Learning Robust Visual Navigation Policies

Jeremy Scott Gerstle

Artificial Intelligence Laboratory
Massachusetts Institue Of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

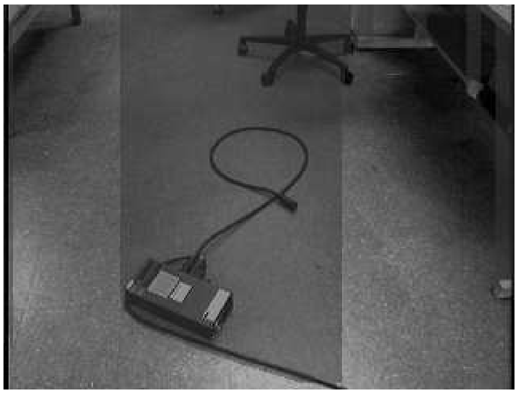

Figure 1: Shaded regions indicate parts of the floor not known to be clear; bottom edge of shaded region indicates occupied parts of the floor.

The Problem: We want to build a framework for learning robust visual navigation policies for robots in an office building environment.

Assuming a simple low-level navigation task like *locate the yellow flower pot and move next to it*, we would like our robot to satisfy its goal by:

1. Looking around a room to acquire its target.
2. Using the initial landscape view to generate a map of free paths or occluded paths which seem to lead toward the goal.
3. Attempt execution of most likely path candidates while keeping map information concurrent with possible environment changes.
4. Formulate or learn better sensory to action mappings based on prior planning using the map.

In order to achieve these sub-goals the robot must incorporate its knowledge of obstacles and approximate metric depth information. In addition, we would like the robot to learn to improve its estimates over time.

Motivation: We want to build robot systems capable of learning about their environment through direct interaction while using vision as the primary environment feedback mechanism. Through this behaviour based approach we intend to demonstrate that low-level navigation routines can be accomplished through simple computations. These routines should make use of the complexity of the environment as it provides a reasonable set of constraints that can help guide the robot agent.

Previous Work: A number of other researchers have experimented with vision-based navigation [3]. Several of these efforts have investigated the construction of occupancy grid maps, a technique which discretizes the environment into a grid mesh and assigns to each grid location a value related to the probability that the location is occupied by an obstacle. See Figure 1. The map can then be used to guide the robot's exploration as it attempts navigation. Nearly all

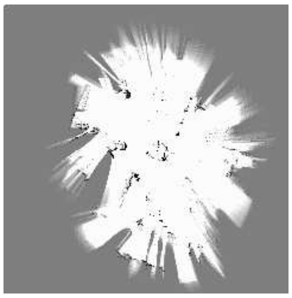

Figure 2: Partial occupancy grid map of our robot lab room in the AI/LCS building.

of these researchers use stereo vision requiring accurately calibrated, complex, and computationally intensive stereo-matching or segmentation algorithms.

Approach: Our approach is to design low level routines that require little computation, allowing for greater computation resources to be spent on higher level task processing. Therefore we have chosen to use a monocular-vision system combined with simple geometric and odometric assumptions for recovery and integration of approximate metric depth information into a 2-D occupancy grid map. This simple processing will lead to less accurate data in each time step, our hope is that slow enough model deterioration and fast updating will allow our robot to maintain enough knowledge about the state of its world to get by just as well. The Figure 2 illustrates an occupancy belief map of the 6x6 meter lab room where our robot Erik resides.

Difficulty: We are faced with three hurdles in our goal of endowing our robot agent with a robust platform for learning low-level navigation policies. Recovery of metric depth information from vision typically requires careful camera calibration. Though general estimation techniques for doing this exist [4], we prefer to find a method which learns the camera's intrinsic parameters directly from the image, perhaps using optical flow. Reducing the reliability of our metric data calculations even further results from the shaft encoders on the robot's wheeled base. As the robot moves around, the values obtained from these encoders, our source of odometric information, tends to drift due to stiction in the drive train. This can lead to actual obstacles becoming smeared and distorted in the robot's occupancy grid map. The difficult question we are then faced with becomes, how do we update the occupancy grid with approximate data sufficient to maintain visual and odometric consistency, and guide navigational search. Using the occupancy grid as a basis for path planning in turn begs the question of how to aggregate routes planned using the map with immediate sensory information to choose goal-oriented actions. This in turn leads to our third and perhaps most difficult question of how to apply reinforcement learning to aid our agent's action selection process, possibly allowing our agent to explore better navigation policies while exploiting efficiencies of the one it currently follows.

Impact: Vision is our most fundamental sensor for acquiring information about our immediate surroundings. Unlike sensors such as sonar or laser range finders, where information arrives in sparse readings that are averaged and integrated over several sensor readings to form statistically meaningful descriptions of the horizon, a single camera image gives us a dense array of meaningful data that we must pick apart and interpret. This separation of vision from more classical robot sensors makes it ideally suited for navigation tasks and ushers the hopes that in the future we might rely on robots to vacuum and clean the office, deliver mail, or chauffer us home at the end of the day. Also, vision is a passive sensor in that it collects information without sending information out. This alone makes the prospect of a visually guided robot attractive in a variety of military applications.

Research Support: This work is supported by DARPA/ITO under contract number DABT 63-99-1-0012.

References:

[1] A. Duchon, W. Warren, L.P. Kaelbling. Ecological Robotics *Adaptive Behavior*, Volume 6, Number 3/4, 1998.

[2] L.P. Kaelbling, M. Littman, A. Cassandra. Planning and Acting in Partially Observeable Stochastic Domains. *Artificial Intelligence*, Vol. 101, 1998.

[3] I. Horswill Visual Collision Avoidance by Segmentation *Proceeding of the 1994 IEEE/RSJ*, May 26, 1994.

[4] L. Faugeras. Three-Dimensional Vision. MIT Press, 1998.