

Computer Simulation of Learning Experiments with Autonomous Mobile Robots

Ten-min Lee, Ulrich Nehmzow and Roger Hubbard

Department of Computer Science,

University of Manchester

Oxford Road, Manchester M13 9PL, U.K.

leet, ulrich, roger@cs.man.ac.uk

26.3.99

Abstract

This paper presents a series of experiments with an autonomous mobile robot to acquire a model of the robot's interaction with a specific environment. In these experiments, an artificial neural network model of the environment is acquired, using data obtained from the robot interacting with the real world. To investigate the feasibility of this approach we have modelled fundamental sensory perception, robot environment interaction in simple, small environments, and the application of this novel approach to a machine learning application in robotics. The behaviour of a learning mobile robot is predicted using our network-based approach and shown to produce more faithful results than classical robot simulation methods.

1 Introduction

The advantages of numerically modelling robot-environment interaction are numerous and widely reported in the literature ([1, 3, 4, 8]). Simulation allows fast and cheap prediction of robot behaviour. It allows simple modifications of the experimental scenario and experimental parameters, and it allows deterministic repetition of experiments, thus facilitating quantitative analysis and replication. However, to obtain these advantages of simulation, the model used has to be reasonably accurate to exhibit the necessary faithfulness in prediction. Furthermore, there is the real risk that unrealistic simulations may mislead the user, for example by causing him to work on artificial problems. As Brooks and Mataric state, "many simulation artifacts do not occur in the real world" ([1]). Likewise, there is the risk of missing significant features and therefore basing the design of a control structure on false assumptions.

Our previous work, which will be mentioned briefly in this paper, has shown that it is possible to obtain models of robot-environment interaction by using data from a real robot to learn models of robot behaviour, and to use that acquired model subsequently for sensor analysis and robot behaviour prediction.

Both sensor modelling and simulation of hard-wired robot control strategies are relatively straightforward, in that the control program itself is fixed and has no parameters that can be modified. A more sensitive test for the validity of our approach would be to apply our robot-environment interaction model to a learning control program, because here the program is more critically dependent upon a faithful model, as small differences between the real robot's behaviour and the model's prediction are amplified through the robot's learning program.

We have, therefore, applied our approach to predicting the behaviour of a robot that learns - by means of an artificial neural network - to follow a wall and found that the acquired model again shows greater behaviour prediction accuracy than a standard industrial simulation environment.

Proc. TIMR 99 "Towards Intelligent Mobile Robots", Bristol 1999.

Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161.

Report number UMCS-99-3-1. <http://www.cs.man.ac.uk/cstechrep/titles99.html>

2 Modelling Robot-Environment Interaction

Robot-environment interaction is characterised by two main features: modelling the robot's actuators, and modelling the robot's sensors. Initial experiments have shown that the fundamental discrepancies between prediction and actually observed robot behaviour are caused by unreliable predictions of the robot's sensors ([5]). We have, therefore, focussed our experiments on modelling robot sensors and using those models subsequently for the prediction of robot behaviour in the target environment.

The most common approach to modelling robot perception is by constructing simplified numerical models of the robot's sensors ([2, 3, 8, 10, 11]). These models attempt to capture the fundamental physical phenomenon underlying the sensor in question, but typically fail to take into account any effects caused by abnormalities in the environment or abnormal properties of the sensor, such as for instance the dependency of sonar sensors on surface structure and incident angle, or the dependency of infrared sensors upon surface colour. These models are usually too generalised to be applicable to any but the most fundamental predictions of robot-environment interaction.

3 Method

Instead of using simplified models, which are pre-installed by a human designer, we use sensor readings obtained from a real robot operating in the target environment. These perceptions are then used to train a set of multi-layer perceptrons, which learn the model that underlies the relationship between the robot's position and its sensory perception at that position. This method is described in detail in ([5, 6]).

An alternative method would be to use look-up tables to store sample data and to retrieve them in a grid-formatted space ([7]). We have chosen not to use this method because the fact that the robot learns the model, allows it to interpolate between sensor readings, and make predictions even for positions which it has never visited.

For the experiments presented in this paper, we have used a set of 16 multi-layer perceptrons, each of which was to model one of our robot's 16 sonar sensors. Each of the 16 networks uses 2 input units, signifying the robot's x and y position in Cartesian space: 40 units on the first hidden layer, 10 units on the second hidden layer. Each network has one output unit specifying the range reading obtained for its corresponding sensor at any given (x,y) position.

We compare the results of our "network simulator" with those obtained from the Nomadic Technologies simulator, provided by the manufacturer of our mobile robot.

3.1 The Mobile Robot and Simulator

All the experiments presented in this paper were performed using a Nomad 200 autonomous mobile robot called *Forty Two* (see the left picture of Figure 1). This robot is equipped with 16 sonar sensors, 16 short-range infrared sensors, 20 tactile sensors and odometry sensors.

Nomadic Technologies supplied a simulator of the Nomad 200 robot (see the right picture of Figure 1), which was used for comparison with the network simulator developed in our research.

The Nomad simulator provides a graphical environment, allowing users to run the simulated robot in a user-defined world. The sensor modelling in Nomad simulator is a conventional simplified mathematical model. Sensory readings are calculated by pre-installed mathematical equations. All objects modelled in this simulator have been simplified to present the same perceptual qualities. Throughout this paper we compare the results produced by our simulator (described as 'network simulator' in this paper) with those of the simulator provided by Nomadic Technologies with the Nomad 200 robot (described as 'Nomad simulator' in this paper).

4 Experiments and Results

4.1 Modelling a Single Wall

In an initial experiment, whose details are given more explicitly in ([5]), we have simulated our robot's perception of a 100-inch long section of a straight wall. This wall was covered with 2 different materials, resulting in a non-uniform sensory perception of the wall.




Figure 1: *Left*: The Manchester mobile robot *Forty Two*. The robot's sonar and infrared sensors are mounted on the turret. In all experiments the turret was kept fixed to provide a "compass sense". Exploration was carried out using the separate translational and rotational motors located in the base of the robot *right*: The Nomad simulator.

The robot's perception of this wall is shown in Figure 2. The network simulator is able to predict a sudden change in perception at the point where the surface structure of the wall changes.

In order to assess the network simulator's ability to predict robot behaviour, we executed a hard-wired wall-following program on the Nomad 200. Due to the irregularities of the wall being modelled, the physical robot actually collides with the wall after about 40 inches of travel (see Figure 2) which was correctly predicted by the network simulator, whereas the Nomad simulator was unable to predict this behaviour, because it is not able to take into account different surface properties.

4.2 Modelling a small-sized room without obstacles

In a second experiment, we then applied our network model to a larger environment with walls of about 3 metres in length. The walls of this environment were covered in 4 different materials: cardboard boxes, bricks, rough wood and smooth wood. Figure 3 shows the sensor reading from one of the 16 sonar sensors obtained along the indicated trajectory.

At position (400,-400) a sudden rise in range reading is observed due to specular reflection. This is correctly predicted by the network simulator whereas, as before, the Nomad simulator is unable to predict this observation due to its inability to model surface structure.

Again, we were interested in observing whether the network simulator was able to predict robot environment interaction correctly. A hard-wired "find free-space" program was used for this purpose. The actual robot trajectory as well as the trajectory predicted by the Nomad simulator and the trajectory predicted by our network simulator, are shown in Figure 3.

As can be seen, the actual robot does not move to the physical centre of the environment, but rather, due to specular reflection, goes towards the edge of the environment and eventually collides with a smooth wall on the right hand side. This behaviour is correctly predicted by our network simulator, whilst the Nomad simulator is unable to predict this phenomenon.

4.3 Modelling the interaction between a learning robot and its environment

The experiments reported so far have used hard-wired control programs, namely wall-following and space-finding to demonstrate that our approach to robot simulation is able to predict the perception and behaviour of a mobile robot interacting with a specific environment. Our experimental results




Figure 2: *top*: The network model successfully predicts the sudden changes in sonar readings between sample 20 and 40, which the Nomad model fails to do. *Bottom*: By assuming a homogeneous wall surface, the Nomad simulator fails to predict the collision of the real robot, caused by specular reflections. The network simulator correctly predicts the collision.




Figure 3: *Top*: Sonar reading predictions. The Nomad simulator is unable to predict the sudden increase in sonar reading due to specular reflection, while the network simulator predicts this correctly. *Bottom*: Simulated and real robot trajectories in response to the “find free-est space” program. The real and network-simulated robot bump into a smooth wall while the Nomad-simulated robot goes to the geometrical centre of the robot’s environment at (450,-450).

demonstrate that perceptual features of the environment can be captured with sufficient detail to allow the prediction of even unusual and unexpected robot behaviour, such as collisions.

A more sensitive evaluation of our approach would be to make predictions of a learning robot's behaviour in its environment, because any discrepancies between the sensor model and actual sensory perception would be amplified through the robot's learning program, and expose disagreements between model and actual observation more clearly. For this reason, we have conducted a further experiment this time using a neural network-based learning program to control our mobile robot. In this experiment, *Forty-Two* was controlled by a learning controller, using an associative-memory with 16 inputs, which learnt a wall-following behaviour through trial and error. More details of the learning controller are given in ([9]). The important point for the purpose of this paper, is that the sensor motor mapping controlling the robot is acquired in this case, not hard wired, and can be changed at any point of the experiment, in an unsupervised way by the robot itself.

The experiment was conducted in the same room as the experiment mentioned in Section 4.2.

4.4 Experimental procedure and experimental results

The experiment was conducted as follows. The simulated robot – simulated both by the Nomad simulator and the network simulator – was left to acquire the necessary weights to achieve a wall-following behaviour. These weights were then downloaded into the physical Nomad 200 robot and the robot was placed in the target environment, executing the control structure as it was encoded by the downloaded weights, but not modifying these weights any more. The resulting trajectory of the physical Nomad 200 robot was then graphically compared with the trajectory predicted by the respective simulator. These trajectories are shown in Figure 4.

It can be seen that the physical robot's trajectory differed considerably from that predicted by the Nomad simulator, whilst the discrepancies between the predictions of our network simulator with the actual robot trace were far less pronounced. This experiment, therefore, serves as a first indication that the network simulator is faithful enough to allow simulation of learning applications.

5 Summary

The advantages of numerical models of robot-environment interaction are widely appreciated in the literature. Robot simulations facilitate rapid prototyping, replication of experiments, fast modification of parameters, and principled investigation of the influence of individual experimental parameters.

However, over-simplified models can lead to two types of error. First, they may result in modelling phenomena which do not actually exist on real robots and secondly, they may over-generalise certain aspects of the environment and therefore miss necessary detail. Both errors are serious impediments for using robot simulators for serious applications.

Instead of using simplified models of the underlying sensor characteristics, we have investigated the use of artificial neural networks for learning a model of robot-environment interaction for a specific robot interacting with a specific target environment.

In our approach, sensory data obtained from a physical mobile robot is used to train a set of 16 artificial neural networks to acquire a model of sensory perception in the environment. This acquired model is subsequently used to predict the robot's behaviour in the target environment.

We have shown that using our approach, it is possible to predict the behaviour of the robot in different environments running different hard-wired control programs. However, it is of specific interest to investigate whether the network model is faithful enough to allow even the investigation of learning control programs. For this purpose, we have conducted experiments in which a Nomad 200 mobile robot was controlled by a learning controller. Our experimental results show that even for this more demanding application, the network controller provides accurate predictions and can be used to model robot-environment interaction faithfully.

6 Conclusions and future work

The environments in which we have tested our approach so far have been very modest. Rooms of sizes of up to 5 x 3 metres relatively uncluttered by furniture were used. To model a robot's interaction with





Figure 4: *Top*: The robot traces predicted by the Nomad simulator and the real robot using the weights generated by the Nomad simulator. *Bottom*: The network predicted robot trace is similar to the real one using the weights generated by the network simulator.

such an environment, even though it is a relatively structured environment, is a hard task. Sonar sensors – the sensors being modelled here – are very dependent upon surface structure and suffer from specular reflections off smooth surfaces, such as polished wood, glass or metal. These effects are widespread and yet difficult to model.

However, if robot simulation is ever to be used in realistic applications, it is necessary to increase the size and complexity of the environment as well as to investigate this approach with more complex robot control programs. From an experimental point of view, it is also necessary that these control programs are obtained from third parties to ensure that control program and simulation environment are truly independent from one another. Both of these aspects are currently under investigation.

References

- [1] R. A. Brooks, M. J. Mataric, Real Robots, Real Learning Problems. In Chapter 8, Robot Learning, Edited by J. H. Connell and S. Mahadevan, Kluwer Academic Publishers, 1993.
- [2] C. Chen, M. M. Trivedi, C. R. Bidlack. Simulation and Animation of Sensor-Driven Robots. In IEEE Transaction on Robotics and Automation, Vol. 10, No. 5, 1994.
- [3] H. R. Everett. Sensors for Mobile robots: Theory and Application. Published by A K Peters, Ltd., 1995.
- [4] K. D. Harris. Closed-Form Probabilistic Models of Time-of-Flight Sonar. In Proc. ESM98, the 12th European Simulation Multiconference, 1998.
- [5] T. Lee, U. Nehmzow, R. Hubbard. Faithful Simulation of Autonomous Mobile Robots by Means of Model Acquisition. Technical Report Series, UMCS-97-12-1, Computer Science Dept., Manchester Univ., ISSN 1361-6161, 1997.
- [6] T. Lee, U. Nehmzow, R. Hubbard. Mobile Robot Simulation by Means of Acquired Neural Network Models. In Proc. ESM98, the 12th European Simulation Multiconference, 1998.
- [7] H. H. Lund and O. Miglino. From Simulated to Real Robots. In Proc. IEEE 3rd Int. Conf. on Evolutionary Computation. IEEE Press, 1996.
- [8] M. M. Trivedi, M. A. Abidi, R. O. Eason, R. C. Gonzalez. Developing Robotic Systems with Multiple Sensors. IEEE Transactions on Systems, Man, and Cybernetics. Vol. 20, No. 6, 1990.
- [9] U. Nehmzow. Autonomous Acquisition of Sensor-Motor Couplings in Robots. Technical Report Series, UMCS-94-11-1, the Department of Computer Science, Manchester Univ., ISSN 1361-6161, 1994.
- [10] M. Pasquier, K. Lim. A Novel Approach to Range Sensor Modelling in Robot Simulation. In Proc. ESM98, the 12th European Simulation Multiconference, 1998.
- [11] K. Song, W. Tang. Environment Perception for a Mobile Robot Using Double Ultrasonic Sensors and a CCD Camera. IEEE Transactions on Industrial Electronics. Vol. 43, No. 3, 1996.