochastic model of the errors, and a good initial estimate. Under these conditions, a Kalman filter can give formal guarantees like convergence and optimality that cannot be provided by current fuzzy techniques. Once again, it appears that the fuzzy and the probabilistic approach to manage uncertainty have complementary merits and demerits: fuzzy sets can adequately accommodate *weak* knowledge, but cannot guarantee strong properties on the result of the inference; while probability require the availability of *strong* knowledge, and can guarantee good formal properties (as long as that knowledge was correct).

4. Using spatial information

Once a spatial representation of the environment is available, an autonomous robot can use it to generate navigation plans, to locate a target, to self-localize, or as an help in object recognition. In this section, we outline two examples of this: self-localization, and the integration of environment models with other sources of spatial information like linguistic descriptions and image data for object recognition.

4.1. *Self-localization*

Perception happens locally, in the egocentric frame of reference of the robot. In order to ensure a correspondence between the local, observer-dependent representations of the environment built by the perceptual processes, and the global, observer-independent representations contained in a map, the robot must be able to estimate its own position with respect to this map. This is often referred to as the *self-localization* problem. Self-localization can use the information from the odometric sensors to update the robot's position as this moves. Unfortunately, odometry has cumulative errors, and the robot's odometric estimate of its position can diverge from reality without bounds. A commonly used technique to correct this problem is to combine odometry with some external information that periodically compensate for the errors: the robot compares the observed position of perceptual landmarks with their expected position, given the information in the map and the current location estimate, and uses the result of the comparison to correct this estimate. Note that this technique requires that prior information about the position of the landmarks is available.

An alternative that does not require the use of landmarks and that works even when the robot accesses a new environment without any prior information, is based on the matching of maps that were built by the robot in different moments (see Gasós and Rosetti (1999) for details). As the robot begins to move in a given environment, it starts building its *initial map* while keeping track of the areas that it has already visited. To do so, we use a coarse binary grid ($2\,\text{m} \times 2\,\text{m}$ cells) where we indicate if the robot has already visited the area covered by the cell. While the robot moves through a non-visited cell, it adds the new segments to the global map and marks the cell as visited. When it arrives to a visited area, the new segments are added to the *local map*. The local and initial maps are now compared until a significant number of local segments match those in the initial map. Then, the transformation that brings both maps together is used to estimate the dead reckoning errors accumulated in the time span between the construction of the maps. Once these errors are corrected, the uncertainty in the robot location is bounded by the uncertainty in the initial map, which depends on the size of the environment and on the precision of the dead reckoning system.

Uncertainty representation is a key aspect of this approach. The map building facility presented in Section 3.1 generates a plausible spatial layout

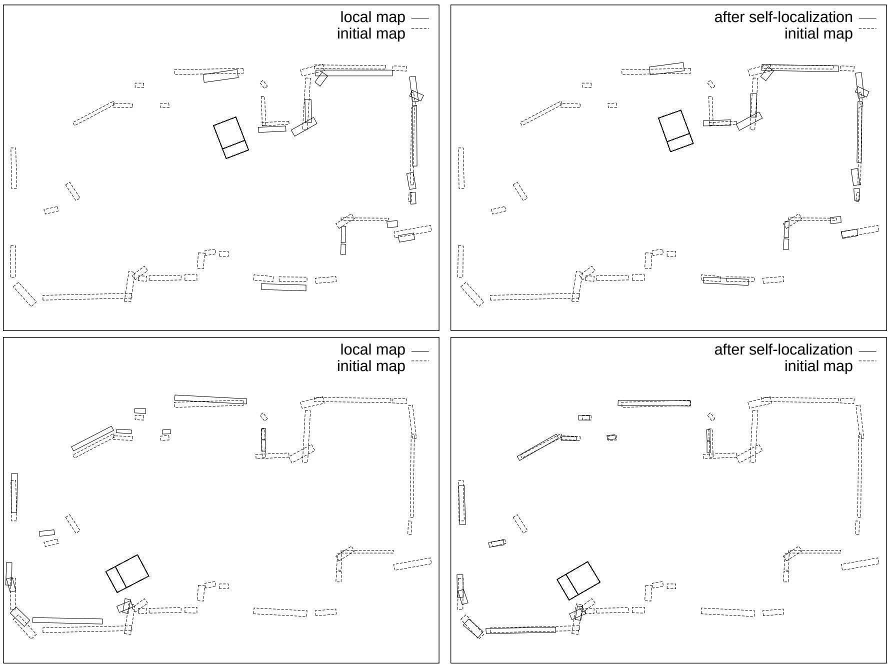

Figure 6. Robot self-localization at two moments during a navigation task. Left side: odometer-based position of the robot and of the local map. Right side: error correction after self-localization.

of the environment using fuzzy sets to represent the uncertainty on the real location of the object boundaries. The interpretation of the degrees of membership to the fuzzy sets as degrees of similarity when two boundaries are compared, facilitates the immediate detection of boundaries of the initial and local maps that come from the same object, thus allowing real time self-localization.

Figure 6 shows two steps of self-localization during a navigation experiment. The local maps are represented with continuous lines, while the dotted lines are used for the initial map. The figures on the left side show the local map generated by the robot using the estimation of the robot position provided by the odometers; the figures on the right side show the same local maps once they have been rotated and translated to match the initial map. This transformation is used to correct the errors in the estimation of the robot location after each self-localization process. In this way, each self-localization step starts from the previoulsy corrected robot position estimate, thus avoiding the effects of cumulative errors over long periods of time. The validity of the approach is guaranteed as long as the real error in the robot localization is smaller than the estimate of the positional error which is included in the constructed fuzzy segments.

Experimental evidence has shown that our continuous re-localization technique allows the robot to navigate for long periods of time without the need of additional positioning systems. Furthermore, this approach also allows to incrementally learn the spatial layout about the environment as new information becomes available in the local maps used for localization, as well as to accommodate changes in this layout due to the dynamics of the environment. Under exceptional circumstances, however, the dead reckoning error may become larger than our estimate of it, and the fuzzy segments in the local map will not match the ones in the initial map. (This may happen, for instance, if the robot goes over a large bump.) To account for these events, we have developed a global localization method that finds the best match between the local and the initial maps searching the full space of possibilities (Gasós and Rosetti, 1999). Note that global localization is relatively time consuming (of the order of seconds for large environments) and it may thus require the robot to stop. By contrast, the continuous localization method, thanks to the representation used for spatial uncertainty, can efficiently detect the segments belonging to the same object boundary and re-localize the robot in a few milliseconds.

The above approach is based on the low-level representation of space discussed in Section 3.1. Self-localization can also be performed with respect to high-level, semantic features. Saffiotti and Wesley (1996) have proposed to use the semantic representation discussed in Section 3.2 to perform approximate self-localization. The main step is to match each perceived feature with one corresponding feature in the approximate map; and to use this match to generate a fuzzy location representing the location where the robot should be in order to perceive the feature in that way. Several perceived features can be taken into account, leading to a fusion-oriented approach to self-localization: each feature is seen as a potential source of information about the robot's position; the pieces of information obtained from different features are then combined, by fuzzy intersection, into a new estimate of the robot's position.

Whether a map representation based on low-level features is better suited to self-localization than one based on high-level features depends on the type of environment and on the available sensors. On the one hand, high-level features may make the map more robust, as they are more stable over time. On the other hand, high-level features may be difficult to extract, and some environments (especially outdoor) may contain only a small number of them. In either case, however, the way in which spatial uncertainty is represented is pivotal to the ability to perform approximate matching, and to weight the contribution of each uncertain perception to the overall localization process.

4.2. *Integrating sonar maps with linguistic descriptions and image data*

Sonar maps, as the ones described above, only provide information on the existence of objects in some positions of the space but do not solve the problem of object identification, particularly when no prior environment knowledge is available. Some help in solving this problem can be obtained by using linguistic descriptions, an important source of information that has been scarcely considered in robotics in spite of its strong potential and availability. In most applications it is possible to obtain linguistic descriptions of the working area from persons who have been there before. These descriptions only provide general knowledge of the objects that might be found and of their global distribution in the environment but, once this information is combined with the sonar maps and the image data, it is enough to help object identification. Three aspects need to be considered in this process: the acquisition, the representation and the use of the linguistic descriptions (see Gasós (2000) for a detailed description).

For object identification the most relevant information is the position and size of the objects, plus their type and color. The acquisition of this information from linguistic descriptions presents some difficulties: translating 3-D spatial relations into the linear structure of speech is not easy for most describers, they may use an unlimited number of linguistic terms, and different describers may ascribe different meanings to the same term. To overcome these problems we divide the process of description in two phases: a first phase of *preparation*, where the person is given guidelines on how to improve the quality of the descriptions; and a second phase of *formulation*, where the knowledge of the environment is structured according to these guidelines.

For the preparation phase, research by linguists and psychologists in the field of spatial cognition (see Denis (1996) for an overview) has provided useful techniques that allow us to improve the quality of the descriptions. For example, the importance of discourse planning, of referential continuity and of hierarchical grouping was explained to the describers. In order to limit the number of linguistic terms and to assign them an uniform meaning, we have given the describers a set of terms that could be used in the descriptions and some indications on the way they will be interpreted. As an example, for the position of an object inside a room, each describer was asked to only use the terms {next, near, far, faraway} with respect to a wall or to a previously described object, and was provided with the sketch in Figure 7 to ensure an uniform interpretation of these terms. [2] The following is an example of the linguistic descriptions provided by one of the describers:

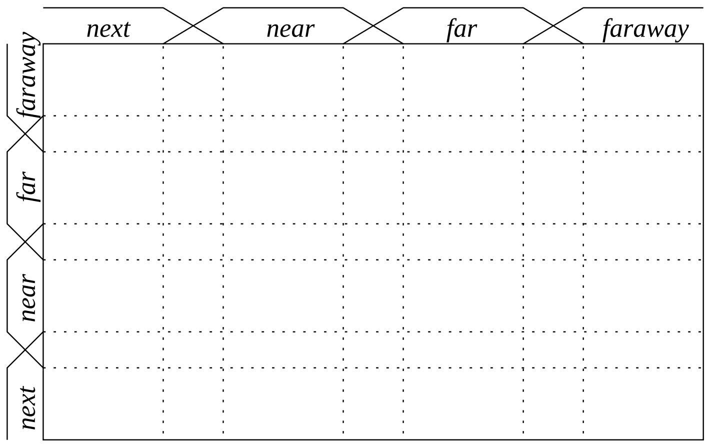

Figure 7. Fuzzy partition of a room.

> Next to the left wall and near the front wall there is a brown low-table. It is rectangular and it is parallel to the left wall.

The representation aspect means to translate the information contained in the linguistic terms into a format that can be used by the robot (metric information in the robot's reference system). To do this, we use fuzzy sets, that provide a convenient representation of the vagueness that is inherent to the linguistic descriptions (Zadeh, 1975). Since the describer can only use a predefined set of linguistic terms, we have built a fuzzy data base that provides a meaning (in the form of a fuzzy set) to each term. In the example of Figure 7, the linguistic terms that define the position of an object are substituted by their corresponding fuzzy sets. In this way, once a linguistic description is available, it can be automatically translated.

As for the use of linguistic descriptions, we can integrate them with sonar maps and image data to assign an identity to each perceived object. The linguistic terms for the object position impose fuzzy restrictions that, once translated to the sonar map, allow one to select a subset of the segments which should contain the projection of the object. Figure 8 illustrates this process for the description of a cabinet: starting from a complete sonar map and according to the linguistic description, only the segments encompassed by the dotted lines may correspond to the object. The linguistic term for the size also imposes a fuzzy restriction on the length and width of the object (height is not represented in the sonar map) that is used as a mask that indicates that any group of segments, in order to be interpreted as the described object, should be contained in an area whose shape and size are determined by the mask. Since there may be more than one possible matching of the mask to

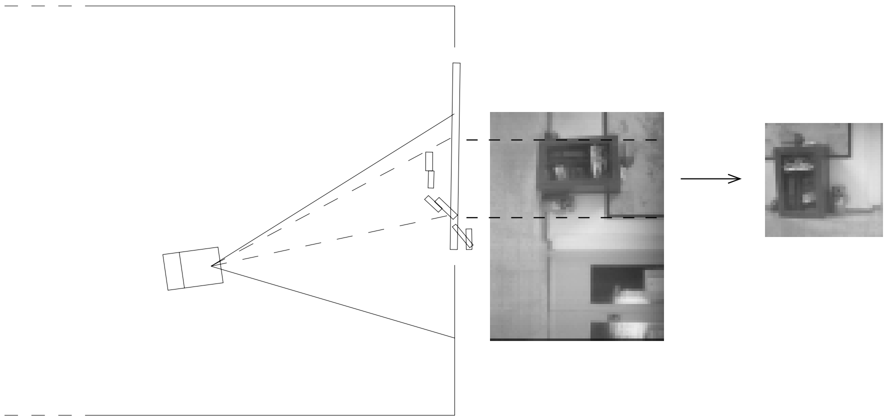

Figure 8. Translating environment information to image coordinates.

the segments, the output of the matching process is a set of hypotheses on the most plausible locations of the object in the map.

These hypotheses can now be checked on the image data. Given the position of the robot (hence of the camera) in the map, we can find the projection on the X axis of the image of any point in the map that is contained in the visual field (dashed lines in Figure 8). A similar analysis can be performed for the projections on the Y axis. However, we have to consider that the position of the robot is only approximately known, and the integration of linguistic descriptions and sonar maps keeps a high degree of uncertainty on the exact position and on the limits of the objects. In a previous work (Gasós and Ralescu, 1997) we have analyzed how the uncertainty on the real position of a point influences its projection on an image, and we have demonstrated that vague knowledge about location and size of the objects can still be translated into relevant information on region allocation of these objects in the image (small image in Figure 8). In this way, the difficult problem of image interpretation without prior environment knowledge is transformed into a more simple one of confirming the existence of a known object in a small part of the image.

The key aspect of this approach is the integration of different sources of information into a coherent representation of the working environment. Linguistic descriptions provide general information about the objects and their distribution, ultrasonic sensors provide metric information about the spatial layout of the environment, and vision is oriented towards object identification and is used to solve conflicts, check expectations and focus on areas of interest. In this way, we count on three independent sources of information, each one of different nature and characteristics, but complementary among them in such a way that from their integration we can obtain the necessary information to solve the recognition tasks.

5. Conclusions

Spatial knowledge about non-instrumented environments is inherently affected by uncertainty. In order to make effective use of this information, a (natural or artificial) autonomous agent must deal with this uncertainty in some way. In this paper, we have focused on artificial agents (mobile robots), and have surveyed the possibilities offered by fuzzy sets to represent and manipulate uncertain spatial information. We have outlined several models and techniques that have been proposed in the technical literature. We have also considered the important issue of the integration of metric models with other sources of spatial knowledge, such as image data and linguistic descriptions.

As it appears, much work remains to be done in order to endow mobile robots with the spatial reasoning capabilities that are necessary for fully autonomous operation. An interesting direction of research is the integration between different spatial reprentations at different levels of abstraction. Some preliminary steps in this direction are reported by Thrun et al. (1998) in a probabilistic framework, and by Gasós and Saffiotti (1999) and in Fabrizi and Saffiotti (2000) in a fuzzy set framework.

The results summarized in this paper suggest that approaches based on fuzzy sets are well suited to represent uncertain spatial information coming from different sources, and to integrate them into an approximate model of the environment. It is our hope that these findings can also contribute to the study and modeling of spatial cognition in animals and humans.

Acknowledgements

The work reported in this paper was performed at IRIDIA, Université Libre de Bruxelles. Financial support was provided by the European Community under the Marie Curie program (first author), and by the *Communauté Française de Belgique* (second author).

Notes

[1] In a sense, B plays a role similar to a probability distribution on X. Fuzzy sets however are essentially different from probability distributions because: (i) semantically, we interpret the value of $B(x)$ as a *degree of possibility*; (ii) formally, fuzzy sets do not impose the additivity requirement, i.e., we may have $\sum_{x \in X} B(x) \neq 1$; and (iii) pragmatically, we use fuzzy intersection as the main update mechanism.

[2] The English terms shown in the figure are translations of the Spanish terms that were actually used in the experiments. Some translations are slightly misleading: for instance, the Spanish term "cerca" (near) is not used when the distance between the objects is very small, in which case the term "al lado" (close) is preferred.

References

Bloch, I. (2000, to appear). Spatial Representation of Spatial Relationship Knowledge, *Int. Conf. on Knowledge Representation and Reasoning*.

Denis, M. (1996). Imagery and the Description of Spatial Configurations. In M. de Vega et al. (ed.), *Models of Visuospatial Cognition*. New York: Oxford University Press.

Elfes, A. (1987). Sonar-Based Real-World Mapping and Navigation, *IEEE Trans. on Robotics and Automation* 3(3): 249–265.

Fabrizi, E. and Saffiotti, A. (2000, to appear). Extracting Topology-Based Maps from Gridmaps. *IEEE Int. Conference on Robotics and Automation*. On-line at http://aass.oru.se/∼asaffio/.

Gasós, J. (2000). Integrating Linguistic Descriptions and Sensor Observations for the Navigation of Autonomous Robots. In D. Driankov and A. Saffiotti (eds.), *Fuzzy Logic Techniques for Autonomous Vehicle Navigation*. Berlin, DE: LNCS. Springer-Verlag. Forthcoming. On-line at ftp://iridia.ulb.ac.be/pub/jgasos/lncs99.ps.gz.

Gasós, J. and Ralescu, A. (1997). Using Imprecise Environment Information for Guiding Scene Interpretation, *Fuzzy Sets and Systems* 88(3): 265–288.

Gasós, J. and Rosetti, A. (1999). Uncertainty Representation for Mobile Robots: Perception, Modeling and Navigation in Unknown Environments, *Fuzzy Sets and Systems* 107(1): 1–24. On-line at ftp://iridia.ulb.ac.be/pub/jgasos/fss99.ps.gz.

Gasós, J. and Saffiotti, A. (1999). Integrating Fuzzy Geometric Maps and Topological Maps for Robot Navigation. *3rd Int. ICSC Symp. on Intell. Industrial Automation*. Genova, IT, June 1999. On-line at http://aass.oru.se/∼asaffio/.

Klir, G. and Folger, T. (1988). *Fuzzy Sets, Uncertainty, and Information*. Prentice-Hall.

Kortenkamp, D. and Weymouth, T. (1994). Topological Mapping for Mobile Robots Using Combination of Sonar and Vision Sensing. In *Proc. of the Twelfth National Conference on Artificial Intelligence (AAAI'94)* (pp. 979–984). Cambridge: Menlo Park.

Kuipers, B. (1978). Modeling Spatial Knowledge, *Cognitive Science* 2: 129–153.

Kuipers, B. and Levitt, T. (1988). Navigation and Mapping in Large Scale Space, *AI Magazine* 9: 25–43.

Kuipers, B. and Byun, Y.T. (1991). A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations, 8: 47–63.

Latombe, J.C. (1991). *Robot Motion Planning*. Boston, MA: Kluwer Academic Publishers.

Leonard, J.J., Durrant-Whyte, H.F. and Cox, I.J. (1992). Dynamic Map Building for an Autonomous Mobile Robot, *Int. J. of Robotics Research* 11(4): 286–298.

López-Sánchez, M., López de Màntaras, R. and Sierra, C. (2000, forthcoming). Map Generation by Cooperative Autonomous Robots Using Possibility Theory. In D. Driankov and A. Saffiotti (eds.), *Fuzzy Logic Techniques for Autonomous Vehicle Navigation*. Berlin, DE: LNCS. Springer-Verlag.

Moravec, H.P. and Elfes, A. (1985). High Resulution Maps from Wide Angle Sonar. In *Procs. of the IEEE Int. Conf. on Robotics and Automation* (pp. 116–121).

Moutarlier, P. and Chatila, R. (1989). Stochastic Multisensory Data Fusion for Mobile Robot Location and Environment Modeling. In *5th Int. Symp. on Robotics Research* (pp. 207–216). Tokyo, JP.

Oriolo, G., Ulivi, G. and Vendittelli, M. (1998). Real-time Map Building and Navigation for Autonomous Robots in Unknown Environments, *IEEE Trans. on Systems, Man, and Cybernetics* 28(3): 316–333.

Ruspini, E.H. (1991). On the Semantics of Fuzzy Logic, *Int. J. of Approximate Reasoning* 5: 45–88.

Saffiotti, A. (1997). The Uses of Fuzzy Logic for Autonomous Robot Navigation, *Soft Computing* 1(4): 180–197. On-line at http://aass.oru.se/Living/FLAR/.

Saffiotti, A. and Wesley, L.P. (1996). Perception-Based Self-Localization Using Fuzzy Locations. In M. van Lambalgen L. Dorst and F. Voorbraak (eds.), *Reasoning with Uncertainty in Robotics*, number 1093 in LNAI (pp. 368–385). Berlin, DE: Springer-Verlag. On-line at http://aass.oru.se/∼asaffio/.

Saffiotti, A., Konolige, K. and Ruspini, E.H. (1995). A Multivalued-Logic Approach to Integrating Planning and Control, *Artificial Intelligence* 76(1–2): 481–526. On-line at http://aass.oru.se/∼asaffio/.

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

Smith, R.C. and Cheeseman, P. (1986). On the Representation and Estimation of Spatial Uncertainty, *Int. J. of Robotics Research* 5(4): 56–68.

Surmann, H. and Peters, L. (2000, forthcoming). The Uses of Fuzzy Control for the Autonomous Robot Moria. In D. Driankov and A. Saffiotti (eds.), *Fuzzy Logic Techniques for Autonomous Vehicle Navigation*. Berlin, DE: LNCS. Springer-Verlag.

Thrun, S., Gutmann, J-S., Fox, D., Burgard, W. and Kuipers, B. (1998). Integrating Topological and Metric Maps for Mobile Robot Navigation: A Statistical Approach. *Proc. of the 15th AAAI Conf.* (pp. 989–996). Wisconsin: Madison.

Thrun, S. (1999). Learning Metric-Topological Maps for Indoor Mobile Robot Navigation, *Artificial Intelligence* 1: 21–71.

Tirumalai, A., Schunck, B. and Jain, R. (1995). Evidential Reasoning for Building Environment Maps, *IEEE Trans. on Systems, Man, and Cybernetics* 25(1): 10–20.

Tunstel, E. (1995). Fuzzy Spatial Map Representation for Mobile Robot Navigation. In *Proc. of the ACM 10th Annual Symp. on Applied Comp.* (pp. 586–589). Nashville, TN.

Zadeh, L.A. (1965). Fuzzy Sets, *Information and Control* 8: 338–353.

Zadeh, L.A. (1975). The Concept of a Linguistic Variable and Its Application to Approximate Reasoning, *Information Sciences* 8: 199–249.

Zadeh, L.A. (1978). Fuzzy Sets as a Basis for a Theory of Possibility, *Fuzzy Sets and Systems* 1: 3–28.