Scavenging with a Laptop Robot

Alan Davidson, Mac Mason, Susanna Ricco, Ben Tribelhorn, and Zachary Dodds

Harvey Mudd College Computer Science Department
1250 N Dartmouth Avenue
Claremont, CA 91711
adavidson, jmason, sricco, btribelhorn, dodds@hmc.edu

Abstract

This synopsis presents Harvey Mudd College's entry into the 2005 AAAI scavenger hunt competition. We are submiting a lap-controlled robot which uses commodity parts and limited sensors to localize itself and perform arrow following and object recognition.

Overview

Harvey Mudd College's entry into the 2005 AAAI robot scavenger hunt contest consists of a laptop robot, an Evolution Robotics' ER1, with added sensors and software (Figure 1). The ER1 is a differential-drive robot with a commodity laptop and web camera, IR sensors, and a home-brew sonar. Rather than use Evolution's ERSP suite of robotic algorithms, we have chosen to implement all the processing we need ourselves.

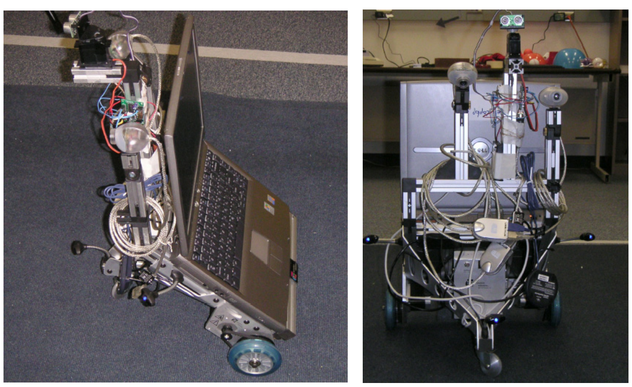

Figure 1 Side and front views of our laptop-controlled ER1

Low-level Control

Developed as part of an undergraduate survey of mobile robotics, this entry's basic operation consists of avoiding obstacles and exploring its environment. Layered atop these routines is an implementation of Monte Carlo Localization (MCL) (Thrun et al., 2001) based on sonar, vision, and empirically derived triangular probability distributions of sensor and motion error. Figure 2 shows our map and the hypotheses being maintained during MCL. We developed a visual routine for finding red moldings that appear "naturally" in our hallways to assist with MCL, but because these landmarks will not be present at the competition, the vision system now runs on top of the directed wandering routines, taking control as objects come into view. In the absence of target objects, the robot wanders according to Figure 3's state machine.

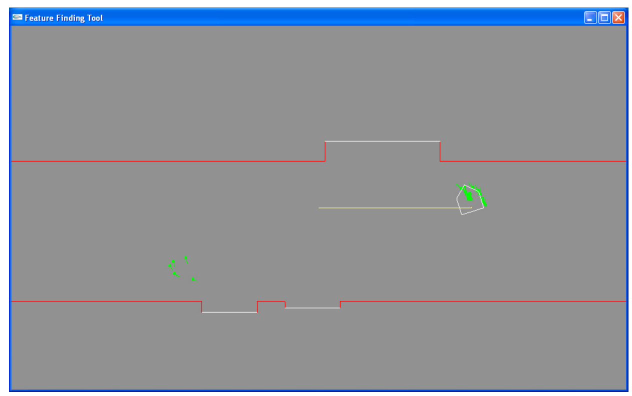

Figure 2 User interface for visualizing the results of MCL. The robot's position hypotheses after an initial wandering leg and adjustment toward the wall are illustrated. Because of the initial ambiguity in pose, several mirror-symmetry hypotheses remain.

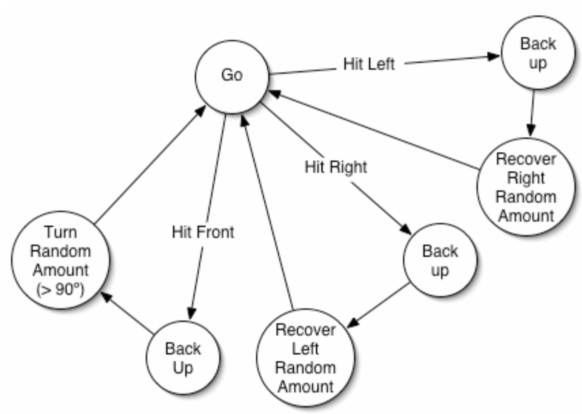

Figure 3 State machine governing robot wandering in the absence of observed scavenger-hunt-related objects

Object Segmentation

Color statistics define distinct regions of interest in the image sequences provided by the cameras. Figure 4, for example, shows the pixels considered part of a red arrow. If an arrow or another object is recognized, the robot breaks out of the wandering behavior to mark the location of the item on the map. It then returns to its wandering behavior, searching for other objects. Additional object-specific processing depends on initial visual cues. For instance, regions thought to be arrows based on color undergo morphological opening and closing to remove small image patches that deviate from the pixel-level definition (Russ 1995). If the largest connected component that results meets some basic criteria, its major axis is extracted. The direction of the arrow is then taken to be the major axis's "heavier" side in pixel count.

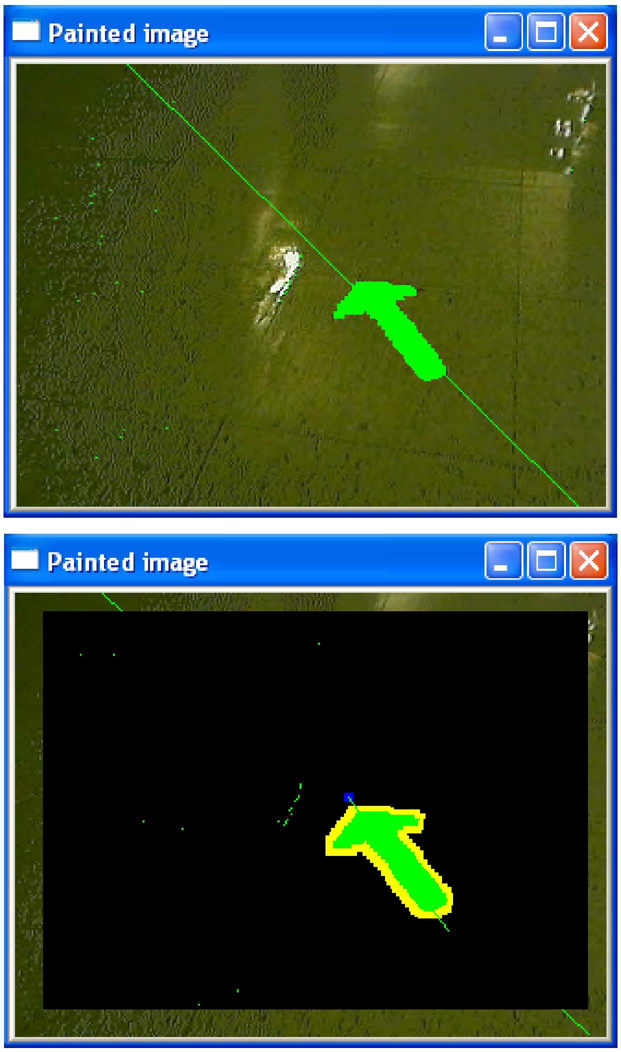

Figure 4 Object identification first uses color thresholds to segment regions of interest (top). Multiple regions are further processed to extract object specific geometry, e.g. an arrow's major axis and direction, shown as the blue dot at the arrow tip.

Arrow Following

The arrow-following routine is currently the most robust of the robot's high-level behaviors. If an arrow is seen, the robot calculates the distance to the arrow and the direction the arrow is pointing. With this information it positions itself on top of the arrow and turns to face the indicated direction. It then returns to a variant of its wandering behavior in search of the next arrow in the sequence. This arrow-chaining strategy has succeeded in directing the robot's search toward scavenger-hunt objects, e.g., around corners and down hallways (Figure 5).

Perspective and Continuing Work

A significant advantage of our robot is its cost effectiveness. It performs multiple tasks with limited actuators and sensors. In its current configuration our robot cost less than $500. Additionally, the processing power of the laptop makes our basic configuration highly extensible. Between this writing and the competition, the team will be further improving the object-recognition and low-level avoidance routines. In particular we intend to add sensors and sensor capabilities, e.g., sonars and a pan/tilt mount for the second camera.

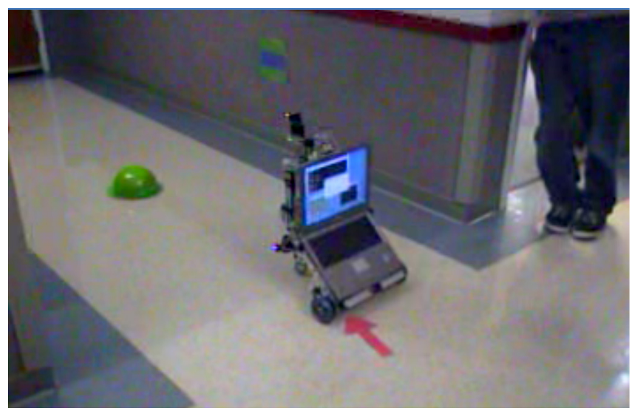

Figure 5 A snapshot taken during successful arrow-following in which a bowl around a corner was located and identified.

Acknowledgments

The authors gratefully acknowledge support from National Science Foundation DUE CCLI #0411176 and funds provided by Harvey Mudd College.

References

S. Thrun, D. Fox, W. Burgard, and F. Dellaert, 2001. Robust Monte Carlo Localization for Mobile Robots, *Artificial Intelligence,* 128(1-2): pp. 99-141.

Russ, John C. *The Image Processing Handbook*, second edition, 1995, CRC Press: Boca Raton, Florida.