Generating Multi-Level Linguistic Spatial Descriptions from Range Sensor Readings Using the Histogram of Forces

Marjorie Skubic, Pascal Matsakis, George Chronis, and James Keller

Dept. of Computer Engineering and Computer Science

University of Missouri-Columbia

{skubicm, matsakisp, chronisg, kellerj}@missouri.edu

Abstract

In this paper, we show how linguistic expressions can be generated to describe the spatial relations between a mobile robot and its environment, using readings from a ring of sonar sensors. Our work is motivated by the study of human-robot communication for novice robot users. The ultimate goal is to exploit these linguistic expressions for navigation of the mobile robot in an unknown environment, where the expressions represent the qualitative state of the robot in terms that are easily understood by humans. The notion of the histogram of forces was presented in previous work and used to generate linguistic descriptions of relative positions in digital images. Here, we demonstrate that it also permits fast processing of vector data and can be applied to a robot with range sensors moving in a dynamic environment. We introduce a new method for detecting partially and completely surrounded conditions, and we show that detailed descriptions can be obtained as well as coarse ones. Numerous examples are included, illustrating a variety of situations.

Keywords

Human-robot communication, robot navigation, histograms of forces, linguistic spatial descriptions, spatial relationships, surroundedness.

1. Introduction

Our work is motivated by the study of human-robot interaction and, in particular, the investigation of human-robot communication. The ultimate goal is to provide easy and intuitive interaction by novice users, so that they can guide, control, and/or program a robot to perform some purposeful task. We consider the communication between the human user and the robot to be crucial to intuitive interaction by non-robotics experts. We further argue that good communication is essential not only from the human to the robot but also from the robot to the human (so that the user can monitor the robot's current state or condition). See (Morik *et al.*, 1999) and (Skubic and Volz, 2000) for further motivation and examples on task-oriented dialogues.

In this paper, we are not attempting to build an exact model of the robot's environment, nor to generate a quantitative map. However, we do want to generate linguistic expressions that represent the state of the robot with respect to its environment in terms that are easily understood by human users. These expressions provide a symbolic link between the robot and the user, thus comprising a navigation language for human-robot interaction, and can be utilized for two-way communications. First, in robot-to-human communication, they provide a qualitative description of the robot's current state (*e.g.*, "there is an object on the left," or "there is an object on the front-right"). Second, in human-to-robot communication, the human can command the robot to perform navigation behaviors based on spatial relations (*e.g.*, "while there is an object on the left, move forward," or "if there is an object on the front-right, turn left," or even a high-level and very human-like directive such as "turn left at the second intersection"). The following examples are representative scenarios in which robot-generated linguistic spatial descriptions may prove useful:

1. Semi-autonomous control of a robot, where control may alternate between varying levels of autonomous robot control and direct human control (*i.e.*, adjustable autonomy).
2. Supervisory control, where the user issues high-level commands, allowing control with a low bandwidth communication channel.
3. Programming the robot to perform some navigation-related task in an unstructured environment with some landmarks.

Cognitive models suggest that people use relative, egocentric spatial reasoning to perform day-to-day navigation tasks (Previc, 1998) (Schunn and Harrison, 2001). Thus, it seems natural to incorporate such spatial language into an interface for a mobile robot. We assume the robot has low-level navigation behaviors (either pre-programmed or pre-learned) that allow it to move safely around its unstructured and dynamic environment. However, task-level navigation is represented and described as a sequence of navigation behaviors that are segmented by qualitative "states" based on spatial relations. That is, each navigation behavior has a corresponding termination state that transitions the robot to the next behavior; the sequence describes the entire navigation task. This model is consistent with hierarchical methods proposed for topological navigation, *e.g.*, (Kuipers, 1998) (Krieg-Brückner *et al.*, 1998). To accomplish this model of navigation and to provide communication with a human user, the robot must be able to recognize its state in terms of egocentric spatial relations between itself and objects in its environment, and it must be able to describe these relations in natural language.

The idea of relying on linguistic spatial expressions to communicate with a semi-autonomous robot has been proposed previously. Gribble *et al.* (1998) use the framework of the Spatial Semantic Hierarchy for an intelligent wheelchair. Perzanowski *et al.* (1999) utilize a combination of gestures and linguistic directives such as "go over there." In (Shibata *et al.*, 1996), positional relations are used to overcome ambiguities in the recognition of landmarks. In (Stopp *et al.*, 1994), the user communicates with a 2-arm mobile robot performing assembly tasks. He selects an object from the robot's environment model by expressing relational spatial descriptions. Finally, Moratz and Fischer (2000) have conducted experiments to investigate the spatial references preferred by users when directing a small mobile robot between obstacles. Results indicated that the robot's point of view was consistently adopted. Furthermore, to obtain the desired trajectory, about half of the study's participants described the trajectory as a sequence of primitive actions (*e.g.*, "move forward, and turn left").

The main focus of this paper is the creation of linguistic expressions that describe the spatial relations between a mobile robot and its environment using readings from a ring of range sensors such as sonar sensors. Our work relies on spatial analysis tools previously applied to image data (Matsakis *et al.*, 2001). It develops and extends a preliminary study presented in (Skubic *et al.*, 2001). Background material as well as new spatial analysis tools are examined in Section 2. In Section 3, we show how these tools can be interfaced with the robot's sonar readings. Specific test cases are shown in Section 4. Concluding remarks along with a discussion of future work are found in Section 5.

2. Using the Histogram of Forces

Freeman (1975) proposed that the relative position of two objects be described in terms of spatial relationships (such as "above," "surrounds," "includes," *etc*.). He also proposed that fuzzy relations be used, because "all-or-nothing" standard mathematical relations are clearly not suited to models of spatial relationships. Moreover, "although the human way of reasoning can deal with qualitative information, computational approaches of spatial reasoning and object recognition can benefit from more quantitative measures" (Bloch, 1999). By introducing the notion of the histogram of angles, Krishnapuram *et al.* (1993) and Miyajima and Ralescu (1994) developed the idea that the relative position between two objects can have a representation of its own and can thus be described in terms other than spatial relationships. However, the representation proposed shows several weaknesses (*e.g.*, requirement for raster data, long processing times, anisotropy).

Matsakis (1998) introduced the histogram of forces. Contrary to the angle histogram, it permits processing of vector data as well as raster data, offers solid theoretical guarantees, allows explicit and variable accounting of metric information, and lends itself, with great flexibility, to the definition of qualitative directional spatial relations such as "to the right of," "in front of," *etc*. (Matsakis and Wendling, 1999) (Matsakis *et al.*, 1999, 2001). For our purposes, the histogram of forces also allows for a low-computational handling of heading changes in the robot's orientation and makes it easy to switch between a world view and an egocentric robot view.

2.1. Handling of Vector Data

In previous work, force histograms were generated using raster image data. The present paper describes the first application of histograms that uses vector data, *i.e.*, a boundary representation of the objects. Here, we show how vector data are handled. In the framework of our experiments with the robot, the objects that we consider are naturally quite simple. In the following, A and B denote two disjoint and polygonal 2D objects, as in Fig. 1a. Note, however, that no theoretical issue prevents them from intersecting, having holes, being non-connected or with uncertain boundaries (*i.e.*, fuzzy). French-speaking readers can find more details in (Matsakis, 1998).

Consider any direction θ. The objects A and B are partitioned by drawing straight lines through their vertices, at the angle of θ (Fig. 1b). Two consecutive lines determine pairs (I,J) of trapezoids (or triangles, which are particular trapezoids) with I included in A and J in B. Seven non-negative values may be used to characterize such a pair: X_1, Y_1, Z_1, X_2, Y_2, Z_2 and ε (Fig. 1c). Now, assume I and J are flat plates of uniform density and constant and negligible thickness (this kind of objects is commonly considered in physics, see, *e.g.*, (Cutnell and Johnson, 2001)). If I is located "after" J in direction θ (which is the case in Fig. 1c), then there exist elementary gravitational forces, exerted by the points of I on those of J, that each tend to move J in direction θ. The scalar resultant Γ_g^{IJ} of these forces is

$$\Gamma_g^{IJ} = \varepsilon[\lambda(X_1+Y_1,X_2+Y_2)-\lambda(Y_1,Y_2)+\lambda(Y_1+Z_1,Y_2+Z_2)-\lambda(X_1+Y_1+Z_1,X_2+Y_2+Z_2)], \tag{1}$$

where λ denotes the function defined by:

for any positive real numbers s and t, if $s \neq t$ then $\lambda(s,t) = [t\ln(t) - s\ln(s)]/(t-s)$,

else $\lambda(s,s) = \lim_{t \to s} \lambda(s,t) = 1 + \ln(s)$.

The sum of all Γ_g^{IJ}, when processing all pairs (I,J), is denoted by $F_g^{AB}(\theta)$. It represents the total weight of the arguments that can be found in order to support the proposition "A is in direction θ of B." The function $F_g^{AB} \mid \theta \rightarrow F_g^{AB}(\theta)$ is called the *histogram of gravitational forces* associated with (A,B). It models the position of A relative to B. The object A is the *argument*, and the object B the *referent*.

In lieu of gravitational forces, constant forces can also be considered (*i.e.*, the elementary force exerted by one material point on another does not depend on the distance between the two points; like gravitational forces, constant forces are commonly encountered in physics). Γ_g^{IJ} is then replaced by Γ_c^{IJ}:

$$\Gamma_c^{IJ} = \varepsilon[(X_1+X_2)(Z_1+Z_2)+X_1Z_1+X_2Z_2]/[6\cos^2(\theta)]. \tag{2}$$

The sum of all Γ_c^{IJ}, when processing all pairs (I,J), is denoted by $F_c^{AB}(\theta)$. The function $F_c^{AB} \mid \theta \rightarrow F_c^{AB}(\theta)$ is the *histogram of constant forces* associated with (A,B). It is another representation of the position of A relative to B. As shown in (Matsakis and Wendling, 1999), the histograms F_c^{AB} and F_g^{AB} satisfy various geometric properties and have interesting characteristics. F_c^{AB} provides a global "view" of the situation. It considers the closest parts and the farthest parts of the objects equally, whereas F_g^{AB} focuses on the closest parts.

The practical computation of a histogram of forces can be summarized by the algorithm below. The computation time is obviously proportional to the number of directions that are considered. For our experiments, in Section 4, we chose 180 directions—an appropriate choice according to (Matsakis and Wendling, 1999). The assessment of each value, $F_c^{AB}(\theta)$ or $F_g^{AB}(\theta)$, is of complexity $O(n\log(n))$, where n denotes the total number of object vertices. Histogram computation is extremely fast. The mobile robot considered in Sections 3 and 4 describes its environment in real time.

Algorithm for computing the histogram of forces:

```
Choose a number d of directions;
FOR each θ in the set {0, 2π/d, 4π/d, … , 2(d-1)π/d} of evenly distributed directions DO
    F^AB(θ) ← 0 ;
    Partition the objects into trapezoids by drawing lines through A and B's vertices at the angle of θ (Fig. 1b);
    FOR each trapezoid (or triangular) region I of object A DO
        FOR each trapezoid (or triangular) region J of object B DO
            IF I and J are delimited by the same lines AND I is "after" J in direction θ THEN
                Compute the scalar resultant Γ^IJ (Fig. 1c and formula (1) or (2));
                F^AB(θ) ← F^AB(θ) + Γ^IJ;
```

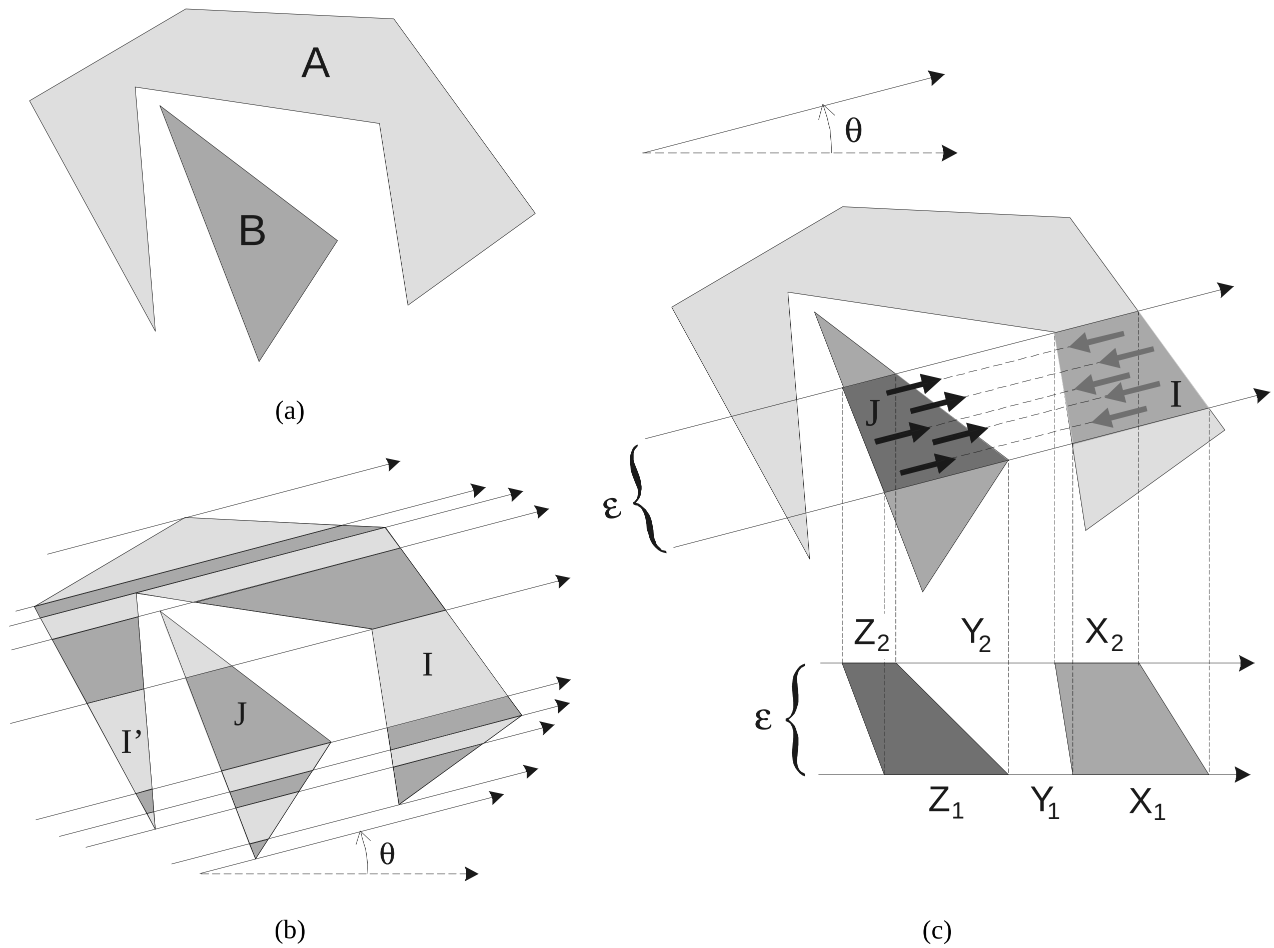

Fig. 1. The computation of a histogram value like $F_g^{AB}(\theta)$ is based on the partitioning of the objects into trapezoids. (a) Two disjoint polygonal objects *A* and *B*. (b) *(I,J)* is one of the five pairs of trapezoids that support the proposition "*A* is in direction θ of *B*": *I* and *J* are delimited by the same parallel lines; *I* is included in *A* and *J* is included in *B*; *I* is located "after" *J* in direction θ. Note that *(I',J)* is not one of these five pairs: since *I'* is located "before" *J* in direction θ, the pair *(I',J)* does not support at all the proposition "*A* is in direction θ of *B*" (it supports, however, the proposition "*A* is in direction $\theta+\pi$ of *B*" and will be considered when computing $F_g^{AB}(\theta+\pi)$). (c) Γ_g^{IJ} is the scalar resultant of elementary forces (black arrows). Each one tends to move *J* in direction θ. The computation of Γ_g^{IJ} requires seven values to be determined.

2.2. Describing Directional Relationships

In previous work on image analysis, Matsakis *et al.* (2001) have presented a system that produces linguistic spatial descriptions. The description of the relative position between any 2D image objects A and B is generated from F_c^{AB} (the histogram of constant forces associated with (A,B)), and F_g^{AB} (the histogram of gravitational forces). The linguistic description produced by the system relies on the four primitive directional relationships: "to the right of," "above," "to the left of" and "below" (imagine that the objects are drawn on a vertical surface). First, eight values are extracted from the analysis of each histogram: α*(RIGHT)*, β*(RIGHT)*, α*(ABOVE)*, β*(ABOVE)*, α*(LEFT)*, β*(LEFT)*, α*(BELOW)* and β*(BELOW)*. They represent the "opinion" given by the considered histogram. For instance, according to F_c^{AB} the degree of truth of the proposition "A is to the right of B" is some value α_c*(RIGHT)*. This value is a real number greater than or equal to 0 (proposition completely false) and less than or equal to 1 (proposition completely true). Moreover, according to F_c^{AB} the maximum degree of truth that can reasonably be attached to the proposition (say, by another source of information, like F_g^{AB}) is $\beta_c(RIGHT)$ (which belongs to the interval $[\alpha_c(RIGHT),1]$). F_c^{AB} and F_g^{AB}'s opinions (*i.e.*, the sixteen values) are then combined. Four numeric and two symbolic features result from this combination. They feed a system of 27 fuzzy rules and meta-rules that outputs the expected linguistic description. The system handles a set of 16 adverbs (like "mostly," "perfectly," *etc*.) that are stored in a dictionary, with other terms, and can be tailored to individual users. A description is generally composed of three parts. The first part involves the primary direction (*e.g.*, "A is mostly to the right of B"). It is often supplemented by a second part that involves a secondary direction (*e.g.*, "but somewhat above"). The third part indicates to what extent the four primitive directional relationships are suited to describing the relative position of the objects (*e.g.*, "the description is satisfactory"). In other words, this assessment indicates to what extent it is necessary to turn to other spatial relations (*e.g.*, "surrounds"). All details can be found in (Matsakis *et al.*, 2001).

2.3. Describing "Surrounds"

2.3.1. Existing Methods

The spatial relationships "surrounds" and "is surrounded by" have an important role in the interpretation of an image scene or an environment. Many quantitative definitions have been proposed. There are two main approaches. The first approach uses the fact that according to most families of directional relations, an object can be in many directions with respect to another. This feature is questionable: usually, people do not combine more than two spatial prepositions when translating visual information into natural language descriptions (Gapp, 1995) (Retz-Schmidt, 1988). However, some authors (Miyajima and Ralescu, 1994) (Bloch, 1999) support the idea that it allows "surrounds" (and "is surrounded by") to be derived. Knowing that A is somewhat above, below, to the right and to the left of B as well, one could conclude that A surrounds B. In fact, drawing such a conclusion is not reasonable, unless it is known that the argument A does not intersect the convex hull of B (Fig. 2). In other words, "A surrounds B" can be assessed only if it is known that B *does not surround* A *at all*. The reason is that the directional relations are tied by the *semantic inverse* principle (Freeman, 1975) (*e.g.*, A is to the left of B as B is to the right of A). Therefore, without constraints on the objects, there is no way to know which one surrounds (or includes!) the other.

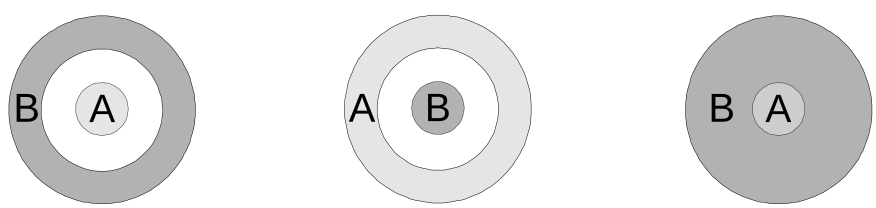

(a) (b) (c)

Fig. 2. The directional relations cannot substitute for the spatial relation "surrounds." In each case, according to most families of directional relations, object A is somewhat above, below, to the right and to the left of object B as well. It does not mean that A surrounds B. (a) A is surrounded by B. (b) A surrounds B. (c) A is included in B.

The second approach derives from Rosenfeld's visual surroundedness (Rosenfeld and Klette, 1985). It is based on the computation of a histogram of angles. It supposes that the argument A is connected and does not intersect B. For any pixel P of B, let θ_P be the angle made by the two tangents from P to A as in Fig. 3. To each element θ of $[0,2\pi]$, the histogram associates the number of pixels P such that θ_P is equal to θ. In (Wang and Keller, 1997), the degree of truth for "A surrounds B" is produced by a multilayer perceptron fed by the histogram values and trained on aggregate responses from a panel of people. Other authors resort to a decreasing membership function μ from $[0,2\pi]$ into $[0,1]$, where μ is chosen such that $\mu(\theta)$ is 1 if θ is 0, and is 0 if θ is greater than π. In (Miyajima and Ralescu, 1994), the histogram of angles is assimilated to a fuzzy set and matched to μ, using the compatibility notion (Dubois and Prade, 1980). The degree of truth for "A surrounds B" is obtained as the center of gravity of the compatibility fuzzy set. In (Krishnapuram *et al.*, 1993), the histogram is used to compute the aggregated value (*e.g.*, the arithmetic mean, or the generalized mean (Klir and Folger, 1988)) of all the $\mu(\theta_P)$. The degree of truth for "A surrounds B" is set to this value. The second approach to defining surroundedness is interesting. Compared with the first one, it gives results which are much more consistent with human perception (Wang and Keller, 1997). However, the computing cost is very high, and vector data cannot be handled.

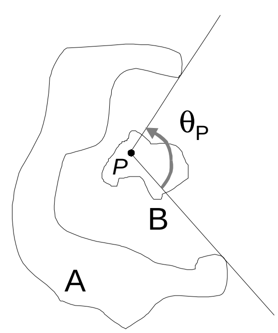

Fig. 3. Defining surroundedness. Second approach (derived from Rosenfeld and Klette, 1985).

2.3.2. New Method

Let μ be a membership function as above, and let (θ_1,θ_2) be the largest open interval on which the histogram F_c^{AB} (or F_g^{AB}) is zero. Fig. 5 illustrates the meaning of angles θ_1 and θ_2. The range for θ_1 is $(-\pi,\pi]$ and the range for θ_2 is $[\theta_1,\theta_1+2\pi]$. The degree of truth for "A surrounds B" is set to $\mu(\theta_2-\theta_1)$. The definition is quite simple, but compares with the others (see Fig. 4 and Table 1). The main advantages of the force histogram-based method are that it ensures much faster processing of raster data, and it is able to handle vector data as well. Moreover, it can be used concurrently to assess the directional relations. In particular, it can easily be incorporated into the system described in (Matsakis *et al.* 2001). Let a and b be two real numbers such that $0 \le b \le a \le 1$, and, for any θ, let $\delta(\theta)$ be the multiple of $\pi/4$ which is the closest to θ (the value $\delta(\theta)$ corresponds to one of the eight directions "right," "above-right," "above," *etc*.). The linguistic description that relies on the four primitive directional relationships is considered first (see Section 2.2). If its self-assessment is not "satisfactory," and if $\theta_2-\theta_1$ is lower than $a\pi$, then the extended system turns to the spatial relation "surrounds." If $\theta_2 - \theta_1$ belongs to $[b\pi, a\pi]$, the output is "A surrounds B on the right side" (or "on the top-right side," "on the top side," *etc*., depending on the value of $\delta(\pi + (\theta_1 + \theta_2)/2)$). If $\theta_2 - \theta_1$ belongs to the open interval $(0, b\pi)$, "A surrounds B, but leaves an opening on the right side" (or "on the top-right side," "on the top side," *etc*., depending on the value of $\delta((\theta_1 + \theta_2)/2)$). If $\theta_2 - \theta_1$ is 0, "A completely surrounds B."

For all our experiments (see Table 1 and Sections 3 and 4), we used the most natural and simplest μ, *i.e.*, a linear membership function: for any θ in $[0,\pi]$, $\mu(\theta)=1-\theta/\pi$. The thresholds a and b were determined empirically, according to our intuition. The linguistic descriptions in Fig. 5 and Sections 3 and 4 have been obtained by setting a to 2/3 and b to 1/3. Note that some non-linear functions, like $\mu(\theta)=\cos(\theta/2)$, produce degrees of truth for "A surrounds B" that are much closer to human average rating than those exhibited in Table 1. However, it can easily be shown that the same linguistic descriptions can be generated with any μ, provided that the thresholds a and b are appropriately chosen.

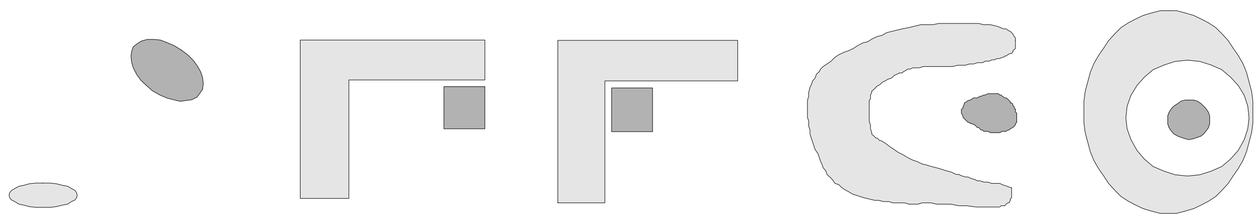

Fig. 4. Some images used for training and testing the neural network-based model of "surrounds" (Wang and Keller, 1997).

TABLE 1

RESULTS OF DIFFERENT METHODS USED TO ASSESS THE "SURROUNDS" RELATION FOR FIG. 4

In (Wang and Keller, 1997), 63 people were shown images with two objects and asked to rate the relation "surrounds." HUMAN represents the average rating. NN is the neural network-based method (Wang and Keller, 1997), AGG the aggregation method (Krisnapuram *et al.*, 1993), and COM the compatibility method (Miyajima and Ralescu, 1994). FH denotes the method that uses force histograms. The associated membership function is linear: for any θ in $[0,\pi]$, $\mu(\theta)=1-\theta/\pi$. The value 1.00 means that the proposition "A surrounds B" is assessed to be completely true, and the value 0.00 that it is completely false.

	Fig. 4a	Fig. 4b	Fig. 4c	Fig. 4d	Fig. 4e
HUMAN	0.01	0.36	0.49	0.60	0.99
NN	0.03	0.41	0.49	0.62	0.99
FH	***0.00***	***0.23***	***0.47***	***0.43***	***1.00***
AGG	0.00	0.03	0.30	0.22	0.90
COM	0.00	0.02	0.20	0.14	1.00

Finally, the histogram of forces can easily be employed to assess surroundedness. However, there is the following assumption: the argument is connected (like in the second approach), and does not intersect the convex hull of the referent (like in the first approach). Also note that the degree of truth for "A surrounds B" does not really use the histogram values but rather considers only the binary conditions of zero or non-zero. This has two important consequences. First, the results do not depend on the choice of the force histogram. Second, the method is not extremely robust. Slight changes in the object shapes (especially at the two “ends” of the argument) may have a noticeable impact on the degrees of truth. In fact, similar comments can be addressed to any method that derives from Rosenfeld’s visual surroundedness; the results are not sensitive to the thickness of the argument, only to tangency points. The above-mentioned limitations (constraints on the objects, low robustness) are not an issue in the present framework of human-robot communication because of the type of objects that are handled (see Section 3.1) and the fact that the degrees of truth are only used for generating linguistic descriptions (whose granularity is much coarser). This may not be the case for other applications (or other sensor configurations). Two promising avenues still remain to be explored: *(i)* redefining the degree of truth for “*A* surrounds *B*” using the histogram values (and not only the fact that these values are either zero or non-zero), and *(ii)* introducing a new type of histogram of forces, dedicated to surroundedness (the idea is to adopt a new set of axiomatic properties, and to change the way the longitudinal sections are handled (Matsakis, 1998) (Matsakis and Wendling, 1999)).

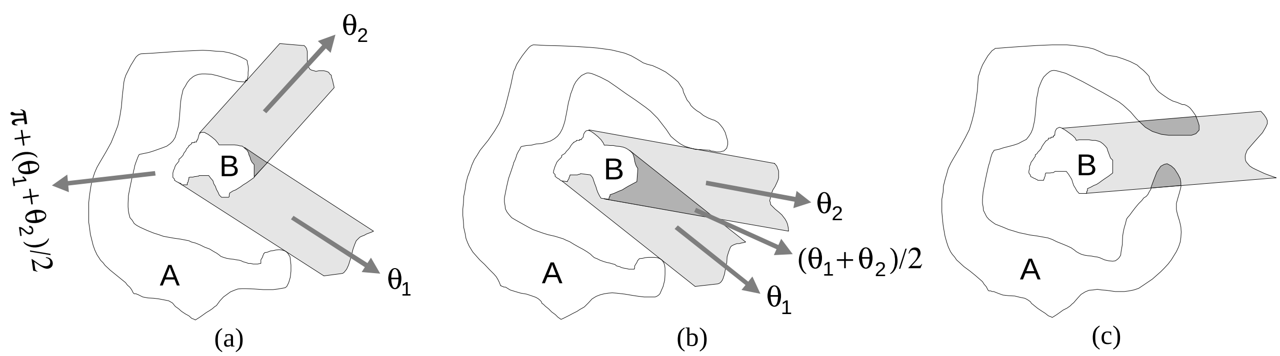

Fig. 5. Linguistic descriptions and surroundedness. (a) “*A* surrounds *B* on the left side.” (b) “*A* surrounds *B*, but leaves an opening on the bottom-right side.” (c) “*A* completely surrounds *B*.” The parameter a was set to 2/3 and the parameter b to 1/3.

3. Egocentric Spatial Relations from Range Sensor Readings

In this section, we describe the application of the histogram of forces for extracting egocentric spatial relations from the range sensor readings of a mobile robot. Here, the referent object (*B*) is always the robot, and a few simplifying assumptions are made, namely that the robot can be modeled as a convex object and is not touching its environment (*i.e.*, it is not allowed to jam itself into obstacles). Our test robot is a Nomad 200 with 16 sonar sensors evenly distributed along its circumference. The sensor readings are used to build an approximate polygonal representation of the surrounding obstacles. The histograms of constant and gravitational forces, F_c and F_g, are then computed as described in Section 2.1, and linguistic descriptions of relative positions between the robot and the environment objects are generated as in Section 2.2. The process is outlined in Fig. 6, and we examine each step in the sections below. Note that, although we illustrate the process using a ring of 16 sonar sensors, there is no theoretical limitation on the sensor type or the configuration. In fact, our implementation is parameterized to support any number of range sensors with varying cone sizes. The sensor type and configuration affect the sensing resolution.

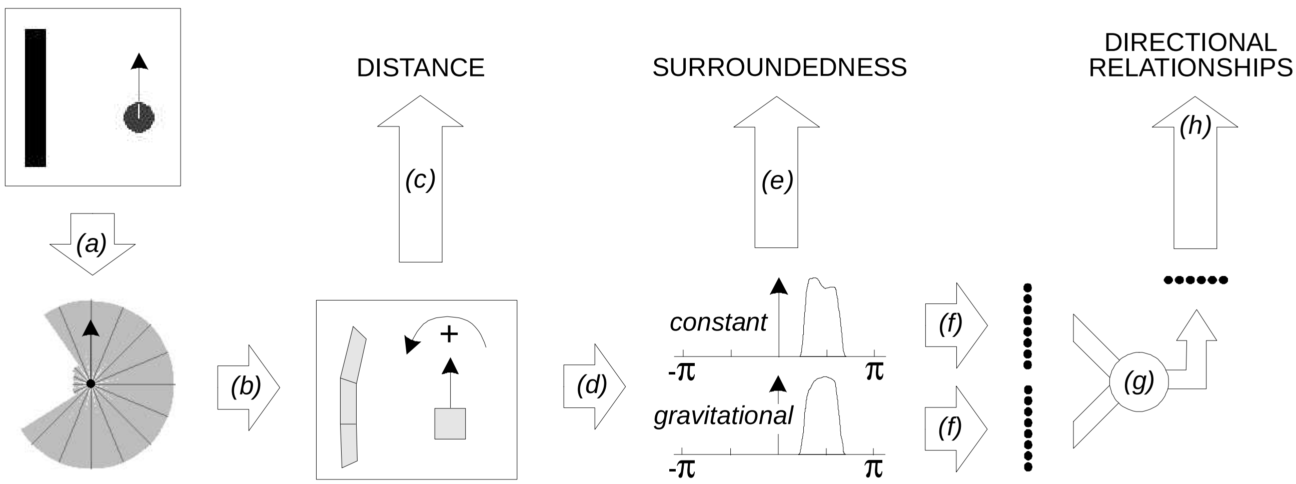

Fig. 6. Synoptic diagram. (a) Sonar readings. (b) Modeling the sensed environment. (c) Describing proximity relationships. (d) Computation of the histograms of forces. (e) Turning to the spatial relation "surrounds." (f) Extraction of numeric features. (g) Fusion of information. (h) Describing directional relationships.

3.1. Modeling the Sensed Environment

The first step in recognizing spatial relations from sensor signals is to model the obstacles around the robot by polygonal objects. Consider the simple case illustrated by Fig. 7. The solid rectangle represents an obstacle. For computational simplicity, our test robot, although circular in shape, is assimilated to a rectangular object with a heading. The sonar sensor *S* is the only one that returns a range value (*i.e.*, a range value less than the maximum), indicating that an obstacle has been detected. Thus, a single trapezoid object is built in the center of cone *S*. Note that the thickness of the obstacles cannot be determined from the sonar readings, so we assign a constant arbitrary thickness when building the polygonal models.

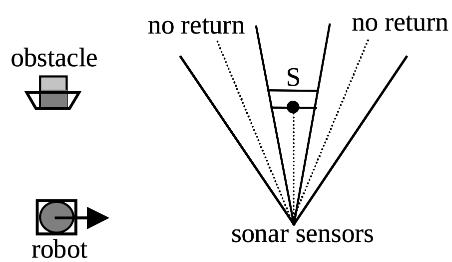

Fig. 7. A single trapezoid object is formed from a single sonar reading.

In the case of adjacent sonar returns, an assessment must be made on whether they are from a single obstacle or multiple obstacles. Our solution to this issue is to compute whether the robot can fit between two adjacent returns. (Here, each sonar reading is assimilated to a point on the axis of the corresponding cone, and its distance to the cone vertex is determined by the returned range value. See Fig. 8.) If the robot cannot fit, then we consider the two readings to be from the same obstacle. Even if there are actually two obstacles, they may be considered as one for navigation purposes. If the robot can fit, we model the environment as separate objects. Whether the robot can fit or not is determined by the following distance measure:

$$\sqrt{s_1^2 + s_2^2 - 2 s_1 s_2 \cos(2\pi / c)} \;,$$

where s_1 is the reading from sonar sensor S_1,

s_2 is the reading from S_2 (adjacent to S_1),

and $\frac{2\pi}{c}$ is the angle between the axes of two adjacent sonar cones.

The above formula simply corresponds to the distance between two points in polar coordinates. Although our test robot has 16 sensors, in our experiments (Section 4) we use c = 24 and not c = 16, *i.e.*, we intentionally underestimate the distance between adjacent sonar returns. This effectively produces extra clearance to make sure that the robot can easily fit between the obstacles that are perceived.

For example, consider the case illustrated by Fig. 8. The distance between the returns is rather large, and we cannot determine whether a single obstacle continues across the three sonar cones, or we have three different obstacles. Therefore, three different objects are built. In the same figure we show the distance computed for c = 16, which is the distance between *A* and *B*, and for c = 24, which is the distance between *C* and *D*. In Fig. 9, the robot is closer to the obstacle. We now have a better resolution, and five adjacent sonar sensors indicate returns. All returns are close together, and the robot cannot pass through. We join the five points and model the environment as one single object.

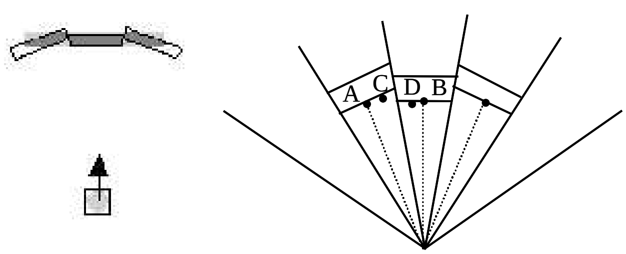

Fig. 8. Three different objects are formed from 3 different sonar readings, if the readings are not "close" enough, according to the distance measure.

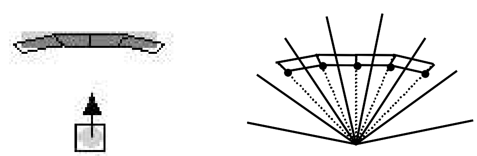

Fig. 9. A single object is formed from 5 different sonar readings, if the readings are "close" enough, as determined by the distance measure.

3.2. Directional Descriptions

After the surrounding obstacles are modeled by polygonal objects, the position of each object relative to the robot is represented by histograms of forces and described by a three-part linguistic expression, as explained in Section 2. Note that the linguistic term set is stored in a dictionary file, thus providing an easy way of modifying the descriptor clauses or even translating them into a different language. In this work, we have chosen a terminology that is appropriate for robot navigation. For instance, instead of using the words "above" and "below" as in Section 2, we use "in front of" and "behind." Moreover, the linguistic descriptions are from the robot's point of view and depend on the robot's heading. This is easily accomplished by shifting the force histograms along their horizontal axis.

Consider the simple case illustrated by Fig. 10. Fig. 10a shows the bounding rectangle that represents the robot, and the trapezoid model built from the sonar readings. The readings are displayed in Fig.10b and the generated description in Fig. 10c. The linguistic expression appears in its three-part form: (1) “An object is mostly in front of the robot” (primary directional relationship), (2) “but somewhat to the left” (secondary directional relationship), and (3) “the description is satisfactory” (self-assessment).

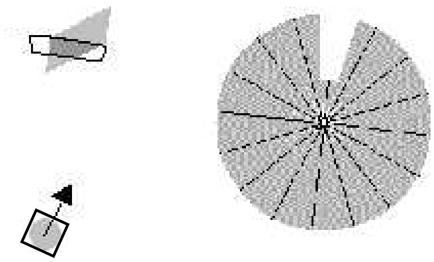

An object is mostly in front of the robot,
but somewhat to the left
(the description is satisfactory).

(a) (b) (c)

Fig. 10. Directional descriptions. (a) The robot detects one object. (b) The corresponding sonar sensor readings. (c) The generated linguistic directional description.

3.3. Distance Descriptions

In addition to directional information, the linguistic descriptions include distance information. For each object detected by the robot, the distance description is based on the range values returned by the sonar sensors. The number of distance relationships and the thresholds used for defining these relationships are also dictionary-driven for flexibility. Let s be the lowest sonar reading for the considered object, expressed in percentage of the maximum range value. In our experiments (Section 4), we use the four following relationships: “very close” if s belongs to [0,25], “close” if s belongs to (25,50], “far” if s belongs to (50,75], and “very far” if s belongs to (75,100].

Fig. 11 shows a more complex example in which the distance descriptions have been added. The obstacle from Fig. 10 remains at the same position (Object 1). A new obstacle has been introduced behind the robot. It is recognized as a single object (Object 2). The obstacle to the right of the robot however, is modeled as three different objects (Objects 3 to 5), as determined by the distance measure (see Section 3.1). Fig. 11c shows the five individual linguistic descriptions. In each case, the self-assessment indicates an adequate description. Using the thresholds above, Object 2 is classified as “very close” and the remaining objects are designated as merely “close.”

3.4. Surrounds Descriptions

In the examples shown thus far, the self-assessment has always been: “the description is satisfactory.” In the case where an assessment is not satisfactory, we consider whether the surrounds description is appropriate. The technique used to determine the possibility of surroundedness is described in Section 2.3.2. Since the robot is modeled as a convex object and is not allowed to jam itself into obstacles, the requirements mentioned in that section are met. We consider all types of surrounds descriptions, as explained in Fig. 5. The three cases are illustrated with the robot in Figs. 12, 13, and 14, where part (a) shows the robot in the environment, part (b) shows the sonar sensor readings, and part (c) shows the generated description. Note that the environment must be close enough so that a single polygonal object is built from the sonar readings. Also, there is no self-assessment for surroundedness.

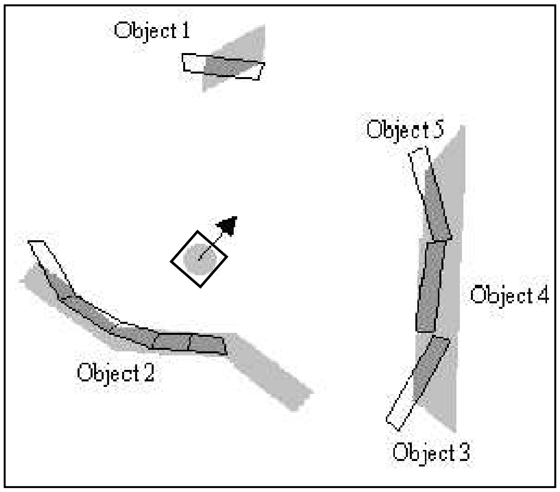

(a)

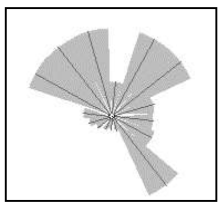

(b)

An object is mostly in front of the robot,
but somewhat to the left
(the description is satisfactory).
The object is close to the robot.

An object is behind-left of the robot
(the description is satisfactory).
The object is very close to the robot.

An object is mostly on the right of the robot,
but extends forward relative to the robot
(the description is satisfactory).
The object is close to the robot.

An object is on the front-right of the robot
(the description is satisfactory).
The object is close to the robot.

An object is mostly in front of the robot,
but somewhat to the right.
(the description is satisfactory)
The object is close to the robot.

(c)

Fig. 11. Distance descriptions. (a) The robot detects five objects. (b) The corresponding sonar sensor readings. (c) The generated detailed linguistic spatial descriptions.

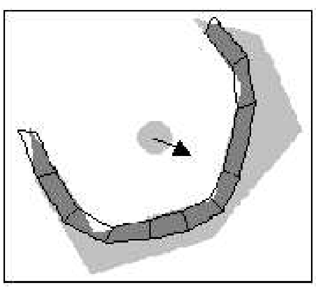

(a)

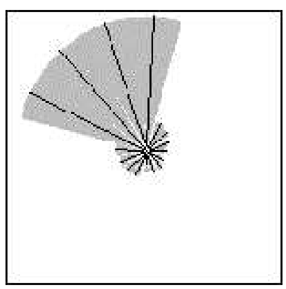

(b)

The robot is surrounded on the front-right.

(c)

Fig. 12. Surrounds descriptions. First case.

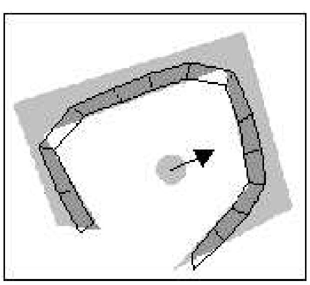

(a)

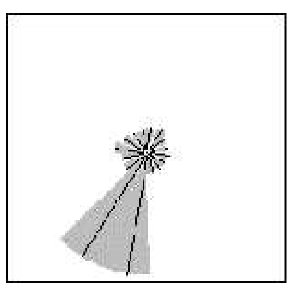

(b)

The robot is surrounded,
but there is an opening on the rear-right.

(c)

Fig. 13. Surrounds descriptions. Second case.

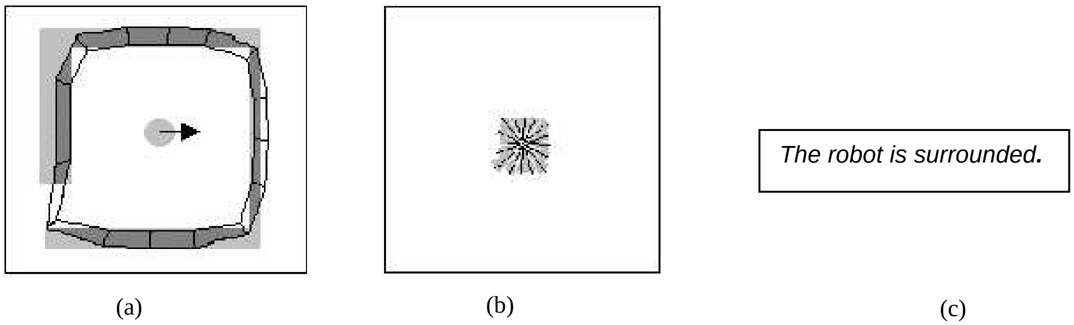

Fig. 14. Surrounds descriptions. Third case.

3.5. Coarse Descriptions

Generating a detailed description for each detected object may not be useful in all cases. Different levels of abstraction may be required for different types of tasks. Detailed, individual descriptions may be necessary for navigating in a cluttered environment (*e.g.*, navigating through a corridor with many obstacles), whereas coarser descriptions may be more appropriate for reasoning about the environment or depicting an orientation within it (*e.g.*, "I am in a corridor, facing a wall"). Moreover, if many objects are detected, the human user may be overwhelmed with the flow of information.

We provide two levels of abstraction. The higher level is obtained by first mapping the objects into 16 symbolic spots evenly distributed around the robot (Fig. 15). Note that an object is not mapped if it surrounds the robot (the associated description is then the same in both levels of abstraction). Spots are assigned according to the primary and secondary directional relationships from each detailed (low-level) linguistic description. Each object will have a corresponding primary and secondary direction depending on its position relative to the robot. For example, an object that is mostly left but somewhat forward will have a primary direction of left and a secondary direction of front. An object that is mostly in front but somewhat to the left will have a primary direction of front and a secondary direction of left. Self-assessments and range information are ignored. As shown in Fig. 15, two or more objects may be assigned to the same spot.

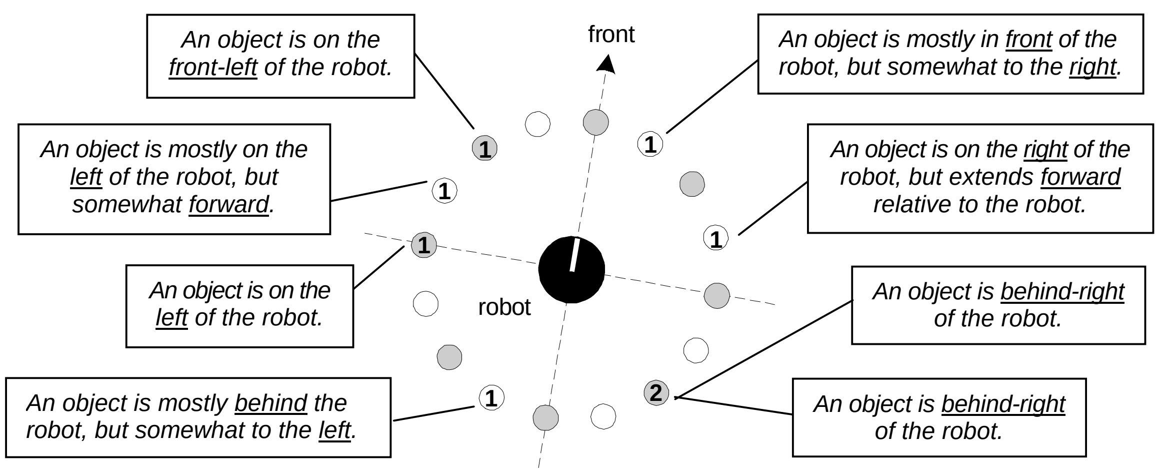

Fig. 15. Object mapping. Sixteen symbolic spots are oriented around the robot. Each one may be "occupied" by one or more objects, depending on the primary and secondary directional relationships. The shaded spots correspond to the primitive and compound directions.

Next, we associate a set of five adjoining spots to each primitive direction (right, front, left, behind) and to each compound direction (front-right, front-left, behind-left, behind-right). The set is centered on the spot that corresponds to the considered primary or compound direction, as shown in Fig. 16. The individual detailed descriptions for all objects that have been assigned to spots in the set are then replaced by a single linguistic description. However, one or both of the following conditions must be met:

1. There is at least one object in the center spot (Fig. 16b).
2. There is at least one object in each one of the spots that are adjacent to the center spot (Fig. 16c).

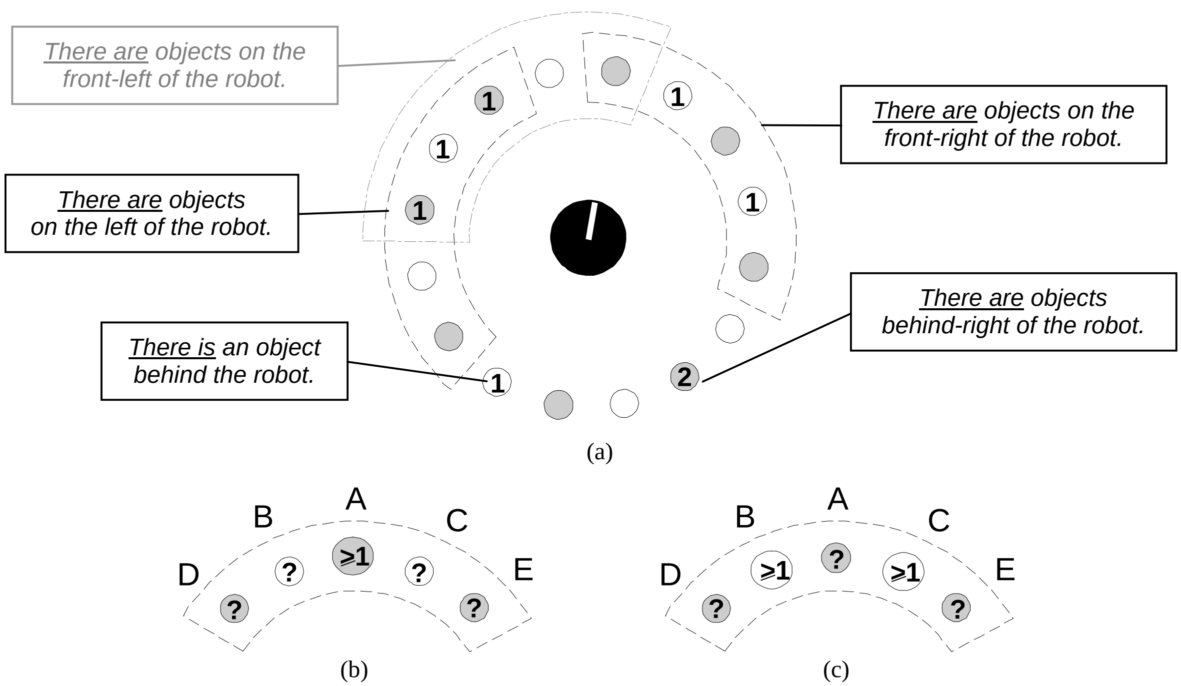

Fig. 16. Object grouping from Figure 15. First, a set of five symbolic spots is associated with each primitive direction and each compound direction. (a) shows three such sets. In a given primitive or compound direction, object grouping is performed if either (b) the center spot A is occupied, or (c) the spots B and C adjacent to A are both occupied.

For each of the occupied spots that is not part of an assigned set after completing the above process, a separate description is produced. If only one object is in the spot, then a simplified description for this object is generated with only the primary directional relationship from the detailed linguistic expressions (*e.g.*, Fig. 16a, “There is an object behind the robot”). If two or more objects are in the spot, then a single description is generated for all of them (*e.g.*, “There are objects behind-right of the robot”). In addition, any description that involves a compound direction is omitted if a description is generated for one of the adjacent primitive directions. For instance, in Fig. 16a, the expression “There are objects on the front-left of the robot” will be ignored since “There are objects on the left of the robot.” Fig. 17 shows an example of high-level linguistic abstraction.

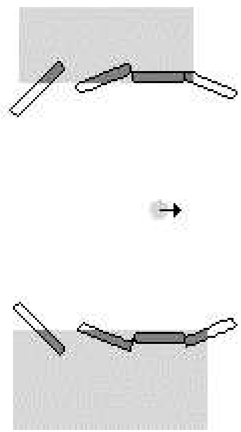

(a)

There are objects on the left of and on the right of the robot.

(b)

Fig. 17. Multi-level linguistic spatial descriptions. (a) A top view of the environment shows that the robot senses four objects on the left and four objects on the right. (b) Coarse description (high-level abstraction).

4. Experimental Results

To test our approach for generating egocentric linguistic spatial descriptions, we imagined a more realistic navigation scenario and created an environment with corridors and doorways, as shown in Fig. 18. For convenience in illustrating a variety of situations, the experiment was done using the Nomadic simulator.

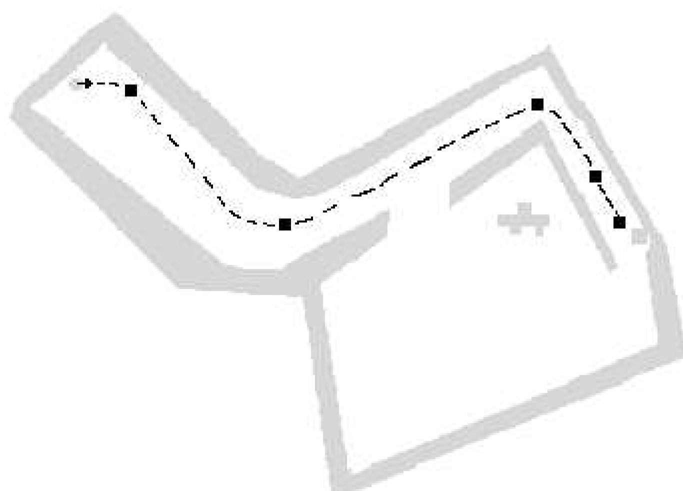

Fig. 18. Test environment. The robot starts in the upper left corner.

The robot was moved along the corridor, and the spatial descriptions were generated at each time step along the trajectory. Snapshots of the results are displayed in Figs. 19 through 26 for selected positions. The figures show the robot position and heading with respect to the environment, the polygonal models built from the sensor readings, and both detailed and coarse descriptions. These are listed according to a counterclockwise scanning around the robot; the objects in front are considered first. The program runs fast enough for interactive, real-time execution, on a Pentium II or Pentium III PC. The tasks of modeling obstacles, computing histograms, and generating expressions in natural language are executed faster than the robot can move, so there is no delay in the linguistic feedback displayed to the user.

The results show reasonable descriptions, given that the robot is trying to find a path through some occupied space. The detailed expressions are realistic for the models built from the sensor readings. For example, in Fig. 19, five objects are detected and each description is assessed as being "satisfactory." One object is detected on the robot's left because the left wall is close, whereas three others are detected on the right because the right wall is farther away. One might argue that the objects are not always an accurate representation of the actual environment, and they do reflect the limitation of using a static view from the sixteen sonar sensors. However, the polygonal models and the resulting descriptions are successful in capturing a static, qualitative image of the environment. More resolution can be obtained by fusing sensor readings over time. One possible technique is the occupancy grid map, in which evidence is accumulated from multiple sensors as the robot moves around the environment (e.g., Martin and Moravec, 1996). Our initial work using evidence grid maps for generating spatial descriptions can be found in (Skubic *et al.*, 2002).

The figures also show the coarse descriptions generated. For example, in Fig. 19, two groups of objects are mentioned: one on the right and one in the rear. The one object on the left is listed by itself. Note that distance information is ignored when grouping the objects, although this might be useful in some situations. In Fig. 21 for instance, the coarse description states that "there are objects on the front-right of the robot." However, the object in front is quite close and the one on the right is quite far.

The ability to detect whether the robot is surrounded illustrates high-level reasoning and provides a useful abstraction for navigation. Figs. 23, 24, and 25 depict the three levels of surroundedness. As explained in Section 3.4, this particular relationship is examined for a single object, and only if the self-assessment of the detailed description is less than satisfactory. The condition is necessary but not sufficient. Fig. 21 shows an example of a "rather satisfactory" assessment where the object does not quite enclose the robot enough to warrant a surrounds description.

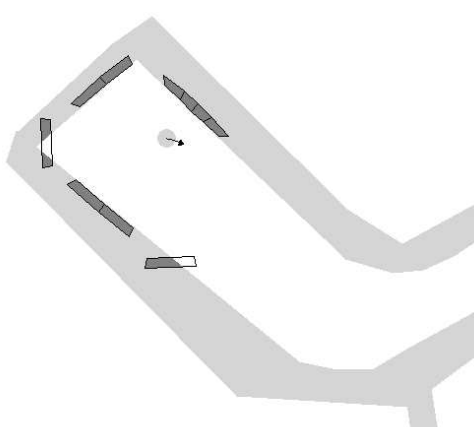

COARSE DESCRIPTIONS:

There are objects behind and on the right of the robot.
An object is on the left of the robot.

DETAILED DESCRIPTIONS:

An object is on the left of the robot,
but extends forward relative to the robot
(the description is satisfactory).
The object is very close to the robot.

An object is mostly behind the robot,
but somewhat to the left
(the description is satisfactory).
The object is close to the robot.

An object is mostly behind the robot,
but somewhat to the right
(the description is satisfactory).
The object is close to the robot.

An object is mostly on the right of the robot,
but somewhat to the rear
(the description is satisfactory).
The object is close to the robot.

An object is mostly on the right of the robot,
but somewhat forward
(the description is satisfactory).
The object is close to the robot.

Fig. 19. Snapshot 1.

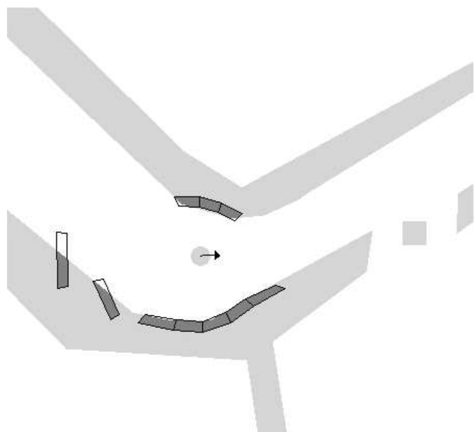

COARSE DESCRIPTIONS:

There are objects behind the robot.
An object is on the left of the robot.
An object is on the right of the robot.

DETAILED DESCRIPTIONS:

An object is on the left of the robot,
but extends forward relative to the robot
(the description is satisfactory).
The object is very close to the robot.

An object is behind the robot
(the description is satisfactory).
The object is close to the robot.

An object is mostly behind the robot,
but somewhat to the right
(the description is satisfactory).
The object is close to the robot.

An object is on the right of the robot,
but extends forward relative to the robot
(the description is satisfactory).
The object is very close to the robot.

Fig. 20. Snapshot 2.

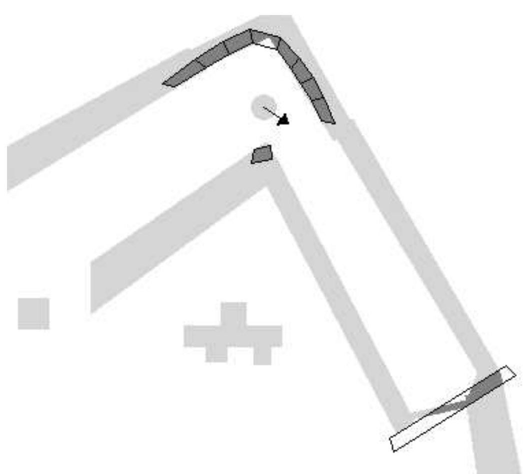

COARSE DESCRIPTIONS:

There are objects on the front-right of the robot.
An object is loosely on the left of the robot.

DETAILED DESCRIPTIONS:

An object is loosely on the left of the robot
and extends to the rear relative to the robot
(the description is rather satisfactory).
The object is very close to the robot.

An object is mostly on the right of the robot,
but somewhat forward
(the description is satisfactory).
The object is very close to the robot.

An object is mostly in front of the robot,
but somewhat to the right
(the description is satisfactory).
The object is very far from the robot.

Fig. 21. Snapshot 3.

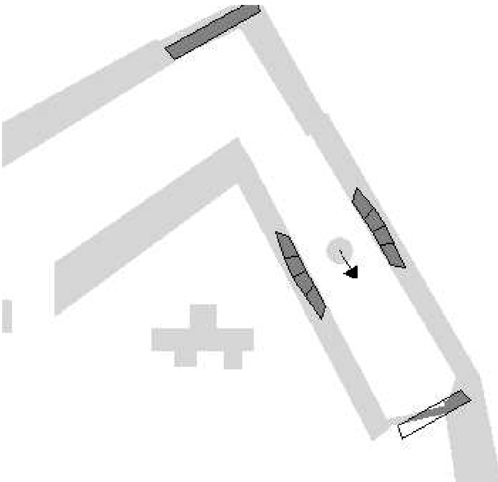

DETAILED DESCRIPTIONS:

An object is in front of the robot
(the description is satisfactory).
The object is close to the robot.

An object is on the left of the robot
(the description is satisfactory).
The object is very close to the robot.

An object is behind the robot
(the description is satisfactory).
The object is far from the robot.

An object is on the right of the robot
(the description is satisfactory).
The object is very close to the robot.

COARSE DESCRIPTIONS:

An object is in front of the robot.
An object is on the left of the robot.
An object is behind the robot.
An object is on the right of the robot.

Fig. 22. Snapshot 4.

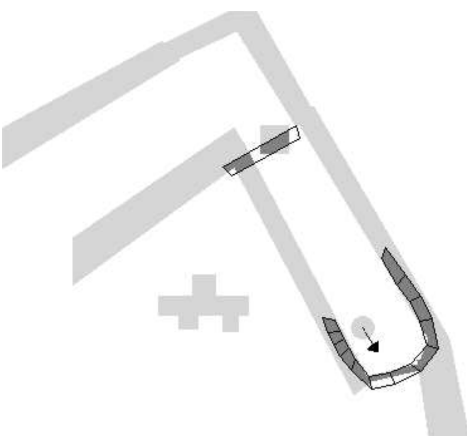

DETAILED DESCRIPTIONS:

The robot is surrounded on the front.

An object is behind the robot
(the description is satisfactory).
The object is far from the robot.

COARSE DESCRIPTIONS:

The robot is surrounded on the front.
An object is behind the robot.

Fig. 23. Snapshot 5.

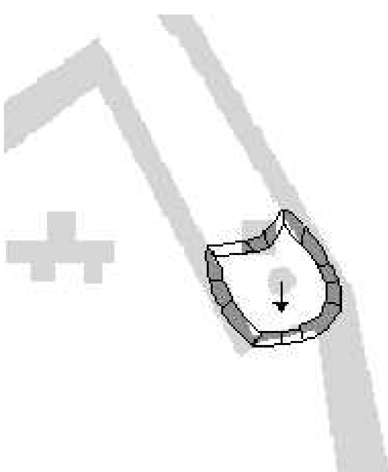

The robot is surrounded.

Fig. 24. Snapshot 6. In this case, the detailed and coarse descriptions are the same.

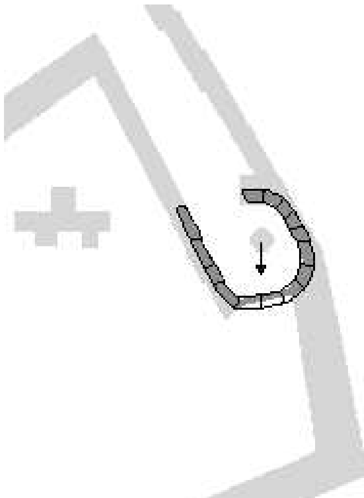

The robot is surrounded,
but there is an opening on the rear-right.

Fig. 25. Snapshot 7. Again, the detailed and coarse descriptions are the same.

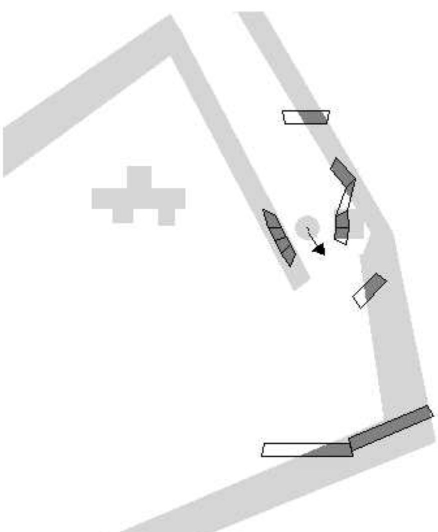

DETAILED DESCRIPTIONS:

An object is in front of the robot (the description is satisfactory). The object is far from the robot.

An object is mostly in front of the robot, but somewhat to the left (the description is satisfactory). The object is close to the robot.

An object is on the left of the robot (the description is satisfactory). The object is very close to the robot.

An object is mostly behind the robot, but somewhat to the left (the description is satisfactory). The object is close to the robot.

An object is on the right of the robot (the description is satisfactory). The object is very close to the robot.

An object is mostly in front of the robot, but somewhat to the right (the description is satisfactory). The object is far from the robot.

COARSE DESCRIPTIONS:

There are objects in front of the robot.
An object is on the left of the robot.
An object is mostly behind the robot.
An object is on the right of the robot.

Fig. 26. Snapshot 8.

5. Concluding Remarks

In this paper, we showed how linguistic spatial descriptions can be obtained from the sonar readings of a mobile robot. First, the robot's sensed environment is modeled by polygonal objects. From the models, force histograms are computed, and then detailed descriptions are generated. Each description includes directional and distance information about the object considered, as well as an assessment of its own adequacy. If the self-assessment shows that the description is not adequate, the surrounded condition is examined. Finally, objects are grouped around the robot, and a coarse description is also provided. To illustrate our approach, we used a mobile robot equipped with a ring of sixteen sonar sensors; however, the methodology is applicable to any ringed configuration of range sensors.

There are three significant contributions to this work. First, we have demonstrated how force histograms can be generated from vector data, allowing real-time output to the user. Second, we have proposed a new model of the spatial relationship "surrounds." It gives us the ability to distinguish between several linguistic, qualitative degrees of surroundedness. Finally, we have shown that different levels of abstraction can be considered in robot-to-human communication. The motivation for this work is to facilitate interactive robot interfaces for novice robot users. The histogram of forces offers an underlying mathematical framework for spatial reasoning and allows a mobile robot with sonar sensors to describe its environment in linguistic terms easily understood by any user, *e.g.*, "there is an object on the right." In the future, we will incorporate human-to-robot communication and our system will support high-level directives like "move forward while there is an object on the right." Cognitive models show that this approach is consistent with human spatial reasoning for navigation.

The experimental results provide some new insight into how spatial information can be utilized for human-robot interfaces. In particular, we have begun investigating how the presented methodology might support an interactive dialog between a human user and a robot (Skubic *et al.*, 2002). Given the current limitations in real-time sensing, object recognition is still a very difficult problem. However, in an interactive interface, the user can augment the robot's sensing and perception capabilities by performing the difficult task of recognition and assigning labels to the objects. The dialog can then utilize spatial references and refer to an object by name. We are also investigating the use of different reference frames, *i.e.*, egocentric *vs.* allocentric ones (Klatzky, 1998). Should the description say "There is an object on the left of the robot" (or "on the left of *me*") or "There is an object *west* of the robot"? Since our system is dictionary-driven, and given the geometric properties of the histogram of forces, the terminology and reference frame can easily be changed. The question is which terminology and frame of reference are easier for a user and under what situations. We are addressing some of these questions in a forthcoming study. Finally, we are also planning to incorporate reasoning about the environment. Rules can be formulated, *e.g.*, if there is something on the left (that extends forward and to the rear), something on the right (that extends forward and to the rear), and nothing (close) in front or to the rear, then we are in a corridor. Another option is to use the force histograms directly, as they represent a rich source of spatial information.

Acknowledgements

The authors wish to acknowledge support from ONR, grant N00014-96-0439, and the IEEE Neural Network Council for a graduate student summer fellowship for Mr. Chronis.

References

Bloch, I. 1999. "Fuzzy Relative Position between Objects in Image Processing: New Definition and Properties Based on a Morphological Approach," *Int. J. of Uncertainty Fuzziness and Knowledge-Based Systems*, vol. 7, no. 2, pp. 99-133.

Cutnell, J. D. and Johnson, K. 2001. *Physics*, Wiley & Sons, 5th Edition.

Dubois, D. and Prade, H. 1980. *Fuzzy Sets and Systems: Theory and Applications*, Academic Press, New York.

Freeman, J. 1975. "The Modelling of Spatial Relations," *Computer Graphics and Image Processing (4)*, pp. 156-171.

Gapp, K. P. 1995. "Angle, Distance, Shape, and their Relationship to Projective Relations," *Proceedings*, *17th Conf. of the Cognitive Science Society*.

Gribble, W., Browning, R., Hewett, M., Remolina, E. and Kuipers, B.. 1998. "Integrating vision and spatial reasoning for assistive navigation," In *Assistive Technology and Artificial Intelligence*, V. Mittal, H. Yanco, J. Aronis and R. Simpson, ed., Springer Verlag, pp. 179-193, Berlin, Germany.

Klatzky, R. 1998. "Allocentric and Egocentric Spatial Representations: Definitions, Distinctions, and Interconnections," in *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*, C. Freksa, C. Habel and K. Wender (ed.), pp. 1-17. Springer-Verlag, Berlin.

Klir, G. and Folger, T. 1988. *Fuzzy Sets, Uncertainty, and Information*, Englewood Cliffs, NJ: Prentice Hall.

Krieg-Brückner, B., Röfer, T., Carmesin, H.-O. and Müller, R. 1998. "A Taxonomy of Spatial Knowledge for Navigation and Its Application to the Bremen Autonomous Wheelchair," in *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*, C. Freksa, C. Habel and K. Wender (ed.), pp. 373-397. Springer-Verlag, Berlin.

Krishnapuram, R., Keller, J. M. and Ma, Y. 1993. "Quantitative Analysis of Properties and Spatial Relations of Fuzzy Image Regions," *IEEE Trans. on Fuzzy Systems*, vol. 1, no. 3, pp. 222-233.

Kuipers, B. 1998. "A Hierarchy of Qualitative Representations for Space," in *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*, C. Freksa, C. Habel and K. Wender (ed.), pp. 337-350. Springer-Verlag, Berlin.

Martin, M.C. and Moravec, H.P. 1996. "Robot Evidence Grids", Technical Report #CMU-RI-TR-96-06, Carnegie Mellon University, Mar.

Matsakis, P. 1998. *Relations spatiales structurelles et interprétation d'images*, Ph. D. Thesis, Institut de Recherche en Informatique de Toulouse, France.

Matsakis, P. , Keller, J. and Wendling, L. 1999. "F-Histograms and Fuzzy Directional Spatial Relations," *Proceedings, LFA'99 (French-Speaking Conference on Fuzzy Logic and Its Applications)*, Valenciennes, France.

Matsakis, P., Keller, J. M., Wendling, L., Marjamaa, J. and Sjahputera, O. 2001. "Linguistic Description of Relative Positions in Images," *IEEE Trans. on Systems, Man and Cybernetics, Part B*, vol. 31, no. 4, pp. 573-588.

Matsakis, P. and Wendling, L. 1999. "A New Way to Represent the Relative Position between Areal Objects," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 21, no. 7, pp. 634-643, 1999.

Miyajima, K. and Ralescu, A. 1994. "Spatial Organization in 2D Segmented Images: Representation and Recognition of Primitive Spatial Relations," *Fuzzy Sets and Systems*, vol. 65, iss. 2/3, pp. 225-236.

Morik, K., Kaiser, M. and Klingspor, V., ed. 1999. *Making Robots Smarter*, Kluwer Academic Publishers, Boston.

Moratz, R. and Fischer, K. 2000. "Cognitively Adequate Modeling of Spatial Reference in Human-Robot Interaction," *Proceedings of the 12th IEEE Int. Conf. On Tools with Artificial Intelligence*, Vancouver, B.C., pp. 222-228.

Perzanowski, D., Schultz, A., Adams, W. and Marsh, E.1999. "Goal Tracking in a Natural Language Interface: Towards Achieving Adjustable Autonomy," In *Proceedings of the 1999 IEEE Int. Symp. on Computational Intelligence in Robotics and Automation*, Monterey, CA, Nov., pp.208-213.

Previc, F.H. 1998. "The Neuropsychology of 3-D Space," *Psychological Review*, vol. 124, No. 2, pp. 123-164.

Retz-Schmidt, G. 1988. "Various Views on Spatial Prepositions," *AI Magazine*, vol. 9, no. 2, pp. 95-105.

Rosenfeld, A. and Klette, R. 1985. “Degree of Adjacency or Surroundedness,” *Pattern Recognition*, vol. 18, no. 2, pp. 169-177.

Schunn, C. and Harrison, T. 2001. Personal communication.

Shibata, F., Ashida, M., Kakusho, K., Babaguchi, N., and Kitahashi, T. 1996. “Mobile Robot Navigation by User-Friendly Goal Specification,” In *Proceedings of the 5th IEEE Int. Workshop on Robot and Human Communication*, Tsukuba, Japan, Nov., pp. 439-444.

Skubic, M., Chronis, G., Matsakis, P. and Keller, J. 2001. “Generating Linguistic Spatial Descriptions from Sonar Readings Using the Histogram of Forces,” in *Proceedings of the IEEE 2001 Int. Conf. on Robotics and Automation*, Seoul, Korea, May, pp. 485-490.

Skubic, M., Perzanowski, D., Schultz, A., and Adams, W. 2002. “Using Spatial Language in a Human-Robot Dialog,” in *Proceedings of the IEEE 2002 Int. Conf. On Robotics and Automation*, Washington, D.C., May.

Skubic, M. and Volz, R.A.. 2000. “Acquiring Robust, Force-Based Assembly Skills from Human Demonstration,” *IEEE Trans. on Robotics and Automation*, vol.16, no.6, Dec., pp.772-781.

Stopp, E., Gapp, K.-P., Herzog, G., Laengle, T. and Lueth, T. 1994. “Utilizing Spatial Relations for Natural Language Access to an Autonomous Mobile Robot,” In *Proceedings of the 18th German Annual Conference on Artificial Intelligence*, Berlin, Germany, pp. 39-50.

Wang, X. and Keller, J. M. 1997. “Fuzzy Surroundedness,” *Proceedings, Sixth IEEE International Conference on Fuzzy Systems*, vol. 2, pp. 1173-8.