Using a Marker to Map an Unknown Environment

Gregory Dudek*, Michael Jenkin+, Evangelos Milios*, David Wilkes*

*Department of Computer Science, University of Toronto
Toronto, Ontario, Canada
+Department of Computer Science, York University
Downsview, Ontario, Canada
(authors listed alphabetically)

Abstract

A fundamental problem in robotics is the exploration of an unknown environment. Current approaches to exploration make use of a global distance metric that is used to relate past sensory experiences to local measurements. Rather than rely on such an assumption we consider the more general problem of exploration without a distance metric: we propose robot exploration as graph building. We demonstrate that it is not possible for a robot to successfully explore an arbitrary graph without additional aids. By augmenting the robot with a distinct marker which can be put down or picked up at will, we show that a robot can successfully map its environment. An algorithm is developed for map building and is applied to a complex environment.

Keywords: Autonomous Navigation, Robot Exploration, Map, Graph

I. Introduction

Many robotic systems exist in a completely known and well specified environment. This is particularly true for industrial robots such as welders. The parts to be welded are always located in exactly the same place, and very little work is required to ensure that the positioning is exact. More complicated systems allow parts to be located within a small window, but again the environment is well specified and rigidly controlled [20]. The situation is similar in current autonomous vehicle research. In many applications it is assumed that a metric map exists for the environment, and that the position of the robot with respect to the map can be maintained [11]. Such maps are often unavailable in practice, however, and accurately establishing the robot's position within one is very difficult. Leaving aside the problems involved in having a robot accurately locate itself within a map, many environments exist which cannot be well-mapped in advance. In addition, many systems rely on the use of a metric to establish spatial relationships between objects in the environment, the robot itself, and the map. This is not easily accomplished in practice, and a number of approaches, e.g. [5], have been suggested in order to more accurately track the motion of a robot in its environment so as to reduce positional errors.

Rather than have a "canned" map of the environment, another approach is to have the robot explore its own environment, and for the robot to build its own internal representation. As the robot moves it builds up a model of its environment by integrating information obtained at different times with the known motion of the robot. Unfortunately, errors in various sensor readings, coupled with errors in estimating the distance that the robot has moved between measurements, can quickly lead to horrendous errors in the robot's estimation of its location. This highlights a fundamental problem with these approaches: in most environments there is no reliable distance metric that can be extracted by a robot. This problem is manifested by errors in knowing how far the robot has moved, by range data errors, and by errors in several other parameter estimation tasks. In this paper we examine how a robot could go about exploring an environment without relying on a distance metric.

Motivated by the need for spatial representations other than ones based on metric information, [15] proposes a four-level spatial semantic hierarchy. The four levels, starting with the lowest, are the sensorimotor (robot sensations and primitive actions), procedural (robot actions to accomplish place-finding and route-following tasks), topological (places and paths and their topological relations), and metric (places and paths and their metric relations). In this paper, we provide precise definitions of what correspond to the sensorimotor, procedural and topological levels associated with the learning and navigation in a graph-like world, assuming that metric information is neither sensed nor stored.

The problem which we address in this paper is this: Given an unknown environment, formulate a series of plans for the robot so that after carrying out those actions, the robot will have an understanding of its environment sufficient for solving navigation tasks with the use of sensors (i.e. it builds a map). We begin by informally discussing the world in which the robot will exist, the robot itself, and the actions and sensations with which it will be equipped. In a later section we define these terms, and the problem we wish to solve, more formally. To solve the problem of determining when the robot has returned to a previously visited location (the "am I there yet" problem) during map build-