AISB Quarterly

No.117

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

The role of emotions in modular intelligent control

Often, artificial intelligence researchers are tempted to look for a single, simple, homogeneous solution to creating intelligent systems. But animal intelligence isn't like that at all. Animal brains have large numbers of different representations for different sorts of information. Over the last 10 years, we have been working on integrating the roles of hierarchical action-selection mechanisms with behaviour modules containing more distributed, dynamic, learning and control processes. Now we have begun working on understanding how different representations and time courses for emotions come together to influence action selection, particularly in social contexts.

We are working on a representational framework for emotions that will integrate with action selection, providing the sort of persistent internal contextual state that animals have found so useful for regulating their behaviour. We are using two mechanisms to test the plausibility of our representations: a facial animation tool—which will allow us to tap human expertise on the believability of emotions—and multi-agent simulations of primate social behaviour, to test the importance of introducing realistic emotional onsets and decays in building stable group social dynamics. This article will concentrate on the former.

An initial draft of an overall emotional and personality architecture can be seen in Figure 1. The lowest level of the emotional architecture consists of *primary emotions*[1] which are generated reactively from the experiences of the agent. *Secondary emotions* map more closely to the common concepts of emotion, e.g. joy or anger, and often have cognitive reference. Moods affect perception as well as expressed behaviour, and typically have longer duration than secondary emotions. The dynamic emotion representation (DER) represents secondary emotions and mood state. It consists of a number of modules, each containing dynamic representations of their intensity. Each module also has a stimulus and decay function—the intensity increases sharply in response to primary emotional stimuli, then decays slowly. The number of secondary emotions can be altered depending on the emotional theories a researcher chooses to represent, but our default values are set for Ekman.[2] More details about the DER and its role in facial expression are available.[3]

We are building a facial animation tool (see Figure 2) that is based on the two-channel concept of facial expressions. This concept assumes two things. First, facial displays are used deliberately by humans to support speech with both redundant and novel visual information. This information is referred to as communicative acts[5] or visible acts of meaning.[6] Second, facial displays also reveal internal emotional state, which could be called emotional facial expressions. This suggests the existence of two different channels creating facial displays: the communicative channel, which is composed of speech and its tightly-synchronised communicative acts, and the emotional channel.

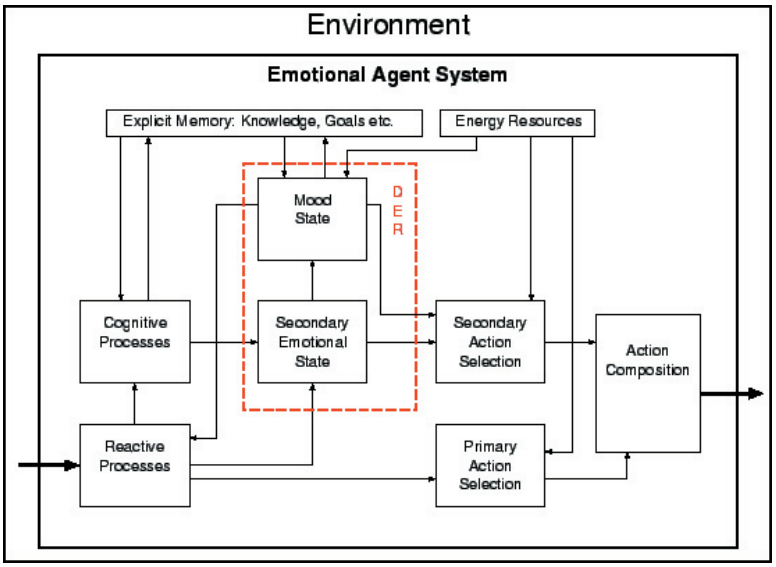

Figure 1. Draft architecture for a complete emotional agent. Arrows indicate the flow of information and boxes the types of processes (which will be further modularised). The dashed box is the location of the dynamic emotion representation (DER).

The emotional channel derives data from the DER. Each secondary emotion has a facial expression associated with it, while the mood impacts other signals such as tension and communication. The

Bryson, *continued on p.6*

In this issue:

Models for mobile context awareness

The **AISB Quarterly** is published by the Society for the Study of Artificial Intelligence and Simulation of Behaviour.

ISSN 0268-4179

Copyright © AISB 2004

Articles may be reproduced as long as the copyright notice is included. The item should be attributed to the AISB Quarterly and contact information should be listed.

Quarterly articles do not necessarily reflect the official AISB position on issues.

Editor:
Sunny Bains
Imperial College London
aisbq@aisb.org.uk
http://www.aisb.org.uk and click on AISBQ.

Editorial Assistant and Layout:
Stuart Barr

Advertising and Administration:
Therie Hendrey-Seabrook
aisbq@aisb.org.uk
School of Science and Technology
University of Sussex
Falmer, Brighton
UK, BN1 9QH
Tel: (01273) 678448
Fax: (01273) 671320

AISB PATRON:
John Barnden
University of Birmingham

Mobile devices are everywhere: PDAs, mobile phones, laptops, and notebooks are in everyday use by everyday people. They are slowly extending their roles as we learn to make a feature of their mobility and connectivity. In order to achieve this effectively we need to get hold of information about the user—their tasks and interests, and the environment they are in—and then put it to good use to allow us to provide timely support. This context is important because it allows us to make use of the environment in a way that supports the user. For example, we can envisage the scenario of a mobile phone that is aware of its user's location and activity and will not disturb an important meeting. But the same phone, being aware of its user's call list and calendar, will permit a call from a pregnant wife. In this way the users themselves form part of the environment they occupy, and we can use information about the user to further enhance our contextual model.

Context is useful in three respects. Firstly, it relates the services to time and location and to the user's needs and interests, thus ensuring that they are useful, learnable, and enjoyable. Secondly, it provides for more effective use of resources—which is especially important in the mobile situation—with many different limitations: these include device processing power, display ability, media capabilities, network bandwidth, connectivity options, and intermittent connections. Other aspects of the situation that are competing for attention can also be taken into account this way. Thirdly, by providing more appropriate information delivered most effectively, context allows the user to focus much less on the technology and more on the actual situation they are in.

We consider context not as a static phenomenon but as a dynamic process: constructed through the user's interactions with the learning materials and the surrounding world over time. All of these domains provide information in themselves, and can interact with the others in a variety of ways. This builds up a rich model of the current world and so allows the system to be more specific in what it offers the user. A simple example clarifies these concepts: environmental information, such as geographical position, allows us to provide location-specific information: for example, detail on exhibits in a museum. Other user information, such as the identification and presence of another person, allows us to create a peer-to-peer network for informal chat. But the combination of the two may allow us to determine that the other user is a curator, and we can provide the mechanisms for one to give a guided tour to the other. The combination of models is potentially richer than each on their own.

Instead of a rigid definition, our intention is to provide a hierarchical description of context as a *dynamic process with historical dependencies*. By this we mean that context is a set of changing relationships that may be shaped by the history of those relationships.

A snapshot of a particular point in the ongoing context process can be captured in a *context state*. This contains all the elements currently present within the ongoing context process that are relevant to a particular learning focus, such as the user's current *project, episode,* or *activity*.[1] A user may, at any one time, be engaged in a number of simultaneous

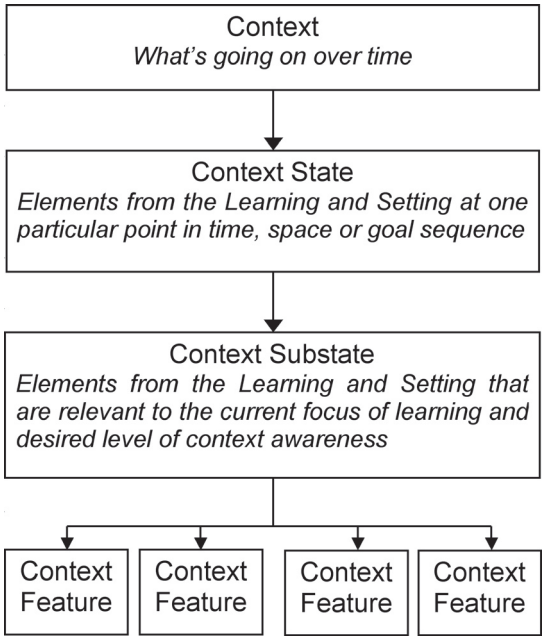

Figure 1. Context hierarchy: This model can be conceptualized using a film metaphor. The context is the whole film: plot, characters and all. The context state is the current scene, with the context substate being a frame in that scene. Context features are the individual actors, props and so on within the scene. This captures both the hierarchical compositionality of context and its temporal nature.

activities and episodes that relate to one project, and they may have several ongoing projects: each of which has its own set of relevant activities and episodes. A *context substate* is the set of those elements from the context state that are directly relevant to the current learning or application focus: that is, those things that are useful and usable. *Context features* are the individual, atomic elements found within a context substate and each refers to one specific item of information about the user or

Beale, *continued on p.10*

Cellularity + Development = Self-Repair?

Long-term robustness may only be gained through some of sort of self-repair, but that property unfortunately remains elusive in computing systems. In the past, this problem has been approached in two main ways: either conventionally, by specifically designing 'fault-tolerant' systems; or in a less-contrived fashion, by using evolutionary techniques in a faulty environment. These two techniques, while showing some success, remain flawed and/or cumbersome in most cases. In our current work, we adopt a different approach. Rather than designing or evolving a system to repair itself, we define a framework that proves sufficient, in most cases, to obtain good self-repair abilities: that is, if one follows the framework, self-repair comes for free. This framework hypothesises that any developmental system over a cellular structure naturally exhibits good-to-excellent self-repair properties. The current preliminary work confirms this hypothesis beyond our expectations.

The main idea behind this work is that a cellular, continuous, developmental process can help to solve most of the problems arising in classical systems. First, reconstruction is all the easier—if the system is able to construct itself in the first place—if it is developmental. Second, if only a fully-developed system is stable, then any faults, whatever their type, should create an unstable state that naturally redevelops into a perfect working state. This last aspect implies a continuous growth process, exactly as in natural systems.

Actually, the original inspiration of our system lies in Nature: most natural organisms are good examples of flexible, self-repairable systems. It can be reasonably said that the growth process is at the heart of these properties and, more specifically, cellular growth. While this cellular aspect has been adopted in some previous studies, its usefulness for self-repair has often been neglected. In earlier work,[1] we used only this cellular, decentralised aspect—without growth—as a means of getting some sort of fault-tolerant computation. Cellularity allows for an easy isolation of the faulty part and—when coupled with a developmental process and each cell being totipotent (able to generate new cells)—allows for easy reconstruction. In addition, the parallel, decentralised nature of cellular systems removes the requirement to have an explicit global plan and reduces the probability that any part is central to proper functioning. These last two aspects are crucial for self-repair.

The idea of having a cellular, developmental system, particularly to increase the robustness of digital circuits, is not new in itself. However, previous works—by Mange, et al.[2] or by Macias and Durbeck[3]—require the active detection of errors, must be hand designed, and do not cope well with bulky error. On the other hand, our system—as demonstrated so far—does not require any sort of error-detection, as such. It is simpler in its deployment as it can be evolved easily and recovers more often than not from bulky errors. Of course, there is a price to pay: in the form of the lack of any guarantee, before testing, of the exact recovery abilities.

This is a development of our early research,[4] inspired by recent work done by Miller.[5] The model is an extension of the well-known model of cellular automata (CA): the state of the cells encode the given system, a reconfigurable digital circuit, for example. Like a classical CA, the result of an internal program execution in each cell will determine its new state. However, unlike the CA model, the cell may also alter the state of its eight neighbours (see Figure 1). This is where the developmental process kicks in. The neighbour's state-rewriting property, which is at heart of the whole process, entails a series of questions on the order of update of the cells. This update is fully parallel, but there exists a priority mechanism when it comes to rewriting. This order, de facto, entails some orientation essentials (up and down, left and right) to develop non-symmetrical organisms. We used a straightforward extension to Cartesian Genetic Programming[6] to evolve the developmental program.

While our work is still at early stage, we have demonstrated on both figurative and functional examples that—while these organisms were not evolved explicitly for self-repair purpose, there was no selection pressure in that way—more than a third

Capcarrere, *continued on p.10*

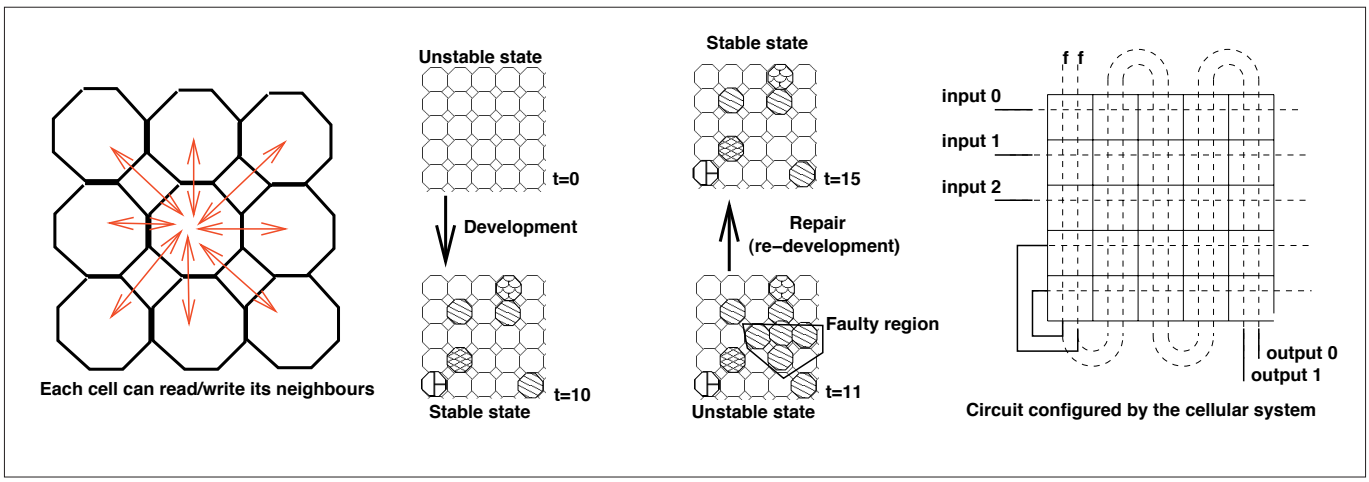

Figure 1. On the left is shown how a cell interacts with its neighbours. In the middle is an illustration of the self-repair process. The right panel shows the circuit encoded by the cellular system.

Integrating situated interaction with mobile awareness

A common problem for many learning organisations is the lack of any formal means to contact staff members rapidly without the disclosure of potentially personal information, such as mobile phone numbers. One 'traditional' method of interaction between staff and students involves the sticking of notes on office doors. A lecturer wanting to leave messages on their office doors from a remote location often requires the help of another member of staff, such as a receptionist, to put up a note on their behalf. This method has worked for a long while, but there are intrinsic issues relating to a lack of security and privacy, and the practical problems of notes falling off doors or the lack of timeliness in posting the note.

To address these issues, a system has been designed and implemented that exploits mobile and other technologies to provide an increased level of interaction between staff and students. The primary focus was on creating a system based on the concept of *situated interaction*[1] in an attempt to bridge the gap that currently exists between the learner and the instructor. This also allows the increased degree of mobility and remote accessibility necessary for modern learning situations.

The system, called IMMS, involves the placement of a number of display units on the office doors of various members of staff in the department, acting as information and messaging terminals for students. Members of staff may write a message and have a picture displayed on the unit; the content displayed on the units is set by its owner via a remote-access web-based management system. For staff users who do not have the time or opportunity to sit down at a computer and log in to the management system, an SMS-based system is available, allowing users to update their display unit content by sending a text message from their mobile phone to the IMMS server. This may be of significant use to staff members who have a great many department-based commitments, but are frequently away from the department for whatever reason.

Student members of the department are not only able to view the image and textual message set by the owner of the unit, but are also able to send messages to the owner via the display unit interface. A student calling in to see a lecturer and finding an empty office may use the web-based interface on the display unit to compose a short message and mark it urgent or non-urgent. In the case of an urgent message being composed by a department member, IMMS will either send a text message to the owner's mobile phone or store the message in the management system for display next time the user logs into the system. The routing of the messages sent via the display unit depends on the configuration settings prescribed in the personal profile of the user. This intelligent routing of messages extends similar work on situated door displays by Cheverst,[2] which had non-intelligent SMS notification.

Russell Beale* and Matthew Jones
University of Birmingham
School of Computer Science
*E-mail: R.Beale@cs.bham.ac.uk

References

1. N. A. Streitz, C. Röcker, Th. Prante, R. Stenzel, D. van Alphen, *Situated Interaction with Ambient Information: Facilitating Awareness and Communication in Ubiquitous Work Environments,* **10th Int'l Conf. on Human-Computer Interaction (HCI International),** June 22-27, 2003.
2. K. Cheverst, D. Fitton, A. Dix, and M. Rouncefield, *Exploring situated interaction with ubiquitous office door display,* **Proc. Int'l Workshop on Situated Interaction at CSCW,** New Orleans, November 2002.

Submit your paper to the Journal

The AISB Journal publishes high-quality papers presenting original and substantial research work in the areas of interest of Artificial Intelligence, the Simulation of Behaviour, Cognitive Science and any related fields. Interdisciplinary submissions are particularly welcome.

Eduardo Alonso and Geraint Wiggins are the AISBJ managing editors. The Editorial Board consists of world-leading researchers in knowledge representation and reasoning, planning, machine learning, natural language processing, robotics, vision, AI industrial applications, cognitive sciences, agents and multi-agent systems, evolutionary computing and computational creativity.

John Barnden, The University of Birmingham, UK
David Bree, University of Manchester, UK
Paul Davidsson, Blekinge Institute of Technology, Sweden
Ramon López de Mántaras, Artificial Intelligence Research Institute, Spain
Nick Jennings, University of Southampton, UK
Rob Milne, Intelligent Applications, UK
Ulrich Nehmzow, University of Essex, UK

Papers must be e-mailed as PDF files to eduardo@soi.city.ac.uk. Please use the LaTeX template for submissions, available at the AISBJ web-page (http://www.aisb.org.uk/aisbj/index.html). Papers are limited to 25 pages and should include only text and black-and-white line diagrams. Authors requiring more complicated diagrams, such as half-tones and photographs, should contact the editors.

Exemplar versus structural knowledge in system control

One area of interest within cognitive psychology is the study of people dealing with dynamic systems. Researchers investigate what types of knowledge are helpful for controlling systems, what type of knowledge people acquire spontaneously when they learn to handle a new system, and how this knowledge can be taught efficiently. In basic research it is common to use simple artificial systems simulated on computers and let subjects explore these under different conditions, observe performance with varying goal states, or ask questions about the structure of the system.

An important classification of knowledge types distinguishes between exemplar knowledge—which represents specific instances of input values and the corresponding output values—and structural knowledge, defined as general knowledge about the variables of a system and their causal relations. Research indicates that subjects largely prefer exemplar knowledge over structural knowledge. The former is most successful when the system has an small problem space, or when the same goal state has to be attained repeatedly: such that only a small part of the problem space is relevant. Simulation studies have shown that small dynamic systems can be controlled by using either declarative representations of specific actions,[1] or learned production rules that also represent specific interventions.[2] In conditions where subjects have to deal with huge problem spaces, the exemplar strategy is no longer useful. Instead, it is more reasonable to use structural knowledge to navigate through the problem space.

In a series of experiments,[3,4] I used a static system of four switches and four lamps whose 16 states can be presented to the subjects as 'system states' or as 'graphical patterns' (without informing them about the patterns being system states). The idea is that if subjects really acquired nothing but specific instances when dealing with the system, there should be no difference between the conditions with or without causal interpretation.

In fact, there were significant effects of causal interpretation: in one experiment, subjects saw all possible system states during the learning phase. In a recognition task, where they had to decide if they had seen the presented patterns in the learning phase, the group with causal interpretation discriminated much better. In another experiment, where subjects saw only a subset of possible system states, correct recognition did not depend on causal interpretation. The assumption that subjects reconstructed system states based on structural knowledge can explain this pattern of results: such a reconstruction is only helpful when all targets in the recognition task are possible states. After being debriefed about the causal nature of the stimuli, subjects were given causal judgment tasks. In two experiments I found significant positive effects of causal interpretation on causal judgment. Marginally-significant effects in favor of causal interpretation were also found in control tasks.

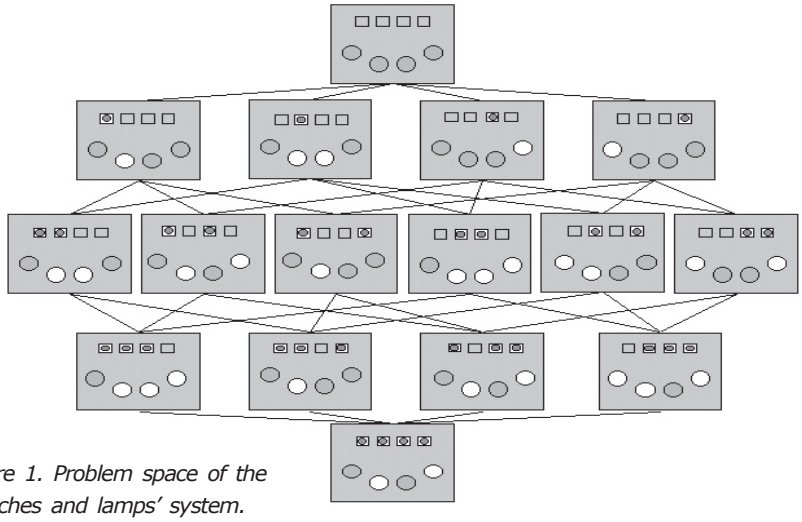

Figure 1. Problem space of the 'switches and lamps' system.

That subjects acquired both exemplar and structural knowledge in the conditions with causal interpretation, but only exemplar knowledge without, is supported by our analyses. These show that contingencies between recognition and completion were below maximum in the former case, close to maximum in the latter. This can be expected when both tasks are tapping the same knowledge base.

I developed a model in the ACT-R cognitive architecture that instantiates the described interpretations. In this model, using exemplar knowledge is the primary strategy. This is backed up by the reconstruction of system states based on structural knowledge in conditions with causal interpretation. Structural knowledge is modelled as associations between pairs of switch states and lamp states. The model reproduces the findings referring to single tasks quite well (discrimination, response times), but not the contingency results. One problem is that the built-in structure for representing instances is too powerful. Currently, I am extending the model so that it can learn this structure autonomously, based on competitive chunking.[5]

Wolfgang Schoppek
University of Bayreuth
E-mail: wolfgang.schoppek@uni-bayreuth.de

References

1. C. Lebiere, D. Wallach, N. Taatgen, *Implicit and explicit learning in ACT-R,* **Proc. 2nd Euro Conf. on Cognitive Modelling,** p. 183, 1998.
2. D. Fum, A. Stocco, *Outcome evaluation and procedural knowledge in implicit learning,* **Proc. 25th Annual Conf. of the Cognitive Science Society,** 2003.
3. W. Schoppek, *The influence of causal interpretation on memory for system states,* **Proc. 23rd Annual Conf. of the Cognitive Science Society,** p. 904, 2001.
4. W. Schoppek, *Stochastic independence between recognition and completion of spatial patterns as a function of causal interpretation,* **Proc. 24th Annual Conf. of the Cognitive Science Society,** p. 804, 2002.
5. W. Schoppek, *Emerging structure in ACT-R: Explorations in competitive chunking,* **Proc. 5th Int'l Conf. on Cognitive Modeling,** 2003.

The role of emotions in modular intelligent control

Continued from p.1

communicative channel's information comes in the form of tagged text, for which techniques already exist for producing speech and facial animation.[7] Since there is a non-unique correspondence between meanings and facial displays, the selection of a facial display should be based on mental and emotional context.[5] What we propose is to use the mood state, delivered by the DER, to discriminate between the facial displays that correspond to a meaning. A second issue will be to merge, mix, or select the facial displays produced by the communicative and emotional channels.

The DER is unique in its support of the different temporal characteristics of onset and decay of different emotions, and in its combination of a variety of time courses (primary emotions, secondary emotions, and moods). These attributes are often overlooked, but presumably serve an adaptive purpose. Some species even seem to display sexually-determined differences in onset curves for some emotions. Our current work involves evaluating both the DER tool and the DER itself. The DER tool will be used to develop communicative agents. Its usability will be judged by the ability of artists to manipulate the perception of the agents: how persuasive they are, for example. The DER will also be evaluated in terms of its utility when incorporated into models of primate social behaviour. This work is so far in its infancy,[8] but we will have a new postgraduate student working on this from September 2004.

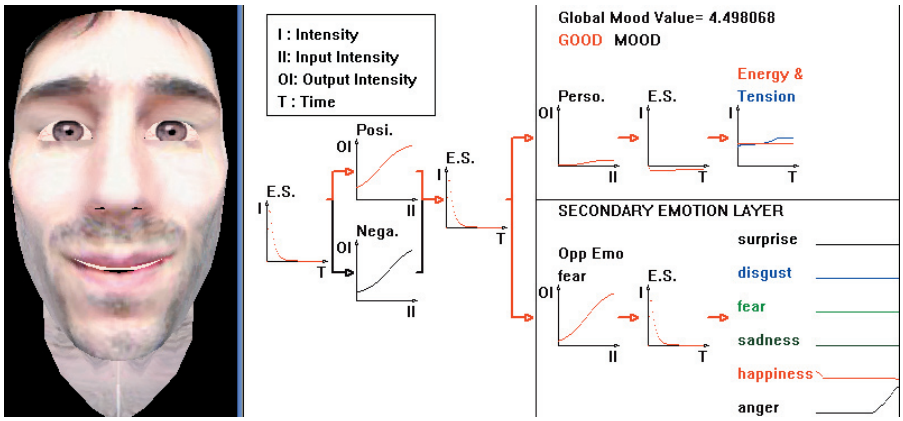

Figure 2. The dynamic emotion representation tool, which allows editing the response curves for secondary emotions and playing the real-time combination of expressions on the model.

Joanna J. Bryson*, Emmanuel A. R. Tanguy and Philip J. Willis
Artificial models of natural Intelligence (AmonI) Group
Department of Computer Science
University of Bath
E-mail: *jjb@cs.bath.ac.uk
http://www.bath.ac.uk/comp-sci/ai/AmonI.html

References

1. R. W. Picard, **Affective Computing,** The MIT Press, 1997.
2. A. Sloman, *Beyond shallow models of emotions,* **Cognitive Processing 2** (1), pp. 177-198, 2001.
3. P. Ekman, *Facial expressions,* in T. Dalgeish and M. Power (eds), **Handbook of Cognition and Emotion,** pp. 301-320, John Wiley, 1999.
4. E. Tanguy, P. Willis, and J. J. Bryson, *A Layered Dynamic Emotion Representation for the Creation of Complex Facial Expressions,* in T. Rist, R. Aylett, D. Balin, and J. Rickel, (eds), **Intelligent Virtual Agents**, pp. 101-105, Springer, 2003.
5. C. Pelachaud and I. Poggi, *Facial performative in a conversational system,* **Proc. Workshop on Embodied Conversational Characters (WECC),** 1998.
6. J. B. Bavelas, and N. Chovil, *Visible acts of meaning: An integrated message model of language in face-to-face dialogue,* **J. of Language and Social Psychology 19** (2), pp. 163-194, 2000.
7. C. Pelachaud, and M. Bilvi, *Computational model of believable conversational agents,* in M. P. Huget, (ed.), **Communication in MAS: background, current trends and future,** Springer-Verlag, 2003.
8. J. J. Bryson, *Where should complexity go? Cooperation in complex agents with minimal communication,* in W. Truszkowski, C. Rouff, and M. Hinchey, (eds.), **Innovative Concepts for Agent-Based Systems,** pp. 298-313, Springer, 2003.

Help make a difference in the AISB

We are always looking for new talent to make the AISB a better organization: offering more services better-suited to the needs of our members.

If you have ideas about what the Society lacks, how it could do better, or the best way to increase membership, then please consider joining the AISB Committee.

To find out more, please contact the AISB Chair, John Barnden, at:

chair@aisb.org.uk

CONFERENCE PREVIEW

RoboCup 2004

Lisbon, Portugal, 27 June – 5 July 2004

RoboCup is an international long-term initiative with the goal of boosting research in artificial intelligence and robotics. A typical RoboCup consists of a series of annual competitions, including robotic soccer and rescue tournaments. These are tied together with a symposium that brings together a broad spectrum of scientific research relevant to achieving success in the various contests.

A central credo of the initiative is that—instead of developing isolated, specialized AI or robotics solutions working only under restricted conditions—the RoboCup competition forces the development of robust solutions which will work even under adversarial conditions. A central motif is the integration of different AI and robotics approaches in a coherent architecture that really works. This requirement weeds out methods that work fine in the lab, but are useless for more demanding scenarios because they are too time-consuming or brittle. RoboCup—with its widespread scenarios, platforms, and challenges—puts well-established methods to the test and encourages creativity in coming up with new approaches.

Progress through the years

Since the first competition in 1997 in Nagoya, Japan, RoboCup has come a long way. In the early competitions, hardware robots had a hard time even finding and identifying the soccer ball, let alone moving towards it and controlling it. In the last few years, however, object identification, self localization, and both robot and ball control have improved to an extraordinary degree. Essential for this success has been the use of Markovian localization techniques and particle filters.

In the middle-size league where teams can freely design and experiment with their own robotic platform, a large variety of sensors have been employed. These range from laser scanners, which provide highly-accurate self-localization in environments with well-defined walls,[1] to specially-designed isometric omni-directional mirrors that minimize the image-reconstruction error of a scene,[2] use specialized vision chips, and minimal infrared-distance detectors.[3] A large variety of architectures have been developed to allow integration of these different sensor modalities.

Particularly striking has been the progress in the legged-robot soccer league, using AIBO robots as players (see Figure 1). In 1998, at the first legged-league competition, the robots could hardly walk: an attempt to kick the ball would often result in their falling over. In contrast, the current teams have an impressively-versatile repertoire of complex behaviours, from which the robots select smoothly, depending on the situation. This is a step towards humanoid robots which, in the long-term grand vision of RoboCup 2050, 'will play a soccer game against the human world champions and win'. RoboCup is still far from that point, and the current humanoid competitions consist mainly of challenges like walking certain patterns or kicking the ball into a goal without falling over. However, the RoboCup competitions and challenges move fast and have also spurred new levels of research effort in material science which is—aside from AI and engineering aspects—believed to be essential to achieve the grand vision of the project.

Because learning and adaptation methods in the hardware leagues are still very limited—mainly concentrating on auto-calibration of sensors—higher-level issues like coordination and cooperation are addressed in the simulation league. The simulated robots mean that there can be less emphasis on controlling physics and hardware, allowing researchers to concentrate on implementing multi-agent AI. Approaches have included reinforcement learning[4]—allowing individual agents to pick up skills and skill combinations, multi-level reinforcement learning[5]—and the use of collaboration graphs[6] to obtain coherent play moves by groups of agents.

This year's competition

At time of writing, RoboCup 2004, chaired by Pedro Lima of the Instituto Superior Técnico, Lisbon, is expected to bring together 350 teams with 2000 active participants, and 600 robots. Most leagues will have to face new challenges: like moving from well-lit fields to those with whatever artificial light is available on site. The simulation league will move from a 2D to 3D scenario: this is expected to add a significant amount of complexity. Our team will also participate in the RoboCup 2004 soccer-simulation league, using a skill-learning approach that estimates the outcome of whole behaviour sequences via pattern matching. This is achieved through a combination of self-organizing maps with information-theoretic decompositions.

Daniel Polani
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
E-mail: d.polani@herts.ac.uk

Figure 1. A scene from the RoboCup Legged League, Padua, 2003. Here, both teams are restricted to the same AIBO robots: their strength therefore arises purely from the software used to control them. © 2003 The RoboCup Federation

References

1. http://www.cs-freiburg.de/
2. http://www.dei.unipd.it/~emg
3. http://www.ais.fraunhofer.de/~ansgar
4. M. Riedmiller, and A. Merke, *Karlsruhe Brainstormers—a reinforcement learning approach to robotic soccer,* **RoboCup 2001: Robot Soccer World Cup V**, LNAI 2377, Springer, 2002.
5. P. Stone, **Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer**, MIT Press, 2000.
6. J. R. Kok, M. T. J. Spaan, and N. Vlassis. *Multi-robot decision making using coordination graphs*, **Proc. 11th Int'l Conf. on Advanced Robotics (ICAR)**, 2003.

BOOK REVIEW

Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics

Paul W. Glimcher

Publisher: *MIT Press.* ***Hardback:*** *Published March 2003, 400pp, £25.50.*

Glimcher's argument, as presented in this book, can be summarised as follows: the main theoretical framework of neuroscience—reflex theory—is inadequate, and should be replaced by the mathematical formalisms offered by economics (e.g., utility theory, game theory). Beyond neuroscience, he argues that the use of these formalisms has implications for questions such as dualism, determinism, consciousness, and free choice.

The first half of the book is devoted to a detailed history of the reflex doctrine. Among other things, it reviews its predecessors from the classical period (Hippocrates, Galen), presents Descartes' dualistic view of the world (which started the idea of a reflex theory), and discusses the theory's modern application, such as by Pavlov and Sherrington. Glimcher concludes this section by arguing that any reflex theory is inadequate for an understanding of neural processes because of two sets of difficulties: those related to explaining empirical evidence, and theoretical problems that, mainly, result from the fact that reflex theories fall prey to Gödel's incompleteness theorem.

In the second half, a 'manifesto' for a new approach—neuroeconomics—is presented. Its review of the history of classical probability theory (including works by Pascal, Laplace, Bernoulli, and Bayes) focuses on the concept of utility, which is central to neo-classical economics. The link between neuroscience and economics is made by a review of recent experiments showing that anticipated reward, and thus utility, can predict the level of activations in lateral intraparietal (LIP) neurons. However, Glimcher argues that classical probability theory is not sufficient to account for all experimental data, and that more recent results in Bayesian probability theory and game theory must be used as well. This leads to a review of these mathematical theories, and then to experiments (mainly taken from the author's research) showing their application in neuroscience. While the argument centres on how formalisms such as probability theory and game theory may help develop more powerful theories in neuroscience, there is also a brief discussion about how progress in neuroscience may improve our understanding of economic behaviour. The final chapter of the book aims to show that neuroeconomics has important implications for classic philosophical questions such as dualism, determinism, and consciousness.

The book has two main strengths. First, it provides an entertaining historical introduction to some important topics in neuroscience and to mathematical formalisms used in economics. Second, it also presents some intriguing results showing how these formalisms can account for data in single-cell experiments that cannot presumably be accounted for by standard 'reflex theories' in neuroscience.

However, it also has a few weaknesses,. We found that it told us a lot about neuroscience and economics, but little about neuroeconomics. Moreover, the historical parts and the analyses they contain are not always accurate. To be fair, this would be difficult to achieve in a single book given the number of fields reviewed. For example, Glimcher argues that Sherrington's reflex theory, etc., are doomed to failure because Gödel's incompleteness theorem shows that formal theories powerful enough to include arithmetic cannot be proven both complete and consistent. This may be true, but Gödel's theorem also applies to the type of formalisms Glimcher recommends (even more perhaps, as the former are mathematically more complicated), and indeed, to most modern formal scientific theories.

Finally, Glimcher argues that neuroeconomics helps solve key questions in philosophy. We found his argument unconvincing. For example, neuroecononomics is shown to overcome dualism, which is (implicitly) presented as dominant in neuroscience. We would argue that mainstream neuroscience is resolutely monist.[1,2] Indeed, even classic symbolic cognitive science is unambiguously anti-dualistic and materialistic.[3] Moreover, we think that Glimcher's term 'indeterminate monism' is inadequate, since it confuses determinism with causality. Bunge[4] pointed out that causality is only one category of determinism. The phenomena that Glimcher wants to capture with the term 'indeterminate monism', in fact, are statistically—but not causally—determined.

In spite of these shortcomings, the book provides a useful introduction to neuroscience and to several formalisms used in (neuro)economics. It raises intriguing questions and highlights the heuristic power of using 'high-level' formalisms in neuroscience, as advocated by Marr.[5]

Fernand Gobet and Guillermo Campitelli
Brunel University

Fernand Gobet is co-author of ***Perception and Memory in Chess*** *(1996),* ***Techniques for Modelling Human Performance in Synthetic Environments*** *(2003), and* ***Moves in Mind*** *(in press). Guillermo Campitelli is a Research Fellow with interests in theoretical psychology and psychology of expertise, consciousness, and memory.*

References

1. M. Bunge, **The mind-body problem.** Pergamon, London, 1980.
2. P. S. Churchland and T. J. Sejnowski, **The computational brain,** MIT Press, Cambridge, MA, 1992.
3. H. A. Simon and A. *Newell, Computer science as empirical inquiry: Symbols and search,* **Communications of the Association for Computing Machinery 19,** pp. 113-126, 1976.
4. M. Bunge, **Causality: The place of the causal principle in modern sciences,** Harvard University Press, Cambridge, MA, 1959.
5. D. Marr, **Vision,** Freeman, San Francisco, 1982.

Let us fund your travel

The AISB offers two student travel grants per year, each of £300. In 2005 one of the two is going to be given preferentially to IJCAI '05. If you'd like to apply for the cash, either for this or another conference, please see the details at:

http://www.aisb.org.uk/awards.html

BOOK REVIEW

Imitation in Animals and Artifacts
K. Dautenhahn and C. L. Nehaniv (eds.)

Publisher: *Bradford Book.* **Hardback:** *Published June 1 2002, 612pp, £38.10.*

It is easy to believe that imitation is an easy form of learning. If we catch someone cheating on an exam, we might say that they hadn't really learned the material, that they were 'just copying'.

Nonetheless, learning to imitate is a complex thing, which is challenging/impossible for many kinds of animals and difficult to implement in artificial systems such as computers and robots. It has been argued that our ability to imitate each other's actions, often in an effortless fashion, is a more defining characteristic of being human than our use of language.

This well-balanced collection of studies on the topic of imitation aims to convince us that imitation is a difficult task, and a fascinating one: this is true both from the perspective of human and animal intelligence, and for people attempting to create systems to cope with complexity. A particular feature of the book is that it draws examples from a number of areas: observations of animal behaviour, experiments with people attempting to imitate each other, and robotic and computer systems that carry out some form of learning by imitation. Notably, these are well linked together, with a number of cross-references between the chapters. This reflects the publication's origins: it began at an AISB symposium and was subsequently developed into a book, allowing authors to introduce such linking material.

In particular, a number of chapters attempt to tease out what is true imitation: in particular, the differences between imitation and other kinds of learning. For example, how can we distinguish imitation—which is based on a 'thoughtful' analysis of the action—from a 'brain-dead' trial-and-error approach, where the imitator compares their attempts with the imitatee and improves their performance through, e.g., a hill-climbing approach? What distinguishes true imitation, where the imitator is attempting to reproduce something of the intention of the action, from 'mere' mimicry, where they are just reproducing the surface behaviour?

Why is imitation hard? Firstly, because the way in which the imitatee carries out the action to be imitated is not typically transparent to the imitator. If I carry out a complex action, such as washing a cup or juggling, an observer needs not only to observe what I am doing but to reconstruct that in terms of the muscular actions that they must carry out to achieve the same result.

This problem permeates many of the chapters: how do people/animals manage to imitate nonetheless? How can we create computer systems that imitate? This provokes further questions that are explored in other chapters. One example of particular interest is how we learn a skill by imitation: i.e. how do we learn to generalize from a number of imitations of individual actions? How do we fill in the gaps?

Secondly, a concern in a number of chapters is what is termed the correspondence problem. This asks how imitation can be carried out where the imitator and imitatee have differences. How can a dolphin imitate its trainer? How can a parrot learn to speak some human words despite having a radically different voice-production system? There would seem to be a lot of scope for links between this and work on analogy-formation such as that of Hofstadter.[1]

Why should we care about imitation? One reason is that imitation is one of the core learning mechanisms used by humans. Thus it is of interest for researchers interested in building systems that work well with humans or operate in a human-like fashion. Furthermore, we would like to understand where imitation arose on an evolutionary timescale: is imitation a development unique to humans, or is it likely to have arisen at an earlier stage in the evolution of mind?

A second reason addressed by a couple of chapters in the book is that of programming computer systems by demonstration of the desired output rather than by specifying how to do the desired behaviour. This offers a third paradigm for describing actions to computers, a 'do what I do' to sit alongside traditional 'do what I tell you to do' programming and 'do what I ask you to do' specification. Again, the idea of generalization is important for this kind of task: the user should be able to demonstrate using a fairly small number of examples and have the computer create a generic behaviour that captures the implicit intention behind the individual cases. There are connections here to Partridge's[2] notion of data-defined problems, i.e. problems that are defined by examples rather than by a formal description of the problem.

It is difficult to give an overall summary to such a diverse book as this. Nonetheless, despite the diversity of the topics covered, there is a rich interconnectivity between the various topics and approaches. Most people interested in intelligence and intelligent behaviour should find something of interest in this book.

Colin G. Johnson
University of Kent

Colin Johnson is a lecturer in computer science at the University of Kent at Canterbury. His research interests include bioinformatics, computing inspired by the natural world, and the application of AI to music and media technologies.

References

1. Douglas R. Hofstadter and the Fluid Analogies Research Group, **Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought**, Penguin, 1998.
2. Derek Partridge, *Non-programmed computation*, **Communications of the ACM 11** (electronic supplement), pp. 293-302, November 2000.

Get free books!

We are always looking for qualified reviewers of books sent in for review. If you're interested in finding out what we have on offer and what we want from you, please go to:

http://www.aisb.org.uk/qbooks.html

Cellularity + Development = Self-Repair?

Continued from p.3

of them exhibited excellent self-repair qualities. Less central to our study, but essential in terms of usability, we also demonstrated the excellent evolvability of the system. We successfully evolved a developmental process that configures a digital circuit based on a cellular structure to do a 5-bit full adder in more than 60% of the runs. Around 35% of these recovered perfectly from all errors concerning less than 20% of the circuit, which compares very favourably with other systems.

Current results allow us to conclude that—while using the framework is not a sufficient condition to gain self-repair ability—it provides a privileged environment in which to gain these, and its high evolvability renders it practical to use.

Dr Mathieu Capcarrere
Computing Laboratory
University of Kent
http://www.capcarrere.org/mathieu

References

1. L. Righetti, S. Shokur, and M. S. Capcarrere, *Evolution of fault-tolerant self-replicating structures,* **Proc. ECAL'03, Lecture Notes in AI 2801,** pp. 278-288, Springer-Verlag, 2003.
2. D. Mange, M. Sipper, A. Stauffer, and G. Tempesti, *Towards robust integrated circuits: The embryonics approach,* **Proc. IEEE, 88** (4), pp. 516-541, 2000.
3. N. J. Macias and L. K. Durbeck, *Self-assembling circuits with autonomous fault handling*, **Proc. 2002 NASA/DoD Conf. on Evolvable Hardware,** pp. 46-55, 2002.
4. M. S. Capcarrere, **Cellular Automata and Other Cellular Systems: Design & Evolution,** Ph.D. Thesis No 2541, Swiss Fed. Inst. of Tech., Lausanne, 2002.
5. J. F. Miller, *Evolving developmental programs for adaptation, morphogenesis, and self-repair,* **Proc. ECAL'03, Lecture Notes in AI 2801,** pp. 256-265, Springer-Verlag, 2003.
6. J. F. Miller and P. Thompson, *Cartesian Genetic Programming,* **EuroGP'00, Lecture Notes in Computer Science 1802,** pp. 121-132, Springer-Verlag, 2000.

Models for mobile context awareness

Continued from p.2

their setting (e.g., their current task or location). In our description of context, context features are indivisible and refer to only one item of relevant information about the user or their setting. The context hierarchy is shown in Figure 1.

This system has been implemented in Java using web services, and forms the basis of the MOBIlearn context architecture, which has undergone preliminary trials[2] and will soon undergo major evaluation in museum, MBA, and medical scenarios.

Russell Beale,* Peter Lonsdale† and Mike Sharples†
*School of Computer Science
†School of Engineering
University of Birmingham
E-mail {r.beale, p.lonsdale, m.sharples}@cs.bham.ac.uk

References

1. G. Vavoula, and M. Sharples, *KLeOS: A personal, mobile, knowledge and learning organisation system,* **Proc. IEEE Int'l Workshop on Wireless and Mobile Technologies in Education,** 2002.
2. R. Beale, and P. Lonsdale, *Mobile Context Aware Systems: the intelligence to support tasks and effectively utilise resources,* **Proc. Mobile HCI 2004,** Glasgow, (to appear).

Don't forget to vote!

Election papers went out with this issue, giving you the chance to elect committee members: you can even vote by e-mail...

Webmaster's Report

The current AISB website has been in existence for several years now. It was originally created by Simon Colton who continues to be an active force in its development. However, the committee has recently divided responsibility for the website in two. Simon is now in charge of developing new content—in particular articles and resources aimed at the general public and schools—while I have taken over the role of webmaster, concerned primarily with technical issues and day-to-day maintenance.

We have recently discussed whether there should be more content of direct use to members (as opposed to the general public) on the website and, in particular, whether there is scope to develop some sort of announcements or news facility. We are thinking of somewhere members could publicise recent publications or grants and where they might ask for collaborators. The most obvious approach would be some sort of noticeboard-based system, possibly with restricted access. If you feel this would be of use to you, please get in touch with me: if enough people are eager for this sort of service then we will look seriously into developing it. Also, if you have any other ideas for ways in which the website can serve you as a member of the Society, then please don't hesitate to contact me.

Louise Dennis, AISB Webmaster
University of Nottingham
E-mail: lad@Cs.Nott.AC.UK
http://www.cs.nott.ac.uk/~lad/

Contacting the AISB

If you have issues related to any of the services that AISB offers, you can contact the person responsible directly.

Areas of responsibility and e-mail addresses are available on page 11 of the Quarterly, and via the web site. Find us at:

http://www.aisb.org.uk/contact.html

Treasurer's Report

The finances of the Society continue to go from strength to strength. The surplus for the financial year 2003 was £3,006, which means that the Society has now accumulated a healthy reserve of £28,568.

Income and *Expenditure* accounts are shown below. Again, as in previous years, I would like to thank the AISB'03 convention organiser, Professor Mark Lee, for making the convention a success: and thus bringing important income to the Society. He did an excellent job in bringing together the programme and managing the finances.

Due to increase workload, at the Society's AGM held at Leeds University in March 2004, I offered my resignation as Treasurer and it was accepted. I am pleased to say that Dr Patrick Olivier has agreed to take on the role. Finally, I would just like to wish him, and the Society, every success.

Paul Chung
(Now former) AISB Treasurer
Loughborough University

Income	**2003**	**2002**
Membership fees	14,861	12,200
Insert/mailshots	0	309
Bank interests	192	256
Proceedings	0	347
AISB '01 Correction	0	-672
AISB Convention	4,500	6,043
Total	**19,553**	**18,483**
Expenditure		
Office		
Service charge	5,358	4,930
Postage	180	471
	5,538	**5,401**
Quarterly and Journal		
Service charge	684	3,409
Editor fees	3,600	0
Printing	3,184	3,560
Postage	797	1,118
	8,265	**8,087**
Miscellaneous		
Committee expenses	310	209
Travlel awards	300	200
ECCAI fees	780	719
Accountant fees	1,117	1,058
Sundry expenses	65	14
Bank Charges	172	106
	2,744	**2,306**
Total	**16,547**	**15,794**
Income - Expenditure	3,006	2,689
Tax	0	-6
Surplus	**3,006**	**2,683**

Interaction Studies Journal discount for AISB members

John Benjamins Publishing Company is now offering a discounted subscription rate to individual AISB members of €79.50 (normal rate €240) for the three issues of Volume 5 (2004). This includes the printed copies, postage, VAT, and full online access (prepaid and for private use only).

Interaction Studies, subtitled *Social Behaviour and Communication in Biological and Artificial Systems,* is edited by Kerstin Dautenhahn University of Hertfordshire and Harold Gouzoules of Emory University, USA. The journal welcomes papers that analyze social behaviour in humans and other animals, as well as research into the design and synthesis of robotic, software, virtual, and other artificial systems. These include applications such as exploiting human-machine interactions for educational or therapeutic purposes. Papers can be experimental, computational, or theoretical studies. A special issue on epigenetic robotics is currently in preparation.

Abstracts and a free-view issue are available online.[1] To get your discount, please e-mail your order to *subscription@benjamins.nl*, with, "AISB Member 2004" in the text.

Reference

1. http://www.benjamins.com/cgi-bin/t_bookview.cgi?bookid=IS%205:1

Editorial

Given that we, unusually, have a little space this quarter, I thought I'd let you know what I've been up to besides bringing out the Q. In June, *Focus* magazine asked me to write a feature on whether robots could take over the world. Though normally I don't like to be given a subject, especially not one like this, I decided it was worth doing. Some other journalist would have written it if I hadn't, and this gave me an opportunity to help shape an argument that, in my mind, has become rather overblown.

The feature—timed to coincide with the release of the new *I Robot* movie, starring Will Smith—should be out in the August issue. (No, I didn't get to preview the film in order to write the article). Anyway, I hope that my piece sheds a little light on the subject for those with a techy but non-scientific background and who have probably only heard hype on this subject in the past. If you see it and feel strongly, feel free to e-mail both bouquets and brickbats.

Sunny Bains
Editor, AISB Quarterly
sunny@sunnybains.com

AISB COMMITTEE MEMBERS
Chair
John Barnden
University of Birmingham
chair@aisb.org.uk
Vice Chair/Co-Managing Editor, Journal
Eduardo Alonso
City University, London
vicechair@aisb.org.uk
aisbj@aisb.org.uk
Secretary/Webmaster/ Schools Liason
Louise Dennis
University of Nottingham
louise.dennis@aisb.org.uk
Treasurer/Travel Awards
Patrick Oliver
U. Newcastle-upon-Tyne
University of York
treasurer@aisb.org.uk
travel@aisb.org.uk
Membership
Gary Jones
University of Derby
membership@aisb.org.uk
Public Understanding of AI
Simon Colton
Imperial College London
sgc@imperial.ac.uk
Industrial Liaison
Vacant
Co-Managing Editor, Journal/Publications
Geraint Wiggins
City University, London
AISBQ
Sunny Bains
Imperial College London
aisbq@aisb.org.uk

AISBQ EDITORIAL BOARD
Eduardo Alonso
City University
Sunny Bains
Imperial College London
John Barnden
University of Birmingham
David Bree
University of Manchester
Lola Canamero
University of Hertfordshire
Simon Colton
Imperial College London
Kerstin Dautenhahn
University of Hertfordshire
Yiannis Demiris
Imperial College London
Louise Dennis
University of Nottingham
Colin Johnson
U. Kent at Canterbury
Gary Jones
University of Derby
Natalio Krasnogor
University of Nottingham
Kia Ng
University of Leeds
Patrick Oliver
U. Newcastle-upon-Tyne
University of York
Frank Ritter
Pennylvania State U.
Barbara Webb
University of Edinburgh
Geraint Wiggins
City University

Father Hacker's Guide
for the Young AI Researcher

Cognitive Divinity Programme
Institute of Applied Epistemology

About the Society
The Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB) is the UK's largest and foremost Artificial Intelligence society. It is also one of the oldest established such organisations in the world.

The Society has an international membership of hundreds drawn from academia and industry. Membership of AISB is open to anyone with interests in artificial intelligence and cognitive and computing sciences.

AISB membership includes the following benefits:

- Quarterly newsletter
- Biannual Journal
- Student travel grants to attend conferences
- Discounted rates at AISB events and conventions
- Discounted rates on various publications
- A weekly e-mail bulletin and web search engine for AI-related events and opportunities

You can join the AISB online via:
http://www.aisb.org.uk

You will know you have arrived when you are first asked to be a consultant. Not only is this a tribute to your growing reputation, but it will also help to supplement your otherwise meagre income as an AI researcher. So, to make the most of this opportunity, it is vital to know....

11. How to be a Consultant

1. Your main goal as a consultant is to ensure that you gain more consultancies. Your every action and recommendation must be driven by the question: will this result in a further invitation to consult? Identify the consulter who makes the hiring decisions and ensure s/he is delighted by your consultancy and will not only rehire you but will recommend you to others.
2. The golden rule of consultancy is that no one hires a consultant for the originality and independence of their opinion. Rather—whether it is to suggest what computational technique will best solve their technical problem, who to appoint to a prestigious position, or the best strategic direction for a research lab—the consulter wants you to provide authority and weight to what s/he intended to do all along. So, find out what this is, and support it.

Read someone else's mind with Hacker's PSYCHIC™ (Prejudices Sensed, Yearnings and Convictions Hypothesised, Inclinations Conjectured).

Using a huge array of sensors, PSYCHIC™ *monitors your target's every nudge and wink then infers and reports back his/her inner feelings.*

3. The experienced consulter will ensure the consultant is fully briefed about what is required of them. With an inexperienced consulter, the consultant will have to infer them from the consulter's body language.
4. Avoid at all costs shouldering any legal responsibility for your recommendations. Litigation is injurious to your time, bank balance, and reputation—the opposite of your main consulting goals. For guidance on formulating a contract that will avoid any liability, read the disclaimer on any commercial software.
5. Your consultant's report must underline your reputation in its weight and authority. Never use one short word when five longer ones will do. Carefully discuss all of the options before rejecting each of the alternatives to your preferred recommendation. The consulter must be able to argue that all opposing arguments have been properly considered, but rejected.

Produce weighty prose with Hacker's TOME™ (Treatise Of a Masterly Expert).

Given a rough outline of your preferred conclusion, TOME™ *will automatically construct an impressive consultant's report.*

6. Don't be afraid to require a substantial fee. A large fee is in the interests of both consultant and consulter: it will add weight to your opinion by underlining your reputation and the consulter's institution will be reluctant to reject a consultant's conclusions if they have cost it a lot of money.

Next AISBQ deadline:
20 August 2004

AISBQ reaches hundreds of Artificial Intelligence and Cognitive Science researchers in the UK, Europe, and beyond. Advertising and other information is available via:

http://www.aisb.org.uk/qhome.html

This page includes full guidelines for the submission of book/conference reviews and technical articles. Books available for AISB members to review are also listed.