Continuous Operation and Perpetual Learning in Mobile Robots*

Ulrich Nehmzow
University of Manchester
Department of Computer Science
Manchester M13 9PL
United Kingdom
ulrich@cs.man.ac.uk

Abstract

We argue that continuous operation of mobile robots over very long periods of time is a) desirable from an applications point of view, and b) a demanding scientific challenge, especially in robot learning tasks, because continuous operation presents the robot with very different perceptual stimuli than learning in a condensed period of time.

The paper discusses the hardware and software components of Manchester's continuously operating mobile robot *Strange*, and proposes a robot learning scenario suitable for continuous operation.

1 Introduction

1.1 Motivation

We are interested in the investigation of continuous operation and perpetual learning in mobile robots for two main reasons:

1. Many mobile robot applications, for instance monitoring, surveillance and inspection tasks, require continuous, uninterrupted operation.
2. "Condensed" learning scenarios over short periods of time differ so drastically from continuous-learning ones, that results obtained in condensed training schemes cannot directly be applied to long-term learning applications.

To date, investigations of the application of machine learning to robot control (robot learning) typically focus on relatively short periods of time, in which the robot is presented with stimuli that are relevant to achieve one specific target behaviour. For instance, typical learning scenarios for achieving an obstacle avoidance competence involve the repeated presentation of an obstacle to the robot over a *condensed* period of time, in which the robot learns to associate those "relevant" sensory perceptions with the appropriate motor response. Examples for such robot learning scenarios are [Nehmzow, 1999, Colombetti and Dorigo, 1993, Nehmzow, 1998, Marjanovic et al., 1996] [Gaudiano et al., 1996, Rao and Fuentes, 1996], but there are many more examples.

There are two major differences between such "condensed" learning scenarios and long-term ones. First, the presentation of predominantly "meaningful"[1]

*International Workshop on Recent Advances in Mobile Robots, Leicester, UK, June 2000.

[1] "Meaningful" in relation to the task that is to be acquired, for example sonar range data in front of the robot for the acquisition of an obstacle avoidance behaviour.

sensory perceptions focuses the learning, and makes it easier. In long-term learning tasks, a much larger amount of “meaningless” (noisy) data has to be dealt with. Secondly, “condensed” learning scenarios typically focus on the acquisition of *one* specific task, whereas in perpetual learning over long periods of time, a range of different tasks has to be acquired, and some kind of staging mechanism is required to control the learning process.

It is therefore probable that learning tasks that are accomplished successfully under a training regime that comprises the *condensed* presentation of relevant stimulus-response pairs, will not succeed under a regime where relevant *and irrelevant* (with respect to some specific task) stimulus-response pairs are presented. The way to find out, of course, is to conduct long-term robot learning experiments, which is why we were interested to build a mobile robot that could operate over extended periods of time.

A third motivation for the construction of such a robot is that the scope of possible robotics applications would be extended. Tasks such as surveillance or monitoring usually require continuous operation, or at least operation over periods of time that exceed a robot’s battery lifespan.

We therefore constructed a mobile robot that can operate continuously, monitor the charge level of its batteries, and find the way to its charging station and dock there without human supervision. This robot, *Strange*[2], is shown in figure 1.

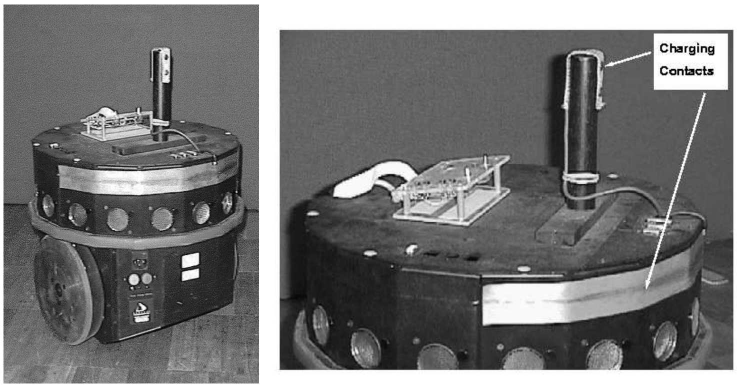

Figure 1: The continuously operating mobile robot *Strange*. Contacts for autonomous recharging are visible at the top and the front of the robot, and shown enlarged in the right image.

Strange is a Nomad Scout robot. It is completely autonomous, carries its own on-board power supply and control computer, and is equipped with sixteen ultrasonic range finders, tactile sensors and a flux gate compass. The custom-built controller controlling *Strange* is discussed in section 2.1.

2 Hardware of the Robot

2.1 Controller

Besides the ability to connect to a charging station, for continuous operation it is especially important to achieve long operating times between charges. Personal computer (PC) controllers with their high power consumption limit the operation time considerably. The power consumption of the controller is usually the determining factor in establishing the operating time of a robot.

[2] *Charm* is a second Manchester robot, while *Up* and *Down* have not yet been built.

We were therefore interested to use a controller with extremely low power consumption, and used a controller based on the asynchronous AMULET2e controller ([Furber et al., 1999]). Compared to a power consumption of about 20 Watt of a laptop, this controller has a power consumption as low as 100 mW when idle.

The Nomad Scout uses two 12 V lead acid batteries with a total of 36 Ah capacity. In practice, with motor currents being between 300 mA and 500 mA, and a power consumption of the robot's motor controller of, say, 250 mW, this means an operating time of perhaps 15 hours using a laptop powered by the robot's batteries, and over 70 hours using the AMULET2e controller.

The controller architecture of the AMULET2e controller is shown in figure 2.

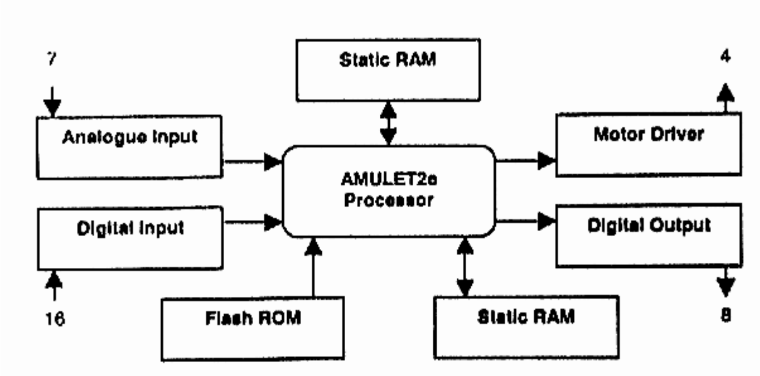

Figure 2: Systems Architecture of the AMULET2e controller.

2.2 The Docking Station

The robot's docking station consists of a copper-plated, hinged roof and a pair of spring-loaded contacts mounted at the wall that provide the second contact. This is shown in figure 3.

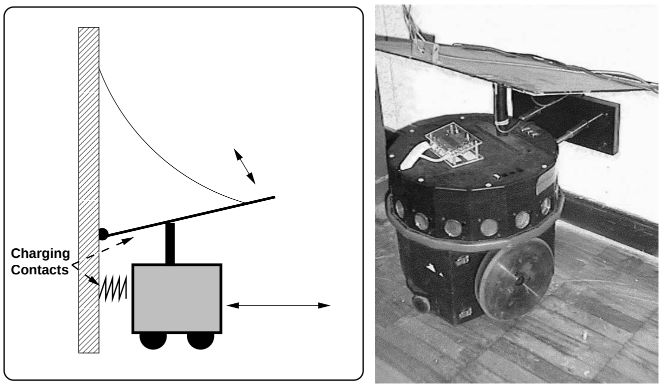

Figure 3: The docking station of the robot.

There are two charging contacts on the robot, one being a contact pole on the top of the robot, the other being a conductive surface mounted at the front of the robot. As the robot enters the charging station, the pole pushes up the roof of the charging station, while the conductive surface eventually makes contact with the spring-loaded contacts on the wall. The presence of a voltage at the charging contacts is detected by the robot's controller, which monitors battery charge and controls the charging process.

3 Control Software of the Robot

Strange operates fully autonomously, i.e. without links to external power, control or computation resources. The on-board controller therefore implements all aspects of the robot's behaviour, in particular safe operation (obstacle avoidance), navigation and docking.

3.1 Obstacle Avoidance

The obstacle avoidance behaviour of the robot is implemented by determining an overall "force" vector $\vec{f}$ (equation 1), given as the vector summation of virtual "forces" $\vec{f_i}$ resulting from the range readings r_i of the sixteen sonar sensors (equation 2).

$$\vec{f} = \sum_i \vec{f_i}, \qquad i = 1 \dots 16; \tag{1}$$

$$||\vec{f_i}|| = \frac{1}{r_i}. \tag{2}$$

Directions of the vectors $\vec{f_i}$ are taken from the respective sensor's orientation with respect to the robot's front.

The steering direction taken by the robot is then chosen to be $-\vec{f}$, the turning speed proportional to $||\vec{f}||$, and the translational speed inversely proportional to $||\vec{f}||$. The resultant obstacle avoidance behaviour is smooth and "measured", as turning speeds are proportional to proximity to obstacles, travelling speed is proportional to distance from obstacles, and the turning directions are continuous over the entire possible range of 360°.

3.2 Navigation, Homing and Docking

3.2.1 Map Building, Localisation and Navigation

During exploration, the robot builds a topological map using a self-organising feature map, or Kohonen network ([Kohonen, 1988]) as shown in figure 4.

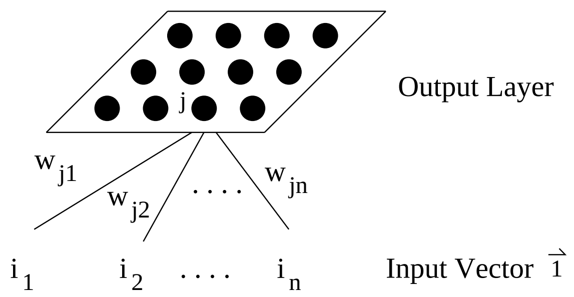

Figure 4: STRUCTURE OF THE SOFM.

This network is presented with raw sensory information obtained from the robot's sixteen sonar sensors, and clusters this data topologically through a process of self organisation. This is described in detail in [Owen and Nehmzow, 1996, Nehmzow, 1999].

After the process of map building is completed — this typically takes a few hundred iterations through the training algorithm — the network's response facilitates robot self-localisation. Every sensory perception clustered by the network represents a location in the robot's perceptual space, which is synonymous with a location in physical space as long as there is no perceptual aliasing (identical sensory perceptions at different physical locations). For a detailed discussion of this method of self-localisation, see [Nehmzow et al., 1991].

Finally, to achieve point-to-point navigation, excitation patterns of the self-organising feature map (i.e. locations) are associated with motor commands in a user-supervised training process as illustrated in figure 5.

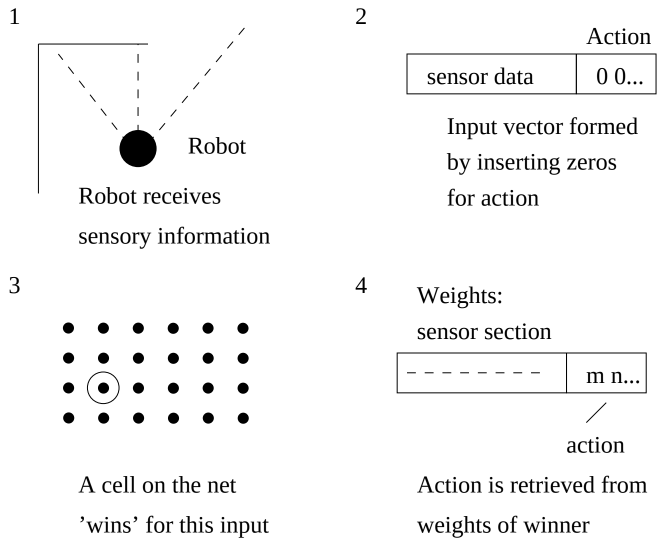

Figure 5: Action retrieval sequence.

During user-supervised route training, the network is presented with an input vector containing the sixteen sonar readings, as well as the required motor action as indicated by the user. This means that the network stores perception-action pairs in the weight vectors of the network's units. On retrieval, i.e. autonomous following of the trained route, the required motor action is determined by presenting the network with an input vector containing the sixteen sonar readings and blank (zero) entries in the action component of the input vector. The winning unit's weight vector then contains the required motor action (figure 5). This route learning algorithm is described in detail in [Owen and Nehmzow, 1996] and [Nehmzow, 1999]. In *Strange's* case, this route learning mechanism enables the robot to learn routes that take it to the immediate vicinity of the charging station.

3.2.2 Docking

The final connection with the charging station is achieved by a two-stage process that is initiated once the robot is placed in the vicinity of the charging station, using the navigation method described in section 3.2.1.

First, a modified force-field obstacle avoidance behaviour is used, in which an "attractive force" in the forward direction is added to the "repulsive forces" exerted by the robot's sonar sensors.

The strength of this attractive "force" was determined through experimentation, and chosen such that the robot has just entered the outer edge of the charging station when all "forces" are in equilibrium.

At this point, a final approach is initiated, which simply consists of a slow forward motion until voltage is detected across the robot's charging terminals.

4 A Proposal for Perpetual Robot Learning

This section presents a proposal about a perpetual robot learning scenario, in which a number of different competences are acquired over a long period of time, in the presence of noisy and "meaningless" sensory perceptions. This learning scenario is the subject of ongoing research at Manchester.

4.1 Stages of Learning

We see at least three distinct stages in *Strange's* learning process, which are initiated at different times of the robot's operation:

1. Leaving and entering the docking station from the "immediate entry area" just outside the charging station. This will be achieved by associating sensory perceptions in the immediate vicinity of the docking station with motor actions that take the robot closer into the station, using a Radial Basis Function network.

2. Learning of fundamental sensor-motor competences such as obstacle avoidance and contour following. This will be achieved by a self-supervised learning mechanisms as described in [Nehmzow, 1999], where sensor-motor pairs are associated in an associative memory by means of trial and error, guided by "instinct rules".

3. Learning to approach the "immediate entry area" from ever increasing distances. This will be achieved by employing the route learning mechanism discussed in section 3.2.1.

4.2 Staging of the Perpetual Learning Process

If perpetual learning throughout the entire operation time of the robot is to be achieved, different learning stages (i.e. the acquisition of different kinds of competences) will have to be activated at different times. For example, the robot will first have to learn to dock at the charging station, then to move competently without collisions, then to navigate, and so on.

To activate the different learning stages at the appropriate time, we propose to control the robot's behaviour and the initiation of the different learning stages by one central variable t, which represents the current duration of operation (the robot's "age"). The activity $a(t)$ of each learning task, such as docking, obstacle avoidance, or navigation, will be governed by two variables λ_1 and λ_2, which are set experimentally for each task, and the global variable t, as indicated in equation 3.

$$a(t) = e^{-\lambda_1 t} + e^{\lambda_2 t}, \tag{3}$$

where $a(t)$ is the activation of a particular learning task at time t, and λ_1 and λ_2 are parameters set individually for each learning component, determined experimentally. Figure 6 shows three example plots for different λ_1 and λ_2. The top curve ($\lambda_1 = 0$, $\lambda_2 = 0.15$) would be suitable for a learning task that becomes more active towards later stages of the robot's operation, the bottom curve ($\lambda_1 = 0.5$, $\lambda_2 = 0$) would be suitable for a learning task that is active only in the early stages of operation, and the middle curve ($\lambda_1 = 0.5$, $\lambda_2 = 0.09$) shows high initial activation which decreases first, but then increases again later.

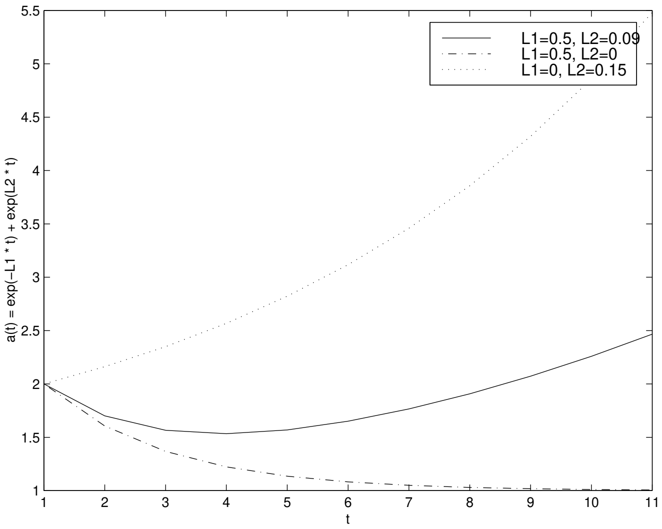

Figure 6: THE ACTIVATION FUNCTION $a(t)$.

5 Summary and Conclusion

The argument for uninterrupted, continuous robot operation is threefold: first, applications such as monitoring or surveillance require continuous operation, secondly, continuous operation and perpetual learning presents a new challenge and a new approach to machine learning applications, because it does not offer the focused time window containing "meaningful" events only that is typical for current machine learning applications in mobile robotics, and thirdly, perpetual learning aims at learning *several* competences over a period of time, rather than one specific competence, as is usually the case in "condensed" learning scenarios.

This paper presents the hardware and software design of the Manchester mobile robot *Strange*. *Strange* is a Nomad Scout, equipped with sixteen sonar sensors, tactile sensors, a flux gate compass and odometry. Furthermore, the robot is equipped with charging contacts that allow fully autonomous docking and charging. The robot's controller is based on an asynchronous AMULET2e microprocessor with extremely low power consumption.

The control architecture of *Strange* consists of two main components. First, a "force" field approach is used to accomplish obstacle avoidance. The advantages of this approach, compared with rule-based methods of obstacle avoidance, is that the resulting motion is smooth and in correspondence to obstacle proximity, and that correct obstacle avoidance is attained in "dense" as well as in "sparse" environments.

Second, a navigation component ensures that the robot is able to return to the vicinity of the charging station, where the robot's docking behaviour is initiated, should docking be required.

Map building, self localisation and path planning are achieved by using a self-organising feature map that clusters the robot's perceptual space. Network responses identify landmarks in perceptual space, which are correlated with physical locations in the world as long as there is no or very little perceptual aliasing.

First experiments show the robot to operate reliably and successfully. Current research addresses the question of operation over extended periods of time, and the application to long-term learning tasks of multiple competences, using a staged learning process governed by one central variable representing the robot's duration of operation.

Acknowledgements

The initial experimental work presented in this paper was conducted by Roy Henderson as his MSc project at the University of Manchester ([Henderson, 1999]). His contribution is gratefully acknowledged.

References

[Colombetti and Dorigo, 1993] Colombetti, M. and Dorigo, M. (1993). Robot shaping: Developing situated agents through learning. Technical Report 40, International Computer Science Institute, Berkeley CA.

[Furber et al., 1999] Furber, S., Garside, J., Riocreux, P., Temple, S., Day, P., Liu, J., and Paver, N. (1999). Amulet2e: An asynchronous embedded controller. *Proceedings of the IEEE*, 87(2):243–256. ISSN 0018-9219.

[Gaudiano et al., 1996] Gaudiano, P., Zalama, E., Chang, C., and Coronado, J. L. (1996). A model of operant conditioning for adaptive obstacle avoidance. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., and Wilson, S., (Eds.), *From Animals to Animats 4*, pages 373–381. MIT Press, Cambridge MA.

[Henderson, 1999] Henderson, R. (1999). A continuously operating mobile robot. Master's thesis, Department of Computer Science, Manchester University.

[Kohonen, 1988] Kohonen, T. (1988). *Self Organization and Associative Memory.* Springer Verlag, Berlin, Heidelberg, New York.

[Marjanovic et al., 1996] Marjanovic, M., Scassellati, B., and Williamson, M. (1996). Self-taught visually-guided pointing for a humanoid robot. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., and Wilson, S., (Eds.), *From Animals to Animats 4*, pages 35–44. MIT Press, Cambridge MA.

[Nehmzow, 1998] Nehmzow, U. (1998). Self-supervised and supervised acquisition of smooth sensory-motor competences in mobile robots. In Cruse, H., Ritter, H., and Dean, J., (Eds.), *Prerational Intelligence in Robotics.* Kluwer Academic Publishers.

[Nehmzow, 1999] Nehmzow, U. (1999). *Mobile Robotics: A Practical Introduction.* Springer Verlag, Berlin, Heidelberg, New York.

[Nehmzow et al., 1991] Nehmzow, U., Smithers, T., and Hallam, J. (1991). Location recognition in a mobile robot using self-organising feature maps. In Schmidt, G., (Ed.), *Information Processing in Autonomous Mobile Robots.* Springer Verlag.

[Owen and Nehmzow, 1996] Owen, C. and Nehmzow, U. (1996). Route learning in mobile robots through self-organisation. In *Proc. Eurobot 96.* IEEE Computer Society, ISBN 0-8186-7695-7.

[Rao and Fuentes, 1996] Rao, R. and Fuentes, O. (1996). Learning navigational behaviors using a predictive sparse distributed memory. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., and Wilson, S., (Eds.), *From Animals to Animats 4*, pages 382–390. MIT Press, Cambridge MA.