Hand-Drawn Maps for Robot Navigation

Marjorie Skubic, Sam Blisard, Andy Carle, and Pascal Matsakis

Dept. of Computer Engineering and Computer Science
University of Missouri-Columbia
Columbia, MO 65211
skubicm@missouri.edu

Abstract

The goal of this work is to create a robot interface that allows a novice user to guide, control, and/or program a robot to perform some task. The assumption is that, although the user may be a domain expert in how the task should be done, he is not an expert in robotics. During the actual robot use, he should focus on the task to be done rather than worrying about the robot or the interaction modality. To address this goal, we have been investigating the use of hand-drawn route maps to transfer navigation tasks to robots. In the paper, we provide an overview and current status of ongoing work with sketches. We discuss what type of information would be useful for directing and controlling a robot and then show how this information can be extracted from a sketched route map, in the form of spatial relationships. An analysis example of a PDA-generated sketch is included. Also, preliminary results are presented which compare the analysis of a sketched map with that of a real map.

Introduction

Being able to interact and communicate with robots in the same way we interact with people has long been a goal of AI and robotics researchers. However, much of the robotics research in the past has emphasized the goal of achieving autonomous agents. In our research, we are less concerned with creating autonomous robots that can plan and reason about tasks, and instead we view them as semi-autonomous tools that can assist a human user. The robot may have some perception capabilities, reactive behaviors, and perhaps limited reasoning abilities that allow it to handle an unstructured and dynamic environment. But the user supplies the high-level and difficult reasoning and strategic planning capabilities.

In this scenario, the interaction and communication between the robot and the human user becomes very important. The user must be able to easily communicate what needs to be done, perhaps at different levels of task abstraction. In particular, we would like to provide an intuitive method of communicating with robots that is easy for users that are not expert robotics engineers. We want domain experts to define their own task use of robots, which may involve controlling them, guiding them, or even programming them.

As one strategy for addressing this goal, we have been investigating the use of hand-drawn route maps, in which the user sketches an approximate representation of the environment and then sketches the desired robot trajectory with respect to that environment. The objective in the sketch interface is to extract spatial information about the map and a qualitative path through the landmarks drawn on the sketch. This information is used to build a task representation for the robot, which operates as a semi-autonomous vehicle. Note that the task representation is based on sensing and *relative* position, not *absolute* position.

Although qualitative navigation presents problems for autonomous robots (*e.g.*, due to perception difficulties; see a discussion in Murphy 2000), we believe that it is a good idea for semi-autonomous robots, where the user interactively observes and directs the robot. Qualitative navigation more closely mimics the human navigation process, and in the context of sketch interfaces, also allows for a more intuitive interface with the human user. Possible applications include the following:

1. Military applications. The user looks at a scene and sketches a route through landmarks. Also, programming strategic behaviors, such as how to search or how to escape.
2. Programming large construction or mining equipment.
3. Guiding planetary rovers.
4. Controlling personal robots.

In the remaining sections of the paper, we first discuss background material on human navigation and the use of sketched route maps. We illustrate how spatial relations can be used to analyze sketched maps and briefly describe our methodology for modeling spatial relationships based on the histogram of forces (Matsakis 1998). In addition, we provide a framework for the robot control architecture, and discuss what kind of information must be extracted from the sketched maps to facilitate the necessary robot control. Finally, sketch interpretation is illustrated with a map sketched on a PDA. We also present preliminary results, which compare the analysis of a sketched map with that of a real map. The conclusion includes a brief discussion on the current status and future directions.

Human Navigation and Sketched Route Maps

A sketched route map is drawn to help someone navigate along a path for the purpose of reaching a goal. An example is shown in Fig. 1a. Although route maps do not generally contain complete map information about a region, they do provide relevant information for the navigation task. People sketch route maps to include landmarks at key points along the path and use spatial relationships to help depict the route, often adding arrows and other notation for clarity (Tversky and Lee 1998). In a study of 29 sketched route maps, each contained the information necessary to complete a complicated navigation task (Tversky and Lee 1998).

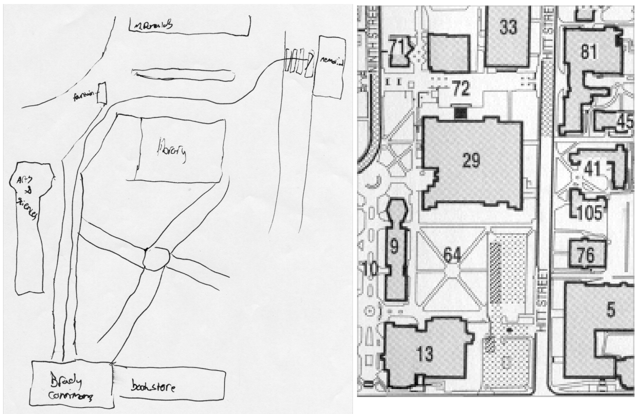

(a) (b)

Figure 1. (a) A map sketched on paper, describing a route through the MU campus. (b) A view of the actual map. Buildings 13, 9, 29 and 81 correspond to major landmarks included on the sketched map.

Research by Michon and Denis (2001) provides insights into how landmarks are used for human navigation and what are considered to be key route points. In studying route directions, they found that landmarks were used more frequently at four types of critical nodes: (1) the starting point, (2) the ending point, (3) at a change in orientation, and (4) at places along the route where errors could easily occur, such as major intersections (Michon and Denis 2001). Thus, people use the relative position of landmarks as cues to keep on track and to determine when to turn left or right.

The work of the researchers noted above and others (*e.g.*, Previc 1998, Schunn and Harrison 2001) indicate the importance of environment landmarks and spatial relationships in human navigation. The work suggests that spatial relationships of landmarks with respect to the desired route may be useful not only for robot control but also as a link between a robot and its human user. In the next section, we describe a tool for modeling spatial relationships that is fast, robust, and handles all object contours using either raster data or vector data (*i.e.*, a boundary representation).

Modeling Spatial Relationships

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. By introducing the histogram of angles, 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).

In the context of image analysis, Matsakis and Wendling (1999) introduced the histogram of forces. Contrary to the angle histogram, it ensures processing of both raster data and vector data. Moreover, it offers solid theoretical guarantees, allows explicit and variable accounting of metric information, and lends itself, with great flexibility, to the definition of fuzzy directional spatial relations (such as “to the right of,” “in front of,” etc.). 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 an allocentric (world) view and an egocentric (robot) view.

The Histogram of Forces

The relative position of a 2D object A with regard to another object B is represented by a function F^{AB} from $\mathbb{R}$ into $\mathbb{R}_+$. For any direction θ, the value $F^{AB}(\theta)$ is the total weight of the arguments that can be found in order to support the proposition “A is in direction θ of B”. More precisely, it is the scalar resultant of elementary forces. These forces are exerted by the points of A on those of B, and each tends to move B in direction θ (Fig. 2). F^{AB} is called the *histogram of forces associated with* (A,B) *via* F, or the F−*histogram associated with* (A,B). The object A is the *argument*, and the object B the *referent*. Actually, the letter F denotes a numerical function. Let r be a real number. If the elementary forces are in inverse ratio to d^r, where d represents the distance between the points considered, then F is denoted by F_r. The F_0–histogram (histogram of constant forces) and F_2–histogram (histogram of gravitational forces) have very different and very interesting characteristics. The former coincides with the angle histogram—without its weaknesses—and provides a global view of the situation. It considers the closest parts and the farthest parts of the objects equally, whereas the F_2–histogram focuses on the closest parts.

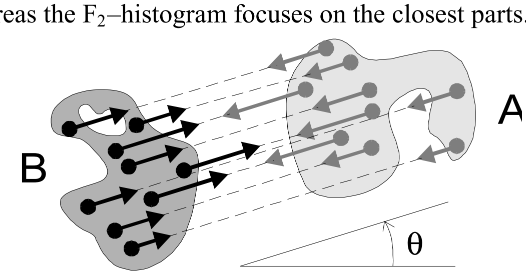

Figure 2. Computation of $F^{AB}(\theta)$. It is the scalar resultant of forces (black arrows). Each one tends to move B in direction θ.

Throughout this paper, the referent B is the robot. The F-histogram associated with (A,B) is represented by a limited number of values (*i.e.*, the set of directions θ is made discrete), and the objects A and B are assimilated to polygons using vector data. The computation of F^{AB} is of complexity O(n log(n)), where n denotes the total number of vertices (Matsakis and Wendling 1999). Details on the handling of vector data can also be found in Skubic *et. al.* (2001a, 2002a).

Linguistic Description of Relative Positions

The histogram of forces provides a tool for modeling spatial relationships; the model can also be used to build qualitative spatial descriptions that provide a linguistic link to the user. Matsakis *et al.* (2001) present such a system that produces linguistic spatial descriptions of images.

The description of the relative position between any 2D objects A and B relies on the sole primitive directional relationships: "to the right of," "above," "to the left of" and "below" (imagine that the objects are drawn on a vertical surface). It is generated from F_0^{AB} (the histogram of constant forces associated with (A,B)) and F_2^{AB} (the histogram of gravitational forces). First, eight values are extracted from the analysis of each histogram: a_r (RIGHT), b_r (RIGHT), a_r (ABOVE), b_r (ABOVE), a_r (LEFT), b_r (LEFT), a_r (BELOW) and b_r (BELOW). They represent the "opinion" given by the considered histogram (*i.e.*, F_0^{AB} if r=0, and F_2^{AB} if r=2).

For instance, according to F_0^{AB} the degree of truth of the proposition "A is to the right of B" is a_0(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_0^{AB} the maximum degree of truth that can reasonably be attached to the proposition (say, by another source of information) is b_0(RIGHT) (which belongs to the interval [a_0(RIGHT),1]).

F_0^{AB} and F_2^{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 fuzzy rules and meta-rules that outputs the expected linguistic description. The system handles a set of adverbs (like "mostly," "perfectly," etc.), which 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"). The second part supplements the description and 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, it indicates to what extent it is necessary to utilize other spatial relations such as "surrounds". When range information is available, a fourth part can also be generated to describe distance (*e.g.*, "A is close to B") (Skubic *et. al.* 2001a).

Framework for Human-Robot Interaction

Robot navigation is modeled as a procedural task *(i.e.*, a sequence of steps) to mimic the human navigation process. The framework for the robot control architecture is shown in Fig. 3. Task structure is represented as a Finite State Automaton (FSA) in the S*upervisory Controller*, following the formalism of the Discrete Event System (Ramadge and Wonham 1989). The FSA models a sequence of moves, each of which is governed by a robot behavior (a local control strategy). The complete sequence comprises a task. The sensor-based qualitative state (QS) is used for task segmentation. A change in QS is an event that corresponds to a change in the type of movement.

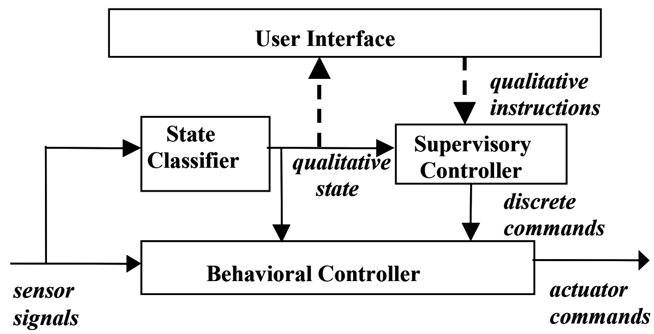

Figure 3. The Robot Control Architecture

For navigation tasks, the QS is formed by the spatial relationships of environment landmarks with respect to the robot. Thus, the robot uses landmarks in the same way that a person would use landmarks. Through the *State Classifier*, the robot is provided with the ability to recognize a set of qualitative states, which are extracted from sensory information, thus reflecting the current environmental condition. For navigation, robot-centered spatial relations provide context (*e.g.*, "there is an object to the left front"). Adding the ability to recognize classes of objects provides additional perception (*e.g.*, "there is a person to the left front").

The robot is also equipped with a set of (reactive) behaviors that are managed by the *Behavioral Controller*. Reactive behaviors allow the robot to respond quickly and safely to dynamic and unexpected conditions, such as avoiding moving obstacles. Output from the behavioral controller is merged with discrete commands issued from the Supervisory Controller. Note that this combination of discrete event control in the Supervisory Controller and the "signal processing" in the Behavioral Controller is consistent with Brockett's (1993) framework of hybrid control systems and has similarities to other approaches used for qualitative robot navigation (*e.g.*, Kuipers 1998).

As shown in Fig. 3, the interface between the robot and the human user relies on the qualitative state for two-way communications. In robot-to-human communications, the QS allows the user to monitor the current state of the robot, ideally in terms easily understood (*e.g.*, "there is an object on the right"). In human-to-robot communications, commands are segmented by the QS, termed qualitative instructions in the figure (*e.g.*, “while there is an object on the right, move forward”). This illustrates the type of information that must be extracted from the sketched route maps, in order to direct the robot along the intended path. We extract a sequence of qualitative states in the form of spatial relationships of environment landmarks with respect to the robot. And we extract the corresponding robot movements for each QS which make up the set of qualitative instructions for the desired task.

Interpreting a PDA-Sketched Map

We now have all of the pieces for analyzing a sketched route map, and in this section we illustrate how navigation information is extracted from a map sketched on a PDA such as a PalmPilot. The stylus interface of the PDA allows the user to sketch a map much as she would on paper for a human colleague. The PDA captures the string of (x,y) coordinates sketched on the screen which forms a digital representation suitable for processing.

The user first draws a representation of the environment by sketching the approximate boundary of each object. During the sketching process, a delimiter is included to separate the string of coordinates for each object in the environment. After all of the environment objects have been drawn, another delimiter is included to indicate the start of the robot trajectory, and the user sketches the desired path of the robot, relative to the sketched environment. An example of a sketch is shown in Fig. 4, where each point represents a captured screen pixel.

The extraction of spatial information from the sketch is summarized in Fig. 5. For each point along the trajectory, a view of the environment is built, using the radius of the sensor range. For each object within the sensory radius, a polygonal region is built using the boundary coordinates of the object as vertices.

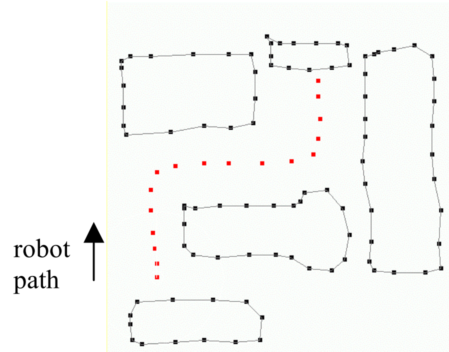

Figure 4. A route map sketched on a PDA.

We have used different strategies for building the polygonal representations of the objects, for example, using only the points of the object that fall within the sensory radius (Skubic *et. al.* 2001b). Here, if any of the object points lie within the sensory radius, we use the entire object boundary. This approach coincides more closely with our recent work using occupancy grid cells (Skubic *et. al.* 2002b).

Once the polygonal region of an object is built, the histograms of constant forces and gravitational forces are computed as described previously. The referent is always the robot, which is modeled as a square for the histogram computations. To capture robot-centered spatial relationships, the robot orientation must also be considered. The robot’s heading is computed using adjacent points along the sketched path to determine an instantaneous orientation. We compensate for the discrete pixels by averaging 5 adjacent points (centered on the considered trajectory point), thereby filtering small perturbations and computing a smooth transition as the orientation changes. After the heading is calculated, it is used to shift the histograms along the horizontal axis to produce an egocentric (robot) view.

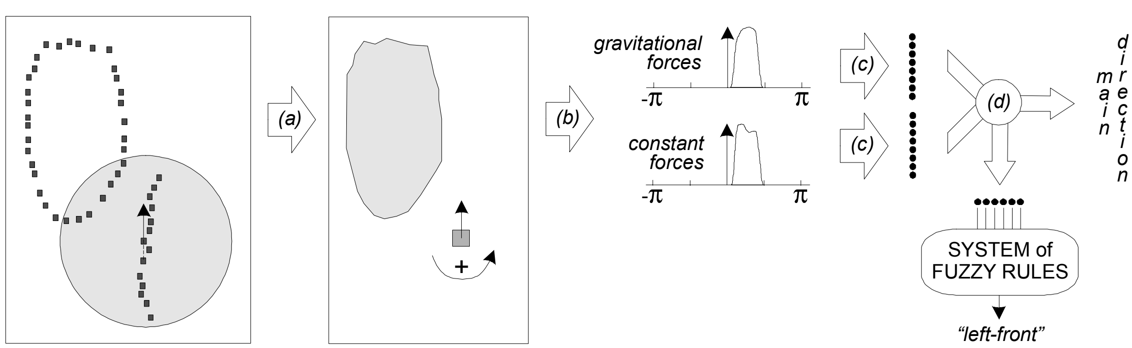

Figure 5. Synoptic diagram showing how spatial information is extracted from the sketch.

The histograms of constant forces and gravitational forces associated with the robot and the polygonal region of each object are used to generate a linguistic description of the relative position. In addition, features can be extracted from the histograms to further represent the spatial relationship. In processing the sketch, we extract what is considered to be the "main direction" of the object with respect to the robot and discretize it into one of 16 possible directions, as shown in Fig. 6. Examples of corresponding linguistic descriptions are also shown in Fig. 6 for a sampling of directions.

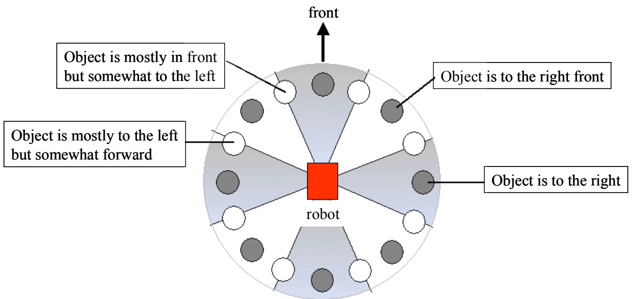

Figure 6. Sixteen directions are situated around the robot (the small circles). The "main direction" of each object is discretized into one of these 16 directions. Examples are included of corresponding linguistic descriptions.

In addition to extracting spatial information on the environment landmarks, we also extract the movement of the robot along the sketched path. The computation of the robot heading, described above, provides an instantaneous orientation. However, we also want to track the change in orientation over time and compute what would correspond to robot commands, *e.g.*, move forward, turn left, make a "hard" left. The turning rate is determined by computing the change in instantaneous heading between two adjacent route points and dividing by the distance between the points to normalize the rate. A positive rate means a turn to the left, and a negative rate means a turn to the right.

The spatial information and robot movement extracted from the sketched map in Fig. 4 is summarized in Fig. 7. In Fig. 7b, the main direction of each object is plotted for the route steps in which the object is "in view"; labels of the corresponding directions are displayed on the graph to show the symbolic connection. The normalized turning rate which tracks the robot movement along the trajectory is also shown in Fig. 7b. For reference, we have included a sampling of linguistic descriptions generated along the route (Fig. 7c), which correspond to the positions shown in Fig. 7a. Note that the descriptions have an egocentric (robot) perspective.

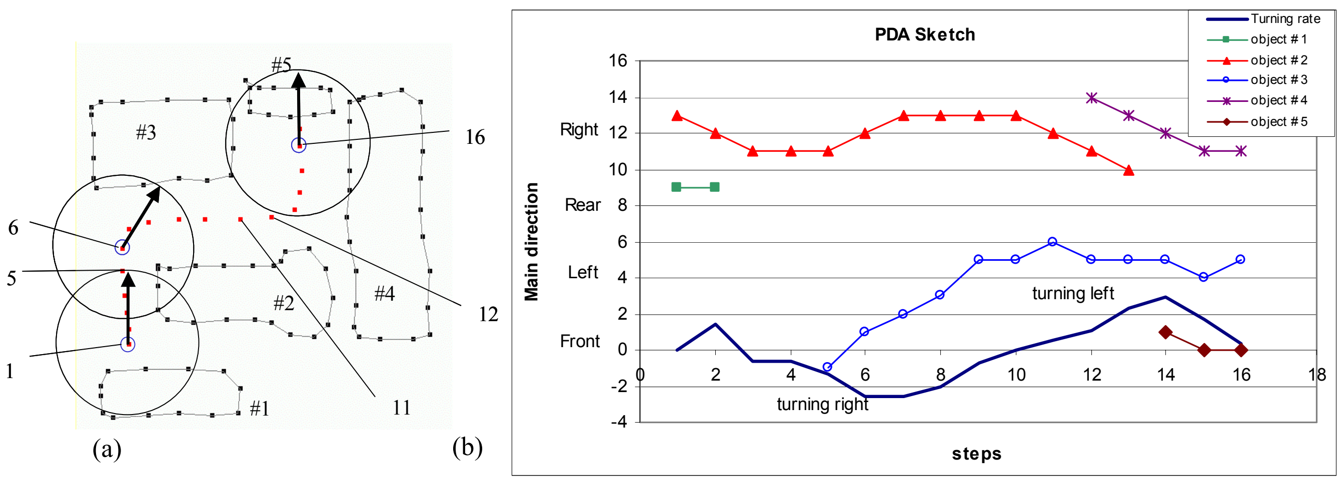

Figure 7. The PDA sketch.
(a) The original sketch with an overlay of the robot's sensory radius for several points along the route.
(b) Normalized turning rate of the robot along the sketched route with the corresponding discrete main directions of the objects.
(c) Generated egocentric linguistic descriptions for the route points shown in (a).

(c)

1. Object #1 is behind the robot but extends to the right relative to the robot
   Object #2 is to the right of the robot but extends forward relative to the robot
5. Object #2 is to the right of the robot but extends to the rear relative to the robot
   Object #3 is in front of the robot but extends to the right relative to the robot
6. Object #2 is to the right of the robot
   Object #3 is in front of the robot but extends to the left relative to the robot
11. Object #2 is to the right of the robot
    Object #3 is behind-left of the robot
12. Object #2 is nearly to the right of the robot but extends to the rear relative to the robot
    Object #3 is mostly to the left of the robot but somewhat to the rear
    Object #4 is loosely to the right-front of the robot
16. Object #3 is to the left of the robot but extends to the rear relative to the robot
    Object #4 is to the right of the robot but extends to the rear relative to the robot
    Object #5 is in front of the robot
    The object is very close to the robot.

The turning rate in Fig. 7b, although not translated into discrete robot commands, shows the general trend in the robot movement along the route and the correlation with relative positions of the environment landmarks. At the beginning of the route, when object #1 is behind the robot, the robot's movement is generally straight ahead (slightly to the left). When object #3 is in view, the robot turns to the right until the object is mostly on the left. When object #4 is in view to the front, the robot turns left and stops when object #5 is in front and very close. In this way, we can extract the key points along the route where a change in direction is made, by capturing the relative positions of the landmarks with respect to the route.

Comparing a Sketched Map to a Real Map

To further investigate the use of sketches, we also compared a sketched (qualitative) map to a real (quantitative) map. In particular, we wanted to test the hypothesis that, although users may not sketch a map to an accurate scale, nor even use accurate shapes, they do tend to use accurate spatial relationships. And we wanted to test how consistent these relationships are, as measured by our tool. In this section, we present preliminary results showing an analysis of the maps in Fig. 1.

Sketches were collected by asking students to draw a map showing the route on the MU campus, from Brady Commons to the Memorial Union, two well-known landmarks. Maps were sketched on paper. For analysis, the sketched maps, as well as the real map, were digitally scanned. A few key boundary points were manually extracted from the scanned maps for 4 building landmarks, to produce a digitized representation as shown in Figures 8a and 9a. The digitized version of the sketch is similar to the type of representation captured from a PDA-sketched map.

The maps have been processed as described in the previous section; the results are shown in Figures 8b and 9b. The one change from the previous analysis is that the main direction of each environment landmark is plotted for the complete route. This eliminates the need for setting a sensory radius that corresponds to both the sketched map and the real map. (We did not ask that the sketches be drawn to an accurate scale and we did not expect that they would be.) Also, note that a main direction of 16 in the graph is equivalent to a direction of 0 (front) because of a circular wrap. The higher numbers are used for convenience to show the continuous trend over the entire route.

In comparing the results, we can see a pattern that is similar although not completely identical. For example, in the first right turn, consider the two closest landmarks. For both maps, the A&S building is slightly left of the rear (main direction = 7) just before the turn and changes to exactly rear (main direction = 8) at the first step in the turn. During these same two steps, Ellis Library is slightly to the rear of right (main direction = 11) in the real map and exactly right (main direction = 12) in the sketched map. At other key steps where the route turns, one can observe similar results.

To further analyze the sketches, we also computed the spatial relationships between pairs of landmarks and compared the results of the real map with those of the sketched maps. This is less important for extracting a qualitative task representation of the navigation, but of interest nonetheless in analyzing the sketches. If the spatial relations between corresponding landmark pairs agree, then the spatial relations might be used as a basis for relating landmarks in a sketched map with those in an accurate map.

Six sketched maps were compared to the real map in Figure 1b. Again, each sketch was scanned and key boundary points were manually extracted as above for the same 4 building landmarks. At this time, we present preliminary results for the sketches, comparing only the main directions of each landmark pair. The sketched route maps represented a broad range in terms of accuracy, scale, and shape of the buildings used as landmarks. Also, the routes sketched did not always follow the same path with respect to the landmarks. However, in spite of these differences, the discrete main directions of two landmarks did not vary by more than +/- 2 values. In most cases, the values agreed or were within +/- 1, especially in the cases where the same route was taken in the sketch. These results are not comprehensive but they do show promise in using spatial relationships to analyze sketched route maps and compare them to accurate maps.

Concluding Remarks

The work on sketched route maps for robot navigation is by no means complete. The work presented here is merely a snapshot of an initial approach in extracting qualitative information from the sketch. We believe the real potential lies in a more interactive approach with the sketch interface. For example, the spatial states and behaviors could be displayed as the route is being sketched so that the user could change them if necessary. Also, editing gestures could be added, as in Landay and Myers (2001), allowing the user to delete landmarks, add labels to landmarks, and to specify qualitative distances (such as how close should the robot get to the landmarks).

Acknowledgements

Several MU colleagues and students have contributed to this overall effort, including Professor Jim Keller and students George Chronis and Ben Forrester. This research has been supported by ONR, the Naval Research Laboratory and the MU Discovery Fellowship Program.

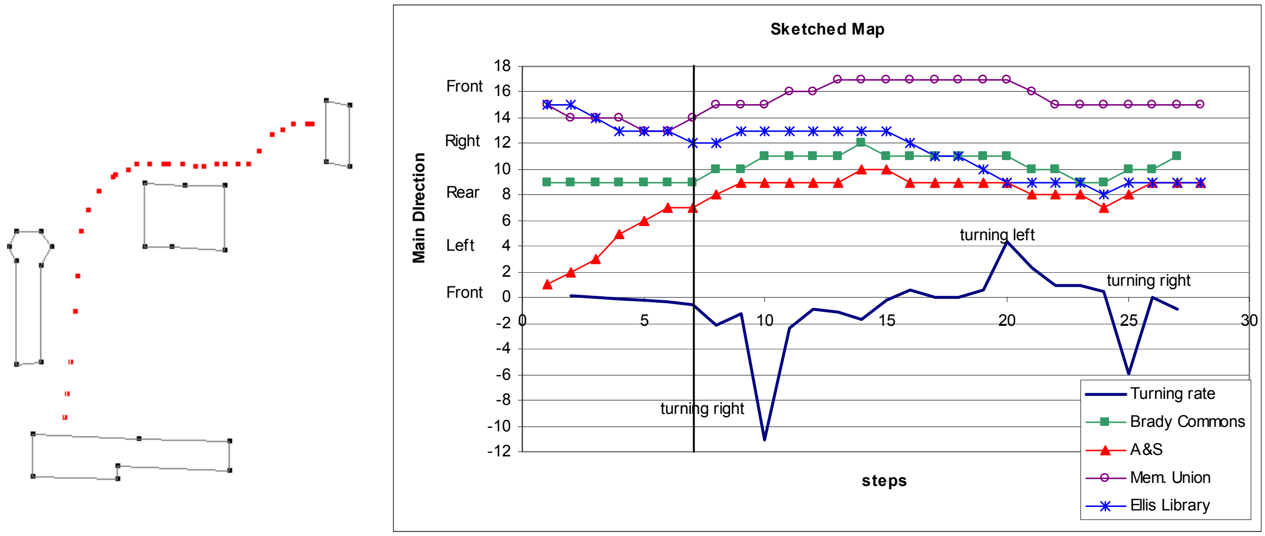

(a) (b)

Figure 8. The sketched route map corresponding to Figure 1a. (a) The digitized representation. (b) Normalized turning rate of the sketched path with the corresponding discrete main directions of the environment landmarks.

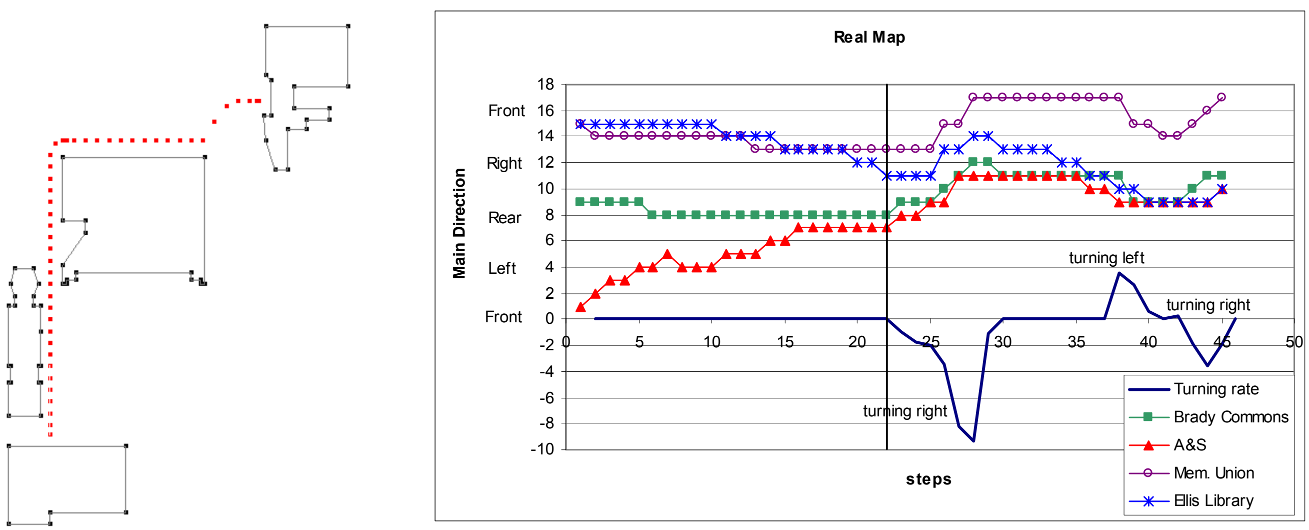

(a) (b)

Figure 9. The route map drawn on the real map in Figure 1b. (a) The digitized representation. (b) Normalized turning rate of the path with the corresponding discrete main directions of the environment landmarks.

References

Brockett, R.W. 1993. "Hybrid models for motion control systems," in *Essays on Control: Perspectives in the Theory and Its Applications*, H.L. Trentelman and J.C. Willems (ed.), pp.29-53. Boston, MA: Birkhauser.

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

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

Landay, J. and Myers, B. 2001. "Sketching Interfaces: Toward More Human Interface Design," *IEEE Computer*, 34(3):56-64

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

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*, 21(7):634-643.

Matsakis, P., Keller, J., Wendling, L., Marjamaa, J. and Sjahputera, O. 2001. "Linguistic Description of Relative Positions in Images", *IEEE Trans. on Systems, Man and Cybernetics*, part B, 31(4):573-588.

Michon, P.-E. and Denis, M. 2001. "When and Why Are Visual Landmarks Used in Giving Directions?" in *Spatial Information Theory: Foundations of Geographic Information Science*, D. Montello (ed.), Berlin: Springer, pp. 292-305.

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

Murphy, R. 2000. *Introduction to AI Robotics*. Cambridge, MA: MIT Press.

Perzanowski, D., Schultz, A., Adams, W., Marsh, E. and Bugajska, M. 2001. "Building a Multimodal Human-Robot Interface," *IEEE Intelligent Systems*, Jan./Feb, pp. 16-20.

Previc, F.H. 1998. "The Neuropsychology of 3-D Space," Psychological Review, 124(2):123-164.

Ramadge, P. and Wonham, W. 1989. "The control of discrete event systems," *Proceedings of the IEEE*, 77(1):81-97.

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

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

Skubic, M., Matsakis, P., Forrester, B and Chronis, G. 2001b. "Extracting Navigation States from a Hand-Drawn Map," in *Proceedings of the IEEE 2001 International Conference on Robotics and Automation*, May, Seoul, Korea, pp. 259-264.

Skubic, M., Chronis, G., Matsakis, P. and Keller, J. 2001c. "Spatial Relations for Tactical Robot Navigation," in *Proceedings of the SPIE, Unmanned Ground Vehicle Technology III*, April, Orlando, FL.

Skubic, M., Matsakis, P., Chronis, G. and Keller, J. 2002a. "Generating Multi-Level Linguistic Spatial Descriptions from Range Sensor Readings Using the Histogram of Forces," Submitted to *Autonomous Robots.*

Skubic, M., Perzanowski, D., Schultz, A. and Adams, W. 2002b. "Using Spatial Language in a Human-Robot Dialog," accepted for the IEEE 2002 International Conference on Robotics and Automation.

Tversky, B. and Lee, P. 1998. "How Space Structures Language," in *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*, C. Freksa, C. Habel, K. Wender (ed.), Berlin: Springer, pp. 157-176.