Robotics and Autonomous Systems 24 (1998) 5–15

The scientific status of mobile robotics: Multi-resolution mapbuilding as a case study

Oliver Lemon, Ulrich Nehmzow *

Department of Computer Science, University of Manchester, Manchester M13 9PL, UK

Received 31 October 1997; accepted 30 April 1998

Abstract

We present a novel approach to mapbuilding and target area identification in mobile robotics – its emphasis being on the efficiency of multiple resolution representations for mapping and goal-area identification (for example, for the purpose of mobile robot self-localisation). After presenting and analysing the experimental results, we examine the scientific status of the research carried out, and of the field of mobile robotics in general. We argue that there is an implicit theory of mobile robotics, but that – for the future health of the discipline – it needs to be extended and made explicit through the development of an appropriate theoretical language and formal models. After exploring the relationship between design, engineering, and science, we argue that mobile robotics still lacks *falsifiable theories* of robot–environment interaction. This leads us to propose some desirable future directions for the discipline of mobile robotics as a maturing science. © 1998 Elsevier Science B.V. All rights reserved.

Keywords: Multiple resolution mapbuilding; Computational efficiency; Philosophy of science

1. An experiment in mobile robotics

Below we describe and analyse a project which develops a novel approach to topological mapping for mobile robotics. The central idea is to develop a hierarchical "multiple resolution" representation from the robot's perceptions of its environment, computing only at a relevant level of detail for a given navigation task. The robot may then perform the task of *efficiently* identifying a specified goal location. The resulting system is efficient since computation occurs only over the simplest representations necessary for each stage of the navigation. "Zooming" to a higher resolution of representation is map-driven, and the different resolutions are achieved by a uniform method using very small artificial neural networks of four nodes. Thus, as well as allowing efficient navigation, the robot's development of its map also uses few computational resources.

1.1. Motivation for the project

In concert with the topological mapping community in mobile robotics (e.g. [4,10]), we claim that a strategy for mapping a robot's environment should not, in the first instance, rely on information about odometry. Such dead reckoning navigation systems rely on proprioception (perception of internal states), and are subject to incorrigible errors due to wheel slippage. Instead, we argue that representations must be anchored in "perceptual landmarks", obtained through exteroception (perception of the external environment). Such perceptual landmarks will represent

* Corresponding author. E-mail: ulrich@cs.man.ac.uk.

PII: S0921-8890(98)00018-9

those features of the environment which are salient with respect to the robot's own perceptual apparatus. The concept of "multiple resolution" in mapping an environment seems a natural choice to us, and we conjecture that it will result in increased efficiency of navigation (both with regard to time and energy resources). Multiple resolution mapping facilitates the allocation of *minimum* resources for the task at hand. To increase resolution (and thus demands on computation and memory) is only required under specific circumstances, such as the identification of a goal location.

1.2. The mapping paradigm

Many researchers (most notably Lee [12]) have asked what constitutes a "good" robot map of an environment. This issue of map quality is, of course, relative to particular tasks (see e.g. [15]). We claim that details of *computational efficiency* of using topological maps have largely been neglected in appraisals of their suitability. Thus we argue for a modification, or enhancement, of the current research paradigm. We consider the following criteria relevant to mapbuilding:

(1) *Spatial consistency*. The map should faithfully represent the topology of some portion of the robot's environment (i.e. the description provided by the map should be consistent with the actual environment).
(2) *Coverage*. The map should represent as much of the environment as is required for completion of the task.
(3) *Relevant detail*. The map should only represent the environment in as much detail as is required for completion of the task.
(4) *Perception-based ontology*. The map should only be constructed on the basis of primitives which the robot is able to perceive reliably.
(5) *Computational efficiency*. The map-building process should use minimal computational resources for completion of the task.

Various proposals in the literature fail to meet at least one of the above criteria. As far as we know, none of them meet our requirement on relevant detail. For instance, if a robot's goal is to navigate to the end of a long uniform corridor, it is not necessary to represent and compute over many separate locations along that corridor. More efficient journeys can be made if more detailed topology is only computed around target locations. Thus we claim that a modification of the existing research paradigm is in order here.

Some theoretical results on topological mapping (e.g. [4]) fail to meet the last criterion, by relying on the information that is difficult, if not impossible, to obtain from a real robot (for example the assumption that there is an enumeration of all path directions leading from each location). In contrast to this, the results presented below are obtained from experimental data.

As far as we are aware, no other multiple-resolution approach to topological mapping exists. However, "hierarchical" representations of space have been proposed before, although in a quite different sense to that which we suggest. For instance Kuipers [9] advocates the use of a hierarchy of spatial representations, moving from topological to geometrical information, but this hierarchy consists of different classes of representation, rather than of different levels of resolution within the same representation scheme.

Various proposals also exist with regard to the use of self-organising feature maps (SOFMs) [5] in mapping strategies. Typically, sonar and infrared readings are fed into (computationally costly) high-dimensional (e.g. 10×10) SOFMs in order to develop a single representation of the environment. In contrast, our approach uses just three linear 4×1 SOFMs to map feature space at three levels of detail. Tree-structured feature mappings are employed by Koikalainen and Oja [6] (although they are not maps in our sense), and are known to facilitate search algorithms of low computational complexity.

1.3. Our hypothesis

Our hypothesis, then, is that a mobile robot can identify a location specified by the user, within an explored and mapped territory, by using a multi-resolution, perception-based mapping algorithm based on self-organisation.

The claim is that the robot's perceptions alone, used in the fashion described above, suffice to achieve target area identification, and that no external frame of reference (Cartesian coordinate system) needs to be used.

Our approach thus has the following features:

- Topological mapping.
- Multiple resolution, determined by the requirements of the task.

- Computational efficiency of representation.
- Perception-based representation.
- Use of data from a real robot.

1.4. Specification

Having generated our hypothesis, we now proceed to specify a particular task for the robot. It must:

(1) be able to develop a multiple-resolution perception-based topological map over a canonical path in an environment containing objects of different shapes and sizes, and made of different materials, and
(2) use this ability to identify to a specified location on the path, using higher resolution mapping only to disambiguate locations or when the target location might lie within its current region.

Now that we have specified the goal of our research, the precise experimental details can be described.

2. Experimental set-up

2.1. Robot and environment

The robot used for these experiments was a Nomad 200 mobile robot (see Fig. 1), named (after its serial number) "FortyTwo".

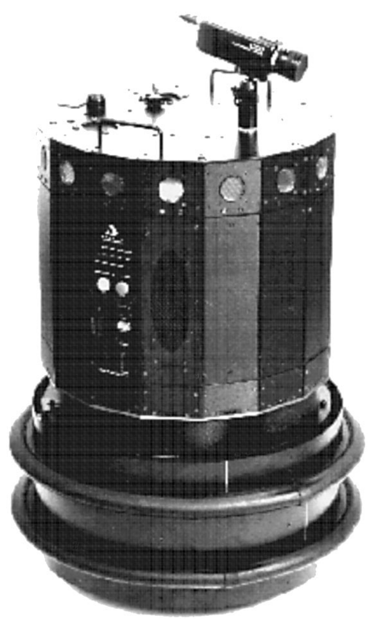

Fig. 1. The Nomad 200 mobile robot "FortyTwo".

This hexagonal robot is equipped with 16 ultrasonic range-finding sensors (range up to 6.5 m), 16 infrared (IR) sensors (range up to 60 cm), 20 tactile sensors and a monochrome CCD camera. The robot is driven by three independent AC motors for translation, steering and turret rotation. A 486 PC is the main controller with several slave processors (Motorola 68HC11) handling the robot's sensors.

As stated above, the robot's task was to recognise a user-specified target region within an environment consisting of walls of varying texture and colour, and a smooth, level floor.

2.2. Experimental procedure

First, "FortyTwo" was left to follow the wall of the environment that was to be mapped, using a wall following behaviour which had been acquired through unsupervised learning earlier, using an artificial neural network to associate perceptions and actions [18]. Sensor signals of all 16 sonar sensors, all 16 infrared sensors and odometric information were logged during this exploration. Sonar and infrared sensors were later used to obtain the multi-resolution mapping, the odometry information was used solely for manual analysis of the results obtained. The robot therefore mapped its environment as perceived along one canonical path, but it should be noted that multiple resolution mapping is independent of canonical paths or particular exploration strategies.

Data from two different environments (a lab and an untreated corridor), for two independent experiments was obtained in this manner, and subsequently used to construct the multi-resolution mapping off line. This experimental procedure can only be applied to investigate open-loop robotics tasks, i.e. tasks where the system responses under consideration do not require specific motor responses – that is the case here, where the reactive wall following behaviour of the robot and the mapbuilding system are independent from one another.

In such cases, this method of processing "live" data off-line has the advantage that it allows the precise replication of experiments under identical circumstances, while using real robot sensor data, and therefore being suitable for subsequent implementation on a mobile robot.

In the case presented here, the experimental method facilitated a comparison between the different levels of resolution, because in each experiment all maps were trained using identical data.

2.3. Preprocessing of data

It is well known that sonar sensor readings are subject to noise. For example, "specular reflections" occur when the sonar pulse hits a smooth surface at too shallow an angle to be reflected directly. The reflection of the sonar burst by some other object in the environment generates a reading that indicates a larger free space than is actually available. Likewise, "crosstalk" (receiving another sensor's sonar burst) can indicate shorter distances to the nearest object than actually present. In Fig. 2 (top left) such freak signals are visible, for example at datapoints 700 and 1700, which should exhibit identical sonar readings to datapoints 200 and 1200 (but don't).

In order to remove this noise, we filtered the sonar readings with a digital low pass filter before starting the mapbuilding. The results for two arbitrary sensors are shown in Fig. 2. As can be seen, the noise has been successfully removed, without throwing away too much of the information content of the signal. The filtering operation causes an end-transient effect in the signal (see Fig. 2, bottom right, datapoint 2000), which can pose problems for the mapping mechanism. We therefore removed the last 20 datapoints.

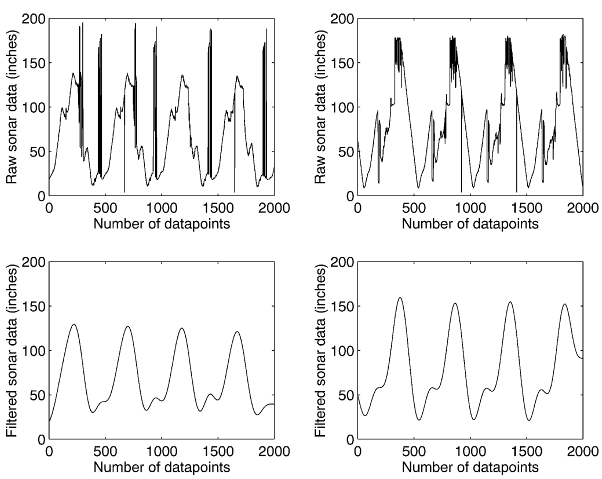

Fig. 2. Row (top) and filtered (bottom) sonar data for two different sonar sensors, before and after smoothing using a low pass filter. Spurious sonar signals, for example at datapoints 700 and 1700 (left) or datapoints 900 and 1400 (right) are removed, as is the high frequency content of the sensor signature.

2.4. The mapping mechanism

The concepts of topological and multi-resolutional maps, and their advantages, have been explained. We now focus on our particular approach. Self-organising feature maps (SOFMs) [5] have previously been used for mapping strategies (e.g. [11,20]): sonar and infrared readings were fed into high-dimensional (e.g. 10×10) SOFMs, requiring a constant high level of computation. Here, we propose three levels of resolution, on each level we use a linear 4×1 SOFM to map feature space,[1] thus providing a possible $4^3 = 64$ different place signatures. (A comparable single resolution mechanism would have to compare over 8×8 nodes.) Each increase in resolution may multiply the number of recognised regions by 4, so the size of a complete high resolution map of an area is bounded by that of the initial low resolution representation. Whenever necessary, we can switch from a lower resolution to a higher one, always increasing the information content of the input signal. The computational cost can therefore be reduced a large amount in certain "uninteresting" regions. Target regions, which the robot is to explore to the highest resolution, are specified externally to the system.

For the neural net training we used predefined "Matlab" functions from the neural net toolbox. The size of the linear (i.e. one-dimensional) self-organising feature map was four units, a neighbourhood size decreased from 3 to 1 over the training period. The learning rate decreased exponentially from 1 to (almost) 0. The SOFMs were trained over 5000 epochs (although fewer would almost certainly have sufficed), and each epoch consisted of an input chosen randomly from the data set.

On the lowest level of resolution, the *whole* of the explored space is mapped into distinct regions (see Fig. 3). As input to the lowest level SOFM we use four infrared sensors only, spaced 90° apart. The trained SOFM (training is completed after 5000

[1] One-dimensional SOFMs were used to achieve computational efficiency, while the fact that four units were used bears no particular significance.

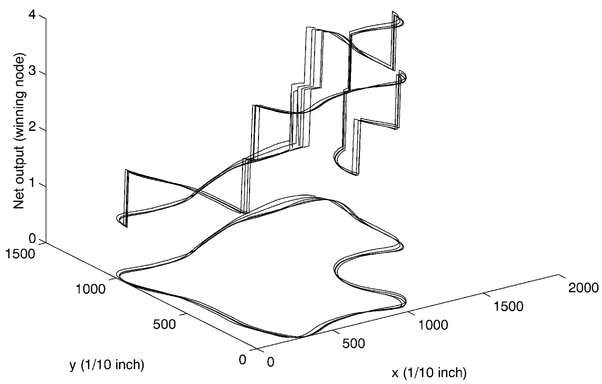

Fig. 3. Coarse resolution mapping of the entire experimental area. The robot's path through the environment is plotted "underneath" the signature map, on the x–y axes. The eventual goal region (specified in a later stage of the experiment) is near location $(0, 800)$, the perceptual signature obtained there is "2".

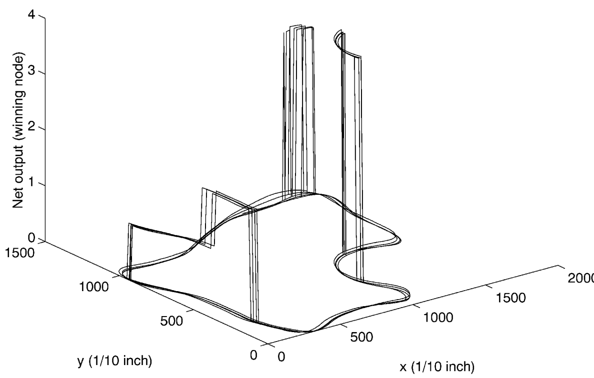

Fig. 4. Medium resolution mapping of all regions giving response "2" at the low resolution level. The goal region now has a perceptual signature of "2/1" and is, in this case, already uniquely identified by this signature.

learning steps) then responds with one of four responses to perceptual stimuli received anywhere along the path.

A goal region is then specified externally by the human operator, and at places that produce the same (low-resolution) output as this goal place a second network is trained at middle-resolution level (e.g. if the specified goal produces node "2" as the winning node at low resolution, sensor signals perceived at all locations generating "2" at the low-resolution level are used to train the middle-resolution network). At the middle-resolution level, four sonar sensors, again spaced 90° apart, are used to train the SOFM. Regions that have been developed in this middle level now each have a signature like for example "2/1" or "2/4".

To differentiate further within those regions, a third 4×1 SOFM is trained, now using all 16 infrared sensor signals and all 16 sonar sensor signals of the robot.

The three levels of mapping resolution are thus developed by:

- *Low level*: four infrared sensors.
- *Medium level*: four sonar sensors.
- *High level*: 16 sonar and 16 infrared sensors, i.e. all range sensors available on a Nomad 200.

Through this process a tree-like topological map is constructed, where each low- or medium-resolution region can be subdivided into smaller regions at a higher level of resolution.

3. Experimental results

3.1. *Experiment 1*

In the first experiment a circular path of 13.5 m was traversed four times, and 2000 sensor readings (sonar, infrared and odometry) obtained in a laboratory environment containing brick walls and cloth covered screens were used to produce the multi-resolution mapping of the environment. Fig. 3 shows the mapping of the entire environment at the lowest resolution level.

The "goal" location of the robot was externally specified as a location near the coordinates (0, 800). At the lowest resolution level, the perceptual signature is "2", therefore all regions having this perceptual signature were subsequently mapped at medium resolution level (see Fig. 4). The goal region then has the perceptual signature "2/1", and is, in this environment, already uniquely identified. However, a further increase in resolution shows that the goal region can be specified even more accurately (see Fig. 5). The complete signature for the goal location is "2/1/3" in this case.

3.1.1. *Discussion of Experiment 1*

Mobile robot self-localisation suffers from two major problems:

(1) Perceptual aliasing, i.e. the fact that several locations in the robot's environment give rise to identical sensory perceptions.

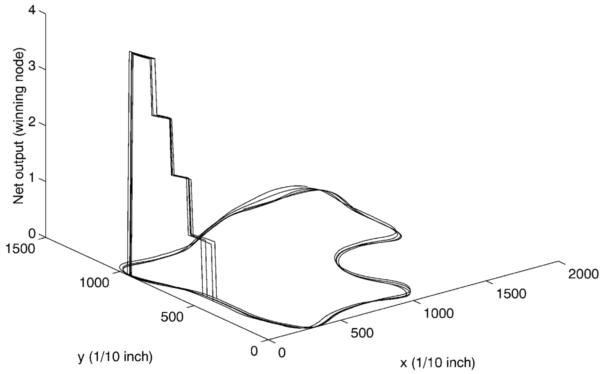

Fig. 5. A further increase in resolution shows that the region "2/1" can be subdivided into four further distinct regions, pinpointing the goal location even more accurately.

(2) Noise, i.e. the fact that different sensor responses can be perceived at the same physical location on subsequent visits, due to effects such as specular reflections or crosstalk. These effects are usually exacerbated by slight changes in robot position or orientation.

Regarding perceptual aliasing, Figs. 3–5 show that locations can indeed be identified uniquely through multi-resolution mapbuilding, even in the presence of perceptual aliasing (that perceptual aliasing *is* present can clearly be seen in Figs. 3 and 4).

Regarding the problem of noise, Figs. 3–5 show that map responses differed slightly at the edges of regions, for each of the four traversals of the route, but that within the centres of regions map responses were consistent.

3.2. Experiment 2

In a second experiment, "FortyTwo" was used to explore an untreated environment outside the laboratory, consisting of corridors, open spaces, doors, and irregularly distributed items of furniture. In total, a path of 35 m was traversed once, and 2000 datapoints were logged and used for subsequent multi-resolution mapping.

Fig. 6 shows the perceptual signature at the lowest resolution level, obtained along the robot's path in this environment. The goal region was specified to be at a distance of 360 in. from the beginning of the path. The perceptual signature at the lowest resolution level is

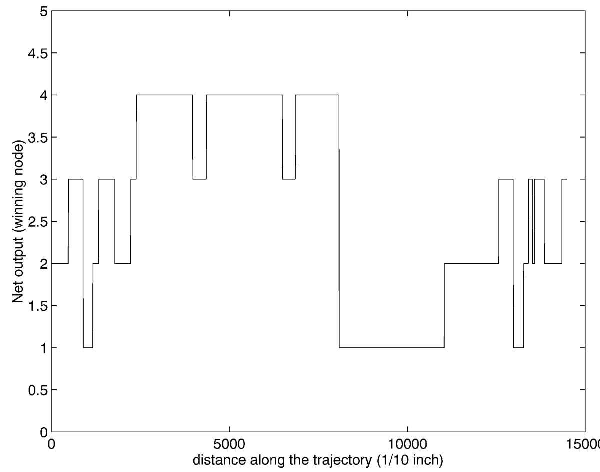

Fig. 6. Coarse resolution mapping of the second environment in its entirety.

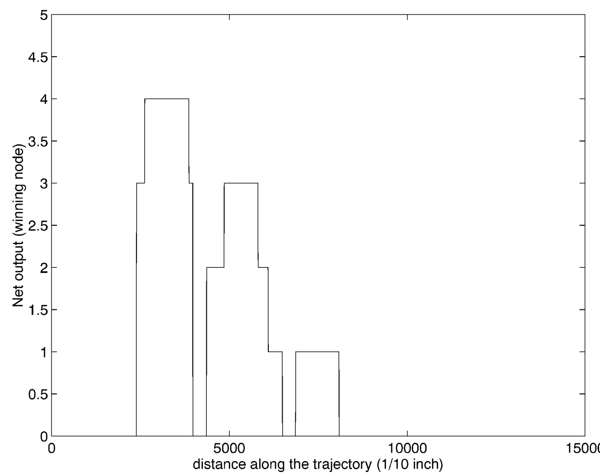

Fig. 7. Medium resolution mapping of all regions of the second environment that have a perceptual signature of "4" at the lowest level. The selected goal area now attains signature "4/3".

"4", so that all regions having that signature were subsequently mapped at a medium resolution level (see Fig. 7).

The perceptual signature of the goal location now becomes "4/3", so that all regions exhibiting that signature were then mapped at the highest resolution level (see Fig. 8). As can be seen from Fig. 8, the goal location is uniquely identified by the perceptual signature "4/3/3".

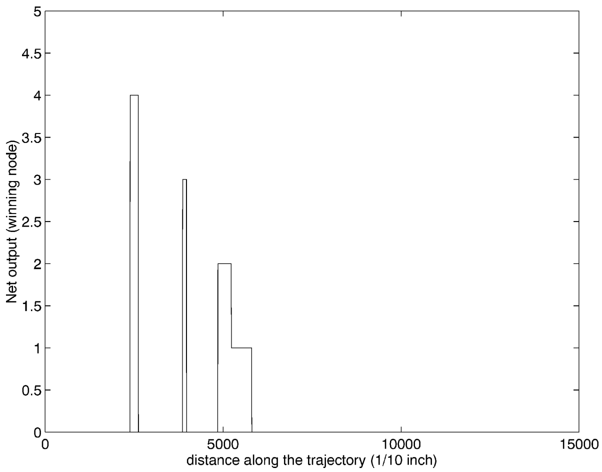

Fig. 8. High resolution mapping of those regions of the second environment that have a perceptual signature of "4/3" at the preceding two levels. The goal area is identified by the signature "4/3/3".

4. Analysis of the experimental results

Various requirements on mapping mechanisms in mobile robotics were presented in Section 1.2. We raised the issue of minimal resource use in the development of perception-based topological maps. Our criterion of "relevant detail" and thus computational efficiency of such a mapping strategy is one which, to the best of our knowledge, has not been noted before. The results presented above show that effective multiple resolution maps can be built from data obtained from a real robot, using a uniform mapping method. Efficiency arises not only in the development of the multiple-resolution map, but also in map-based navigation.

Thus we have, to some extent, satisfied our requirements (Section 1.2) and confirmed our initial hypothesis, in two experimental environments. This confirmation serves to reinforce the paradigm in which we are working.

As mentioned above, the computational cost of developing each level of representation is small (using only 4-node SOFMs) in comparison to other approaches. Further, the computational cost of planning routes over a small number of locations, in the low-resolution representation, is clearly less than that of the same problem over a more detailed representation. One might expect that this benefit could be obscured by the cost of developing a hierarchical representation, but since the more complex representations are developed only when required in order to complete the task (i.e. when the robot is in a target area), it is only in quite "uninteresting" environments (where perceptual aliasing is common, so that there are many target areas to be explored and disambiguated at higher resolutions) that computational costs might outweigh benefits. Thus, navigation should be faster, and more energy efficient, than in existing topological navigation methods.

Secondly, all landmark-based mapping algorithms suffer from perceptual ambiguities, largely due to the limited resolution of robot sensors. Focussing the high-resolution perceptions on areas of interest, rather than the whole environment, will reduce the impact of these perceptual ambiguities, leading to more robust localisation systems. However, such comparative studies are for future research at the present time.

4.1. Discussion of the experimental procedure

Our main concern in the experimental work presented here was that of *representation*, rather than *robustness*. Our goal was to develop an efficient mapping paradigm that would meet the criteria outlined in Section 1.2, but not, at this stage, to test its robustness in the presence of sensor noise.

Consequently, we did not divide our data into training set and test set, but trained the SOFM using all available data. Our experiments suggest that the proposed multi-resolution mapping paradigm indeed facilitates target area identification at low computational cost, and it is subject to future research to establish how sensitive the method is to noise, by using separate test data sets for evaluation.

Judging from Figs. 3–5, however, we expect the results of such a study to be similar to the ones reported here. From these figures it can be seen that sensory perception at the boundaries of regions varies slightly, but that it is robust in the centres of regions, therefore resulting in consistent map responses.

5. Analysis of the project

Our experimental results have demonstrated that a multi-resolution mapbuilding process can achieve reliable target area detection. But how did we arrive at this particular solution? Was it the result of good luck and clever engineering, or was the final system the result of applied systematic, methodological principles? In this project, as in current mobile robotics research in general, much credit has to be given to the former, but we argue that both our debt and our responsibility to scientific theory must also be acknowledged.

In this final section of the paper we make two claims:

(1) That every type of engineering (i.e. robot engineering, too) benefits from scientific theory.
(2) That an *implicit* theory of robot–environment interaction exists already, and that some research effort in mobile robotics should be directed at making this theory *explicit*.

5.1. The benefits of scientific theory

Readers may need to be persuaded as to the advantages of establishing a scientific theory of mobile robotics. Designing robots that are capable of performing a specified task in a target environment is, in the first instance, dependent upon competent engineering – but engineering itself rests upon, and applies, scientific knowledge. Science therefore lies at the heart of good engineering, but there are other motivations. In particular, having a good theory of some domain allows us to make accurate predictions about behaviour without having to go into the costly process of doing experimental work. Thus, having a scientific theory of some domain often dramatically reduces the search space that must be explored for possible solutions to a problem. Further, if we are interested in understanding *how* and *why* mobile robots function well, when they do so, then we are in the business of *explaining* their behaviour – and a scientific theory is an explanation *par excellence*.

It is difficult to see how any major engineering design ever accomplished could have been achieved without some scientific knowledge – be it knowledge of aerodynamics and mechanics in the case of aeronautical engineering, or of mechanics, physics, and material science in the case of civil engineering, to name but two examples. We claim that a similar situation exists in robotics – while it is certainly possible to design working robots purely on an engineering basis, it is a necessity for the progress of the field that deeper general theories of robot–environment interaction are developed. Such an understanding can only lead to the better design and better engineering of better robots. In our view, scientific theorising, of the most general nature, is central to good engineering and design, not separate from it. Good engineering is applied science.

5.2. Evaluation of the project as a scientific venture

This project, as well as research in robotics in general, exhibits the following stages of development:

(1) *Specification of the problem or question*. That is the environment, the robot, and the task, within a body of background literature and common practice (see Section 1.2).
(2) *Synthesis*. Building a robot that is expected to succeed in the specified task (prediction!).
(3) *Experiment*. Conducting relevant empirical studies.
(4) *Analysis*. Interpretation of the experimental results.
(5) *Conclusion*. Judgement regarding success or failure of the synthesis, and its possible modification.

Synthesising a possible approach to a particular problem (step (2)) implies that some underlying theory must exist, even if it cannot be expressed explicitly (yet). The roboticist designs the robot with the expectation (or prediction) that this particular robot, thus designed, is able to solve the particular problem in the target environment. We also observe that the project took place within a research paradigm which has its own vocabulary and methodology, and some care was taken with respect to precision and repeatability of the experiments. Note, however, that the following hallmarks of a scientific theory were lacking:

(1) Explicit universal generalisations and general predictions concerning robot-environment interaction.
(2) Explicit formal models of robot–environment interaction.

Thus, it is worth asking what, precisely, we learned about mobile robotics by way of the project? Did we establish any general principles, or refute any theoretical claim? No. We showed how to build a robot conforming to a particular design specification. Thus the status of the claim established by the project is existential, rather than universal. However, the results do fit into some broadly theoretical context – the current paradigm of mobile robotics research.

Similar research efforts can be observed in fields such as cognitive science, where researchers will specify a task which they find interesting for various theoretical reasons, and then implement a computer program or neural network, or whatever, such that it produces the desired behaviour. But this functionalist approach to explanation (that reproduction of input–output behaviour counts as an explanation) is unconvincing unless it takes place within a wider theoretical context. No physicist would claim to have explained motion by constructing a vehicle which moves in a particular way. But if the observed motion can be predicted precisely, and linked to other, well understood theoretical concepts (friction, gravitation, and thermodynamics, for example) then the physicist's experiment will be seen to have deeper theoretical import. Theoretical isolation and lack of predictive precision, we claim, is observable in experimental mobile robotics, as well as in other new sciences.

6. Towards a science of mobile robotics

It is clear that robotics cannot be considered a *natural* science – an inquiry into natural phenomena – for robots are human-built artifacts, and the goal of robotics research is to create functioning robots rather than to understand some pre-existing system. However, in order to construct better machines in general, we argue, one must understand them at some level of generality and abstraction. Not only that, but in order to create a working system, we have to be able to *predict*, at least to some degree, how its various components function and interact. Thus, while robotics is certainly not a natural science, we can sensibly place it amongst the sciences of the artificial, by which we mean theories of artifacts and technologies. Like Lelas [13], we believe that creating a functioning robot requires more than the description and understanding of the system provided by physics or engineering.

The research project presented and analysed here illustrates each of these claims. We claim that the community of researchers knows how to 'do robotics' – how to carry it out as a practice, but not how to articulate fully what it is they are achieving, or how their experiments bear on general theoretical issues. Indeed, it seems that there is currently little articulation of what these "general theoretical issues" actually are.

Because of observable design methods within the mobile robotics community (e.g. choice of sensors, use of sensors, choice of actuators, choice of control algorithm) we believe that an *implicit* theory of mobile robotics already exists, influencing the way mobile robots are built and experiments are conducted. However, as yet the theory has not been made *explicit*, such that specific predictions can be derived from a theoretical framework. In other words, mobile robotics still lacks *falsifiable theories* of robot–environment interaction [22].

Firstly, science any falsifiable theory has to be expressed in some language, we believe that the next step in a maturing science of mobile robotics has to be the establishment of a formal language in which to discuss different architectures, environments, and theories of robot–environment interaction. Notable attempts at establishing such a language and at providing design criteria and methodological principles for mobile robotics can be found in [21,23,25], and a theoretical foundation of "motor programming" in terms of neural representations is discussed in [7]. A more formal, information-theoretic approach to robot–environment interaction can be found in the field of "cognitive robotics" [24]. It seems that a future science of mobile robotics may be founded on a general theory which incorporates elements of each of these approaches – the fundamental task being to develop valid generalisations about robot–environment interaction from which specific details of the behaviour of particular systems can be predicted. Finding the right levels of abstraction and generalisation, and a sound theoretical vocabulary for the discussion, are the central issues here. Formalism, perhaps along the lines of [24], is developing in this regard.

Secondly, following Popper's notation that significant advances in knowledge are achieved either through the falsification of weak hypotheses, or the confirmation of bold ones, we believe that mobile robotics will advance as a science by careful analysis of our failures – judging from the number of reported failures, current practice appears to be not to report unsuccessful experiments at all – and by proposing (and evaluating) bold conjectures. Such a bold conjecture would be a general theory of robot–environment interaction.

7. Conclusion

By presenting and analysing an efficient multi-resolution mapbuilding and goal-area identification process, we argue that (a) successful robot controller designs are not the result of luck, but that an implicitly and partially known theory of robot–environment interaction guides the researcher, and that (b) future advances in the science of mobile robotics, both in terms of efficiency and general competence, require a theoretical understanding of the processes that govern robot–environment interaction, in other words a falsifiable, scientific theory of robot–environment interaction.

Acknowledgements

The authors thank Kris de Meyer for his help with the experimental work, and the referees for their careful and constructive criticism of this paper. Oliver Lemon is supported by the Leverhulme Trust (grant no. F/120/AQ).

References

[1] R. Beer, The dynamics of adaptive behavior: A research program, Robotics and Autonomous Systems 20 (1997) 257–289.
[2] A.F. Chalmers, What is This Thing Called Science?, 2nd Ed., Open University Press, 1982.
[3] T. Duckett, U. Nehmzow, Experiments in evidence based localisation for a mobile robot, Proceedings of the AISB Workshop on Spatial Reasoning in Animals and Robots, Department of Computer Science, Technical Report Series, ISSN 1361-6153. Report UMCS-97-4-1, Manchester, 1987.
[4] G. Dudek, P. Freedman, S. Hadjres, Using local information in a non-local way for mapping graph-like words, Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI-93), Morgan Kaufmann, Los Altos, CA, 1993.
[5] T. Kohonen, The self-organizing map, Proc. IEEE 78 (9) (1990) 1464–1480.
[6] P. Koikalainen, E. Oja, Self-organising hierarchical feature maps, IJCNN International Joint Conference and Neural Networks, vol. II, IEEE Press, New York, 1990, pp. 279–284.
[7] K. Kopecz, W. Erlhagen, G. Schöner, Dynamic representations provided the gradual specification of movement parameters, in: B. Kappen, S. Gielen (Eds.), Neural Networks: Artificial Intelligence and Industrial Applications, Proceeding of the SNN Symposium on Neural Networks, Nijmegen, 1995.
[8] T. Kuhn, The Structure of Scientific Revolutions, University of Chicago Press, Chicago, 1970.
[9] B. Kuipers, Modeling spatial knowledge, Cognitive Science 2 (2) (1978).
[10] B. Kuipers, Y. Byun, A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations, Robotics and Autonomous Systems 8 (1993) 47–63.
[11] A. Kurz, Constructing maps for mobile robot navigation based on ultrasonic range data, IEEE Transactions on Systems, Man, and Cybernetics – Part B: Cybernetics 26 (2) (1996) 233–242.
[12] D. Lee, The map-building and exploration strategies of a simple, sonar-equipped, mobile robot: An experiments, quantitative evaluation, Ph.D. Dissertation, UCL, 1995.
[13] S. Lelas, Science as technology, British Journal for the Philosophy of Science 44 (1993) 423–442.
[14] O. Lemon, I. Pratt, Spatial logic and the complexity of diagrammatic reasoning, in: Diagrammatic Reasoning (special issue), Machine Graphics and Vision 6 (1) (1997) 89–108.
[15] O. Lemon, I. Pratt, Putting channels on the map: A channel-theoretic semantics of maps, in: L. Moss (Ed.), Logic, Language, and Computation, vol. 2, CSLI Publications, Stanford, to appear.
[16] C. Malcolm, T. Smithers, J. Hallam, An emerging paradigm in robot architecture, in: T. Kanade, F.C.A. Groen, L.O. Hertzberger (Eds.), Intelligent Autonomous Systems, vol. 2, Amsterdam, 1989.
[17] K. de Meyer, O. Lemon, U. Nehmzow, Multiple resolution mapping for efficient mobile robot navigation, Proceedings “Towards Intelligent Mobile Robots 97”, Technical Report UMCS-97-9-1, University of Manchester, 1997.
[18] U. Nehmzow, Flexible control of mobile robots through autonomous competence acquisition, Measurement and Control 28 (2) (1995) 48–54.
[19] U. Nehmzow, Scientific methods in mobile robotics?!, Proceedings “Towards Intelligent Mobile Robots 97”, Technical Report UMCS-97-9-1, University of Manchester, 1997.
[20] U. Nehmzow, T. Smithers, Mapbuilding using self-organising networks, in: J.-A. Meyer, S. Wilson (Eds.) From Animals to Animats, MIT Press, Cambridge, MA, 1991, pp. 152–159.
[21] R. Pfeifer, Building “fungus eaters”: Design principles of autonomous agents, in: From Animals to Animats 4, MIT Press, Cambridge, MA, 1996, pp. 1–12.
[22] K. Popper, The Logic of Scientific Discovery, Hutchinson, London, 1968.

[23] G. Schöner, M. Dose, C. Engels, Dynamics of behavior: Theory and applications for autonomous robot architectures, Robotics and Autonomous Systems 16 (1995) 213–245.
[24] M. Shanahan, A logical account of the commonsense informatic situation for a mobile robot, Proceedings of Commonsense 96, 1996.
[25] T. Smithers, Taking eliminative materialism seriously: A methodology for autonomous systems research, in: F. Varela, P. Bourgine (Eds.), Proceedings of the First Conference on Artificial Life, MIT Press, Cambridge, MA, 1992, pp. 31–40.
[26] S. Thrun, A. Bucken, Integrating grid-based and topological maps for mobile robot navigation, Proceedings of the 13th National Conference on Artificial Intelligence (AAAI-96), Portland, OR, AAAI Press, Menlo Park, CA, 1996.
[27] B. Yamauchi, R. Beer, Spatial learning for navigation in dynamic environments, IEEE Transactions on Systems, Man, and Cybernetics – Part B: Cybernetics 26 (3) (1996) 496–505.

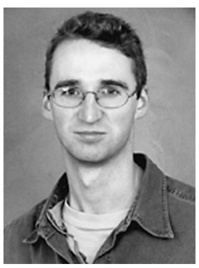

Oliver Lemon is a Research Associate in the Department of Computer Science at the University of Manchester. He studied mathematics and philosophy at Manchester University and received a Ph.D. from the Centre for Cognitive Science, Edinburgh, in 1996. He has taught logic, semantics, and philosophy to postgraduate level. His current research focuses on the formal properties of graphical representations – maps in particular. Recently he has also worked on modal logic, diagrammatic reasoning, robot navigation, and spatial logics. More generally, he writes on theories of representation and information, and philosophical issues concerning mind and language.

Ulrich Nehmzow is a lecturer in Artificial Intelligence and Robotics at the Department of Computer Science of the University of Manchester, leading the Robotics Research Group. He graduated in 1988 from the Technical University of Munich with a Diploma in Electrical Engineering and Information Science, and with a Ph.D. in Artificial Intelligence from Edinburgh University in 1992. Following appointments as research associate at the Department of Artificial Intelligence, Edinburgh University, and postdoctoral research associate at the Department of Psychology, Edinburgh University, he became lecturer at Manchester in 1994. He is a member of the Institution of Electrical Engineers, member of the IEE Northwest Center committee, member of the Society for Artificial Intelligence and Simulation of Behaviour, and has worked as Visiting Research Scientist at Carnegie Mellon University, Bremen University, and, most recently, as a Royal Society/STA Fellow at the Electrotechnical Laboratory in Tsukuba (Japan). His research interests are autonomous mobile robotics, navigation, robot learning, and artificial neural networks.