


AD-A255 809



MENTATION PAGE

Form Approved
GMB No. 0704-0188

Estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

REPORT DATE

1991

3. REPORT TYPE AND DATES COVERED

01 Sep 90 TO 31 Aug 91 FINAL

4. TITLE AND SUBTITLE

CONFERENCE ON SIMULATION OF ADAPTIVE BEHAVIOR : FROM
ANIMALS TO ANIMATS.

5. FUNDING NUMBERS

G: AFOSR -90-0313

PR: 2313/A9

61102F

②

6. AUTHOR(S)

JEAN-ARCADY MEYER & STEWART W. WILSON

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

ECOLE NORMALE SUPERIEURE (ENS). URA 686 CNRS
GROUPE DE BIOINFORMATIQUE - 46 rue d'Ulm
75230 PARIS cedex 05 FRANCE8. PERFORMING ORGANIZATION
REPORT NUMBER

AFOSR-90-0313

6

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Dr John F. Tangney
AFOSR / NL
BUILDING 410, BOLLING AFB DC 20332 -644810. SPONSORING / MONITORING
AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

Published by the MIT Press/ Bradford books

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release;
distribution is unlimited.

12b. DISTRIBUTION CODE

DTIC
ELECTE
OCT 01 1992
S A D

13. ABSTRACT (Maximum 200 words)

Sixty contributions from researchers in ethology, ecology, cybernetics, artificial intelligence, robotics and related fields delve into the behaviors and underlying mechanisms that allow animals and, potentially, robots to adapt and survive in uncertain environments. They focus in particular on simulation models in order to help characterize and compare various organizational principles on architectures capable of inducing adaptive behavior in real or artificial animals. This report includes the program and abstracts from the conference and the table of contents from the resulting book.

This document has been approved
for public release and sale; its
distribution is unlimited.

92 9 30 079

92-26267



425579

2448

14. SUBJECT TERMS

ETHOLOGY - ROBOTICS -

ANIMAL BEHAVIOR - SIMULATION METHODS - ANIMATS -

15. NUMBER OF PAGES

549

16. PRICE CODE

9 - 780262631389

17. SECURITY CLASSIFICATION
OF REPORT

(U)

18. SECURITY CLASSIFICATION
OF THIS PAGE

(U)

19. SECURITY CLASSIFICATION
OF ABSTRACT

(U)

20. LIMITATION OF ABSTRACT

UL

24 JAN 1992

FROM ANIMALS TO ANIMATS
Proceedings of the First International Conference on
Simulation of Adaptive Behavior

edited by Jean-Arcady Meyer and Stewart W. Wilson

A Bradford Book

The MIT Press
Cambridge, Massachusetts
London, England

PREFACE

An international conference entitled "Simulation of Adaptive Behavior: From Animals to Animats" took place in Paris on September 24-28, 1990. The object of the conference was to bring together researchers in ethology, ecology, cybernetics, artificial intelligence, robotics, and related fields so as to further our understanding of the behaviors and underlying mechanisms that allow animals and, potentially, robots to adapt and survive in uncertain environments.

"SAB90", as we called it, was the first major conference to test the hypothesis that people interested in understanding animal behavior, and people interested in simulating or constructing autonomous robots, would have important common interests and would welcome the chance to listen to and learn from each other. The conference further tested the somewhat more radical hypothesis that its focus constituted not only an intersection but a growing new *field* concerned, in both animals and "animats", with adaptive behavior.

By a variety of measures including size and international range of attendance, intellectual enthusiasm, quality and diversity of contributions, and degree of interaction among the participants, SAB90 offered strong support for both these hypotheses. Furthermore, the emergence of a field in its own right was signaled by the fact that while there was lively debate along many axes—e.g., top-down vs. bottom-up, learning vs. reflexes, hierarchical vs. flat, simulate vs. build, to mention a few—it was striking how people everywhere along the *animals to animats* axis were thinking about similar sets of problems.

These proceedings contain 62 papers, 59 that were actually presented at the conference, plus three whose authors could not attend. The book is divided into sections corresponding to the conference sessions. In each section, papers presented as talks are followed by related papers that were presented as posters.

The first section, The Animat Approach, contains papers on artificial animal research as a tool for understanding adaptive behavior and, indeed, as a new approach to artificial intelligence. The next sections—Perception and Motor Control, Cognitive Maps and Internal World Models, Motivation and Emotion, Action Selection and Behavioral Sequences, Ontology and Learning, Collective Behaviors, and Evolution of Behavior—contain papers on these themes from both the animal and animat perspectives. There follows a large section on Architectures, Organizational Principles, and Functional Approaches, containing several strong—and differing—theses on how to understand or achieve natural or artificial systems with adaptive behavior. The book concludes with a two-paper section, Animats in Education, that describes novel and uncomplicated robot and simulation technologies designed for teaching and research.

SAB90 could not have taken place without the assistance of many people and organizations. We are especially grateful to members of the Program Committee, whose conscientious reviewing selected the papers here from the more than 90 submitted, and who ably chaired the conference sessions. The Committee members were

Lashon Booker, MITRE Corporation, USA
 Rodney Brooks, MIT Artificial Intelligence Lab, USA
 Patrick Colgan, Queen's University at Kingston, Canada
 Patrick Greussay, Université Paris VIII, France
 David McFarland, Balliol College, Oxford, UK
 Luc Steels, VUB AI Lab, Belgium
 Richard Sutton, GTE Laboratories, USA
 Frederick Thates, The Open University, UK
 David Waltz, Thinking Machines Corp. and Brandeis University, USA

We thank each of the following sponsors of the conference, and mention particularly AFOSR which enabled us to provide substantial needed travel assistance, and the anonymous Corporate Donor whose early confidence in our project was a great boost.

Ecole Normale Supérieure
 The Rowland Institute for Science
 Ministère de la Recherche et de la Technologie
 AFOSR (U. S. Air Force Office of Scientific Research)
 C2V
 Electricité de France
 IBM France
 OFFILIB
 A Corporate Donor

The members of the Groupe de BioInformatique, URA 686, Ecole Normale Supérieure, worked long and hard to make the conference a success. We thank Anne Brelet, Eric Granjeon, Agnès Guillot, Jean-Louis Pennetier, Philippe Tarroux, and Pierre Vincens.

We wish to express our gratitude to the Ministère de la Recherche et de la Technologie for having generously placed at our disposal the Amphithéâtre Poincaré in which the conference sessions were held. We also express our particular thanks to Josiane Serre (Administrator of the ENS) and René Lafont (Director of URA 686) for their kind help in solving various administrative problems.

Finally, we are particularly indebted to Louis Bec and Wang Xueqing for the artistic conception of the SAB90 poster and the Proceedings cover.

We invite readers to enjoy and profit from the papers in this book, and look forward to the next SAB conference!

Jean-Arcady Meyer and Stewart W. Wilson
 Conference Co-Chairs

Accession For	
NTIS CP821	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution	
Availability	
Not	Available
A-1	

DTIC QUALITY INSPECTED 3

CONTENTS

<i>Preface</i>	ix
----------------------	----

THE ANIMAT APPROACH

Simulation of Adaptive Behavior in Animats: Review and Prospect	2
<i>Jean-Arcady Meyer and Agnes Guillot</i>	
The Animat Path to AI	15
<i>Stewart W. Wilson</i>	
What It Means for Robot Behaviour to be Adaptive	22
<i>David McFarland</i>	
Computational Neuroethology: A Provisional Manifesto	29
<i>Dave Cliff</i>	
On the Feasibility of Computational Artificial Life—A Reply to Critics	40
<i>Harold C. Morris</i>	
The Animat and the Physician	50
<i>Alexandre Parodi and Pierre Bonelli</i>	
Extended Classifiers for Simulation of Adaptive Behavior	58
<i>Renaud Dumeur</i>	

PERCEPTION AND MOTOR CONTROL

Biomimetic Sonar Processing: From Dolphin Echolocation to Artificial Neural Networks	66
<i>H. L. Roitblat, P. W. B. Moore, P. E. Nachtigall, and R. H. Penner</i>	
Hierarchical Dishabituation of Visual Discrimination in Toads	77
<i>DeLiang Wang and Michael A. Arbib</i>	
The Computational Hoverfly: A Study in Computational Neuroethology	87
<i>Dave Cliff</i>	
Can Allied Reflexes Promote the Integration of a Robot's Behavior?	97
<i>Philip Teitelbaum, Vivien C. Pellis, and Sergio M. Pellis</i>	
Coordination of Leg Movement in Walking Animals	105
<i>Holk Cruse</i>	
A Biological Visuo-Motor System: How Dissimilar Maps Interact to Produce Behavior	120
<i>Gerhard Manteuffel</i>	
Biological and Computational Stereo Vision	127
<i>Stephen T. Barnard</i>	
Modeling and Simulation of Animals' Movements	135
<i>Simon Benhamou and Pierre Bovet</i>	

COGNITIVE MAPS AND INTERNAL WORLD MODELS

Schemas for Prey-Catching in Frog and Toad.....	142
<i>Michael A. Arbib and Alberto Cobas</i>	
Mapbuilding Using Self-Organising Networks in "Really Useful Robots"	152
<i>Ulrich Nehmzow and Tim Smithers</i>	
The Contribution of Quantitative Models to The Long Distance Orientation Problems	160
<i>Marc Jamon</i>	
Navigating with a Rat Brain: A Neurobiologically-Inspired Model for Robot Spatial Representation	169
<i>Maja J. Mataric</i>	
Four Important Issues in Cognitive Mapping	176
<i>W. K. Yeap and C. C. Handley</i>	
Attracting Similar Shapes Towards Each Other	184
<i>Mitch R. Harris</i>	

MOTIVATION AND EMOTION

Ethological and Psychological Models of Motivation—Towards a Synthesis	194
<i>Frederick Toates and Per Jensen</i>	
Pleasure: The Answer to Conflicting Motivations	206
<i>Michel Cabanac</i>	
Machine Motivation.....	213
<i>J. R. P. Halperin</i>	
A Possibility for Implementing Curiosity and Boredom in Model-Building Neural Controllers.....	222
<i>Juergen Schmidhuber</i>	

ACTION SELECTION AND BEHAVIORAL SEQUENCES

Instinct as an Inductive Bias for Learning Behavioral Sequences.....	230
<i>Lashon B. Booker</i>	
A Bottom-Up Mechanism for Behavior Selection in an Artificial Creature	238
<i>Pattie Maes</i>	
The Neural Basis of Behavioral Choice in an Artificial Insect.....	247
<i>Randall D. Beer and Hillel J. Chiel</i>	
An Investigation of Two Mediation Strategies Suitable for Behavioural Control in Animals and Animats	255
<i>Martin Snaith and Owen Holland</i>	
Computer Simulation of an Animal Environment.....	263
<i>Toby Tyrrell and John E. W. Mayhew</i>	

ONTOGENY AND LEARNING

Ontogeny of Preferences in Guppies (<i>Poecilia reticulata</i>) Exposed to Food and Conspecifics....	274
<i>Patrick W. Colgan, T. Lynne Jamieson, Janice E. Frame, and J. Terry Smith</i>	
Simulation Studies of Song Learning in Birds.....	281
<i>James M. Williams and P. J. B. Slater</i>	
Reinforcement Learning Architectures for Animats.....	288
<i>Richard S. Sutton</i>	
Self-Improving Reactive Agents: Case Studies of Reinforcement Learning Frameworks.....	297
<i>Long-Ji Lin</i>	
Exploring Adaptive Agency II: Simulating the Evolution of Associative Learning.....	306
<i>Peter M. Todd and Geoffrey F. Miller</i>	
Lookahead Planning and Latent Learning in a Classifier System.....	316
<i>Rick L. Riolo</i>	
Learning Hill-Climbing Functions as a Strategy for Generating Behaviors in a Mobile Robot.....	327
<i>David Pierce and Benjamin Kuipers</i>	
Some Parallels Between Associative Learning and Object Classification	337
<i>David R. Shanks</i>	

COLLECTIVE BEHAVIORS

Task Differentiation in Polistes Wasp Colonies: A Model for Self-Organizing Groups of Robots.....	346
<i>Guy Theraulaz, Simon Goss, Jacques Gervet, and Jean-Louis Deneubourg</i>	
The Dynamics of Collective Sorting: Robot-Like Ants and Ant-Like Robots.....	356
<i>J. L. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain, and L. Chrétien</i>	

EVOLUTION OF BEHAVIOR

Evolution and Co-Evolution of Computer Programs to Control Independently-Acting Agents	366
<i>John R. Koza</i>	
The Evolution of Information Gathering: Operational Constraints.....	376
<i>Cynthia Kurtz</i>	
Representations for Artificial Organisms.....	382
<i>Robert J. Collins and David R. Jefferson</i>	
Evolving Organisms That Can Reach for Objects.....	391
<i>Federico Cecconi and Domenico Parisi</i>	
The Artificial Evolution of Behaviour.....	400
<i>Inman Harvey</i>	
The Rise of Interaction: Intrinsic Simulation Modeling of the Onset of Interacting Behaviour.....	409
<i>Fabio De Luigi and Vittorio Maniezzo</i>	

The Evolution of Behavior: Some Experiments.....	419
<i>Jan Paredis</i>	
A Von Neumann Approach to a Genotype Expression in a Neural Animat	427
<i>Dan Wood</i>	

ARCHITECTURES, ORGANIZATIONAL PRINCIPLES, AND FUNCTIONAL APPROACHES

Challenges for Complete Creature Architectures	434
<i>Rodney A. Brooks</i>	
Cognitive Action Theory as a Control Architecture	444
<i>H. L. Roitblat</i>	
Towards a Theory of Emergent Functionality	451
<i>Luc Steels</i>	
Eight Principles for Building an Intelligent Robot.....	462
<i>David L. Waltz</i>	
Robot Ethology: A Proposal for the Research into Intelligent Autonomous Systems.....	465
<i>Uwe Schnepf</i>	
A Model of the Mechanisms Underlying Exploratory Behaviour.....	475
<i>Liane M. Gabora and Patrick W. Colgan</i>	
The Causal Analysis of an Adaptive System: Sex-Ratio Decisions as Observed in a Parasitic Wasp and Simulated by a Network Model.....	485
<i>Marijke Vonk, Felix Putters, and Berend-Jan Velthuis</i>	
Evolution as Pattern Processing: TODO as Substrate for Evolution	492
<i>P. Hogeweg and B. Hesper</i>	
Matching, Maximizing and Melioration as Alternative Descriptions of Behaviour.....	498
<i>Alasdair I. Houston</i>	
An Adaptation Anomaly of a Genetic Algorithm.....	510
<i>Yuval Davidor</i>	
Autonomous Agents, AI and Chaos Theory.....	518
<i>George Kiss</i>	
Incrementing Intelligent Systems by Design	525
<i>Brendan McGonigle</i>	

ANIMATS IN EDUCATION

Animal Simulations with *Logo: Massive Parallelism for the Masses.....	534
<i>Mitchel Resnick</i>	
Lego Vehicles: A Technology for Studying Intelligent Systems.....	540
<i>Jim Donnett and Tim Smithers</i>	
Author Index.....	551

Acknowledgments

The organizers acknowledge financial support from the following public and private institutions to which they are especially grateful:

Ecole Normale Supérieure
AFOSR (US Air Force)
Electricité de France
IBM FRANCE
OFFIUB

The Rowland Institute for Science
C2V
Ministère de la Recherche
et de la Technologie
a Corporate Donor

They wish to thank the Ministère de la Recherche et de la Technologie for having generously placed at their disposal the Amphithéâtre Poincaré in which the conference sessions will be held.

They also express their particular thanks to Josiane Serre (Administrator of the ENS) and René Lafont (Director of the URA 686) for having kindly help to solve various administrative problems.

They are particularly indebted to Louis Bec and Xueqing Wang for the artistic conception of the SAB90 poster.

They are very grateful to the MIT Press/Bradford Books company for publishing the conference's Proceedings.

Abstracts

Schemas for Prey-Catching in Frog and Toad¹

Michael A. Arbib and Alberto Cobas²

Center for Neural Engineering
University of Southern California
Los Angeles, CA 90089-2520, U.S.A.
arbib@pollux.usc.edu

Abstract

This paper first outlines the methodology of schema theory (Arbib 1981), which integrates perception and action by decomposing an overall behavior into the interaction of functional, neurally explicable, units called *schemas*. It offers comparisons with other methodologies from Artificial Intelligence (AI) and Brain Theory (BT), and reviews the RS (Robot Schema) language and the Arbib and House (1987) model of detour behavior in *Rana computatrix* which associates potential fields with objects — an attractant for the prey; a repulsor for the fencepost; and a forward field for the toad itself — which are then combined to create more complex fields which determine the trajectory of the animal. However, rather than analyze detour behavior here, the remainder of the paper presents a schema-theoretic model for the decision-making mechanisms which control prey-catching behavior in frog and toad; the extension of the work to model predator-avoidance is discussed elsewhere. It thus contributes to *Rana computatrix*, an evolving set of models of anuran visuomotor coordination (e.g., Arbib 1987). Our new model of prey-catching is rooted in recent experimental data on the behavior of animals with and without brain lesions. These data motivate the model's use of independent processing of the different

parameters that define the stimulus position (horizontal eccentricity, elevation and distance). The model, which emphasizes action generated by the concurrent activity of multiple motor schemas rather than the serial activity of such schemas, predict new behaviors for experimental test.

1. An Introduction to Schema Theory

Schema theory (Arbib 1975, 1981) provides a way to tame the complexity of large systems that are to function in the real world, offering an approach explicitly designed to bridge between cognitive science and brain theory (BT), as well as to contribute to distributed artificial intelligence (DAI). Schemas are *active* modular entities, each involving data structures and control:

a) Schemas serve to represent, at least, perceptual structures and distributed motor control. Schemas are ultimately defined by interaction with a physical environment rather than (as in most AI systems) by cross-references in some logical formalism.

b) Schema theory provides a distributed model of computation. The brain can support many concurrent activities for recognition of different objects, and the planning and control of different activities. Thus schema theory views the use, representation, and recall of knowledge as mediated through the activity of a network of interacting computing agents, schema instances. This activity may involve passing of messages, changes of state (including activity level), instantiation to add new schema instances to the network, and deinstantiation to remove instances.

c) The activity level of an instance of a perceptual schema represents a "confidence level" that the object represented by the schema is indeed present; while that of a motor schema may signal its "degree of readiness" to control some course of action. A schema network does not, in general, need a top-level executor since schema instances can combine their effects by distributed processes of competition and cooperation

¹ The research described in this paper was supported in part by grant no. 1RO1 NS 24926 from the National Institutes of Health (M.A.Arbib, Principal Investigator) and Fulbright/MEC fellowship FU88-3501116 (Spain) to A.C.

² Present address: Center for the Neurobiology of Learning and Memory, University of California at Irvine, Irvine, CA 92717, USA.

Biological and Computational Stereo Vision

Stephen T. Barnard

Artificial Intelligence Center
SRI International
333 Ravenswood Ave., Menlo Park, CA 94025

Abstract

A computational model of stereo fusion is examined in the light of biological and psychophysical knowledge of stereo vision in humans and other mammals. Several analogies are suggested, including the use of independent spatial-frequency channels with one-octave separation, the role of vergence and the limits of fusion, sensitivity to vertical disparity, and the use of 3 pools of disparity detectors. It is argued the similarity between the morphology of the visual cortex and the fine-grained, SIMD architecture exploited by the computational model leads to similar constraints on the computation of stereo disparity in both milieus, and therefore naturally leads to processes with similar properties.

1 Introduction

Scientific investigation of stereo vision in humans and other animals has an extensive history in neurobiology and psychology, dating from Wheatstone's discovery of the phenomenon in 1838 [19]. Recently computational modelers have made substantial progress in simulating the process of stereo fusion on the computer. This paper examines one such model in detail, and in particular points out some striking similarities between the model and current knowledge of stereo vision in higher mammals.

1.1 Stereo geometry

Stereo vision is a way of interpreting and exploiting visual information that is relatively well understood, in animals as well as machines. The reason is clear: compared to other perceptual cues for depth, the problem is well defined. Once the images are brought into point-to-point correspondence, recovering the third dimension is a straightforward application of trigonometry.

*The work described in this article was supported under DARPA contracts MDA903-86-C-0084, DACA76-85-C-0004, and 89F737300. Use of the Connection Machine was provided by DARPA.




Figure 1: Basic Stereo Geometry

The geometrical principle behind stereo vision, illustrated in Figure 1, is quite simple. Assume that two cameras form images through left and right centers of perspective l and r , onto planes L and R . (In practice these would be imperfect optical lens systems, but for this discussion we assume ideal "pinhole" projections.) Furthermore, assume that the cameras are fixed upon point v , which is to say that the two rays perpendicular to the image planes passing through the centers of perspective (the *principle rays*) intersect at v . Let θ_v be the angle between these principle rays. We say that the *absolute disparity* of v is θ_v . Now consider another point p projected onto image planes L and R as shown, and let the angle between these rays be θ_p . We say that the *relative disparity* of p with respect to v is $\delta_{p|v} = \theta_p - \theta_v$. Relative disparity is the more commonly used definition.

The circle through l , r , and v (actually a sphere) is called the *Vieth-Müller circle* (closely related to the

Louis B E C

" MACHINATIONS ZOOLOGIQUES "

Parallèlement à un enseignement artistique dans plusieurs écoles d'art françaises et des responsabilités pédagogiques au sein du Ministère de la Culture, Louis Bec développe un travail qui interroge perfidement les relations entre les domaines artistiques, scientifiques et technologiques.

Seul zoosystémicien en titre, il propose une épistémologie fabulatoire s'appuyant d'une part sur les aspects méthodologiques (modélisation, théories des systèmes, taxonomie, attitudes comportementales, etc...) et d'autre part sur la pertinence des modèles de la représentation du vivant et de l'animalité.

Il élabore des systèmes zoologiques arbitraires et imaginaires dans lesquels des zoomorphies singulières, des biologies curieuses, des zoosémiotiques aberrantes se développent.

En fondant en 1970 l'Institut Scientifique de Recherche Paranaturaliste, il se dote d'un instrument efficace pour interroger de toute part l'incapacité du vivant à saisir le vivant par les méthodes conventionnelles et propose une stratégie "hypocritique" par la construction de leurres heuristiques et par l'élaboration d'une métazologie biaisée.

Il participe à de nombreuses expositions (Aix en Provence, Marseille, Paris, Chalon sur Saône, Fondation du Futur/Nicolas Ledoux, Arc et Senans, Biennale de Sao Paulo, Frasso Telesino, Hanovre, Vienne, Zurich, etc...) et à de nombreux colloques et séminaires (Muséum National d'Histoire Naturelle de Paris, Conseil de l'Europe de Strasbourg, Toulouse, Lyon, Sao Paulo, Graz, Santa Fe...)

Il organise de nombreuses manifestations dont "Le Vivant et l'Artificiel" au Festival d'Avignon en 1984 et obtient une bourse de recherche du Ministère de la Culture en 1986 sur le thème des relations entre la Biologie et les Arts plastiques.

Il développe actuellement un certain nombre de projets dans les laboratoires de recherche en image de synthèse de l'Ecole Normale Supérieure de Physique de l'Université de Strasbourg et à EIKON/CERISE/RTL production (Centre Européen de Recherche en Image de Synthèse).

The Neural Basis of Behavioral Choice in an Artificial Insect

Randall D. Beer^{1,2} and Hillel J. Chiel²

Departments of ¹Computer Engineering and Science, and ²Biology
and the Center for Automation and Intelligent Systems Research
Case Western Reserve University
Cleveland, OH 44106 USA

Abstract

As its internal state and external environment continuously change, any truly autonomous agent must choose actions which are most appropriate to its immediate circumstances. This paper explores the idea that neurobiological design principles can be applied to the flexible control of autonomous agents. We describe a simulated insect whose behavior is controlled by an artificial nervous system. In particular, we focus on the neural basis of two different examples of behavioral choice in this artificial insect.

1 Introduction

One of the most fundamental problems faced by any agent, either natural or artificial, which must function autonomously in the real world is deciding what to do next. As both its external environment and internal state continuously change, an autonomous agent must constantly choose actions which lead to global behavior most appropriate for the current situation. Broadly speaking, the problem of behavioral choice encompasses the entire spectrum from minor adjustments of ongoing behavior to discrete switches between different behaviors. In addition, it involves the generation of groups of related behaviors with the appropriate timing and sequencing to accomplish specific objectives (McFarland, 1981, pp. 118-121).

How should the control system of an autonomous agent be organized to support such behavioral choice? Recently, there has been a trend toward more distributed approaches. For example, Brooks (1986) has been exploring the *subsumption architecture* for autonomous agent control. This architecture consists of layers of task-achieving behaviors each of which is implemented as a network of finite state machines augmented by timers and registers. Interactions between behaviors are handled by allowing machines in one layer to suppress interactions between machines in lower lev-

els. In a similar vein, Maes (1989) has proposed an approach to action selection in which a collection of simple agents interact by passing activation along a variety of special-purpose links.

Our own approach is grounded in the study of the neuronal mechanisms underlying the behavior of simpler natural animals, a field known as *neuroethology* (Camhi, 1984). We have been exploring the idea that neural network control architectures for autonomous agents can be designed using principles drawn directly from biological nervous systems. This approach has the advantage that more direct interactions between biological and artificial mechanisms for autonomous behavior are possible. This style of modeling was first proposed by Braitenberg (1984). In this paper, we describe an artificial nervous system we have designed for controlling the behavior of a simulated insect. The design of this insect is based in part upon specific behaviors and neural circuits drawn from several natural animals. We focus here on two different examples of behavioral choice in this artificial insect.

2 The Artificial Insect Project

The artificial insect project is aimed at exploring the use of neuroethological principles to design artificial nervous systems for controlling the behavior of complete autonomous agents, an endeavor which we have termed *computational neuroethology* (Beer, 1990). We have developed a simulated insect, a simulated environment with which it must cope, and an artificial nervous system for controlling its behavior. The insect is capable of locomotion, wandering, edge-following, and feeding, as well as properly managing the interactions between its various behaviors. In order to understand behavioral choice in this artificial insect, it is essential to understand the details of its design, which we briefly review below.

The artificial insect is a two-dimensional abstraction of a biological insect (see Figure 5). Its body consists

MODELING AND SIMULATION OF ANIMALS' MOVEMENTS (*)

Simon BENHAMOU & Pierre BOVET

Centre National de la Recherche Scientifique

Laboratoire de Neurosciences Fonctionnelles

F-13402 Marseille Cedex 9 (BOVET@FRMOP11)

Abstract

Probabilistic models were developed to represent animals' movements. The simplest one makes it possible to compute the sinuosity of an animal's search path and to determine some basic properties such as its diffusion. Applying this model in the framework of optimal foraging theory led us to determine the sinuosity value which minimizes the path length of a central place forager. More complex models, integrating cybernetic controls of the sinuosity and the velocity as a function of environmental stimulations, show how animals can orient themselves in relation to a stimulation gradient or exploit patchy environments using simple klino- and ortho-kinetic mechanisms. Another type of movement model was developed to study orientation mechanisms based on an egocentric spatial memory.

1. Introduction

Animals often exhibit random search paths: take for example the paths of foraging ants, which anybody can observe. This intrinsic randomness does not however prevent the animals from orienting efficiently towards specific goals and/or aggregating in the the most suitable areas of their environment. To understand space-use mechanisms (those whereby an animal regulates the time

it spends in the various areas of the environment) and orientation mechanisms (those whereby an animal moves towards a specific goal), it is necessary to first model the search paths. It is afterwards important to determine which environmental cues are relevant to animals and which kinetic parameters they have to regulate to be efficient. Using modeling and computer simulation of animals' movements, we have attempted to formalize some of the mechanisms involved in movement control. In this context, animals have been taken to be probabilistic self-directed mobile agents.

Here we present a general overview of the theoretical studies we have published over recent years in the field of modeling animals' movements. These models deal with a large range of natural spatial behaviours, from random foraging to oriented movements based on spatial memory. Some of these models link up with the optimal foraging theory: they are an attempt to determine which movement strategies maximize the efficiency of food searching in a stochastic environment. Other models were devised with a view to explaining orientational performances on the basis of elementary sensorimotor probabilistic mechanisms. Since a mathematical approach would be too complex to be practicable here, the properties of our models were established using computer simulations.

(*) This text is a revised version of a communication previously presented at a workshop on Modeling, Analysis and Simulation of Biological Motion, Bonn 1989.

Instinct as an Inductive Bias for Learning Behavioral Sequences

Lashon B. Booker*

Navy Center for Applied Research in Artificial Intelligence
Naval Research Laboratory, Code 5510
Washington, DC 20375-5000

Abstract

Ethologists have identified many ways that innate behavioral primitives and predispositions facilitate the learning of complex behaviors. This paper applies some of these insights to *classifier systems*. We show how certain built-in mechanisms for generating behavior in classifier systems provide an *inductive bias* that can be manipulated to improve learning. The effects of this bias on learning are illustrated by a classifier system that learns to solve a simple navigation task.

1. Introduction

Animals are born with a large repertoire of innate, coordinated patterns of muscle movement and behavior commonly referred to as *instincts*. Instincts are often thought of as inflexible, low-level "motor programs". However, ethologists have discovered that innate behaviors in fact often use learning as a strategy for filling in details that are too complex to specify completely in advance. While some instinctive behaviors may be rigid and stereotyped, many others are remarkably plastic.

The relationship between instinct and learning is not just relevant to ethologists however. Computational models of adaptive behavior can also benefit from understanding the important influences of prior structure on learning and behavior. Unfortunately, the role of prior structure and innate rules of behavior is ignored in most computational models of adaptive behavior. External reinforcement is usually viewed as the primary, if not the only, influence on learning.

This paper examines how prior structure or "instinct" can work together with reinforcement in *classifier systems*, a rule-based framework for studying adaptive behavior. The next section briefly reviews a few of the relationships between instinct and learning that

are evident in animal behavior. In subsequent sections we describe the classifier system framework and show how it can be used to implement computational models of some of these relationships. In particular, we show how certain built-in mechanisms for generating behavior provide an *inductive bias* that can be manipulated to improve learning in classifier systems. The effects of this bias on learning are illustrated by a classifier system that learns to solve a simple navigation task.

2. Instinct and Learning

At one extreme, instinctive behaviors can be sufficiently preordained and inflexible that they proceed to completion automatically once they are triggered by an appropriate stimulus. These motor programs are called *fixed action patterns* and, once initiated, they often require little or no external feedback. A classic example of a fixed action pattern is the egg-rolling behavior of geese, which a goose will complete even if the egg is taken away (Tinbergen, 1951).

Other kinds of innate behavior patterns are rigidly programmed, yet exhibit a great deal of "run-time" flexibility (Gould, 1982). Examples include the construction of bird nests, beaver dams, and spider webs. These structures have fixed, species-specific characteristics, but the building behavior can adapt itself to a wide variety of both predictable and unpredictable contingencies in the immediate environment. Even more sophisticated examples of plasticity are evident in the way human infants learn to crawl and walk. One innate component of this behavior seems to be a goal-directed specification of what to learn — infants seem to have a built-in sense of which movements "feel right" — coupled with a motivational drive for repeated experimentation. Moreover, the learned movements show a further flexibility in the way they recalibrate themselves to accommodate the growth of the body. These are just a few examples of the many ways instinctive mechanisms and predispositions can facilitate the learning of complex behaviors.

* Author's current address is Artificial Intelligence Technical Center, The MITRE Corporation, 7525 Colshire Drive, McLean, VA 22102, booker@ai.mitre.org

Challenges for Complete Creature Architectures

Rodney A. Brooks
MIT Artificial Intelligence Laboratory
545 Technology Square
Cambridge, MA 02139, USA
brooks@ai.mit.edu

Abstract

In recent years there has been a move within the artificial intelligence and robotics communities towards building complete autonomous creatures that operate in the physical world. Certain approaches have proven quite successful, and have caused a re-analysis within the field of artificial intelligence of what components are necessary in the intellectual architecture of such creatures. However nothing built thus far yet comes close the dreams that many people hold dearly. Furthermore there has been quite some criticism of the new approaches for lacking adequate theoretical justification. In this paper we outline some of the more obvious challenges that remain for these new approaches, and suggest new ways of thinking about the tasks ahead, in order to decompose the field into a number of manageable sub-areas that can be used to shape further research.

1 Introduction

There is a growing interest in building artificial creatures of some sort. One example is the recent boom in a field known as Artificial Life (see [Langston 87] and [Langston 90]). While much of the emphasis is on building forms resident in computers, which are agents acting in an information domain, there has also been some interest in physical embodiments of artificial creatures.

This author, at the MIT AI Lab, introduced the subsumption architecture ([Brooks 86] and extended in [Brooks 90]) with the explicit goal of building mobile robots with long term autonomy. Later the word *creature* crept into the language of the MIT group (e.g., [Connell 87]). The goal is to build autonomous mobile robots which operate over long periods of time, completely autonomously, in dynamic worlds. It is envisioned that these worlds are worlds which already exist for some other purpose—not worlds specially built to house the robots. Further, it is envisioned that these robots carry out some task which has some utility for whoever wanted the robots to exist and live in this

world.

As [Flynn 87] points out, there are many components to such creatures, including sensors, actuators, power sources, and intelligence. Over the last five years we have found that all these components are intimately related as we have tried to build prototype creatures ([Flynn and Brooks 89]). Choices in any part of the system architecture (e.g., sensor characteristics) have major impacts upon other parts of the system. In general it is very dangerous to think that any one component (such as intelligence) can be isolated and studied by itself.

Our experience with the subtleties of such interactions has led us to our current construction of a very complex robot, named Attila ([Angle and Brooks 90]) pictured in figure 1 (in fact we are building multiple copies of Attila). It has six legs, each with three degrees of freedom, an active whisker, a gyro stabilized pan-tilt head carrying a range finder and a CCD camera, 10 on-board processors, and over 150 sensors. We built an earlier six legged robot named Genghis ([Angle 89], [Brooks 89]), but its complexity pales in comparison to that of Attila. Many of the issues raised in this paper were brought to our attention as we have tried to work out how to program this complex robot Attila to be an artificial creature.

The bulk of this paper is devoted to the problems and challenges in designing and building the computational architectures for such creatures. However, the reader should not forget that the other aspects of a creature's architecture cannot be considered in isolation from intelligence. In a complete design, all aspects greatly influence each of the others.

We first argue that there are multiple levels of analysis or abstraction with which we must be concerned in designing and building complete creature architectures. There can be no single magic bullet or theory which will tell us all we need to know. Some problems within these levels are well circumscribed and so can be worked on in isolation. However, in order to build complete creatures we need to bridge the gaps between these levels also.

The bulk of the paper then goes on to examine each

PLEASURE: THE ANSWER TO CONFLICTING MOTIVATIONS

Michel CABANAC
Département de physiologie
Faculté de médecine
Université Laval
Québec, G1K 7P4
Canada
email: CABANAC@LAVALVM1

Abstract

An experiment has been set up to explore the hypothesis according to which the solution of conflicts of motivations is reached by the trend to maximize pleasure. Subjects were placed in a situation of conflict where the pleasure of playing a videogame clashed with the increasing discomfort of a cold environment. The time lapse tolerated by each of the subjects could be predicted from the algebraic sum of the rating of displeasure aroused by the cold environment and the rating of pleasure aroused by the videogame, obtained in different sessions. This result supports the working hypothesis and permits the conclusion that pleasure is the common currency which allows tradeoffs among various motivations.

1. Models.

One may identify two types of models those describing the behavior of a given system and those describing the system itself. The first type of model is pragmatic. A good behavioral model is adequate enough when it replicates the behavior, and is able to predict future behavioral responses of the system. Yet this may be achieved without any knowledge of the intimate mechanisms that produce the behavior. In the same way as a given envelope function can be approximated with various summations of different functions, a model may ignore the

inside of the black-box system whose behavior it nevertheless replicates adequately. The model may lump several functions into one, incorporate approximations, ignore some rare or extreme conditions but remain nevertheless a good model.

The second type of model is more ambitious and aims at theoretical description of the system. This type of model is more difficult to achieve because it includes the intimate laws of the system and because it implies the knowledge of the structures and functions of the constituent elements of the system. Eventually this second type also reaches the same goal as models of the first type.

In an enlightening chapter on instinct and motivation Epstein (1982) has recently reflected that there are behaviors which are simple reflexes in the Cartesian and Sherringtonian acceptation of the word, but there are also complex behaviors. The latter occur with the cooperative action of an endogenous component, an acquired component, and a reactive component. He stressed that «we need concepts that take account of the complexities of behaviors that are not reflexive». Theoretical models of complex behavior must include these three components. This implies the recognition and the incorporation of emergent properties and functions within the central nervous system, *i.e.* functions which cannot be predicted from the sum of properties of the elements of the central nervous system itself. Excellent theoretization on this point will be found in Toates (1986a).

Evolving organisms that can reach for objects

Federico Cecconi

Domenico Parisi

Institute of Psychology, C.N.R.

15, Viale Marx - 00137 Rome

Tel.: 0039-6-8894596

Fax: 0039-6-824737

E-mail: cecconi at irmkant.bitnet

domenico at irmkant.bitnet

Abstract

An evolutionary method based on selective reproduction and random mutation was used to evolve neural networks that control two types of simple organisms which can reach for objects using their 2-segment arm. One kind of organism does not move and can only capture an object if it is at reaching distance; the other can displace itself and therefore it first approaches an object and then captures it. Individual learning during lifetime to predict changes in the position of an object or of the hand relative to the organism's body helps in the evolution of the object reaching capacity, although inheritance of the weight matrix is strictly Darwinian. Finally, a more sophisticated fitness criterion which penalizes arm movements causes the more complex organism to move its arm only when an object is at reaching distance.

1. Introduction

Our purpose in the present paper is to describe an attempt at evolving simple simulated organisms that have the capacity to reach for objects using their single 2-segment arm. We will describe two such organisms. One does not move and can only reach for objects if the objects are located at reaching distance from the organism. The second organism can displace itself in space and therefore can approach objects and, when they are at reaching distance, it can reach for them with its arm.

The methodology we have used to develop such organisms is an evolutionary technique based on selective reproduction and random mutation (Holland, 1975; Goldberg, 1989). (For other approaches to developing similar behaviors see Jordan, 1989; Booker, 1988; Patamello and Carnevali, 1989.) An initial population of organisms each randomly different from all others is created. Each organism lives in its individual environment which contains a number of objects. At the end of their life they are rank ordered based on their performance on the object reaching tasks. There is no learning of this task during life but performances vary because of chance. Only the best performing organisms are allowed to reproduce by generating a number of copies of themselves while the others estinguish without leaving any offspring. A small amount of random variation is added to the copies so that some offspring will result in a better individual and some in a less good individual than their common parent. However, selective reproduction will insure that a better offspring is more likely to reproduce than a less good one. The net result of this evolutionary process is that the capacity to reach for objects gradually increases across a number of generations.

Another purpose of this research is to examine how the performance that emerges evolutionarily can be controlled and shaped by appropriately manipulating the fitness criterion, that is, the criterion in terms of which individuals are rank ordered and which therefore dictates who will reproduce and who won't. In the second simulation with the organism which both displaces itself in space and moves its arm, we will show

Computational Neuroethology: A Provisional Manifesto

Dave Cliff

University of Sussex School of Cognitive and Computing Sciences

Brighton BN1 9QN, England, U.K.

E-mail: davec@uk.ac.sussex.cogs

Abstract

This paper questions approaches to computational modeling of neural mechanisms underlying behaviour. It examines the "simplifying" (connectionist) models used in computational neuroscience and concludes that, unless embedded within a sensorimotor system, they are meaningless. The implication is that future models should be situated within closed-environment simulation systems: output of the simulated nervous system is then expressed as observable behaviour. This approach is referred to as "computational neuroethology". Computational neuroethology offers a firmer grounding for the semantics of the model, eliminating subjectivity from the result-interpretation process. A number of more fundamental implications of the approach are also discussed, chief of which is that insect cognition should be studied in preference to mammalian cognition.

1 Introduction

This paper questions approaches to computational modeling of the neural mechanisms underlying behaviour. It examines the relationship between computational neuroscience [26] and that style of modeling popularly referred to as "connectionism", "parallel distributed processing", or "neural networks"¹ which has recently been subject to renewed attention in the fields of cognitive science and artificial intelligence (see e.g. [25, 23]).

Connectionist models are characterised by their simplified nature and concomitant inattention to biological data, and it is argued here that such "simplifying" computational neuroscience has serious inadequacies. A different approach is suggested which pays far more attention to the sensorimotor system and hence to behavioural interactions with the external environment. This approach involves computational modeling of the neural mechanisms underlying behaviour, in a manner akin to that used in connectionism. Such an analysis of behaviour as a product of neural activity is properly the

domain of the field of neuroethology, and the new approach is therefore referred to as "computational neuroethology". Meaning is supplied to the models by embedding them in simulated environments which supply visual feedback without human intervention, that is they close the external feedback loop from motor output to sensory input.

The advantage of computational neuroethology is that the semantics of the network are well grounded, and thus results are generated by observation rather than by interpretation. That is, the fruits of computational neuroethology simulations are "hard" objective measurements rather than "soft" subjective ones. At a metatheoretical level, it is argued that the computational network simulation of cognitive processing should pay much more attention to the evolutionary history of those faculties it wishes to replicate. In particular, a conclusion of this paper is that the study of linguistic processes using network models is wildly premature. The study of insects is advocated as the most fruitful path for future research.

As the reader will probably already have detected, this paper is intentionally polemic. It is aimed at an interdisciplinary audience, and the author is no polymath. For that reason, this paper is offered as a provisional manifesto in the hope that it provokes some interesting discussion. The argument is based on previous work by a number of authors. Because of its disputatious nature, there are more direct quotes in this paper than is common. There is no denying that this is a *selective* review of the literature. This paper is abridged from [9].

The paper opens with a discussion of computational neuroscience, distinguishing it from neural engineering, and identifying two classes of model: realistic and simplifying. Following this, the connectionist paradigm is briefly summarised. Next, criticisms of connectionism are discussed, with particular attention to the argument that connectionist models have no semantic grounding without behavioural linkage to a sensorimotor system. Then, a remedy to this objection is proposed: the adoption of the computational neuroethology approach. Computational neuroethology is defined, and a specific technique for providing a behavioural linkage is discussed. This approach has some important implications for future research, the most significant of which

¹In this paper, these three terms will be treated as synonymous, and referred to collectively as "connectionist" models.

The Computational Hoverfly; a Study in Computational Neuroethology

Dave Cliff

University of Sussex School of Cognitive and Computing Sciences
Brighton BN1 9QN, England, U.K.

E-mail: davec@uk.ac.sussex.cogs

Abstract

Studies in computer vision have only recently realised the advantage of adding a behavioural component to vision systems, enabling them to make programmed 'eye movements'. Such an *animate vision* capability allows the system to employ a nonuniform or *foveal* sampling strategy, with gaze-control mechanisms repositioning the limited high-resolution area of the visual field. The hoverfly, *Syritta pipiens* is an insect that exhibits foveal animate vision behaviour highly similar to the corresponding activity in humans. This paper discusses a simulation model of *Syritta* created for studying the neural processes underlying such visually guided behaviour. The approach differs from standard "neural network" modeling techniques in that the simulated *Syritta* exists within a closed simulated environment, i.e. there is no need for human intervention: such an approach is an example of computational neuroethology.

1 Introduction

For an animal (or an autonomous robot) to adapt and survive in uncertain environments, a sense of sight is undoubtedly a useful thing to have. The creation of seeing machines has been the topic of much research in artificial intelligence and computer vision. Unfortunately, most such research has ignored the behavioural contexts in which natural vision occurs. Recently, a research paradigm known as *animate vision* has emerged, where the seeing machine is given the ability to make programmed 'eye movements', allowing it to look around: animate vision acknowledges the behavioural contexts of natural vision.

One such context is the need to control gaze when the image sampling strategy uses nonuniform resolution, i.e. where only a restricted area of the field of view is high-acuity, as in foveal vision commonly found in predatory animals. Foveal vision offers a number of advantages for any real-time visual system, whether artificial or natural (robot or animal). Animate foveal vision is not an ex-

clusively mammalian trait. The hoverfly *Syritta pipiens* exhibits animate foveal vision behaviour which is remarkably similar to corresponding behaviour in humans.

The research project described here is a 'neural network' simulation-model study of the mechanisms underlying animate vision in *Syritta*. The study of the neural basis of behaviour is properly the domain of the field of neuroethology, and this project (relying as it does on computer simulation) is thus a form of "computational neuroethology".

The simulation models the hoverfly living in a closed dynamic environment: activity in the model nervous system is expressed as flight behaviour of the model fly, which in turn generates new visual input for the simulated eyes, which feed the model nervous system, thus completing a visual feedback loop from sensorimotor output to photoreceptor input. The model nervous system is not 'hard-wired' but is created using techniques from current 'connectionist' network learning theory and from adaptive filter theory. The method by which the network is created is incidental: it is the capability of the mature network that is of interest. A number of different strategies for creating the network are being explored; the optimal solution (in the engineering sense) is not necessarily the most biologically interesting.

This paper gives a brief review of the past research on which this project builds, and then presents an overview of the simulator system. The project draws on literature in a number of fields: a more complete account of the background literature and the simulator is given in [13], which this paper is abridged from. Computational neuroethology [12] is discussed here only in passing.

The work is at an early stage. So far the most significant results are in the design of the simulator, especially the generation of the view through *Syritta*'s eyes.

2 Rationale

Work of this nature does not have a long academic pedigree. Therefore, the notes below concentrate on a few papers in some depth, rather than superficially skimming many. The animate vision paradigm is described, followed by a discussion of nonuniform sampling. The argument underlying the interest in insects is then re-

Ontogeny of Preferences in Guppies (*Poecilia reticulata*) Exposed to Food and Conspecifics

Patrick W. Colgan, T. Lynne Jamieson, Janice E. Frame,
and J. Terry Smith

Queen's University at Kingston
Kingston, Ontario, Canada, K7L 3N6
Bitnet COLGANP@QUCDN
(613) 545-6144

Abstract

A statistical method for quantifying, summarizing, and evaluating information about the behaviour of a living system is illustrated using data from a study of the ontogenetic development of preferences in guppies (*Poecilia reticulata*). This method permits comparisons of groups of individuals in different treatments when the data on each individual consists of a long sequence of behavioural states together with the entry time into those states. This tool is readily adaptable to the study of a simulation of an organism, set of organisms, or the comparison of robotic models with living systems.

1. Introduction

The use of robotic systems to model living organisms has had varying degrees of success in expanding our knowledge of living systems. Generally, the exercise is fruitful when relevant aspects of the living system are modelled faithfully. It is important, therefore, to have tools for comparing the behaviour of robotic systems with the living systems they model.

We present here one example of such a tool: a statistical method for quantifying, summarizing, and evaluating information about the behaviour of a living system. We illustrate this statistical tool by a study of preferences in guppies (*Poecilia reticulata*) for food and conspecifics. The tool is readily adaptable to the study of a simulation of an organism or set of organisms, or to the comparison of robotic models with living systems.

Guppies (*Poecilia reticulata*) are an intensively studied species for many aspects of behaviour. Pilot

studies in our laboratory indicated significant variation among both male and female guppies in their preferences for stimuli of food and conspecifics. Such individual variation has increasingly attracted the attention of behaviourists (e.g. for fish see Magurran, 1986; Gotceitas & Colgan, 1988) with respect to its basis and function. From an ontogenetic perspective, such variation raises questions concerning the role of social isolation in its development, sexual differences in this ontogeny, and the extent to which adult behaviour is predicted by juvenile behaviour.

Ontogenetic research on guppies has included studies on the development of stimulus preferences (Candland & Milne, 1966), avoidance of predators (Goodey & Liley, 1986), temperature preferences (Johansen & Cross, 1980), sexual activities (Liley & Wishlow, 1974), and the role of stress and rearing conditions (Pinckney & Anderson, 1967; Newton, 1982). Germane to the present work are the following findings from these studies. Isolated fish were less active than controls and that activity decreased over the duration of the observational trial (Newton, 1982). The isolates scored higher frequencies of social display and lower levels of sexual activity. Liley (1966) found that as experience was gained with females, previously isolated male guppies showed decreased levels of display. Pinckney and Anderson (1967) observed that group-reared fish spent a decreasing amount of time near the stimulus fish in contrast to the isolated fish that spent an increasing amount of time near the stimulus fish. The group-reared fish showed a preference for the stimulus fish of the same sex while the isolated fish showed no significant preference for either sex of stimulus fish.

All of these studies examined the behaviour of mature adult fish as the outcome of various manipulations. The objective of the present study was to monitor at frequent intervals the stimulus preferences and activity of guppies throughout the juvenile and early adult periods, and to determine the extent and ontogeny of individual differences in each sex. Individuals were raised as experimental fish

Representations for Artificial Organisms

Robert J. Collins and David R. Jefferson

rjc@cs.ucla.edu jefferso@cs.ucla.edu

Department of Computer Science

University of California, Los Angeles

Los Angeles, CA 90024

Abstract

We are interested in simulations of biological evolution, i.e. simulations of populations of organisms over many generations living in a complex and dynamic environment. Our simulations are *microanalytic*, meaning that each individual organism and gene is separately represented, and the biologically significant events in an organism's life are all separately simulated in detail. Although we have been successful with simple models, we have encountered fundamental difficulties when scaling up the complexity of the organisms and the complexity of behaviors we expect of them. These difficulties all lead to a single question: What is an appropriate representation for an organism, i.e. what is an appropriate programming paradigm in which to express the complex behavior of organisms, and how should such programs be encoded into strings so that genetic algorithms will be successful over them?

The project that brought these issues to the surface is a complex evolutionary simulation called AntFarm, in which we are attempting to evolve cooperative foraging behavior in a population of colonies of artificial "ants." In this paper we survey a number of candidate representations for organisms, that we have considered for AntFarm, all of which have been used in the past for simple evolution models. We show that none of the representations are well-suited for AntFarm. From their inadequacies we abstract a number of principles that we believe are necessary for successful evolution of complex artificial life. Finding a representation that has all of the properties we identify is still an open problem.

1 Introduction

The simulation of evolving populations of artificial organisms is very important in the study of ecological, adaptive,

and evolutionary systems whose dynamics are too complex to study analytically or experimentally [11]. In this paper we consider simulating organisms that live and reproduce in relatively complex environments, with many sensors (external and internal), and many possible actions at each moment. In addition the organisms possess some amount of internal memory, allowing their behavior to be history sensitive. In the course of its life each organism is born, makes thousands of decisions (eat, move, mate, etc.), and eventually dies. The reproductive success of a particular organism is affected by its behavior throughout its lifetime.

Our simulations are *microanalytic*, meaning that each individual organism and gene is separately represented, and the biologically significant events in an organism's life are all separately simulated in detail. Each organism in the evolving population is separately represented as a program. Each organism's life is represented as a process, a detailed sequence of events including its birth, its interaction with a dynamic environment, its competition with other organisms, its mating and reproduction (if any), and its death.

The structural representation of an organism consists of the following parts:

- interpreter: used to execute organism behavior functions (programs);
- phenotype: the behavior function (program), that maps sensory inputs and memories into memories and effector outputs;
- genotype: a bitstring that encodes the behavior function;
- development function: the mapping that decodes the genotype to produce the phenotype.

In all of our experiments, the development function is fixed for all organisms and for all time; it is not subject to evolution. The genotype, of course, differs from animal to animal, but is static throughout the life of the organism; it is the genetic material used in reproduction. At the time of reproduction, recombination and mutation operators are applied to a pair of parent genotypes to produce an offspring genotype. The phenotype of an

COORDINATION OF LEG MOVEMENT IN WALKING ANIMALS

Holk Cruse

Dept. of Biol. Cybernetics, Fac. of Biology, and Center of Interdisciplinary
Research, University of Bielefeld, Fed. Rep. of Germany
E-mail: UBIOF140 at DBIUNI111

ABSTRACT

When a robot has to move in a locally uncertain environment, propulsion by means of walking legs is advantageous compared to a wheel driven system. However, the control of walking legs is more complicated. The question of how the movement of the different legs is coordinated will be investigated here. Three different solutions which have been developed during natural evolution will be compared. These are the walking systems of an insect, of a crustacea, and of a mammal. The results show that coordinating mechanisms differ considerably in these animals.

1. INTRODUCTION

An autonomous robot which has to move in an uncertain environment has to deal with the problem of how to perform a goal-oriented behaviour. This problem is a global one, meaning that the robot has to deal with the detection of possible paths and to decide which of these it should take. This includes the problem of obstacle avoidance. When the environment is "locally certain", as for example a semi-artificial environment with flat surfaces, the technical problem of how to move the body forward is comparatively simple and can be solved by driving the robot with wheels. If, however, the local structure of the environment is uncertain, i.e., consists of irregular terrain, a robot with walking legs is advantageous. Several attempts have been made to construct such a walking machine. Nevertheless, comparing the walk of a six-legged robot with that of an animal such as an insect, immediately reveals differences. The walking of an animal is much more versatile, and it appears to be more efficient and elegant. Thus it is useful to consider biological control mechanisms in order to apply these or

similar mechanisms to the control of walking legs in machines. In the past, engineers have pointed out that little information is available on the biological control mechanisms, but this situation has changed recently. This paper provides a summary of the recent biological results focussing on data obtained from insects, crustaceans, and cat by means of behavioural methods.

2. CONTROL OF THE INDIVIDUAL LEG

The results to be described here mainly concern the coupling mechanisms between legs, i.e. those mechanisms that produce a proper coordination of the walking legs even when walking is disturbed. However, for this purpose it is necessary also to consider briefly how the movement of an individual leg is controlled. The mechano-neuronal system that produces this movement might be called the step pattern generator. To avoid a possible misunderstanding it should be stressed that it is completely open whether this step pattern generator contains an endogenous central pattern generator.

For simplicity, only forward walking will be discussed. The cyclic movement

AN ADAPTATION ANOMALY OF A GENETIC ALGORITHM

Yuval Davidor

Department of Applied Mathematics and Computer Science

The Weizmann Institute of Science

Rehovot 76100, Israel

yuval@wisdom.weizmann.ac.il

Abstract

The purpose of this paper is threefold. To analyze an adaptation anomaly observed in a specific genetic algorithm designed to optimize robot trajectories, to propose an explanation for this unusual adaptive behaviour by drawing an analogy with some elementary mechanisms in nature, and to suggest that a much more robust adaptive strategy is to allow concurrent adaptation of both the information content and the representation structure by the genetic plan. This latter strategy results in the optimization of the ability to adapt.

This paper suggests that when artificial life or machine learning applications attempt to capture the essence of natural adaptation, it is important they allow selection to operate on all levels of the system. Furthermore, it is essential to expose both the structure of the system and its information content to selection.

1 A Genetic Algorithm Designed to Optimize Robot Trajectories

Most robot applications are based on a motion trajectory composed of a movement sequence of a robot arm. Mechanically, a robot arm is an open kinematic chain comprising relatively stiff links with a joint between adjacent links. Each link represents one degree of freedom and can be commanded to move independently of all other links. Standard systems have six degrees of freedom to obtain full spatial flexibility. Since a robot arm performs a task through the motion of its end-effector attached to the last link, the last link is the primary component of the whole structure.

An *arm-configuration* is a unique arm structure defined by a set of link positions (Fig. 1). Given the positions, the end-effector's position is uniquely determined. A robot trajectory is defined by specifying a sequence of spatial positions the end-effector is required to visit. Thus, any

given path is approximated with a finite number of intervals specified by the motion vectors joining the discrete end-effector positions. In other words, any sequence of end-effector positions (>2) defines a valid trajectory which will approximate the desired path with a measurable error.




Fig. 1 - 3-link planar robot arm commanded to a fully stretched position, but the end-effector exhibits a steady-state positioning deviation from the horizontal due to limited positioning control accuracy.

Therefore, the optimization of a robot trajectory means the identification of both the optimum combination and number of end-effector positions, and that means a great many alternative trajectories that should be considered. The size of the trajectory space grows substantially even further when the robot used is of a redundant structure and most end-effector positions instantiate a multitude of different arm-configurations.

The complexity of programming a trajectory can be appreciated by examining the vertical plane in which the end-effector is required to follow the straight line connecting points A and B (Fig. 2a). One robot trajectory may be specified by sites 1 and 2 the end-effector should visit (Fig. 2b), while another trajectory might consider sites 3, 4, 5 and 6 as an alternative specification (Fig. 2c). The performance resulting from the different trajectories might be quite

THE DYNAMICS OF COLLECTIVE SORTING ROBOT - LIKE ANTS AND ANT - LIKE ROBOTS

J.L. Deneubourg¹, S. Goss¹, N. Franks², A. Sendova-Franks², C. Detrain¹ and L. Chrétien¹

¹ Unit of Behavioural Ecology, CP 231
Université Libre de Bruxelles, 1050 Bruxelles

² School of Biological Sciences
University of Bath, Bath BA2 7AY

Abstract

A distributed sorting algorithm, inspired by how ant colonies sort their brood, is presented for use by robot teams. The robots move randomly, do not communicate, have no hierarchical organisation, have no global representation, can only perceive objects just in front of them, but can distinguish between objects of two or more types with a certain degree of error. The probability that they pick up or put down an object is modulated as a function of how many of the same objects they have met in the recent past. This generates a positive feed-back that is sufficient to coordinate the robots' activity, resulting in their sorting the objects into common clusters. While less efficient than a hierarchically controlled sorting, this decentralised organisation offers the advantages of simplicity, flexibility and robustness.

1. Introduction

What is the common point between a shopkeeper and an ant colony? Each of these organisms is able to sort similar but different objects. When one examines an ant nest it is clear that neither the workers, the brood nor the food are randomly distributed. For example the eggs are arranged in a pile next to a pile of larvae and a further pile of cocoons, or else the three categories are placed in entirely different parts of the nest. The same is true in a shop. There is, however, an essential difference. The shopkeeper decides where he is going to put his different goods, and if he has assistants he

tells them where to place what. Ants work in parallel but do not, as far as we can tell, have the capacity to communicate like the shopkeeper, nor do they have a hierarchical organisation whereby one individual makes the necessary decisions and the others follow. Nevertheless, if you tip the contents of a nest out onto a surface, very rapidly the workers will gather the brood into a place of shelter and then sort it into different piles as before.

This article describes a simple behavioural algorithm, to be followed by each worker, that generates a sorting process. Sorting is achieved without requiring either external heterogeneities (e.g. temperature or humidity), hierarchical decision-making, communication between the individuals or any global representation of the environment. We also stress that the ants/robots have only very local information about the environment and a very short-term memory, and furthermore move randomly, no oriented movement being necessary. They can't see far off nor move directly towards objects or piles of objects.

Our aim in this article is not to prove that the model proposed is actually how the ants behave, but to show that such an algorithm both works and could be used by a team of robots. Inspired from our knowledge of the importance of functional self-organisation or distributed intelligence in ant colonies (Deneubourg, 1977; Deneubourg et al., 1984, 1986, 1987; Deneubourg and Goss, 1989; Goss et al., 1990; Aron et al., 1990), our idea presents a working illustration of how such a distributed system can have practical applications in robotics, in accordance with ideas developed by ourselves (e.g. Deneubourg et al., 1984, 1990; Deneubourg and Goss, 1989), and others (e.g. Beni, 1988; Brooks and Flynn, 1989; Sandini and Dario, 1989; Fukuda and Kawachi, 1989; Brooks et al., 1990; Steels, 1990). The

Lego Vehicles: A Technology for Studying Intelligent Systems

Jim Donnett and Tim Smithers[†]

Department of Artificial Intelligence
University of Edinburgh
5 Forrest Hill
Edinburgh EH1 2QL
SCOTLAND
Telephone: 031 667 1011 ext 2517
Fax: 031 225 9370
E-mail: donnett@uk.ac.ed.edai tim@uk.ac.ed.edai

Abstract

During the summer of 1989 we established the *Intelligent Sensing and Control Laboratory* in the Department of Artificial Intelligence at Edinburgh. This laboratory is designed to support both post-graduate teaching and basic research into intelligent sensing and control. In this paper we present the motives for setting up the Intelligent Sensing and Control Laboratory and the design and implementation of a Lego TechnicTM based technology used to build simple autonomous vehicles intended to support the teaching and research activities of the laboratory. We report on some of the experiences gained from its first year of operation and relate these to the requirements of our biologically oriented research programme into intelligent behaviour and its development in autonomous artificial mobile systems.

1 Introduction

During the summer of 1989 we established the *Intelligent Sensing and Control laboratory* (ISC lab) in the Department of Artificial Intelligence at Edinburgh. This laboratory is designed to support both post-graduate teaching and basic research into intelligent

sensing and control. In the autumn term (October to December, 1989) it was used for the practical sessions of a new ten week module called "Intelligent Sensing and Control"¹ during which students worked in pairs to build a series of Braitenberg-like ([Braitenberg 1986]) mobile vehicles of increasing motor-sensory sophistication. It is now being used for a number of other projects in a long term research programme to investigate the acquisition, maintenance, and adaptation of task achieving competences in autonomous mobile robots.

In this paper we present the motivation for setting up the ISC lab and the design and implementation of the technology used to build the autonomous vehicles. We also report on some of the experiences gained from its first year of operation in both teaching and research.

2 Background and Motivation

There are two different motives for setting up the Intelligent Sensing and Control Laboratory within the Artificial Intelligence Department at Edinburgh. The first concerns the nature of intelligent behaviour and the way we want to investigate it. The second concerns the kind of sensing and control we want to teach to our intelligent robotics students.

[†]Names appear in alphabetical order, with both being principal authors on this occasion.

¹This forms part of the Intelligent Robotics theme of the Department's MSc in Knowledge-based Systems.

Extended classifiers for simulation of adaptive behavior

Renaud Dumeur
Département Informatique,
Université de Paris-8,
2, rue de la Liberté,
93256 St Denis cedex 02
red@litp.ibp.fr,
C²V,
82, Boulevard Haussmann,
75008 Paris,
renaud@ccv.fr

August 28, 1990

Abstract

In this paper, we propose the use of the animat approach to automatically generate animation scripts for computer synthesized movies. We want to design a system where animats are actors able to improvise from a few contextual informations. We plan to implement such animats with extended classifiers allowing a compact encoding of behavior rules while preserving their ability to be modified by inductive genetic operations.

1 Introduction

When computer animated characters become more and more realistic in their rendering, the problem of specifying more realistic individual and collective behaviors also appears. We think that animats can be used to solve this new problem and that behavior simulation based systems can change the way computer animations are designed. We propose a "programming by environment" approach using animats and evolution simulation which can drastically reduce the work of animation script writing.

Animats are computer simulated entities, exhibiting animal-like autonomous individual or collective behavior. We want to use them as low-cost credible crowd artists in order to ease the writing of animation scripts. In the world of film making, crowd artists are employed to give the audience background informations about the time and location where the action of the movie takes place. The movie director gives to these actors fuzzy indications on how to behave, then they are left nearly without control during the filming. On the other hand, as the central character behaviors may not easily be inferred from the movie context, the corresponding actors are more precisely directed and have less freedom than crowd artists.

2 Animation automation

Consider traditional animation movies: drawn by hand, they require a frame-by-frame precision level script for each character of the animation. A first level of automation consists in drawing a few frame images by hand and apply an interpolation procedure to generate the missing frames required for a smooth animation. Improvements from the traditional paper based technique only reside in greater flexibility of the drawing tools and reduction of involved manpower.

By using computers, we can automate the generation of the script by writing procedures in a computer language. Object oriented [6] languages seem to match the requirements of animation programming. In such languages, an object is specified by a local state (a set of state variables) and a set of procedures (the object's methods). An object's behavior is implemented by its methods. As its methods process the object's local state, it is easy to obtain a wide range of behaviors from a few procedural specifications, by varying the content of the entities local states.

If object oriented programming systems facilitate the behavioral specification of large groups of simulated actors, they do not give any assurance that the resulting acting will meet the desired goals. This work is still under the responsibility of the designer who must keep in mind a model of the potential interactions between actors.

A few fixed procedures can simulate complex behaviors. Even a cellular automaton, Conway's *game of life* [2], can generate primitive animats – *gliders* and *glider guns*, for instance – with interesting behaviors. In the frame of computer animated graphics, we are moreover helped by spectators who will tend to interpret the events occurring in the movie. This fact is highlighted in V. Braitenberg's book [3] where simple animats (called *vehicles*) are involved. These animats are carts, endowed with captors (photosensible cells) and effectors (propellers such as wheels), which are wandering in an environment containing light sources. By changing the disposition and the properties of the different vehicle brain compo-

A Model of the Mechanisms Underlying Exploratory Behaviour

Liane M. Gabora

Center for Research on Concepts and Cognition
Indiana University, Bloomington, IN, U.S.A., 47405
Email: liane@cogsci.indiana.edu

Patrick W. Colgan

Department of Biology
Kingston, Ontario, Canada, K7L-3N6
Email: Colganp@qucdn.bitnet

Abstract

A model of the mechanisms underlying exploratory behaviour, based on empirical research and refined using a computer simulation, is presented. The behaviour of killifish from two lakes, one with killifish predators and one without, was compared in the laboratory. Plotting average activity in a novel environment versus time resulted in an inverted-U-shaped curve for both groups; however, the curve for killifish from the lake without predators was (1) steeper, (2) reached a peak value earlier, (3) reached a higher peak value, and (4) subsumed less area than the curve for killifish from the lake with predators. We hypothesize that the shape of the exploration curve reflects a competition between motivational subsystems that excite and inhibit exploratory behaviour in a way that is tuned to match the affordance probabilities of the animal's environment. A computer implementation of this model produced curves which differed along the same four dimensions as differentiate the two killifish curves. All four differences were reproduced in the model by tuning a single parameter: the time-dependent component of the decay-rate of the exploration-inhibiting subsystem.

1. Introduction

Selection tends to favour the evolution of systems whose organization enables more efficient ways of perceiving and interacting with the environment, and greater capacity to cope flexibly with environmental change. This often entails progressive differentiation of the parts of a system into subsystems that are specialized to take care of

different aspects of survival. For example, the parts of an animal involved in the detection and avoidance of predators can be thought of as comprising one subsystem, the parts involved in obtaining and metabolizing energy another, and so on.

Thus, an adaptive system can be thought of as a set of specialized subsystems working together to maintain the integrity of the whole. However, when we divide the system into subsystems according to ultimate causal goals such as avoiding predators and obtaining energy, every subsystem that relies on behaviour includes skeletomusculature. Much of the skeletomusculature plays a role in the functioning of many subsystems. This makes sense; subsystem-specific limbs are redundant, to the extent that (1) subsystems require only periodic control of skeletomusculature to function effectively, and (2) the skeletomusculature required to meet the needs of one subsystem could also be used to meet the needs of another. The upshot: skeletomusculature is shared amongst subsystems, and though subsystems have to work together cooperatively, they must also *compete* for control over what McFarland and Sibly (1975) refer to as the system's "behavioural final common path" (see for example, Miller, 1971; Baerends and Drent, 1982; Colgan, 1989). Which subsystem wins the competition depends upon the relative need (deviation from homeostasis) of the subsystems, the opportunities and dangers currently afforded by its environment, and the pros and cons of engaging in behaviour that has only indirect or long-term effects. Dawkins (1976) has suggested that subsystem competition is lessened somewhat by the fact that the behaviours associated with various subsystems occupy different positions on an established behavioural hierarchy; given equal need to control behaviour, the behaviour that is higher

MACHINE MOTIVATION

J. R. P. Halperin

Zoology Dept., Univ. of Toronto, Toronto, Can.,
fax: (416)-978-8532; email: janh@zoo.toronto.edu / uunet!attcan!utzooljanh

Abstract

This paper presents a design for a sensory-motor interface to generate motivated behaviour and some capacity to learn, which may be useful in producing task-oriented autonomous robots. There is a single design for the basic sensory-motor interface. By varying sensor bias, parameter settings and type of sensory input one can generate many motivational phenomena observed in lower vertebrates. Chains of instinctive behaviours can be refined by the experience of whether a response to a particular stimulus is successful. The property of incentive motivation can be mimicked, as well as learning by punishment and drive-reduction learning. Throughout the discussion we attempt to point out some of the difficulties which might be encountered in designing bio-mimic autonomous robots.

There are three broad classes of behaviour which are generally referred to as motivated. In the first class, motivation is based on the reduction of physiological deficits. Hungers for calories, proteins and minerals are the obvious examples of these "drive-reduction" motivations. No deficit-correction is involved in the second category, where the motivations are "incentive"; however these incentive motivations involve stimuli which are inherently hedonic - either attractive or aversive. An example is the motivation to consume tasty non-caloric sugar substitutes like saccharin, which can reward behaviour quite effectively in hungry animals. Finally, there is motivation to perform sequences of displays or "fixed action patterns" with neither physiological deficit-correction nor obvious hedonic responses involved - for example the motivation of territorial animals to fight intruders.

This paper presents a model which was devised to simulate fighting behaviour in Siamese fighting fish (the last, least obvious class of motivation), and which turned out to model the other types of motivation as well. This model is so simple and explicit that it could be used as a design principle for building sensory-motor interfaces to control a range of motivated behaviours in autonomous robots.

The model has two basic elements: (1) a circuit diagram showing "neurons" connected by

"synapses", some of which have variable strength (fig. 1), and (2) a connection strength change rule related to Hebb's rule, which controls the strengths of the synapses as a function of activity in the connected neurons (fig. 2).

We will describe the model in terms of an example rather than in abstracto, although a brief mathematical description is included for completeness at the end of the paper (fig. 4). A detailed mathematical model has been constructed (Halperin and Halperin, in prep.) and it forms the basis of a successful computer simulation (Halperin, Halperin, Rutherford and Dunham, in prep.). Since the fighting fish social behaviour which led to the development of the model would be difficult to visualize for anyone but ethologists, we will instead present the entirely hypothetical example of a scrap-collecting robot. The goal is to illustrate the principles and also to illustrate the problems which would have to be dealt with when using these ideas in a practical design context. Many of these problems are of the type which have been dealt with so successfully in the biological mimic robots built at MIT (e.g. Brooks and Flynn, 1989), and the success of these robots creates optimism that extensions to motivated bio-mimic robots can succeed.

The hypothetical example we will consider is a litter-collecting system for a park. There is a large, heavy collecting machine (or an expensive human to do the collecting) and a small army of

Attracting Similar Shapes Towards Each Other

Mitch R. Harris

Department of Artificial Intelligence
University of Edinburgh
Forrest Hill, Edinburgh EH1 2QL, Scotland
email m.harris@uk.ac.edinburgh

Abstract

This paper draws an analogy between the way insects use vision to move *themselves* with respect to local landmarks, and the problem of moving *objects* relative to each other in vision-guided robotic assembly. In particular, an algorithm is presented for attracting similar shapes together which was directly inspired by a model of navigation in honeybees, and which shares the same characteristics of robustness and immunity to noise. In particular, the algorithm can rotate, translate and scale one 2D shape to align it with another despite the presence of significant distortion and missing or extraneous features.

1 Introduction

There are several reports¹ that insects are able to store retinal images and to later compare these stored images with a current one in order to facilitate homing with respect to a set of (proximal) landmarks. For example digger wasps can accurately fly to their tiny nest entrance by using local landmarks such as pine cones scattered up to a metre or more from the entrance itself (Tinbergen, 1932). Similarly, a honeybee can reliably relocate a food source such as a flower even though it is too small to see from a distance or can only be seen from above due to surrounding vegetation. Again, once they are in the approximate vicinity of the food it has been shown that bees can home into the exact location by using local landmarks (Anderson, 1977).

One of the most striking features of this type of homing mechanism is its robustness in the face of missing, additional or moved landmarks between visits. Cartwright and Collet proposed an elegant computational model of landmark learning in honeybees in which the bee stores an image (retinal snapshot) of its surroundings when it finds a food source (Cartwright & Collett, 1983). To re-find the this location the bee must first find the approximate vicinity (ie by some other form of navigation) so that the most of the

relevant landmarks are visible. Given these starting conditions, the authors show how comparison of the snapshot and a current image can lead to the computation of a flight vector which, if followed, tends to bring the bee into a position where snapshot and new image are more similar. Thus as the (simulated) bee flies around it continually compares its snapshot with the image of its current surroundings and computes new flight vector and so on, thus bringing it ever closer to the 'memorised' location.

The main reason for the robustness of Cartwright and Collet's algorithm is that it does not need to solve the correspondence problem of which features in the image match which features in the snapshot. Computation of the flight vector is independent of the goodness of fit between image and snapshot and so no search is involved. The effect of the computation is that the bee moves towards things (or gaps between things) which appear too small and away from things which appear too big.

For simplicity, it is assumed that image and snapshot each consist of a 360° wide ring of dark and light segments. Each light or dark region in the image is paired with the nearest region of the same sign in the snapshot and the angular difference is found. These differences are simply averaged to find an overall rotation for the bee. Translational vectors are computed by comparing the angular sizes of dark and light areas. For example a dark area in the image is compared with the nearest dark area in the snapshot; if the image area is *smaller* then this suggest moving radially *towards* that part of the image. As with rotation, radial vectors generated in this way are simply averaged, giving an overall translation vector. Even if individual pairings of edges are wrong in terms of the correspondence, the overall flight vector tends to be roughly right and small errors are compensated for by the iterative nature of the process.

1.1 A robot guidance problem

It was decided to take a similar approach to the bee model to help solve an analogous problem which arises when guiding a hand-held object to bring it in line with a stationary object under camera guid-

¹ See (Cartwright & Collett, 1983) for review

The Artificial Evolution of Behaviour

Inman Harvey

School of Cognitive and Computing Sciences, University of Sussex

Brighton BN1 9QH, England

email: inmanh@cogs.susx.ac.uk

Abstract

Darwinian evolution produces intelligent behaviour without a designer, and this can be used to evolve behaviour in simulated organisms. The problems associated with using genetic algorithms to evolve programs in a conventional language, and to evolve the architecture of connectionist networks, are compared. For an open-ended space of network architectures, a developmental language to model the production of a network from the genotype is necessary, as is a new theoretical analysis of genetic algorithms not limited to genotypes of a fixed length.

1 Introduction

Artificial Intelligence arose as a field of study from the belief that intelligent human behaviour could be formalised, and hence could be mechanised. Problem solving must be done according to rules, so this approach went; put the rules in a machine, and we will have an intelligent machine. The greatest successes have been, understandably, in just those fields of human intelligence where the problems can be formally defined, e.g. chess-playing, expert systems in simple domains; even here, success has been limited. Connectionism (McClelland and Rumelhart 1987) uses a different approach; investigating network models which are based on a very simplified model of the brain, namely large numbers of simple processing nodes with many wires connecting them, passing activations throughout the network. A major insight to come from this approach is that the behaviour of the whole can look as though it is obeying explicitly programmed rules, even though one can see that this is just an emergent property of the underlying mechanisms.

The conventional A.I. approach tries to design into a program intelligent behaviour; the connectionist tries to design networks that will produce intelligent or adaptive behaviour. Yet Darwinian evolution shows us that there can be intelligent behaviour without a designer.

Hence as one can consider the intelligence and adaptability of humans and other animals to be an emergent property of their evolutionary history in their environment, one can also consider the possibility of the emergence of intelligent and adaptive behaviour in simulated

organisms, in simulated environments, with some form of evolutionary algorithm.

Approaches on these lines have been the subject of much recent interest, and conferences on similar themes have been held under the title of 'Artificial Life' in Santa Fe in 1987 and 1990 (Langton 1989, Langton *et al.* 1991), and with the titles 'Evolution, Games and Learning' and 'Emergent Computation' by the Center for Nonlinear Studies, Los Alamos, in 1985 and 1989 (Farmer *et al.* 1986, Forrest 1990).

This paper begins with an introduction to the most developed work on evolutionary algorithms, Genetic Algorithms (GAs). There follows a discussion of the problems involved in using an evolutionary approach to developing conventionally structured programs, such as might be required by a conventional, symbol-processing, A.I. approach to simulating behaviour; a program evolution system is described that has a novel evaluation function which sidesteps one of these problems, at the expense of others. Classifier systems are mentioned, leading to a discussion of recent work on applying evolutionary techniques to developing connectionist networks with architectures appropriate for the cognitive task they are facing. Finally, the necessity is stressed for a new theoretical backing to GAs, and the underlying pitfalls of evolving computational systems are looked at.

2 Genetic Algorithms

GAs were developed initially by John Holland in the 1960's (Holland 1975) as a form of search technique modeled on Darwinian evolution. The most accessible introduction is by Goldberg (1989); other sources are Davis (1987), and the Proceedings of the first three GA conferences (Grefenstette 1985, Grefenstette 1987, Schaffer 1989).

Given a search problem, with a multi-dimensional space of possible solutions, a 'genetic-code' representation is chosen such that each point in the search space is represented by a string of symbols, the genotype. A number of initial random genotypes are produced, typically by a random number generator, which form the initial population. Each of the corresponding points in the search space (which can be considered as representing

Evolution as pattern processing: TODO as substrate for evolution

P. Hogeweg & B. Hesper
Bioinformatica Padualaan 8,
3584 CH Utrecht the Netherlands
EMAIL: HESPER@RULCRI.Leiden.Univ.nl

Abstract

Evolution is most often viewed (and formalised) as an optimisation process. In this paper we suggest that a useful alternative heuristic may be to view (and formalise) evolution as a self-enhancing pattern transformation, pattern detection and pattern generation process. This suggestion is based on the growing awareness of the emergence of complex behaviour from simple environmentally dependent action rules when they operate in a structured environment, and the structuring of the environment by such behavioural patterns. We describe paradigm worlds, which suggest that such emergent behaviour may underlie behaviour patterns observed in various species. We show how the emergence of macro-behaviour patterns can be interpreted as a form of pattern processing by the action rules on the environment. We suggest that these emergent patterns function as a prepattern for evolutionary processes: evolution fixates and enhances these patterns.

1. Introduction

Ever since Darwin's profound insight in equating 'survival' and 'fitness' for self-replicating entities, evolutionary theory has had a strong footing in terms of optimisation processes. This optimisation viewpoint pervades most biological thinking about evolution and adaptation. Population genetics is entirely formulated in such terms, and traits of organisms are customarily 'explained' in terms of their 'fitness', recently in particular in sociobiology and behavioural ecology. Genetic Algorithms (Holland 1976, Goldberg 1989) have used basic 'genetic mechanisms' for solving general optimisation problems. Only relatively recently have quantitative studies begun to expose the constraints on a 'mutation selection' process leading to appreciable optimisation (after all: 'optimisation' by 'survival of the fittest' is not a tautology). Eigen and Shuster (1979) exposed the 'error threshold', i.e. they showed that only a limited amount of mutation is compatible with evolutionary optimisation. Kauffman (Kauffman & Smith 1986 Kauffman 1989a,b) stressed that optimisation is only possible in not too rugged fitness landscapes, i.e. only if similar genotypes have, in general, similar fitnesses. Rugged fitness landscapes result from extensive coupling between genes, by which the system exhibits strong selforganising properties. Such selforganisation thus seems to be a constraint on evolutionary optimisation.

All these approaches use an external, apriori, 'user' imposed fitness criterion. Only a few models are studied, in which only survival determines the evolutionary process in a

coevolutionary context. (Conrad & Rizki 1989, Packard 1988, Holland 1990, Kauffman 1989b). In such models 'fitness' is not clearly defined, and fitness landscapes are wildly dynamic entities, if they can be visualised at all.

In this paper we propose that it may be worthwhile to view evolutionary processes not as primarily an optimisation process, but instead as pattern processing (i.e. as pattern detection, pattern transformation or pattern generation). Such a view is of course entirely compatible with the optimisation viewpoint, but provides different heuristics for studying such processes. In particular we should like to study what may be called the generation of 'fitness' dimensions, rather than walks through fixed fitness landscapes, although we prefer a terminology like 'pattern of survival' rather than fitness. The insight that pattern recognition and pattern detection can be studied in terms of energy minimisation ('optimisation' as used in evolutionary theory) (Hopfield 1984, Ackley et al. 1985 and the extensive literature which followed this conceptualisation) has led to important new models and machines for pattern recognition/ pattern detection. We hope that a similar, but reversed, change of viewpoint with respect to evolution (i.e. from optimisation to pattern detection) will likewise lead to new models elucidating (and possibly machines, exhibiting) innovative evolutionary adaptation.

2. TODO and emergent behaviour

2.1 Introduction to the TODO principle.

The potential of local rules to generate complex behaviour in interaction with a structured environment was first hinted at by Simon (1969) in his phrase: "an Ant viewed as a behaving system is quite simple, the apparent complexity of its behaviour in time is largely a reflexion of the complexity of the environment in which it finds itself..."; "a Man viewed as a behaving system is quite simple, the apparent complexity of his behaviour in time is largely a reflexion of the complexity of the environment in which he finds himself...". Simon, and with him most Artificial Intelligence research, have concentrated entirely on humans, and have in practice dismissed this phrase as an irony. Nor have those studying animal behaviour taken the hint seriously, they have continued to study behaviour virtually independent of the environment or they have paid attention to the environment as a constraint on (optimising) behaviour only. By contrast our own research has been in the direction of Simon's pointer, but has gone beyond it by

Matching, maximizing and melloration as alternative descriptions of behaviour

Alasdair I. Houston
NERC Unit of Behavioral Ecology
Department of Zoology
University of Oxford

Abstract

This paper reviews some issues concerning the use of optimality principles in the context of animal behaviour. After a brief discussion of constraints, I look at matching and maximizing as accounts of behaviour. As long as maximization is used descriptively there is no conflict between these approaches. They are alternative forms of description; in general it is possible to convert one form of description into the other. Melloration has been suggested as a principle that underlies matching. When faced with two alternatives melloration requires that the animal increases its allocation of time to the alternative with the higher local rate of reinforcement. When the local rates become equal, an equilibrium is attained at which matching occurs. I present a simple model of melloration in which this equilibrium is not necessarily stable. It is shown that time allocations can be periodic or chaotic, and that matching does not necessarily occur.

1. Introduction

The use of optimality principles in biology remains controversial (see for example Maynard Smith 1978, Gould and Lewontin 1979, Williams 1985, Ollason 1987, Dupré 1987). In this paper I make no attempt to review all the issues. After some remarks on the problem of constraints, I concentrate on alternatives to optimality that have emerged in operant

psychology. It is argued that some of the confusion in this area arises because of ambiguities concerning the status of optimality as an account of behaviour. I then consider a dynamic principle called melloration. This is a local optimization principle that may not always result in optimal behaviour. I show that under some circumstances a form of melloration can result in chaotic behaviour.

2. Constraints

Gould and Lewontin (1979) use the central dome of St Mark's, Venice as a starting point for their attack on optimality. The dome rests on four pendentives, which Gould and Lewontin refer to as spandrels. (The term "spandrel" is usually restricted to 2-dimensional surfaces of the sort shown in Figure 1.) Each pendentive is decorated with a mosaic of one of the four evangelists, together with his associated river. Gould and Lewontin suggest that the design could result in the view that the architecture is secondary to it. They contrast this with the proper path of analysis's that 'begins with an architectural constraint: the necessary four spandrels and their tapering triangular form'. (p582)

It is important to realise that this is not an argument against an optimality analysis, but merely an illustration of the importance of asking the right question. As Gould and Lewontin (1979) state, the mosaics clearly serve the function of expressing the Christian faith. We can look at the decoration as a whole in this light, and we can compare the decoration of restricted triangular spaces in various domed churches. The architecture specifies the constraints within which an "optimality" analysis of the decoration operates.

THE CONTRIBUTION OF QUANTITATIVE MODELS TO THE LONG DISTANCE ORIENTATION PROBLEMS

Marc Jamon

Centre National de la Recherche Scientifique
Laboratoire de Neurosciences Fonctionnelles
France-13402 Marseille Cedex 9
e-mail: BOVET@FRMOP11

Abstract

Examples of simulation models applied to the long distance orientation of rodents, bees, salmon and pigeons are reviewed. In each case the models provide new insights which are contradictory to the current map hypothesis. Salmon and rodent performances are assumed to be consistent with a random process. Bees supposedly use non cognitive processes to perform map-like orientation procedures, and pigeons' homing is found to involve mainly stochastic processes. The efficiency of the simulated models as well as the fact that they point to hitherto underestimated aspects of orientation make them useful new tools for solving spatial problems.

1. Introduction

Authors using the theoretical approach to the long distance orientation in animals have consistently neglected the development of quantitative models. Consequently, the properties of random processes have been rarely formulated and their possible contribution to observed performances have been underestimated (Jamon, 1987). Nor have performances been attributed to any distinguishable orientation mechanisms. Indeed, long distance orientation has been largely interpreted in terms of analogies with human navigation techniques, leading to the explicitly formulated idea that "routine animals' movements are governed by a navigational process closely analogous to every day marine practice" (Gallistel, 1989). The reference to man made concepts of orientation is based on the "map" metaphor, which implies that animals can build up a mental representation of their environment. Two sorts of maps have been envisaged. When an animal moves in a familiar area, it is supposed to develop a mental representation of the landmarks (Wiltschko and Wiltschko, 1987, talk about animals' mental picture of the spatial distribution of the factors used), which has been variously called the mosaic map, familiar area map, topographical map or cognitive map. When the animal has to orient in an unfamiliar territory, it is supposed to use some sort of bico-ordinate map,

formed by the intersection of at least two gradients along association with a compass, exactly as human navigators do. In this case, one talks about a grid map, gradient map or navigational map. This map and compass hypothesis was first formulated by Kramer (1953), but considerable difficulties have since been encountered in determining the map.

Recent developments in the theoretical approach to orientation have led to more attention being paid to quantitative models, owing to the development of computers and simulation techniques. These have thrown new light on orientation by predicting the performances of orientation mechanisms which were considered *a priori* to be implausible.

To show how these new techniques can change our view of animal orientation, I propose to review four examples where models have provided new insights contradicting the current map hypothesis. Each of these involves species suspected of being capable of long distance orientation. The first two are alternative models proposed in situations where homing by means of familiar area map was first envisaged: they concern rodents and bees. The last two deal with situations where the use of a navigational map has been hypothesized: these concern salmon and pigeons.

2. The case of rodents' homing ability.

Numerous experiments have shown that rodents transported a relatively long distance from their home can return to their previous territory with a probability of success which is higher than that predicted by randomness (Joslin, 1977) (fig. 1).

Various orientation mechanisms have been suggested to explain this good successful homing performance level. The animals have been said to rely on some specific orientation mechanism such as route-based navigation involving the use of magnetic cues (Mather and Baker, 1980, 1981), or to pilot by means of a large familiar area map (Furrer, 1973). In the latter case, they would have to build up a topographical representation of a large territory, extending beyond the limit of their actual home range during the course of exploratory trips.

Autonomous Agents, AI and Chaos Theory

George Kiss

Human Cognition Research Laboratory
The Open University
Milton Keynes
England, MK7 6AA
email: gr_kiss@vax.acs.open.ac.uk

Abstract

Agent theory in AI and related disciplines deals with the structure and behaviour of autonomous, intelligent systems, capable of adaptive action to pursue their interests. In this paper it is proposed that a natural reinterpretation of agent-theoretic intentional concepts like knowing, wanting, liking, etc., can be found in process dynamics. This reinterpretation of agent theory serves two purposes. On the one hand we gain a well established mathematical theory which can be used as the formal mathematical interpretation (semantics) of the abstract agent theory. On the other hand, since process dynamics is a theory that can also be applied to physical systems of various kinds, we gain an implementation route for the construction of artificial agents as bundles of processes in machines. The paper is intended as a basis for dialogue with workers in dynamics, AI, ethology and cognitive science.

1 Introduction

Agent theory is a branch of artificial intelligence (Kiss, 1988). Its domain is the theory, design and implementation of artificial systems, similar to animals or people, that are capable of autonomous, rational actions through which to pursue their interests and goals. Aspects of this theory cover, among other things, how actions are related to knowledge, how plans for actions to reach goals can be formed, how goals are formed, what the role of intentions for action is, how the state of the world is perceived, and many others.

The abstract formulation of agent theory can be stated in many different languages, both informal and formal. Much current work in this field has made use of formal logical languages (Georgeff and Lansky, 1986). Although these specialised logics are convenient and expressive, often it is difficult to formalise their semantics, or the semantics that have been offered have undesirable properties. An example

of this is the possible-world semantics of epistemic logics which unfortunately makes agents omniscient.

The implementation of theories expressed in such formal languages has additional problems. When agent implementation is done by direct mechanisation of the logic, for example as a theorem-prover, the resulting systems turn out to be inefficient. This is a natural consequence of the expressiveness of the language. On the other hand, the languages are sometimes not expressive enough to deal with some concepts that seem needed to describe agents. An example is the expression of quantitative magnitudes for describing strength of belief in an agent.

Refinement of these logics and their formal semantics, and their efficient implementation, is of course an ongoing enterprise. This paper is intended as an informal preliminary to such work, offering some intuitions about the interpretation of agent theory through the general theory of process dynamics.

Such an interpretation can also provide a strategy for implementation. The situation is analogous to the relationship between the abstract Boolean algebra of classes, the propositional calculus, and hardware logic circuits. The abstract algebra is defined in terms of classes and operations on them; intersection, union, complementation, etc. One interpretation of the Boolean algebra is propositional logic, where the variables range over propositions and the operations are truth-functional manipulations, etc. The possibility of implementation arises from the fact that another interpretation of Boolean algebra can be found in the operation of physical electrical circuits. Because of this, the operation of the circuits can thus be described by propositional logic, or stated conversely, the circuits are an implementation of the logic.

Let us represent this by the following schema:

Propositional logic \rightarrow Abstract Boolean algebra \rightarrow
Electrical circuits

EVOLUTION AND CO-EVOLUTION OF COMPUTER PROGRAMS TO CONTROL INDEPENDENTLY-ACTING AGENTS

John R. Koza

Computer Science Department

Stanford University

Stanford, CA 94305 USA

Koza@Sunburn.Stanford.Edu

415-941-0336

ABSTRACT

This paper describes the recently developed "genetic programming" paradigm which genetically breeds populations of computer programs to solve problems. In genetic programming, the individuals in the population are hierarchical computer programs of various sizes and shapes. This paper also extends the genetic programming paradigm to a "co-evolution" algorithm which operates simultaneously on two populations of independently-acting hierarchical computer programs of various sizes and shapes.

1. INTRODUCTION AND OVERVIEW

This paper describes the recently developed "genetic programming" paradigm which genetically breeds populations of computer programs to solve problems. In genetic programming, the individuals in the population are hierarchical compositions of functions and arguments of various sizes and shapes. Each of these individual computer programs is evaluated for its fitness in handling the problem environment and a simulated evolutionary process is driven by this measure of fitness.

This paper also extends the genetic programming paradigm to a "co-evolution" algorithm which operates simultaneously on two (or more) populations of independently-acting hierarchical computer programs of various sizes and shapes. In co-evolution, each population acts as the environment for the other population. In particular, each individual of the first population is evaluated for "relative fitness" by testing it against each individual in the second population, and, simultaneously, each individual in the second population is evaluated for "relative fitness" by testing it against each individual in the first population. Over a period of many generations, individuals with high "absolute fitness" tend to evolve as the two populations mutually bootstrap each other to increasingly high levels of fitness.

In this paper, the genetic programming paradigm is illustrated with three problems. The first problem involves genetically breeding a population of computer programs to allow an "artificial ant" to traverse an irregular trail. The second problem involves genetically breeding a minimax control strategy in a differential game with an independently-acting pursuer and evader. The third problem illustrates the "co-evolution" and involves genetically breeding an optimal strategy for a player of a simple discrete two-person game represented by a game tree in extensive form.

2. BACKGROUND ON GENETIC ALGORITHMS

Genetic algorithms are highly parallel mathematical algorithms that transform populations of individual mathematical objects (typically fixed-length binary character strings) into new populations using operations patterned after (1) natural genetic operations such as sexual recombination (crossover) and (2) fitness proportionate reproduction (Darwinian survival of the fittest). Genetic algorithms begin with an initial population of individuals (typically randomly generated) and then iteratively (1) evaluate the individuals in the population for fitness with respect to the problem environment and (2) perform genetic operations on various individuals in the population to produce a new population. John Holland of the University of Michigan presented the pioneering formulation of genetic algorithms for fixed-length character strings in *Adaptation in Natural and Artificial Systems* (Holland 1975). Holland established, among other things, that the genetic algorithm is a mathematically near optimal approach to adaptation in that it maximizes expected overall average payoff when the adaptive process is viewed as a multi-armed slot machine problem requiring an optimal allocation of future trials given currently available information. Recent work in genetic algorithms and genetic classifier systems can be surveyed in Goldberg (1989), Davis (1987), and Schaffer (1989).

3. BACKGROUND ON GENETIC PROGRAMMING PARADIGM

Representation is a key issue in genetic algorithm work because genetic algorithms directly manipulate the coded representation of the problem and because the representation scheme can severely limit the window by which the system observes its world. Fixed length character strings present difficulties for some problems — particularly problems in artificial intelligence where the desired solution is hierarchical and where the size and shape of the solution is unknown in advance. The need for more powerful representations has been recognized for some time (De Jong 1985, 1988).

The structure of the individual mathematical objects that are manipulated by the genetic algorithm can be more complex than the fixed length character strings. Smith (1980) departed from the early fixed-length character strings by introducing variable length strings, including strings whose elements were if-then rules (rather than single characters). Holland's introduction of the classifier system (1986) continued the trend towards increasing the complexity of the

The evolution of information gathering: operational constraints

Cynthia Kurtz
Dept. of Ecology & Evolution
State University of New York at Stony Brook
Stony Brook, NY 11794 USA
e-mail: kurtz@sbbiov.m.bitnet
ckurtz@sbbiov.m.bitnet

Abstract

I present a scheme for partitioning information used in decision making. Three types of information are recognized: internal information, or an individual's internal state; external information, or environmental factors; and relational information, or rules for predicting transformations of internal state. A genetic simulation model is described which tracks the evolution of alleles for high and low information access in each information type in a population with density dependence. Stable polymorphisms result. Interactions between the three genes are explored. The relevance of the model to foraging situations is discussed.

1. Introduction

A major criticism of classical optimal foraging models is that they assume complete information (Stephens & Krebs 1986). The problem of incomplete information has received considerable attention in the past twenty years, and the basic prey and patch models have been expanded to consider incomplete information in prey recognition (Houston et al. 1980, Getty & Krebs 1985), patch sampling (McNamara 1982, Lima 1984, Bernstein et al. 1988), and tracking a changing environment (Stephens 1987, Shettleworth 1988). In each of these cases the information studied is "about" the environment; that is, animals sense the states of pertinent environmental parameters.

Is environmental information the only sort of information an animal needs to make its decisions? How might one partition information in a way that is useful in thinking about the evolution of behavior?

Let us begin with a simplified look at the process of survival and reproduction. Any individual in a population can be described at a

given time by an internal state which largely determines its current reproductive potential. That state will be a complex of many factors, including stored energy, health, fertility, and (in a sexual population) attractiveness to mates. The lifetime of any individual can be thought of as a time series of internal state transformations. Transformation will be influenced by the individual's own behavior and by factors in the environment which influence that individual, and will be governed by a set of rules which we may collectively describe as a transforming function.

From this basic scheme, there appear three types of information to which an individual making decisions may have access. First, the individual may be aware of its own state. Call this internal information. Second, it may be aware of environmental conditions, or external information. Third, it may be aware of the form of the transforming function, or relational information.

Now consider foraging models with these three information types in mind. The classical models assume complete external and relational information; but since they are static, or state-independent, ignore internal state. Their refinements deal with deficiencies in external information only. Dynamic optimization (Mangel & Clark 1988) considers state-dependent decision making; but it assumes complete internal information. Finally, studies of rules of thumb (Janetos & Cole 1981, Green 1984) implicitly consider reductions in relational information, mainly as satisficing strategies (Simon 1956) due to functional constraints, but assume complete internal and external information.

This paper presents a model that examines the interaction of internal, external and relational information genes in an evolving population. I chose simulation methods to model this situation for two reasons. First, including enough processes to generate information of each type makes the model sufficiently complex to prohibit an easy analytical analysis. Second, I was interested in

Self-improving Reactive Agents: Case Studies of Reinforcement Learning Frameworks

Long-Ji Lin

School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
e-mail: ljl@a.gp.cs.cmu.edu

Abstract

The purpose of this work is to investigate and evaluate different reinforcement learning frameworks using connectionist networks. I study four frameworks, which are adopted from the ideas developed in [Barto, Sutton & Watkins, 1989; Watkins, 1989; Sutton, 1990]. The four frameworks are based on two learning procedures: the Temporal Difference methods for solving the credit assignment problem, and the backpropagation algorithm for developing appropriate internal representations. Two of them also involve learning a world model and using it to speed learning. To evaluate their performance, I design a dynamic environment and implement different learning agents, using the different frameworks, to survive in it. The environment is nontrivial and nondeterministic. Surprisingly, all of the agents can learn to survive fairly well in a reasonable time frame. This paper describes the learning agents and their performance, and summarizes the learning algorithms and the lessons I learned from this study.

1. Introduction

Reinforcement learning is an interesting learning problem. It requires only a scalar reinforcement signal as a performance feedback from the environment. Reinforcement learning often involves two difficult subproblems. The first is called the *credit assignment problem*. Suppose the learning agent performs a sequence of actions and finally obtains certain outcomes. It must figure out how to assign credit or blame to each individual situation (or situation-action pair) to adjust its decision making and improve its performance. The second subproblem arises from the need to develop the appropriate internal representations required to achieve the target learning tasks. In the course of learning, both subproblems must be solved.

Several reinforcement learning frameworks or algorithms have been proposed in the literature (e.g., [Sutton, 1984; Williams, 1987; Barto, Sutton & Watkins, 1989; Watkins, 1989; Kaelbling, 1989; Sutton, 1990]). However, most have only been studied solving simple learning problems (e.g., [Anderson, 1989]). In addition, no

serious comparison of different frameworks has been done. This work is thus intended to be a first step towards the investigation and evaluation of different reinforcement learning frameworks in solving nontrivial learning tasks. In particular, I am interested in reinforcement learning using connectionist networks.

In the paper I study four reinforcement learning frameworks, which are adopted from the ideas developed in [Barto, Sutton & Watkins, 1989; Watkins, 1989; Sutton, 1990]. All of these frameworks are based on two learning procedures: the Temporal Difference (TD) methods [Sutton, 1988] for solving the credit assignment problem and the error backpropagation algorithm [Rumelhart, et al., 1986a] for developing appropriate internal representations.

Generally speaking, reinforcement learning based solely on the TD methods is a slow process. In domains where reinforcements are sparse, the learning rate is slow, and if the cost of mistakes (e.g., physical damage) is also high, the agent would make more mistakes than allowed. A solution to these problems is to learn a world model, and practice with the model. This idea is embodied in two of the frameworks studied here.

My approach to evaluating different learning frameworks is to design a dynamic environment, implement learning agents to survive in it using different frameworks, and evaluate the performance of the agents. Four kinds of objects are involved in this environment: the agent, fixed food and obstacles, and moving enemies. Although survival in this environment is easy for humans, it is by no means trivial for knowledge-poor agents.

The remaining of this paper is organized as follows. Section 2 discusses the four learning frameworks. Section 3 describes the rules of the environment. Sections 4 and 5 present the implementation and performance of the learning agents. Section 6 assesses the merits of the agents. Finally, Section 7 concludes the paper by summarizing the lessons I learned from this study.

2. Reinforcement Learning Frameworks

Learning to survive in an unknown environment can be characterized as a kind of *reinforcement learning*. In reinforcement learning, the learning agent continually receives sensory inputs from the environment, selects and

The Rise of Interaction

Intrinsic simulation modeling of the onset of interacting behaviour

Fabio De Luigi

Town Municipality of Ferrara, C.E.D, 44100 Ferrara, Italy

Vittorio Maniezzo

MP-AI Project Dipartimento di Elettronica
Politecnico di Milano, 20133 Milano, Italy

E-mail: MANI%ipmell.infn.it@iboinfn.bitnet

Abstract

A simple model is presented, aimed at capturing the essential traits of the interacting behaviour of adapting species. The modelling paradigm followed is centered on *intrinsic adaptation*, with no explicit fitness function, and is implemented using a probabilistic cellular automaton. Some simulation results are shown, regarding the outset of a predator species and prey/predator population dynamics, also with respect to environmental structural changes.

1. Introduction

Complex systems science aims at capturing the fundamental characteristic of adaptive systems composed of multitudes of interacting entities. The fundamental concept around which all its approaches hinge is that of *self-organization* [Nicolis, Prigogine, 1977], that is the emergence of organized behaviour in systems which were not designed explicitly to manage entities at the level of the outputs. A common feature of all such models in fact regards the input specification and the system description, which are defined at an aggregation level quite far - both in terms of object structuring and of characteristic time scale - from that of the output of interest.

This approach allows simulation of realities unapproachable with traditional analytic techniques. A paradigmatic example is the evolution of an ecosystem where several species coexist. In this situation every species affects with its presence the environment it lives in, consisting both in the world and in its other occupants: its survival probability, along with its fitness to the environment, should therefore continuously be readjusted. In the paper we present a model where fitnesses are only implicitly dealt with, in that they emerge from environment self-organizing evolution and can be computed only *a posteriori*. Given the description of an environment and an initial uniform population, we simulate the adaptation of the population to the environment, along with the possible outset of new species.

The paper is organized as follows: in section 2 we introduce the simulation approach followed in our research along with the essential features of our model, in section 3 we describe the methodology followed to introduce interaction among the basic individuals, in section 4 we give a more detailed description of our implemented system and propose some simulation result. Finally, in section 5, we briefly outline our current activities regarding extensions of the model.

A BOTTOM-UP MECHANISM FOR BEHAVIOR SELECTION IN AN ARTIFICIAL CREATURE

Pattie Maes

AI-Laboratory
Vrije Universiteit Brussel
Pleinlaan 2
B-1050 Brussels
Belgium
pattie@arti.vub.ac.be

AI-Laboratory
Massachusetts Institute of Technology
545 Technology Square
Cambridge, MA 02139
pattie@ai.mit.edu

Abstract.

In this paper we propose a mechanism for motivational competition and selection of behavior. One important characteristic of this mechanism is that the selection of behavior is modelled as an emergent property of a parallel process. This in contrast with mechanisms for behavior selection and motivational competition proposed earlier, which are based on a hierarchical, preprogrammed control structure. We show that selection of behavior can be modeled in a bottom-up way using an activation/inhibition dynamics among the different behaviors that can be selected. There is no weighing up of behaviors in a cognitive manner and neither are hierarchical or bureaucratic structures imposed. The paper elaborates upon the results we obtained with simulated creatures based on this mechanism. It draws parallels between characteristics observed in animal behavior and characteristics demonstrated by our artificial creatures. Examples are: displacement behavior, opportunistic behavior, fatigue, selective attention, and so on.

1. Introduction

This paper is concerned with the problem of behavior selection for an artificial creature. The context in which we discuss this problem is that of the behavior-based systems (Brooks, 1986) (Brooks, 1990), which embody a new philosophy for building artificial creatures, inspired by the field of Ethology (Mc Farland, 1981) and not unrelated to the Society of Mind theory (Minsky, 1986).

A creature is viewed as consisting of a set of behaviors. Examples of behaviors are: the feeding behavior, sleeping behavior, drinking behavior, etc. Only a few -- or often only one -- of these behaviors can be active at a time. However, a creature at every moment is probably motivated towards a variety of them. This means that there has to be some mechanism which decides which behavior "wins" and as such gets control over the "muscles" or actuators of the artificial creature.

In the case of simple animals, and also in the case of simple artificial creatures, the optimal strategy can be hard-wired respectively by nature (natural selection) or the programmer. The changes in behavior can be entirely preprogrammed, and selection of behavior is a matter of routine, showing very regular, rhythmic patterns. Although such a preprogrammed decision strategy may be useful for creatures living in a very stable and predictable environment, it does not suffice for creatures that have many jobs to do in an environment in which the opportunities to perform a job vary considerably (see Mc Farland, 1981, for the case of a natural creature) (see Maes, 1990b for the case of an artificial creature).

Complex creatures need a more flexible behavior selection mechanism which bases selection on the internal motivational state of the creature as well as on external circumstances. It is clear from observation of animal behavior that a change in the external environment may override the current behavior, for example, with some alarm response behavior. But animals also demonstrate changes in behavior without a change in the external situation, which suggests that behavior selection is also determined by internal motivation. E.g. a domestic hen, when presented with an egg may eat it on certain occasions or brood on it on other occasions (Mc Farland,

A biological visuo-motor system: how dissimilar maps interact to produce behavior

Gerhard Manteuffel
Brain Research Institute (FB 2)
University of Bremen
POB 33 04 40
D-2800 Bremen, FRG

Abstract

Results from electrophysiology, anatomy, and simulations show that the tectal sensori-motor system in salamanders can be understood as a system that preserves retinal coordinates on both tectal hemispheres and establishes a head-centered three-dimensional coordinate system by a combination of the bilateral visual maps. Thus, the outputs of the sensory maps are results of retinal inputs and intertectal signal transfer. The sensory maps are connected to respective motor maps organized in the coordinates of the neck muscles. The bilateral distribution of excitation in the whole system enables the animal to perform head saccades directed toward a stimulus at an arbitrary location within its egocentric frame of reference.

1. The biological system

Research on the system controlling saccades in amphibians has largely concentrated on the visual input side. It is well established that moving stimuli, like squares of moderate size or bars elongated in the direction of movement, elicit saccades of the head with a high probability. The probability is reduced if stimuli of other shapes are used.

As in other vertebrates, saccades are triggered by the optic tectum and can be released by electrical stimulation of this area. Tectal output units are well activated by visual stimuli like squares or horizontal rectangles which are moved across their receptive fields (c.f. Himstedt et al., 1987). The tectum receives an orderly arranged input from the eyes, establishing a retinotopic map on which the nasal visual field is represented rostrally and the temporal visual field caudally; the superior visual field is represented close to the dorsal midline, whereas the ventro-lateral margin is stimulated by visual objects in the lower part of the visual field.

Recently we were able to show that an ipsilateral visual map exists as well. It is established by an information transfer from the other tectal hemisphere (that receives input from the contralateral eye). The ipsilateral map represents only the binocular part of the visual field and, therefore, covers only the rostral half of the optic tectum. The ipsilateral map is pointsymmetrical to the contralateral one

and, thus, constitutes the ability to calculate binocular disparity (Manteuffel et al., 1989).

In comparison to the amount of knowledge on visual properties of tectal neurons, there are only few data on sensori-motor coupling in the saccade control system. It is known that electrical stimulation of the toad's tectum evokes saccades toward a location that roughly corresponds to the retinotopic map established by afferents from the contralateral eye (Ewert, 1967). However, recent stimulation experiments have shown that there exists no simple correlation between stimulation site, stimulation strength and evoked saccade (Jordan et al., 1990). It is likely that the motor map of the tectum in salamanders is not arranged in the coordinates of external space. The bimodal distribution of the populations of efferent neurons in a dorsal and a ventrolateral group with increasing density toward caudal tectal levels (Fig. 1) rather indicates a recruitment system arranged in the coordinates of the neck muscle system (Manteuffel, 1990; Naujoks-Manteuffel and Manteuffel, 1990). According to this principle, larger saccades would be evoked when more premotor neurons become activated. In fact large saccades are necessary toward temporal goals which are stimulating more posterior sites of the optic tectum in the visual domain.

The motor system executing the saccades in salamanders is comparatively simple. The basic structure of the arrangement of the muscles responsible for head movements has most likely been inherited from fish ancestors with a basic equipment of paired epaxial and hypaxial muscles (i.e. above and below the axis of the spinal column). With a joint intercalated between the posterior pole of the head and the first vertebra, vertical and horizontal articulations can occur. Therefore, the resulting head movements can be described best in the coordinates given by the directions of the forces of the two pairs of muscles. The bilateral epaxial muscles (*m. intertransversarius capitis superior*) insert at the ear capsule and the hypaxial *m. rectus cervicis* at the *os triangulare*. Both pairs of muscles insert caudally at transverse processes of the second and third vertebrae.

Saccades are largely ballistic movements in salamanders (Werner and Himstedt, 1985), in general falling short with larger horizontal angles. Therefore two or more saccades are often needed to bring a target into the center of the visual field. The animals approach a prey by executing a saccade followed by a few steps of straight walk. If necessary, this sequence can occur repetitively

Navigating With a Rat Brain: A Neurobiologically-Inspired Model for Robot Spatial Representation

Maja J Mataric
MIT Artificial Intelligence Laboratory
545 Technology Square, #721
Cambridge, MA 02139, USA
maja@ai.mit.edu

Abstract

This paper presents a neurobiologically-feasible spatial representation model. The model was implemented and tested on a physical autonomous mobile robot. It was shown to be both computationally simple and physically robust.

The described model is a possible interpretation of the organization and function of the rat hippocampus. The paper presents relevant biological, psychological, and neurobiological data, and gives a detailed set of comparisons between the physical hippocampus and our "synthetic" rat implementation. The implications of the many similarities are described. Finally, areas for future study in both biology and robotics are suggested.

1 Introduction

Most animals, including humans, spend much of their waking time in transit from one place to another [Waterman 89]. Purposefully moving about requires a system for spatial modeling integrated with the mechanisms for handling navigation, locomotion, and motivation. These systems have evolved to perform with impressive robustness. Understanding their function has long been a goal of cognitive scientists, biologists, and neuroscientists. More recently, this goal has been adopted by members of the Artificial Intelligence and robotics communities interested both in simulating biological systems and designing better artificial ones.

The question asked by both communities is: "What kind of spatial information is stored?" In order to answer it, experiments are designed to test where on the qualitative-to-quantitative scale the representation lies, and whether it is centralized or distributed. This paper describes a qualitative, distributed spatial representation tested empirically on a mobile robot.

2 Cognitive Maps

A *cognitive map* is a generic term for an internal representation of spatial information. The term has come to connote a very analytical, centralized representation. In this paper, we will use the term in its generic meaning, and analyze its variants.

A cognitive map is usually assumed to represent space with a set of landmarks, each of which is an element (object or feature) serving as a point of reference [Presson and Montello 88]. According to Piaget, a landmark is a spatial primitive, and thus a basic building block of spatial representations [Piaget and Inhelder 67]. Although most landmark studies concentrate on visual cues, the concept generalizes to any perceptible feature. Animals construct landmarks from auditory, olfactory, and tactile cues as well [Gould 82], taking advantage of their different characteristics [O'Keefe 89].

In this paper, we will analyze cognitive maps along two dimensions: 1) what information they encode and 2) how they encode it. The "what" dimension can vary from completely qualitative or topological to very quantitative or metric. The "how" dimension varies from totally global or centralized to entirely distributed or decentralized.

2.1 How Qualitative?

The nature of the representation determines the type and number of landmarks required for localizing. In a qualitative representation, an object can be remembered as being proximate to a landmark, defined within a radius around it. On the other end of the spectrum, the position of an object can be computed precisely from the known locations of three landmarks [Pick, Montello and Somerville 88]. The question is how much metric information is recorded.

The psychological literature is divided on this issue. Studies testing response times in object position recall

WHAT IT MEANS FOR ROBOT BEHAVIOUR TO BE ADAPTIVE

BY DAVID MCFARLAND

Dept. Zoology. University of Oxford, U.K.

ABSTRACT

There is an analogy between animal and product design that can be formulated as an exact mathematical analogy. The success of a biological design is measured by the success of the genes that produce it, and this depends upon the ability of those genes to increase their representation in the population in the face of competition from rival genes. Similarly, when a variety of products is under consideration, they vary in the period required for product development, in the chance of failure in the market place, and in the expected returns from sales if the product is successful. The development period refers to the period before any return is achieved on investment. For animals this is the period between birth and reproduction, and for products it is the period prior to time that financial return accrues to the investor. The success of a design is evaluated by the net rate of increase of the genes coding for it (i.e. the return on investment) in the animal case, or, in the case of a product launched into the marketplace, of the money invested in it. If we are to take the biological approach to robot design seriously, then we should first consider the ecological (or market) niche that a proposed robot is to occupy. Is the robot to be a toy, a brick-laying robot, or a bomb-disposal robot? Just as there are no general-purpose animals, so there should be no general-purpose robots. For robot behaviour to be adaptive, in terms of the analogy, it must optimise with respect to the selective pressures of the market place. Other forms of adaptation, such as acclimatisation and learning, are subject to the same criteria.

The term adaptation, as used in biology has a number of meanings: Biologists usually distinguish between (1) evolutionary adaptation, which concerns the ways in which species adjust genetically to changed environmental conditions in the very long term; (2) physiological adaptation, which has to do with the physiological processes involved in the adjustment by the individual to climatic changes, changes in food quality, etc.; (3) sensory adaptation, by which the sense organs adjust to changes in the strength of the particular stimulation which they are designed to detect; and (4) adaptation by learning, which is a process by which animals are able to adjust to a wide variety of different types of environmental change.

Acclimatization to altitude. Adaptive changes in a man breathing rarefied air for 4 days, followed by 6 days at sea level.

V = lung ventilation, E = serum erythropoietin, H = rate of hemoglobin synthesis, R = fraction of red blood cells (after Adolph, 1972)




Fig. 1 Adaption by acclimatisation. The physiological changes that occur in acclimatisation to altitude run through a spectrum, ranging from fast but costly processes to slow-acting processes that are cheap in energetic terms.

Adaptation implies cost reduction, as can be seen from the example in Figure 1. In animal behaviour, real costs relate to Darwinian fitness. So, in considering the usage of the term adaptation in robotics, we should ask if there a concept equivalent to the fitness of a robot?

I will argue that there is such an equivalent concept, and that the analogy between animal and product design can be formulated as an exact mathematical analogy. Briefly, the success of a biological design is measured by the success of the genes that produce it, and this depends upon the ability of those genes to increase their representation in the population in the face of competition from rival genes. How does this

INCREMENTING INTELLIGENT SYSTEMS BY DESIGN

Brendan McGonigle
Laboratory for Cognitive Neuroscience
Edinburgh University
Scotland

CHARACTERISATION

SUMMARY

Our operating hypothesis can be stated simply enough. It is that the class of non-trivial solutions to the problem of information management is small (see McGonigle, 1987). If true, it follows that intelligent systems (however instantiated) are solution constrained *a priori* and must converge on a similar design logic if they are to succeed. In the case of biologically instantiated intelligence, this convergence is seen more as a primary consequence of such solution constraint (for example, the power-generalality trade-off), and less one of genetic affinity (although it may be secondarily a factor).

Our general goal is the characterisation of intelligent systems in the broadest sense, informed in particular by converging research in comparative, developmental and cognitive psychology situated within a neuroscience framework (see McGonigle and Chalmers, 1990a). More recently, we have opened a dialogue with roboticists at Edinburgh (Smithers, Malcolm and Donnett, in particular) and are delighted to see the growing basis for productive dialogue. The position outlined here also intersects with one espoused by Brooks (1986) but has an independent origin and rationale.

The goal of this paper is to summarise some of the characterisations which emerge from work on biological systems in an attempt both to cross-check with designers of artificial systems and to exchange concepts of possible mutual benefit. To limit the vast area under review, I shall concentrate on the incremental aspect of intelligent systems as this is, as I see it, the key issue.

Although there are (and have been) many documented approaches to the study of evolutionary intelligence, most have failed on the key issue of how systems 'invest in complexity' or grow (ontogenetically or phylogenetically) from 'weak to strong'. Classical ethology, for example, has worked best with simpler 'reactive' systems which 'wear their adaptation' on their sleeve¹. Certainly, their domain of inquiry has precluded the study of human problem solving and intelligence either from a comparative, a cognitive or a developmental standpoint. As a consequence, ethologists may offer some solace to designers interested in making simple reactive agents as a first step. However, and crucially for our current agenda, they offer few general characterisations of intelligent systems that could afford strong clues as to the ways and means

¹ A psychologist looking at these reactive systems as a group, however, might characterise their group adaptation as an example of the power-generality trade-off, both in terms of niche/habitat selection and in the specialisation of subsets within the group as in (say) bees. However, a study of learning mechanisms *per se* as an alternative to reactive agents, fares little better. Motivated by black box behaviourism, the search for universal laws of learning and memory has failed to provide any stable correlates of species differentiation according to phylogeny or brain architectures. Although learning appears early in the evolutionary process, it is not the fact of learning, but WHAT is learned that differentiates animals. And if simple habit formation is mainly what has been achieved in most of these experiments, it is not surprising that they have not picked up direct implications of advanced nervous systems, for, as Mishkin (1985) has claimed, such habit mechanisms are mediated by the motor cortex (an 'old' area) of the brain.

Simulation of Adaptive Behavior in Animats: Review and Prospect

Jean-Arcady Meyer * and Agnès Guillot **, **

* Groupe de BioInformatique
Ecole Normale Supérieure, 46 rue d'Ulm
75230 Paris Cedex 05 - France
Université Paris 6 - URA 686 CNRS
meyer@frulm63.bitnet

** Université Paris-X
200 avenue de la République
92000 Nanterre Cedex 01 - France
guillot@frulm63.bitnet

ABSTRACT

Following a general presentation of the numerous means whereby animats -i.e. simulated animals or autonomous robots- are enabled to display adaptive behaviors, various works making use of such means are discussed. This review cites 172 references and is organized into three parts dealing respectively with preprogrammed adaptive behaviors, with learned adaptive behaviors, and with the evolution of these behaviors. A closing section addresses directions in which it would be desirable to see future research oriented, so as to provide something other than proofs of principle or ad hoc solutions to specific problems, however interesting such proofs or solutions may be in their own right.

1. INTRODUCTION

In a changing, unpredictable, and more or less threatening environment, the behavior of an animal is adaptive as long as the behavior allows the animal to survive. Under the same conditions, the behavior of a robot is considered to be adaptive as long as the robot can continue to perform the functions for which it was built. Now, the survival of an animal is intimately involved with its physiological state and the successful operation of a robot depends upon its mechanical condition. Under these circumstances, it is obvious that one can associate with an *animat* - whether the term indicates a simulated animal or an autonomous robot (Wilson, 1985, 1987a) - a certain number of state variables upon which its survival or successful operation depend, and that each of these state variables is characterized by a range of variation within which the animat's continued survival or operation are preserved. Such variables were referred to as *essential variables* by Ashby (1952) long ago. Their variation ranges describe a viability zone inside the given state space, and the animat can be referenced at any instant by a point within this zone (Figure 1). Under the influence of environmental or behavioral variations affecting the animat, the corresponding reference point moves and may at times come close to the limits of the viability zone. In this case, the animat's behavior can be called adaptive so

long as it avoids transgressing the viability boundary (Ashby, 1952; Sibly & McFarland, 1976).




Figure 1. Viability zone associated with two essential variables, V1 and V2. The animat's behavior is adaptive because corrective action has been taken at point B, so as to avoid crossing out the corresponding viability zone at point A.

Such behavior can be generated by means of several different or complementary abilities and architectures. For example, the laws governing the animat's operation may rely upon various homeostatic mechanisms thanks to which, if the reference point alluded to earlier moves away from an adapted point of equilibrium -adapted because it is suitably located within the viability zone-, this process tends to return it to its original position, thereby decreasing the risk that it will pass outside the limits of the zone. Other ways in which to lower this risk involve the use of high-quality sensory organs or motor apparatus that allow the animat to detect as early as possible that it is approaching these limits and/or to move away from them quickly and effectively. In this line of reasoning, it is obvious that the equivalent of a nervous system is mandatory in order to connect the animat's perceptions with its actions and that reflex circuits activated as quickly as possible increase the adaptive nature of its behavior. It is

ON THE FEASIBILITY OF COMPUTATIONAL ARTIFICIAL LIFE

A REPLY TO CRITICS

Harold C. Morris
P.O. Box 657
Los Alamos, NM
87544 USA

Abstract

The thrust of this paper is to meet the objections of what may be termed a philosophical school, whose principals are Robert Rosen, Howard Pattee, and Peter Cariani. The objection that a computational universe is a flat "pseudo-world," because it is "all syntax and no semantics," is inquired into and refuted, as is the claim that nothing really new can evolve within such an artificial universe. It is concluded that no persuasive reasons have been advanced as to why computational artificial life is not feasible.

A convergence of several fields has resulted in the new discipline of Artificial Life ("AL") research. And just as AL science has several sources, it is moving in several distinct directions. Life-like entities are being developed as biochem-

ical "wetware," robotic hardware, and as computer software. This paper shall, however, deal only with controversies surrounding the third variety, namely the computational AL form.

Generally stated, the AL program is to develop life-like organisms in the medium of choice. For myself and some other AL researchers, the computer is our medium of choice. Our objective is to implant or evolve individuals or colonies in automaton universes, to observe instances of propagation, adaptation, or communication, such as one usually associates with life forms.

Since the pioneering work of von Neumann (1966), cellular automata have been much used as computational media for AL research. In recent years, more sophisticated systems, for example

Mapbuilding using Self-Organising Networks in "Really Useful Robots"*

Ulrich Nehmsow and Tim Smithers†
Department of Artificial Intelligence
University of Edinburgh
5 Forrest Hill
Edinburgh EH1 2QL
Scotland
Telephone 031 - 667 1011 ext 2529/2517
Telex 727442 UNIVED G
E-mail:ulrich@uk.ac.ed.edai tim@uk.ac.ed.edai

September 1990

Abstract

The Really Useful Robots (RUR) project is seeking to understand how robots can be built that develop and maintain the task achieving competences they require for flexible and robust behaviour in variable and unforeseen situations, as opposed to these being installed by their designers. In this paper we present an experimental autonomous robot with a map building competence which uses a self-organising network. Map building forms a necessary step on the way to development of a navigational competence. Some encouraging initial test results are also presented.

1 Introduction

The traditional approach to control in (mobile) robots is to decompose the task into separate components, and implement these using standard control techniques, see [Levi 1987], for example. This we call an *analytical approach*. Alternatively, a control structure can be built 'bottom up', first building foundational competences (such as 'move around and avoid obstacles'), and later on top of these more complicated competences (such as

'explore', 'map building', and 'map using'). This we call a *synthetic approach*, see [Brooks 1986], for example.

At Edinburgh we have adopted a synthetic approach in what we call the 'Really Useful Robots' project (RUR) [Nehmsow et al 1989]. This project is attempting to develop a control architecture which supports the development of task achieving competences by the robot. In other words, we are trying to understand how a robot can sequentially acquire and maintain the behavioural competences it requires, rather than have them 'installed' by us as its designers. We believe that this *autonomous acquisition* of task achieving competences will lead to greater flexibility and robustness in the behaviour of robots with respect to variable and unforeseen situations. In investigating this idea we are motivated and informed by the adaptive control mechanisms we see in simple animals which result in them having flexible, reliable, and robust competences well matched to the tasks they are responsible for achieving and to the environment in which they are exercised.

Trying to get a robot to acquire the skills it needs means that as many decisions as possible are left to the robot, rather than being predefined by the designer. Alder, the first of the 'Really Useful Robots' (see figure 6)¹, is able to adapt to a changing environment, and to acquire useful competences. It uses what we call *fixed* and *plastic* components in its control architecture

¹Alder is a mobile robot whose base is built with a Fischertechnik kit. It is about 25cm long, has an 8082 based microcomputer on board (16k RAM) and is equipped with up to seven tactile sensors plus odometer. In addition to this a sonar sensor is available, but has not been used to obtain the results presented in this paper. More information about Alder and the 'Really Useful Robots' approach can be found in [Nehmsow et al 1989].

*©U. Nehmsow and T. Smithers, May 1990

†Names appear in alphabetical order, with both being principal authors on this occasion.

THE EVOLUTION OF BEHAVIOR: SOME EXPERIMENTS.

JAN PAREDIS

RIKS

Postbus 463

NL-6200 AL Maastricht,

The Netherlands

ABSTRACT.

This paper investigates the evolutionary development of (problem solving) behavior. Through evolution, artificial animals learn to survive in a given world.

We use layered neural networks (NNs) as the substrate on which evolutionary learning operates. The fault tolerance of neural networks allows for a genotype / phenotype distinction which maintains the variation in the genetic pool. Furthermore, we define building blocks which take into account the functionality of the NNs.

The result of the algorithm can be inspected at two levels. First, there is the behavior of the individual animals. A description of their behavior is obtained through the induction of decision trees which describe the functionality of the NN. Second, the behavior of the population as a whole can be described. The distribution of the animals over the world often provides an analogical representation of a problem solution.

Keywords: autonomous agents, evolutionary learning, genetic algorithms, inductive learning, machine learning, neural networks.

1. Introduction.

Natural evolution continues to intrigue mankind. Particularly, the complexity to which it leads often surprises us. Examples of this complexity are abundant. Animals in a prey - predator relation, for example, develop complex defensive and offensive behavior, such as camouflage and

imitation. Insect colonies (e.g. colonies of ants, bees, wasps, termites etc.) are another instance of the complexity to which evolution can give rise. Such complexity is obvious if one examines the nests termites build or the social organisation within insect colonies [Wilson 85]. Many other examples can be found in [Dawkins 86] or [Tinbergen 65], amongst others.

Computational methodologies based on the evolutionary metaphor have been developed for a wide-range of problems, such as search, optimization and machine learning. One of the most notable methodologies are genetic algorithms (a good overview of GAs can be found in [Goldberg 89]). Another, related, methodology is evolutionary learning, which searches through a space of behaviors using the principles of variation and selection. In contrast with GAs, evolutionary learning does not require an explicit, domain specific fitness measure. We only need to specify the characteristics of the environment (e.g. the amount of food present) and the effects of an animal's actions on itself and on the environment.

In this paper we propose an evolutionary framework in which successive generations of animals learn to improve their chance of survival in a given environment. Or, in other words, successive generations adaptively develop behavior (such as: look for food, avoid predators ...) in correspondence with their natural needs. In order to be successful, an animal has to find an answer to the following question: When should I perform which actions? This question is answered through evolutionary learning on NNs. Our system learns only through evolution, no life-time learning mechanism is incorporated. We should stress that our primary goal is to obtain adequate (problem solving) behavior. The

The Animat and the Physician

Alexandre Parodi

Indodyne SARL

10, rue de la Paix

75002 Paris

Fax: (33) (1) 40 20 98 98

Tel: (33) (1) 69 81 88 41

E-mail: (uucp) parodi@lri.lri.fr

Pierre Bonelli (LRI)

Equipe I&A

LRI, bâtiment 490

Domaine Universitaire de Paris-Sud

91405 Orsay cedex

Tel: (33) (1) 69 41 64 07

E-mail: (uucp) bonelli@lri.lri.fr

Abstract

Animat research has already produced interesting concepts and algorithms. In this paper, we analyze how this research can be applied to human intelligence understanding and to reproducing of some expert behaviors. To support these ideas, we experiment with an improvement of Boole, a Genetic Based Learning algorithm from animat research, in a medical domain of expertise. We experimentally demonstrate that our system obtains good results on a well known realistic medical diagnosis task, and we analyze its potential ability to solve more complicated problems.

Introduction

There has been much debate about how one can consider that a system is intelligent, most of the time according to how it processes information (rules, neurons etc.) in connection with the human brain. However, in [Wilson, 85] Wilson developed the idea that we could probably learn more from ethology, and he introduced the concept of animats, which are autonomous systems which learn how to survive and expand in a given environment.

We propose to discuss how this research can be profitable to the understanding of human intelligence, and how animat algorithms can be used to reproduce some intelligent human behavior.

1. Intelligence hierarchy

Let us consider the following intelligent systems hierarchy based on how explicit the input knowledge from the environment must be: systems that learn by

being told, systems that learn by complete examples, and animats, i.e. systems that learn by reward.

1st level: Systems that learn by being told

Most computers get their knowledge by being given programs, i.e. a list of instructions to be executed in a specified order; the processing is completely explicit in the input knowledge.

Production systems without learning ability get their knowledge from rules which are used to reason about the input data and conclude about the output data to provide. They are somehow more intelligent because the order in which rules are executed depends on the data: control is data driven. This means that a few rules implicitly specify many different reasoning traces.

Specifically built neural networks get their knowledge from a set of predetermined weights which indicate how a formulated hypothesis (micro-feature represented by one neuron) influences another. However, without learning, they are not really much more intelligent than production systems, but some experiments tend to show that they are less brittle and noise sensitive and can exhibit even richer behavior than can be expected because they use parallel analog processing.

2nd level: Systems that learn by complete examples

Learning systems can manage with even less formalized knowledge: they only need examples which contain both the input and the corresponding desired output; they take care of extracting the appropriate knowledge that is needed to generalize the sampled behavior to new inputs. Such systems can be either rule based systems, connectionist networks, or classifier systems. A number of learning algorithms

Learning Hill-Climbing Functions as a Strategy for Generating Behaviors in a Mobile Robot

David Pierce

Department of Computer Sciences
University of Texas at Austin
Austin, TX 78712
(dmpierce@cs.utexas.edu)

Benjamin Kuipers

Department of Computer Sciences
University of Texas at Austin
Austin, TX 78712
(kuipers@cs.utexas.edu)

Abstract

We consider the problem of a robot with uninterpreted sensors and effectors which must learn, in an unknown environment, behaviors (i.e., sequences of actions) which can be taken to achieve a given goal. This general problem involves a learning agent interacting with a reactive environment: the agent produces actions that affect the environment and in turn receives sensory feedback from the environment. The agent must learn, through experimentation, behaviors that consistently achieve the goal. The difficulty lies in the fact that the robot does not know *a priori* what its sensors mean, nor what effects its motor apparatus has on the world.

We propose a method by which the robot may analyze its sensory information in order to derive (when possible) a function defined in terms of the sensory data which is maximized at the goal and which is suitable for hill-climbing. Given this function, the robot solves its problem by learning a behavior that maximizes the function thereby resulting in motion to the goal.

1 The credit assignment problem

The learning problem addressed in this paper is illustrated in Figure 1. The learning agent, which we are calling a "critter," receives sensory input (vector s) from the world and acts on the world via motor outputs (represented by a , the action vector). In addition, the critter has access to a reward signal, r , by which it knows when it has achieved its goal. (In the experiments discussed later, the reward signal is incorporated into the sense vector for simplicity.) The critter's task is to learn a behavior which reliably achieves the goal. This behavior is a sequence of actions (most likely dependent on the concomitant sequence of sense vectors) which takes the




Figure 1: The general problem is for the learning agent, the "critter," to learn sequences of actions which produce rewards. The critter is rewarded when it is in a goal state.

critter from its present state to the goal state. The problem is difficult because the reward signal does not provide feedback for every action. The critter only knows that it has done the right thing when it stumbles onto the goal and is rewarded. It is then faced with the credit assignment problem - the problem of deciding which actions led to the goal.

2 A solution

In this paper, we propose the following solution to this problem:

1. Derive a function defined in terms of the sense vector (which is itself a function of the state of the world) such that this function is maximized at the goal state and is suitable for hill-climbing. It may in some cases be impossible to find such a function, in which cases, the method fails.
2. Learn a behavior that does gradient ascent on this hill-climbing function.

The problem explored in this paper can be viewed as the problem of learning a hill-climbing function to replace an *a priori* function where the latter is not appropriate for gradient ascent. In general, this problem can be described as follows: There is some function, given a

Animal Simulations with *Logo: Massive Parallelism for the Masses

Mitchel Resnick

Massachusetts Institute of Technology
Cambridge, MA 02139 USA
mres@altdorf.ai.mit.edu

Abstract

Experimentation with animal simulations is limited, in large part, by the difficulty of converting ethological ideas into computer programs. *Logo is a new programming language that aims to make it easier for non-expert programmers (researchers as well as students) to develop and modify their own simulations. *Logo is designed specially for simulating colony-level behaviors—that is, group behaviors that emerge from interactions among hundreds or thousands of individual creatures. Unlike most simulation languages, *Logo gives the creatures' environment an equal computational status to the creatures themselves. Users write rules for creatures and for "patches" of the environment, then observe the higher-level behaviors that result. A sample *Logo program shows how local, parallel actions among ants can lead to spatial-extended and temporally-sequential patterns in the colony-level behavior.

1. Introduction

During the past several years, a growing number of researchers have begun creating computer-based simulations of animal behavior. Some are motivated by ethological goals: they hope to gain a better understanding of the mechanisms underlying the behaviors of real animals. Others are motivated by engineering goals: they hope that simulations of animals will provide ideas (or at least inspiration) for building computers and robots that function more effectively in the world.

Unfortunately, designing and programming animal simulations typically requires significant programming expertise. Most animal simulations are developed as customized programs, by experienced programmers. Although there are several new

languages and tools designed to facilitate the development of animal simulations, even these tools are meant primarily for experienced programmers.

As a result, it is difficult for non-experienced programmers to convert ethological ideas into computer simulations. Certainly, novice programmers can change parameters or initial conditions on existing simulations. But they are not able to make more serious modifications or create entirely new simulations. In short, animal simulations are still not for the masses.

What is needed is a new type of programming language that allows people to more easily create and experiment with animal simulations. This paper describes a language that aims to do just that. The language, called *Logo (pronounced *star-logo*), is designed especially for simulating "colony-level behaviors"—that is, group behaviors that emerge as large numbers of individual animals interact with one another, as in bird flocking or ant foraging. These types of simulations are particularly difficult to construct using traditional, sequential programming languages. Indeed, simulations of colony-level behaviors highlight the need for a new "massively parallel" approach to programming, in which many "computational creatures" act in parallel (at least conceptually, if not in reality).

Section 2 describes the audience for *Logo. Although *Logo is designed primarily to help students explore self-organizing phenomena, it could serve equally well as a tool for ethologists. Section 3 discusses the computational requirements for programming colony-level animal simulations. Section 4 discusses the central ideas underlying the design of *Logo, including the decision to treat the creatures' "world" as an active computational actor, equal in status to the creatures themselves. Section 5 presents examples of *Logo simulations. A simulation of ant foraging, for example, shows how local, parallel actions by hundreds of individual ants can result in spatially-extended and temporally-sequential behaviors by the colony as a whole. Section 6 describes future directions for *Logo.

Lookahead Planning and Latent Learning in a Classifier System

Rick L. Riolo
University of Michigan
Ann Arbor, MI 48109 USA
rick_riolo@um.cc.umich.edu

Abstract

Classifier systems (CSs) have been used to simulate and describe the behavior of adaptive organisms, animats, and robots. However, classifier system implementations to date have all been reactive systems, which use simple S-R rules and which base their learning algorithms on trial-and-error reinforcement techniques similar to the Hullian Law of Effect. While these systems have exhibited interesting behavior and good adaptive capacity, they cannot do other types of learning which require having explicit internal models of the external world, e.g., using complex plans as humans do, or doing "latent learning" of the type observed in rats. This paper describes a classifier system that is able to learn and use internal models both to greatly decrease the time to learn general sequential decision tasks and to enable the system to exhibit latent learning.

1 Introduction

Classifier systems (CSs) have been used to understand, through metaphor and simulation, the behavior of adaptive organisms and robots from animats ([Holland and Reitman, 1978], [Booker, 1982], [Wilson, 1985]) to rabbits [Holyoak et al., 1990] to humans [Holland et al., 1986]. However, all CSs implemented to date have been reactive systems, i.e., they store and use knowledge as rules of the form "If the situation is X, do A" (where X may describe a set of world states). In the terms of animal psychology, the system stores S-R associations [Walker, 1987]. These systems all have used a trial and error learning technique, the *bucket brigade algorithm (BBA)*,¹ to assign priorities to rules; the priorities determine which rules will fire and what the system will do in a given situation. The BBA is reminiscent of the Hullian "Law of Effect," i.e., rules active when reward is received (from the environment or from

subsequently active rules) have their priorities modified in proportion to the reward. (They use a more sophisticated learning algorithm, the *genetic algorithm*, to form generalizations over the space of situations, i.e., to form concepts.) These CSs use only a very simple *model* of the world, in which a rule's priority in effect *predicts* the reward expected if that rule is fired. Despite using such simple representational and learning techniques, CSs have shown surprisingly interesting behavior when controlling animats which must learn and adapt to survive in simple environments; they have also been used to solve concept learning [Wilson, 1987a], dynamic control [Goldberg, 1988], and sequential decision problems ([Grefenstette, 1988],[Booker, 1989]).

A second type of learning involves building more complex models which not only predict rewards, but also *predict world states*.² These models can be implemented as rules of the form "If the situation is X, and I do A, then expect situation Y," i.e., S-R-S associations. Systems can create and update these models continuously, even when no rewards are being received, by predicting the (non-reward) outcomes of actions and then modifying the model when the predictions are incorrect. That is, rather than just using the usually infrequent feedback provided by rewards or punishments to build simple S-R models, systems can exploit the flood of non-reward experiences they have to build much more complete models of the world. Predictions of expected states then can be integrated with motivations and predictions of rewards to choose actions that lead to goals. Using internal models enables systems to reduce the number of trials required to learn tasks ([Sutton, 1990], [Whitehead and Ballard, 1989]). Further, a model enables a system to easily integrate newly acquired knowledge about the world or about changes in goals or motivations ([Holland et al., 1986], [Dickinson, 1980]). Internal models also have been used to simulate Piagetian cognitive development during infancy ([Drescher, 1986], [Drescher, 1989]).

¹The bucket brigade algorithm is a temporal difference method [Sutton, 1988].

²Unless stated otherwise "internal model" and "model" will refer to this more complex type of model.

Cognitive Action Theory as a Control Architecture

H. L. Roitblat

Department of Psychology
University of Hawaii at Manoa
Honolulu, Hawaii, USA 96822
herbert@uhccux.uhcc.hawaii.edu

The three laws of robotics: (1) No robot may harm a human being, or through inaction, allow a human being to come to harm. (2) A robot must obey the orders given it by human beings except where such orders would conflict with the First Law. (3) A robot must protect its own existence as long as such protection does not conflict with the First or Second laws. Asimov, 1950/1977, p. 40.

Abstract

Standard versions of control theory approach their limits in autonomous robotic control because of their conceptualization in terms of a fixed mapping between environmental variables and behavior. Cognitive action theory, in contrast, views behavior as hierarchically organized by a network of interacting nodes. Nodes at different levels represent action with different degrees of abstraction. Activation of nodes in the hierarchy is controlled by potentiation and inhibition received from other nodes and by environmental stimulus information. Learning consists of the formation of connections between nodes as the result of locally available information about the satisfaction of cybernetic feedback functions. Such networks are capable of planning and executing highly flexible behaviors in complex, uncertain environments.

1. Introduction

Among the difficult problems for designers of autonomous robot systems are control of (a) nonlinear

systems, (b) stochastic systems, (c) systems requiring sensor fusion and feedback, and (d) systems with incomplete information about the environment and its structure (Meystel, 1988). Many of these problems are exacerbated by the control-theory-dominated approach that robotic designers have taken. Although control theory has been remarkably successful in many domains, it approaches its limits in autonomous robotic control, because in its standard versions it assumes that behavior can be adequately conceptualized as a system with a fixed mapping between environmental variables and behaviors. It typically characterizes the behavior of the robot in terms of movements or operations that are precisely defined as specific responses to specific environmental conditions (Brooks, 1986). This approach works well in situations in which variability is limited, goals are simple and noncompetitive, uncertainty is minimal (or at least statistically characterizable), and the environment and actions appropriate to it are fairly exhaustively known (e.g., a factory environment). The approach is likely to prove ultimately inadequate, however, in highly variable environments, in the face of complex competing demands, and in situations in which no complete set of situation-response rules is available a priori.

Biomimetic Sonar Processing: From Dolphin Echolocation to Artificial Neural Networks

H. L. Roitblat¹, P. W. B Moore², P. E. Nachtigall², & R. H. Penner²

¹Department of Psychology, University of Hawaii, Honolulu, HI 96822 USA
herbert@uhccux.uhcc.hawaii.edu

²Naval Ocean Systems Center, Hawaii Laboratory, Kailua, Hawaii, 96734 USA
pmoore@nosc.mil

Abstract

Analysis of animal performance can provide important cues about the design of automated artificial biomimetic systems. On the basis of behavioral observations, we have been developing models of dolphin echolocation ability that have applicability to the design of biomimetic sonar systems. A dolphin was trained to perform an echolocation delayed matching-to-sample task. The clicks the animal generated during task performance were recorded and digitized along with the echoes returned by the stimulus objects. The dolphin's performance was then modeled using artificial neural networks.

1. Biomimetics

The study of animals can provide a very important adjunct to formal analyses in the design of automated systems such as robots and autonomous vehicles. Animals have evolved in a real world, solving real problems, such as gathering and interpreting essential information. Evolution supports the emergence of solutions that are well adapted to the animal's ecological niche, but provides no guarantee that the evolutionary solutions an animal derives are the best possible solution to a given problem (see Gould & Lewontin, 1979). Evolution merely asserts that, in light of the competing demands presented by the animal's evolutionary history, its ecology, and its other needs, a solution (vis a vis the whole organism and all its adaptations and constraints) is

better, or at least no worse, than any other that its ancestors had achieved (Roitblat, 1987).

Although formal analyses have undoubtedly been successful in developing solutions to many automatic process problems such as those encountered in designing robots, many other problems have resisted solution. Solutions to scientific and engineering problems are inspired by many sources, but are ultimately derived from the intuition of the engineer, as formalizations of folk physics, folk psychophysics, folk psychology, etc. Folk science is the set of generally held beliefs that people employ in their ordinary activities. For example, many college students recognize that a ball rolled out of an inclined tube will fall some distance in front of the tube. Most of these same students, however, often mistakenly believe that if they drop a ball while walking, the ball will fall directly under the point at which it was released (McCloskey, 1983; see also Holland, Holyoak, Nisbett, & Thagard, 1986).

Systems involving falling balls are well analyzed so anyone with training in mechanics can see clearly the difference between the folk beliefs concerning falling balls and formal scientific beliefs (we may call these "facts"). Our intuitions have been trained to correspond with the analyzed facts, rather than with unanalyzed apparent perceptions. In situations in which formal scientific analyses have not yet been fully applied, we have no assurance that our scientific intuitions similarly avoid the pitfalls of our naivete. A creative scientist or engineer will apply his or her folk-science intuition to a difficult

A Possibility for Implementing Curiosity and Boredom in Model-Building Neural Controllers

Jürgen Schmidhuber*

Institut für Informatik

Technische Universität München

Arcisstr. 21, 8000 München 2, Germany

schmidhu@tumult.informatik.tu-muenchen.de

Abstract

This paper introduces a framework for 'curious neural controllers' which employ an adaptive world model for goal directed on-line learning.

First an on-line reinforcement learning algorithm for autonomous 'animats' is described. The algorithm is based on two fully recurrent 'self-supervised' continually running networks which learn in parallel. One of the networks learns to represent a complete model of the environmental dynamics and is called the 'model network'. It provides complete 'credit assignment paths' into the past for the second network which controls the animats physical actions in a possibly reactive environment. The animats goal is to maximize cumulative reinforcement and minimize cumulative 'pain'.

The algorithm has properties which allow to implement something like *the desire to improve the model network's knowledge about the world*. This is related to *curiosity*. It is described how the particular algorithm (as well as similar model-building algorithms) may be augmented by dynamic *curiosity* and *boredom* in a natural manner. This may be done by introducing (delayed) reinforcement for actions that increase the model network's knowledge about the world. This in turn requires the model network *to model its own ignorance*, thus showing a rudimentary form of *self-introspective behavior*.

1. Introduction

In the sequel first an on-line algorithm for reinforcement learning in non-stationary reactive environments is described. The algorithm heavily relies on an *adaptive model of the environmental dynamics*. The main contribution of this paper (see the second section) is to demonstrate how the algorithm may be naturally augmented by *curiosity* and *boredom*, in order to improve the world model in an on-line manner.

Consider an 'animat' whose movements are controlled by the output units of a neural network, called the control

network, which also receives the animat's sensory perception by means of its input units. The animat potentially is able to produce actions that may change the environmental input (external feedback caused by the 'reactive' environment). By means of recurrent connections in the network the animat is also potentially able to internally represent past events (internal feedback).

The animat sometimes experiences different types of reinforcement by means of so-called *reinforcement units* or *pain units* that become activated in moments of reinforcement or 'pain' (e.g. the experience of bumping against an obstacle with an extremity). The animat's only goal is to minimize cumulative pain and maximize cumulative reinforcement. The animat is autonomous in the sense that no intelligent external teacher is required to provide additional goals or subgoals for it.

Reinforcement units and pain units are similar to other input units in the sense that they possess conventional outgoing connections to other units. However, unlike normal input units they can have *desired activation values* at every time. For the purpose of this paper we say that the desired activation of a pain unit is zero for all times, other reinforcement units may have positive desired values. In the sequel we assume a discrete time environment with 'time ticks'. At a given time the quantity to be minimized by the learning algorithm is $\sum_{t,i} (c_i - y_i(t))^2$ where $y_i(t)$ is the activation of the i th pain or reinforcement unit at time t , t ranges over all remaining time ticks still to come, and c_i is the desired activation of the i th reinforcement or pain unit for all times.

The reinforcement learning animat faces a very general spatio-temporal credit assignment task: No external teacher provides knowledge about e.g. desired outputs or 'episode boundaries' (externally defined temporal boundaries of training intervals). In the sequel it is demonstrated how the animat may employ a combination of two recurrent *self-supervised* learning networks in order to satisfy its goal.

Munro [2], Jordan [1], Werbos [12], Robinson and Fallside [6], and Nguyen and Widrow [4] used 'model networks' for constructing a mapping from output actions

*This work was supported by a scholarship from SIEMENS AG

Robot Ethology: A Proposal for the Research into Intelligent Autonomous Systems

Uwe Schnepf

Expert Systems Research Group

German National Research Center for Computer Science (GMD)

P.O.Box 1240

5205 Sankt Augustin 1, FRG

usc@gmdzi.uucp

Abstract

In this paper it is argued that studying the behaviour of (even simple) animals and human beings from an Ethologist's point of view will provide a basis for understanding human cognition. A simple model of the organisation of behavioural sequences in animals as described by Niko Tinbergen [1] and Konrad Lorenz [2] is presented as an starting point to develop intelligent autonomous systems. The relation to current research into behaviour-based Robotics is shown (cf. Brooks [3]) and essential extensions to the behavioural model such as optimisation procedures based on genetic algorithms and evolution technology, a framework to link basic sensory-motor skills to higher-order categorical perception as described by the symbol grounding problem and more advanced models of the organisation of behaviour are presented. In particular, self-organisational processes are advocated as a key feature in achieving intelligent behaviour of autonomous robots.

1 Introduction

Initially, Robotics has been the most attractive field for researchers in Artificial Intelligence (AI) in order to study the whole range of cognitive capabilities of human beings. Defined as the intelligent connection between sensing and acting, Robotics was supposed to naturally pose the questions one has to answer in order to understand human intelligence.

Researchers focussed on various aspects of the sensing-to-acting chain using a knowledge-based approach to specify the information necessary for the robot to perform various tasks. In general, this information consists of a world model and the knowledge about possible transitions from one world state into another due to the actions of such an intelligent agent.

Since a robot is usually located in an unstructured environment, world modelling requires us to impose structure on this environment using an explicit description of the objects surrounding the robot. The problem of which aspect of the environment should be modelled and which aspects of the robot-world-interactions should remain constant or not became one of the most important questions in this domain, well-known as the Frame problem [4].

Another problem which arises when using a world model is the precision of the robot's internal representation of the surrounding world. Since the control algorithms of today's robots are based on the transformation of three-dimensional coordinates in order to perform the required actions, the internal representation of the world which serves to interpret sensor data and to choose the appropriate actions (the representation of which in turn has to be translated into sequences of three-dimensional trajectories of the robot's actuators) must be precise in order to allow an exact mapping from the external world to the internal representation and vice versa.

The data received from the sensors is used to update the internal world model. But as it is difficult to gain precise three-dimensional description from currently available sensors, it was claimed that the current sensor technology was not sufficient to deliver precise updates of the model.

Because of this lack of information, it is argued, robots are not able to perform appropriate and flexible actions. Once more precise sensors have been developed, the problem of incorporating more precise world knowledge into the internal representation could be solved more easily.

We do not agree with this opinion and are convinced that the inflexibility and subtlety of current robot controllers are inherently based on the approach chosen. Making a decision in favour of such a knowledge-based approach to world modelling naturally leads to the problems discov-

Some Parallels Between Associative Learning and Object Classification

David R. Shanks

MRC Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EF
England

Abstract

Certain parallels are explored between the mechanism of associative learning in animals - and its description by contemporary learning theories such as that of Rescorla and Wagner (1972) - and object classification, which can be construed as the problem of learning to associate visual features or micro-features with a category. A specific associative learning theory of classification is presented. While *prototype* effects can be easily accommodated within associative theories of classification, *exemplar* effects appear to be fatal for such accounts, since no explicit representations of exemplars are stored in associative networks. An experiment is reported which attempted to see whether the model would be able to reproduce exemplar effects. Surprisingly, it could. In examining why this was the case, a new interpretation of exemplar effects emerged.

1. Introduction

Some years ago, it was very fashionable to assume that people represented categories of visual objects in terms of their *prototypes*. As a result of learning that a large number of different visual stimuli all belong to a certain category - such as the category *dog* - people were assumed to have extracted a prototype representing the central tendency of the stimuli on a variety of feature dimensions.

Figure 1 (left-hand panel) illustrates how this abstraction is supposed to work. Imagine a number of exemplars (marked by X's) of a category which vary on, say, two dimensions. For instance, the exemplars could be dogs varying in colour and size. According to prototype theories, what is actually extracted and mentally represented of this category is the central tendency of the exemplars within the feature space. The prototype (marked by a dot) has a value on each dimension corresponding to the modal value of the actual exemplars on that dimension.

What happens when a new stimulus is presented for

classification? According to prototype theories, the similarity of the test stimulus to the prototype is determined - this being just the inverse of the distance between them - and the stimulus is classified as being a member of this category if its similarity to the *dog* prototype is greater than its similarity to any other category prototypes. The greater the similarity to the prototype, the faster, more confident, or more accurate the classification decision.

The crucial commitment, then, of prototype theories is that during the process of classification, no representations of specific training exemplars play a role in classification. It is only the abstracted prototype which determines the course of classification. Of course, representations of specific training exemplars clearly do exist: thus each of us has memory traces of specific faces that we are familiar with - our mother's face, for example. Prototype theories maintain, however, that when a new physical stimulus appears before me and I decide that it is a face, the trace that exists in memory of my mother's face plays no role in this process. It is only the prototypical face that I have abstracted from all the faces that I have ever experienced that determines my classification decision.

Now, there is one piece of empirical evidence which appears to provide strong encouragement for the prototype view, and this is the fact that subjects often respond more accurately, or with greater confidence, to the prototype of a category than they do to the specific training exemplars, even though they may never have seen the prototype before. A clear example of this occurs in an experiment by Knapp and Anderson (1984). They generated training stimuli, which were dot patterns, by distorting a particular prototype pattern. Subjects saw a number of distortions of the prototype, and learned correctly to classify all of these patterns. When they were subsequently tested either with new distortions, the original training patterns, or with the prototype itself, they responded most accurately to the prototype (I shall call this the *prototype effect*). This was most evident when the number of different training stimuli used had been large. Of course, this result is precisely what one would expect on a prototype account, since the prototype pattern corresponds to what the

An investigation of two mediation strategies suitable for behavioural control in animals and animats

Martin Snaith and Owen Holland

Technology Applications Group
Bolam's Mill
ALNWICK
NE66 1LN
United Kingdom
+ 44 665 604895 (phone)
+ 44 665 510624 (telefax)

Abstract

This paper gives details of recent experiments to determine the characteristics of two simple mediation strategies. We have chosen to compare two mutual inhibition strategies; the standard 'n-flop' model [Ludlow 76][Shackelford 89], and a scheme suggested by McFarland [McFarland 65]. To study the comparative features of each network we use TAG's 'non-spiking' neuron model (SNF block model [Snaith 89a][Snaith 89b][Holland & Snaith 90a]) both in simulation and on one of our Hilda series of mobile robots to show the mediation between simple 'fight' and 'flight' behaviours set up in conflict.

1. Introduction

A number of problems in animal behaviour centre around the situation where there are a number of stimuli present and a number of appropriate responses possible, only one of which can take place at any time. Which response (or behaviour) will occur? and when, and under what circumstances, will it be supplanted by another? In behavioural analysis variations in response have been attributed to set, fatigue, habituation, attention, displacement, and so on. At the physiological level, mechanisms have been proposed for mediating the interactions of a number of individual neurons, or a number of pools of neurons, each of which controls a behaviour, so that one or another is temporarily dominant.

This paper examines the consequences of making some very simple assumptions about the requirements which the mechanism must satisfy, and examines two possible implementations on paper, on the bench, and in a mobile robot.

2. Conflict and Resolution

We have already assured that the situation is one in which several stimuli are present. Let us further assume that each stimulus is linked to a separate behaviour, which is appropriate to the stimulus. For example, there may be food and water present, and behaviours for eating and drinking to which they are linked. The first step in considering any restrictions on the nature of any controlling mechanism might usefully be to imagine that no control mechanism is present. What will happen if the behaviours occur simultaneously? There appear to be three main possibilities:

- a). The behaviours operate satisfactorily together eg: I can walk and chew gum at the same time.
- b). The behaviours fail to operate satisfactorily together because they require effectors to carry out incompatible actions - eg: I cannot eat and roar at an enemy simultaneously.
- c). The behaviours fail to operate satisfactorily together because the effects of one behaviour - eg: flight - make it impossible to continue with another - eg: mating.

Let us confine our attention to situations (b) and (c). While it is true that (c) may bring a sort of resolution to the conflict between the two behaviours, it is worth noting that there is no guarantee of this. An oscillation between the two behaviours might occur on an arbitrary time scale (flee for a second, mate for a second). This could also occur in (b), and might even in some circumstances be an effective strategy, amounting to time sharing [McFarland 73 [McFarland 74]. However, both the robot designer

TOWARDS A THEORY OF EMERGENT FUNCTIONALITY

Luc Steels

Artificial Intelligence Laboratory

Vrije Universiteit Brussel

Pleinlaan 2, B-1050 Brussels Belgium.

email: steels@arti.vub.ac.be

Abstract

The paper sets the first steps towards a theory of emergent functionality. The theory tries to make explicit what emergent functionality is by contrasting it with hierarchical functionality. It analyzes the principle advantages of this approach and conjectures a formal structure common to systems with emergent functionality.

1. Introduction

Emergent functionality means that a function is not achieved directly by a component or a hierarchical system of components, but indirectly by the interaction of more primitive components among themselves and with the world. Emergent functionality has become one of the main themes in research on Artificial Life (Langton, 1988) and autonomous agents (Brooks, 1989). So far, engineers and scientists have used their intuitions to build systems that exhibit emergent functionality but there is no explicit theory yet on what emergent functionality is, how it can be achieved, when it is appropriate and why. This paper reports on research to understand the principles underlying emergent functionality and how it can be used for designing and building systems. We first discuss hierarchical systems to make the specific properties of systems with emergent functionality stand out more clearly. Then we turn to emergent functionality itself. Examples are discussed and the advantages of emergent functionality are analyzed. The final part of the paper conjectures a formal structure that seems common to systems with emergent functionality.

2. Hierarchical systems

2.1. Characterization of hierarchical systems

In hierarchical systems there is a direct relationship between structure and function. The system consists of a set of components. Each of these components has three aspects: (1) inputs and outputs, (2) a control element to turn the component on or off, (3) a well-determined functionality. A component stands on its own in the sense that its functionality can be tested independently from the other components. Moreover this functionality is a recognizable subfunction of the global functionality of the system. For example, the motor of a car needs fuel. The tank realizes a subfunction of fuel supply, namely to hold the fuel. The tank realizes this function independently from the other components like the pipes that transfer fuel to the motor or the accelerator that regulates the flow.

Some components are specialized in obtaining input from the environment. Others are concerned with output and actions in the environment. There are also components whose major role is the control of the operation of other components. So the different components interact in two ways: (1) There is flow of data through input/output relations between components. (2) There is flow of control when one component turns another component on or off.

Because the components function independently of each other they can be constructed and put together in a modular fashion: The global functionality is decomposed into different subfunctions. A subfunction is either directly realized by a particular component or it is further decom-

Reinforcement Learning Architectures for Animats

Richard S. Sutton
GTE Laboratories Incorporated
Waltham, MA 02254
sutton@gte.com

Abstract

In the first part of this paper I argue that the learning problem facing animats is essentially that which has been studied as the reinforcement learning problem—the learning of behavior by trial and error without an explicit teacher. A brief overview is presented of the development of reinforcement learning architectures over the past decade, with references to the literature.

The second part of this paper presents Dyna, a class of architectures based on reinforcement learning but which go beyond trial-and-error learning. Dyna architectures include a learned internal model of the world. By intermixing conventional trial and error with hypothetical trial and error using the world model, Dyna systems can plan and learn optimal behavior very rapidly. Results are shown for simple Dyna systems that learn from trial and error while they simultaneously learn a world model and use it to plan optimal action sequences. We also show that Dyna architectures are easy to adapt for use in changing environments.

1 Animats and the Reinforcement Learning Problem

What is an Animat? An animat is an adaptive system designed to operate in a tight, closed-loop interaction with its environment. An animat need not be a learning system, but often it is; some sense of adaptation of behavior to variations in the environment is required.

Figure 1 is a representation of the animat problem as I see it. On some short time cycle, the animat receives sensory information from the environment and chooses an action to send to the environment. In addition, the animat receives a special signal from the environment called the *reward*. Unlike the sensory information, which may be a large feature vector, or the action, which may also have many components, the reward is a single real-valued scalar, a number. The goal




Figure 1. The Reinforcement Learning Problem facing an Animat. The goal is to maximize cumulative reward.

of adaptation is the maximization of the cumulative reward received over time.

This formulation of the animat problem is the same as that widely used in the study of *reinforcement learning*. In fact, reinforcement learning systems can be defined as learning systems designed for and that perform well on the animat problem as described above. Informally, we define reinforcement learning as learning by trial and error from performance feedback—i.e., from feedback that evaluates the behavior generated by the animat but does not indicate correct behavior. In the next section we briefly survey reinforcement learning architectures.

One might object to the problem formulation in Figure 1 on the grounds that all possible goals have been reduced to a scalar reward. Although this appears limiting, in practice it has proved to be a useful way of structuring the problem. Some examples of goals formulated in this way are:

- Foraging: Reward is positive for finding food objects, negative for energetic motion, slightly negative for standing still.
- Pole-balancing (balancing a pole by applying forces to its base): The reward is zero while the pole is balanced, and then becomes -1 if the pole falls over or if the base moves too far out of bounds.
- Towers of Hanoi: Reward is positive for reaching the goal state.

Can Allied Reflexes Promote the Integration of a Robot's Behavior?

by

Philip Teitelbaum, Vivien C. Pellis, and Sergio M. Pellis

Psychology Department, University of Florida, Gainesville, FL 32611

If we wish a robot to behave adaptively, it may be useful to design it according to principles enabling animals and people to do so. We use brain damage and drugs in animals to break behavior down into its subcomponents, and stages of recovery to reveal levels of reintegration. By partial transection in the brain, or by drugs, we appear temporarily to inactivate central motor programs involved in spontaneous behavior (20). As in earlier work using complete transection (19), behavior breaks down into reflexes; spontaneous environmentally-directed behavior is absent. But by our procedures, allied reflexes operate as intermediate-level submodules. By studying how they interact in recovery, and by analyzing controls over individual reflexes, some principles emerge, perhaps useful for robot design. In addition, approaching the subject via phenomena produced by pathology may yield insight into imperfections in a robot that may arise from imbalance in adaptive systems.

(1) Allied Reflexes can Act as Isolated, Adaptive Submodules

Extensive lateral hypothalamic (LH) damage produces a simplification of motivated behavior, that we have called, after Magnus (10), a "zero-condition". All spontaneous environmentally oriented behavior is temporarily abolished. The animal lies motionless, virtually comatose. However, its autonomic system remains relatively intact -- if such an animal is tube-fed, it lives and recovers (20). Within a couple of days, somnolence usually disappears, but the animal remains for several days in a state of catalepsy and akinesia, symptoms often seen in Parkinson's disease (16).

For instance, an LH-damaged cat remains

for many minutes with its forelimbs spread widely apart, or one foreleg placed up on its back at quite an extreme angle. Similarly, a cataleptic rat remains with its hindlegs on a raised platform, forelegs on the floor, in an awkward downward-tilted posture. These symptoms support the generally held view that dopamine-deficiency produces an inability to initiate movement. It therefore seems paradoxical when, as shown in fig. 1, a rat made cataleptic by the dopamine receptor blocker haloperidol will, if pushed from behind, leap vigorously into the air (2, 12). This paradox is resolved by detailed analysis of the rat's responses leading to the jump. As the cataleptic animal is pushed forward (Fig. 1A), it braces against such displacement by shifting its weight backwards. When its hind legs begin to slip (B), a leap is triggered (C and D), away from the surface where it is unstable. When it lands on the horizontal table top, it immediately resumes immobility. Thus, cataleptic leaping is merely an allied postural support defensive reflex, triggered by postural instability. The animal does not suffer from a general inability to initiate movement -- its isolated support submodule simply does not do so when the animal is in a stable posture, even if awkward.

The static postural support submodule consists of an aggregate of allied reflexes (including standing, crouching, bracing, clinging, stability-related stepping, righting, and jumping), all of which homeostatically maintain support or regain upright unmoving stability. It is isolated from other submodules, individually involved in locomotion, turning, head-scanning, orienting, and ingestion, which are inactivated (20). When a submodule is isolated, inhibitory controls are

Task differentiation in *Polistes* wasp colonies : a model for self-organizing groups of robots

Guy THERAULAZ *, Simon GOSS **,
Jacques GERVET * and Jean-Louis DENEUBOURG **

* CNRS - UPR 38, 31 chemin Joseph Aiguier, 13402 Marseille cedex 09, France

** Unit of Theoretical Behavioural Ecology, CP 231, Université Libre de Bruxelles, 1050 Bruxelles, Belgium

Abstract

In this paper we described a simplified model for a case of functional self-organization. It deals with the emergence of a particular form of task assignment and parallel hierarchical organization within a social group which depend basically on the interactions occurring between individuals and with their immediate local surroundings. The task organization within the colony appeared to be a distributed function which does not require the presence of an individualized central organizer. We discussed how such elementary processes could potentially be applied in the coordination and self-organization of groups of interacting robots with simple local computational properties to perform a wide range of tasks.

Introduction

In eusocial insect societies, all the individuals must cooperate to perform a certain number of tasks the nature of which depends on the internal needs of the colony, as well as on the particular environmental conditions. At any time each individual can act and interact either with other individuals or with their environment and thus causes changes in the state of the group. The group is nevertheless the focal point for a stable, self-regulated organization of individual behaviors. The study of the processes leading to the emergence of a stable collective order in insect societies has recently emphasized the importance of individual interaction dynamics (Deneubourg et al., 1987; Deneubourg and Goss, 1989; Goss et al., 1990; Beckers et al., 1990). This research has demonstrated that quite simple elementary rules of individual behavior often make it possible for the society to create surprisingly complicated patterns and to make efficient decisions when certain types of external constraints are encountered.

Our biological study examined the processes involved in task assignment in primitive *Polistes* wasp colonies (Theraulaz et al., 1990 a, b and c). *Polistes dominulus* is a species which is to be found in temperate, northerly regions. These

wasps had two advantages from the point of view of our study, since they build their nest with no envelope and their colonies contain a relatively small number of individuals ($n = 20$), which makes it possible to observe all the members of a colony over a period of time. These primitively eusocial species have little individual differentiation and no morphological differences between castes or predetermined control of activities depending on age or on any other known physiological predetermination. The integration and coordination of individual activities therefore depend largely on the interactions which take place among the members of these societies and on the immediate relationships between a society and its environment.

A two fold morphogenetic process thus occurs within the colony :

- Each individual acquires its own behavioral profile which is characterized by all the observable stable behavioral items in which it takes part. All the profiles can be described by a reduced number of behavioral forms to which the various individual profiles belong (see Theraulaz et coll., 1990 a).

- In a society at a given moment in time, the whole set of individual behavioral profiles does not constitute a random sample of all the possible profiles; they constitute a profile configuration which can be defined by the proportion of individuals belonging to each of the behavioral forms.

The way this configuration is controlled by internal and external constraints on the colony constitutes the task assignment process. The model presented here aims at describing the task assignment process in a hierarchically structured society.

ETHOLOGICAL AND PSYCHOLOGICAL MODELS OF MOTIVATION- TOWARDS A SYNTHESIS

Frederick Toates,

Biology Department,
The Open University,
Milton Keynes
MK7 6AA,
United Kingdom.

Per Jensen,

Section of Ethology,
Swedish University of Agricultural Sciences,
POB 345,
Skara, S-53227,
Sweden.

Abstract

By looking at a variety of individual motivational systems, the first steps towards reconciliation of models in the psychological and Lorenz traditions is made. A model that contains features of both emerges. The relevance of this model to animal welfare issues is discussed.

1.0 Introduction

It seemed once that study of motivation theory was in serious decline; in psychology and ethology, the days of the grand theories (e.g. Hull, Lorenz) appeared to be over (Toates, 1986). Psychologists lost interest in the topic and ethologists had moved in their droves to the greener pastures of optimal foraging and sociobiology. Then applied ethology came to the rescue; a renewal of interest in motivation theory was prompted by considerations of animal welfare, suffering and the associated recommendations for legislation (Dawkins, 1980). To assess when an animal is suffering or why animals in captivity might spend much time in various bizarre activities, not observed in wild-living conspecifics, one needs motivation theory. However, applied ethologists encountered a major difficulty: fragmentation and contradictions in the literature. To some extent, psychological and ethological approaches differ in their assumptions and areas of interest. However, differences do not divide neatly along party lines. Thus, ironically, some writers in the psychological tradition (e.g. Gallistel, 1980; Glickman and Schiff, 1967; Herrnstein, 1977; Hogan, 1967; Toates, 1986) argued for the applicability of the best-known ethological model, that of Lorenz (1950), whereas, to some ethologists, it has been used for little more than to illustrate how motivation doesn't work (Dawkins, 1986; Archer, 1988). The issue is more than of academic interest; recommendations for animal husbandry depend upon which model one believes (Baxter, 1983; Hughes and Duncan, 1988a,b).

Some of the conflict seems to arise from semantic confusion, loose use of language and the assumptions of different perspectives, and we attempt to resolve some of this. We suggest where common ground can be found. First, it is necessary to try to sanitize the vocabulary.

1.1 The meaning of motivation

One reason for confusion seems to be the meaning different authors assign to the concept of "motivation". Even within psychology, the concept is used in very different contexts and meanings by different researchers (Toates, 1986). In psychology, "drive" and "motivation" are overlapping and sometimes identical concepts, while in ethology, since the influential work of Hinde (1959; 1960), "drive" has assumed a role much like a swear-word in church, while "motivation" has survived as a well-respected term. The term "motivation" is found in most ethology books, but rarely defined in a way that enables one to know what is meant by claiming that an animal is more or less motivated for certain behaviour. Although mainly used to discuss how responses may vary in different contexts, "motivation" is usually only a heading for the chapter in which this is found. For example, Slater (1985) explicitly takes this broad view, in saying, while discussing the variability of responses in animals: "...discovering just what it is that leads them to behave differently from one time to another presents some very interesting problems...These are the problems of motivation or, in other words, the mechanisms leading animals to do what they do when they do it." McFarland and Sibly (1975) provided a formal representation of motivation, the state-space approach. Although primarily introduced as a powerful alternative to unitary drive concepts, they also tried to tidy up the terminology. Thus, "motivational state" is the state value of all causal factors influencing a setup of functionally related behaviour patterns. The motivational state maps onto a "tendency" which is the strength of the behaviour in the competition for the motor apparatus (the "final common path").

This may be a convenient and fairly unambiguous way of talking about motivation, but for many psychologists, it doesn't relate to their conception of motivational problems. In the terminology of McFarland and Sibly, all behaviour of any animal will by definition be guided by a motivational state, determining its tendency. This leads to the danger of a circular argument, since by the definition, an animal will always be under the control of the behaviour for which the tendency is strongest, i.e. for which it is most motivated. So, if all behaviour reflects a hypothetical motivational state, which can only be deduced from observations of the behaviour performed, we run into

Exploring Adaptive Agency II: Simulating the Evolution of Associative Learning

Peter M. Todd

Geoffrey F. Miller

Department of Psychology

Jordan Hall Building 420

Stanford University

Stanford, CA 94305

todd@psych.stanford.edu

geoffrey@psych.stanford.edu

Abstract

We consider psychology as the study of *adaptive agency*, investigating the processes and mechanisms resulting in fitness-increasing behavior in the world. A central issue in psychology so construed becomes: what are the relations between the primary adaptive process of evolution by natural selection, and the adaptive processes psychologists call 'learning'? In particular, under what conditions would learning evolve? To explore this issue, we use genetic algorithms to simulate the evolution by natural selection of neural networks, which in turn control the behavior of simple creatures in virtual environments. We have developed what we consider the simplest possible environmental challenge in which unsupervised associative learning could prove adaptive: 'bootstrapping' the learned use of one highly accurate, but individually varying, sensory modality by another less accurate, but evolutionarily stable, modality. We have found a possibly quite general U-shaped curve relating the time (in number of generations) to evolve the use of unsupervised learning on the varying 'bootstrapped' modality, to the accuracy of perception in the stable modality which guides this learning. This U-shaped curve appears to represent a trade-off between the adaptive pressure to evolve learning (which peaks when perception accuracy in the stable guiding modality is at chance) and the ease of learning during a given lifespan (which peaks when this accuracy is perfect.)

1 Introduction

Natural selection has constructed animals' minds and behaviors for adaptive fit to the environmental problems they must face. As the study of such minds and behaviors, psychology should focus on the notion of *adaptive agency* -- the generation of action in the world in response to challenges to individual fitness. This framework encompasses many approaches, including (1) the elucidation of complex species-typical adaptations (as in human and animal experimental psychology and cognitive ethology), (2) the comparison of psychological adaptations across species and

consideration of their phylogenetic origins (as in comparative psychology), and (3) the general exploration of the adaptive processes themselves that yield adaptive agency (e.g. by simulation methods, including those in the field of artificial life -- see Langton, 1989). In the current paper, we consider the phenomenon of 'learning' as an aspect of adaptive agency, by investigating via evolutionary simulations some conditions under which the ability to learn may prove adaptive and so spread through a population. The theoretical, historical, and methodological background for this work is presented more extensively in Miller and Todd (1990), and further extensions of this method applied to habituation and sensitization as adaptations to short-term environmental dynamics appear in Todd and Miller (in press).

Evolution as an adaptive process has itself undergone changes: "survival of the stable" probably preceded "survival of the fittest" (Dawkins, 1976). Evolution in the earth's early environment is likely to have selected for replicating systems with relative stability in the shifting primordial soup. After stability came replication and metabolism: the ability to turn external material into copies and extensions of oneself. The evolution of larger, more complex phenotypes then allowed the evolution of behavior-generating systems that could produce innately programmed sequences of activity and movement. Sensory systems could then evolve to guide these behavior-generators more adaptively, based on sensitivity to particular environmental cues. Thus, blind activity may have preceded reactivity -- the ability to adaptively adjust to the current changing environment on a moment-by-moment basis. Only after these first two stages had evolved could a further adaptive process evolve -- 'learning,' defined as the ability to make long(ish)-term adaptive changes in behavior-generators in response to the environment.¹ In this theoretical framework, learning emerges not as the primary adaptive force some have assumed it to be, but rather as a tertiary one, following genotypic evolution and short-term environmental reactivity (see also Shepard, 1987, 1988). Once we re-conceptualize 'learning' as merely one process among several that generate adaptive agency, the questions we might ask about this process begin to change as well.

¹ By this definition, learning includes such processes as experience-guided development not commonly included in this category. For examples of such processes, see Knudsen, 1988; Singer, 1984, 1988; and Stein et. al., 1989.

Computer Simulation of an Animal Environment

Toby Tyrrell and John E. W. Mayhew
(e-mail: toby@aivru.uucp, jewm@aivru.uucp)
AIVRU, University of Sheffield.

Abstract

This paper describes a computer simulation of an animal environment which has been created as a tool for investigating the mechanisms behind 'behavioural choice' in animals. The simulated environment has been designed to provide sufficient complexity and realism for meaningful behavioural experiments.

1 The Problem

"All brains, even those of the tiniest insects, generate and control behaviour." [Albus, 81]. This basic task of animal brains can be split into three subtasks (e.g. [Brooks, 86]) as shown in Fig. 1:

- 1 - Sensing of the environment so as to be able to perceive what is going on at each moment in time (perception).
- 2 - Taking the interpretation of the environmental situation and using it to decide which of the animal's repertoire of behaviours is the most appropriate (behavioural choice).
- 3 - Transforming the chosen behaviour into a pattern of movements of parts of the body (motor control).

Our aim is to investigate the mechanisms underlying the second of the three stages outlined. This stage is perhaps the least well understood of the three due to the fact that the processes involved are internal and cannot be directly observed, only inferred from the resulting behaviour. The behavioural parts of a brain do not interact directly with the outside world, but rather through the interfacing systems of perception and motor control [Halliday, 83].

It should be noted that the use in this paper of the term 'behaviour' (e.g. in 'behavioural choice') refers to a pattern of actions such as eating, mating, avoiding predators, etc. This should be distinguished from the usual use of the term 'behaviour' in the AI literature, where it more often denotes the coordinated move-




Figure 1: Three functions of an animal brain.

ment of a set of limbs (e.g. in order to grasp an object, navigate around an obstacle, etc). We want to look at high-level decision-making (e.g. should the animal obtain food from the nearby fruit bush or else flee from the predator that has just appeared in the distance). We do not want to address the problem of how the animal should move its limbs in order to pick the fruit and transfer it to its mouth or how the animal should best move its legs so as to be able to run away from the predator.

In short, we want to examine the second of the three functions of Fig. 1, while ignoring the other two as much as possible.

2 Why a Simulated Environment?

Given that we want to examine behavioural choice mechanisms, what is the best method of going about it?

One approach to examining behaviour has involved the building of robots which can navigate in a

**THE CAUSAL ANALYSIS OF AN ADAPTIVE SYSTEM:
SEX-RATIO DECISIONS AS OBSERVED IN A PARASITIC WASP
AND SIMULATED BY A NETWORK MODEL**

Martijke Vonk, Felix Putters & Berend-Jan Velthuis
Department of Zoology, Div. of Ethology, University of Leiden,
P.O. Box 9516, 2300 RA Leiden, The Netherlands,
e-mail: SBW1MV@HLERUL57

ABSTRACT

The computer metaphor, based on rule and symbol manipulation, is challenged by our connectionist model. Neural nets, Parallel Distributed Processing systems, etc. are capable to achieve goals by way of cooperative activity and can ignore all gadgets that are required in the computational models based on the digital computer metaphor. The consequences of our structure-oriented modelling approach for causal and functional explanations in ethology are discussed.

KEY WORDS

Neural nets, PDP, motivation, adaptation, parasitic wasps, sex allocation

1. INTRODUCTION

Ethologists use the modelling approach in order to reveal the mechanisms underlying overt behaviour. Relevant principles of organization are tested by embodying them in a mathematical model, the behaviour of which is compared to experimental data. The serial computer metaphor and systems theory have set the fashion for some time in modelling information processing, such as involved in problems of decision-making. Models based on the computer metaphor manipulate bits of data in a formal way, according to preset rules and operations which are specified in programs that were designed for a given task (van Rhijn 1977; van Rhijn & Westerterp-Plantinga 1989; Coderre 1989; Travers 1989). This way of modelling involves a distinction between system hardware and software, and postulates a central processor that operates on the data and drives the system. These procedures are attractive from a methodological point of view, in that one knows the complete set of assumptions necessary to make these constructs work as desired. But, what is contributed to causal analysis when we apply this approach? In practice, one postulates homunculi with the self-same capacities that the theory sets out to explain. Functionally defined concepts take up key positions in the

computer model, but these concepts lack a proper (albeit potential) backing of causal mechanisms that could perform the assumed tasks. The approach of cybernetics and system theory (Toates 1986) may provide appropriate solutions when dealing with simple phenomena (e.g., orientation), but their theoretical framework is not suited for handling more complex phenomena. Moreover, many models based on this conventional modelling approach happen to be goal-directed.

The condition of an 'explicit goal-representation' makes conventional modelling an inappropriate technique for the assessment of decisions that are made by our experimental animals: parasitic wasps. Many parasitic wasps are known to adjust the offspring sex ratio to characteristics of the environment, to maximize reproductive success. They are capable to optimize the F1 offspring sex ratio: the actual ratio sons - daughters is such that the maximum number of gene copies will be present in the F2 progeny. In order to explain the actual decisions of our wasps (but we believe that this point can be generalized, see McFarland 1989), we require systems which adapt to the environment without an explicit representation of the world, and which can achieve goals without a representation of these goals (i.e., without goal directedness). Conventional explanatory concepts are therefore of no use when applied to the decisions made by our experimental animals.

Recent developments in the field of Artificial Intelligence have expanded the possibilities to develop ethologically relevant models: systems of semi-autonomous information-processing entities, whose local interactions with one another are controlled by a set of simple rules. Such systems do not contain rules for their behaviour at the global level. The observable behavioural output and its complex dynamics are emergent properties, which develop from the local interactions of the low-level entities (Langton 1989). Connectionist models perform without an explicit representation of a goal or an environment. This kind of modelling is the tool for exorcising the homunculi from ethological theory.

2. BIOLOGICAL BACKGROUND

Sex allocation theory presents one of the finest oppor-

Eight Principles for Building an Intelligent Robot*

David L. Waltz

Thinking Machines Corporation
245 First Street
Cambridge, MA 02142
USA

Brandeis University
Computer Science Department
& Center for Complex Systems
Waltham, MA 02254 USA

Abstract

We cannot expect to know the detailed "wiring diagram" of the nervous system for any intelligent creature for quite a long time. Even then, the true organization is likely to be incredibly complex and tangled. However, in order to build intelligent robots now, we need a plausible interim architecture. A functional model for robot organization is proposed, starting with a basic, first order model, which is gradually refined. In particular, it is proposed that associative memory provides a useful -- and perhaps plausible -- basis for an intelligent system.

0. Introduction

While remarkable progress is being made by neuroscientists in unraveling portions of the nervous system (see, for example, [Kosslyn 89] or [Halgren 87] for insights into the visual system and memory systems, respectively), we are still far from being able to map the wellsprings of action, intention, and decisions. Other researchers have investigated abstract models of adaptation and learning, such as genetic algorithms and classifier systems [Holland 77], or the SOAR system [Newell 87]; abstract models have been used to build explicit models of creatures (e.g. the Animat [Wilson 87]). [Drescher 89] has introduced the "schema mechanism," and his ideas have much in common with the proposals below, especially in his views on chaining, and in his key ideas on identifying and learning reliable schemas, using large amounts of statistical analysis. "Subsumption architecture" researchers in AI (e.g. [Brooks 86], [Maes 90]) hope to arrive at intelligent systems by first building a (layered) system with the abilities of, say, a cockroach, and adding yet more control layers to eventually reach greater and greater intelligence. This work is broadly within a "Society of Mind"-type theory that views intelligence as composed of a very large number of independent agents and "bureaucracies" of agents, each responsive to specific situations or patterns [Minsky 87]. While I subscribe in general to the Society of Mind view, I believe that it is both possible and valuable to divide up the model of mind somewhat differently than is done within subsumption architectures.

I propose here a model of a robot's "mind" whose components are divided up along very different lines, somewhat analogous to principal components analysis: the first component is a general associative memory model that captures general patterns and principles of behavior, while successive components add refinements, culminating in society of mind-like demons that recognize very specific situations or patterns, and then override (by priming or inhibiting) more general behaviors. Intermediate refinements include control structures that allow search and chaining of actions, as well as rote learning and generalization. Such a model fits neatly on any massively parallel computer architecture (e.g. [Hillis 85]), but can also be simulated on serial computers (though perhaps not fast enough to allow real-time performance, except in the simplest of environments).

1. Principle One

Use associative memory as the overall organizing conception.

Basic associative memory operations can capture the essence of what intelligent entities do: select relevant precedents in any situation, and act on them. "Precedents" can be actions, options, reminders, etc. This type of operation, akin to case-based reasoning (CBR) [DARPA 88, 89] and memory-based reasoning (MBR) [Stanfill & Waltz 86] is easily programmed on a massively parallel machine, and has found useful applications [Waltz 90]. A number of techniques can be used to find "relevant" items, including nearest-neighbor algorithms, and majority votes of n nearest neighbors.¹

If only a single precedent is close to the current situation (as when the robot is operating in a familiar environment on a familiar task), then little more than an associative memory is needed in order to act intelligently. Only when two or

*This work was supported in part by the Defense Advanced Research Projects Agency, administered by the U.S. Air Force Office of Scientific Research under contract #F49620-88-C-0058.

¹What makes a neighbor "near" is a very subtle issue, and the key open problem in CBR and MBR.

Hierarchical Dishabituation of Visual Discrimination in Toads¹

DeLiang Wang

Michael A. Arbib

Center for Neural Engineering
University of Southern California
Los Angeles, CA 90089-2520, U.S.A.
wang@rana.usc.edu, arbib@pollux.usc.edu

Abstract

Habituation is a basic form of learning in which animals come to respond less and less to repeated presentation of a given stimulus. Studies of habituation in *Aplysia* have yielded important insights into basic mechanisms of synaptic plasticity, while studies in humans reveal stimulus-specific habituation with mutual dishabituation by pairs of different stimuli. Studies in toads suggest a new phenomenon which leads us to new models of vertebrate learning which are subject to experimental test. Instead of mutual dishabituation for different worm-like stimuli, toads exhibit a dishabituation hierarchy, in which stimulus A may dishabituate B, but not vice versa. We offer a model of this hierarchy in which the toad's visual discrimination is reflected in different firing rates in some higher visual center, hypothetically anterior thalamus. This theory, developed through neural simulation based on an extensive model of toad retina, predicts that retinal R2 cells play a primary role in the discrimination while R3 cells refine the feature analysis by inhibition.

The theory predicts new dishabituation hierarchies based on reversing stimulus-background

ground contrast and shrinking stimulus size. After the predictions were made, several were tested by behavioral experiments. In particular, we selected a pair of stimuli whose ordering in the dishabituation hierarchy we predict to be changed by contrast reversal, and the experimental result is as predicted. A size shrinking prediction failed to be validated, and further experiments suggest that visual pattern discrimination in toads is relatively unaffected by stimulus size. Finally, we discuss new insights into a network, offered as a preliminary model of the medial pallium, that can express the dishabituation hierarchy of toads.

1. Background

Habituation is an elementary form of learning in which response to a stimulus will diminish with repeated presentation of the stimulus if there is no punishment or reward associated with the presentations. In the marine mollusc *Aplysia*, habituation has been much studied for insight into molecular mechanisms of synaptic plasticity (see Bailey and Kandel, 1985, for a review). Here, habituation seems to be independent of the specific patterning of stimuli, whereas habituation in mammals is stimulus specific. Given two different patterns A and B (e.g., two tones of different pitch, volume, or duration), the animal can exhibit this specificity by the phenomenon of *dishabituation* (i.e., stimulus B can release behavior despite habituation to A). Moreover, this dishabituation is mutual in mammals: If stimulus A can dishabituate stimulus B, then stimulus B can dishabituate stimulus A (Thompson and Spencer 1966; Sokolov 1975).

¹ The research described in this paper was supported in part by grant no. 1R01 NS 24926 from the National Institutes of Health (M.A.Arbib, Principal Investigator). We wish to express our sincere thanks to Prof. J.-P. Ewert with whom one of us (W.-D.L.) conducted the experiments that greatly furthered the modeling reported here.

SIMULATION STUDIES OF SONG LEARNING IN BIRDS

James M. Williams & P.J.B. Slater

Department of Biology & Preclinical Medicine,
University of St Andrews,
Fife KY16 9TS, U.K.

Email : ZORJW@uk.ac.st-andrews.sava
& ZOSPS@uk.ac.st-andrews.sava

Abstract

Song birds learn their songs, but the accuracy of the copying process varies. As a result the songs present in an area change with time.

Our evidence to date suggests that various aspects of chaffinch song in the wild can be accounted for by a simple random model in which individuals learn the songs that they sing from various adults. The distribution of songs between repertoires of different birds is best matched by simulations with a 15% copy-error rate. This rate of error, combined with a realistic mortality rate of 40%, also gives a good approximation to the changes in the song types present in a population with time.

Simulations have been used to examine the distribution of song types between individuals in a population. When the simple situation, where all birds had a single song type and four neighbours, from one of whom new birds copied their song was examined, small groups of birds sharing a song type were found, as in some dialect species. An extension of this approach to simulate variations in repertoire size or in numbers of neighbours has recently shown that both these factors can have a strong effect on the sharing of song types and their distribution in the population. If a bird chooses the commonest song type sung by its neighbours rather than one of them at random very large groups of birds can occur.

These simulations suggest that the complex distributions of song types often found in wild bird populations may result simply from random copying processes which are not always exact.

1. Introduction

Learning plays a role in the song development of all songbirds studied to date (see review by Slater 1989). In many cases the copying of song takes place from neighbours when young birds first set up their territories so that birds on adjacent territories tend to share songs, while the songs of those further apart are less similar. While the learning can be remarkably accurate, so that the songs of two individuals are often identical, there is good evidence from a number of studies that inaccuracies of copying may lead to new forms of song arising (e.g. Jenkins 1978, Slater & Ince 1979). These "cultural mutations" may be the reason why the songs present in a given area change with time, and why there are also differences in song between localities. Much of the geographical variation in song is complex, especially in cases where each individual has a repertoire of several different song types and these are not all learnt as a package from one other bird. However, in some species dialect areas have been described in which groups of birds share the same song type or types and are separated from each other by more or less sharp boundaries. Whether or not this mosaic pattern has any functional significance is a matter of a good deal of controversy (see Baker & Cunningham 1985).

2. Song in the Chaffinch

We have used computer simulation mainly to supplement our studies of song distribution in the chaffinch (*Fringilla coelebs*), a small European songbird. Its song has been extensively studied over many years, starting most notably with the work of Marler (1952) on song in the wild and the laboratory studies of Thorpe (1958) on song

The Animat Path to AI

Stewart W. Wilson
The Rowland Institute for Science
100 Cambridge Parkway
Cambridge, MA 02142 USA
(wilson@think.com)

Abstract

A research methodology is proposed for understanding intelligence through simulation of artificial animals ("animats") in progressively more challenging environments while retaining characteristics of holism, pragmatism, perception, categorization, and adaptation that are often underrepresented in standard AI approaches to intelligence. It is suggested that basic elements of the methodology should include a theory/taxonomy of environments by which they can be ordered in difficulty—one is offered—and a theory of animat efficiency. It is also suggested that the methodology offers a new approach to the problem of perception.

1. Introduction

There are two broad approaches to the scientific understanding of intelligence, or how mind arises from brain. One is the natural science approach, analyzing and experimenting with phenomena of life, mind, and intelligence as they exist in nature. In this there are two main branches: physiology and especially neurophysiology, in which living systems are subject to detailed internal investigation; and experimental psychology, including studies of animals, in which living systems are studied through their external behavior. Related to the latter, but more observational, are fields such as linguistics and anthropology.

In contrast, the second broad approach to intelligence may be termed synthetic and computational, in which the objects studied are constructed imitations of living systems or their behavior. In "Computing machinery and intelligence", Turing (1950) suggested two possible directions for the computational approach:

We may hope that machines will eventually compete with men in all purely intellectual fields. But which are the best ones to start with? Even this is a difficult decision. Many people think that a very abstract activity, like the playing of chess, would be best. It can also be maintained that it is best to provide the machine

with the best sense organs that money can buy, and then teach it to understand and speak English. This process could follow the normal teaching of a child. Things would be pointed out and named, etc.

Turing's first proposed direction led to "standard AI" or computational cognitive science. Standard AI is basically competence-oriented, modelling specific human abilities, often quite advanced ones. However, while many AI programs exhibit impressive performance, their relevance for the understanding of natural intelligence is, in several respects, limited.

In addressing isolated competences, AI systems typically ignore the fact that real creatures are always situated in sensory environments and experience varying degrees of need satisfaction. Furthermore, the systems attach less importance to such basic natural abilities as perception, categorization, and adaptation than they do to algorithmic processes like search and exact reasoning. This leads eventually to problems connecting the arbitrary symbols used in internal reasoning with external physical stimuli ("symbol grounding" (Harnad, 1990)), and "brittleness" (Holland, 1986), the tendency for AI systems to fail utterly in domains that differ even slightly from the domain for which they were programmed.

AI systems also have an arbitrariness: it is often not clear why one program that exhibits a certain intellectual competence is to be preferred over some other one exhibiting the same competence, especially since the field has not agreed on—or too much sought—a clear definition of intelligence. In a sense, the programmer's facility for imitating a high-level fragment of human competence is a kind of trap, since from a natural science perspective there is usually no strong relation to nature.

Turing's second proposal, for a "child machine", received, over forty years, little attention or resources, perhaps because it seemed fantastic. Yet the child machine was to be situated from the start in a real sensory environment and was to learn through experience. It would have emphasized precisely the abilities that standard AI minimized. Turing's proposal is in fact

A Von Neumann Approach to a Genotype Expression in a Neural Animat

Dan Wood

ThinkAlong Software, Inc., PO Box 359, Brownsville California 95919 USA
thinkalong@applelink.apple.com (916) 675-2478

Abstract

This paper describes an approach to evolution upon an artificial insect existing in a simulated two-dimensional world, comprised of a "spreading-activation" neural network. Study of the difficulty of hand-building neural networks has resulted in a genetic expression, based on the notion of the Von Neumann computer architecture and the biological principles of DNA, which incorporates both operations and data in a simulated DNA strand. Work in progress indicates that by embodying behavior parameters and actual neural connections in a genotypic "language," and by expressing that phenotypically as an animat, the computer simulation appears able to evolve a better species of animat through mutations upon the genotype.

Introduction

A domain of "artificial life" explores limited computer simulations of animal behavior in an artificial environment. This class of simulated animal is coming to be known as the *animat*, as first coined by Stewart Wilson in [Wilson, 1985]. Research on the notion of an artificial insect has proceeded in many diverse headings (see for instance [Travers, 1988], [Maes, 1990], [Park, 1988], [Wilson, 1987]). The implementation described here is a direct extension of Jack Park's animat as described in [Park, 1988].

The existing animat system upon which this program has extended utilizes a spreading-activation [Collins and Loftus, 1975, Anderson, 1983] neural network which was originally conceived as an exercise to ascertain the capabilities of a neural net in *controlling* some sort of process. This neural implementation has been applied not only to the animat described here, but to process control in the manufacturing domain [Park, et al., in prep.], part of a scientific discovery system [Wood and Park, in press], and to the study of the Piagetian development of an infant brain [Wood, manuscript].

The "wetware" implementation controlling this animat consists of a few dozen "neuron" nodes, each connected to several other nodes with varying levels of positive (excitatory) or negative (inhibitory) strength. Sensors (input of sight, pain, hunger, taste, and satiation) stimulate neurons which in turn dissipate activation to other neurons; muscle neurons that have reached a certain threshold of activation will cause appropriate routines to execute (to cause the bug to move, turn, eat, or move randomly). Additionally, the system applies a constant decay, to keep the overall activation steady, and random noise injected into the activation levels of all neurons. The system does not employ learning methods during the "lifetime" of the animat; this project studies only the non-plastic neurons found in the lower animals.

The animat's environment, implemented on a Macintosh II, is a bound arena (figure 1) with morsels




Figure 1

Four Important Issues in Cognitive Mapping

W.K. Yeap and C.C. Handley
AI Laboratory,
Department of Computer Science,
University of Otago, New Zealand

coscwky@otago.ac.nz
chandley@otago.ac.nz

Abstract

We introduce the notions of Raw and Full Cognitive Maps and Absolute Space Representations (ASRs) and discuss (i) the value of computing a representation of the local environment (ASR), (ii) the need for a global representation of one's immediate surroundings, (iii) categorisation and the formation of concepts and (iv) the problem of planning a route.

0. Introduction

The problem of computing a cognitive map is fundamental to any autonomous mobile system, be it a rat, a human or a robot. When Tolman (1948) first suggested the idea of a cognitive map, he was probably referring to a "map" of the spatial layout of the environment (mazes in his case) but later it became clear, especially after Lynch's (1960) work, that the notion of a cognitive map is a complex one. In the early 70's, there was an outcry from geographers, urban planners and designers that a cognitive map is not a map (see [Downs and Stea, 1973]). A cognitive map is tied to our spatial behaviour and it is therefore affected by a wide variety of factors ranging from our mode of travel and past experiences to our preferences and attitude. One problem with this view is that it leads to a confusing use of the term and one is often left with the impression that a cognitive map holds one's entire knowledge. However, if the system is to adapt and survive in a (hostile) world with other agents in it, such factors must be considered.

Cognitive mapping is therefore a complex process which involves both one's perception and conception of the outside world. Studies which emphasised only one level, either the perceptual (e.g. work on autonomous mobile robots) or the conceptual (e.g. early models), were at best incomplete and very often asked many questions inappropriate at that level. For example, robotics researchers were concerned with how to partition the environment in terms of spaces large enough for the robots to plan a collision free path, but planning a collision free path is a local problem and an attempt to solve it at the path planning level is inappropriate (for more detail, see [Yeap et al., 1990]).

Our past work has been the development of a computational theory of cognitive maps to explain what needs to be computed and why [Yeap, 1988, 1990; Yeap and Robertson 1990; Yeap et al., 1990]. We stress the importance of studying the process as a whole, from perception to cognition and generally in that order. Our investigation of the process begins with Marr's (1982) computational theory of vision and, as in Marr's work, the notion of a *representation* is central to our study. The next section presents a brief overview of the theory, the main idea being that a cognitive mapping process should first compute a *raw* map and then a *full* map. Using this theory we discuss four important issues in cognitive mapping, two at each level: (i) the significance of computing a representation of the local environment (an Absolute Space Representation or ASR) and (ii) the need for a global representation of one's immediate surroundings; (iii) categorisation and concept formation and (iv) the problem of planning a route. These issues arise from the insights gained from the implementation of the theory and from further consideration of the nature of the cognitive mapping process.

1. A Computational Theory of Cognitive Maps

Although there are many factors which influence our conceptual view of the world, our representation of the world must begin from what we perceive. This observation suggests that the first step in a cognitive mapping process is to compute a representation of the physical environment. Since the different conceptual views of the world are but different ways of looking at what is already computed, this representation of the physical environment should be fairly independent of the conceptual representations that are developed later. Our theory therefore suggests that a cognitive mapping process should be studied as a process consisting of two loosely-coupled modules. An early cognitive mapping process computes a representation of the physical world as perceived by our senses. We call this representation a *Raw Cognitive Map*; indicating that the computed representation is not interpreted. A later cognitive mapping process computes a representation of the conceptual world. We call this representation a *Full Cognitive Map*; the word "full" indicates the full richness of the map as a cognitive representation.

**List
of
participants**

ARBIB Michael
Center for Neural Engineering
U. of Southern California
CA 90089-2520 Los Angeles
USA
Fax: 0 Email: arbib@poliux.usc.edu

ATLAN Laurent
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: atlan@frulm63.bitnet

BARRE Robert
ISRP/ENSPR/IAR
14 avenue du Griffon
84700 Sorgues
France
Fax: 0 Email: 0

BEC Louis
ISRP/ENSPR/IAR
14 avenue du Griffon
84700 Sorgues
France
Fax: 0 Email: 0

BERSINI Hughes
ULB & CREA
Laboratoire Iridia, 50 av F. Roosevelt- CP 194/6
1050 Brussels
Belgium
Fax: 0 Email: 0

BESSIERE Pierre
IMAG/LGI
BP 53 X
38041 Grenoble cedex
France
Fax: 0 Email: bessiere@imag.fr

BONELLI Pierre
Equipe Inférence et Apprentissage, LRI
Bat. 490, Université de Paris Sud
91405 Orsay cedex
France
Fax: 0 Email: 0

BOVET Pierre
CNRS Laboratoire de Neurosciences fonctionnelles
BP 71, 31 chemin Joseph Aiguier
13402 Marseille cedex 9
France
Fax: 0 Email: bovet@frmop11

BRIN David
2130 S. Vermont avenue
CA 90007 Los Angeles
USA
Fax: 0 Email: 0

ARNAUD Daniel
Palais de la Découverte
Av. Franklin D. Roosevelt
75000 Paris
France
Fax: 0 Email: 0

BARNARD Stephen
AI Center, SRI International
333 Ravenswood Ave.
CA 94025 Menlo Park
USA
Fax: 0 Email: barnard@iu.sri.com

BAUCHAU Vincent
Unité de Biologie animale, U. de Louvain
5 Croix du Sud
1348 Louvain la neuve
Belgium
Fax: 10 47 35 15 Email: vincent@buclln11

BEER Randall
Dept. of Computer Engineering and Science
Case Western Reserve U.
OH 44106 Cleveland
USA
Fax: 0 Email: 0

BERTIN René
Neuro-ethology Group, Dept. of Comparative Physiology
Utrecht U., Trans III, Padualaan 8
3584 CH Utrecht
The Netherlands
Fax: 0 Email: bertin@neuretb.biol.ruu.nl

BOERLYST Maarten
Bioinformatics Group
Padualaan 8
3584 CH Utrecht
The Netherlands
Fax: 313 051 36 55 Email: awbpbia@cc.ruu.nl

BOOKER Lashon B.
Navy Center for Applied Research in Artificial Intelligence
Naval Research Laboratory
DC 20375-5000 Washington
USA
Fax: 0 Email: booker@aic.nrl.navy.mil

BRELET Anne
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: brelet@frulm63.bitnet

BROOKS Rodney
MIT Artificial Intelligence Lab.
545 Technology square
MA 02139 Cambridge
USA
Fax: 0 Email: brooks@ai.mit.edu

CABANAC Michel
Dept. of Physiology, Faculty of Medicine
Laval University
G1K 7P4 Quebec
Canada
Fax: 418 656 34 42 Email: cabanac@lavalvm1

CECCONI Federico
Institute of Psychology, CNR
15 viale Marx
R 00137 Rome
Italy
Fax: 6 8244737 Email: cecconi@irmkant.bitnet

CHATILA Raja
LAAS/CNRS
7 avenue Colonel Roche
31077 Toulouse cedex
France
Fax: 61 55 35 77 Email: chatila@laas.fr

CLIFF David
University of Sussex
School of Cognitive and Computing Sciences
BN1 9QN Brighton
UK
Fax: 273 678 335 Email: davecc@cogs.sussex.ac.uk

COLLINS Robert
Dept of Computer Science
U. of California
CA 90024 Los Angeles
USA
Fax: 213 825 22 73 Email: rjc@cs.ucla.edu

CORBARA Bruno
Laboratoire d'Ethologie
Université Paris XIII
93430 Villetaneuse
France
Fax: 494 33 33 Email: vety1@frors31

DANDRIEUX Jean-Pierre
1 rue Championnet
75020 Paris
France
Fax: 0 Email: 0

DE GARIS Hugo
CADEPS AI & AL Research Unit, U. libre de Bruxelles
av. F. Roosevelt, 50 CP 194/7
B 1050 Brussels
Belgium
Fax: 32 2 642 2715 Email: cadeps@bbrnsf11.bitnet

DE LUIGI Fabio
CED
Town Municipality of Ferrara
44100 Ferrara
Italy
Fax: 5 32419389 Email: 0

CAHIER Jean-Pierre
Le Monde Informatique
185 avenue Ch. de Gaulle
92521 Neuilly sur Seine
France
Fax: 49047800 Email: 0

CERF Raphaël
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: cerf@frulm63.bitnet

CHUNG TO SANG Bernard
SUN Microsystem
13 avenue Morane Saulnier
78140 Vélizy
France
Fax: 0 Email: 0

COLGAN Patrick
Dept. of Biology
Queen's University
K7L 3N6 Kingston
Canada
Fax: 0 Email: colganp@qucdn

CONKIE Alistair
Dept of Artificial Intelligence, U. of Edinburgh
5 Forrest Hill
EH1 2QL Edinburgh
UK
Fax: 312 259 370 Email: adc@edai.ed.ac.uk

CRUSE Holk
Fac. of Biology, U. of Bielefeld
Postfach 8640
D-4800 Bielefeld 1
Germany
Fax: 0 Email: ubiof140@dbiuni11

DAVIDOR Yuval
Dept of applied mathematics & computer science
The Weismann Institute of Science
76100 Rehovot
Israel
Fax: 0 Email: yuval@wisdom.weismann.ac.il

DE JONG Kennet
Computer Science Dept.
George Mason University
VA 22030 Fairfax
USA
Fax: 0 Email: kdejong@aic.gmu.edu

DE PRACONTAL Michel
Le Nouvel Observateur
14 rue Dussoubs
75081 Paris cedex 02
France
Fax: 42361963 Email: 0

DE VRIES Han
Ethology & Socio-ecology RUU
Post box 80.086
3508 TB Utrecht
The Netherlands
Fax: 313 052 11 05 Email: wbqkojv@hutruu0

DEFFUANT Guillaume
LEMAGREF
64 av du Gal Leclerc
75014 Paris
France
Fax: 0 Email: 0

DERIX Ruud
Ethology & Socio-ecology RUU
Post box 80.086
3508 TB Utrecht
The Netherlands
Fax: 313 052 11 05 Email: wbqkoyv@hutruu0

DUCHATEAU Guillemette
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: duchat@frulm63.bitnet

DWORETZKY Tom
Omni magazine
232 W, 16th St.
NY 10011 New-York
USA
Fax: 0 Email: 0

FORSTER Peter
Dept of Artificial Intelligence, Edinburgh University
5 Forrest Hill
EH1 2QL Edinburgh
UK
Fax: 312 259 370 Email: pmf@edai.ed.ac.uk

GABORA Liane
Indiana University
CRCC, 510 N. Fess St
IN 47405 Bloomington
USA
Fax: 0 Email: liane@cogsci.indiana.edu

GLOTIN Hervé
78 avenue du Bois Guinier
94100 Saint Maur
France
Fax: 0 Email: 0

GRAFF Christian
Dept. de Psychologie, U. des Sciences Sociales
BP 47X
38040 Grenoble
France
Fax: 33 76825654 Email: 0

DEDIEU Eric
LIFIA
46 av Felix Viallet
38000 Grenoble
France
Fax: 0 Email: dedieu@lifia.imag.fr

DENEUBOURG Jean-Louis
Unit of theoretical behavioural ecology CP 231, U. Libre de Bruxelles
Campus Plaine, Bld du Triomphe
1050 Brussels
Belgium
Fax: 322 649 87 74 Email: 0

DORMOY Jean-Luc
DER-EDF IMA-TIEM
1 avenue du Général de Gaulle
92141 Clamart cedex
France
Fax: 0 Email: 0

DUMEUR Renaud
Dept. Informatique, U. de Paris 8
2 rue de la Liberté
93256 St Denis cedex 02
France
Fax: 0 Email: red@litp.ibp.fr/renaud@ccv.fr

FERRIERE Regis
Laboratoire d'Ecologie, ENS
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: 0

FRANCESCHINI Nicolas
CNRS, Laboratoire de Neurobiologie
31 chemin J. Auguier
13009 Marseille
France
Fax: 91 22 875 Email: frchini@frmop11

GIAVITTO Jean-Louis
LRI
Bat 490 Université Paris Sud- Orsay
91405 Orsay cedex
France
Fax: 0 Email: giavitto@lri.lri.fr

GOSS Simon
Unit of theoretical behavioural ecology CP 231, U. Libre de Bruxelles
Campus Plaine, Bld du Triomphe
1050 Brussels
Belgium
Fax: 322 649 87 74 Email: 0

GRANJEON Eric
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: granjeon@frulm63.bitnet

GREUSSAY Patrick
Université Paris 8 Dept. Informatique URA 248/LITP
2 rue de la Liberté
93526 Saint-Denis
France
Fax: 0 Email: 0

GUILLAUME Marc
Ministère de la Recherche et de la Technologie
1 rue Descartes
75005 Paris
France
Fax: 0 Email: 0

HALPERIN Janet
Dept. of Zoology, U. of Toronto
24 Harbord Street
M5S 1A1 Toronto
Canada
Fax: 416 978 85 32 Email: janh@zoo.toronto.edu

HARRIS Mitch
Dept of Artificial Intelligence, U. of Edinburgh
5 Forrest Hill
EH1 2QL Edinburgh
UK
Fax: 312 259 370 Email: mitch@edai.edinburgh.ac.uk

HESPER Ben
Bioinformatics Group
Padualaan 8
3584 CH Utrecht
The Netherlands
Fax: 313 051 36 55 Email: hesper@hlerul5.bitnet

HOLLAND Owen
Technology Applications Group
Bolam's Mill
NE66 1LN Alnwick
UK
Fax: 665 00510624 Email: 0

HUYNEN Martijn
Bioinformatics Group
Padualaan 8
3584 CH Utrecht
The Netherlands
Fax: 313 051 36 55 Email: awbpbia@cc.ruu.nl

JAMON Marc
CNRS Laboratoire de Neurosciences fonctionnelles
BP 71, 31 chemin Joseph Aiguier
13402 Marseille cedex 9
France
Fax: 0 Email: bovet@frmop11

KELLY Kevin
Whole Earth Review
149 Amapola Ave
CA 94044 Pacifica
USA
Fax: 0 Email: kk@well.sf.ca.us

GRUMBACH Alaina
Telecom
46 rue Barrault
75634 Paris cedex 13
France
Fax: 0 Email: grumbach@inf.enst.fr

GUILLOT Agnès
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: guillot@frulm63.bitnet

HARDY Noah
Hopital St Louis
1 avenue C. Vellefaux
75475 Paris cedex 10
France
Fax: 0 Email: 0

HARVEY Inman
School of cognitive and computing sciences
Sussex University
BN1 9QH Brighton
UK
Fax: 273678188 Email: inmanh@cogs.susx.ac.uk

HOGEWEG Pauline
Bioinformatics Group
Padualaan 8
3584 CH Utrecht
The Netherlands
Fax: 313 051 36 55 Email: hesper@rulcri.leiden.univ.nl

HOUSTON Allisdair
NERC Unit of Behavioral Ecology, Dept of Zoology
U. of Oxford
OX1 3BJ Oxford
UK
Fax: 0 Email: 0

IDAN Yishak
ENST Telecom
46 rue Barrault
75634 Paris cedex 13
France
Fax: 45 81 37 94 Email: idan@liszt.enst.fr

JOHNSON Michael
Computer graphic and animation group, MIT Media Lab
20 Ames Street
MA 02139 Cambridge
USA
Fax: 0 Email: wave@media-lab.media.mit.edu

KHOUAS Sallha
LRI
Bat. 490, Université de Paris Sud
91405 Orsay cedex
France
Fax: 0 Email: khouas@lri.lri.fr

KISS George
Human Cognition Research Lab.
The Open University
MK7 6AA Milton Keynes
UK
Fax: 0 Email: gr_kiss@vax.acs.open.ac.uk

KURTZ Cynthia
Dept. of Ecology & Evolution
State University of New York @ Stony Brook
NY11794 Stony Brook
USA
Fax: 0 Email: ckurtz@sbbiovm.bitnet

LESTEL Dominique
Ecole des Hautes Etudes, GERS
10 rue Monsieur le Prince
75006 Paris
France
Fax: 0 Email: 0

LIN Long-Ji
School of Computer Science, Carnegie Mellon University
5000 Forbes Avenue
PA 15213 Pittsburgh
USA
Fax: 412 681 57 39 Email: ljl@a.gp.cs.cmu.edu

MALER Oded
IRISA
Campus de Beaulieu
35042 Rennes
France
Fax: 0 Email: maler@allegro.irisa.fr

MANTEUFFEL Gerhard
Brain Research Institute FB2
U. of Bremen, POB 33 04 40
D-2800 Bremen
FRG
Fax: 421 219 663 Email: 0

MATARIC Maja
MIT Artificial Intelligence Lab
545 Technology Square
MA 02139 Cambridge
USA
Fax: 617 253 00 39 Email: maja@ai.mit.edu

MAZER Emmanuel
LIFIA
46 av Felix Viallet
38000 Grenoble
France
Fax: 0 Email: manu@lifia.imag.fr

McGONIGLE Brendan
Lab. for Cognitive Neuroscience
Edinburgh University
EH 12QL Edinburgh
UK
Fax: 0 Email: M.A.Chalmers@edinburgh.ac.uk

KOZA John
Computer Science Dept., Stanford University
Box 1439
CA 94305 Stanford
USA
Fax: 415 941 00 36 Email: koza@sunburn.stanford.edu

LECOMTE Jane
Laboratoire d'Ecologie, ENS
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: 0

LEVY Steven
Viking-Penguin
24 Fifth Ave.
NY 10011 New-York
USA
Fax: 0 Email: steven@well.sf.ca.us

MAES Pattie
AI-LAB Vrije Universiteit Brussel
Dienst Arti, Pleinlaan 2
B1050 Brussels
Belgium
Fax: 322 641 28 70 Email: pattie@arti.vub.ac.be

MANIEZZO Vittorio
Dept. di elettronica - Politecnico di Milano
Piazza Leonardo da Vinci 32
20133 Milano
Italy
Fax: Email: maniezzo%ipmell1.infn.it@iboinfn.bitnet

MARTIN Serge
Laboratoire d'Ecologie, ENS
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 0 Email: 0

MATHIEU Philippe
Swedish Institute of Computer Science
PO Box 1263
S - 164 28 Kista
Sweden
Fax: 8 751 72 30 Email: philippe@sics.se

McFARLAND David
Dept. of Animal Behaviour, Bailliol College
Oxford University
OX1 3BJ Oxford
UK
Fax: 0 Email: 0

MEYER Jean-Arcady
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: meyer@frulm63.bitnet

MEYER Christophe

3 rue des Saulx
28130 St Martin de Nigelles
France
Fax: 0 Email: 0

MILLER Geoffrey

Psychology Department
Build. 420, Stanford University
CA 94305 Stanford
USA
Fax: 0 Email: geoffrey@psych.stanford.edu

MORENO BERGARECHE Alvaro

Dept of Logic & Philosophy of Science
University of the Basque Country Apartado 1249
20080 San Sebastian
Spain
Fax: 43 471 699 Email: alvaro@fil.ehu.es

NADAL Jean-Pierre

Laboratoire de Physique Statistique ENS
24, rue Lhomond
75230 Paris cedex 05
France
Fax: 0 Email: nadal@frulm62.bitnet

NEHMZOW Ulrich

Dept of Artificial Intelligence, Edinburgh University
5 Forrest Hill
EH1 2QL Edinburgh
UK
Fax: 312 259 370 Email: ulrich@edai.ed.ac.uk

NOLDUS Lucas

Noldus Information Technology b.v.
University Business & Technology Center, Vadaring 51
6702 EA Wageningen
The Netherlands
Fax: 31 8370 24496 Email: noldus@rcl.wau.nl

PAILLETTE Madeleine

Laboratoire de Biologie et Génétique évolutives
CNRS
91198 Gif sur Yvette
France
Fax: 69 07 04 21 Email:

PARENTHOEN Marc

ENS Cachan F317
61 avenue Président Wilson
94235 Cachan
France
Fax: 0 Email: 0

PARODI Alexandre

Infodyne SARL
10 rue de la Paix
75002 Paris
France
Fax: 40209898 Email: (uucp)parodi@lri.lri.fr

MILLER Thomas

Dept of Electrical and Computer Engr.
Kingsbury Hall, U. of New Hampshire
NH 03827 Durham
USA
Fax: 603 862 20 30 Email: t_miller@unhh.bitnet

MOORE Patrick

Naval Ocean Systems Center
PO Box 997
HI 96734 Kailua
USA
Fax: 808 257 5236 Email: pmoore@nosc.mil

NACHTIGALL Paul

Naval Ocean Systems Center
PO Box 997
HI 96734 Kailua
USA
Fax: 808 257 52 36 Email: nachtig@nosc.mil

NAILLON Martine

Dassault
12 Av de Verdun
92330 Sceaux
France
Fax: 0 Email: 0

NITSCHKE Volker

Zoologisches Institut der Universität München
Luisen Strasse 14
D-800 München
FRG
Fax: 895 902 450 Email: 0

NOLFI Stefano

Institute of Psychology, CNR
15 viale Marx
00137 Rome
Italy
Fax: 6 8244737 Email:

PAREDIS Jan

RIKS
Postbus 463
NL-6200 Al Maastricht
The Netherlands
Fax: 43 25 31 55 Email: jan@riksnl.uucp

PARISI Domenico

Institute of Psychology, CNR
15 viale Marx
00137 Rome
Italy
Fax: 6 8244737 Email: domenico@irmkant.bitnet

PENNETIER Jean-Louis

Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: pennetier@frulm63.bitnet

PERRIER Jean-Jacques
La Recherche
57 rue de Seine
75006 Paris
France
Fax: 46347508 Email: 0

PEYRAUD Jean-Jacques
FR3
28 cours Albert 1er
75008 Paris
France
Fax: 0 Email: 0

PIERCE David
Dept. of Computer Sciences
U. of Texas at Austin, Taylor 2.124
TX 78712 Austin
USA
Fax: 0 Email: dmpierce@cs.utexas.edu

RAPPAPORT Alain
Neuron Data / Robotics Institute, CMU
444 High St.
CA 94301 Palo Alto
USA
Fax: 321 74 34 Email: atr@ml.ri.cmu.edu

REYNOLDS Craig
Symbolics Inc.
1401 Westwood Boulevard
CA 90024 Los Angeles
USA
Fax: 213 477 13 46 Email: cwr@symbolics.com

RING Mark
U. of Texas at Austin
Dept of Computer Sciences, Taylor 2.124
TX 78712 Austin
USA
Fax: 0 Email: ring@cs.utexas.edu

ROBERTS Gary
Dept. of Artificial Intelligence, U. of Edinburgh
80 South Bridge
EH1 1HN Edinburgh
UK
Fax: 312 26 27 30 Email: gary@aipna.ed.ac.uk

SALAME Naoum
INRP
rue d'Ulm
75005 Paris
France
Fax: 0 Email: 0

SCHALIT Emmanuel
Laboratoire de Marcoussis CGE
route de Nozay
91460 Marcoussis
France
Fax: 64 49 06 95 Email: schalit@crcge1.cge.fr

PETROWSKI Alain
Institut National des Télécommunications/Paris 6
DIT, INT, 9 rue Charles Fournier
91011 Evry cedex
France
Fax: 60 76 47 80 Email: petrowski@frint51.bitnet

PFEIFER Rolf
AI-LAB Vrije Universiteit Brussel
Dienst Arti, Pleinlaan 2
B1050 Brussels
Belgium
Fax: 322 641 28 70 Email: 0

PUTTERS Felix
Ethology Research Group, U. of Leiden
Kaiserstraat 63 PO Box 9516
2300 RA Leiden
The Netherlands
Fax: 71 27 49 00 Email: putters@hlerul57

RESNICK Mitchel
MIT Media Lab
20 Ames Street
MA 02139 Cambridge
USA
Fax: 617 253 62 15 Email: mres@altdorf.ai.mit.edu

RICCI Nicola
Dept Scienze dell'Ambiente e del Territorio
via A Volta 4
56100 Pisa
Italy
Fax: 0 Email: 0

RIOLO Rick
U. of Michigan
Logic of Computers Group
MI 48109 An Arbor
USA
Fax: 0 Email: rick_riolo@um.cc.umich.edu

ROITBLAT Herbert L.
Dept of Psychology, U. of Hawaii at Manoa
2430 Campus Road
HI 96822 Honolulu
USA
Fax: 808 956 47 00 Email: herbert@uhccux.bitnet

SANSONNET Jean-Paul
Laboratoire de Recherches en Informatique
Bat 490 Université Paris Sud- Orsay
91405 Orsay cedex
France
Fax: 0 Email: 0

SCHMIDHUBER Jürgen
Technische Universität München, Institut für Informatik
Arcisstr, 21
8000 München 2
FRG
Fax: 921 058 207 Email: schmidhu@tumult.informatik.tu-muench

SCHMIDT Sabine
Zoologisches Institut der Universität München
Luise Strasse 14
D-800 München
FRG
Fax: 895 902 450 Email: 0

SCHULL Jonathan
Haverford College

PA 19041 Haverford
USA
Fax: 0 Email: 0

SMITHERS Tim
Dept of Artificial Intelligence, Edinburgh University
5 Forrest Hill
EH1 2QL Edinburgh
UK
Fax: 312 259 370 Email: tim@edai.ed.ac.uk

SPIESSENS Piet
AI-LAB Vrije Universiteit Brussel
Dienst Arti, Pleinlaan 2
B1050 Brussels
Belgium
Fax: 322 641 28 70 Email: piet@arti.vub.ac.be

STANTON Harry
The MIT Press
55 Hayward Street
Ma 02142 Cambridge
USA
Fax: 6172535627 Email: stanton@uvmvm.bitnet

STRETTA Jean Michel
ORSTOM
BP 5045
34032 Montpellier cedex 1
France
Fax: stretta@orstom.orstom.fr Email: 0

TANGNEY John F.
Life Sciences Directorate
AFOSR/NL Building 410
DC 20332-6448 Washington
USA
Fax: 0 Email: jtangney@nas.bitnet

TEITELBAUM Phillip
Psychology Dept.
U. of Florida
FL 32611 Gainesville
USA
Fax: 904 392 79 85 Email: teitelb@webb.psych.ufl.edu

TOATES Frederick
Biology Dept., The Open University
Walton Hall
MK7 6AA Milton Keynes
UK
Fax: 908 653 744 Email: f_toates@uk.ac.open.acs.vax

SCHNEPF Uwe
Expert Systems Research Group, GMD
PO Box 1240
5205 Sankt Augustin 1
FRG
Fax: 2241 142 618 Email: usc@gmdzi.uucp

SHANKS David
MRC Applied Psychology Unit
15 Chaucer Rd
CB2 ZEF Cambridge
UK
Fax: 22 33 59 62 Email: davids@uk.ac.cam.mrc-apu

SNAITH Martin
Technology Applications Group
Bolam's Mill
NE66 1LN Alnwick
UK
Fax: 665 00510624 Email: 0

SPRATT Michael
Applying Evolving Technologies
11 rue Daru
34000 Montpellier
France
Fax: 0 Email: 0

STEELS Luc
AI LAB Vrije Universiteit Brussel
Pleinlaan 2
B-1050 Brussels
Belgium
Fax: 0 Email: steels@arti.vub.ac.be

SUTTON Richard
GTE Lab. Incorporated
MS-44
MA 02254 Waltham
USA
Fax: 0 Email: sutton@gte.com

TARROUX Philippe
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: tarroux@frulm63.bitnet

THERAULAZ Guy
CNRS- UPR 38
31 chemin Joseph Aiguier
13402 Marseille cedex 09
France
Fax: 0 Email: 0

TODD Peter
Psychology Dept., Stanford University
Jordan Hall Build. 420
CA 94305 Stanford
USA
Fax: 415 725 5699 Email: todd@psych.stanford.edu

TORREELE Jan
AI-LAB Vrije Universiteit Brussel
Dienst Arti, Pleinlaan 2
B1050 Brussels
Belgium
Fax: 322 641 28 70 Email: jant@arti.vub.ac.be

TRAVERS Mike
MIT Media Lab E15 494
20 Ames Street
MA 02139 Cambridge
USA
Fax: 0 Email: mt@media-lab.media.mit.edu

TUNLEY Hilary
University of Sussex
School of Cognitive and Computing Sciences
BN1 9QN Brighton
UK
Fax: 273 678 335 Email: hilary@cogs.sussex.ac.uk

VAN AEKEN Francis
AI-LAB Vrije Universiteit Brussel
Dienst Arti, Pleinlaan 2
B1050 Brussels
Belgium
Fax: 322 641 28 70 Email: francis@arti.vub.ac.be

VAN DER LAAN Jan
Bioinformatics Group
Padualaan 8
3584 CH Utrecht
The Netherlands
Fax: 313 051 36 55 Email: awbpbia@cc.ruu.nl

VEISSIER Isabelle
U.R. Adaptation comportementale
INRA
63122 Theix
France
Fax: 0 Email: 0

VINCENS Pierre
Groupe de BioInformatique
46 rue d'Ulm
75230 Paris cedex 05
France
Fax: 43 29 81 72 Email: vincens@frulm63.bitnet

WALDRON Ronan
Dept. of Computer Science
Trinity College
0 Dublin 2
Ireland
Fax: 353 1 772 204 Email: rwaldron@cs.tcd.ie

WEISBUCH Gérard
Laboratoire de Physique Statistique ENS
24, rue Lhomond
75230 Paris cedex 05
France
Fax: 0 Email: weisbuch@frulm62.bitnet

TOULOUSE Gérard
Laboratoire de Physique Statistique ENS
24 rue Lhomond
75231 Paris cedex 05
France
Fax: 45 87 34 89 Email: toulouse@frulm11

TROXELL Wade
Mechanical Engineering
Colorado State University
Colorado 90523 Fort Collins
USA
Fax: 303 491 1055 Email: wade@longs.lance.colostate.edu

TYRRELL Toby
Dept. of Psychology, Sheffield U.
Western Bank
S102TN Sheffield
UK
Fax: 0 Email: toby@aivru.uucp

Van AMERONGEN Anton
Ryles universiteit Utrecht
PO 81086
3508 TB Utrecht
The Netherlands
Fax: 313 052 11 05 Email: wbqkojv@hutruu0

VANNIMENUS Jean
Laboratoire de Physique Statistique ENS
24, rue Lhomond
75230 Paris
France
Fax: 0 Email: weisbuch@frulm62.bitnet

VELTHUIS Berend-Jan
Ethology Research Group, U. of Leiden
Kaiserstraat 63 PO Box 9516
2300 RA Leiden
The Netherlands
Fax: 71 27 49 00 Email: sbw1bv@hleru157

VONK Maryke
Ethology Research Group, U. of Leiden
Kaiserstraat 63 PO Box 9516
2300 RA Leiden
The Netherlands
Fax: 71 27 49 00 Email: sbw1mv@hleru157

WALTZ David
Thinking Machines Corporation
245 First St.
MA 02142 Cambridge
USA
Fax: 0 Email: 0

WESTERDALE Thomas
Birkbeck College, U. of London
Malet Street
WC1E7HX London
UK
Fax: 716 36 49 71 Email: tom@cs.bbk.ac.uk

WILLIAMS James

Dept. of Biology and Preclinical Medicine
U. of St. Andrews, Bute building, Westburn Lane
Fife KY16 9TS St. Andrews
UK

Fax: 33 47 83 99 Email: zorjw@uk.ac.st-andrews.sava

WOOD Dan

ThinkAlong Software, Inc.
PO Box 359
95919 Brownsville
USA

Fax: 0 Email: thinkalong@applelink.apple.com

YEAP Wai-Kiang

Dept. of Computer Science
U. of Otago
Dunedin Box 56
New-Zealand

Fax: 64 24 798 577 Email: coscwky@otago.ac.nz

ZELMANS Bernard

SUN Microsystem
13 avenue Morane Saulnier
78140 Vélizy
France
Fax: 0 Email: 0

WILSON Stewart W.

The Rowland Institute for Science
100 Cambridge Parkway
2142 Cambridge
USA

Fax: 617 497 46 27 Email: wilson@think.com

YAEGER Larry

Apple Computer
8733 Lookout Mountain Ave
90046 Los Angeles
USA

Fax: 0 Email: larryy@apple.com

YEDIDIA Jonathan

Laboratoire de Physique Théorique ENS
24, rue Lhomond
75230 Paris cedex 05
France

Fax: 0 Email: 0