

AISB '03
Cognition
in Machines
& Animals

**Proceedings of the
AISB '03 Symposium on
Scientific Methods for the Analysis
of Agent-Environment Interaction**

7th – 11th April 2003

The University of Wales, Aberystwyth

EPSRC



Prifysgol Cymru
Aberystwyth
The University of Wales



The AISB'03 Convention
COGNITION in MACHINES and ANIMALS

7th - 11th April 2003
University of Wales, Aberystwyth

www.aish.org.uk

Proceedings of the Symposium on
Scientific Methods for the Analysis of
Agent-Environment Interaction

Printed at the University of Wales, Aberystwyth, SY23 4DD, Wales

Published by

**The Society for the Study of Artificial
Intelligence and the Simulation of
Behaviour**

www.aisb.org.uk

ISBN 1 902956 32 3

Printed at the University of Wales, Aberystwyth, SY23 3DB, Wales

The proceedings of the five symposia in the AISB'03 Convention are available from SSAISB. They are entitled as follows:

The Second International Symposium on Imitation in Animals and Artifacts
ISBN 1 902956 30 7

The Third Symposium on Adaptive Agents and Multi-Agent Systems
ISBN 1 902956 31 5

Scientific Methods for the Analysis of Agent-Environment Interaction
ISBN 1 902956 32 3

Biologically-Inspired Machine Vision, Theory and Application
ISBN 1 902956 33 1

Artificial Intelligence and Creativity in Arts and Science
ISBN 1 902956 34 X

The AISB'03 Convention

Cognition in Machines and Animals

The Society for the Study of Artificial Intelligence and the Simulation of Behaviour (SSAISB) is the UK's national society in the field of Artificial Intelligence. It is one of the oldest AI societies in Europe and has a long track record of supporting the UK research community. The society has a tradition of holding annual meetings and a convention format has proved successful, most recently at the Universities of Edinburgh (1999), Birmingham (2000), York (2001), and Imperial College, London (2002). We are delighted to host the 2003 convention at the University of Wales, Aberystwyth.

Recent advances in artificial intelligence, cognitive science, neuroscience and robotics have stimulated new interest in the relationships between biologically-based and computer-based cognitive systems. There is a current upsurge of activity both in biologically-targeted computer modelling and in adapting and experimenting with biological concepts for the design and development of novel computing technologies. For example, the BBSRC and EPSRC recently started a joint programme on Adaptive and Interactive Behaviour of Animals and Computational Systems (AIBACS). These topics are providing much inspiration for new research and innovation and the general theme of the 2003 convention reflects this stimulus and has attracted a range of high quality paper submissions and presentations.

In the early years of the SSAISB there was considerable interdisciplinary excitement about the general phenomenon of cognition, and lively meetings were attended by both computer scientists and psychologists. Perhaps the success of AI technology across a wide range of hard applications in the '80s and '90s caused the research focus to drift away somewhat from that most fascinating and astounding source of inspiration: cognition in the human brain. Consequently, it is very pleasing that we are now in the throes of a swing back towards those original concerns and I hope this convention will encourage further interaction between psychology and computer science, not least through the exciting selection of topics and invited speakers.

The convention is structured around five independently organized symposia, each with their own Programme Chairs and Programme Committee. The proceedings of all symposia are published and several of the best papers have been selected for edited publication in the AISB Journal. The unifying theme of the convention is reflected in a series of plenary sessions which feature invited keynote speakers. The six speakers, all international authorities in their field, are: Professor Peter Hobson, from the Tavistock Clinic and University College London; Professor Rajesh Rao, from the Center for Mind, Brain and Learning, University of Washington; Professor Giulio Sandini, from LIRA-Lab, University of Genova; Professor Sorin Solomon, from Theoretical Physics, Hebrew University of Jerusalem; Professor Yiannis Aloimonos, from the Computer Vision Laboratory, Institute for Advanced Computer Studies, University of Maryland; and Professor Gregor Schöner, from the Institut für Neuroinformatik, Ruhr-Universität Bochum.

We are extremely grateful for the support of the EPSRC, who provided funds for both the invited speakers' travel costs and for a large number of bursaries for postgraduate research students. The bursaries cover most of the cost of attending the convention and are particularly valuable as they enable the next generation of scientists to interact and to discuss issues with some of the key players in their research field. Support and assistance has been given by the Department of Computer Science at UWA, Meinir Thomas and the Conference Office at UWA, the IEE, and the Committee of the SSAISB, especially Geraint Wiggins. Special thanks are

due to all the symposia Chairs for all their time and hard work: Kerstin Dautenhahn, Chrystopher Nehaniv, Dimitar Kazakov, Daniel Kudenko, Eduardo Alonso, Ulrich Nehmzow, Gregor Schöner, Horst Holstein, Fred Labrosse, Pablo Gervas and Simon Colton. Special thanks also to our Technical Programme Assistant, Joanne Walker, and, for typesetting expertise, to Fred Labrosse. Finally, thanks to all the authors for their contributions, which form the essential core of the event. I am looking forward to the programme and I wish everyone a stimulating and productive meeting.

Mark Lee (Convention Host and Programme Chair)

Scientific Methods for the Analysis of Agent-Environment Interaction

The motivation behind this symposium on scientific methods for the analysis of agent-environment interaction is the desire to finally start with the “mopping up” activities that characterise “normal science” (Kuhn, *The structure of scientific revolutions*), the attempt “to force nature into the pre-formed and relatively inflexible box that a paradigm supplies.” The current paradigms underlying autonomous agents research still seem too weak to support “normal science”. When I first identified this as a problem within the discipline of autonomous agents in general, and autonomous mobile robotics specifically, I hadn’t realised that this concern was so widely shared within the community; but the call for papers for this symposium elicited numerous encouraging comments and sufficient contributions for a viable meeting, so that this first symposium (ever?) on scientific methods received the go-ahead from the AISB committee.

Arguably, research in agent-environment interaction hasn’t even reached the stage of “normal science” yet, because paradigms forming the foundation of our research are still emerging and formulated in vague terms. Still little “mopping up” is to be done, as we lack theories with predictive powers of agent-environment interaction. We have few precise instruments (such as for instance quantitative descriptions of agent-environment interaction) at our disposal, and are consequently left to present qualitative descriptions of experiments and existence proofs. These are by no means futile activities, but perhaps the time has come to aim for more rigorous experimental methods that allow independent replication and verification of experimental results.

The fact that this symposium is happening is very encouraging. Slowly, but inexorably the emergence of a new branch of research in behaving agents, namely “scientific methods”, is gathering momentum, and for the first time it will be discussed within the research community what we actually mean by “scientific methods for the analysis of agent-environment interaction”.

It will, on the one hand, mean characterisation through *measurement* of behaviour, and contributions in these proceedings will look at the evaluation of navigating robots (Hafner) and quantitative description of agent-environment interaction (Nehmzow and Walker).

“Scientific methods” will also entail aspects of computer modelling (Evans, Heuvelink and Nettle) and system identification (Singh and Singh), as well as the transfer of robotics research to related disciplines such as embodied computational neuroscience (Prescott, Gonzalez, Humphries and Gurney) and theoretical cognitive science (Ziemke).

This symposium would not have happened without the overall positive response by the scientific community, but neither would it have been possible without the dedicated work by Mark Lee, who masterminded the entire AISB convention, Gregor Schöner, who worked admirably as a programme chair for this symposium, and Joanne Walker, who kept the whole machinery running smoothly. Thank you to you all.

Ulrich Nehmzow, University of Essex
Symposium Chair

AISB Contents Papers

AISB Keynote Papers	1
Human Babies and Robot Cubs. <i>Giulio Sandini</i>	3
Dynamic field theory and embodied cognition. <i>Gregor Schöner</i>	4
Contributed Papers	5
Evaluating cognitive maps for mobile robot navigation behaviour. <i>V. V. Hafner</i>	7
Is the Behaviour of a Mobile Robot Chaotic? <i>U. Nehmzow and K. Walker</i>	12
The evolution of optimism: a multi-agent based model of adaptive bias in human judgement. <i>D. Evans, A. Heuvelink and D. Nettle</i>	20
A fuzzy modeling & identification of a dynamical system. <i>J. Singh and C. Singh</i>	26
On the analysis of robot behavior and its contribution to theoretical cognitive science. <i>T. Ziemke</i>	31
Towards a methodology for embodied computational neuroscience. <i>T. J. Prescott, F. M. Gonzalez, M. Humphries and K. Gurney</i>	38

Harvard Business School

Executive Education

Executive Education

AISB Keynote Papers

The AISB Keynote Papers are a series of papers presented at the AISB Annual Conference. The papers are presented by leading experts in the field of Artificial Intelligence and are a valuable resource for researchers and practitioners alike. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic.

The AISB Keynote Papers are a series of papers presented at the AISB Annual Conference. The papers are presented by leading experts in the field of Artificial Intelligence and are a valuable resource for researchers and practitioners alike. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic.

The AISB Keynote Papers are a series of papers presented at the AISB Annual Conference. The papers are presented by leading experts in the field of Artificial Intelligence and are a valuable resource for researchers and practitioners alike. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic. The papers are presented in a series of sessions, each focusing on a different topic.

Human Babies and Robot Cubs

Giulio Sandini

Faculty of Engineering
University of Genova

The role of technology for the study of brain functions has always been fundamental in providing new tools for the acquisition/analysis of biological data. However the increasingly complex picture of brain functions emerging from neuroscience research is now posing a new challenge: how to extend our knowledge beyond the scope of specific experiments and methodologies? Is it possible to find new tools enabling neuroscientists to verify new theories and to guide new experiments beyond the, now established, methods of mathematical modeling and system's theory? The scientific goal of the LIRA-Lab is to investigate if the implementation of artificial systems through physical models is a useful tool to help understanding complex brain functions. The reasons why we believe that this is indeed the case are, essentially, two. The first stems from the very high complexity and non-homogeneity of our current knowledge of brain functions. The second is that the physical world (in a general sense) is far too complicated to be "simulated" realistically preventing adequate testing of new theories and ideas to be performed.

In the context of human cognition, implementing "artificial systems" as explicit physical models of biological ones means implementing humanoid systems. Not with the aim of building an "artificial human being" or more efficient robots but to test assumptions and hypothesis more explicitly. In particular the use of humanoids as tools to understand human cognition, is focused, in our lab, on trying to explain how adaptation develops through interaction with the external environment. Our reference framework is human sensorimotor and cognitive development and we approach the problems by trying to implement motor and cognitive abilities in an artificial system. Is it possible to "program" a system to "have cognition" as we program a robot to assemble a car? Is cognition similar to motor control and sensorimotor coordination? Do we know enough about our own cognitive abilities to transfer them into an artificial being? How do we interact safely and "intelligently" with other humans (and machines)? How do we predict the effects of our actions? How do we adapt our behavior to unpredictable situations? How do we anticipate what other humans are doing? Can all (or even some) of these abilities be hard coded into a humanoid robot? Looking at natural systems it seems that pre-coding cognitive and adaptive behaviors is not possible (we believe that adaptive behaviors cannot be pre-programmed).

In this talk I will claim that if future robots have to have cognitive abilities, they will have to go through developmental phases similar to those found in human babies. I will do that from a multidisciplinary perspective by presenting findings derived from studies of human motor and cognitive development as well as a robotic implementation of the first few months of "existence" of a robot cub (Babybot). In doing so I will stress the consequences that this multidisciplinary approach has in discovering new technologies and the relevance that robotics research will continue to have as a research tool to understand human cognition.

Acknowledgements: Research described here is supported by the EU project COGVIS (IST-2000-29375) and MIRROR (IST-2000-28159) and by the Italian Space Agency.

Dynamic field theory and embodied cognition

Gregor Schöner

Institut für Neuroinformatik,
Ruhr-Universität,
Bochum
Germany

Overt behavior is always based on a direct link between sensory and motor surfaces. Moreover, behavior is typically adjusted to the perceived environment, a current "task" setting, and to longer term goals. These two boundary conditions of behavior are emphasized in the "embodied cognition" approach to cognition.

This talk introduces the concepts of dynamic field theory, in which simple cognitive properties such as decision making and working memory emerge from a mathematical description of behavior that remains close to sensori-motor processes. The mathematical framework is provided by dynamical systems theory. Autonomous robots designed in terms of these concepts will be used to exemplify the concepts. How these ideas can be used to analyze human behavior will be illustrated drawing on examples from motor control and the development of action planning.

Evaluating cognitive maps for mobile robot navigation behaviour

Verena V. Hafner*

* Artificial Intelligence Laboratory
Department of Information Technology
University of Zurich, Switzerland
vhafner@ifi.unizh.ch

Abstract

This paper examines biologically inspired cognitive map models, which provide an artificial navigating agent with a topological map of places after an exploration and learning phase in a previously unknown environment. The evaluation requires analytical methods that either rate the cognitive map on the basis of the map's features, or that rate the observed behaviour of the navigating agent. Some possible solutions to this problem are investigated, including map optimisation using artificial evolution, and alternative navigation behaviours.

1 Introduction

The navigation strategies of animals can be categorised according to their complexity, ranging from systematic search at the simplest end of the spectrum to complex wayfinding, with intermediate strategies such as local visual homing (Trullier et al. (1997), Franz and Mallot (2000)). The same categories can be applied to the behaviour of mobile robots or simulated agents navigating in their environments, where the choice of navigation strategy is usually both task dependent and dependent on the sensory information available to the agent.

More complex navigation strategies require an internal (neural) representation of the environment and are often referred to as 'cognitive maps' (O'Keefe and Nadel (1978)). Neurophysiological and behavioural experiments suggest that rats and primates use cognitive maps for navigation, whereas most insects have to rely on more basic navigation strategies. The cognitive maps referred to here are based on the following principles: *Place cells* are neurons found mainly in the hippocampus of rats, their activation is dependent on the location of the rat within its environment (O'Keefe and Nadel (1978)). *Head direction cells* are neurons whose activity is dependent on the orientation of the rat's head within the environment (Taube et al. (1990)). *Place fields* are areas in the environment where a particular place cell has the strongest activation. The evidence for cognitive maps in animals has inspired roboticists to implement those biological findings to allow for stable and adaptive navigation behaviour in mobile robots coping with (partly) unknown and ever-changing environments.

Here, we consider biologically-inspired cognitive map models, which provide an artificial navigating agent with a topological map of places (possibly enhanced with additional metric information) after an exploration and

learning phase in a previously unknown environment. Some models and implementations can be found in Schölkopf and Mallot (1995), Hafner (2000), Gaussier et al. (2002), Filliat and Meyer (2002). We are now faced with the problem of how to evaluate these representations with respect to their usefulness for navigation behaviour. Evaluating the cognitive map in a quantitative way is difficult, and there are several levels of analysis that may be applied. The evaluation requires either analytical methods that rate the cognitive map on the basis of the map's features, or that rate the observed behaviour of the navigating agent. Both methods require a large number of different exploration runs in different environments to make a general statement on the usefulness of the navigation strategy.

In section 2, some possible methods are described. Section 3 shows the difficulties of evaluating cognitive maps for navigation on a practical example of a robot navigating in a simulated environment, and shows how these measures can be used for a fitness function to optimise the agent's exploration and learning strategy. In section 4, analysing the behaviour of the agent will be discussed and further possible approaches will be outlined.

2 Methods for Analysis

Methods for the analysis of simulated cognitive maps include the use of statistical properties such as density, standard deviation, shape, and number of place fields; number, covered space, and activity of place cells; number of connections per place cell, graph distance vs. metric distance of the cognitive map. These data can be collected easily in simulation. In principle, the same can be done for a real world mobile robot, but various problems increase the difficulty of the task: The exact robot position has to be known at any time, a sensor information

database with place data for a large number of positions may be necessary, and the number of runs has to be higher than in simulation to make up for 'real world' noise. Since two different maps of the same environment can result in similar navigation behaviour, it is necessary to consider the navigation behaviour of the robot in addition to the analysis of the map. One possibility of evaluating the behaviour of a robot with a particular navigation strategy is to answer the following question: How long does it take to get from position A to position B after having explored the environment? To have any statistical significance, this has to be done for a large number of different positions, runs, and environments.

3 Evaluating Cognitive Maps: An Example.

3.1 Exploration and Learning

In this section, an example of an agent learning how to navigate in a virtual environment is shown. During an exploration phase, the agent has the task of learning a map of the environment using its internal neural structure such as place cells and connecting synapses. The usefulness of the cognitive map for a specific environment depends crucially on the exploration and map learning strategy.

The environment consists of a plane with twelve cylinders of different diameter (see figure 1 for a top view). The agent is equipped with a compass and omnidirectional 1D view with a resolution of 90 pixels. The cylinders are black, walls are white. The agent starts




Figure 1: Virtual environment for the robot navigation experiments. The environment consists of a plane with twelve cylinders (shown from above) of different size. The recorded views are taken at the crossing points of a 50×50 grid (indicated by tick marks on the frame). The ring shows an example of an omnidirectional view in this environment.

with a given number of place cells which are initialised at the beginning with random weights to the visual

input. It performs random exploration tours within the environment, avoiding bumping into the obstacles. At each step, the current view is taken, resulting in a certain neural activation in the place cell layer. Since two views at different positions could look similar, the activation of the place cells is in addition dependent on the previous place cell activation and the movement direction of the agent. Information on the connection and the connection heading of two place cells representing adjacent places is - for simplicity - directly encoded, instead of being available through head direction cells.

During exploration and learning, the neural weights are adjusted as follows:

1. The activation a is calculated for the current view v for each place cell, $a = f(vw + c + o)$, where w are the weights connecting the visual input and the place cells. c is a value which enforces place cells with a lateral connection to the previous winning cell to win, o accounts for the heading of the agent, which is also stored in the lateral connections. f is a sigmoid function.
2. The weights w of the cell with the highest activation a are updated and moved towards the input:
 $\Delta w = \eta_1(v - w)$
3. The weights w_- of the cell with the previously highest activation are updated and moved towards the input:
 $\Delta w_- = \alpha \eta_1(v - w_-)$
4. w are normalised for each place cell
5. The lateral connections w_l between the winning cell and the previous winner are updated:
 $\Delta w_l = \eta_2 a_1 a_2 (1 - w_l)$
6. The orientation information γ_l in the lateral connections between the winning cell and the previous winner is updated.
7. A decay factor is applied to all lateral connections:
 $w_l = \delta w_l$

In figure 2, the place field representations of the agent are projected onto the 2D environment. Areas of same colour are represented by the same place cell. The centres of mass of these place fields are shown as small white circles, with variable sized circles indicating the variance of the specific place field. Connections are drawn between the centres of mass of connected place cells. A map is well suited for navigation purposes when it contains many, roughly circular place fields with connections between regions that are adjacent and can be traversed.

One problem which can be seen in figure 2 is that of multiple place fields. Here, one place cell fires at




Figure 2: Map of the environment with black areas indicating obstacles, and grey coloured shadings indicating the different place fields after the agent has explored the environment. The centres of the place fields are interconnected as specified by the weights.

two spatially separated regions within the environment. Being unaware of this can result in strange behaviour. We know however that rats have to cope with the same problem. Multiple place fields have also been measured in the rat's hippocampus.




Figure 3: Average activation of the winner neuron for each grid position over time during the exploration and learning phase for one agent.

During the exploration and learning phase, some parameters introduced in section 2 which may be relevant to the quality of the resulting map are plotted over time

for 10,000 steps shown every 100th step. In figure 3, the mean activation of the winner place cell for each grid position is shown. Here, the place cell activation a is calculated purely from the visual input v and the neural weights w .

$$a = wv$$

Figure 4 shows the average variance v of the place fields over time for each place cell with a place field in the environment.

$$v_p = \frac{1}{n} \sum_i (x_i - \mu_p)^2$$

where μ_p is the centre of mass of the place field for place cell p , x_i a position on the grid where p is the winner cell (place cell with highest activation), and n the number of these positions. Figure 5 shows the number of connections between place cells as well as the average connection length using the metric distance between centres of mass of place fields from the plane in the simulation environment.




Figure 4: Average variance of the place fields over time during the exploration and learning phase for one agent.

3.2 Optimising Learning Parameters with Evolution Strategies

The learning strategy described in section 3.1 has been optimised to result in 'better' (to be defined) maps for the agent by evolving some of the following parameters:

- η_1 learning rate for weights w between visual input v and place cells
- α factor to learning rate of previous winning cells
- η_2 learning rate for lateral connections w_l between place cells
- β_1 contribution of lateral connections w_l to activation
- β_2 contribution of orientation γ_l to activation
- δ decay factor of map connections




Figure 5: Number of connections and average connections length (10*metric distance of centres of mass of place fields) during the exploration and learning phase for one agent.

Artificial evolution strategies have been used with the following fitness function:

$$f = f_1 * f_2$$

where f_1 is a measure for the regions (the average variance of the place fields), and f_2 is a measure for the connections (between adjacent place fields). f_1 simply contains the mean variance of all place fields, whereas f_2 is a mixture of positive fitness for adjacent place fields whose place cells are connected and negative fitness for connected place cells whose place fields are not adjacent. An alternative fitness function would be the correlation between metric and graph distance for all connections.

The evolution ran on the same randomly generated path through the environment for each individual in a generation. The number of map neurons (place cells) was chosen to be 50, the evaluation ran for 3000 steps each. A detailed description of the evolution strategy and its parameters can be found in Schwitter (2003). The evolved learning parameters of the described cognitive maps are: $\eta_1 = 0.045$, $\alpha = 0.62$, $\eta_2 = 0.5$, $\delta = 0.998$, $\beta_1 = 0$, $\beta_2 = 0$. Figures 2-5 are produced using these parameters.

Surprisingly, the evolutionary algorithm came up with a solution for map learning, which only takes into account the current view for the activation of the place fields, and rejects an influence of lateral connections and orientation information to the activation calculation ($\beta_{1,2} = 0$). A possible explanation for this is that the connections are not available at the beginning of the exploration tour, and the orientation error being high. The high error of the orientation information of the connections is due to the

size of place fields. Traversing from one to another place field can be done in a range of different directions.

4 Analysing the Behaviour

We saw from the previous section that evaluating the topological map in a quantitative way is difficult and often needs an external observer, which cannot be the agent itself but someone with a metric map of the environment as well as access to the agent's brain. Analysing the navigation behaviour after exploration and learning is more straightforward, but it is very time consuming and dependent on the specific environment the agent interacts with. One possibility is to assess the behaviour for a high number of randomly chosen place pairs, measuring how often and how fast an agent can navigate from one to another place given the previously learned cognitive map. The success rates in our environment were varying strongly between places. One interesting effect that could be observed is asymmetry: The success rate of navigating from place A to place B can be different than when navigating from place B to place A . The reason is the higher importance of correct connections to the goal place.

Falling back on other navigation strategies like visual homing facilitates the success of the agent. For example, a version of the average landmark vector (ALV) model (Lambrinos et al. (2000)) could be used to navigate between places. The catchment areas in this specific environment are rather large and reliable. An example is shown in figure 6.




Figure 6: Trajectories of an agent using a learned version of the ALV model to navigate between two places in the environment used for the cognitive map learning.

Although the agent does not have a metric map itself, it is able to extract some metric information from the topological map when it contains some additional information like orientation of the connections. Examples of the application of such algorithms are the spring force model (Hafner (2000)) or the relaxation algorithm (Duckett et al. (2002)).

One has to be careful when drawing conclusions from observations of the behaviour or the cognitive map structure. An interesting example is the observation in rats, that their place field density is higher around obstacles. An outside observer would conclude that these places are 'more interesting' and therefore get a better neural representation. However, when the same experiment was conducted with an artificial agent, a higher place field density could be observed around obstacles without the robot having any concept of 'interesting' (compare figure 2 or Hafner (2000)). The effect can simply be explained by the sensory information changing more rapidly in the vicinity of obstacles.

5 Discussion

We have shown on the example of learning cognitive maps optimal for navigation that evaluating these is often difficult because of the lack of suitable analysis methods. This applies even stronger for real world mobile robots, and leads us back to the well-known discussion on simulation vs. real-world experiments (Jakobi (1998)). We are either restricted using an unnatural, reproducible toy-world (or simulation), or have a robot in a natural environment (if an office room can be called natural) where the experiments cannot be compared with experiments in other environments. It may be more conclusive to compare the achieved navigation behaviour of a robot not just to other robots and other environments, but directly to the behaviour of navigating animals and humans in the same environment. What remains tricky is the effect of the difference in sensory modalities (e.g. vision, olfactory and somatosensory cues for rat navigation compared with camera and touch sensors on a mobile robot). Therefore, virtual environments would be ideally suited to compare human navigation behaviour with autonomous agent behaviour under the same conditions. There is even a virtual-reality setup for rats in preparation (Dahmen (2003)). A closer collaboration between behavioural robotics and fields like biology, psychology and mathematics would be a first step to create comparable and reproducible results, without having to resort to artificial quantification methods.

Acknowledgements

Many thanks to Christoph Schwitter for implementing the artificial evolution strategy as part of his Semesterarbeit. This research has been supported by grant #20-61372.00 of the Swiss National Science Foundation.

References

- H. Dahmen. Neural Correlates for Spatial Cognition in Virtual Environments. 2003. URL http://www.uni-tuebingen.de/cog/forschung/rodents_place/correlates_spatial_cognition.htm
- T. Duckett, S. Marsland, and J. Shapiro. Fast, On-Line Learning of Globally Consistent Maps. *Autonomous Robots*, 12:297–300, 2002.
- D. Filliat and J. A. Meyer. Global localization and topological map-learning for robot navigation. *Proceedings of the 7th international conference on simulation of adaptive behavior (SAB)*, 7:131–140, 2002.
- M. O. Franz and H. A. Mallot. Biomimetic robot navigation. *Robotics and Autonomous Systems*, 30:133–153, 2000.
- P. Gaussier, A. Revel, J. P. Banquet, and V. Babeau. From view cells and place cells to cognitive map learning: processing stages of the hippocampal system. *Biological Cybernetics*, 86:15–28, 2002.
- V. V. Hafner. Cognitive Maps for Navigation in Open Environments. *Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6)*, 2000.
- N. Jakobi. The Minimal Simulation Approach to Evolutionary Robotics. *Proceedings of ER'98, AAI Books, Kanata*, 1998.
- D. Lambrinos, R. Möller, T. Labhart, R. Pfeifer, and R. Wehner. A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems, special issue on Biomimetic Robots*, 30:39–64, 2000.
- J. O'Keefe and L. Nadel. *The hippocampus as a cognitive map*. Oxford University Press, 1978.
- B. Schölkopf and H. A. Mallot. View-Based Cognitive Mapping and Path Planning. *Adaptive Behavior*, 3(3): 311–348, 1995.
- C. Schwitter. Learning cognitive maps for navigation - Evolving parameters. *Semesterarbeit AI Lab, Univ. of Zurich*, 2003.
- J. S. Taube, R. U. Muller, and J. B. Ranck. Head-direction cells recorded from the postsubiculum in freely moving rats. *Journal of Neuroscience*, 10:420–447, 1990.
- O. Trullier, S. I. Wiener, A. Berthoz, and J.-A. Meyer. Biologically based artificial navigation systems: review and prospects. *Progress in Neurobiology*, 51:483–544, 1997.

Is the Behaviour of a Mobile Robot Chaotic?

Ulrich Nehmzow*

*Department of Computer Science
University of Essex
Colchester CO4 3SQ
udfn@essex.ac.uk

Keith Walker†

†Department of Physics and Engineering
Point Loma Nazarene University
San Diego CA 92106
kwalker@ptloma.edu

Abstract

We present a method to describe the behaviour of a mobile robot *quantitatively*, using methods from dynamical systems theory, time series analysis and deterministic chaos theory.

Experimental results obtained with a Pioneer II mobile robot demonstrate the use of the method, and show that robot behaviour exhibits deterministic chaos, and is substantially influenced by the control program executed by the robot, while changes to the environment have far less influence.

1 Background

1.1 Motivation

Research in mobile robotics to date has, with very few exceptions, been based on trial-and-error experimentation and the presentation of existence proofs. Task-achieving robot control programs are obtained through a process of iterative refinement, typically involving the use of computer models of the robot, the robot itself, and program refinements based on observations made using the model and the robot. This process is iterated until the robot's behaviour resembles the desired behaviour to a sufficient degree of accuracy. Typically, the results of these iterative refinement processes are valid within a very narrow band of application scenarios, they constitute "existence proofs". As such, they demonstrate that a particular behaviour *can* be achieved, but not, how that particular behaviour can *in general* be achieved for *any* experimental scenario.

The purpose of this paper is to introduce quantitative measures of robot behaviour, as components of a theory of robot-environment interaction. Using dynamical systems theory and methods of analysis derived from chaos theory, we investigate quantitatively in what way the behaviour of a mobile robot changes if a) the robot's environment is modified, and b) the robot's control code is modified. Underlying this research, however, is the fundamental question: How can the interaction of a mobile robot with its environment be characterised quantitatively?

1.2 Robot-Environment Interaction

The behaviour of a mobile robot cannot be discussed in isolation: it is the result of properties of the robot itself (physical aspects — the "embodiment"), the environment ("situatedness"), and the control program (the

"task") the robot is executing (see figure 1). This triangle of robot, task and environment constitutes a non-linear system, whose analysis is the purpose of any theory of robot-environment interaction.




Figure 1: THE FUNDAMENTAL TRIANGLE OF ROBOT-ENVIRONMENT INTERACTION.

Rather than speaking solely of a robot's behaviour, it is therefore necessary to speak of *robot-environment interaction*, and the robot's behaviour resulting thereof.

1.3 The Role of Quantitative Performance Measures

1.3.1 Measurement: The Backbone of Science

Measurement is the backbone of science, and supports

- the precise documentation of experimental setups and experimental results,
- the principled modification of experimental parameters,
- independent verification of experimental results,
- theoretical design of artefacts without experimental development, and

- predictions about the behaviour of the system under investigation.

The experimental scenarios within mobile robotics, using *quantitative* performance measures, are manifold and expressed in figure 1. Basically, if any two of the three components shown in figure 1 remain unaltered, then the quantitative performance measure will characterise the third, modified component. This would allow the investigation of, for instance, i) the effect of modifications of the robot, ii) the influence of the robot control program on robot behaviour, or iii) the effect of modifications to the environment on the overall behaviour of the robot.

The purpose of the experiments reported in this paper was to demonstrate how robot-environment interaction could be characterised quantitatively, and how such quantitative measures could be used to determine the influence of i) a change in the robot controller, and ii) a change of environment.

1.3.2 Three Theses

The experimental work reported in this paper is based on three theses:

Thesis 1 A mobile robot, interacting with its environment, is essentially an analog computer that “computes” *behaviour* (the output) from the three inputs *robot morphology*, *environmental characteristics* and *executed task*.

Thesis 2 The behaviour exhibited by a robot, interacting with its environment while executing some particular task, is encapsulated, as a first approximation, in the robot’s *trajectory*.

Thesis 3. One suitable method to describe a robot’s trajectory quantitatively is to analyse the trajectory taken, using methods of time-series analysis and dynamical systems theory.

Based on these theses, we analysed the (x, y) components of logged robot trajectories, using time-series analysis methods from chaos theory.

1.4 Related Work

So far, quantitative measures of robot-environment interaction are not widely used in mobile robotics, and the amount of related work is consequently limited. The most important references are probably the work by Schöner et al. and Smithers.

Schöner et al. (1995) used dynamical systems theory to investigate robot-environment interaction, and Smithers (1995) discussed the use of quantitative performance measures as a tool of scientific mobile robotics research.

Regarding the methods of characterisation used, the most important references are found in the area of time series analysis using chaos theory. Computing Lyapunov

exponents, we used Arbarbanel’s [Arbarbanel (1996)] and Wolf’s et al. [Wolf et al. (1995)] methods. We computed correlation dimensions again using Arbarbanel’s method, as well as the method discussed by Kaplan and Glass [Kaplan and Glass (1995)]. Background reading to the methods applied are Peitgen et al. (1992) and Kantz and Schreiber (1997).

1.5 Experimental Setup

To investigate the two questions posed in section 1.3.1, we conducted experiments with a Pioneer II autonomous mobile robot (figure 2), executing a number of different control programs in a number of different environments.




Figure 2: THE PIONEER II MOBILE ROBOT USED.

The robot’s trajectory was logged, using an overhead camera. Every 250 ms the robot’s x and y -coordinate were recorded. Figure 3 gives an example of the kind of trajectory obtained by this method, the dataset in that figure contains just over 26000 data points, recording 109 minutes of robot operation.

Figure 4 shows part of the x and y coordinate of that trajectory — trajectories like these were used for subsequent computation of Lyapunov exponent and correlation dimension (see section 2).

2 Quantitative Measures of Behaviour using Chaos Theory

2.1 Reconstructing the Attractor

The behaviour of dynamical systems, such as a mobile robot interacting with its environment, is characterised by velocity and position along each degree of freedom of the system — the phase space. Because a direct representation of the phase space of our robot-environment system is not available to us, the phase space has to be reconstructed from an observed time series such as the one shown in figure 4.




Figure 3: ROBOTIC BILLIARD BALL (OBSTACLE AVOIDANCE) BEHAVIOUR — (ENTIRE TRAJECTORY (TOP) AND 150 DATA POINTS (BOTTOM)).




Figure 4: x (TOP) AND y COORDINATE (BOTTOM) OF THE DATA SHOWN IN FIGURE 3 (PART).

To achieve this, we used the time-lag embedding method (discussed, for instance, in Peitgen et al. (1992) or (Kaplan and Glass, 1995, p. 309)), embedding the attractor in p -dimensional space, using the coordinates given in equation 1.

$$\mathbf{D}_t = [D_t, D_{t-h}, D_{t-2h}, \dots, D_{t-(p-1)h}], \quad (1)$$

where $\mathbf{D}_t$ is the p -dimensional embedding of the attractor at time t , p the embedding dimension, and h the embedding lag.

As an example, figure 5 shows a three-dimensional phase space reconstruction of the attractor underlying the robotic billiard ball behaviour shown in figure 3.




Figure 5: PHASE SPACE RECONSTRUCTION OF ROBOTIC BILLIARD BALL EXPERIMENT, USING THE x COORDINATE, EMBEDDING DIMENSION 3 AND AN EMBEDDING LAG OF 30.

2.1.1 Determining Embedding Lag and Embedding Dimension

In order to reconstruct the attractor, using the embedding method given by equation 1, embedding dimension p and embedding time lag h have to be determined. To determine a suitable embedding lag, we used two methods:

1. The autocorrelation method presented in (Kaplan and Glass, 1995, p. 353): according to Kaplan and Glass “a good choice for the embedding lag is [that value] at which the autocorrelation function falls to $e^{-1} \approx 0.37$ ”. Figure 6 shows the autocorrelation function for the x -coordinate of the data shown in figure 4, it can be seen that for $t = 31$ the autocorrelation has fallen below 0.37. Because $t = 1$ represents an embedding lag of $h = 0$, we select an embedding lag of $h = 30$ to reconstruct the attractor.
2. The Mutual Information Method: This procedure is associated with non-linear statistics while the autocorrelation method is based on linear statistics. In essence, mutual information is analogously a non-linear autocorrelation function. Mutual Information has its origin from the treatment of chaos as a source of information [Gallager (1968)]. The mutual information essentially measures the degree to which knowledge of one measurement determines another

measurement h seconds later. Fraser and Swinney (1986) argue that one should take the first minimum that occurs in the mutual information function as the ideal time lag for constructing an embedded attractor. This minimum in the mutual information function corresponds to the time separation between two measured values so that there is a minimum of redundancy of information connecting the two measurements in the time series. Figure 7 gives the Mutual Information of the x-coordinate data shown in figure 3. The first minimum of this curve is obviously at a time lag of 30, thereby giving the same result as the autocorrelation function.




Figure 6: AUTOCORRELATION OF THE X-COORDINATE OF THE DATA SHOWN IN FIGURE 4. FOR A LAG OF $h = 30$ THE AUTOCORRELATION HAS DROPPED BELOW e^{-1} .




Figure 7: MUTUAL INFORMATION OF THE X-COORDINATE OF THE DATA SHOWN IN FIGURE 4. THE FIRST MINIMUM OCCURS FOR A LAG OF $h = 30$.

It is known [Takens (1981)] that the dynamics of a reconstructed attractor are geometrically identical to the original dynamics (for both continuous and discrete systems), if the relationship $p = 2\nu + 1$ holds (ν is the dimension of the original attractor). In order to establish a

suitable embedding dimension p to reconstruct the attractor, it is therefore necessary to determine the dimension ν of the original attractor. One common measure of the dimension of an attractor is the correlation dimension ν .

In order to determine the correlation dimension, one uses the *correlation integral* $C(r)$, defined in equation 2.

$$C(r) = \frac{\theta}{N(N-1)}, \quad (2)$$

where θ is the number of times that $|\mathbf{D}_i - \mathbf{D}_j| < r$, i and j are two different times at which the embedding $\mathbf{D}$ is taken, and r is an arbitrary distance measure, the "correlation distance". $N(N-1)$ the maximum number of cases where $|\mathbf{D}_i - \mathbf{D}_j| < r$ is theoretically possible (the trivial case $i = j$ is excluded).

The important measure is not the correlation integral $C(r)$ itself, but how this integral changes with respect to r , i.e. the slope of the curve $C(r)$ versus r . This slope is the correlation dimension ν .

The mechanism to determine a suitable embedding dimension to reconstruct the attractor is based on these considerations. Obviously, the computation of the correlation dimension is dependent upon the chosen embedding dimension p and the correlation distance r . To compute both p and ν from the same process is an ill-defined problem, and the goal therefore is to find a range of parameters p and r for which ν is computed identically (a so-called "scaling region"). In other words, one aims to find a region where the computation of the correlation dimension ν is not critically dependent upon the choice of embedding dimension p and correlation distance r .

To find a scaling region, we plot the correlation dimension ν as a function of correlation distance r for all embedding dimensions p between 1 and 10 (see (Kaplan and Glass, 1995, page 323)). Figure 8 shows ν versus p and r for the x-coordinate shown in figure 4.




Figure 8: CORRELATION DIMENSION VERSUS CORRELATION DISTANCE FOR THE DATA SHOWN IN FIGURE 4, FOR VARIOUS EMBEDDING DIMENSIONS AND VARIOUS CORRELATION DISTANCES. $h = 30$.

Choosing a large correlation distance r is somewhat like looking at an object from a great distance: it will have dimension zero (Kaplan and Glass, 1995, p. 323). On the other extreme, choosing too small an r will result in signal noise being amplified, which is equally undesirable. Figure 8 reveals that for $r \approx 110$ the computed correlation dimension is similar for a large number of embedding dimensions, in other words, increasing the embedding dimension no longer changes the computed correlation dimension. The correlation dimension for $r = 110$ is just above 2, meaning that an embedding dimension of 5 is a good choice. This finding, incidentally, is confirmed when the ‘false nearest neighbours’ method is used.

2.2 Characterising the Attractor

Systems that exhibit deterministic chaos are characterised by four properties:

1. They are stationary,
2. they are deterministic,
3. they are aperiodic, and
4. they show sensitivity to initial conditions.

The first two properties are easily established using mechanisms discussed, for instance, in (Kaplan and Glass, 1995, p.314) and (Kaplan and Glass, 1995, p.324ff). It is the latter two properties that are interesting for the purpose of this paper, as they supply a *quantitative* description of the attractor.

2.2.1 Aperiodicity

One main characteristic of a dynamical system exhibiting chaos is that the state variables never return to their exact previous values: The trajectory in phase space lies on a strange attractor. There is, however, variation from system to system in how close state variables return to previous values, and it is therefore desirable to quantify this degree of ‘proximity’.

The measure to quantify the degree of aperiodicity is the correlation dimension ν , which was discussed in detail above. Besides being a quantitative description of the attractor, the correlation dimension indicates whether data is aperiodic or not, and to what degree: Periodic data has an integer correlation dimension, while chaotic attractors have a non-integer correlation dimension.

2.2.2 Sensitivity to Initial Conditions

Another distinctive characteristic of a chaotic system is its sensitivity to a variation in the system’s variables: Two trajectories in phase space that started close each other will diverge from one another as time progresses, the more chaotic the system, the greater the divergence.

Consider some state S_0 of a deterministic dynamical system and its corresponding location in phase space. As time progresses the state of the system follows a deterministic trajectory in phase space. Let another state S_1 of the system lie arbitrarily close to S_0 , and follow a different trajectory, again fully deterministic. If d_0 is the initial separation of these two states in phase space at time $t = 0$, then their separation d_t after t seconds can be expressed as $d = d_0 e^{\lambda t}$.

λ is known as the Lyapunov exponent. For a m -dimensional phase space, λ will have m -components. If any one or more of those components are positive, then the trajectories of nearby states diverge exponentially from each other in phase space and the system is deemed chaotic. Since any system’s variables of state are subject to uncertainty, a knowledge of what state the system is in can quickly become unknown if chaos is present. The larger the positive Lyapunov exponent, the quicker knowledge about the system is lost. One only knows that the state of the system lies somewhere on one of the trajectories traced out in phase space, i.e., somewhere on the strange attractor.

The Lyapunov exponent is one of the most useful quantitative measures of chaos, since it will reflect directly whether the system is indeed chaotic, and will quantify the degree of that chaos. Also, knowledge of the Lyapunov exponents becomes imperative for any analysis on prediction of future states.

Determination of the Lyapunov exponents from a time series such as the robot trajectories used in this research or, more importantly, the existence and value of any positive exponents, has been discussed extensively in the literature [Wolf et al. (1995); Peitgen et al. (1992); Kantz and Schreiber (1997); Abarbanel (1996); Kaplan and Glass (1995); Abarbanel et al. (1993)]. Several academic and commercial software packages have been made available to compute Lyapunov exponents from time series, for example Wolf et al. (1995); Applied Nonlinear Sciences LLC and Randle Inc (2003); H. and Schreiber (2003). In our analyses we have used the method proposed by Wolf et al. (1995), as well as Abarbanel’s commercially available software package Applied Nonlinear Sciences LLC and Randle Inc (2003) to determine the largest Lyapunov exponent.

Figure 9 shows the result of the computation of the Lyapunov exponent for the data shown in figure 4, using the algorithm described in Wolf et al. (1995).

It can be seen that the algorithm converges to λ just above 0.3, which confirms our earlier finding this data exhibits deterministic chaos.

3 Experiments: Analysing Robot-Environment Interaction

We stated above that robot-environment interaction is governed by the triangle of robot, task and environment




Figure 9: COMPUTATION OF THE LYAPUNOV EXPONENT OF THE X-COORDINATE OF THE DATA SHOWN IN FIGURE 4, USING THE METHOD OUTLINED BY WOLF ET AL. (1995).

(figure 1). If, therefore, two of the three components remain constant, Lyapunov exponent and correlation dimension will characterise the third component and its influence on the overall behaviour of the robot.

Of the many experiments we conducted, we would like to discuss two experiments in this paper:

1. In experiment 1 we kept the environment constant, and altered the robot control code.
2. In experiment 2, we kept the control code constant, but altered the environment.

In both experiments the same robot was used.

3.1 Changing the Robot Task

In the first experiment, the Pioneer II operated in an empty environment, surrounded by plasterboard walls. The robot executed two different behaviours: wall following (13000 data points, 54 minutes of operation, see figure 10) and “billiard-ball” obstacle avoidance (26000 data points, 109 minutes of operation, see figure 3).

For the wall following behaviour, we computed a correlation dimension of $\nu \approx 1.5$ and a Lyapunov exponent of $\lambda \approx 0.15$. For the billiard ball behaviour noticeably different values were computed, i.e. $\nu \approx 1.8$ and $\lambda \approx 0.29$. In other words: The billiard ball behaviour exhibited a considerably larger degree of deterministic chaos than the wall following behaviour.

3.2 Changing the Environment

In the second environment, we kept both robot and task constant (using the billiard ball behaviour, but changed the environment by introducing a central obstruction to the arena). The trajectory obtained for that case is shown in figure 11.




Figure 10: WALL FOLLOWING BEHAVIOUR OF THE PIONEER ROBOT — ENTIRE TRAJECTORY (TOP) AND 100 CONSECUTIVE POSITIONS (BOTTOM).

Here, both the correlation dimension ($\nu \approx 1.7$) and the Lyapunov exponent ($\lambda \approx 0.3$) remained essentially the same to the values obtained when no central obstruction was present. In other words: The same degree of deterministic chaos was present, irrespective of whether the central obstruction was in place or not.

4 Conclusion

4.1 Towards a Theory of Robot-Environment Interaction

Mobile robotics research to date is still largely dependent upon trial-and-error development, and results often are existence proofs, rather than generalisable, fundamental results. Furthermore, independent replication and verification of experimental results are largely unknown in mobile robotics research practice.

One reason for this is the dearth of quantitative measures of behaviour. We have therefore investigated the application of dynamical systems theory — specifically deterministic chaos theory — to the analysis of robot environment interaction.




Figure 11: ENTIRE TRAJECTORY OF BILLIARD BALL BEHAVIOUR IN AN ENVIRONMENT WITH A CENTRAL OBSTRUCTION (TOP), AND 200 CONSECUTIVE POSITIONS (BOTTOM).

Using Lyapunov exponent and correlation dimension of the (strange) attractor underlying our Pioneer robot's behaviour, we found that the robot's behaviour indeed exhibited deterministic chaos. Moreover, we found that the degree of deterministic chaos increased noticeably when the robot's control program was changed from wall following to billiard ball obstacle avoidance, while it remained essentially constant when the environment was changed, but the control program remained the same.

However, we see the main contribution of this paper not in the specific analysis of our Pioneer II mobile robot, but in presenting a method of analysing robot-environment interaction quantitatively. Quantitative descriptions form the backbone of any scientific theory, and are therefore the first step towards a theory of robot-environment interaction, that allows the formation of hypotheses, precise reporting and independent verification of results, and the theoretical design of robot systems.

4.2 Interpretation of Results

The specific interpretation of our experimental results is that in the experimental environment we have used, the robot control code has far more influence on the behaviour of the robot than the environment itself. The wall following behaviour exhibits a far lower degree of deterministic chaos than the billiard ball behaviour, while the billiard ball behaviour shows a similar degree of chaos in two different environments.

4.3 Future Work

The obvious extension to the experiments reported here is to conduct the wall following behaviour in an environment that is different to the one that is shown in figure 10 — for example an environment similar to the one shown in figure 12.




Figure 12: POSSIBLE ALTERNATIVE ENVIRONMENT TO INVESTIGATE WALL FOLLOWING BEHAVIOUR.

Our hypothesis is that in such an environment we would obtain a correlation dimension of $\nu \approx 1.5$ and a Lyapunov exponent of $\lambda \approx 0.1$.

Acknowledgement

This project was supported by a grant from the Research Promotion Fund of the University of Essex. We acknowledge this support gratefully.

References

- Applied Nonlinear Sciences LLC and Randle Inc. Csp time series analysis software package, 2003. URL <http://www.zweb.com/apnonlin/csp.html>.
- H.D.I. Abarbanel, R. Brown, J.J. Sidorowich, and L.S. Tsimring. The analysis of observed chaotic data in physical systems. *Rev. Mod. Phys.*, 65:1331–1392, 1993.
- Henry Abarbanel. *Analysis of Observed Chaotic Data*. New York: Springer-Verlag, 1996.
- A.M. Fraser and H.L. Swinney. Independent coordinates for strange attractors from mutual information. *Phys. Rev. A.*, 33:1134–1140, 1986.
- R.G. Gallager. *Information Theory and Reliable Communication*. New York: John Wiley and Sons, 1968.
- Kantz H. and T. Schreiber. Tisean — nonlinear time series analysis, 2003. URL <http://ls11-www.cs.uni-dortmund.de/people/hermes/NLDdocs/docs/>.
- H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge University Press, 1997.
- D. Kaplan and L. Glass. *Understanding Nonlinear Dynamics*. Springer Verlag London, 1995.
- H.-O. Peitgen, H. Jürgens, and D. Saupe. *Chaos and Fractals*. Springer Verlag, 1992.
- G. Schöner, M. Dose, and C. Engels. Dynamics of behavior: theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems*, 16, 1995.
- T. Smithers. On quantitative performance measures of robot behaviour. *Robotics and Autonomous Systems*, 15:107–133, 1995.
- F. Takens. Detecting Strange Attractors in Turbulence, volume 898 of *Lecture Notes in Mathematics*. Berlin: Springer-Verlag, 1981.
- A Wolf, J. Swift, H. Swinney, and J. Vastano. Determining lyapunov exponents from a time series. *Physica 16D*, 1995.

The evolution of optimism: A multi-agent based model of adaptive bias in human judgement

Dylan Evans*

* Biomimetics Group
Department of Mechanical
Engineering
University of Bath
Bath BA2 7AY
UK
D.Evans@bath.ac.uk

Annerieke Heuvelink*~

~ Department of Philosophy
University of Utrecht
Heidelberglaan 8
3584 CS Utrecht
The Netherlands
A.Heuvelink@phil.uu.nl

Daniel Nettle°

° Departments of Psychology
and Biological Sciences
The Open University
Walton Hall
Milton Keynes MK7 6AA
UK
D.Nettle@open.ac.uk

Abstract

Human judgement under uncertainty has been shown to involve consistent departures from normative rationality, but it is hard to tell whether these biases are adaptive or not. We have attempted to throw light on this question by constructing a multi-agent based computer model. In the model, a variety of agents with different decision rules are allowed to compete in a variety of environments. We find that under certain environmental conditions biased agents who behave in ways similar to those observed in humans outperform the classically rational agent, who acts purely to maximise expected utility. These conditions are similar to those in which humans find themselves.

1 Introduction

Human judgement under uncertainty has been shown to involve consistent departures from normative rationality. In particular, people show 'motivational biases' in judgements of probability, over-estimating the probability of events with a positive return to the self and under-estimating the probability of events with a negative return (Miller & Ross, 1975; Zuckerman, 1979).

These biases disappear when people are depressed ('depressive realism') and when people are asked to estimate the probability of the same events happening to *others*. Apart from these situations, however, the optimistic bias is remarkably resilient, persisting despite repeated disappointment and evidence to the contrary.

From the standpoint of rational choice theory, these biases are clearly maladaptive. Some psychologists, however, have argued that they are adaptive (Taylor & Brown, 1988). We have attempted to adjudicate between these two possibilities by constructing a multi-agent based computer model.

In the model, a variety of agents with different decision rules are allowed to compete in a variety of environments. We find that under certain environmental conditions biased agents who behave in ways similar to those observed in humans outperform the classically rational agent, who acts purely to maximise expected utility. Moreover, it is plausible that these conditions are similar to those that humans find themselves in.

Our findings therefore add support to the view that motivational biases are in fact adaptive.

2 Methods

In order to test our hypothesis we designed a multi-agent based simulation using NetLogo. NetLogo is an agent-based parallel modelling and simulation environment produced by the Center for Connected Learning and Computer-Based Modelling at Northwestern University.

In our model, patches represent 'opportunities'. Each opportunity has a 'probability of success' (p , ranging from 0-1), a benefit for success (b , ranging from 0.0001 to 10 energy points) and a cost of failure (c , ranging from 0.0001 to 10 energy points). The colour of the patch is determined by the probability of success, with darker patches representing more difficult opportunities.

Agents have only one goal - to maximise their energy points. In other words, their utility function is a linear function of their energy level. Agents have some knowledge of the cost of failure (c), the benefit for success (b), and the probability of success (p), for each opportunity they face.

The values of c , b and p are properties of the patch that the agents find themselves on at any given moment. The level of noise affecting the agents' knowledge of these values can be set by means of the error-sliders on the interface (see Figure 1).

The error-sliders can vary from 0 (perfect information) to 10 (great uncertainty). This error determines the standard deviation used for the normal distribution of which the mean is the true value of c, b or p of the patch. A random number drawn from this distribution determines the agents' guesses about the values of c, b and p. There are two error sliders; one affects the agents' knowledge of p, the other affects the agents' knowledge of c and b.


Agents do not move, but since each patch is updated each turn, each agent is presented with a different opportunity at every time step. Each turn, every agent must decide whether or not to 'play' that opportunity or not. This decision is made according to the agent's 'decision rule'. There are three types of agent, each with a different decision rule:

1. The RATIONAL agent uses the principle of expected utility; i.e. it only plays when the expected utility of playing is greater than that of not playing.

2. The OPTIMISTIC agent also uses the principle of expected utility, but uses a biased estimate of p (its estimate of p is multiplied by its estimate of b divided by its estimate of c).
3. The EMOTIONAL agent always plays if it estimates b to be more than twice c, and never plays when it estimates b to be less than half c. When it estimates that b and c are between these limits, its chance of playing is proportional to its estimate of p.

If an agent decides to play an opportunity, its chance of success is determined by the probability of success associated with that opportunity. If it plays and succeeds, its energy level is increased by the benefit for success associated with the opportunity. If it plays and fails, its energy level is decreased by the cost of failure associated with that opportunity. If an agent does not play, its energy level remains the same for that turn. Agents start with zero energy. Agents never die, and there is no reproduction. Figure 1 shows the interface of the program.²

Figure 1: Interface Program



3 Results

We let the program run 10 times for every possible combination of the error of p and the error of c and b . One run lasts for 500 time steps. For each run we recorded the final average energy of each type of agent. We used this data to calculate the means and 95% confidence interval for all those runs for each kind of agents.

As can be seen from the graphs (Figures 2 – 5) there are significant differences (the error bars do not overlap). We show four specific 2D graphs, and an overall 3D graph seen from two different angles (Figures 6 and 7). The 3D graph does not show error bars but gives an overall view of our results.

Figure 2:

Graph of the data with ErrorCB = 2

Shown are the means and error bars using 95% Confidence Interval.





Figure 3:

Graph of the data with ErrorCB = 4

Shown are the means and error bars using 95% Confidence Interval.



Unsurprisingly, the rational agents do better than all other agents under most conditions. More interesting is the fact that there are conditions under which the rational agent is outperformed. This occurs when the error affecting agents' knowledge of c and b is quite small but the error affecting agents' knowledge of p is high. Both the emotional and optimistic agents do better then.

When the error affecting agents' knowledge of all three variables (c , b , and p) is high, all the agents perform at similar levels. Under these conditions there is great uncertainty and so they all perform quite badly.

Figure 4:

Graph of the data with ErrorCB = 6

Shown are the means and error bars using 95% Confidence Interval.




Figure 5:

Graph of the data with ErrorCB = 8

Shown are the means and error bars using 95% Confidence Interval.





Figure 6:




4 Discussion

Dealing with uncertainty is a common problem for agent systems. The ability to reason with uncertain information is an indispensable requirement for modeling intelligent behavior in a complex and dynamic environment. This is why uncertain reasoning has become a major research topic of AI with many important applications. These circumstances of uncertainty are the ones we investigated.

Psychologists have shown that human judgement under uncertainty involves consistent departures from normative rationality (Nettle, 2003; Kahneman & Tversky, 1973; Kahneman et al, 1982). In particular, people show motivational biases in judgements of probability, over-estimating the probability of events with a positive return to the self and under-estimating the probability of events with a negative return (Harris & Middleton, 1994; Larwood & Whitaker, 1977; Weinstein, 1980;1982).

These biases disappear when people are depressed ('depressive realism') (Alloy & Ahrens, 1987; Alloy & Abramson, 1979) and when people are asked to estimate the probability of the same events happening to *others* (for example Mirels, 1980). Apart from these situations, however, the optimistic bias is remarkably resilient, persisting despite repeated disappointment and evidence to the contrary.

Figure 7:



From the standpoint of rational choice theory, these biases are clearly maladaptive. Some psychologists, however, have argued that they are adaptive (Taylor & Brown, 1988).

Our results show that there are environmental conditions under which biased agents outperform the classical rational agent, who acts purely to maximise expected utility. Our findings therefore add support to the view that motivational biases are in fact adaptive.

It is interesting to see that the biased agents, which we have designed to have biases similar to those shown by humans, do best when the error for b and c is low, but the error for p is high.¹ This is arguably the situation that people mostly encounter in the real world. It is plausible to think that people can estimate costs and benefits of an opportunity quite accurately, by observing other people faced with similar opportunities and by memories of past experiences. However, the chances of success of any specific opportunity depend on the interaction with other human beings and many other imponderable factors. Hence our ability to estimate probably is much poorer than our ability to estimate the cost of failure or the benefit of success.

¹ We have based the decision rules used by our biased agents necessarily on our interpretation of the empirical literature, but this literature is complex and clearly open to different interpretations.

Our emotional agents are based on the observation that when the benefits of a certain action are high people tend to play, independent of what the probability of success is.

The lottery is a case in point. The chance of success here is so small that nobody with a rational view will buy a lottery ticket. However, the benefit of success is so much high and the cost of failure so low that people do play. On the other hand, if there is a certain decision to make and there are high costs associated with failure but low benefits associated with success, then it is not likely that people will play, even if the probability of success is quite high. Only when the difference between costs and benefits is not that big do people seem to attend to the probability of success (Rottenstreich & Hsee, 2001).

We translated this empirical data to the rule for the emotional agents, such that they always play if the benefit of success is equal to or greater than twice the cost of failure, and never play if the cost of failure is equal to or greater than twice the benefit of success. Between those two extremes they only pay attention to the probability of success.

There is a certain degree of arbitrariness associated with the cut-off points in this rule, but the point here is not to provide a detailed model of the psychological process in real humans. Rather, we are interested merely in showing that biased agents that resemble humans to some first approximation can do just as well, or better than the purely rational agent.

Our optimistic agents are inspired by the empirical data about the situations in which people's judgement is biased.

A recurrent finding in the empirical literature is that when the benefits of an opportunity are higher than the costs, people tend to be more optimistic about their chances of success. They estimate the probability of success as higher than it really is and hence increase their change of playing. Conversely, when the costs of failure are higher than the benefits of success, people are pessimistic in the sense that they estimate their chances of success to be smaller than they actually are, so decreasing the change of playing.

Thus while the normative model of expected utility theory suggests that behavioural decisions should be based on the simple product of the probability of success and the benefit, people's weightings of probabilities in fact follow an inverse S-shaped function (Kahneman & Tversky, 1984).

We modelled this by giving our optimistic agents a biased estimate of p ; in their case, the estimate of p is multiplied by the estimate of b divided by the estimate of c . The result is that the agents play more than would be rational if the benefits are higher

than the costs, and play less if it is the other way around.

Our rational agents are clearly inspired by the classical decision theory (Von Neuman & Morgenstern, 1944). They only play if the expected utility of playing is greater than that of not playing.

In this program we did not investigate the possibility of giving different computational costs to different decision rules. This would be interesting, since then we could introduce an evaluation function that favoured cheaper rules.

The problem with such an experiment is that we do not have any empirical data about the computational costs associated with different rules, at least insofar as they are implemented in the human brain.

It would, however, be possible to calculate the computational cost of implementing the various decision rules in a specific robot. This could be done by measuring how long it takes the robot to make a certain decision following one rule and compare that to the other rules. In such circumstances we might find that a cheaper rule, that performs a little worse than the more expensive rule, will perform better in total.

5 Conclusion

From the standpoint of classical decision theory, motivational biases are clearly irrational. We have found that under certain environmental conditions biased agents, who behave in ways similar to those observed in humans, outperform the classically rational agent, who acts purely to maximise expected utility. Moreover, it is plausible that these conditions are similar to those that humans find themselves in.

Our findings therefore add support to the view that motivational biases are in fact adaptive.²

Acknowledgements

A Platform Grant from the Engineering and Physical Sciences Research Council, EPSRC GR/M97503 supported Dylan Evans' work.

² The code of our program can be found on:
<http://www.dylan.org.uk/optimismAISB.nlogo>

References

- Alloy, L. B. & Abramson, L. Y. (1979) Judgement of contingency in depressed and non-depressed subjects: Sadder but wiser? *Journal of Experimental Psychology: General*, **108**, 443-479.
- Alloy, L. B. & Ahrens, A. H. (1987) Depression and pessimism for the future: Biased use of statistically relevant information in predictions for self versus others. *Journal of Personality and Social Psychology*, **41**, 1129-1140.
- The Center for Connected Learning and Computer-Based Modelling
Northwestern University
Annenberg Hall 311
2120 Campus Drive
Evanston, IL 60208
USA
Homepage: <http://ccl.northwestern.edu/>
- Harris, P. & Middleton, W. (1994) The illusion of control and optimism about health: On being less at risk but no more in control than others. *British Journal of Social Psychology*, **33**, 369-386.
- Kahneman, D., Slovic, P. & Tversky, A. (eds) (1982) *Judgement under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kahneman, D. & Tversky, A. (1973) On the psychology of prediction. *Psychological Review*, **80**, 237-251.
- Kahneman, D. & Tversky, A. (1984) Choices, values and frames. *American Psychologist*, **39**, 341-350.
- Larwood, L. & Whitaker, W. (1977) Managerial Myopia: Self-serving biases in organisational planning. *Journal of Applied Psychology*, **62**, 194-198.
- Miller, D. T. & Ross, M. (1975) Self-serving biases in attribution of causality: Fact or fiction? *Psychological Bulletin*, **82**, 213-225.
- Mirels, H. L. (1980) The avowal of responsibility for good and bad outcomes: The effects of generalized self-serving biases. *Personality and Social Psychology Bulletin*, **6**, 299-306.
- Von Neuman, J., Morgenstern, O. (1944) *Theory of Games and Economic Behaviour*. Princeton: Princeton University Press.
- Nettle, D. (forthcoming) Adaptive illusions: optimism, control and human rationality. In *Emotion, Evolution and Rationality*, eds. Dylan Evans & Pierre Cruse, Oxford: Oxford University Press.
- Rottenstreich, Y. & Hsee, K. (2001) Money, kisses and electric shocks: On the affective psychology of risk. *Psychological Science*, **12**, 185-190.
- Taylor, S. E. & Brown, J. D. (1988) Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin*, **103**, 193-201.
- Weinstein, N. D. (1980) Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, **39**, 806-820.
- Weinstein, N. D. (1982) Unrealistic optimism about susceptibility to health problems. *Journal of Behavioural Medicine*, **5**, 441-460.
- Zuckerman, M. (1979) Attribution of success and failure revisited: The motivational bias is alive and well in attribution theory. *Journal of Personality*, **47**, 245-287.

A Fuzzy Modeling & Identification of a Dynamical System

Jagmohan Singh
BUC College, Batala-143521 PB India
Ph: 0091-1874-234162 ®
E-mail: jagmohanbajwa@yahoo.co.in

Charandeep Singh
Govt. Polytechnic, Batala-143521 PB India
Ph: 0091-1874-235025
E-mail: cdeep Singh@yahoo.com

Abstract

Identification of dynamic systems from input-output measurements is an important topic of scientific research with a wide range of practical applications. Many real-world systems are inherently nonlinear and cannot be represented by linear models used in conventional system identification. In this paper, fuzzy dynamic system modeling and identification method is presented and discussed with NARX (Non linear Auto Regressive with Exogenous input) model. The method identifies the fuzzy rules describing the dynamic system behavior from input-output data using TS (Takagi- Sugeno) models. In order to show the procedure performance, the fuzzy modeling & identification of a dynamic system is presented with an example. The simulation of this example is also reported using MATLAB.

Keywords: Fuzzy modeling, dynamic system, identification, simulation, validation

1 Introduction

The problem of developing a general methodology for intelligent system design has always been demanding. Fuzzy modeling and identification methodologies have been successfully used in a number of real-world applications [1]. The Takagi-Sugeno model has often been employed in the fuzzy modeling and identification of nonlinear technical systems from input-output data. Examples are the modeling of a human operator's control action of water cleaning process and a converter in a steel-making process [2]. The approach to qualitative modeling based on fuzzy logic has been proposed with fuzzy c-mean clustering method for the structure identification of a gas furnace, chemical plant, water purification process, trend of stock prices etc. are described in [3]. Fuzzy identifier is constructed to estimate the engine signals necessary to calculate the deviation from nominal engine behavior for engine fault detection and isolation of internal combustion engine based on c-mean clustering in [4]. A study is described for identification of models for the temperature within the melter portion of a glass-melting furnace [5].

Science and engineering have typical examples of areas where conventional modeling techniques do not give satisfactory results. Fuzzy models can serve as decision support systems to assist operators, or can be used to clone the operators based on traces of their behavior. Fuzzy modeling is a framework in which different modeling and identification methods are combined, providing, on the one hand, a transparent interface with the designer or the operator and, on the other hand, a flexible tool for nonlinear system modeling and control, comparable with other

nonlinear black-box techniques. The rule-based character of fuzzy models allows for a model interpretation in a way that is similar to the one humans use.

This work addresses a fuzzy modeling and identification of a dynamic system with NARX model. The method identifies the fuzzy rules describing the system behavior from input-output data using TS models. With this approach, the domains of the antecedent variables are simply partitioned into a specified number of equally spaced and shaped membership functions. The rule base is then established to cover all the combinations of the antecedent terms. The consequent parameters are estimated by the least-squares method.

2 TS Model

The TS fuzzy model, which uses crisp functions in the consequents, can be seen as a combination of linguistic modeling and mathematical regression, in the sense that the antecedents describe fuzzy regions in the input space in which consequent functions are valid [3]. The TS rules have the following form:

$$\begin{aligned} R_i : & \text{If } x \text{ is } A_i \\ \text{then } & y_i = f_i(x), \quad i = 1, 2, \dots, K. \end{aligned} \quad (1)$$

Variable $x \in X \subset R^p$ and $y \in Y \subset R^q$ is the domain variable in which fuzzy sets are defined. The membership functions of the antecedent and consequent fuzzy sets are the mappings: $\mu(x) : X \rightarrow [0,1]$ and $\mu(y) : Y \rightarrow [0,1]$,

respectively. Fuzzy set A_i define fuzzy regions in the antecedent space, for which the respective consequent propositions hold. The linguistic terms A_i are usually selected from sets of predefined terms, such as small, medium, etc. The rule base is $R = \{R_i, i = 1, \dots, K\}$

Contrary to the linguistic model, the input x is a crisp variable (linguistic inputs are in principle possible, but would require the use of the extension principle (Zadeh) to compute the fuzzy value of y_i).

The functions f_i are typically of the same structure; only the parameters in each rule are different. Generally, f_i is a vector-valued function. A simple and practically useful parameterization is the affine (linear in parameters) form, yielding the rules:

$$\begin{aligned} R_i : & \text{If } x \text{ is } A_i \\ \text{then } & y_i = a_i^T x + b_i, \quad i = 1, 2, \dots, K. \end{aligned} \quad (2)$$

Where a_i is a parameter vector and b_i is a scalar offset. If $a_i = 0$ for each i , the singleton model is obtained. The inference formula of the TS model is:

$$y = \frac{\sum_{i=1}^K \beta_i y_i}{\sum_{i=1}^K \beta_i} = \frac{\sum_{i=1}^K \beta_i (a_i^T x + b_i)}{\sum_{i=1}^K \beta_i} \quad (3)$$

where β_i , the degree of fulfillment of the i th rule's antecedent. When the antecedent fuzzy sets define distinct but overlapping regions in the antecedent space and the parameters a_i and b_i correspond to a local linearization of a nonlinear function, the TS model can be regarded as a smoothed piece-wise linear approximation of that function. The affine TS model can be regarded as a quasi-linear system (i.e., a linear system with input-dependent parameters). To see this, denote the normalized degree of fulfillment by:

$$\gamma_i(x) = \beta_i(x) / \sum_{j=1}^K \beta_j(x) \quad (4)$$

Here we write $\beta_i(x)$ explicitly as a function x to stress that the TS model is a quasi-linear model of the following form:

$$\begin{aligned} y &= \left(\sum_{i=1}^K \gamma_i(x) a_i^T \right) x + \\ &\sum_{i=1}^K \gamma_i(x) b_i = a^T(x) x + b(x) \end{aligned} \quad (5)$$

The parameters $a(x)$, $b(x)$ are convex linear combinations of the consequent parameters a_i and b_i i.e.:

$$a(x) = \sum_{i=1}^K \gamma_i(x) a_i, \quad b(x) = \sum_{i=1}^K \gamma_i(x) b_i \quad (6)$$

In this sense, a TS model can be regarded as a mapping from the antecedent (input) space to a convex region (polytope) in the space of the parameters of a quasi-linear system. This property facilitates the analysis of TS models in a framework similar to that of linear systems [3].

3 Modeling of Dynamic Systems

The time-invariant dynamic systems are in general modeled by static functions, by using the concept of the system's state [1]. Given the state of a system and given its input, we can determine what the next state will be. In the discrete-time setting we can write:

$$y(k+1) = f(y(k), u(k)) \quad (7)$$

Where $y(k)$ and $u(k)$ are the state and the input at time k , respectively, and f is a static function, called the state-transition function. Fuzzy models of different types can be used to approximate the state-transition function. As the state of a process is often not measured, input-output modeling is often applied. The most common is the NARX model:

$$\begin{aligned} y(k+1) &= f(y(k), y(k-1), \dots, \\ &y(k-n_y+1), u(k), u(k-1), \dots, \\ &u(k-n_u+1)) \end{aligned} \quad (8)$$

Here $y(k), \dots, y(k-n_y+1)$ and $u(k), \dots, u(k-n_u+1)$ denote the past model outputs and inputs respectively and n_y, n_u are integers related to the model order. For example, a

singleton fuzzy model of a dynamic system may consist of rules of the following form:

$$\begin{aligned}
 R_i: & \text{ If } y(k) \text{ is } A_{i1} \\
 & \text{ and } y(k-1) \text{ is } A_{i2} \\
 & \text{ and } \dots y(k-n+1) \text{ is } A_{in} \\
 & \text{ and } u(k) \text{ is } B_{i1} \\
 & \text{ and } u(k-1) \text{ is } B_{i2} \\
 & \text{ and } \dots u(k-m+1) \text{ is } B_{im} \\
 & \text{ then } y(k+1) \text{ is } C_i
 \end{aligned} \tag{9}$$

Since fuzzy models can approximate any smooth function to any degree of accuracy, models of type (9) can approximate any observable and controllable modes of a large class of discrete-time nonlinear systems.

4 Knowledge-Based Design

To design a linguistic fuzzy model based on available expert knowledge, the following steps can be followed: (1) Select the input and output variables, the structure of the rules, and the inference and defuzzification methods. (2) Decide on the number of linguistic terms for each variable and define the corresponding membership functions. (3) Formulate the available knowledge in terms of fuzzy if-then rules. (4) Validate the model by using data. If the model does not meet the expected performance, iterate on the above design steps.

It should be noted that the success of the modeling and identification heavily depends on the problem at hand, and the extent and quality of the available knowledge. For some problems, the knowledge-based design may lead fast to useful models, while for others it may be a very time-consuming and inefficient procedure, especially manual fine-tuning of the model parameters. Therefore, it is useful to combine the knowledge-based design with a data-driven tuning of the model parameters. The following sections review method for the adjustment of fuzzy model parameters by means of data.

4.1 Data-Driven Acquisition and Tuning of Fuzzy Models

In this section, we assume that a set of N input-output data pairs (x_i, y_i) $i = 1, \dots, N$ is available. Recall that $x_i \in R^p$ are input vectors and y_i are output

scalars. Denote $X \in R^{N \times p}$ a matrix having the vectors x_k^T in its rows, and $y \in R^N$ a vector containing the outputs y_k :

$$X = [x_1, \dots, x_N]^T, Y = [y_1, \dots, y_N]^T. \tag{10}$$

4.2 Least-Squares Estimation of Consequents

The defuzzification formulas of the singleton and TS models are linear in the consequent parameters, a_i, b_i as shown in equation (2). Hence, these parameters can be estimated from the available data by least-squares techniques. Denote $T_i \in R^{N \times N}$ the diagonal matrix having the normalized membership degree $\gamma_i(x_k)$ of (4) as its k th diagonal element. By appending a unitary column to X , the extended matrix $X_e = [X, 1]$ is created. Further, denote X' the matrix in $R^{N \times KN}$ composed of the products of matrices T_i and X_e :

$$X' = [T_1 X_e, T_2 X_e, \dots, T_K X_e]. \tag{11}$$

The consequent parameters a_i and b_i are lumped into a single parameter vector $\theta \in R^{K(p+1)}$:


$$\theta = [a_1^T, b_1, a_2^T, b_2, \dots, a_K^T, b_K]^T \tag{12}$$

Given the data X, y , equation (3) now can be written in a matrix form, $y = X' \theta + e$. It is well known that this set of equations can be solved for the parameter θ by:


$$\theta = [(X')^T X']^{-1} (X')^T y. \tag{13}$$

This is an optimal least-squares solution, which gives the minimal prediction error, and as such is suitable for prediction models. At the same time, however, it may bias the estimates of the consequent parameters as parameters of local models. If an accurate estimate of local model parameters is desired, a weighted least-squares approach applied per rule may be used:


$$[a_i^T, b_i]^T = [X_e^T T_i X_e]^{-1} X_e^T T_i y. \tag{14}$$




(a) Membership function



(b) Validation




(b) Estimated consequent parameters



(c) System and model

Figure 1. (a) Membership function designed for the output $y(k)$; (b) System nonlinearity (solid line) and its approx. in terms of the estimated consequent parameters (dashed line).

Figure 2. (a) Identification of data set : (b) performance of the model on a validation data set: (c) System (Solid line), model (dashed-dotted line)



(a) Identification of data set




Figure 3. Surfaces of original & identified model

5 Example

We considered a nonlinear dynamic system described by a first-order difference equation:

$$y(k+1) = y(k) + u(k)e^{-3[y(k)]} \quad (15)$$

We use a stepwise input signal to generate with this equation a set of 300 input-output data pairs as shown in Fig. 2(a). We also considered that it is known that the system is first order and that the nonlinearity of the system is only caused by y , the TS rule structure is:

$$\begin{aligned} \text{If } y(k) \text{ is } A_i \\ \text{then } y(k+1) = a_i y(k) + b_i u(k) \end{aligned} \quad (16)$$

Assuming that no further prior knowledge is available, seven equally spaced exponential membership functions, are defined in the domain of $y(k)$, as shown in Fig. 1(a). The consequent parameters are estimated by the least-squares method. Figure 1(b) shows a plot of the parameters a_i, b_i against the cores of the antecedent fuzzy sets A_i . The interpolation between a_i and b_i is linear, since the membership functions are piece-wise linear (exponential). It is observed that the dependence of the consequent parameters on the antecedent variable approximates quite accurately the system's nonlinearity, which gives the model a certain transparency. Validation of the model in simulation using a different data set is presented in Fig. 2(b), 2(c). Surfaces of original & identified models are shown in Fig. 3.

6 Summery

In this paper, fuzzy modeling and identification of a dynamic system with NARX model is presented and discussed. The presentation identifies the fuzzy rules describing the system behavior from input-output data using TS models. In order to show the procedure performance, the fuzzy modeling & identification of a dynamic system is presented with an example. The simulation of this dynamic system is also reported using MATLAB. If no knowledge is available as to which variables cause the nonlinearity of the system, all the antecedent variables are usually partitioned uniformly. In order to obtain an efficient representation with as few rules as possible, the membership functions must be placed such that they

capture the non-uniform behavior of the system. This dynamic fuzzy modeling and identification can be used for various aims: analysis, design, control, monitoring, supervision, etc. Modeling of complex systems will always remain an interactive approach. Intuition and experience of the team will always play a major role in this process. Special software tools will also need to be developed for this purpose.

7 References

1. Robert Babuska, "Fuzzy Modeling for Control", *Kluwer Academic Publishers*, ISBN 0-7923-8154-8 (1998)
2. Takagi T., Sugeno M., "Fuzzy Identification of Systems and Its Application to Modeling and Control", *IEEE Trans. Systems, Man and Cybernetics*, 15(1), 116-131, 1985
3. Sugeno M., Yasukawa T., "A Fuzzy-Logic-Based Approach to Qualitative Modeling", *IEEE Transaction on Fuzzy Systems*, Vol. 1(1), 7-31, 1993
4. Laukonen E.G., Passino K.M., Krishnaswami V., Luh G.C., Rizzoni G., "Fault Detection and Isolation for an Experiment Internal Combustion Engine via Fuzzy Identification", *IEEE Trans. On Control Systems Technology*, 3(3), 347-355, 1995
5. Hadjili M., Lendasse A., Wertz V., Yukovich S., "Identification of Fuzzy Models for a Glass Furnace Process", *Proceeding of the IEEE International Conference on Control Applications*, Trieste, Italy, 964-968, 1998.

On the analysis of robot behavior and its contribution to theoretical cognitive science

Tom Ziemke

Dept. of Computer Science, University of Skövde

Box 408, SE -54128 Skövde, Sweden

tom@ida.his.se

Abstract

This paper discusses how neuro-robotic simulation studies of “minimally cognitive behavior” can contribute to the development and validation of cognitive-scientific theories. An example is provided of an analysis of the behavior of a simulated robot solving delayed response tasks. Combining qualitative and quantitative descriptions, based on detailed correlation of sensory inputs, motor outputs, and neural activation values as well as synaptic connection weights at each point in time, the analysis allows us to understand the interaction between minimal internal mechanisms and the behavioral dynamics they exploit.

1 Introduction

As several authors have pointed out, mobile robotics, or autonomous agents research in general, can be viewed from at least two different, though intertwined perspectives: that of engineering, mostly concerned with the design of artefacts, and that of science, mostly concerned with the understanding of natural systems. Furthermore the latter can of course be broken down according to exactly which science it is that is contributed to: e.g. cognitive science (e.g. Pfeifer, 1995), biology (e.g. Webb, 2000), etc.

While these distinctions appear fairly obvious, they seem to receive surprisingly little attention in discussions of autonomous robotics methodology, where general statements such as “simulations are useless” or “Khepera robots are not real robots” or “existence proofs just don’t do it” often can be heard. The last point, for example, is certainly true from an engineering point of view, whereas in cognitive science existence proofs can be highly valuable in the development of theories. Similarly, simulations might have limited value in robot engineering, and might only be second choice in the modeling of animal behavior in cases where a real robot can made to interact with (roughly) the same physical environment as the modeled animal (e.g. Lambrinos et al, 1997; Webb, 2000). From the perspective of cognitive science, however, robot simulations are very useful tools in agent-

based modeling (e.g. Schlesinger & Parisi, 2001), paying more attention to the interaction of agents and environments than traditional computational cognitive modeling of mostly internal processes.

Much of our own experimental work has been concerned with what Beer (1996) referred to as “*minimally cognitive behavior*”, i.e. “the simplest behavior that raises cognitive interesting issues”, and the detailed analysis of the mechanisms underlying such behavior (e.g. Biró & Ziemke, 1998; Ziemke, 1999, 2000; Ziemke & Thieme, in press). This paper provides an example of how we combine qualitative and quantitative descriptions, based on a detailed correlation of sensory inputs, motor outputs, resulting robot behavior/motion, and neural activation values as well as synaptic connection weights at each point in time. This allows us to analyze the interaction between minimal internal mechanisms and the environmental/behavioral dynamics they exploit.

2 An Example¹

In the example discussed in this paper a simple ANN-controlled robotic agent has been evolved to solve a variation of standard *delayed response tasks* as they are commonly used in experimental psychology and neuroscientific studies of short-

¹ Parts of Section 2 appear in similar form in Ziemke & Thieme (in press).

term memory (STM) (cf. Ziemke & Thieme, in press). The experiment and results documented here are part of larger set of experiments with different architectures and tasks (Thieme, 2002; Thieme & Ziemke, 2002; Ziemke & Thieme, in press). The following subsections first briefly overview task, agent, and control system, and then present an analysis of successful behaviors and their underlying mechanisms. For further details on this example the reader is referred to Ziemke and Thieme (in press).

2.1 Task

The agent's task, a type of delayed response task, is illustrated in Figure 1. Its starting position and goal are indicated by an arrow and a circle, respectively. Two lights appear in the first corridor, and the sides on which they appear indicate the way to the goal: the first one indicates the *correct* turning direction in the first junction, the second one indicates the *wrong* turning direction in the second junction.² Exact starting position, orientation and light locations vary randomly; goal locations vary accordingly. It should be noted that Figure 1 only shows one particular case as an example of the general setup; other examples can be found in Figure 4.




Figure 1: An example situation in a double T-maze. Arrow and circle indicate starting position and goal respectively.

2.2 Agent

The experiments discussed in this paper have been carried out with the simulated Khepera robot (Mondada, Franzi & Ienne, 1994) shown in Figure 2, using an extended version of Michel's (1996) Khepera simulator. As part of the simulator's functionality, random noise was added to simulated

sensor and motor values as follows: $\pm 10\%$ to distance sensor values, $\pm 5\%$ to light sensor values, $\pm 10\%$ to motor speed amplitudes, and $\pm 5\%$ to the direction of motion resulting from the speed differences (Michel, 1996).




Figure 2: Simulated Khepera robot. From Ziemke & Thieme (in press).

In Figure 2, the labels beginning with 'D' refer to the infrared distance/proximity sensors: back left (DBL), left-left (DLL), left (DL), front left (DFL), front right (DFR), right (DR), right-right (DRR), and back right (DBR). The figure also shows left and right wheel/motor (ML and MR), which are controlled independently, as well as the left and right ambient light sensors (LL and LR) used in the experiments.

2.3 Neurocontroller

The control architecture used here is fairly simple, although it might look complicated in Figure 3. The extended sequential cascaded network (ESCN), a variation of Pollack's (1991) SCN, consists of:

- a simple feed-forward *function network*, mapping sensory inputs to motor outputs and state units through a single layer of synaptic connection weights, and
- a second simple network, the *context network*, which maps state units to new function network weights through a single layer of weights, *only* when the so-called decision unit is activated (above a certain threshold).

In simpler terms, the function net constitutes a simple, reactive sensorimotor mapping, and the context net provides a mechanism for changing that sensorimotor mapping when necessary. In many cases, as demonstrated in the next subsection, these ESCN controllers are relatively simple to analyze, since they combine simple reactive mechanisms with occasional switches between different sets of reactive behaviors.

² This is the most difficult of six different setups/tasks discussed by Ziemke and Thieme (in press).




Figure 3: Extended sequential cascaded network (ESCN) as a robot controller. Solid arrows indicate fully connected layers of synaptic connection weights. The dotted arrow indicates the selectivity of feedback/modulation, depending on the activation of the decision unit. From Ziemke & Thieme (in press). See text for a detailed explanation.

More specifically, given an input vector $i_j(t)$, $j = 1 \dots n$, state unit vector $s(t)$ and output vector $o(t)$ are calculated as follows by the *function network*:³

$$o(t) = f(W_o(t) \cdot (i_1(t), \dots, i_n(t), 1))$$

$$s(t) = f(W_s(t) \cdot (i_1(t), \dots, i_n(t), 1))$$

Where f is the logistic activation function, and $W_o(t)$ and $W_s(t)$, together referred to as *function network weights*, are two-dimensional connection weight matrices dynamically computed by the *context network*. In the ESCN the use of the context net is conditional upon the activation of the *decision unit* in the ESCN (cf. Figure 3). The idea behind this is that the robot should be able to make use of memory/feedback selectively and decide actively when to change its sensorimotor mapping. More specifically, given the state vector $s_k(t)$, $k = 1 \dots m$, the decision unit activation $d(t)$ and the function network weight matrices $W_o(t)$ and $W_s(t)$, are dynamically computed in every time step t as follows by the ESCN's *context network*:

$$d(t) = f(V_d \cdot (s_1(t), \dots, s_m(t), 1))$$

if $d(t) \geq 0.5$

$$\text{then } W_o(t+1) = V_o \cdot (s_1(t), \dots, s_m(t), 1)$$

$$\text{else } W_o(t+1) = W_o(t)$$

if $d(t) \geq 0.5$

$$\text{then } W_s(t+1) = V_s \cdot (s_1(t), \dots, s_m(t), 1)$$

$$\text{else } W_s(t+1) = W_s(t)$$

Where f is the logistic activation function, V_d is a one-dimensional connection weight matrix, mapping current state to decision unit activation, and V_o and V_s are two-dimensional connection weight matrices mapping the current internal state $s(t)$ to the next time step's function network weights, if the decision unit is active, i.e. has an activation value of at least 0.5.

The control networks were trained with an evolutionary algorithm the details of which are omitted here since they are irrelevant to the analysis of successful controllers in the next subsection. For the details of the training procedure the reader is referred to Ziemke and Thieme (in press).

2.4 Behavior Analysis

The successful behavior of ESCN-controlled agents in four different scenarios, corresponding to different light combinations (left-left, left right, etc.), is illustrated in Figure 4. The figure shows the robot's position as well as all sensor inputs, motor outputs and neural activation values in each time step which together with the sensorimotor (i.e. function network) weights shown in Figure 5 allow a detailed analysis of what happens at each point in time.

The first thing to note in this case is that the use of the decision unit (labeled 'SEL' in Figure 4), and consequently the use of the context network for modulation of sensorimotor weights in the function network, was highly selective and only occurred once in each scenario. This makes the underlying mechanisms relatively easy to analyze, since the robot is basically controlled by a simple reactive controller at each point in time, except that "memory" is realized in each scenario through a single 'switch' from one set of reactive behaviors to another.

³ Output and state units have biases (or bias weights), therefore 1 is appended to the input vector in these equations.




Figure 4. ESCN-controlled agents' behavior and activation values in four different double T-maze scenarios. Empty circles indicate goal locations, whereas gray/striped circles indicate areas in which the agent 'dies' immediately. C0/C1 denote the context/state units, SEL the decision unit, for other abbreviations see Figure 3. Activation values are illustrated as vertical black lines whose height corresponds to the represented value (for each time step). All unit activation values lie between 0 (black dot) and 1 (full height black line). A filled circle for each time step indicates the agent's position. To simplify the matching of the agent's position/behavior and its network activation values the trajectories and the time lines have been segmented into five equally long time segments (labeled 0 to 4). From Ziemke & Thieme (in press).

The ESCN's sensorimotor (or function network) connection weight settings are illustrated in Figure 5. The initial weight configuration (cf. Figure 5a) was such that in the absence of other stimuli, the approximately equally large positive motor biases resulted in straight forward motion. In the cases where the first light source was on the left side (cf. Figures 4a and 4b) the agent initiated a right-circling behavior (facilitated through the large negative weight between the left light sensor (LL)

and the right motor (MR), combined with a weaker inhibitory connection between LL and ML). This temporarily actually made the agent move away from the first junction. When facing the wall in the dead end this led to activation of the left frontal proximity sensor (DFL), which in turn briefly activated context unit C1 (through the large positive weight between DFL and C1; cf. Figure 15a), which in turn activated the decision unit.




Figure 5: Hinton diagrams for the ESCN sensorimotor (i.e. function network) synaptic connection weights. Weights are illustrated as black and white squares (black for negative/inhibitory weights, white for positive/excitatory weights), whose size reflects the weight magnitude. The columns labeled '1.0' contain the biases (or 'bias weights'). C0 and C1 values indicate the context unit activation values that led to the respective weight setting. For other abbreviations see Figure 2. From Ziemke & Thieme (in press).

As a result the sensorimotor mapping was re-set to the one shown in Figure 5b, which realized a left-hand wall-following behavior. This was facilitated through a large positive bias for the right motor (MR) and a small negative bias for the left one (ML), giving the agent a tendency to turn left in the absence of stimuli, combined with excitatory connections between the proximity sensors on the left (DLL, DL, DFL) and the left motor. Hence, the agent would move forward as long as there was a wall to the left, and turn left as soon as that wall 'disappeared' in the junction. In the cases illustrated in Figures 4a and 4b this took the agent all the way to the goal, in both cases more or less unaffected by the second light source on the left and right respectively.

In the case illustrated in Figure 4c the agent encountered the first light source on its right hand side, which also activated context unit C0 temporarily (through the large positive weight between left light sensor LL and C0; cf. Figure 5a),

but the behavior was not influenced at all. When sensing the second light source to the left, the agent, as in the above cases, turned around to temporarily move away from the first junction. During that turn the agent sensed the first light source again with the right light sensor (LR) which also activated context unit C0 again. Now, however, context unit C1 was activated at the same time, when facing the wall during the turn (through the large positive weight between left frontal proximity sensor DFL and C1; cf. Figure 5a). The joint activation of both context units also triggered the decision unit (SEL), which resulted in the new sensorimotor weight setting illustrated in Figure 5c. This sensorimotor mapping made the agent move forward by default (through two positive motor biases), and turn right in junctions (both frontal proximity sensors, DFL and DFR, have excitatory connections to the left motor (ML) and inhibitory connection to the right one (MR); cf. Figure 5c). In this case these behaviors took the agent all the way to the goal, as illustrated in Figure 4c.

In the case where both light sources were placed on the right side, as illustrated in Figure 4d, the agent never turned away from the first junction. But as in the cases illustrated in Figures 4a and 4b, the first time left front sensor DFL became activated, context unit C1 was activated temporarily as well. This in turn triggered the decision unit (SEL), which resulted in the sensorimotor mapping illustrated in Figure 4b, which even in this case took the agent to the goal through left-hand wall-following. Hence, the first two scenarios and the last one were all solved using the same combination of behaviors/sensorimotor mappings, but in the latter case the switch did not occur until after (or in) the first junction.

As mentioned briefly above, it should be noted that in all four cases the decision unit became active only once, i.e. the sensorimotor mapping was adapted only once. Hence, the agent is basically controlled by purely reactive sensorimotor mappings all the time, except that at one point there is a switch from one reactive mechanism to another. It should also be noted that at all other points in time the context unit activation values do not influence the sensorimotor mapping at all. That means, STM is here not realized through the sustenance of neuronal activation patterns, but only through the dynamic modulation of reactive sensorimotor mappings, as embodied in the sensorimotor connection weights. Consequently, in three out of four cases the modulation was actually

not triggered by the light stimuli. Hence, the new sensorimotor weights correspond to the behavior that will take the agent to the goal, but they could hardly be said to 'represent' the light stimuli in the traditional sense of the term (as adopted in most studies of the neural mechanisms underlying STM; cf. Ziemke & Thieme, in press).

3 Discussion

As mentioned in the introduction, much of our experimental work is concerned with neuro-robotic studies of "minimally cognitive behavior" and the analysis of the underlying mechanisms. This paper provides an example of how qualitative and quantitative descriptions can be combined in the analysis of robot behavior, based on a correlation of sensory inputs, motor outputs, and neural activation values as well as synaptic connection weights at each point in time. This allows us to analyze in detail the interaction between minimal internal mechanisms and the environmental and behavioral dynamics they exploit.

Simulation studies of the type discussed here certainly have many flaws from an engineering point of view, and typically they cannot directly be transferred to real robots. They can, however, make valuable contributions to cognitive scientific theories and discussions, not least due to the fact that they are relatively easy to analyze in detail. For example, the analyses presented here illustrate how STM can be realized through synaptic plasticity. Thus they provide an alternative to more traditional ANN and neuroscientific models of STM which usually assume sustenance of neural activation patterns as the (main) underlying mechanism. Furthermore, analyses of this type allow making cognitive-scientific discussions of issues such as "representation" more concrete than they usually are, by complementing philosophical and psychological arguments with concrete case studies of "minimally cognitive behavior" and the mechanisms underlying it. Finally, it should of course be mentioned that agent-based models of this type also have obvious limitations. Schlesinger and Parisi (2001), for example, pointed out that such models have similar problems with "brittleness" as their more traditional, computational counterparts, i.e. solutions are usually highly task- and context-specific.

In sum, despite some obvious limitations, simulation studies of situated robotic agents allow extensive, systematic variations of environments

and agents, as well as detailed observation and analysis of underlying mechanisms, and thus they constitute a useful complement to possibly more realistic, physical robot experiments as well as theoretical discussions.

Acknowledgements

The experiments analyzed in this paper have been carried out by Mikael Thieme (2002). Furthermore, this paper benefited from valuable comments that Ezequiel Di Paolo and four anonymous reviewers had on an earlier paper (Ziemke & Thieme, in press). This work has been supported by a grant (1507/97) from the Knowledge Foundation, Stockholm, Sweden.

References

- Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In *From Animals to Animats 4*, pp. 421-429. Cambridge, MA: MIT Press.
- Biró, Z. & Ziemke, T. (1998). Evolving visually guided approach behaviour in recurrent artificial neural network robot controllers. In *From animals to animats 5*, pp. 73-76. Cambridge, MA: MIT Press.
- Lambrinos, D., Maris, M., Kobayashi, H., Labhart, T., Pfeifer, R., & Wehner, R. (1997). An autonomous agent navigating with a polarized light compass. *Adaptive Behavior*, 6(1), 131-161.
- Michel, O. (1996) Khepera simulator package version 2.0. Freeware mobile robot simulator. (<http://diwww.epfl.ch/lami/team/michel/khep-sim/>).
- Mondada, F., Franzi, E. & lenne, P. (1994). Mobile robot miniaturization: A tool for investigation in control algorithms. In Yoshikawa, T. & Miyazaki, F. (eds.) *Proceedings of the Third International Symposium on Experimental Robotics*, pp. 501-513). Berlin: Springer Verlag.
- Pfeifer, R. (1995). Cognition – Perspectives from autonomous agents. *Robotics and Autonomous Systems*, 15(1-2), 47-70.
- Schlesinger, M. & Parisi, D. (2001). The agent-based approach: A new direction for computational models of development. *Developmental Review*, 21, 121-146.

- Thieme, M. (2002). Intelligence without hesitation. Masters Dissertation HS-IDA-MD-02-001, Dept. of Computer Science, University of Skövde, Sweden.
- Thieme, M. & Ziemke, T. (2002). The Road Sign Problem Revisited: Handling Delayed Response Tasks with Neural Robot Controllers. In *From animals to animats 7*, pp. 228-229. Cambridge, MA: MIT Press.
- Webb, B. (2000). What does robotics offer animal behavior? *Animal Behavior*, 60, 545-558.
- Ziemke, T. (1999). Remembering how to behave: Recurrent neural networks for adaptive robot behavior. In Medsker & Jain (eds.) *Recurrent Neural Networks: Design and Applications*, pp. 355-389. New York: CRC Press.
- Ziemke, T. (2000). On 'Parts' and 'Wholes' of Adaptive Behavior: Functional Modularity and Diachronic Structure in Recurrent Neural Robot Controllers. In *From animals to animats 6*, pp. 115-124. Cambridge, MA: MIT Press.
- Ziemke, T. & Thieme, M. (in press). Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response tasks. *Adaptive Behavior*, 10(3/4), to appear.

Towards a methodology for embodied computational neuroscience

Tony J. Prescott, Fernando M. Gonzalez, M. Humphries, and Kevin Gurney

Adaptive Behaviour Research Group,

Department of Psychology,

University of Sheffield.

t.j.prescott@shef.ac.uk

Introduction

Two of the general goals of our research group are to use robotics in order to further the understanding of animal nervous systems—*embodied computational neuroscience* if you will—and to endow robots with some of the useful control capabilities of animal nervous systems—a brand of *biorobots*. Both of these activities require a research programme that bridges the substantial gap between neurobiology and robot control. The research activity of our group therefore includes, to some degree at least, each of the following elements:

- (i) empirical and theoretical research in the neurosciences,
- (ii) computer simulation of target neural systems at several levels of abstraction,
- (iii) embedded (robotic) models of target neural systems,
- (iv) the evaluation of candidate biomimetic robot control mechanisms with respect to biological constraints (embodied computational neuroscience) and engineering requirements (biorobotics).

The current paper is specifically concerned with robotics research in categories (iii) and (iv), however, as this builds upon the past and ongoing work by our group in more conventional computational neuroscience it is worth briefly describing our style of computational neuroscience research to show the context in which the methodology of our robot research has been developed.

Computational Neuroscience

Research in the neurosciences provides a mass of detailed descriptions of interesting phenomena concerning target neural structures, but gives little guidance as to the correct functional understanding of this data. To make progress requires computational hypotheses that can be formulated as simulation models and whose outputs can be tested against available neurobiological findings. In recent years, our group has developed several computational models^{1, 2} of a group of central structures in the vertebrate brain called the *basal ganglia* based on the hypothesis that these nuclei play an important functional role in vertebrate *action selection*^{3, 4}. Figure 1 shows the main basal ganglia nuclei and some of their intrinsic and extrinsic connections within the mammalian brain. The principle structures are the *striatum* and *pallidum* found at the base of the forebrain and the *substantia nigra* in the midbrain. These structures appear to have homologues in the nervous systems of all classes of jawed vertebrates and possibly in all vertebrates.




Figure 1: Principle structures of the basal ganglia and some of their intrinsic and extrinsic connections in the mammalian brain. Abbreviations: SN—substantia nigra, STN—subthalamic nucleus.

The proposal that the basal ganglia performs action selection in the vertebrate brain is not a radical perspective on their function but rather derives from a growing consensus that a key function of these structures is to enable desired actions and to inhibit undesired, potentially competing, actions (for review see refs 3 and 4). Key elements of this hypothesis include that (i) the input structures of the basal ganglia (principally the striatum) receive convergent input from many regions of the brain indicating the urgency (or *saliency*) of competing ‘bids’ for motor resources, (ii) that the continuous inhibitory output from the substantia nigra projects to motor systems throughout the brain and therefore serves as a break on all voluntary movement, (iii) that action selection occurs when a ‘winning’ bid in the striatum generates inhibition of neurons in substantia nigra thereby causing them to release their inhibition of target motor systems (thus action selection through *disinhibition*).

The structures of the basal ganglia are highly interconnected with feedback loops both between and within the different nuclei. To achieve a deep insight into basal ganglia function therefore demands a computational approach involving quantitatively specified models, delimited by anatomical and physiological constraints. Our computational neuroscience research on this topic includes both relatively low-level models, focusing on some of the properties of identified populations of neurons, and ‘systems-level’ models of basal ganglia nuclei and their inter-connections. This research has resulted in computational hypotheses concerning the function of specific basal ganglia structures and pathways (see ref. 1, 2), and has produced some simulation results that show a good match with neurobiological data. For instance, as shown in figure 2 below, we are able to generate signal characteristics in the component circuits of our basal ganglia model (the insets in the top right of each graph) that follow similar temporal patterns to electrophysiological recordings of neural firing (larger graphs- a, b, c) in corresponding nuclei in the rat brain.




Figure 2. Model output compared with recordings in the globus pallidus (aka pallidum) (a, b)⁵ and substantia nigra (c)⁶.

Embodied Computational Neuroscience

The above work provides the basis on which we have begun to develop embodied models of these same neural circuits, and to evaluate the usefulness of these mechanisms for robot control^{2, 7}. Embodiment should not only provide a strong test of the sufficiency of our model of the basal ganglia (since it brings with it all the demands of real-world, real-time behaviour), but should also allow us to make contact with a whole new tier of neurobehavioural data. Having developed candidate robot controllers with biomimetic elements our current approach is to evaluate their behaviour on a benchmark task loosely modelled on some of the activities of a hungry laboratory rat placed in a novel arena. Benchmark tasks are vital if one control architecture is to be effectively compared/contrasted with another, and our hope is that selecting a benchmark that approximates a model biological situation will allow comparisons with neuroethological observations of behaving animals.

Figure 3 shows a screen-shot of the activity of our current robot model. The bottom-right corner of the figure shows the current location and activity of the robot as captured by a video camera, while other windows show the current internal activity of the model. In our benchmark task the robot selects between behaviours related either to object acquisition (collecting cylinders from the ‘open field’ and depositing them in the corners of the arena) or to object-carrying and threat avoidance (avoiding open spaces). There are five behaviours in all—*cylinder seek* (1), *cylinder pickup* (2), *wall seek* (3), *corner seek* (4), and *cylinder deposit* (5).




Figure 3. The benchmark task for evaluating the action selection architecture. In the figure, the current saliencies are shown as histograms in the centre-top window, larger bars indicating higher saliency. Each saliency value is computed by the model basal ganglia as a function of several inputs. These include: (bottom-left window) inputs from sensory systems that detect the presence or absence of an object in the robot gripper, or of a nearby corner, wall, or cylinder; (top-left window) artificial motivations simulating levels of 'fear' (dark-grey line) and 'hunger' (light-grey line); (centre-left window) positive feedback, or other direct input, from the current selected behaviour (shown by a darkened bar).

Figure 4 shows the role of the embedded basal ganglia model in the robot controller for this benchmark task. Inputs from the perceptual, motivational, and action sub-systems converge on the basal ganglia input structures. Intrinsic basal ganglia circuitry then resolves the competition between the action sub-systems and disinhibits the motor outputs (to the gripper and wheel motors) for the winning sub-system. The basal ganglia also provides a positive feedback signal, via the *thalamus*, to the winning competitor.




Figure 4. Architecture of the embodied basal ganglia model. Note that only a small sub-set of actual connections are shown.

In biological research a standard technique for analysing and displaying behavioural data is the *ethogram*. The graph in Figure 5 below shows an ethogram of the activity of our robot in an example run, indicating which of the five available behaviours (or a sixth ‘no activity’ behaviour) is selected at any given moment in time. In this example the initial behaviour is driven by fear of open spaces—*wall seek* followed by *corner seek*. This is followed by four successful bouts of foraging behaviour—*cylinder seek*, *cylinder pickup*, *wall seek* (this time carrying a cylinder), *corner seek* (also cylinder-carrying), and *cylinder deposit*.




Figure 5. Ethogram of robot behaviour in a single run for the benchmark task.

Although this model task is a very straightforward one from a robotics perspective it contains many elements that can test the effectiveness of candidate action selection mechanisms. These include multiple competing behaviours; varying and noisy sensory data; motivational drives that vary with time, circumstances, and behaviour; and complex sequential behaviours that permit a secondary, and finer-grained, decomposition of action selection. By using this benchmark we therefore expect to be able to make direct comparisons between competing models. However, we also expect to formulate more demanding benchmark tests, to establish the versatility of the biomimetic controllers we develop.

Evaluating candidate biomimetic control systems for action selection

We have previously described^{3, 4} several criteria for an effective action selection mechanism in addition to the primary requirement of selecting appropriate behaviour for any given circumstance. Specifically, an effective selection mechanisms should provide clean and rapid switching between competing behaviours; lack of distortion (by losing competitors) of the execution of a selected behaviour; appropriate behavioural persistence or *maintenance* (the behaviour should not be terminated too early or too late); and the flexibility to allow re-prioritisation when circumstances change unexpectedly. We have demonstrated analytically, and in computer simulation, that our prototype model of the vertebrate basal ganglia possesses suitable control properties to serve as an effective action selection device. In our current work we plan to conduct a detailed analysis of synchronised records of the behaviour of the robot (in video) and the intrinsic activity of its control system, in order to evaluate the effectiveness of candidate action selection mechanisms according to specific action selection metrics. In pilot studies with our model we have also begun to investigate similarities between robot action selection, and video recordings of behaviour switching in laboratory rats. This work has found disruptive effects on robot behaviour of altered levels of the simulated neuromodulator *dopamine* that show parallels with similarly-treated animals, and could have implications for understanding brain disorders that involve dopamine such as Parkinson’s Disease. In the following we summarise two of the interesting results from our research to date, that have involved developing novel metrics for measuring/classifying robot behaviour.

Behavioural maintenance

An effective action selection system must be able to select a behaviour that is appropriate to the current circumstances. However, a further key issue concerns what happens once a behaviour has become selected and the animal/robot is engaged in performing it. We can identify two specific requirements for action selection with regard to an ongoing behaviour. First the behaviour should be maintained for an appropriate length of time or until completion. Failure to do so will lead to an error of behaviour maintenance that we term *early interrupt*. Second, the behaviour should be terminated quickly and cleanly once it ceases to be appropriate. Failure to achieve this can lead to a second type of maintenance error termed *perseveration* (a term frequently used in ethological experiments with behaving animals). A specific situation in which errors of behaviour maintenance can arise is where the salience for the behaviour falls after the time it is first selected but before it has been effectively completed. For instance, as an animals feeds, the energy deficit that caused it to be hungry will

reduce leading to potential inefficiencies in behaviour maintenance—the animal may stop feeding before its belly is full. A second situation in which maintenance errors occur is where a behaviour involves a sequence of actions that must be completed in order to maintain the physical integrity or safety of the animal, but where the original motivation for the behaviour is satisfied part-way through the sequence. For instance, a grooming behaviour could involve licking an itchy limb. Once the itch is gone, so is the motivation for grooming, however, the animal must still retract its tongue and rearrange its limbs/head into a stable posture. In our robot model we have examined a similar type of situation in the *cylinder deposit* task. Here the robot must lower its gripper arm, open the gripper to release the cylinder (an event which leads to a big drop in the salience for cylinder-deposit), and then close the gripper and return the gripper arm to a safe, vertical position. We have found that the two types of maintenance error can occur here: early interrupt—the robot switches to a new behaviour (e.g. looking for another cylinder) before resetting its gripper/arm to the safe position; and perseveration—the robot repeats the cycle of arm-lowering, gripper-opening, gripper-closing, arm-lifting after it has already released the cylinder.

The problem of behaviour maintenance has, of course, been the subject of considerable research. Much of this has focused on the use of *positive feedback* to maintain a selection. Here the selection/output of a behaviour leads, via a feedback loop, to a boost in the salience for that behaviour. Positive feedback can be used to organise behaviours such as feeding and drinking into efficient bout lengths⁸, and is incorporated into our basal ganglia model via a loop that returns to the action sub-systems via the *thalamus* (see figure 4). However, we have discovered through our robot model that positive feedback is a relatively crude mechanism for providing accurate timing of the maintenance and completion of a behaviour such as *cylinder deposit*. The level of feedback has to be set just right to avoid both early interrupt and perseveration errors, and variation in the salience of the selected behaviour, or that of its competitors, easily upsets this balance leading to maintenance errors of either type. To avoid this problem we have discovered that there is a need for a behaviour sub-system, such as *cylinder deposit*, to be able to make a direct contribution to its own salience at critical moments of behaviour execution. Thus after releasing a cylinder, the *cylinder deposit* sub-system sends a signal of a specified size and duration to the action selection system that boosts its own salience sufficiently for the behaviour to be carried through to completion. Since the behaviour sub-system has control over the size and timing of this '*busy signal*' it can provide a much better mechanism for effective behaviour maintenance than positive feedback alone.

In the graphs below we illustrate the effects of using three different levels of positive feedback, or the combination of positive feedback and a '*busy signal*', on maintenance errors in the cylinder deposit behaviour. In top half of figure 6 the salience of the behaviour is subject to minor variation due to sensor noise, whilst in the bottom half there is both sensor noise and major variation in salience signals caused by large fluctuations in the two '*motivational*' variables intended to loosely model varying levels of '*hunger*' or '*fearfulness*' in a rat.




Figure 6. Comparison of outcomes for cylinder deposit behaviour with different levels of 'busy signal' and positive feedback (+f) in a model with *fixed* motivations (top), and with *varying* motivations (bottom).

When the motivational inputs are removed from the model (top graph) we see that cylinder-deposit is effectively maintained and completed (good switching) either where there is no busy signal with medium levels of positive feedback (+f), or where there is a 'busy signal' together with low positive feedback. In the two case where positive feedback is low or high, and there is no busy signal, this leads to either many errors of early interrupt (low +f) or of perseveration (high +f). In the bottom graph we see the effect of including in the model the large variations in salience caused by motivational variables. Here we see that positive feedback alone is rarely effective in producing good-switching, showing primarily early interrupt errors where the strength of the feedback is at low or medium levels, a mixture of early interrupt and perseveration errors at high positive feedback levels, and primarily perseveration errors at very high levels. In contrast the system has a much wider operating range when the 'busy signal' is engaged together with low positive feedback. Here most attempts at cylinder-deposit result in good switching and the few errors that occur happen when the motivational variables reach extreme values (for instance, errors of early interrupt typically occur when the level of simulated 'hunger' is particularly high).

Interestingly, electrophysiological studies have noted that activity changes in the basal ganglia often occur after a behaviour has been initiated. One interpretation of this post-selection activity is that it may reflect 'efference copy' of motor commands. The hypothesis generated from the research outlined above is that such signals to the BG may act to maintain selected actions until completion, in other words, that they may convey information similar to that encoded by the 'busy signal' in our model.

Behavioural effects of changes in baseline dopamine

Manipulations of tonic (baseline) dopamine levels in the basal ganglia are known to affect action selection (or behaviour switching). For instance, the dopamine blocker (antagonist) haloperidol can induce akinesia (no movement) or bradykinesia (slowed movement), whilst the dopamine enhancer (agonist) apomorphine can lead to stereotyped behaviour in which the animals continuously repeat a particular pattern of movement. In our robot model we have observed similar outcomes from changes in simulated dopamine to those seen in animal studies. First, lowering simulated dopamine below the normal baseline level causes slowed movement and difficulty in initiating behaviours. Second, an increase in simulated dopamine can cause the simultaneous selection of two behaviours resulting in an inappropriate mixture of two activity patterns; at the same time other candidate behaviours are usually excluded by a property of the basal ganglia model we term 'selection limiting' (this makes it more difficult for a behaviour to be selected when there are already active behaviours). Where the behavioural interference caused by dual selection prevents one of the active behaviours from completing, the robot can become stuck in a 'behavioural trap'. So for instance, the robot may become trapped in a cycle of repeatedly trying to pick-up a cylinder where success in this action is prevented by interference from a simultaneously active cylinder seek behaviour. This result suggests that one source of behavioural stereotypy in animals could be the failure to deselect or interrupt the activity of two simultaneously selected behaviours. More generally the robot's inability, in the excess-dopamine condition, to limit selection to one active behaviour has interesting similarities to the difficulty in suppressing inappropriate actions or thoughts that is characteristic of some patients with Tourette's syndrome or schizophrenia.

In order to more precisely characterize the errors observed in low and high dopamine conditions we have developed two specific action selection metrics. First, we quantify the *efficiency* with which a behaviour is selected by measuring how fully (0–100%) basal ganglia inhibition is removed from the motor system. Second, we quantify the *distortion* (0–100%) of the currently selected behaviour, due to other, partially selected, behaviours, as the ratio of the total efficiency of all 'losing' actions to the sum of efficiency scores for all possible actions (including the winning one). Typically, the low dopamine 'Parkinsonian' condition described above produces robot behaviour with little or no distortion but also low efficiency—winning behaviours are only partially disinhibited. Conversely, the high dopamine condition can be characterized as high efficiency, but relatively high distortion—winning behaviours are fully disinhibited but the partial expression of some losing behaviours may result in inappropriate actions.

Concluding comments

We hope that the better characterization of such failures and errors in action selection systems will help both to create stable metrics for comparing alternative selection mechanisms for robots, and to improve our understanding of the root causes of behavioural deficits in patients with basal ganglia-related disease. More generally, we hope that the research program we have embarked on, and have partially outlined above, will contribute towards a developing methodology for *embodied* computational neuroscience that will allow biomimetic robotics to make a significant contribution to the brain sciences.

Acknowledgements

This work is part-funded by EPSRC grant no. GR/R95722/01.

References

1. Gurney, K., Prescott, T.J. & Redgrave, P. Biological Cybernetics 84, 401-423 (2001).
2. Prescott, T.J., Montes Gonzalez, et al. The Basal Ganglia VII, Kluwer Academic, 349–356.
3. Redgrave, P., Prescott, T. & Gurney, K.N. Neuroscience 89, 1009-1023 (1999).
4. Prescott, T.J., Redgrave, P. & Gurney, K.N. Adaptive Behavior 7, 99-127 (1999).
5. Ryan, L.J. & Clark, K.B. Experimental Brain Research 86, 641-651 (1991).
6. Shultz, W. Journal of Neurophysiology 55, 660-677 (1986).
7. Montes Gonzalez, F., Prescott, T.J., et al. Proc. 6th Int. Conf. Simulation of Adaptive Behavior (ed. Meyer, J.A.) 157-166 (MIT Press, Cambridge, MA, 2000).
8. Houston, A. and Sumida, B. (1985). Animal Behaviour, 33(FEB), 315-325.