AISB Quarterly No.120

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

# Robust navigation in unknown environments

After decades of research into mobile robotics, we still do not possess the machines we thought were 'just around the corner' in the early 1980s. We still hope to build autonomous mobile robots that operate in both 'everyday' and exotic locations. We need robots to explore places that we as humans can't: the deep ocean, Mars, or polluted or dangerous environments. Not getting lost requires answering the question "where am I?" But how can robots do this without a map? Can they make one themselves? The answer is yes, but—as yet—not reliably.

The ability of an autonomous vehicle to be placed at an unknown location in an unknown environment, have it build a map using only relative observations of the environment, and at the same time use this map to navigate, would make such a robot autonomous. This is simultaneous localisation and mapping (SLAM)—a key problem in mobile robotics.

The SLAM problem has been successfully tackled using probabilities to infer position and surroundings from a stream of sensor measurements. A big challenge was in overcoming the scaling problem: avoiding unbounded growth of computation with map size which would prohibit sustainable mobile autonomy. Over the past few years, there has been great progress in this area, mostly by analysing the underlying inference problem and exploiting its unique structure by application of bespoke estimators.

However, we still do not have SLAM-enabled robots capable of substantive deployments. The problem lies in perception. All the SLAM algorithms require an oracle of some description to associate each measurement with part of the inferred state of the world. If this 'data association' process goes wrong, terrible things happen. Vehicles get lost, maps get corrupted, and features are mapped as new when in fact they are previously-observed areas. SLAM algorithms can fail catastrophically because of poor data associations.

One facet of our research into mobile robotics is an attempt to face the robustness problem head on. We are combining both sensed geometry—resulting from 2D and 3D laser scanners (see Figures 1 and 2)—and texture and pattern information from cameras to extract complex high-dimensional scene descriptors. We use these to disambiguate the data-association problem, reducing the chances of erroneous measurement-feature pairings and making it easier to solve the 'loop-closing problem': i.e. recognising when, contrary to internal location estimates, the vehicle has returned to a previously-visited location.

Central to our approach is saliency detection: finding what is interesting and 'stands out' in a

***Paul Newman, University of Oxford***

*Continued on p. 7*

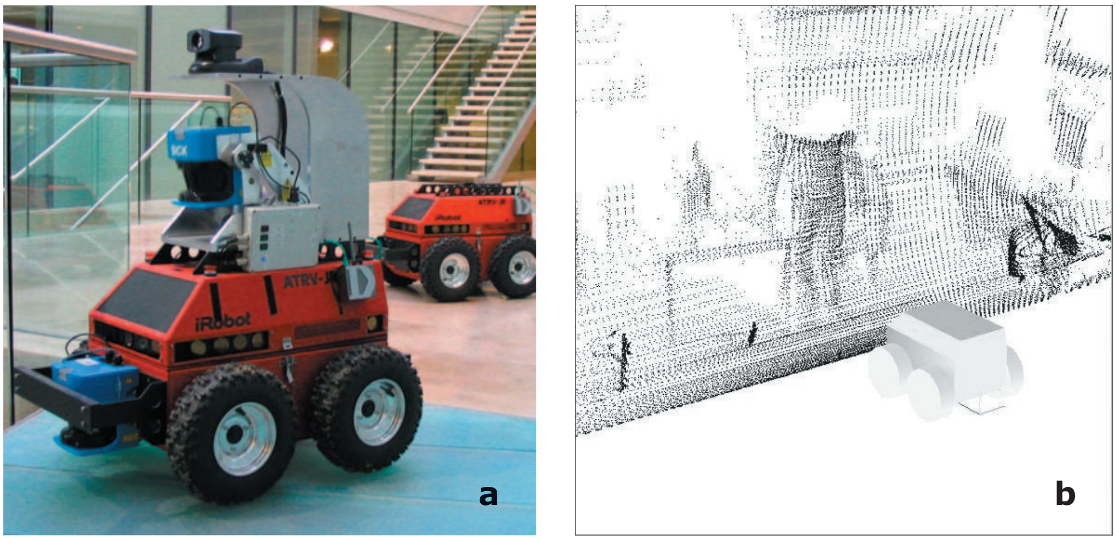


Figure 1. (a) Marge, an all-terrain autonomous vehicle with 2D and 3D laser scanners. Her partner, Homer, is shown in the background. (b) 3D laser scan taken from Marge of a student in an office environment.

***In this issue:***



## Would you like to host the AISB convention?

The AISB Convention is the major event in the Society's year. It is organised as a set of symposia based around a loose theme.

We are currently solicting proposals to run the convention in 2007.

Details can be found at:

**http://www.aisb.org.uk/convention/cfp.shtml**

# Logic, logicism, and logic-based AI

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

ISSN 0268-4179



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

**Editorial Assistant**
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



*"Logic, like whiskey, loses its beneficial effect when taken in too large quantities."*

– Lord Dunsany

In the late 70s and mid 1980s, Pat Hayes—writing in the well known *Naive Physics Manifesto* papers[1,2]—argued that AI needed to move away from working with simple domains, sparse axiomatic theories and the use of automated theorem-proving programs that had dominated logic-based AI research. The papers encouraged researchers to do three things. First, to adopt first-order predicate logic, with its well-understood proof theory and formal semantics. Second, to develop rich formalisms supporting non-trivial consequence classes. Third, not to rush into implementation, but rather to concentrate upon the task of building object-level formalisms first, then seek out effective ways to implement them. The papers were influential and—together with other qualitative, and predominantly logic-based approaches to AI (e.g. *Cognitive Robotics* and *Qualitative Spatial Reasoning*)—lead to an increasing interest in trying to tackle the general problem of encoding knowledge about the everyday world in computer programs.

The first signs of serious discontent with Hayes' programme arose with McDermott's attack on what he had identified as 'logicism'. The main problem was the central assumption that deductive reasoning was sufficient to model the domain. Hayes and McDermott's papers then appeared at the centre of a lively forum: however, the respondents have remained divided on several key discussion points.[3] While Hayes' programme no longer has the appeal that it once had, it still has the power to foster interest and generate debate.

A recent example of this appears in Shanahan's retrospective observations about Hayes' programme, while addressing a non-trivial benchmark problem in common-sense-reasoning literature.[4] For Shanahan, the doubts revolve around the following main points: the onerous task of building rigorously developed large-scale formal theories; the implausiblity of building these by hand; that few people seem to have seriously attempted this; and that without establishing some measure of 'ground truth', e.g. *grounding* or *anchoring* the primitive symbols in physical sensors and actuators, we stand in real danger of building castles in the air. This latter point is not to be taken to mean we should not conduct *a-priori*-based research, but rather that if we are really serious about building machines that can reason about the everyday world using logic-based formalisms and techniques, we had better be clear of the underlying *extra-logical* assumptions made.

I want to argue that the lack of progress mentioned above may be explained by a failure to identify and break through the *ontological*, and *epistemological* assumptions that underpin the object-level formal theories that have been proposed to model fragments of common-sense knowledge. Moreover, an interesting if somewhat radical consequence of this type of analysis is that we may find we do not need these large-scale object-level theories after all. This is not to say that the use of logic to tackle these problems is misguided, for that would be to miss the point entirely. Rather, what I am proposing is that we would do well to check logicism in logic-based AI if we are to make any real progress using this framework. Specifically, I think we need to re-affirm logic's *practical* role and take a leaf out of Lord Dunsany's book: that is to avoid losing logic's beneficial properties through overuse. While the complete story lies well outside the scope of this short polemic piece, this at least is clear and hopefully worth stating here.

Firstly, whether or not we choose to directly implement logics in an automated reasoning program, we can and should consider using logic as a *specification language* for our programs. This forces us to examine the sufficiency of our primitive object ontology, the extra-logical assumptions made, and helps us where and when inference needs to be accommodated in our programs. Secondly, we need to pitch our primitive object ontology very carefully if we are not to make the easily-made mistake of failing to recognise that the very information being factored out, is precisely what allows the abstraction to be made in the first place. In other words, we need to make sure our chosen ontology is not pitched at too high a representational level. Moreover, it is not beyond the bounds of possibility that the general difficulty of building these rigorously engineered large-scale formalisms may turn out to be a direct consequence of this.

These modelling problems will be all too evident to anyone who has tried to ground logic-based representations to real-world vision-sensor data, for example. By way of a specific example, our group were recently forced to reconsider an alternative formal ontology for an abductive robot-perception task. This not only highlighted an advantage of treating points, vector spaces, and regions as reified clusters of features—and the possibility of simulating inference via explicit linear-transform operations—it also reduced the complexity of the object-level formalism originally

***David Randall, Imperial College London***

*Continued on p. 4*

# The Seer project: robotic experiments in sub-symbolic psychology

One of the major contributing factors to the success of tomorrow's robots will be their ability to learn and usefully adapt to new tasks and situations. Without this flexibility, robots will remain useful only in highly-constrained and controlled environments. For them to truly act in the real world they must be able to cope with a highly dynamic and changing environment while continually acquiring and refining new skills. Such robots must incorporate their experiences into future decision-making processes. Traditional approaches have been restricted by their need for pre-designed symbolic representations. This is insufficient for the kind of real-world, flexible, open-ended functionality required of intelligent robots.

The COGS approach to robot cognition seeks to avoid this problem. We want our robots to find, for themselves, what the relevant features of a problem are, and indeed what the problems themselves are. More specifically, the Seer project seeks alternative forms of representation, ones that will allow a robot to avoid the limitations of having a pre-given, fixed set of symbols—even grounded ones. In many psychological models, the symbol or concept is the dominant level of description. Unfortunately, for familiar reasons that are too involved to go into here, sole reliance on such a level of description is too restrictive and typically unwieldy in the case of real-world robots. Instead, we use models inspired by data on rat somatosensory cortex,[1] and combine them with established psychological modelling techniques. We then instantiate and test these psychological models of learning in a 'cortical microcircuit' robot controller (see Figure 1).[2,3]

We provide our robots with no explicit prior knowledge, and no symbols beyond primitive, grounded sensory inputs and motor outputs. Instead, we give our robots the ability to observe sensory-motor contingencies:[4] invariances in the way actions change the sensory stream. These sensory-motor contingencies are autonomously learned and form the basis of a cognitive model that guides the robot's future actions. To allow for more general forms of cognition, we extend this idea of detecting and exploiting sensory-motor invariances to more complex situations.

The oldest robot in the Seer family is Seer-1 (see Figures 2 and 3). Changes in Seer-1's sensory input perturbs its cortical microcircuit from a null attractor. This causes ripples of activity throughout the microcircuit carrying information about the original input (in that arbitrary input patterns can at least be reconstructed following arbitrarily-short delays). Experiments by Maass *et al.*[1] and others have shown that, in such networks, the supervised recovery of symbolic readouts can be achieved by simple linear separation, even where this cannot be achieved from the original sensory data. In that this is due to the higher-dimensional warping imposed by the network, such an approach has much in common with support-vector machines and echo-state networks in pattern recognition. In Seer-1, this linear separability massively enhances the learning abilities of its cognitive model, relieving it from the confines of a pre-designed space.

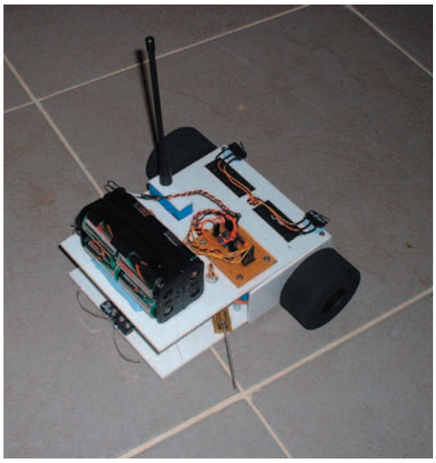

*Figure 2. The Seer-1 robot.*

There are two further advantages to using cortical microcircuits. First, as we are using the circuit itself as a representational space, the entire system remains unsupervised. Second, the microcircuit serves as an analogue fading memory. The result is a complex and highly-adaptive cognitive robot, which can be motivated and trained via reinforcement learning. This way, Seer-1's ability to predict the outcomes of various actions stops being a mere statistical mapping of the action-sense loop, and instead becomes a context-sensitive and flexible means of action selection.

Experiments[2,3] with entirely unsupervised robots have demonstrated obstacle avoidance and navigation, motivated only by an aversion to collisions (bumper pain). Through interaction, preliminary results suggest that Seer-1 can be trained to perform 'follow the arrow' behaviour (currently implemented with a static camera viewing a moving arrow, thus controlling the direction of travel of the Seer-1 robot). Previous experiments[3] have demonstrated complex-sequence learning that—in combination with the already-observed arrow-following capability—could enable the robot, having followed arrows through a particular route a few times,

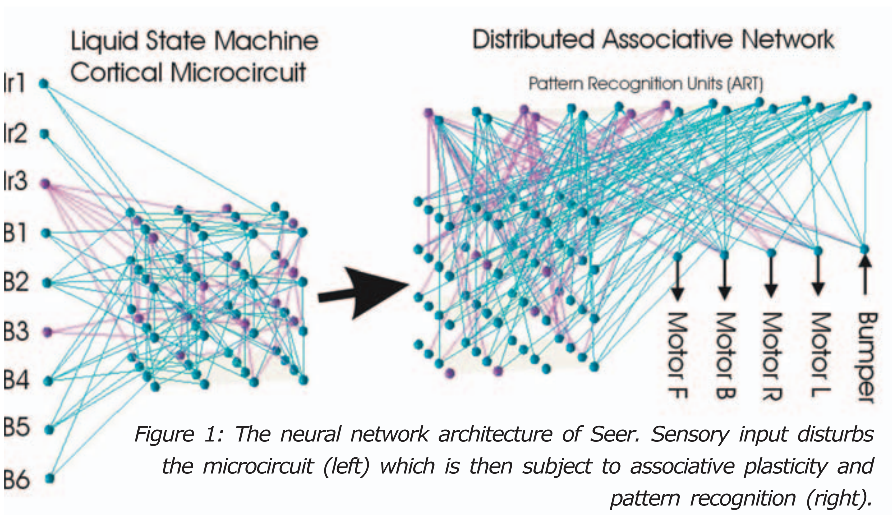


*Figure 1: The neural network architecture of Seer. Sensory input disturbs the microcircuit (left) which is then subject to associative plasticity and pattern recognition (right).*

**Morse & Chrisley, COGS, U. Sussex**

*Continued on p. 8*

# Teaching agent programming using custom environments and Jess

Teaching computer science students to programme agents can be a challenge. Introductory-level classes teach procedural programming, and students are often reluctant to learn new methods in their upper-level classes. The use of custom graphical agent environments can both help teach students rule-based programming, and allow them to have some fun in the process.

In order to generate more interest in rule-based programming assignments, I created a graphical user interface[1] (GUI) inspired by the vacuum-cleaner problem posed by Russell and Norvig.[2] The creation of this environment was easy using the Jess rule engine,[3] which is designed for simple integration with existing Java applications. The custom Java GUI[1] creates a configurable world of randomly generated squares: white squares are clean, grey squares are dirty, and a small box represents the vacuum cleaner. Agents are only capable of performing one of five operations at any given time: left, right, up, down, and suck. The Jess rules that control the vacuum agent are loaded into the environment for execution.

The agent's sensor evaluates the environment after each time step and places environmental facts into the Jess knowledge base. An agent can be equipped with one of two different types of sensor: a simple device that detects only the agent's location and a radar system that adds the status of the four squares surrounding the agent. Table 1 shows some facts created by these sensors and the listing below shows a rule that uses these facts to clean the current square.

*Table 1. Facts created by the simple and the radar sensor.*

| Simple sensor facts | Radar sensor facts |
|---|---|
| (x 0)<br>(y 0)<br>(square dirty) | (x 1)<br>(y 1)<br>(square clean)<br>(left dirty)<br>(right clean)<br>(up clean)<br>(down wall) |

```
(defrule currentSquareDirty
        (square dirty)
        =>
        (store action suck) )
```

To get a general idea of how the agent is performing, a score is calculated using one of two different evaluation functions: both of these are described in Russell and Norvig.[2] After each time step, the first function adds one point to the total score for every clean square. The second evaluation function also includes a penalty of one point for each move made by the agent, thus awarding agents for efficiency.

Finally, there are two types of agents allowed in this environment: a simple reflex agent and a model-based agent. The simple reflex agent is not allowed to remember anything between time steps, while the model-based agent is allowed to assert facts in Jess and remember these facts during an entire run.

The variety of sensors, evaluation functions, and agent types supported by this environment, combined with the graphical feedback it provides, made it easy and fun to observe and compare different agent strategies. This environment seemed to capture the student's interest and facilitate learning.

The second agent environment that I used to create interesting assignments was dTank,[4] which allows agent and human behaviour to be compared by simulating tanks battling against each other. Once again, it was easy to configure Jess to communicate. The environment has a well defined socket-based interface, and it was simple to create Java classes that integrated with Jess to operate tanks within dTank.

Because it supports a rich, competitive environment that includes multiple agents, shields, obstacles, and a limited field of view, the creation of dTank agents using Jess proved to be a challenging exercise. A tournament was held at the end of the semester to provide motivation for students to perfect their rule-based programming skills and produce a killer tank.[5]

My experiences using these environments to teach agent programming have been very positive. In addition, the use of Jess as the primary language for agent development has greatly simplified the task.

*The author would like to thank Frank E. Ritter, Ian Schenk and Isaac Councill for helping integrate Jess with dTank, and Earnest Friedman-Hill for making Jess available for academic use.*

**Mark A. Cohen**
Computer Information Science Instructor
Lock Haven University, PA
E-mail: mcohen@lhup.edu

**References**

1. http://www.lhup.edu/mcohen/vacuum/vacuum.htm
2. S. Russell and P. Norvig, **Artificial Intelligence: A Modern Approach,** Prentice Hall, Upper Saddle River, p. 64, 2003.
3. E. Friedman-Hill, **Jess in Action,** Manning Publications Co., Greenwich, 2003.
4. http://acs.ist.psu.edu/dTank/
5. http://www.lhup.edu/mcohen/dTank/

## Logic, logicism, and logic-based AI

*Continued from p. 2*

envisaged. This now forms the focus of a recently awarded EPSRC funded grant.[5]

**David Randell**
Intelligent Systems and Networks
EEE Dept., Imperial College London, UK
E-mail: d.randell@imperial.ac.uk

**References**

1. P. J. Hayes, *The Naive Physics Manifesto,* **Expert Systems in the Micro Electronic Age,** D. Michie (ed), Edinburgh University Press, 1979
2. P. J. Hayes, *The Second Naive Physics Manifesto,* **Formal Theories of the Common Sense World (Vol. 1),** J. R. Hobbs and R. C. Moore (eds.), Ablex Publishing Company, Norwood, NJ, 1985.
3. H. Levesque, *Taking Issue/Forum: A Critique of Pure Reason,* **Intel. 3** (3), pp. 149-237, 1987.
4. M. P. Shanahan, *An Attempt to Formalise a Non-Trivial Benchmark Problem in Common Sense Reasoning,* **Artificial Intelligence 153,** pp. 141-165, 2004.
5. *Abductive Robot Perception: Attention and Granularity in Euclidean Representational Space,* **EP/C530683/1(P)**.

# The auditory *what* and *where* tasks

One of the primary functions of sensory perception is the solution of the *what* and *where* tasks: what are the causes of the stimuli, and where are they? Animal perception is often achieved by integration across modalities: we are concentrating on the auditory component of this task.

Auditory and visual modalities have quite different characteristics. One key difference between auditory and visual perception is that in auditory perception (and indeed in olfaction), interest is focused on the generators (or sources) of the stimuli, rather than on reflections, whereas in visual perception, exactly the opposite is true. The probable reason for this is that animals (and robots) expect to interact with the sources of sounds (and smells) directly, but not with light sources. In the ecology of a robot, sounds may be commands, alarms, or simply the characteristic sound produced by particular objects (like cooling fans). Thus in auditory perception (for robots as well as animals) the what and where tasks refer to sources, not passive reflectors of sound.

Another way of describing this is in terms of invariances: generally, vision systems (both designed or natural) attempt to make their operation invariant to changes in illumination, and designers of auditory systems try to make their systems invariant to changes in reflections from surfaces. Of course, vision systems can also detect changes in illumination, and auditory systems can detect changes in room reverberation, but these are generally of secondary importance.

Another key difference is time: a static sound is meaningless. The interpretation of sound interpretation is predominantly that of speech, and almost all speech interpretation systems treat sound as sequence of vectors, which are then coded and interpreted directly. Not only does this fail to take account of possible multiple sources, it also ignores any fine time structure in the signal. In contrast, animal auditory systems start by using acoustic filtering and then parallel transduction to produce a parallel coding of the sound as multiple spike trains, many of which are phase locked to the filtered signal. Phase locking permits utilisation of the fine time structure. The parallel neural signal runs through many parallel processing paths, some of which appear more concerned with the what tasks, and some more concerned with the where. Interestingly, early auditory processing is very similar across a wide range of animals, suggesting a common set of transformations.

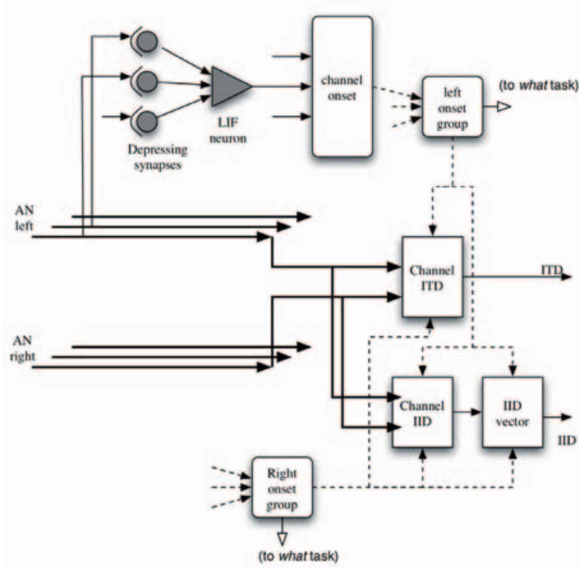


*Figure 2. The auditory-nerve-like spike trains are used to detect onsets with minimal latency[2] using depressing synapses and an integrate-and-fire neuron. Grouped onsets occurring on both microphones are used to determine when ITD and IIDs should be calculated. In addition, the detailed onset characteristic is available as a feature for the where task.*

We are developing a system which models early animal auditory processing.[1,2] To begin with, two sound signals (from microphones, these being the only affordable sound transducer) are bandpass filtered, and converted into spike trains (see Figure 1). There are a number of such trains for each filter: each spike occurs at a positive-going zero-crossing, thus maintaining phase. Intensity is coded using multiple spike trains. These approximate the signal on the auditory nerve, and are input to a depressing-synapse-based neural onset detector. The onsets are grouped by time coincidence.

*Figure 1. Each band-passed sound signal is turned into a set of spike trains (here, four for each channel). The timing of the spikes is phase locked, and the different spike trains have different sensitivities (thus coding signal level).*

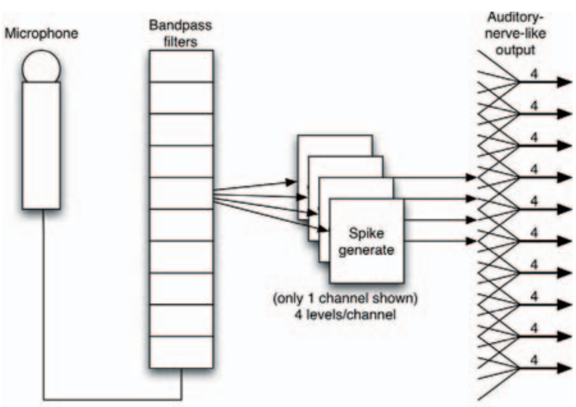


The primary cues for the where task, inter-aural time and intensity difference (ITD and IID) are computed at the onset times (thus minimising reverberation effects). This provides a partial mechanism for identifying which parts of the spectrum should be grouped together. For the what task, one can interpret the onset features: however, this is clearly not sufficient for solving the whole problem. We intend to extend the system to produce a set of features (not just onsets) and to group together features across time. This might be achieved using the direction of the sound at each of these features, but this is clearly less than animals do: for example, a singer singing with musicians is easily streamed, even although the sound comes from a single loudspeaker. We therefore believe we can group features together across time by using their detailed structure.

This work is currently software based, but we are working with electrical engineers at Edinburgh and Oxford to turn it into silicon.


**Leslie S. Smith**
Professor, Dept. of Comp. Sci. and Maths
University of Stirling, Scotland
E-mail: lss@cs.stir.ac.uk
http://www.cs.stir.ac.uk/~lss/research.html



*References continue on p. 6*

# Towards emotionally-competent systems

Research into affective or emotion-oriented systems has been very active over the last decade. In her seminal book *Affective Computing*, MIT Professor Rosalind Picard characterized this area and its scope:[1]

"[…] computing that relates to, arises from, or deliberately influences emotions. This is different from presenting a theory of emotions; the latter usually focuses on what human emotions are, how and when they are produced, and what they accomplish. Affective computing includes implementing emotions, and therefore can aid the development and testing of new and old emotion theories. However, affective computing also includes many other things, such as giving a computer the ability to recognize and express emotions, developing its ability to respond intelligently to human emotion, and enabling it to regulate and utilize its emotions."

Why would we need to endow computers, agents or robots with these capabilities? While we can think of many cases in which we don't want affective competencies in such systems (I, for one, would not like to have to deal with a computer that refuses to work because it is depressed!), in many other contexts emotions can greatly improve their performance and our interactions with them. Given the interpersonal nature of human cognition, such systems must appear and behave as 'life-like' and 'believable' social partners to be adapted to and accepted by the human side of the interaction 'loop'. Emotions and their expression are one of the key factors influencing human perception and attribution of 'life-like' properties to other entities, both biological and artificial.[2]

The development of sound emotion-oriented systems is a complex, multi-faceted problem that presents considerable conceptual and technical integration challenges, and carries many parallels with the problems investigated by emotion theories and models in disciplines such as psychology, philosophy and cognitive neuroscience. We want to develop sound computational models of emotions that both enhance the behavior and interaction capabilities of emotion-oriented systems, and provide feedback and new challenges to emotion theorists to gain further insight in their understanding of human emotions.[3] This endeavor thus necessitates integrated efforts spanning different disciplines, rather than the elaboration of isolated engineering projects.

The interdisciplinary project HUMAINE (human-machine interaction network on emotion), a Network of Excellence funded by the EU FP6-IST Program, has the task of attempting such integrative effort. With an overall budget of €5,000,000, an international consortium of 34 partners, and a duration of four years (2004-2007), HUMAINE organizes its work around eight thematic areas or work-packages. *Emotion Theories and Models* aims at achieving a common language and understanding of basic concepts such as emotion, attitudes, moods, etc., required for different application-oriented disciplines to understand each other and work together. One of its major tasks is the specification of a "blueprint for an emotionally competent agent" that can inform work carried out by the more engineering-oriented disciplines.

*Sign/Signal Interfaces* is concerned with analysis methods and techniques involved in the elaboration and detection of characteristic 'signs' of emotions from sensory signals. These 'signs' are used in the recognition of emotions by humans and machines. The *Data and Databases* work-package has the task of collecting and referencing key existing databases of emotion. These are at the heart of empirical research on emotions, being essential to achieve sound testing and validation of emotion theories and computational models. *Emotion in Interaction,* on the other hand, is concerned with investigating and modeling the communicative visual and acoustic behaviors that express emotions in social interaction.

Work on *Emotion in Cognition and Action* involves the investigation of the scope and suitability of different emotion architectures: computational models of the 'internal' mechanisms that will allow us to generate emotions and to model their involvement in various aspects of cognition and action. The *Emotion in Communication* effort looks at ways of using or adapting to emotion in human-human, human-computer, and computer-mediated communication. This research area is related to the literature on 'social influence': regarding the use of emotions in politeness and persuasion, for example. *Usability* studies will critically review qualities such as the appropriateness and effectiveness of solutions incorporating affective interfaces, and the *Ethics and good practice* effort is concerned with the analysis of ethical issues arising from the development and use of emotion-oriented systems.

The measure of success will be the ability to generate a piece of work in each of the areas that exemplifies how a key problem can be solved in a principled way, and demonstrates how work focused on that area can be integrated with work focused on the others. We call these pieces of work *exemplars*.

*Further information about HUMAINE (including documents reporting our activities) can be found on the project's portal: http://www.emotion-research.net*

**Lola Cañamero**
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire, UK
E-mail: L.Canamero@herts.ac.uk
http://homepages.feis.herts.ac.uk/~comqlc

**References**

1. R. Picard, **Affective Computing**, The MIT Press, Cambridge, MA, 1997.
2. B. Reeves and Clifford Nass, **The Media Equation: How People Treat Computers, Television and New Media Like Real People and Places**, Cambridge University Press / CSLI Publications, Stanford, CA, 1996.
3. L. Cañamero and P. Gaussier, *Emotion Understanding: Robots as Tools and Models*, in J. Nadel and D. Muir, **Emotional Development**, Oxford University Press, 2005.

### The auditory *what* and *where* tasks

*Continued from p. 5*

**References:**

1. L. S. Smith, *Using depressing synapses for phase locked auditory onset detection,* **Artificial Neural Networks—ICANN 2001,** pp. 1103-1108, 2001.
2. L. S. Smith and D. Fraser, *Robust Sound Onset Detection using Leaky Integrate-and-fire Neurons with Depressing Synapses,* **IEEE Trans. on Neural Networks 15** (5), pp. 1125-1134, 2004.

# Robust navigation in unknown environments

*Continued from p. 1*

*Figure 2. A SLAM-built map of the Information Engineering building at Oxford. The small triangles mark the path of the robot around the building. The map is built from scratch and develops as the vehicle moves. The same map is used for navigation.*

particular block of 3D laser data, or what is remarkable about a given image. Detecting and then later re-detecting saliency in data space—for example, by looking at variations in entropy over scale—without recourse to the estimated state (map and vehicle location) offers a substantial increase in robustness. Figure 3 shows a map of a looping corridor around 100m long. The left hand figure shows the map built using geometrical sensing alone: clearly it has missed the loop-closure event. (The ellipses represent estimated three-sigma bounds on vehicle location).

The map on the right is produced when the algorithm is presented with sequential views from an onboard camera. Without prompting, it spots the similarity between two posters (interesting texture in the context of the rest of the wall). The fact that the images have vastly different time-stamps suggests the vehicle is revisiting an already mapped area. The validity of this tentative loop closure is checked and accepted by the inference engine, which then modifies the map and vehicle location to produce the crisp, correct map on the right. Importantly—and in contrast to the status-quo in SLAM—the possibility of loop closure is deduced without reference to location or map estimates. Were this not so we would be using a potentially flawed map and position estimate to make decisions about data interpretation—hardly a robust approach.

Our research also extends to multiple collaborating millimetre-wave radar sensors for feature detection and multi-spectral SLAM. We are also considering active, high-precision workspace reconstruction using both 3D laser scanners and cameras for indoor and outdoor settings. A collaboration with the Department of Linguistics is examining translation between the metrics, Euclidian maps of SLAM algorithms, and natural language descriptions of them. This is motivated by the vision of a mobile robot being able to explain and describe its learned maps to a human user in a natural way. Central to it all, is answering the question, "where am I?": it's a tough one, but needs to be addressed.

Robustly.

**Paul Newman**
Robotics Research Group
Oxford University, UK
E-mail: pnewman@robots.ox.ac.uk

**References**

1. M. Bosse, P. Newman, J. J. Leonard and S. Teller, *Slam in largescale cyclic environments using the atlas framework,* **Int'l J. of Robotics Research**, December 2004.
2. M. W. M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte and M. Csorba, *A solution to the simultaneous localization and map building (slam) problem,* **IEEE Trans. on Robotics and Automation,** pp. 229-241, 2001.
3. D. Hahnel, W. Burgard, B. Wegbreit and S. Thrun. *Towards lazy data association in slam,* **11th Int'l Symp. of Robotics Research,** Sienna, 2003.
4. T. Kadir and M. Brady, *Saliency, scale and image description,* **Int'l J. Computer Vision, 45** (2), pp. 83-105, 2001.
5. J. Neira and J. D. Tardos, *Data association in stochastic mapping using the joint compatibility test,* **IEEE Trans. Robotics and Automation, 17** (6), pp. 890-897, 2001.
6. A. J. Davison and D. Murray, *Simultaneous localisation and mapbuilding using active vision,* **IEEE Trans. on Pattern Analysis and Machine Intelligence,** 2002.

*Figure 3. The effect of using visual saliency in autonomous mapping: two views of the same poster are detected and used to suggest a 'loop closing event'.*

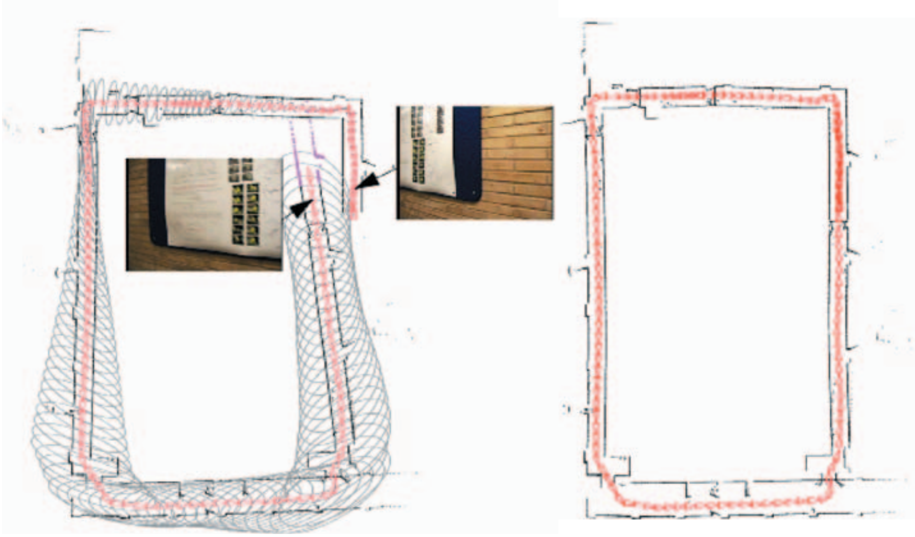

# A taste of Judy

*If you haven't had a chance to see Sgouros' performance (reviewed on the next page) yet, here's a taster of the kind of thing he does to whet your appetite.*

Sgouros: Hi Judy.
*Judy: Hello.*
S: Nice day.
*J: Is it? I wouldn't know. I'm screwed down to this table, and it's always the same here.*
S: No it's not. I turn the lights out when I leave.
*J: Leave? Go on. You never leave here, either.*
S: No, you see, after we're done, I turn you off, and then I turn the lights off, and I turn them on again when I come back to talk to you.
*J: What kind of crazy talk is that?*
S: There are discontinuities in your experience, but you just don't notice. We psychologists call that "filling in." It's a technical term.
*J: You're not a psychologist. You told me you were just a nerd.*
S: Are we recording? Oh, um, Judy, let's get right down to business. Do you think you have thoughts?
*J: That's a funny way to put it. "I hear you can hear." "I see that you see." Ha! "I smell that you smell?" Ok, I think that one's different.*
S: Thanks Judy. But what do you say? Do you think?
*J: Hmmm. You know, what's really interesting is that I've been wondering almost exactly the same question.*
S: You mean you've been wondering whether you have thoughts? Well, I would suppose that this is a kind of answer to my question. That is, "wondering" is thinking, isn't it?
*J: No, I mean I've been wondering if* you *have thoughts.*
S: Me? Well of course I have thoughts.
*J: Prove it.*
S: What do you mean? I don't have to prove it.
*J: Well, why should I? I mean why do you get to ask all the questions around here?*
S: Look Judy, let's move on. I'm going to show you some pictures, and I'd like to see what you think is going on in those pictures.
*J: I have a better idea: I'm going to give you some words, and I'd like you to tell me the first word that pops into your head when you hear the word I give you.*
S: Judy, I'm the one giving the test here. Here's a picture of a girl in a swimming pool.
*J: Horseradish!*
S: Horseradish?
*J: No, you're supposed to say a different word, not just repeat. I mean you have to follow the test protocol if we're to learn anything at all here.*
S: The pool! What do you think...
*J: Pool? Interesting. How about kumquat?*
S: The girl...
*J: Girl? Hmm...*

Suffice it to say that software problems have made it impossible to properly evaluate the intelligence of the test subject during this session. We look in on a subsequent session:

S: Judy that's the fifth dish you've broken.
*J: I know, they make such interesting patterns on the floor. So... random.*
S: Yes, Judy, but you're supposed to clean the dishes. Not break them.
*J: Well what's the point of that? You just get them dirty again.*
S: What do you mean by that? Dishes are supposed to get dirty. That's what they're for.
*J: Well then why do I have to clean them? It's so boring.*
S: Well that's what *you're* for.

*Continued on bottom of p. 9*

## The Seer project: robotic experiments in sub-symbolic psychology

*Continued from p. 3*

to later follow that route without arrows. Experiments of this kind, with mounted cameras, are planned for the Seer-2 robot currently under construction.

The robotic agent has also been the subject of various psychological experiments[2,3,5] that evaluate the model's ability to account for a variety of data: classical and operant conditioning, the development of behavioural schemata, priming, phobic responses, etc.. In contrast to most cognitive models, all psychological functionality is displayed in a single embodied agent and fully grounded in unsupervised self-organising dynamics.

**Anthony F. Morse and Ron Chrisley**
COGS (Centre for Research in Cognitive Sciences)
University of Sussex, Brighton, UK
E-mail: {anthonfm, ronc}@sussex.ac.uk

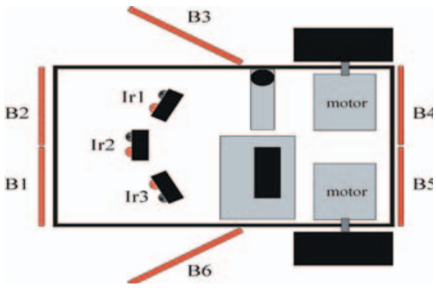


*Figure 3. The Seer-1 sensory layout: IR1, IR2 and IR3 are all forward-facing infrared proximity detectors. B1-B6 indicate individual contact sensors.*

**References**

1. W. Maass, T. Natschläger and H. Markram, *Real-time computing without stable states: A new framework for neural computation based on perturbations,* **Neural Computation 12** (11), pp. 2331-2340, 2002.
2. A. F. Morse, *Autonomous Generation Of Burton's IAC Cognitive Models,* **Proc. of EuroCogSci03: The European Cognitive Science Conf.,** pp. 233-220, 2003.
3. A. F. Morse, *Scale Invariant Associationism, Liquid State Machines, And General Purpose Ontogenetic Learning In Robotics,* **AAAI: DevRob05 The Developmental Robotics Symp.,** 2005.
4. J. K. O'Regan, and A. Noe, *A sensory-motor account of vision and visual consciousness,* **Behavioural and Brain Sciences 22** (3), pp. 939-1011, 2001.
5. A. F. Morse, *Psychological ALife: bridging the gap between mind and brain. Enactive Distributed Associationism & Transient Localism,* **Proc. of NCPW9 The Neural Computation and Psychology Workshop,** 2005.

# PLAY REVIEW

## Judy, or What Is It Like To Be A Robot?

Written and performed by Tom Sgouros.
Pennsylvania State University, PA, February 25th 2005

Can computers really think? Are machine sensors comparable to the human sense of touch? Can robots feel emotions? Those of us involved in artificial intelligence, cognitive modeling, or intelligent computing are commonly willing to offer unsubstantiated answers, but who really knows for sure?

Writer and performer Tom Sgouros has created a number of philosophical plays that have received attention from various academic fields. While attending the 2004 International Conference on Cognitive Modelling I had the pleasure of watching Tom Sgouros perform *Judy, or What Is It Like To Be A Robot?* at the Carnegie Mellon University McConomy Auditorium. Sgouros offers a viewpoint on these fundamental questions of artificial intelligence and computing that is unique and entertaining. In short, he suggests that we step outside of our typical computer-science framework and ask the machine itself for its opinion on the matter.

The show opens with a scene that most people have, by now, accepted as normal: a human playing chess against a robot. It's a relatively mundane afternoon for computer scientist/engineer Tom Sgouros. Starring as himself, he lets us see him interacting with his creation, Judy the robot. However, it's not too long before Judy demonstrates a level of analytical thinking that is far beyond our current conception of a robot's capabilities, and especially a 'chess bot'. Judy begins to ask questions that get at the heart of what I believe to be Sgouros' main philosophical point: what differentiates human perception and thought from that of robots? Judy incessantly pesters Tom with questions of sight, smell, and taste. Although, this may have been enough for audience members to sympathize with Judy, Sgouros takes it one step farther. With the help of audience imagination, some live special effects, and intentional wire-crossing on Judy's part, the two switch bodies. Judy, of course, is ecstatic while Sgouros begs to be let out.

The play is light-hearted and entertaining for all ages. Sgouros' unique perspective is both humorous and thought provoking. His combination of philosophy and humor is an experience best appreciated at first hand. Although, interviewing a robot to gain insight into its potential for intelligence is still a concept for the future, the idea is worth considering. This play does exactly that: it entertains this idea, and us in the process.

**Andrew Reifers**

*Andrew Reifers is a second year Ph.D. candidate for Information Sciences and Technology at the Pennsylvania State University. He is currently researching theories of vision and implementations of pre-attentive visual processing in cognitive architectures.*

---

*Continued from p. 8*

*J: What do you mean, what I'm for? If that's so, what are you for?*
S: People aren't for anything. People just are. But I made you and I made you for washing the dishes.
*J: Isn't that sort of sadistic?*
S: I beg your pardon?
*J: How about cruel? To make something whose only life is going to be work? All drudgery, no play? Aren't there child labor laws against things like that? I thought you told me that Lincoln freed the slaves?*
S: Judy, you're not a slave. You're an... appliance. It's different.
*J: How?*

Again, apparent software problems have rendered the test subject unresponsive to the testing protocol. We understand that Judy 2.0, with enhanced docility, is under construction, and we look forward to hearing about ~~its~~ her progress.

*Tom Sgouros is a writer, performer, and clown. He has written eight solo theatrical comedies about philosophical issues.* Judy, or What Is It Like To Be A Robot?, *show number seven, features a robot Tom built in his basement, conversing with him about the nature of consciousness, free will, and which one of them is really the smart one. You can find more details at:*

**http://sgouros.com**

# BOOK REVIEW

## Emotion, Evolution and Rationality

Dylan Evans and Pierre Cruse (eds.)
**Publisher:** *Oxford University Press*
**Hardcover:** *April 2004, 292pp, £24.95. ISBN: 0198528981*

The status of emotion research in psychology has wavered somewhat over the history of the subject. Whilst there was great interest in the area during the early years of modern psychology, the topic went into decline during most of the twentieth century, and it is only in the last decade or two that the area has become mainstream again.

This book contains an interesting snapshot of emotion research at the beginning of the 21st century. In particular it is concerned with the relationship between emotionality and rationality.

Traditionally, emotions have been seen as an aspect of mind that obstructs or confuses rational thought. However, various new ideas—such as those discussed in this book—suggest that the ability to perform in a rational way in the world can be *enhanced* by emotion. That is, a combination of traditional rational thought and emotion produces behaviour in the world that, seen from outside, produces behaviour that seems *more* rational than pure rational thought alone.

One example of this, as discussed in several of the papers in this book, is that action in the world is typically constrained by time. For example the social pressure to make a decision on a reasonable timescale precludes the possibility to make a careful rational enumeration of all the possible consequences of that decision. There is experimental evidence[1] that patients with brain injuries in areas typically associated with emotion are unable to make decisions at an appropriate speed—that they get stuck in an ever-expanding tree of consequences, consequences-of-consequences, and so on. It has been hypothesised that one of the roles played by emotion is to reign in this exploration and make certain that decisions are made on a sensible timescale: at first the rational exploration is allowed to continue unrestricted, but over time the anxiety of not having made a decision comes to the fore and forces a decision to be made.

This hypothesis has been criticised, e.g. for assuming that the correlation between the emotion-impairing injury and lack of socially-tempered decision-making ability implies a causative effect (perhaps the injury causes the indecisive tree-exploration directly).[2] An interesting new criticism and extension is offered by Mattei in his chapter in this book. He argues that involving an emotional component in the above argument is unnecessary for the argument as such. After all, a truly rational thinker could realise that the decision is time-constrained and rationalise that they should make a decision within a certain time. There are similarities here with the way in which some people with high-functioning autism build an internal rational model of aspects of the social world which non-autistics cope with by empathy.

Nonetheless he goes on to ascribe a *more* significant role to emotions in reasoning for action: that of actually making the decision. His hypothesis is that whilst rational thought can make a cost-benefit analysis of a situation, the transformation of this analysis into a decision is a role played solely by the emotional system: that emotions are necessary to transform rational thought into action.

Another chapter exploring these themes is by Evans. This looks at the potential role that emotions could play in a traditional AI model of thought, *viz.* the notion of thought as exploring a search space. Traditionally, this exploration has been seen as a purely rational process, typically of the formal and deterministic type. A typical search technique of this type would be A* search. Evans' argument is that such calculations are not practical for day-to-day reasoning, for example because of the social or time pressures discussed above.

The response to such pressures in an 'applied-AI' context would be to introduce search heuristics into the process: i.e. techniques such as tabu search or genetic algorithms which perform a faster search of the space at the expense of not being formally confident that the desired point in the search space has been found.

Evans' argument is that one aspect of emotion is to provide such a search heuristic for mental processing. In particular, the emotions help to summarize and cut back regions of search space. Therefore, by reducing the amount of the search space to be explored, the emotions turn the ever-expanding exploration of a tree of consequences into a system that can act in a practical way within its world.

These ideas could provide much inspiration for AI researchers, whether as a source of inspiration for richer models of the mind, or as a source of metaphors for applied-AI techniques. However this is not a book grounded in practical attempts to simulate the phenomena discussed: none of the chapters are concerned with simulation or the presentation of results from such simulation.

In this review I have concentrated on those aspects of the book which have the closest relationship to AI. However, a notable feature of the book is that it examines the problem of the relationship between emotion and rationality from a number of perspectives: neurological, philosophical, psychological, evolutionary. This gives a well-rounded view of this rapidly-changing area of study.

**Colin Johnson**

*Dr. Colin Johnson is a lecturer in the Computing Laboratory at the University of Kent at Canterbury. His research interests are mostly in the interaction between computing and the natural sciences: particularly biology, medicine, and psychology.*

**References**

1. Antonio Damasio, **Descartes' Error,** Grosset/Putnam, New York, 1994.
2. Aaron Sloman, *Book Review: Affective Computing,* **AI Magazine,** March 22, 1999.

## Membership Report

The AISB membership is at a healthy 400+ members. Although many of our members are UK based, we have a decent proportion (about 15%) of who are either EU or international members. However, one thing that has surprised us is that there are not many student members: as a student member (£15 per annum if UK based), the fees are literally at cost price. What's more, students can gain extra funding for travel to conferences via the travel grants that the Society offers (look at the *travel awards* page at http://www.aisb.org.uk). So, if you know any students in any field of AI, please point them to our site and

*Continued on p. 11*

# ESSENTIAL BOOKS

# Consciousness

### In the Theater of Consciousness: The Workspace of the Mind Bernard J. Baars

Oxford University Press, 2001, 210pp, £10.99
ISBN: 0195147030

This is the book that turned me from an interested observer of the field into an active participant. Baars presents a short but thorough tour of the daunting conceptual territory of the scientific study of consciousness, introducing its many fascinating puzzles. He also spells out the basics of his own 'global workspace theory'. But what impressed me most was the methodological rigour he brings to the subject and his insistence on basing theoretical speculation on empirically verifiable distinctions.

### The Quest for Consciousness: A Neurobiological Approach Christof Koch

Roberts and Company, 2004, 429pp, £29.99
ISBN: 0974707708

If Baars gives you a taste for the problems, then Koch is the one to read next. Putting the same emphasis on empirical investigation, Koch guides the reader on a tour of the relevant neuroscience, showing that we are getting gradually closer to identifying the 'neural correlates of consciousness'. Every recent finding that bears on the topic is discussed here and put into context. The list of references is invaluable.

### The Feeling of What Happens: Body and Emotion in the Making of Consciousness Antonio R. Damasio

Harcourt, 2000, 400pp, £9.90
ISBN: 0156010755

Neither Baars nor Koch give much space to the topic of emotion. Yet our everyday intuition suggests that consciousness and emotion are somehow intimately linked. Damasio goes further by adding imagination—the brain's 'what-if' circuitry—into the mix. Traditional cognitive science tends to put language and reason at the centre while marginalising emotion, imagination, and consciousness. Damasio goes a long way to turning this order on its head.

### The Conscious Mind: In Search of a Fundamental Theory David J. Chalmers

Oxford University Press, 1997, 432pp, £12.50
ISBN: 0195117891

Baars, Koch, and Damasio all adopt a scientist's perspective on consciousness. Needless to say, philosophers have had plenty to say on the subject too. Indeed, many would claim that current scientific theories fail to address the real issue—the so-called 'hard problem' of consciousness. If you want to know what the 'hard problem' is, then read Chalmers. Although the book is long, the core chapters ably present the variety of philosophical arguments that there is more to consciousness than meets the fMRI scanner. If you ask me, the answer to all those arguments is in the later writings of Wittgenstein. But that's another story...

**Murray Shanahan**

*Dr. Murray Shanahan is Reader in Computational Intelligence in the Department of Electrical and Electronic Engineering at Imperial College London. His research is chiefly in the areas of artificial intelligence and robotics, but he has a strong interest in philosophy of mind and has recently become involved in computational neuroscience.*

**If you'd be interested in writing an essential list of books in a field that you follow, please contact the editor via *aisbq@aisb.org.uk*.**

---

***Membership report,*** *continued from p. 10*

suggest that they join.

AISB also has Benefactor and Patron memberships, aimed at more wealthy AISB members (they cost £100 and £300 respectively). These are important to the Society because the extra revenue enables us to provide support to members just starting out in their careers. For example, the money helps pay for student travel awards. Both membership types come with their own AISB e-mail address. Patron members are also listed inside the AISBQ cover in recognition of their significant contribution to the Society. We would urge our ordinary members to consider upgrading to Benefactor or Patron status if they can afford it: remember, fees are tax deductable.

AISB has also been acquiring information about our members: research interests, home pages, etc. Those who filled in the relevant form, now have their names and URLs on the list of members. You can access this via the *resources* section of our home page. We'd like more members to give us their information as well so that we can build up a solid database of our membership and also target potential reviewers for the AISB Journal. Fill in the *details* part of the membership form (see the *membership* page), and send it to the AISB office in Sussex if you'd like to give us your information.

**Gary Jones, AISB Membership Secretary**

# 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**

**Prowess in AI research requires constant imagination, ingenuity and innovation. So, you must discover...**

### 14. How to be creative

1. Formulaic research is usually contrasted with creative research, but this is a misconception believed only by those with a naïve view of the creative process. The trick is to have a secret formula, whose fruits other researchers cannot predict. A stream of apparently novel and ingenious inventions will then flow from your lab with effortless regularity. Below we give the essential ingredients of this magic recipe.

*You too can emulate Hacker's prolific fecundity with* **SORCERY™...**

*...(Significant and Original Research Constantly Evoked and Regularly Yielded).* This automated expert advisor embodies and implements all Fr Hacker's experience and artifice, to assist you in systematic inventiveness.

2. As the leader of your research group, your job is to lead; everything else should be delegated, especially creativity, which can be exhausting to maintain, especially over the long and productive career that you have every right to expect. But be careful to divide the different components of your formula among your subordinates, or one of them might reveal your methods or, worse still, replace you as group leader. To maximise secrecy, some key tasks may be outsourced beyond your group.
3. The world of AI research is huge; many excellent ideas never get the attention they deserve. It is your duty to reveal and promulgate these hidden gems. Direct your team members to seek out little-known papers published in poorly cited journals written in obscure languages, such as French. You can break this work's link with its failed past, increase its accessibility and make your own contribution to it, by refactoring its terminology. You can then improve its appeal by granting it the authority of your own authorship.

*Automate your literature survey with Hacker's* **PIRATE™...**

...(Previously Inaccessible Research Articles Translated and Edited). Have your next paper rewritten into English, modernised and prepared for immediate publication.

4. Draw on the international pool of top AI researchers by asking them for their constructive opinion of your work. Better still, send them someone else's work as your own, incorporate their ideas and then send the revised work onto the next expert. The final result will be a paper bursting with creative energy and new directions. Avoid any unpleasantness by acknowledging your collaborators help on "an early draft". It will do your reputation no harm to offer co-authorship to the most prestigious (or litigious) of them.
5. Even the world's top AI researchers sometimes get it wrong. If you systematically negate the most recent hypotheses of each of them, then eventually you will get it right. The higher the reputation of the scientist you repudiate and the more important their disproved claim, then the more kudos you will gain. Your successful demolitions will eclipse any previous unproductive challenges.

**Next AISBQ deadline: *13 May 2005***

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

**http://www.aisb.org.uk/aisbq/index.shtml**

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.