

Robot Motion Planning Integrating Planning Strategies and Learning Methods

Luca Maria Gambardella and Cristina Versino
IDSIA - Istituto Dalle Molle di Studi sull'Intelligenza Artificiale
Corso Elvezia 36 - 6900 - Lugano - CH
email: luca@idsia.ch - cristina@idsia.ch

Abstract

Robot motion planning in a dynamic cluttered workspace requires the capability of dealing with obstacles and deadlock situations.

The paper analyzes situations where the robot is considered with its shape and size and it can only perceive the space through its local sensors. The robot explores the space using a planner based on an artificial potential field and incrementally learns a fast way to escape or prevent deadlock situations using a combination of sensor perceptions, field information and planner experience.

The knowledge acquired is a high-level network useful for avoiding deadlock areas consisting of local minimum nodes, backtracking nodes and sub-goal nodes.

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 goal is to find a free path from any starting position to any final position in the workspace.

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). This workspace is very complex with respect to robot capability because it requires the ability to move around obstacles, to escape from narrow corridors and to avoid closed doors without any description of the environment. We briefly summarize the robot environment as involving:

- dynamic space
- cluttered workspace
- local minima
- lack of global knowledge
- robot with shape and size

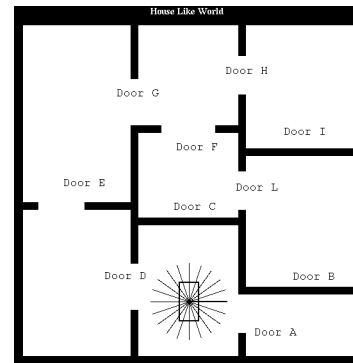


Figure 1: Workspace and robot sensors.

- local sensor perceptions only

The main difficulty to deal with during navigation are deadlock situations that change dynamically. We have investigated different solutions, some based on planning mechanisms and others based on reactive systems supported by learning methods.

In particular we have studied a planning solution based on a global artificial potential field [Latombe 91] [Gambardella and Haex 92] that moves a robot with a given shape and size in a cluttered workspace. The system is based on a best-first search supported by a backtracking policy where robot motion is executed step by step computing a subset of the robot configuration space.

We have investigated solutions based on reactive systems that learn relations between robot positions and robot actions during navigation [Millan 93] [Sutton 92]. These systems show interesting capabilities for avoiding local minima because they learn how to foresee areas where deadlock situations are detected. In these systems the robot starts without any navigation knowledge, and learns how to move in the workspace, acquiring relations between robot perceptions and robot actions. The robot associates perceptions close to the local minimum area with actions that prevent the robot from being attracted there. The final result, achieved after expensive computation, shows some limitations

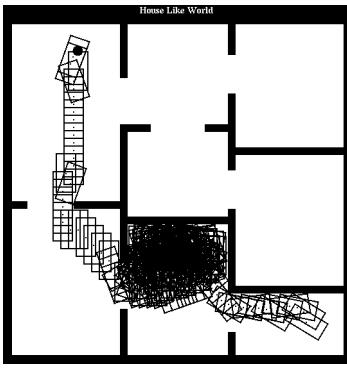


Figure 2: Robot path using the planner.

because the mapping between locations and actions is only valid for static obstacle configurations and predefined goal positions.

Another solution to this problem is proposed by Kaelbling [Kaelbling 93] who uses a high-level description of the space in terms of landmark networks in a stochastic domain. Each landmark is related to a workspace area and the system learns how to navigate between adjacent landmarks minimizing the distance between the actual robot landmark and the goal landmark.

This solution shows an interesting space representation but the problem of landmark definition is not deeply analyzed and the related navigation strategies are still only valid for a static environment.

We propose to overcome the previously mentioned limitations by introducing a planner able to reach a subgoal in a cluttered dynamic workspace moving robots with shape and size. The planner is a modified version of Latombe's planner and is based on an artificial potential field that does not require any global computation because the field is dynamically calculated considering the robot and the goal position.

The real problem during navigation is related to the computation for avoiding local minima for which the planner performs an exhaustive search before escaping from them.

To solve this problem we propose a new planning algorithm that integrates the potential field approach with a high-level network that describes local minimum areas.

The system incrementally learns this high-level network combining input from robot sensors, field information and planner experience. The network consists of local minimum nodes, backtracking nodes and subgoal nodes.

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

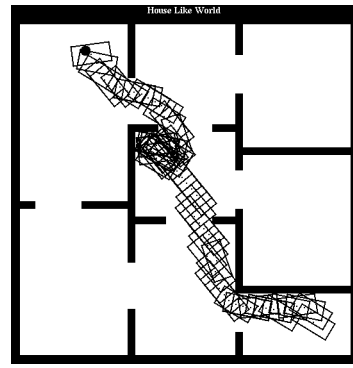


Figure 3: Path opening door C.

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.

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

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 [Latombe 91] [Gambardella and Haex 92] 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

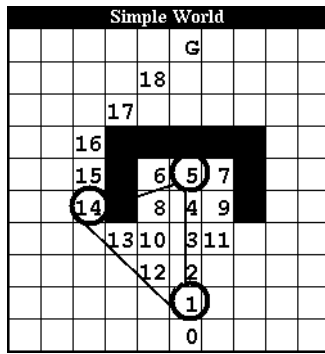


Figure 4: Simple workspace and high-level network.

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 2).

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.

A way to help the system to avoid a complete search is to take advantage of planning experience associating robot perceptions with robot actions through a learning module.

Local Minimum

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 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 [Millan 93] [Sutton 92].

This solution is only suitable for static environments where local minima are always in the same locations. If we open door *C* in Figure 1 (Figure 3), a better path can be achieved because the previous local minimum vanishes. From this example it is clear that in a dynamic environment a local minimum situation cannot be foreseen from a local perception of a distant region. 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.

High-Level Network

For this reason we have built a high-level model of local minimum situations that allows the robot to avoid deadlock areas or to escape quickly from them according to its navigation strategy. See also [Gambardella

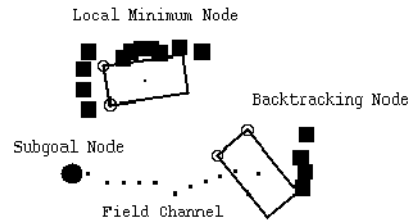


Figure 5: High-level network of figure 2.

and Haex 93] for a different solution.

The model is based on a high-level network consisting of local minimum nodes, backtracking nodes and subgoal nodes. This knowledge is acquired combining input from sensor perceptions, potential field information and planner experiences.

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.

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.

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.

Under some conditions, that we will illustrate later, when the robot detects a local minimum configuration, first backtracks to the backtracking node and after instantiates the goal useful to avoid the deadlock area.

In the next section we explain the mechanism that allows the system to learn the high-level network.

Learning

The system must be able to identify three kinds of node, local minimum nodes, backtracking nodes and subgoal nodes. To accomplish this task the system analyses the path of the robot and learns this information combining robot configurations and robot perceptions.

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 (we have identified eight possible directions between robot configuration and goal locations). In this way we are able to use the same knowledge in different workspace positions.

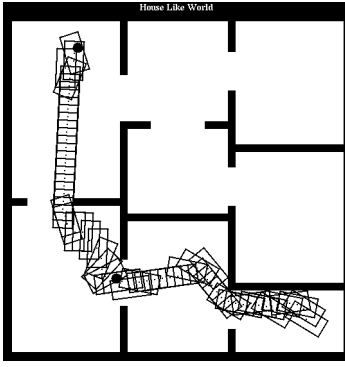


Figure 6: Prudent robot path.

We show our approach using a simplified grid environment with a cell robot having one control point. As we show in the other figures the same solution can be applied to robots with shape and size in any kind of workspace.

In Figure 4 we show the workspace where the starting configuration is labelled with 0 and the goal with G . The planner uses a 8 neighbouring cells for navigation.

Following the goal potential field the robot explores the path $Ep = \{0, 1, \dots, 18, G\}$. The exploration is made using a step by step backtracking mechanism that navigates along the possible robot configurations. For example, starting from location 3 the path of the robot is $\{3, 4, 5, 4, 6, 4, 7, 4, 3, 8, 3, 9, 3\}$ and the robot does not jump directly from location 7 to location 8.

The final best path is $Bp = \{0, 1, 12, 13, 14, \dots, G\}$

From Ep we can identify 5 as the *local minimum node*, because all the neighbour configurations that start from 5 and improve the total potential field are not collision free.

The *backtracking node* is the node labelled with 1 because it is the first backtracking point, after the local minimum, that has a neighbour configuration that belongs to Bp .

We can now identify the *potential field channel* that brings the robot out of the local minimum area. This channel belongs to Bp and starts from 12 and ends in 14. We recognize 14 as the end of the potential field channel because 14 is the first configuration, after the backtracking node, without a possible neighbour that belongs to $(Ep - Bp)$. For example 13 has the possible neighbour 8 that belongs to $(Ep - Bp)$.

At this point we recognize 14 as a *subgoal node*. In other words the subgoal is the first node that belongs to the potential field channel, after which the robot cannot be attracted again in the deadlock area.

The nodes are linked together considering directions and distances between them. This *high-level network* represents the knowledge that describes the local minimum area.

Another piece of information we need to identify

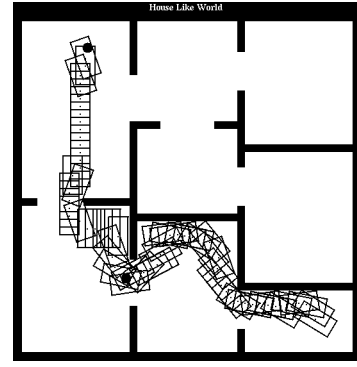


Figure 7: Exploratory robot path.

a subgoal is its survival time. As illustrated before, in a dynamic environment obstacle configurations may change. It is then possible that a given subgoal cannot be reached starting from the backtracking or the local minimum node. We associate with each subgoal a number of planner steps after which the subgoal is removed and the previous goal is activated.

The survival time for the subgoal related to the backtracking node is $\alpha \times \text{potential-field-channel-length}$.

The survival time for the subgoal linked to the local minimum node is $\beta \times (\text{potential-field-channel-length} + \text{steps-between-backtracking-node-and-local-minimum-node})$, with $(1 \leq \alpha, \beta \geq 2)$.

In Figure 5 we show the acquired network related to Figure 2. Figure 6 and Figure 7 illustrate the two paths generated with the prudent policy and the exploratory policy respectively.

The last point we want to illustrate is related to the use of the acquired knowledge in different situations.

The situation in Figure 8 is new and never explored by the planner. In this example the planner finds a robot perception that corresponds to a previous local minimum and a previous backtracking point. In this case the distance between the two nodes is bigger than the previous distance.

The planner, instead of generating the subgoal directly from the local minimum node, backtracks to the backtracking node and instantiates the subgoal useful for avoiding the deadlock area.

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

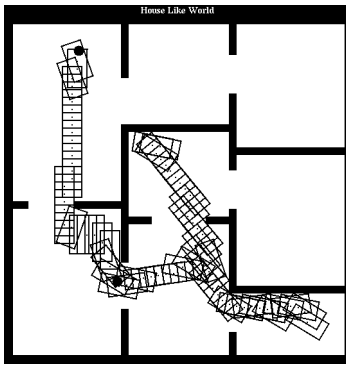


Figure 8: New robot path.

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.

Acknowledgments

This research is partially supported by the Swiss National Research Fund Project "Integrating Planning and Learning in Robotics" a collaborative effort between Istituto Dalle Molle di Studi sull'Intelligenza Artificiale and the Institute of Microtechnique of the Federal Institute of Technology, Lausanne.

References

- [Gambardella and Haex 92] L.M.Gambardella and M.Haex, Grasp Planning for Automatic Assembly Tasks Using Artificial Fields, In *Proceedings ECAI*, August 3-7th 1992, Vienna.
- [Gambardella and Haex 93] L.M.Gambardella and M.Haex, Incorporating Learning in Motion Planning Techniques, In *Proceedings IROS-93, IEEE/RSJ International Conference on Intelligent Robot and Systems*, July 26th-30th, Yokohama, Japan, 1993.
- [Kaelbling 93] L.Pack Kaelbling, Hierarchical Learning in Stochastic Domains: Preliminary Results. In *Proceeding of the Tenth International Conference on Machine Learning*, Amherst, Massachusetts, 1993.
- [Latombe 91] *Robot Motion Planning*, J-C. Latombe, Kluwer Academic Publishers, 1991.

[Millan 93] J.Millan, Teseo:An Autonomous Mobile Robot Controlled By a Neural Network, In *Proceedings of the Second European Workshop on Learning Robots*, Torino, October 25th, 1993.

[Sutton 1992] *Reinforcement Learning*, Editor R.S., Sutton, Kluwer Academic Publishers, 1992.