Leading the Way: An Efficient Multi-Robot Guidance System

Piyush Khandelwal and Peter Stone
Department of Computer Science
The University of Texas at Austin
2317 Speedway, Stop D9500
Austin, TX 78712, USA
{piyushk,pstone}@cs.utexas.edu

Abstract

Prior approaches to human guidance using robots inside a building have typically been limited to a single robot guide that navigates a human from start to goal. However, due to their limited mobility, the robot is often unable to keep up with the human's natural speed. In contrast, this paper addresses this difference in mobility between robots and people by presenting an approach that uses multiple robots to guide a human. Our approach uses a compact topological graph representation to formulate the multi-robot guidance problem as a Markov Decision Process (MDP). Using a model of human motion in the presence of guiding robots, we define the transition function for this MDP. We solve the MDP using Value Iteration to obtain an optimal policy for placing robots and evaluate this policy.

Indoor environments such as airports, shopping malls, hospitals, and warehouse stores are characteristically full of people hurrying towards a destination or trying to locate a particular item. Often, they are unfamiliar with the environment and spend a fair amount of time locating these resources. With recent advancements in service robotics, it is becoming far more feasible to deploy a large number of robots to aid humans in these environments. Once multiple robots are pervasive in indoor environments, can we we use these robots effectively to interact with and serve humans?

In this work, we specifically study the problem of deciding where to place robots in an environment to guide a human who is not familiar with that environment. A multi-robot solution to the guidance problem can make use of a human's ease of navigation by proactively placing robots where the human is likely to need help in the future, instead of using a single robot to guide the human from start to goal. Whenever the system needs to guide a human at a specific location, it can commission a nearby robot to direct the human towards the next objective, whether it be another guide robot or the goal.

Such a system also needs to account for the inherent non-determinism in the human's actions, as the human may not interpret a guide robot's intentions as expected and deviate from the path the system has set out for the human to follow.

Consequently, it becomes necessary to have a model of human behavior in the presence of guide robots, and use this model to plan out an optimal path to the goal. To solve this problem, we formulate the multi-robot guidance problem as a Markov Decision Process (MDP), and use a hand-coded model of human motion in the presence of guide robots to define the transition function for this MDP. We then use Value Iteration (Sutton and Barto 1998) to solve this MDP, generating an optimal solution for placing robots. Such a solution can take the uncertainty in a human's movement into account and avoid actions that have a significant probability of failure. Finally, we evaluate the generated policy by comparing it against a heuristic solution for deciding robot placements. Experiments are run using the model of human motion, as well as with avatars controlled by real humans in a simulation environment.

Related Work

Single robot guides have been used over the past two decades to guide humans on a tour (Thrun et al. 1999), or provide navigation assistance to the elderly and visually impaired (Montemerlo et al. 2002; Lacey and Dawson-Howe 1998). In contrast, the goal of our work is to navigate able-bodied humans inside a building by using multiple robots.

While past work has looked into building models of how humans interpret natural language navigation instructions (Chen and Mooney 2011), and how such instructions can be used by robots to get to a human-directed goal (Tellex et al. 2011), little attention has been paid to modeling how humans follow navigation instructions given by robots. In this work, the robot's display directional arrows to guide a person, and we use this simple interface to build a model of human motion in the presence of guiding robots.

Problem Statement and MDP Formulation

We assume that the multi-robot guidance system is given a graph representation of the environment (see Figure 1b), the current location of the human, the locations of all the robots, and the human's goal. The system also has access to a model of human motion in the presence of guiding robots. Given these inputs, the system attempts to minimize the overall distance traveled by the human before reaching the goal. To do so, the system can assign robots to specific locations to guide

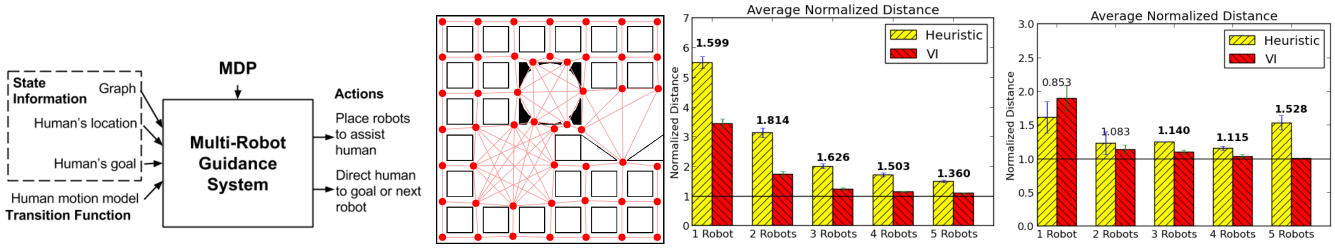

(a) MDP Formulation (b) Graph (c) Simulated Agent Experiment (d) Human Avatar Experiment

Figure 1: Fig. 1a shows how the multi-robot guidance problem is framed as an MDP. Fig. 1b shows the environment and graph used in experiments. The remaining figures show average performance of the heuristic and VI approaches on this map with standard error bounds. The value above each set of bars is bold when the difference between the two is significant using a two sample t-test with 95% confidence.

the person, and if a human walks up to one of these robots, select the direction in which the robot points the human towards. Furthermore, we assume the following:

1. The maximum number of robots that the solution can place is pre-specified.
2. A robot directs a human by displaying an arrow.
3. Whenever a robot is commissioned to guide the human, and is assigned to a specific location, the robot gets there instantaneously.

Given this system specification, we formulate the problem as an MDP $M = \langle S, A_s, P^a_{ss'}, R^a_{ss'} \rangle$, where all the terms have their usual meanings. Locations of the human and the robots are used to construct the state representation S. The human motion model defines the transition function T, The action set A allows the system to place a robot, direct the human or wait for the human to make a move. The distance traveled by the human is used to compute the reward function R. Once the problem has been framed as an MDP, standard solvers such as VI can be used to solve this problem.

Experiments

In this section, we present two experiments that evaluate the effectiveness of the VI-based solution for guiding a human, and compare it against a heuristic solution that attempts to balance shortest path with number of robot placements.

Simulated Human Agent

The first experiment compares the performance of both solutions through an agent exactly following the hand-coded human motion model described in this paper (Fig. 1c). Since the same model was used while generating the VI solution, the policy computed by the VI solution is optimal for this agent. Results were averaged across 1000 trials. The heuristic approach cannot account for human behavior in open spaces correctly, and the VI solution significantly outperforms this heuristic approach.

Human experiments

The second experiment evaluates both the VI and heuristic solution with real humans inside a 3D simulation environment constructed using Gazebo (Koenig and Howard 2004), and based on the graph in Fig. 1b. The simulator allows a human to control an avatar inside the environment, and provides sensory information through a simulated camera. Human position inside the environment is mapped to a graph node to compute the system state inside the MDP, and the simulator queries the policy being evaluated for the action at that state to move robots into their assigned locations as requested. In this experiment, since real humans control what happens in the environment, their behavior may significantly differ from the hand-coded human motion model, and the VI solution is no longer guaranteed to be optimal.

As Fig. 1d shows, the difference between the 2 policies is either insignificant, or the VI solution significantly outperforms the heuristic. Of particular note is the case with 1 robot, where the heuristic performs better than the VI solution (not significantly). The heuristic places its single robot immediately to point the human in the general vicinity of the goal, and the human executes a search strategy in the correct region to find the goal. Since the hand-coded human motion model does not account for this strategy, the VI solution is quite conservative about placing its single robot unless it can place a robot on a location directly visible from the goal, and the human wanders aimlessly for some time. A more realistic human model may be able to improve the VI solution in such cases.

Discussion

In this paper, we have introduced the multi-robot human guidance problem, formulated the problem as an MDP using a model of human motion and solved the MDP to generate an optimal policy for placing multiple robots in an environment to efficiently guide a human. Our MDP formulation uses a topological graph representation of the environment. Our ongoing research agenda includes extending this work by collecting data from real humans interacting with real guide robots to construct realistic models of human motion, as well as accounting for robot navigation time and costs.

References

Chen, D. L., and Mooney, R. J. 2011. Learning to interpret natural language navigation instructions from observations. In *AAAI*.

Koenig, N., and Howard, A. 2004. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *IROS*.

Lacey, G., and Dawson-Howe, K. M. 1998. The application of robotics to a mobility aid for the elderly blind. *Robotics and Autonomous Systems*.

Montemerlo, M.; Pineau, J.; Roy, N.; et al. 2002. Experiences with a mobile robotic guide for the elderly. In *IAAI*.

Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: An introduction*. Cambridge Univ Press.

Tellex, S.; Kollar, T.; Dickerson, S.; Walter, M. R.; Banerjee, A. G.; et al. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *AAAI*.

Thrun, S.; Bennewitz, M.; Burgard, W.; Cremers, A. B.; Dellaert, F.; Fox, D.; et al. 1999. MINERVA: A second-generation museum tour-guide robot. In *ICRA*.