AISB Quarterly No.122

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

In this issue:

## Features



# A fully-automated system for facial muscle action

Machine understanding of facial expressions could revolutionize human-computer interaction technologies and fields as diverse as security, behavioural science, medicine, and education. In security contexts, facial expressions play a crucial role in establishing or detracting from credibility. In medicine, facial expressions are the direct means of identifying when specific mental processes are occurring. In education, pupils' facial expressions inform the teacher of the need to adjust the instructional message. The human ability to read social signals and emotions from someone else's facial expressions is the basis of facial affect processing. Thus, understanding this ability may lead to expanding interfaces with emotional communication and, in turn, to obtaining a more flexible, adaptable, and natural interaction between humans and computers/robots/machines.[1]

Most systems for automatic analysis of facial expressions attempt to recognize a small set of prototypic facial expressions of basic emotions, such as happiness and anger.[1] In everyday life, however, prototypic expressions of emotions occur infrequently. Typically, facial expressions convey conversational signals and signs of attitudinal states—such as interest and boredom—that are usually displayed as one or few facial muscle actions, such as raising the eyebrows in disbelief.

Instead of classifying facial expressions into a few basic emotion categories, we attempt to measure a large range of facial behaviour by recognizing facial muscle action units (AUs) that produce expressions. As described in the FACS,[2] all visually-distinguishable facial activity can be described in terms of 44 AUs. Examples include the inner brow raise (AU1), outer brow raise (AU2), brows pulling together (AU4), and so forth. Thus, if a computer system were able to detect these 44 AUs automatically, it could identify each and every facial expression that the human face can possibly display.

Figure 1 outlines our system for recognition of 27

***Maja Pantic***
***Delft University of Technology***
*Continued on p. 8*

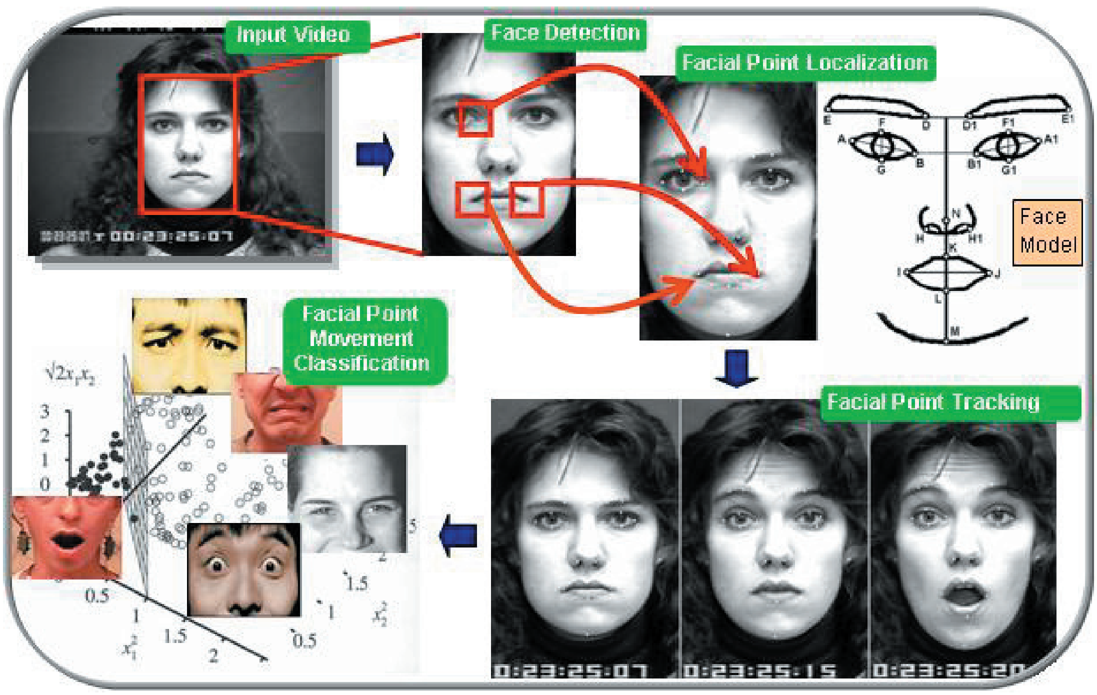


***Figure 1.*** Automatic detection of facial muscle actions (AUs).

# A test for irrational thinking

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

ISSN 0268-4179



One of the fundamental concepts of decision-making is that of utility and its expected value (EU). Theory concerning utility and rational decision-making in conditions of uncertainty has been largely developed by von Neumann and Morgenstern,[1] Savage,[2] and Anscombe and Aumann.[3] Despite differences in the treatment of uncertainty (i.e. objective or subjective), the concept that an agent's preferences will reflect utility is widely accepted, and the *maximum expected utility* principle (maxEU) is commonly used.

According to this principle, the choice between several alternative actions $Z = \{z_1,\ldots,z_n\}$ is the $z_i$ whose $E\{U_i\}$ (expected utility) is the largest. Thus, the principle ensures that an agent makes the optimal choice with respect to its immediate knowledge of the world. The maxEU principle has been successfully used in different areas of science such as Markov decision processes, dynamic programming, and reinforcement learning. It is considered by many to be the foundation of *rational* behaviour.

However, critics immediately questioned the fully-rational approach by noting that some degree of irrationality may be required in certain situations (e.g. the rational donkey paradox when the unfortunate is placed between two identical haystacks). Moreover, more fundamental flaws have been spotted by Allais[4] and Ellsberg,[5] who suggested examples where the maxEU principle failed to predict preferences expressed by people. This has been confirmed in numerous experiments.[6]

The ACT-R cognitive architecture,[7] which is used widely to model psychological data, also uses the maxEU principle in its subsymbolic mechanism for conflict resolution between alternative production rules (i.e. decisions). However, the expected utilities in ACT-R are corrupted by noise of small variance. Thus, ACT-R is not entirely rational. This feature has allowed many researchers to model the behaviour of human and animal subjects in various experiments. Our work with the ACT-R architecture has been to investigate the role of noise and its impact on the decision-making process.[8] It has been suggested that dynamic control of noise variance (i.e. how much the EUs are corrupted) can both dramatically increase the adaptation and performance of agents and make the models produce a better fit with the data. Elsewhere, the author has suggested two ways of controlling noise variance: by use of entropy and through variances in probability distributions.

In the more recent work, the rational choice (i.e. the maxEU principle) has been compared with 'irrational' strategies using simple agents placed in the same environments. The agents used Bayesian learning of transitional probabilities between states in the world, actions, and their utilities (i.e. learning the Markov decision models). The only difference between these agents was the way they used the transitional probability models to make choice: either by using the maxEU principle or by making choices randomly from these distributions (i.e. in a Monte-Carlo fashion). In these experiments, agents were placed in a simple world where they tried to collect as many rewards as possible, and the number of rewards collected was compared.

A few words have to be said about the way rewards appeared in the environments. There was no definite location in the world containing rewards. Instead, they appeared stochastically in different places, with some having higher chance of a reward than others. Thus, the environments are stochastic. Different patterns of reward appearance were used, including completely random (i.e. uniform), and we could also control the frequency at which rewards could re-occur.

The results of these tests showed that random choice was not only as good as the maxEU, but often the irrational agents outperformed the rational ones by almost two to one. Figure 1 compares the percentage of rewards collected by different agents as function of rewards frequency in a stochastic world with some regular pattern of rewards. One can see that as the number of rewards in the world increases (due to higher re-occurrence frequency), the percentage of rewards collected decreases. However, the maxEU agent performs worse than two random agents, especially at a moderate rate of rewards.

This result can be explained as follows. First, recall that expected value is the first moment characterising the probability distribution

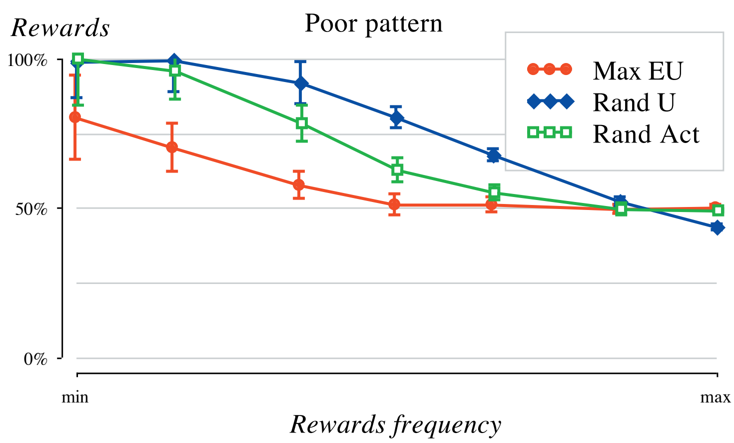


**Figure 1.** Shown is the proportion of rewards collected by agents in a stochastic environment with a regular pattern of rewards as function of rewards frequency.

***Roman Belavkin***
***Middlesex University***

*Continued on p. 8*

# Brain-machine interfaces: a novel paradigm for the study of learning and adaptation

In recent years, there has been a dramatic increase in brain-machine interface (BMI) research at institutions around the world. The idea is that users can perceive sensory information and enact voluntary motor intentions through a direct interface between their brains and an artificial actuator. Through this mediated process, users can see, walk or grab an object in virtually the same way as they would with their own, natural limbs. The rationale behind proficient brain-control is predicated on strong coupling between the brain and the machine: coupling that can be achieved through training with some combination of visual, tactile, and auditory feedback. As a result of long-term use of the BMI, the brain will adapt to the artificial actuator by incorporating its dynamic and physical properties into a somatosensory representation. This is equivalent of saying that the brain can incorporate an artificial device as an extension of its own body.

There have been several proof-of-concept studies in this area. These have demonstrated that different species—including rodents,[1] primates[2-5] and humans[6-8]—could use a BMI to control artificial devices such as cursors and robots. The main difference between the human and animal studies has been that the former mainly involve non-invasive interfaces: mainly using electroencephalographic and electrocorticographic techniques. However, these systems have a relatively low bandwidth—in terms of communicating and decoding neuronal signals—than more invasive techniques such as multi-electrode arrays. Putting aside the debate between invasive vs. non-invasive BMIs, the primary motivation for BMI research has been clinical: the development of neuroprostheses for restoring sensory and motor function after traumatic lesion of the central nervous system (CNS). If the challenge is met, BMI technology will dramatically improve the quality of life for people with sensory-motor disabilities.

However, the potential of BMI research goes beyond neuroprosthetics. The combination of the BMI paradigm with the multi-electrode technique provides a powerful tool for modern systems neuroscientists to study learning and adaptation in the brain. This technique allows simultaneous extracellular recordings from many neurons in different areas of the brain to be made in awake, behaving animals over long periods of time ranging from months to years. The main advantage of using this technique within the BMI context is that it allows neural circuit function to be visualized while subjects perform behavioural tasks in both manual and brain control modes of operation. This means we can compare the response, throughout training, of large neuronal ensembles when the natural limb is being controlled manually, and when the brain is controlling the artificial device via the interface.

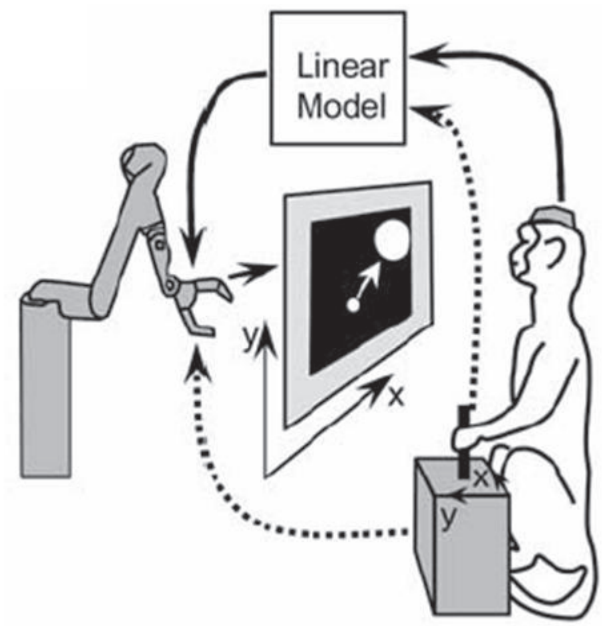


*Figure 1. Experimental design for a primate brain-machine interface (BMI). The monkey is seated in front of a computer monitor where visual stimuli are shown. The monkey must pursue a visual target (large circle) with a cursor (small circle) by moving a hand-held pole (manual control, dotted line). The pole actually controlled a robotic arm invisible to the monkey, and the cursor position on the screen reflected the robot's position. A linear model was trained to predict hand/robot velocity from neuronal ensemble activity recorded from the monkey's cortex. Then, the pole was disconnected, and the robot was directly controlled by the model's output (brain control, solid line). [Extracted from Reference 9].*

Among the different systems within the CNS, the motor system has been targeted by most of the groups doing BMI research. In addition to its direct implications to neuroprosthetics, the strong causal relationship between neuronal activity and the forces that generate movement makes the motor system ideal for the study of learning and adaptation. In a recent study,[4] we demonstrated that monkeys implanted with arrays of microelectrodes in fronto-parietal cortical areas can learn to reach and grasp virtual objects. They did this by controlling a robot arm through a BMI using visual feedback, even in the absence of overt arm movements (see the experimental setup illustrated in Figure 1).

Recording from hundreds of neurons simultaneously provided enough resolution for the plastic changes in the neuronal population to be quantified. Specifically, learning to operate the BMI was paralleled by functional reorganization in multiple cortical areas, suggesting that the dynamic properties of the BMI were incorporated into motor and sensory cortical representations. Further analysis[9] showed how the activity of individual neurons and neuronal populations became less representative of the animal's hand movements while favouring the movements of the actuator. These results demonstrate that during BMI control, cortical ensembles represent behaviourally-significant motor parameters, even if these are not associated with movements of the animal's own limb.

Despite these exciting results, BMI research is still at an embryonic stage. Further work is needed to characterize the dynamics of cortical plasticity associated with BMI control and to elucidate the


***Jose M. Carmena***
***University of California at Berkeley***


*Continued on p. 9*

# Decentralised decision making in ants

Drawing inspiration from nature to tackle engineering problems has a long and fruitful history, pre-dating even Holland's[4] pioneering work with genetic algorithms. In more recent years there has been increasing interest in deriving algorithms from social insects that would be suitable for application to both benchmark computer-science and real-world optimisation problems. The most well known and ubiquitous of these algorithms is often referred to as ant colony optimisation (ACO). This was inspired by the foraging behaviour exhibited by many species of ant, in which the laying and following of pheromone trails enables the colony to exploit available food sources efficiently. It has been applied to such shortest-path problems as the Travelling-Salesman Problem, and routing optimisation in telecommunication networks.[1]

Recently, another decision-making system has been described in the ant *Temnothorax albipennis*, which must emigrate periodically from its current nest to one of several potential new sites, each of differing quality. In this case, the colony is faced with the classic dilemma of balancing speed against the accuracy of decision making. In some circumstances, such as the original nest having been destroyed and the presence of predators nearby, it may be better to make a quick—but relatively inaccurate—decision. In other circumstances, such as the original nest being semi-intact and no predators being present, it may be better to take longer over the decision but be sure to choose the best available alternative site. The ability of *T. albipennis* colonies to make exactly this kind of compromise has been shown by both Franks *et al.*[3] and Dornhaus *et al.*[2]

The key to this flexibility is in the 'quorum threshold' used by the ants when they emigrate. Having discovered and evaluated a potential site, an ant scout will delay in inverse proportion to the quality it assesses the site to have. Ant scouts then use a two-stage recruitment process during emigration. First, they recruit other individuals by a slow process in which the recruit is led to a potential site, and made to learn the route. When the number of ants found in the potential site exceeds the quorum threshold, however, scouts change to bring in nestmates by carrying them: this is much faster, but it means the new recruits don't learn the route. By increasing this quorum threshold from one, the colony is able to move from quick-but-inaccurate individual decision making, to taking slower collective decisions that are more accurate because they allow more quality assessments of a site to be taken into account.

The flexibility of this strategy in exploiting the trade-off between speed and accuracy makes it an interesting candidate use in engineering tasks as a kind of anytime algorithm. However, greater understanding of its characteristics are required. We have examined the nature of this trade-off, and found it to be dependent on both the noise and the time cost inherent in assessing the available alternatives.[5] Without these features, the algorithm could be easily parameterised to collapse a best-of-many decision problem into a best-of-one decision problem (see Figure 1).

The importance of noise and the cost when assessing alternatives give useful information about the kind of applications the new algorithm may be suitable for. Future work will involve the investigation of such applications, as well as the relationship between the new algorithm and existing ant-inspired approaches such as ACO.

*I am grateful for the collaboration of Tim Kovacs, Nigel Franks and Anna Dornhaus on the work described here, which was supported by EPSRC grant number GR/S78674/01.*

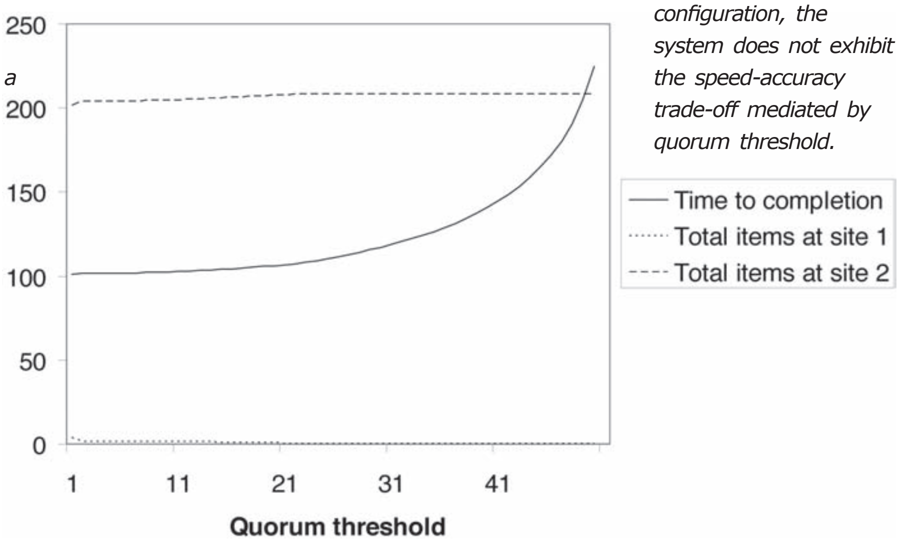


***Figure 1.*** *Output from a mathematical model*[6] *of emigration behaviour in* T. albipennis, *in which the increased willingness of scouts to switch from an inferior to superior nest site results in a perfectly-accurate decision at quorum threshold 1. Under this configuration, the system does not exhibit the speed-accuracy trade-off mediated by quorum threshold.*

**James Marshall**
Department of Computer Science
University of Bristol
E-mail: marshall@cs.bris.ac.uk

**References**

1. E. Bonabeau, M. Dorigo and G. Theraulaz, **Swarm Intelligence,** Oxford University Press, New York, 1999.
2. A. Dornhaus, N. R. Franks, R. M. Hawkins and H. N. S. Shere, *Ants move to improve: colonies of Leptothorax albipennis emigrate whenever they find a superior nest site,* **Anim. Behav. 67,** pp. 959-963, May 2004.
3. N. R. Franks, A. Dornhaus, J. P. Fitzsimmons and M. Stevens, *Speed vs. accuracy in collective decision making,* **Proc. R. Soc. Lond. B 270,** pp. 2457-2463, 7 October 2003.
4. J. H. Holland, **Adaptation in Natural and Artificial Systems**, University of Michigan Press, Michigan, 1975.
5. J. A. R. Marshall, A. Dornhaus, N. R. Franks and T. Kovacs, *Noise, cost and speed-accuracy trade-offs: decision making in a decentralised system,* **J. R. Soc. Interface,** in press.
6. S. C. Pratt, E. B. Mallon, D. J. T. Sumpter and N. R. Franks, *Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant Leptothorax albipennis,* **Behav. Ecol. Sociobiol. 52,** pp. 117-127, 9 May 2002.

# Creating an artificial rocky-shore community

The rocky shore is an extremely harsh environment, and yet has ecological communities that are perhaps the most studied and best understood on the planet. The ease of access to these communities, the sedentary or sessile nature of the organisms present, and the two-dimensional structure of the habitat make the study of these communities relatively easy. They are, therefore, important communities both for research and the teaching of ecological theory.

Artificial-life-based models of rocky shores, which model behaviours of individual agents and investigate emergent properties of their interactions, are extremely rare. This is surprising when one considers the wealth of studies on insect communities, which are inherently more difficult to study, and the historical impact of rocky shores in shaping ecological theory in general.

The ecology of this type of environment has been at the forefront of manipulative experimentation for several decades. However, this scientific strategy can take a long time to achieve results, often in the order of several months, and not all factors can be manipulated easily. As with any community, it is also difficult to manipulate the behaviours or physiology of the organisms directly. It is in these areas that artificial life simulations can help to test theories that would otherwise be very difficult, or indeed impossible, to verify on the shore.

Aggregation between individuals of small snail species that live high on the shore (in the family *Littorinidae*) is known to protect the animals from desiccation (see Figure 1). They move to feed, leaving a mucus trail on the rock surface, and then become inactive, often in groups in cracks and crevices. How the snails organise themselves into these aggregations during retreating tides from a dispersed feeding pattern is unknown, despite a range of manipulative experiments designed to address the problem.[1] Our recent artificial life simulations were designed to address this behaviour.[2]

We originally modelled a simulated movement pattern of each snail and provided a time-dependent decision-making step in which the snails might stop moving if they encountered other individuals. No part of the simulation allowed any individual snail to know the location of the other snails. The simulation did allow aggregations to form, but there were fewer and smaller aggregations formed than in equivalent laboratory experiments. To address this we added a simple location mechanism, where a snail could follow the previously-laid mucus trail of another snail. With trail following in place, the aggregations closely matched the patterns found in real snails. This study shows the potential ability of the snails to self-organise into aggregations, and illustrates the importance of information in the mucus trails, which appears to be vital to establishing these aggregations.

***Figure 1.*** *Littorinidae aggregating in a crevice.*


***Figure 2.*** *A simulation of snail aggregation on a realistic rocky shore. Light grey areas indicate the presence of crevices. Snails move into the grazing area to feed, then retreat down the shore as the tide goes out. The simulation is still under development to include the persistence of mucus trails between tidal cycles and interactions with primary producers and other species.*

The value of this study, outside of the immediate concern of rocky shore scientists, may seem trivial. After all, self-organising patterns are common, well understood, and have already been put to a variety of technological uses. However, we believe that artificial life simulations are an important tool for understanding rocky shore ecosystems and, indeed, to contributing towards general ecological theory.

We are currently looking at simulating the aggregation patterns of snails that may occur on real shores (see Figure 2). A simulation of such a shore allows many important environmental variables—such as the structural complexity of the shore in the form of crevices and the spatial pattern of food supply—to be manipulated. By incorporating a 'fitness' function into each simulated individual, the long term effects of structural complexity and food supply can be assessed rapidly. Importantly, although it will take a much longer time period, these simulations can be validated by manipulative experiments on real rocky shores. If successfully validated, the simulations

***Stafford, Davies, and Williams,***
***U. Newcastle upon Tyne/Hong Kong***

*Continued on p. 9*

# Artificial Life music: an adaptive distributed-agent systems approach

AI research has produced some excellent techniques for designing software that can 'compose' music of convincing quality automatically: Cope's work[1] is a good example. Such systems are only good, however, at mimicking well-known musical styles. They are either hard-wired to produce music in a given style or able to learn it by looking at patterns in training examples. Conversely, the issue of whether computers can *create* new styles of music is much harder to address. Here, the computer should neither be embedded with particular musical knowledge at the outset, nor learn from training examples.

One approach is to associate abstract models or mathematical formulae with compositional processes. Xenakis' work[2] is often cited as the main inspiration for composers working in this way. In fact, this approach to musical composition goes back to the 1960s. A large number of pieces have been composed since then using a variety of abstract models and/or mathematical formulae such as combinatorics, stochastics, chaos, and fractals.

Artificial Life (or A-life) has recently emerged as a promising means for the composition of computer music. The concept of 'A-life' is taken here in its broader sense to refer to algorithms that display some form of emergent behaviour resembling natural and/or biological phenomena. A growing number of composers have been using A-life algorithms, such as L-systems, cellular automata and genetic algorithms. Although such algorithms can sometimes produce interesting musical results, the Computer Music Research Group at Plymouth University is focusing on a slightly different approach: we are interested in modelling aspects of musical creativity using adaptive distributed-agent systems. Our hypothesis is that if we furnish the agents with proper cognitive and physical abilities, combined with appropriate interaction dynamics and adequate environmental conditions, then they should be able to evolve realistic musical cultures. Or, paraphrasing Casti,[3] 'would-be' music.

As a preliminary proof of concept, we implemented a model with a small community of agents furnished with a vocal tract, an auditory apparatus and basic cognitive skills. Together, the agents evolved from scratch a shared repertoire of short tunes (3 to 6 notes long) after a period of spontaneous creation, imitation, adjustment and memory reinforcement.[4]

Currently, a team of post-graduate students is taking this work forward. Marcelo Gimenes is developing a model inspired by Dawkins' notion of memes as basic units of cultural transmission.[5] Gimenes' agents extract musical memes from given examples and interact by generating short musical sequences to each other. They operate in two stages: learning and production. In the learning stage, the agents are trained with existing pieces of music in order to develop a memory of basic musical memes—each of which is tagged with a fitness measure—and meme sequencing rules. In the production stage, agents learn from each other's 'compositions' and generate new memes (using a genetic-algorithm-like procedure) and new meme sequencing rules. The 'biological' substrate for representing these memes more realistically is being addressed by Joao Martins.

Martins is studying the evolution of rhythms by furnishing the agents with categorisation and memorisation mechanisms using various neural network architectures. He is addressing two main issues: the definition of the 'biological' substrate of the agents, and the interaction dynamics that will enable learning and consequent evolution of rhythms. Where Martins is looking at imitation as the main form of interaction, Qijun Zhang is studying the potential of other forms of interaction, particularly those related to games.[6]

From a slightly different angle, Eduardo Coutinho is looking at the possibility of programming the agents to associate emotions with musical preferences. His hypothesis is that emotions are driven by a process of 'homeostatic regulation'. He is furnishing his agents with a virtual physical body (with metaphorical variables referring to levels of 'adrenaline', 'blood sugar', 'endorphin', etc.) in order to allow their behaviour to be affected by physical interactions with the environment. Vadim Tikhanoff has initiated preliminary work towards the robotic embodiment of these agents.[7]

As these projects are less than 12 months old, results are clearly not mature enough for publication. Nevertheless, the results from our preliminary proof-of-concept implementation[4] indicate that the use of adaptive distributed agents to simulate the evolution of music in surrogate worlds is a promising approach to the development of intelligent music systems. Interactive intelligent systems with the ability to evolve their own musical culture autonomously should soon become a tangible reality.

**Eduardo R Miranda**
Computer Music Research Group
University of Plymouth, UK
E-mail: eduardo.miranda@plymouth.ac.uk
http://cmr.soc.plymouth.ac.uk/

**References**

1. D. Cope, **Computers and Music Style**, A-R Editions, Madison, 1991.
2. I. Xenakis, **Formalized Music: Thought and Mathematics in Composition**, Indiana University Press, Bloomington, 1971.
3. J. Casti, **Would-be Worlds,** John Wiley & Sons, New York, 1997.
4. E. R. Miranda, *Emergent Sound Repertoires in Virtual Societies,* **Computer Music Journal 26** (2), pp. 77-90, 2002.
5. R. Dawkins, **The Blind Watchmaker,** Penguin Books, London, 1991.
6. E. R. Miranda and Q. Zhang, *Composition as Game Strategy: Making Music by Playing Board Games Against Evolved Artificial Neural Networks,* **Proc. Int'l Computer Music Conf. (ICMC 2005),** Barcelona, 2005.
7. E. R. Miranda and V. Tikhanoff, *Musical Composition by Autonomous Robots: A Case Study with AIBO,* **Proc. Towards Autonomous Robotic Systems (TAROS 05),** London, 2005.

# Live algorithms

The EPSRC-funded Live Algorithms for Music (LAM) research network is establishing an inter-disciplinary community of musicians, software engineers, and cognitive scientists. Our aim is to investigate autonomous computers in music.

The use of computers in live music is not new: the fields of generative (algorithmic) composition and live electronics are of particular interest to LAM. A key discriminator between these is the degree of interaction with the performer. Interaction is intrinsic to *live electronics.* A performer may jam with commercial or custom software: a 'laptop-as-instrument' paradigm in which the computer is controlled directly. Another approach links players of traditional instruments with computers: incoming sound or data is analysed by software, and a resultant reaction (e.g. a new sound event) is determined by pre-arranged processes. Such 'reflex systems' can accompany performance, but might also use stochasticity to effect surprise, as determined by organizational decisions made by the composer/designer. We would term such a system 'weakly interactive' because there is only an illusion of integrated performer-machine interaction, feigned by the designer.

*Algorithmic composition* generates music off-line, although it can be used in real-time. Algorithms based on fractals, chaos theory, neural networks, and evolutionary computing have been exploited by composers for their patterning properties.[1] Such systems are not interactive, since all the parameters needed for sound generation are predetermined.

In contrast, strong interaction is exemplified in the human-only practice of *'free' improvisation.* This music rejects top-down organisation (*a priori* agreements, explicit or tacit) in favour of open, developing patterns of behaviour.[2] Social theories describe experiences with a sense of certainty, and with a unified artistic intent, as 'becoming situated'. An 'interactional semiotics' has been proposed, stemming from Meade's idea of emergence: the ensemble as a single entity exhibiting self-organising behaviours (see Reference 1 for references).

LAM is interested in computer systems that might interact strongly with musicians in both a supportive and a creative capacity, and the research agenda is a marrying of algorithmic music, live electronics, and free improvisation. Properties of human performance—and therefore of a live algorithm (LA)—include strong interactivity, autonomy, innovation, idiosyncrasy and comprehensibility.

*Strong interactivity* depends on instigation and surprise, as well as response. Individual decision-making is immediate, necessary and basic: when to play or not, when to modify activity in any number of parameters (loudness, pitch, tone quality), when to imitate or ignore another participant, when to 'agree' the performance is concluding. When to make a decision. And why. Without the capacity to *innovate*, listeners would lose the belief that the LA was truly engaged with the performance instead of merely accompanying it. The iterative, generative, *idiosyncratic* world of algorithmic organisation must be accessed, but the mechanical and the predictable must be avoided. It is the ability to innovate that distinguishes automation from *autonomy*. It is not hard to generate music of great complexity. Harder, though, is to ensure that these contributions are *comprehensible* to fellow performers who might not have heard the ideas before, and to do this in real-time. (But an incomprehensible, opaque system can be contrasted with a transparent one where the association between input and output is too trivial.)

Such considerations show the research goal is prescient, but there are reasons to believe that it is imminent too. Our own Swarm Music/Granulator systems implement a model of interactivity derived from the organisation of social insects.[3] These systems embody our idea of a proxy environment that holds meaningless sonic events. The system (human or machine) explores the environment, discovering and manipulating found sonic objects. Long-term organisation can

**Figure 1.** *The modular structure of a self-organising system. Analysis parameters* **p** *obtained from audio (Y) are mapped into the patterning space (F). In this example we see swarm particles drawn towards a new attractor, creating new synthesis data* **q** *for the resulting sound (X).*

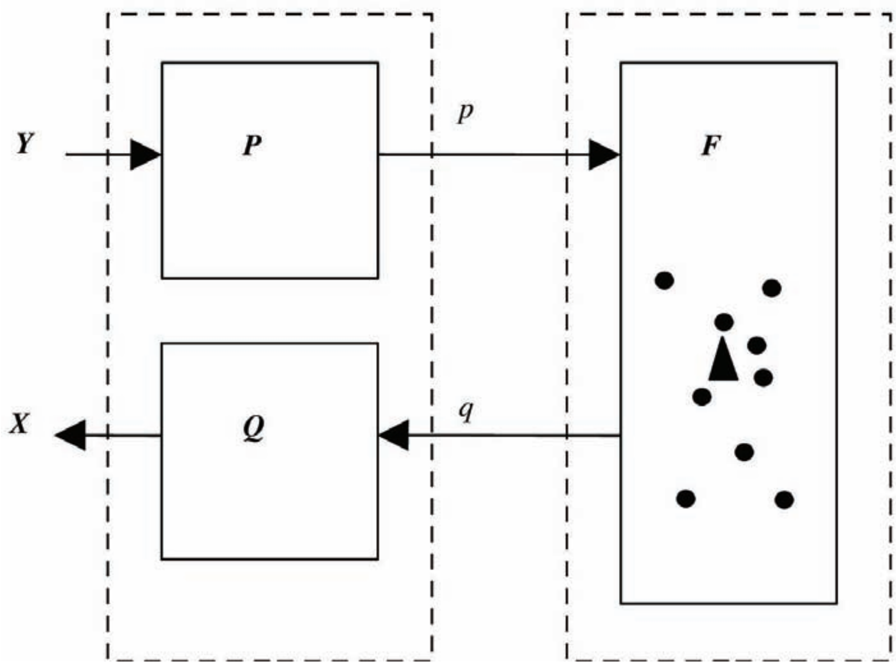



***Blackwell and Young***
***Goldsmiths College***

*Continued on p. 9*

## A fully automated system for facial muscle action detection

*Continued from p. 1*

AUs occurring alone or in combination in an input face video. There are four main steps: face detection,[3] facial fiducial-point localization,[3] point tracking,[4] and AU coding.[4,5] To detect the face region in the first frame of an input video, we use an adapted version of the original Viola-Jones face detector. The face detector consists of a cascade of Haar-feature-based 'GentleBoost' classifiers. Then, the detected face region is divided into 20 relevant regions of interest (ROIs), each corresponding to one facial point to be detected. Individual feature patch templates are further used to detect points in the relevant ROI. These are GentleBoost templates built from both grey level intensities and Gabor-wavelet features.

After 20 fiducial points are localized in the first frame, windows positioned around each of the facial points define the colour templates that we subsequently track for the rest of the image sequence with the auxiliary particle-filtering scheme. Based upon the tracked changes in the position of the fiducial points, we measure changes in facial expression. These changes are described as upward and downward movements of different points and as the increase or decrease of distances between certain points. Based upon the temporal consistency of the changes in question, a rule-based method encodes temporal segments (onset, apex, offset) of 27 AUs occurring alone or in combination in input face videos. An alternative is to use probabilistic actively-learned support-vector machines for AU recognition.[5] When tested on both the Cohn-Kanade and the MMI Facial Expression Database, our method achieved an overall correct recognition rate of 90%.[4,5]

To demonstrate the utility of this AU detection technology in understanding human facial behaviour, it was used as the basis for a case-based reasoning system capable of classifying facial expressions (coded in terms of AUs) into the emotion categories learned from the user.[6] The main advantage here is that the system can learn the user's association of any given facial display (a set of AUs) with any given interpretation such as an attitude, emotion, or an alternative to a traditional keyboard/mouse command. The latter is crucial for enabling interfaces driven by facial-expression (e.g., for disabled users).

**Maja Pantic**
Delft University of Technology
The Netherlands
E-mail: mpantic@ieee.org
http://mmi.tudelft.nl/~maja/

**References**

1. M. Pantic and L. J. M. Rothkrantz, *Toward an affect-sensitive multimodal HCI,* **Proc. IEEE 91** (9), pp. 1370-1390, 2003.
2. P. Ekman and W. Friesen, **Facial Action Coding System,** Consulting Psychologist Press, 1978.
3. D. Vukadinovic and M. Pantic, *Fully automatic facial feature point detection using Gabor feature based boosted classifiers,* **Proc. IEEE Int'l Conf. SMC,** 2005.
4. M. Pantic and I. Patras, *Detecting facial actions and their temporal segments in nearly frontal view face image sequences,* **Proc. IEEE Int'l Conf. SMC,** 2005.
5. M. F. Valstar, I. Patras and M. Pantic, *Facial Action Unit Detection using Probabilistic Actively Learned Support Vector Machines on Tracked Facial Point Data,* **Proc. IEEE Int'l Conf. CVPR,** 2005.
6. M. Pantic and L. J. M. Rothkrantz, *Case-based reasoning for user-profiled recognition of emotions from face images,* **Proc. IEEE Int'l Conf. ICME,** 2004.

## A test for irrational thinking

*Continued from p. 2*

(in this case, the distribution of utilities conditional on a certain action $z_i$). Thus, the maxEU principle only uses the most basic characteristics of the distributions to make choices. However, with Bayesian learning, agents have a better picture of these distributions (e.g. variance can also be estimated). Monte-Carlo methods use all known information about distributions, and thus irrational agents make better use of their knowledge. Moreover, the concept of optimality itself, at which the maxEU principle is aiming, has no clear meaning in stochastic worlds. This weakness of a totally rational approach is also visible from the behaviour of agents: the maxEU agent tends to over-exploit some places in the world, while irrational agents explore the whole environment.

One of the random decision-making methods was recently implemented as a subsymbolic mechanism for the ACT-R cognitive architecture. It was presented by the author at the ACT-R workshop in Trieste, 2005, and has been used to model some of the paradoxes of decision-making.

**Roman Belavkin**
School of Computing Science
Middlesex University, London, UK
E-mail: R.Belavkin@mdx.ac.uk
http://www.cs.mdx.ac.uk/staffpages/rvb/

**References**

1. J. von Neumann and O. Morgenstern, **Theory of games and economic behaviour** (First Edition), Princeton University Press, Princeton, 1944.
2. L. Savage, **The foundations of statistics,** John Wiley & Sons, New York, 1954.
3. F. J. Anscombe and R. J. Aumann, *A definition of subjective probability,* **Annals of Mathematical Statistics 34,** pp. 199-205, 1963.
4. M. Allais, *Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'Ecole americaine,* **Econometrica 21,** pp. 503-546, 1953.
5. D. Ellsberg, *Risk, Ambiguity, and the Savage axioms,* **The Quarterly J. of Economics 75** (4), pp. 643-669, 1961.
6. A. Tversky and D. Kahneman, *The framing of decisions and the psychology of choice,* **Science 211,** pp. 453-458, 1981.
7. J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere and Y. Qin, *An integrated theory of the mind,* **Psychological Review 111** (4), pp. 1036-1060, 2004.
8. R. V. Belavkin and F. E. Ritter, *The use of entropy for analysis and control of cognitive models,* **Proc. of the Fifth Int'l Conf. on Cognitive Modelling,** pp. 21-26, Universitats-Verlag Bamberg, 2003.

## *Brain-machine interfaces*

*Continued from p. 3*

exact mechanisms of cortical adaptation to artificial actuators. Progress in this endeavour will allow us to gain a deeper understanding of the mechanisms of learning and adaptation, and to build better prosthetic devices for the disabled.

**Jose M. Carmena**
Dept. of Electrical Engineering & Computer Sciences
Helen Wills Neuroscience Institute
University of California
Berkeley, CA
E-mail: carmena@eecs.berkeley.edu
http://www.eecs.berkeley.edu/~carmena

**References**

1. J. K. Chapin, K. A. Moxon, R. S. Markowitz and M. A. L. Nicolelis, *Real time control of a robot arm using simultaneously recorded neurons in the motor cortex,* **Nature Neruoscience 2** (7), pp. 664-670, 1999.
2. M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows and J. P. Donoghue. *Instant neural control of a movement signal,* **Nature 416,** pp. 141-142, 2002.
3. D. M. Taylor, S. I. Tillery and A. B. Schwartz. *Direct cortical control of 3D neuroprosthetic devices,* **Science 296,** pp. 1829-1832, 2002.
4. J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez and M. A. L. Nicolelis, *Learning to control a brain-machine interface for reaching and grasping by primates*, **PLoS Biology 1** (2), pp. 193-208, 2003.
5. S. Musallam, B. D. Corneil, B. Greger, H. Scherberger and R. Andersen, *Cognitive Control Signals for Neural Prosthetics,* **Science 305,** pp. 258-262, 2004.
6. J. R. Wolpaw, N. Birbaumer, D. J. McFarland, G. Pfurtscheller and T. M. Vaughan, *Brain-computer interfaces for communication and control,* **Clin. Neurophysiol. 113,** pp. 767-791, 2002.
7. E. C. Leuthardt, G. Schalk, J. R. Wolpaw, J. G. Ojemann and D. W. Moran, *A brain-computer interface using electrocorticographic signals in humans,* **J. of Neural Engineering 1,** pp. 63-71, 2004.
8. P. G. Patil, J. M. Carmena, M. A. L. Nicolelis and D. A. Turner, *Ensembles of human subcortical neurons as a source of motor control signals for a brain-machine interface,* **Neurosurgery 55,** pp. 1-10, 2004.
9. M. A. Lebedev, J. M. Carmena, J. E. O'Doherty, M. Zacksenhouse, C. S. Henriquez, J. C. Principe and M .A. L. Nicolelis, *Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface,* **J. of Neuroscience 25** (19), pp. 4681-4693, 2005.

## *Creating an artificial rocky-shore community*

*Continued from p. 5*

can be extended to more environmentally-sensitive ecosystems such as coral reefs, sea-grass beds, and mangroves. These are environments where habitat complexity is thought to be extremely important in maintaining biodiversity, but the large-scale experiments required to demonstrate such theories may cause damage to the habitats themselves.

**Richard Stafford, Mark S. Davies,* and Gray A. Williams[†]**
School of Biology
University of Newcastle-upon-Tyne
*Ecology Centre
University of Sunderland
[†]Swire Institute of Marine Sciences, Department of Ecology & Biodiversity, The University of Hong Kong
E-mail: richard.stafford@ncl.ac.uk

**References**

1. M. G. Chapman, *Variability in trail-following and aggregation in Nodilittorina unifasciata Gray,* **J. Exp. Mar. Biol. Ecol. 224** (1), p. 48, 1 June 1998.
2. R. Stafford and M. S. Davies, *Examining Refuge Location Mechanisms in Intertidal Snails using Artificial Life Simulation Techniques,* **Lecture Notes in Artificial Intelligence** (in press).



## *Live algorithms*

*Continued from p. 7*

develop, just as it does in termite nest construction. Within this framework, we envisage a modular system comprising of analysis (P) and synthesis (Q) functions that interface and interpret the sonic environment and relay parameters to a hidden patterning algorithm (F), analogous to listening, playing, and musical thinking enjoyed by a human performer. This picture integrates interaction with algorithmic composition and exploits recent developments in real-time music analysis/synthesis.

The network has some 70 members, including representatives from France, Portugal, the USA and Australia. Activities include an open meeting and two network workshops each year. Each event features invited speakers, contributions from LAM project teams and performances. The next meeting will be 19-20 December 2005, with an international conference in December 2006. LAM warmly encourages AISB readers to participate: please see our web site.

**Tim Blackwell and Michael Young**
Departments of Computing and Music
Goldsmiths College, London
E-mail: t.blackwell@gold.ac.uk

**References**

1. E. Miranda, **Composing Music With Computers,** Focal Press, 2001.
2. T. M. Blackwell and M. Young, *Self-Organised Music,* **Organised Sound 9** (2), pp. 123-136, 2004.
3. T. M. Blackwell and M. Young, *Swarm Granulator,* **Applications of Evolutionary Computing EuroWorkshops 2004, Proc. LNCS 3005,** Springer-Verlag, pp. 399-408, 2004).
4. http://www.livealgorithms.org

# Secretary and Webmaster's Report

You will have seen in a recent bulletin that the society is hoping to put together a response to the RAE consultation exercise. We are also hoping, at some point in the near future, to make a delegation to EPSRC. We would therefore be interested to hear of any issues that you feel we should raise with them. Obviously, we are particularly interested in issues that are connected with Artificial Intelligence and Cognitive Science. Please email secretary@aisb.org.uk with any responses.

The new procedures for annual elections are now working well, and you should be receiving the annual call for committee nominations with this issue. In fact, you will have two calls for nominations, one for regular committee members and one for editor of the Quarterly itself. As announced at the AGM, the committee feels it is desirable if the Quarterly editor has a specific term of office (three years), and is elected specifically to that office, rather than being drawn from the committee.

A side effect of the move to annual elections is that we are likely to go slightly over the constitutional maximum for committee numbers for about six months. This was probably inevitable in any attempt to move to regular elections from the situation in 2004, when we had to run two elections and co-opt several people all at once simply to maintain a reasonably sized committee.

With my other hat on, I can report that it has been quite a busy year for the website. It received a complete overhaul over the winter. Much of the content has remained the same, but it has been reorganised and given a face-lift. Underneath, we are now using server-side includes and style sheets in the hope that it will be easier to maintain a consistent look and feel in future. The sharp-eyed will have noticed that the Weekly Bulletin archive has disappeared. This was previously maintained entirely by hand, and was an unreasonably time-consuming activity. I'm currently looking into solutions for automating this process. As usual, I'm happy to receive any feedback or suggestions for extending or improving the site.

**Louise Dennis**
AISB Secretary and Webmaster

# BOOK REVIEW

## Sweet Dreams: Philosophical Obstacles to a Science of Consciousness

### Daniel C. Dennett

**Publisher:** MIT Press **Hardcover:** April 2005, 216pp, £18.95 **ISBN:** 0262042258

**Dennett in a Nutshell**

Like Bertrand Russell, Daniel Dennett is appreciated by philosophers as an incisive debunker. Where Russell attacked the inconsistencies of 19th and 20th century philosophy, Dennett is seen as the pit-bull terrier at the heels of those who have attempted to retain inexplicability as a feature of the philosophy of consciousness of the last twenty years. Debunking is fine, but, again, like Russell, the debunker's own philosophy may not emerge as clearly as the criticism of others. In *Sweet Dreams: Philosophical Obstacles to a Science of Consciousness* (based on Dennett's Jean Nicoud lectures given to philosophers in Paris in 2001) Dennett takes the opportunity of crystallizing not only his objections to obstacles to a science of consciousness as perceived by others, but also of spelling out his favoured 'explanation' of consciousness which he has dubbed *heterophenomenology.*

At a recent encounter, US philosopher Alva Noë, on being reprimanded by Dennett for not having embraced heterophenomenology, simply suggested that had a shorter word been used, the world might have recognised Dennett's brilliance in having solved the problem of consciousness. As explained in the book, heterophenomenology is not a solution to any problem, just a sane way of being able to talk with scientific rigour about first person experience with the support of physiological and psychological data.

From this standpoint, Dennett severely slates David Chalmers' contention that data only leads to solutions of the 'easy' problem, as Zombies could produce the same data without being conscious. Dennett points out that to believe in the Zombie position (the Zombie Hunch) one transgresses scientific principles and then argues that this hunch *allows* one to perpetrate such transgressions. Round one to Dennett.

In round two he aims a knock-out punch at Chalmers' very idea that there exists a 'hard problem', as whatever scientifically-based explanations are put forward can be lobbed into the 'easy' basket leaving the 'hard' notion without definition. Qualia (the felt qualities of experience) are knocked out of the ring on similar grounds: they are a concept invented by individuals who had no need to invent it. On the constructive side, Dennett cites some neuro-functional models such as Baars' 'global workspace' ideas, but points to their fragility and questionable assumptions.

Those familiar with the literature of explanations of consciousness will possibly feel that they have heard all this before. They have, and Dennett clearly says so: this is his opportunity to step back a bit and streamline his earlier thought. The result is that the reader gets a punchy and amusing compendium of Dennett's ideas, neatly compressed into 178 pages.

**Igor Aleksander**

*Igor Aleksander FREng is Emeritus Professor and Leverhulme Research Fellow in Neural Systems Engineering at Imperial College London. Year 2000 IEE Outstanding Achievement Medallist for Informatics, he has written 13 books and 200 papers on neural approaches to intelligence, cognition and consciousness. His latest book—The World in My Mind, My Mind in the World—is published by Imprint Academic, 2005. He plays tennis and jazz drums.*

# The Life of A. Hacker

by Fr. Aloysius Hacker

Cognitive Divinity Programme
Institute of Applied Epistemology

## Episode 1: The Early Years

I was born to humble parents in a Mayfair slum in London in 1938. Neighbours would remark on how much I took after my father, but I saw little of him during my childhood, since he was a frequent guest at a high-security Government establishment at Wormwood Scrubs, undertaking what I was assured was top-secret work for His Majesty. Between these absences, he had a wide range of lucrative schemes providing luxuries for the rationed population of war-torn London: silk underwear, chocolate, petrol—all were available from his company SPIV™ (Sumptuous Provisions Indulging Vanity).

My father's entrepreneurial spirit soon manifested itself. At age five I set up a neighbourhood automobile protection association targeted at visiting motorists. For a modest contribution to our favourite charity, ANGELS™ (Automobile Nursemaid, Guardian, Escort, and Lookout Service) would guard their parked cars. The winning feature of my scheme was the practical demonstration, to those declining our services, of the kind of threat we protected against.

This period also saw the birth of my lifelong interest in advanced technology. Wartime created unparalleled opportunities for SPIV to divert army 'surplus' electronic equipment to domestic use. Large boxes of valves, resistors, capacitors, etc., would regularly arrive at our house in the middle of the night. I enjoyed helping to construct wireless sets that my father would then sell in street markets. By the time I was ten I was experimenting with a variety of electronic gadgets, including small robots and bugging devices. Monitoring my neighbours created a new business opportunity for ANGELS. I received sizable supplements to my pocket money from errant housewives anxious to ensure that their husbands remained unaware of their extracurricular activities. The vision sensors also enhanced my education in adult matters.

During my father's long absences on government business, my mother had found a lucrative hobby that both supplemented her income and supplied male companionship. Unfortunately, she was not willing to contribute to ANGELS's privacy protection service. Her refusal initiated a chain of events that culminated in her confinement in a government institution and the divorce of my parents. My parents both blamed me for these outcomes and I was estranged from them: they did not seem to understand that it was nothing personal, just business. Consequently, at age twelve, I became a functional orphan and entered religious and scholastic life at ABUSE™ (Academy for Belief in and the Upholding of Spiritual Education). This experience initiated my lifetime's association with organised religion, especially that organised by myself. But more of this in Part 2 of my autobiography.

*Don't miss the next exciting episode next quarter!*

## Want more say?

If you have lots of ideas about what we should have in the Quarterly, there are two positions you should consider: *Editorial Board Member* and *Editor*. Contact the editor or see the elections information accompanying this issue.

## Want to advertise?

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

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

## Want to write something?

The link above gives access to full guidelines for the submission of reviews and technical articles. Books available for AISB members to review are also listed.

The deadline for the next issue is:

***Friday 4 November 2005***

**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**