

In Proceedings of PerAc'94 *From Perception to Action*.
Lausanne, CH, September 5-9, 1994

Learning High-Level Navigation Strategies from Sensor Information and Planner Experience ^{*}

Luca Maria Gambardella and Cristina Versino

IDSIA Istituto Dalle Molle di Studi sull'Intelligenza Artificiale

Corso Elvezia 36 - CH - 6900 Lugano

Phone: +41 91 22 88 81 Fax: +41 91 22 89 94

E-mail: luca@idsia.ch - cristina@idsia.ch

Abstract

Moving a robot with shape and size in a cluttered dynamic workspace requires the capability of dealing with obstacles and local minima. The research analyzes situations where no global knowledge about the environment exists, and where the robot can only perceive the space through its local sensors. The system explores a dynamic space using a planner based on local artificial potential fields, and incrementally learns a fast way to escape from dead lock situations using a combination of sensor perceptions and field information. As main result the system learns and uses an high-level description of the workspace consisting of local minimum nodes, backtracking nodes and subgoal nodes.

Keywords

ROBOT MOTION PLANNING - LEARNING - PERCEPTION

1 Introduction

The purpose of this paper is to describe a way to combine motion planning techniques with learning mechanisms in order to move a robot in a dynamic workspace.

The robot moves in a cluttered environment where a lot of deadlock situations can be found. The robot does not have any global knowledge about obstacle configurations and it is able to locally perceive the workspace through external sensors.

We made simulated experiments in house-like environment with doors that can be dynamically opened and closed (Figure 1 a).

The main difficulty to deal with during navigation are deadlock situations that change dynamically. One way to solve the problem of local minimum exploration is to prevent the robot entering the area where a deadlock situation is detected. The goal of this approach is to give the robot the knowledge to foresee a deadlock area from distant perception of the workspace. This result is achieved by creating a relation between robot local perceptions far from the local minimum with actions that allow the robot to avoid the deadlock area [SUTTON92]. As consequence the robot will avoid to explore areas where previous local minima were detected.

This solution is only suitable for static environments where local minima are always in the same locations and usually is very expensive to compute.

^{*}This research is partially supported by Swiss National Fund Project "Integrating Planning and Learning in Robotics" a collaboration with IDSIA and the Institute of Microtechnique of the Ecole Polytechnique the Lausanne.

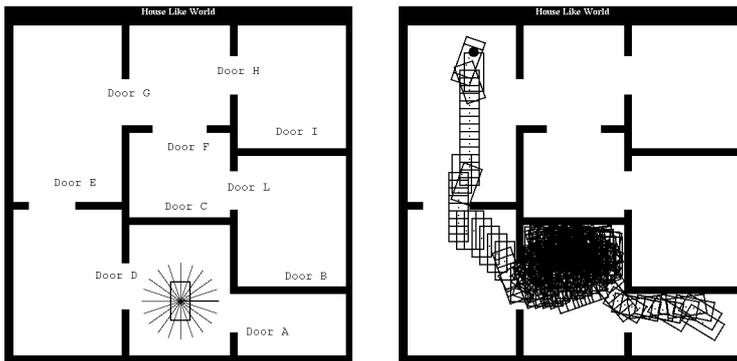


Figure 1. Robot workspace (a) and Robot path using the planner (b)

We therefore believe that local minima must not be avoided a priori. The robot must explore areas where previous local minima have been detected and the most important knowledge it needs is related to the fastest means of escape. To achieve this goal our system incrementally learns an high-level description of the workspace in term of a network that combines robot sensors, field information and planner experience. The network consist of local minimum nodes, backtracking nodes and subgoal nodes.

In the first section we will investigate planner solutions; we will then describe local minima and the high-level network showing some examples.

2 Planning

To move a robot with shape and size in a dynamic cluttered environment we need a planner able to backtrack from deadlock situations that does not model the robot as a point. The existing planner [LATOMBE91] [GAMBARDELLA92] places a set of control points C_j on the robot to model its shape and size. A global artificial field based on a Voronoi Graph is defined over the free space. The decision concerning the next robot motion is driven by a combination of the artificial field in the control points.

The robot makes small changes in all its degrees of freedom and computes the total potential P of all neighbouring configurations S_i .

$$P(S_i) = \sum \rho_j V(C_j) \quad (1)$$

The position with the smallest potential is chosen and the motion is executed if it does not generate a collision or has not been already explored. This navigation process is supported by a backtracking policy and by some special heuristics to overcome local minima.

The main problems with this approach is that a global potential field requires complete knowledge about space and obstacles and that the planner does not take advantage of its experience. In addition, when the environment changes, the potential field becomes invalid and a new field computation must be performed. For these reasons the method cannot be applied to a dynamic environment.

To overcome these limitations we have decided to substitute the global field based on a Voronoi Graph with a field computed considering the distance between the robot control points and the goal. In this way the potential field is computed dynamically and we are able to take advantage of planner navigation strategies avoiding complex computations (Figure 1 b).

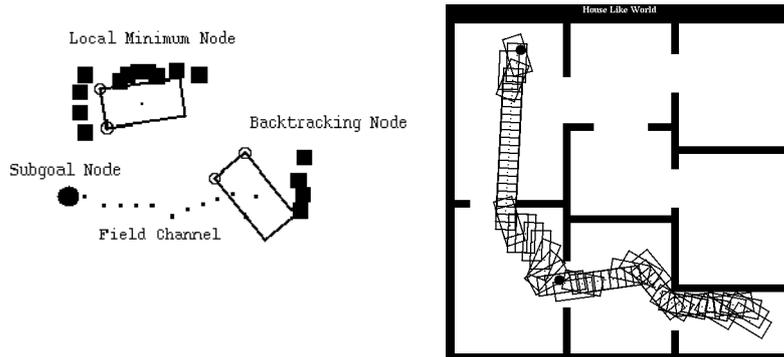


Figure 2. Acquired knowledge (a) and Prudent robot path (b)

The other problem is related to the inability of the planner to deduce new strategies from errors. A robot in the same situation always explores the same positions because it has no knowledge about previous path and deadlock situations.

3 High-Level Network

We propose a solution to manage local minimum by leaning a high-level description of the workspace based on a network that allows the robot to avoid deadlock areas or to escape quickly from them. See [GAMBARDELLA94] for a detailed description. Another solution to this problem is proposed by Kaelbling [KAELBLING93] who uses a high-level description of the space in terms of landmark networks in a stochastic domain. This solution shows an interesting space representation but the problem of landmark definition is not deeply analyzed.

The high-level network is learned during navigation and consists of local minimum nodes, backtracking nodes and subgoal nodes. This knowledge is acquired analyzing the path of the robot and combining sensor perceptions and potential field information (Figure 2 a).

The backtracking node identifies the configuration where the local minimum area starts. The local minimum node is related to the local minimum configuration. The subgoal node is the location where to instantiate a planner subgoal in order to avoid or to escape from the deadlock situation.

The planner, according to its local perceptions, proposes the best actions to reach the goal following the artificial field. When a backtracking node is found, according to its navigation strategy, the planner chooses between generating a subgoal useful for avoiding the local minimum area or continuing exploration with the risk of reaching a deadlock situation. When a local minimum node is detected the planner generates the related subgoal that allows the robot to escape quickly from the deadlock area.

We have defined two different robot navigation strategies. With the *prudent policy* as soon as the robot finds a backtracking configuration it instantiates the subgoal useful for avoiding the deadlock area (Figure 2 b). With the *exploratory policy* the robot ignores backtracking nodes but, when it finds a local minimum node it instantiates the planner subgoal that drives the robot out of the deadlock area (Figure 3 a). We want to point out that all the knowledge we acquire is associated with the robot perceptions and with the robot's relative direction towards the goal. In this way we are able to use the same knowledge in different workspace positions (Figure 3 b).

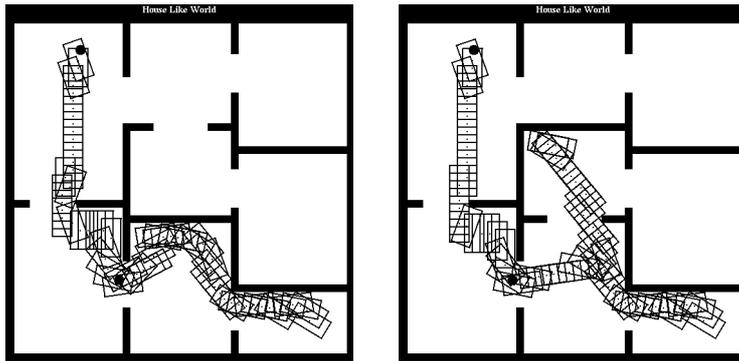


Figure 3. Exploratory robot path (a) and New robot path (b).

4 Conclusions

In this paper we have proposed a solution to the problem of moving a robot in a cluttered workspace where dynamic local minima can be detected. The system takes advantage of planner's ability to move a robot with shape and size.

The robot uses a backtracking policy and a potential field based on goal distance that does not require global computations.

To avoid repetition of the same errors the system uses local robot perceptions, field information and planner experience to model local minima and to learn strategies to escape from them. This knowledge is organized as a high level network that supports planner navigation.

The integrated system shows good performance in moving the simulated robot in the workspace and we are now starting to experiment with a real robot. We are also investigating the possibility of deducing the high-level network using reinforcement learning methods and of introducing more sophisticated methods for generalizing the learning knowledge through clustering techniques.

References

- [GAMBARDELLA92] Luca Maria Gambardella and Marc Haex. Grasps planning for automatic assembly tasks using artificial fields. In *Proc. European Conference on Artificial Intelligence - ECAI '92*, Vienna, Austria, August 3-7 1992.
- [GAMBARDELLA94] Luca Maria Gambardella and Cristina Versino. Robot motion planning integrating planning strategies and learning methods. In *Proceedings of AIPS 94 - The Second International Conference on AI Planning Systems*, Chicago, June 1994.
- [KAEHLING93] L.Pack Kaelbling. Hierarchical learning in stochastic domains: Preliminary results. In *Proceeding of the Tenth International Conference on Machine Learning*, Amherst, Massachusetts, 1993.
- [LATOMBE91] Jean-Claude Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, 1991.
- [SUTTON92] Richard S. Sutton, editor. *Reinforcement Learning*. Kluwer Academic Publishers, 1992.