Visually Augmented POMDP for Indoor Robot Navigation

LÓPEZ M.E., BAREA R., BERGASA L.M., ESCUDERO M.S.
Electronics Department
University of Alcalá
Campus Universitario. 28871 Alcalá de Henares (Madrid)
SPAIN

Abstract: - This paper presents a new approach to robustly track a robot's location in indoor environments using a partially observable Markov model. This model is constructed from topological representation of the environment, plus actuator an sensor characteristics. The system takes into account various sources of uncertainty to maintain a probability distribution over all possible locations of the robot. A novel feature of our approach is the integration of visual information to augment the robustness of the system. We show the first results of this approach in localizing an actual mobile robot navigating corridors.

Key-Words: - Robot Localization, Partially Observable Markov Decision Processes, Topological Maps, Natural Landmarks and Artificial Vision.

1   Introduction

While the state of the art in autonomous indoor navigation is fairly advanced [1], there are not good solutions to overcome errors in perception and action that permit the robot to traverse corridors for long periods of time without getting lost. For this reason, robot localization techniques explicitly dealing with perceptual and motor uncertainty have received considerable attention in the last few years [2],[3].

Research in this area can be broadly categorized with respect to the environment representation technique adopted. The two most popular representation schemes are the occupancy grid and the topological map. Both of them can be formulated with a probabilistic state transition network in order to cope with uncertainty in the real sensory data. Into the problem of using metric maps, Thrun et al. [2] introduced this kind of probabilistic formulation with good results. However, this kind of metric representation, in which a state consist of a precise location (x,y) and a heading θ of the robot, is still time consuming. On the other hand, for many tasks in many environments, it is not necessary to know the robot's pose in detail. Given robust low-level routines that can, for example, use local sensors to drive along a corridor, it is only necessary to know that the robot is in some region to allow the navigation task. In such cases, a more coarse-grained uncertainty model may be appropriate, using a topological representation of the connectivity of the environment.

In the framework of probabilistic navigation with topological maps, the Markov foundations for robot localization adopt the form of Partially Observable Markov Decision Processes (POMDP's). The Dervish project at Stanford University [4] and the Xavier project at Carnagie-Mellon University [5] used these kind of navigation strategies for localization and path planning. The major contribution of this paper is to augment the Markov model with visual information of natural landmarks in order to increase state recognizability.

The paper is organized as follows. After a brief overview of POMDP's foundations (section 2), we describe the Markov model construction adding visual information (section 3). In section 4 we resume the steps to perform robot localization using Markov state estimation. Finally, we show some experimental results (section 5), whereas a final conclusion summarizes the paper (section 6).

2   POMDPs Review

Before describing how we construct the Markov model of an indoor environment, we introduce some terminology and foundations.

A Markov Decision Process (MDP) is formally defined as a triple $\{S,A,T\}$, where S is a finite set of states, A is a finite set of actions, and T are transition matrixes that contain the transition probabilities $p(s'|s,a)$ for all $s,s' \in S$ and $a \in A$ (the probability that new state is s' if action a is executed in state s). For each state $s \in S,$ it's possible that only a subset of all actions are feasible.

Additionally, Partially Observable Markov Decision Processes (POMDPs) are used under domains where there is not certainty about the actual state of the

system. Instead, the agent can do observations and use them to compute a probabilistic distribution over all possible states. So, a POMDP includes a finite set of observations O and observation probabilities $p(o|s)$ for all $s \hat{I} S$ and $o \hat{I} O$ (the probability of doing an observation o when the system is in state s). The Markov models assume that the transition and observation probabilities are determined only by the current state of the system (this is the "Markov assumption").

To maintain a belief of the current state of the system in form of a belief distribution $Bel(S)$ over the set of states, the distribution must be updated whenever a new action or perception is carried out [6],[7].

When an action a is executed, the new probabilities become:

$$Bel_{posterior}(S = s') = K \cdot \sum_{s \in S} p(s'|s,a) \cdot Bel_{prior}(s) \qquad (1)$$

where K is a normalization factor to ensure that the probabilities all sum one.

When a sensor report o is received, the probabilities become:

$$Bel_{posterior}(S = s) = K \cdot p(o|s) \cdot Bel_{prior}(s) \qquad (2)$$

In the context of robot navigation, the states of the Markov model are the locations (or nodes) of a topological representation of the environment. Actions are the behaviors that the robot can execute to move from one state to another, and observations are all kind of environment information the robot can extract from its sensors. In this case, the Markov model is partially observable because the robot may never know exactly which state it is in. On the other hand, this approach of localization and navigation using probabilistic and topological representations is attractive because it does not depend on geometric accuracy and is reactive to sensed features of the environment.

3 Markov Model Construction

In the following sections we describe the Markov model used for corridor navigation in indoor environments, and how it can be improved using a camera as additional sensor.

3.1 The states
States of the Markov model are directly related to the environment representation used. Taking into account


Fig. 1. Example of Markov states for a corridor environment.

that the final objective of the navigation system will be to direct the robot from one room to another, we use coarse-grained "regions" of variable size in accordance with the topology of the environment. These regions are small enough to permit the planning task, but not as small as to increase the number of states without adding functionality to the system.

As it can be seen in figure 1, only one state is assigned to each room, while the corridor is discretized into thinner regions. The limits of these regions correspond to any change in lateral features of the corridor (such as a new door, opening or piece of wall). For each one of these regions there are four states, one for each of the four orientations the robot can adopt.

3.2 The actions
The actions selected to produce transitions from one state to another correspond to local navigation behaviors of the robot. We assume imperfect actions, so the effect of an action can be different of the expected one (this is modeled by the transition model T).

The selected actions are:
- "Go out room". This action is defined only in room states. The robot uses the camera to localize the door. Then it gets towards the door and traverses it using vision and sonar information. The robot must get the corridor perpendicularly to the door.
- "Enter room". This action is defined only in corridor states oriented to a door. As in the previous case, vision and sonar are used to perform this task.
- "Turn right". The robot turns 90 degrees to the right using odometry sensors to detect the end of the action.
- "Turn left". The robot turns 90 degrees to the left.
- "Follow corridor". The robot continues through the corridor to the next state.

The last action ("Follow Corridor") is the more complex one. To ensure that the robot only adopts the

four allowed directions without large errors, it's very important that, during the execution of this action, the robot becomes aligned with the corridor longitudinal axis. This is made using sonar buffers to detect the walls and constructing a local model of the corridor. Besides, an individual "Follow Corridor" action terminates when the robot reaches a new state of the corridor. Detecting these transitions only with sonar readings is very difficult when doors are closed. Although the transition model can contemplate this as an imperfect action, the system becomes much more robust if we add visual information to detect state transitions.

3.2.1 Visual landmarks to improve state transition detection

During corridor following, the robot must detect state transitions (end condition of a "Follow Corridor" action) to update the belief distribution after them. When doors are opened these can be easily detected using side sonar readings, but this sensor can't detect transitions to states with closed doors.

To solve this problem, we add visual information from a camera to detect door frames as natural landmarks of state transitions (using color segmentation and some geometrical restrictions). The advantage of this method is that the image processing step is fast and easy, being only necessary to process two lateral windows of the image as it's shown in figure 2. Whenever a vertical transition from wall to door color (or vice versa) is detected in a lateral window, the distance to travel as far as that new state is obtained from the following formula (see figure 3):

$$d = \frac{l}{tg(\alpha)} = K \cdot l \qquad (3)$$

where l is the distance of the robot to the wall in the same side as the detected door frame (obtained from

Fig. 2. Lateral windows to extract door frames and predict inmediate state transitions with "Follow Corridor" action.


Fig.3 Estimation of next state distance from visual and sonar information.

sonar reading) and α is the visual angle of the door frame. As the detected frame is always in the edge of the image, the angle α only depends on the zoom of the camera, that is constant. After covering distance d (measured with relative odometry readings), the robot reaches the new state. This transition can be confirmed with sonar if the door is opened.

Another advantage of this transition detection approach is that no assumptions are made about doors or corridor widths. The only initial knowledge needed by the system is wall and door colors, that can be easily trained. Besides, we only use doors as landmarks, that always are present in indoor environments, so the robot can easily be installed in a new building or house.

3.3 The observations

At each state, the robot can get some observations of the environment using sensor information, and use them to update the belief using observation probabilities $p(o|s)$. Besides sonar observations, we also introduce visual observations to increase the robustness of the system.

3.3.1. Abstract Sonar Observations

In each state, the robot is able to make an "abstract sonar observation". It can perceive, in each of three nominal directions (left, front and right), whether it's "free" or "occupied". An abstract sonar observation is the combination of the percepts in each direction. Thus, there are 8 possible abstract sonar observations, as it's shown in figure 4.


Fig. 4. Abstract Sonar Observations (A.S.O)

3.3.2. Landmark Visual Observations

Abstract sonar observations don't distinguish walls from closed doors, neither openings from opened doors. So, a lot of states in a corridor can produce the same abstract sonar observation.

Again, we introduce visual information to solve this problem and add more informative observations to the system. For this task, we use again doors as landmarks, being the number of doors in the image the "landmark visual observation". This observation is easily obtained with a color segmentation of the image, using the same trained color that in door frame transition detection.

The final probability *p(o|s)* is obtained from the combination of abstract sonar observations and landmark visual observations. The last ones increase the state recognizability, making possible to distinguish states at the beginning of the corridor from states at the end of the same.

4 POMDP Estimation for Robot Localization

The problem of acting in a partially observable environment can be decomposed into two components [8]: a *state estimator*, which takes as input the last belief state, the most recent action and the most recent observation, and returns an updated belief state using equations (1) and (2), and a *policy*, which maps belief states into actions. In robotics context, the first component is *robot localization* and the last one is *task planning*. In this paper we only deal with the robot localization problem, but currently we are already working in planning.

Next, we resume the steps needed to have the robot localization module working in a new indoor environment:

1. To learn door and wall colors using a training setup. The user must click on these objects in several captured images and colors and tolerances are extracted.
2. To build, manually, a topological model of the environment, with nodes corresponding to regions defined in section 3.1. As an example, figure 5 shows the topological model corresponding to the environment shown in figure 2. This step will be eliminated when we develop the automatic environment learning module.
3. To compute the Markov model using the last graph and the following steps: (a) Assignment of
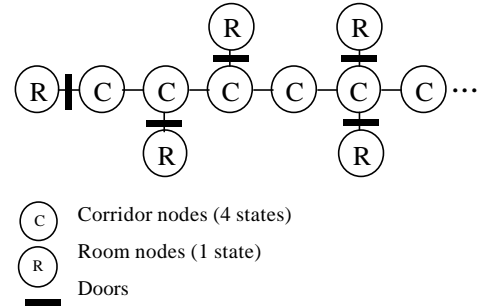

Fig. 5. Topological model of figure 2 environment.

states to the nodes (1 state for room nodes and 4 for corridor nodes); (b) Assignment of possible actions to states; (c) Transition model calculation from the connectivity of the map and an action-error model obtained through informal experimentation; (d) Observation model calculation using a perception-error model also obtained through experimentation.

4. To initialize the belief distribution in one of the two following ways: (a) If initial state of the robot is known, that state is assigned probability 1.0 and the rest 0.0. (b) If initial state is unknown, a uniform distribution is calculated over all states.
5. The movement of the robot is always executed as a sequence of individual actions. Whenever an action is terminated (end of action condition detected), the robot obtains the new observations, and the belief distribution is updated using equations (1) and (2).

5 Experimental Results

In order to verify the behavior of the probabilistic localization module, several experiments have been performed using the commercial robot PeopleBot (ActivMedia Robotics) shown in figure 6.

Fig. 6. The robot used in our experiments.

This robot is built on an Pioneer base and includes bump sensors, encoders, two sonar rings and a color camera on a pan-tilt head. Control, perception and localization are all carried out on an on-board PC.

All the experiments were carried out in a corridor of the Electronics Department which topological map has 71 states, being 11 of them room states. First we made several experiments using only sonar sensor for transition detection and observations. Due to the high symmetry of the environment, the robot took about 20 actions (depending on the actual initial state) to global localize itself when the initial belief was uniform. Adding visual information the number of steps to global localize the robot is reduced to 5. So, visual information eliminates observation symmetry, improving state recognizability.

6 Conclusions and Future Work

We have described the integration of visual information to a probabilistic localization module based on a Partially Observable Markov Decision Process (POMDP). This new sensor provides better information to state transition and observation models, making possible a faster global localization when the initial position of the robot is unknown.

We are extending this work in several directions. We intend to pursue planning an action selection algorithms to efficiently direct robot movements towards a target state (generally a room). Another line of future work is to learn the world model (topology and probabilities) from experience using techniques adapted from hidden Markov models.

References:
[1] Borenstein J., Everett B. and Feng L., "*Navigating Mobile Robots: Systems and Techniques*". A.K. Peters, Ltd. 1996.
[2] Thrun S., "Probabilistic Algorithms in Robotics". *Technical Report CMU-CS-00-126.* 2000.
[3] Cassandra A.R., Kaelbling L.P. and Kurien J.A, "Acting under Uncertainty: Discrete Bayesian Models for Mobile-Robot Navigation". *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems.* 1996.
[4] Nourbakhsh I., Powers R. and Birchfield S., "DERVISH: An Office-Navigaing Robot". *Artificial Intelligence Magazine*, 16(2). 1995.
[5] Simmons R. and Koenig S., "Probabilistic Robot Navigation in Partially Observable Environments", *Proc. of the International Joint Conference on Artificial* Intelligence. 1995.
[6] Asoh H., Motomura Y., Hara I., Akaho S., Hayamizu S. and Matsui T., "Combining Probabilistic Map and Dialog for Robust Life-long Office Navigation". *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'96), 807-812.* 1996.
[7] Tomatis N., Nourbakhsh I. and Siegwart R., "Combining Topological and Metric: a Natural Integration for Simultaneous Localization and Map Building". *Proceedings of the Fourth European Workshop on Advanced Mobile Robots (Eurobot 2001).* 2001.
[8] Zanichelli F., "Topological Maps and Robust Localization for Autonomous Navigation". *IJCAI Workshop on Adaptive spatial representations of dynamic environments.* 1999.