

Locust neurons inspire tech to prevent car crashes

Car crashes are responsible for over forty thousand deaths and over five million injuries each year in the US alone. Mechanisms and sensors to detect, avoid, or lessen the impact of collisions are therefore an area of important industrial research. The problem with traditional approaches to making collision avoidance mechanisms for cars lies in the huge amount of information processing needed to successfully determine if a collision will occur. This is especially important when both the car and the object with which it is colliding are moving, and involves calculations of trajectories, speeds, and many other characteristics of both the car and colliding object.

A solution to this information overload can be found in the visual system of a locust. The lobula-giant-movement-detector (LGMD) neuron (see Figure 1) is highly selective, registering images of colliding objects and ignoring those that would miss the locust by only a few centimetres. The circuitry of the LGMD is well understood and involves the parallel processing of visual information while maintaining the spatial arrangement found in the insect's retina. Using only simple mathematical operations and local interactions between the retinotopic arrangements of the neuron's inputs, computer models can be built that show the same selectivity for approaching objects.¹

Our current work brings together experimental and computational neuroscientists at Newcastle University, VLSI (very large silicon integration) microchip designers, and electronic engineers at both the Microelectronics Institute of Seville and the Analogical and Neural Computing Laboratory in Budapest, and an industrial partner: Volvo Car Corporation. Together, we are developing and testing a VLSI-based vision system for automotive use based on the locust's LGMD.

Recently, computer simulations of the LGMD neuron were stimulated with video sequences of road scenes and interfaced with a driving simulation game on a Sony Playstation. The simulated neuron was able to predict collisions up to one second before impact, allowing numerous collision mitigation responses to be performed by the car. This included a significant reduction in the speed of the car prior to collision. The excitation that triggers the collision alert is caused by the expansion of edges of the colliding car exceeding the spread of lateral inhibition (see Figure 2a). In

contrast, most non-collision scenes are fully filtered by the LGMD's inputs and little, if any, excitation can be seen (see Figure 2b).

Currently, a VLSI chip is being fabricated that exploits the key benefits of the locust's LGMD neuron. This chip will be tested in a number of robots and model cars, and in a real car later this year. However, our research into the use of insect-inspired collision avoidance has not finished. Combining the LGMD model with other insect-inspired visual neurons—such as directionally-sensitive motion-detecting neurons—can improve the information obtained about the colliding objects. Information can be obtained about the direction of the approaching objects, allowing modifications to collision mitigation responses to be made. Alternatively, the information can be used to provide even greater selectivity for colliding, versus non-colliding, objects.²

Although considerable research has been conducted into artificially implementing insect neural networks in robots, this project is one of the first to attempt a transition from laboratory research into industrial technology. Because of the relative simplicity of insect neural networks, compared with conventional engineering approaches, it may not be long before a large amount of everyday technology is inspired by insects.

This work is funded by an EU Future and Emerging technology grant (IST-2001-38097: LOCUST). Further information can be found at the project website listed

Stafford & Rind
University of Newcastle-upon-Tyne
Continued on p. 7




Figure 1. The LGMD model and its input afferents. Excitation is captured from each of 100x100 pixels of an input image (only a small sample of the processing units are shown). The excitation, caused by change in luminance of each pixel, is calculated by the P units and is passed directly to the S units. Excitation is also passed to the I units where it is converted into inhibition. The inhibition is able to affect neighbouring S units after a short synaptic delay. Resultant excitation, where the change in excitation caused by the change in image exceeds the spread of inhibition, passes to the LGMD (lobula giant movement detector).

In this issue:

Features

- **Yorick Wilks** 2
Companion agents on the web
- **Sellers & Paul** 3
Tyrannosaur speed
- **Giorgio Metta** 4
RobotCub approach to cognition
- **Aaron Sloman** 5-7
Altricial systems
- **James Anderson** 7
Perspex Machine IV

Reviews

- **Essential books** 8
Johnson on bio-inspired computing
- **Book** 9
Web on 'Mechanical Bodies, Computational Minds' collection

Society

- **Officers** 10
Chair's report

Obituary

- **Robert W. Milne** 11
U. of Edinburgh

Father Hacker

- **How to edit a journal** 12

Companions: personal agents to help with your whole digital life on the web

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

ISSN 0268-4179

Copyright © AISB 2005

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

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

Editor

Sunny Bains
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

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

AISB PATRON

John Barnden
University of Birmingham

AISB FELLOWS

Prof Harry Barrow
Schlumberger

Prof Margaret Boden
University of Sussex

Prof Mike Brady
University of Oxford

Prof Alan Bundy
University of Edinburgh

Prof Tony Cohn
University of Leeds

Prof John Fox
Imperial Cancer Research Fund

Prof Jim Howe
University of Edinburgh

Prof Aaron Sloman
University of Birmingham

Dr Richard Young
University of Hertfordshire

One part of my current research is based on the belief that the Internet does not serve less socially-competent citizens well, and certainly not excluded groups like the old: although of course it serves researchers very well indeed, since they created it! As the engines and structures powering the Internet become more complex, it will also be harder for the average citizen to get what they can from an increasingly-complex structure, by which I mean what Berners-Lee calls the future semantic web. (This is a much-overused term, but I intend it closer to home in the sense of knowledge management and ontology construction and use in a project¹ in which Sheffield participates). This structure will be used chiefly by artificial agents, as well as by people. My claim is that citizens will need personalised, permanent, companion agents to interface with the future Internet for them: both to interact with it and in some ways to protect them from its torrent of information about them and the world.

Artificial intelligence and computer dialogue technologies now make simple companions feasible: they will not be robots but can be thought of as objects as small as mobile phones that will, within ten years, become the main access to the Internet for everyone as standard non-Internet voice technology disappears. But the companion will be more than an Internet agent: it will offer company, memory, diversion, and a way of making sense of the older citizen's whole life of text and images that will held on the Internet. It can best be thought of as an intelligent aide from assembling one's autobiography from one's Internet store of text, sound and image for one's relatives, colleagues, and successors. The companion, if it is to make citizen journalists/authors of us all, must also have conceptions of time and events in order to do this, and this is now feasible, too. There will be a wealth of issues to be explored surrounding the trust in, fate of, and legal and social status of such long-term artificial companions.

The technical base for such a companion is largely to be found in current dialogue technologies. In particular, speech and (my concern) language processing will be important. The latter involves both the deployment of machine learning techniques to detect dialogue acts and semantic content in utterances and models of whole dialogues that reach beyond responses to particular input. However, it can reach back and forward in time to see what has been achieved or left undone in a conversation.² Dialogue agents, to be plausible, must not only learn. They must also learn individual points-of-view: their own, and—in the case of companions—those of their owners.^{3,4}

I ran a project in 1997 that won the Loebner Prize in New York for the best conversationalist of the year (CONVERSE⁵) and we learned a lot from that about

what makes conversation plausible. It is closely tied into appropriate politeness and social relations, for which there is very little theory for automata.

Such a project clearly has implications for the online society in a broad sense, but it can also be seen as a general technological investigation. On the one hand it involves the nature and construction of companion agents from available and new technology in speech and language research. On the other it is concerned with information/knowledge management of Internet content. The former requires technologies for conducting dialogues between agents and humans, and is relatively well developed. The latter rests on techniques to do with ontologies and information extraction that enable very large amounts of personal Internet information to be organised, preferably by their 'owner'.

I am a member of two Networks of Excellence relevant to this research. First, HUMAINE⁶ is an European Commission network concerned with making agents personal and believable. Second, M4L (Memories for Life) is an EPSRC network⁷ concerned with the idea that all information about an 80-year human life could now be put on an Internet server in 28 Terabytes. Its premise is that we need to be concerned not only with how to protect such information, but how an owner can use companion agents to make sense of it and derive a coherent time-line for their own life within the torrent of information. The last notion is very close to separate work, with my student Angelo Dalli, to time-stamp large parts of the Internet and to time-organize them both around individuals and to separate individuals. George W. and George H. W. Bush—whose Internet presence shares many features and roles—would be good candidates for the latter treatment.

Prof. Yorick Wilks

Dept. of Computer Science
The University of Sheffield, Sheffield, UK
E-mail: Y.Wilks@dcs.shef.ac.uk
<http://www.dcs.shef.ac.uk/~yorick>

References

1. <http://www.aktors.org>
2. Y. Wilks, N. Webb, A. Setzer, M. Hepple and R. Catzone, *Human dialogue modelling using machine learning*, **Recent Advances in Natural Language Processing III**, Nicolov, Bontcheva, Angelova, and Mitkov (eds), John Benjamins, Amsterdam, 2004.
3. A. Ballim and Y. Wilks, **Artificial Believers**, Erlbaum, New Jersey, 1991.
4. Y. Wilks, *Artificial Companions*, **Machine Learning for Multimodal Interaction, LNCS 3361**, Bengio & Boulard (eds.), Springer Verlag, Berlin, 2004.
5. R. Catzone, Y. Wilks and D. Levy, *Life after CONVERSE*, **Proc. 2nd Int'l Workshop on Human-Computer Conversation**, Bellagio, Italy, 1998.
6. <http://emotion-research.net>
7. <http://www.memoriesforlife.org>

Speed potential of giant tyrannosaurs

The locomotive capabilities of giant tyrannosaurs have been the source of much debate over the past several decades: some authors describe the animal as a slow scavenger, whilst others attribute to it an ostrich-like swiftness.¹ Recent work using biomechanical scaling from small bipedal birds suggested that, at best, an animal of this size could only manage a slow run.² To further investigate this problem we have been using evolutionary robotics to create biorealistic simulations and directly assessing their locomotive performance.

The field of evolutionary robotics was developed for the automatic creation of autonomous robots.³ Researchers have primarily concentrated on navigation and learning issues. However, the evolutionary technique also enables the spontaneous generation of locomotion in legged robots: it is therefore an excellent tool for investigating the locomotion of extinct animals. In this situation, the basic mechanics (limb segments, joints, muscles) of the model are based directly on fossil evidence, and a central pattern generator is used to activate the muscles and produce the required locomotion. The activation pattern is produced using a genetic algorithm optimisation procedure so that the locomotive performance achieved maximises some fundamental parameter such as locomotive economy or top speed. This technique was previously used to investigate bipedal walking in early hominids,⁴ where it accurately predicted the cost of locomotion: by altering the optimization to maximise distance travelled in a given time, rather than for a given amount of energy, it could be used to predict maximum speeds.

Our model is based on the reconstruction of *Tyrannosaurus (T.) rex* using an estimated body mass of 5700kg.¹ The linear dimensions are based directly on the reconstruction and inertial properties derived from a 3D extrusion CAD model. Muscle groups are modelled as point-to-point force generators acting around the hip, knee, and ankle joints with physiological parameters derived from their estimated masses (see Figure 1). However, as soft-tissues are only very rarely preserved in fossils, we use a range varying from 25% to 40% of body mass. Similarly, muscles vary in internal geometry depending on force and velocity requirements. To accommodate this uncertainty, the

simulation was repeated using different assumptions of fibre length and force generation capability to identify a range of possible speeds for a given limb mass.

The model was able to produce stable running after a few thousand iterations of the optimisation procedure (see Figure 1) with speeds from 6ms⁻¹ to over 15ms⁻¹ (see Figure 2). To check the validity of the reconstruction, the average muscle power output was calculated (see Figure 3). For comparison, the maximum power output of horse muscle is about 90Wkg⁻¹ but, by using elastic energy stored in tendons, an instantaneous power output of 4400Wkg⁻¹ can be achieved.⁵ If *T. rex* could also use such a mechanism, even the highest values are entirely plausible: this suggests that *T. rex* was indeed a fast-moving animal.

This approach to fossil reconstruction is still in its infancy and considerably more work needs to be done to validate its predictions based on experimental work with living animals. However, with more sophisticated simulations, we should be able to calculate the biomechanical performance limits (and produce realistic animated output) for a whole range of fossil animals.

Further information about this and other simulations can be found on the Animal Simulation Laboratory website.⁶

W. I. Sellers and G. S. Paul*

Department of Human Sciences
Loughborough University, UK

E-mail: W.I.Sellers@lboro.ac.uk

*Baltimore, MD, USA

References

1. G. S. Paul, *Limb design, function and running performance in ostrich-mimics and tyrannosaurs*, **Gaia** **15**, pp. 257-270, 1998.
2. J. R. Hutchinson and M. Garcia, *Tyrannosaurus was not a fast runner*, **Nature** **415**, pp. 1018-1021, 2002.
3. S. Nolfi and D. Floreano, **Evolutionary Robotics**, MIT Press, Cambridge, MA, 2000.
4. W. I. Sellers, L. A. Dennis, W. Wang and R. H. Crompton, *Evaluating alternative gait strategies using evolutionary robotics*, **J. Anat.** **204**, pp. 343-351, 2004.
5. S. W. Wilson and J. C. Watson, *A catapult action for rapid limb protraction*, **Nature** **421**, pp. 35-36, 2003.
6. <http://mac-huwis.lut.ac.uk/ASL/>




Figure 1. Still frames taken from the animation sequence at 60ms intervals. These are from a run with a forward velocity of 10.7ms⁻¹.

Figure 2. Graph showing the maximum forward velocity achieved by the model as the proportion of body mass in the legs is increased. The two lines represent different assumptions of the maximum contraction velocity of the muscle.

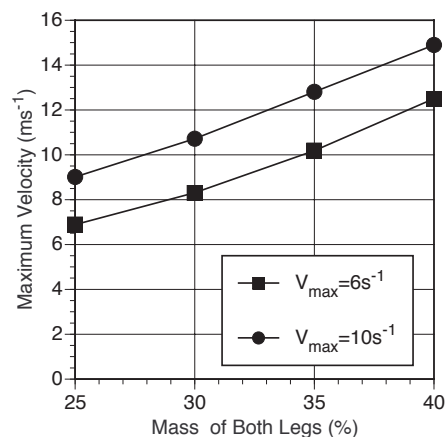
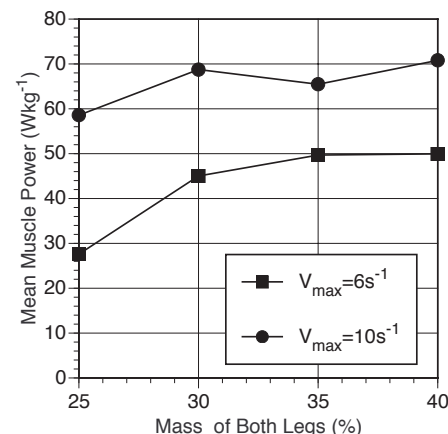


Figure 3. Graph showing the mean power output per kilogram of muscle for the model as the proportion of body mass in the legs is increased. The two lines represent different assumptions about the maximum contraction velocity of the muscle.



The RobotCub approach to the development of cognition

Two different stances dominate the study of cognition:¹ the cognitivist and the emergent. Greatly simplifying, the cognitivist approach is forged around the physical symbol hypothesis of Newell and Simon² and the emergent approaches are based, to various extents, on principles of self-organization and development.³⁻⁴ The RobotCub approach to the study of cognition falls within the latter category. Our previous work⁵ developed this argument in greater detail.

Cognitivist systems, because they are based on the assumption that cognition is simply the manipulation of symbols, need not be embodied: although they can benefit from the tuning of the symbolic engine by real-world learning. For emergent systems however, embodiment is fundamentally tied to the approach itself: emergence is obtained through the interaction with the environment, the shape of the body, and by means of sensorimotor coordination.

For this reason the first aim of RobotCub is to build a fully-fledged humanoid robotic platform shaped like a two-year-old child. The robot will have about 55 degrees of freedom. It will have sophisticated hands to manipulate objects, an oculo-motor system, and arms and legs for crawling, sitting, and interacting with the environment. The platform design, both hardware and software, will be an open system: it will be distributed under General Public/Free Documentation Licenses and shared with scientists interested in the study of embodied cognition.

The rationale is that, by creating a common platform, we will enable many laboratories to join this effort without having to invest themselves in developing yet another robotic platform. Our hope is that, over time, RobotCub will accumulate a substantial number of different skills: all learned and implemented by different research groups.

The second aim of RobotCub is thus to investigate the development of these cognitive skills. The project will carry out a plan of empirical research including neuroscience, developmental psychology, and robotics. This plan is centered on manipulation, ranging from the direct aspects of reaching and grasping for objects to the use of gestures for communication. Aspects that will be touched along the way are—for instance—looking and overt attention, reaching, the detection and

discovery of affordances, learning through imitation, and interaction.

The emergent approach naturally encompasses the study of ontogenic development and, in fact, a comparatively large effort will be devoted to its study. Our roadmap for this investigation⁷ includes the study of the *starting point* in terms of core abilities, the motive of the system to explore and gather data, and a few research areas such as looking, reaching and manipulation, posture, locomotion, and social interaction. For each of these areas, issues of prospective use of information, motivation, and the mechanisms of exploration have to be experimentally investigated. The RobotCub agenda aims at covering—through targeted empirical investigation—most if not all of these issues. Table 1 summarizes this agenda.

Finally, we wish to emphasize again that the principal motivation for this initiative is to help foster the study of embodied cognition throughout the global research community by making the RobotCub humanoid and cognitive software freely available. Representatives of this international community have been involved with RobotCub from the outset. Our goal is to increase this involvement as much as possible over the coming years and we welcome potential collaborators.

This work is funded by the European Commission's Cognition Unit, Directorate-General Information Society, as part of project no. IST-2004-004370: RobotCub—ROBotic Open-architecture Technology for Cognition, Understanding, and Behaviour.

The authors gratefully acknowledge Claes von Hofsten and Luciano Fadiga for the useful discussions about development and motor control in humans, and the whole RobotCub consortium for the preparation of this research agenda.

Giorgio Metta, David Vernon, and Giulio Sandini

Lab. for Integrated Advanced Robotics
University of Genoa, Italy
E-mail: pasa@liralab.it
<http://www.robotcub.org>

References

1. D. Vernon et al., *A research roadmap on cognitive vision*, http://www.ecvision.org/research_planning/ECVisionRoadmapv4.2.pdf
2. A. Newell and H. A. Simon, *Computer science as empirical inquiry: Symbols and search*, **Communications of the Association for Computing Machinery** **19**, pp. 113-126, 1975.
3. R. D. Beer, *Dynamical approaches to cognitive science*, **Trends in cognitive sciences** **4** (3), pp. 91-99, 2000.
4. T. van Gelder and R. F. Port, *It's about time: An overview of the dynamical approach to cognition*, **Mind as motion—Explorations in the Dynamics of Cognition**, MIT Press, Cambridge, MA, 1995.
5. G. Sandini, G. Metta and D. Vernon, *RobotCub: An Open Framework for Research in Embodied Cognition*, **IEEE-RAS/RJS Int'l Conf. on Humanoid Robotics**, Santa Monica, CA, 2004.
6. C. von Hofsten, *An action perspective on motor development*, **Trends in Cognitive Sciences** **8** (6), pp. 266-272, 2004.
7. C. von Hofsten, **Personal Communication**, 2005.

Table 1. The roadmap for RobotCub experimental research.

Development		Learning what	Perception/Action exploitation	Component of Social Interaction	Goal of system
Immediate	No reaching yet	Head-eye coordination	Gazing, smooth pursuit	Shared attention	Look around
	Pre-reaching	Approach an object	Controlling arm and hand movements in space	Pointing	Touch
Delay between action onset and consequences	Power grasp	Eye-hand coordination based on object position and object motion	Anticipatory dosing of the hand	Reaching for object held by other person "I got it!"	Grasp (become "owner")
	Differentiated grasping	Adjustment to object shape and size	Eye-arm-hand coordination based on object's shape	Take and give	Grasp appropriately (geometric)
	Object manipulation	Objects' affordances	Eye-arm-hand coordination based on actions to be executed on objects	Play games	Handle objects appropriately (use)
Long delays	Imitate acts on objects	Associate what is seen with what the system can do	What I do looks like what I see	Play games	Action's interpretation
	Act to communicate	Associate what is seen (perceived) with meaning	What I do generates some reactions	Communication	Action's meaning

Altricial self-organising information-processing systems

It is often thought that there is one key design principle—or, at best, a small set of design principles—underlying the success of biological organisms. Candidates include neural nets, ‘swarm intelligence’, evolutionary computation, dynamical systems, particular types of architecture, or use of a powerful uniform learning mechanism such as reinforcement learning. All of these support types of self-organising, self-modifying behaviours. But we are nowhere near understanding the full variety of powerful information-processing principles ‘discovered’ by evolution. By attending closely to the *diversity* of biological phenomena we may gain key insights into the following: how evolution happens; what sorts of mechanisms, forms of representation, types of learning and development, and types of architectures have evolved; how to explain ill-understood aspects of human and animal intelligence; and new useful mechanisms for artificial systems.

The precocial–altricial spectrum

Consider the relative influence of nature and nurture during development. The vast majority of species (e.g. grazing mammals, chickens, fish, reptiles, insects, ...) are ‘precocial’: the young are born or hatched relatively well developed and competent, with most behaviours genetically pre-programmed. Conversely, ‘altricial’ species start physically helpless and generally incompetent, requiring a period of support, including feeding, by parents. Paradoxically, some of these—e.g. humans, primates, hunting mammals and nest-building birds—exhibit cognitive capabilities of far greater sophistication in adult life than precocial species.

What can in principle be achieved by genetic pre-programming is shown by ‘cathedrals’ produced by termites, and by cognitive systems that are sufficiently powerful within hours of birth that they enable animals such as deer to stand up, run with the herd, find a nipple, and suck. In contrast, the fact that those animals that require the more complex and varied skills as adults tend to start life helpless and incompetent suggests that evolution discovered limits to pre-programming. It has added something else, something of great power, and something that is appar-

ently required for human intelligence.

The two labels ‘precocial’ and ‘altricial’ suggest a simple dichotomy between species whose behaviours are all innate and species whose behaviours result from learning and development. This is an oversimplification: there is a spectrum of cases. At every stage, in all animals, there are combinations of capabilities determined to varying degrees by the genome, embryonic development, maturation and kinds of learning. Precocial behaviours are largely developmentally fixed, but often allow calibration and gradual re-shaping through maturation or processes like reinforcement learning. Conversely, those species labelled ‘altricial’ because individuals start helpless and under-developed, nevertheless have some well-developed precocial skills at birth: e.g. those related to suckling in mammals and begging in birds. They also have some developmentally-fixed capabilities manifested later, e.g. flying at first attempt, and migration skills in altricial birds.

Known mechanisms for learning and self-organisation explain neither the genetically determined sophistication at the precocial end, nor the richness and diversity of achievements of individuals of the same species at the altricial end. In particular, nobody knows how to design a robot with the precocial capabilities of a new-born deer, and no known learning mechanisms could transform a helpless infant-like robot placed in an any country into a lively talkative child.

Altricial bootstrapping architectures

Analysis of nature/nurture trade-offs, and variation in requirements for ‘adult’ information processing systems, reveals a need for previously-unnoticed varieties of designs for artificial self-organising systems. Application domains where tasks and environments are fairly static and machines need to be functional quickly, require precocial skills: possibly including some adaptation and self-calibration. Others require altricial capabilities: e.g. where tasks and environments vary widely and change in complex ways over time, and where machines need to learn how to cope without being sent for re-programming. Architectures, mechanisms, forms

of representation, and types of learning may differ sharply between the two extremes. And the end results of altricial learning by the same initial architecture may differ widely.

Many species require rich cognitive structures closely adapted to complex features of the environment. Sometimes those requirements change rapidly: e.g. because individuals migrate to new terrain; because climate patterns or geological catastrophes produce rapid environmental changes; or because other species, whether prey or predators, learn new behaviours, or new varieties arrive from elsewhere. If requirements change too fast for natural selection to keep up, and too fast for the forms of self-modification produced by mechanisms like reinforcement learning, a more powerful form of learning is needed: as evolution seems to have ‘discovered’.

Where learned capabilities involve collaboration with conspecifics (others of the same species), rapid cultural changes can cause additional pressures favouring mechanisms capable of rapidly acquiring complex non-innate knowledge. These include novel ontologies, as demonstrated by very young human children picking up concepts their parents never had to learn at that age. Such mechanisms, in turn, can speed up cultural change: a form of positive feedback. A special case is language learning: where phonology, syntax and vocabulary learnt by a child born in one country may be very different from what the parents learned as children in other countries. What could produce so much structural variation in knowledge and behaviours within a species? Perhaps a new type of self-bootstrapping information-processing architecture evolved to enhance and complement both innate mechanisms and slow forms of individual learning.

Towards altricial architectures

We conjecture that altricial bootstrapping mechanisms, instead of being driven only by reward and punishment, also spontaneously discover discrete, re-usable, and

Aaron Sloman, U. of Birmingham

Continued on next page

Altricial self-organising information-processing systems

Continued from p. 6

(recursively) recombinable chunks of information. Perceptual or action patterns found during spontaneous play and exploration are selected for storage according to very general criteria e.g. symmetry, frequency, and production of complex effects through simple actions. Stored chunks can be used as components of larger chunks, which can be used in still larger chunks, using syntactic combination mechanisms forming conjunctions, sequences, loops, and conditional tests, enabling larger behavioural units to be formed, explored, and—if found ‘interesting’—stored as new units.

Such mechanisms might discover ever more complex re-usable structures, in percepts and in actions, and store them for future use both separately and in varied combinations. They might start from (implicit) innate knowledge not about the specific features of the environment but about generic (meta-level) features that can be instantiated in different ways in different environments. There would have to be innate mechanisms for combining structures to form new, more complex, concepts, actions, strategies, percepts.

More specifically, altricial learning may be based on genetically-determined mechanisms that have a number of important traits. For instance, they may have implicit meta-level knowledge. This could be about kinds of information chunks that might be learned, including perceptual chunks (using concepts of space and time), and action chunks (using a concept of causation). The knowledge could also be about the kinds of associations that might be perceived, and knowledge about how to investigate which are causal and which spurious.

As well as meta-knowledge, mechanisms may exist to combine old chunks into more complex wholes (e.g. complex goals, or action sequences) and discover new complex wholes that occur in the environment. The latter could include: enduring objects that have persistent features, parts, and patterns of behaviour; processes extended in time in which objects endure even when not perceived; and more and more complex actions produced and controlled by the individual. The creation and manipulation of hypo-

thetical structures, which might describe unobserved portions of reality or possible future complex actions, could be another important ability. Further, a mechanism for deriving consequences from complex information structures, and for comparing and selecting between complex structures with different consequences, might also be required.

Variants of such altricial mechanisms tailored to communication might support development of languages with combinatorial syntax and semantics.

In humans, and perhaps some other species, altricial capabilities that were originally *outwardly* directed (e.g. perceiving and acting on external objects and processes) might, after suitable architectural extensions, also be *inwardly* directed. This would allow individuals to develop more and more complex chunks of information not only about the environment, but also about their own internal processes of perception, reasoning, learning, problem solving, motivation, choosing, and so forth.

Ontologies used for such internal ‘meta-management’ could also be used in mechanisms for perceiving, reasoning about, and behaving towards others (e.g. conspecifics, prey and predators). Both the inward-directed and outward-directed cases require *meta-semantic* competence: the ability to represent and reason about entities which themselves process information. Animals and machines with such mechanisms can, for example, try to produce, change, or prevent beliefs, plans, or desires in others.

In humans, one aspect of growth of the architecture seems to be acquisition of new sub-ontologies, new forms of representation, and new collections of skills required for particular domains. The latter would include learning new languages, to read music and play instruments, programming, academic disciplines, athletic or dancing skills, mathematics, or quantum physics. Later growth enriches the architecture by growing new links between such domains: including using some as ‘metaphors’ for others.

Syntactic operations in such an altricial learner could themselves be either precocial

(genetically determined) or altricial, i.e. made of more basic building blocks that are assembled into larger units by learning during ‘playful’ thinking.

Individuals with such architectures are not limited to combinations of action units available at birth, but acquire more complex chunks indexed by their preconditions and effects. Searching for a combination that solves a complex problem may be very much faster than if the search either had to use more primitive units or had to use gradual modification of existing units.

The sources of meaning

The existence of mainly ‘precocial’ species shows that sophisticated visual and other apparatus can develop without individual learning. This means that the semantic content of the information structures is somehow determined by unlearned structures and how they are applied, refuting theories that require all symbol-users to base their concepts on ‘symbol-grounding’ using processes of abstraction from experiences of instances. Instead, meaning can be largely determined by formal structures that limit possible (e.g. Tarskian) models, combined with ‘symbol attachment’ to reduce residual ambiguity. This helps to explain how altricial systems can develop theories about the unknown and unobservable as humans do, e.g. in science.¹

Conclusion

These ideas seem to be close to Piaget’s theories about a child’s construction of reality. Which chunks an altricial individual learns will be influenced by physical actions possible for its body, the environment and its affordances, and the individual’s history. These factors could produce different kinds of understanding and representation of space, time, motion, causality and social relations in different species, or in similar individuals in different environments. If all this is correct, after evolution discovered how to make physical bodies that grow themselves, it discovered how to make virtual machines that grow themselves.

Aaron Sloman, U. of Birmingham

Continued on next page

Perspex Machine IV: Spatial properties of computation

Altricial systems

Continued from previous page

Whether computers as we know them can provide the infrastructure for such systems is a separate question.

The authors would like to thank their Birmingham colleagues on the CoSy project.² Please note that a longer draft paper on this topic is available online.³

Aaron Sloman and Jackie Chappell*

School of Computer Science

*School of Biosciences

The University of Birmingham

Birmingham, UK

<http://www.cs.bham.ac.uk/?axs/>

*<http://www.biosciences.bham.ac.uk/staff/staff.htm?ID=90>

References

1. <http://www.cs.bham.ac.uk/research/cogaff/talks/#meanings>
2. <http://www.cs.bham.ac.uk/research/projects/cosy/>
3. <http://www.cs.bham.ac.uk/research/cogaff/altricial-precocial.pdf>

Locust neurons: continued from p. 1

below.

Richard Stafford and F. Claire Rind

School of Biology

University of Newcastle-upon-Tyne, UK

{richard.stafford, claire.rind}@ncl.ac.uk

<http://www.imse.cnm.es/locust>

References

1. F. C. Rind and D. I. Bramwell, *Neural network based on the input organisation of an identified neuron signaling impending collision*, **J. Neurophysiol.** **75** (3), p. 967, 1996.
2. R. Stafford, R. D. Santer and F. C. Rind, *A bio-inspired visual collision detection mechanism for cars: Combining insect inspired neurons to create a robust system*, **Biosystems** (submitted).

Figure 2. Top: excitation passing into the LGMD (light or dark areas) after filtering by lateral inhibition is high when an object is about to collide. Bottom: when no collision threat is present there is little excitation passing into the LGMD.



AI has its many debates about the nature of computation, about the role of symbols and analog signals, about the relationship between the physical world and computational representations. The perspex machine aims to cut through the Gordian knot of these debates by combining geometry with computation, and relating these directly to the world. This provides, at least, a virtual machine that can exploit the geometrical properties of any computation, and which gives quantitative predictions about how the geometry of space limits any form of symbolic computation. This is a radical departure for AI in which all computations are spatial. It throws up new ways of computing things, and challenges accepted ideas about computation.

The perspex machine arose from the unification of the Turing machine with projective geometry.¹ In essence, certain geometrical objects were identified with the program tape and the states of a finite state machine, and certain geometrical transformations were identified with the operations of the Turing machine. This gave a constructive proof of how to make a Turing machine out of geometrical stuff.

The constructive proof guarantees that any Turing program can be compiled into a neural network. A C-source-to-perspex compiler has been implemented in Pop11. The compiler does a lexical analysis of the C source, performs a recursive-descent parse, then generates perspexes that are the data and operations specified by the C source. Initially, the compiler templates for data and operations were exactly the templates provided by the constructive proof, but these were soon adapted to provide a more convenient implementation of C's arithmetic operations, conditionals, loops, and function call and return.

A compilation of a C implementation of the Fibonacci series has the interesting property that the position of one neuron controls the number of Fibonacci terms computed. Thus, an important parameter of the program is mapped onto a spatial analogue automatically by the compiler. So far, the most complex program that has been compiled into a perspex, neural network is a C implementation of Dijkstra's solution to the Travelling Salesman Problem. The C source is about two pages long and the

compiled network has about 600 neurons. If the compiler were extended to cover the whole of C then it would be possible to compile any C source. It would be possible, for example, to compile the whole of Linux into a neural network.

There is a moral here for other researchers. Develop a constructive proof of the equivalence of your favourite kind of neural network with the Turing machine then implement it as a compiler. At a stroke, this will help AI deliver massive neural networks for use in all manner of software applications. This might be useful in itself, but the perspex machine does much more than this.

For a start, the perspex machine corrects a bug in the Turing machine. The Turing machine can enter a non-deterministic state where the current symbol on its tape instructs it to enter more than one state. In this condition the Turing machine stalls until an external agency, or oracle, decides which one state to enter. By contrast, the perspex machine is always deterministic, though it can emulate this Turing non-determinism, say, by raising flags to indicate that the Turing non-determinism has been encountered. This property of the perspex machine arises from the connectivity of geometrical space and its underlying, total arithmetic.² This arithmetic can be used on its own to remove division by zero errors from all numerical programs, thereby creating safer and more robust software.

There is a moral here, too. If you choose not to use a total arithmetic, you leave all of your software open to Turing's bug and risk your code crashing.

More profoundly, the perspex machine maps all Turing computations into geometrical stuff so that geometrical operations can be applied to them. For example, programs can be Fourier transformed and filtered so that the broadest filter band is a single neuron that approximates an entire program, and successively finer bands contain more and more neurons that, ultimately, reproduce the original program exactly.³ This makes it theoretically possible for a compiler to construct global-to-fine processing threads for any Turing program. In other

Anderson et al., U. of Reading

Continued on p. 10

Biologically-inspired computing

One important strand of AI research has been looking at the natural world, particularly the biological world, as a source of inspiration for computing. By looking at how biological systems generate and handle complexity, we can develop artificial systems that abstract the complexity-generating ideas from biosystems and use these ideas to handle complex computational problems in many domains.

Perhaps the earliest such metaphor can be found in neural networks. This is a vast field in its own right, worthy of a book list of its own; so I shall not mention any books on neural networks here.

A good overview of the whole area is provided by Peter Bentley's popular-science book **Digital Biology**. This gives a quick and easy-to-read tour of the important ideas in biologically-inspired computing, considering neural systems, insect behaviour, immune systems, growth and development, and many other topics. This is a good starting point if you want an overview of the area, for students wanting coverage of a number of topics before looking into the technical literature, or to explain this area of research to the general public: unlike some popular-science books, this is genuinely accessible to the general reader.

Perhaps the best developed of the biologically-inspired areas is that which draws its inspiration from evolution. There are many different algorithms that draw upon the basic ideas of evolutionary computing, but they are all grounded in a few basic ideas: maintaining a population of solutions to a problem then—over the course of time—selecting better solutions and producing new solutions by some kind of mutation and crossover of members of the population.

A good overview of the basic ideas of evolutionary computing is given in Melanie Mitchell's book **An Introduction to Genetic Algorithms**. This is a short book which nonetheless manages to explain how to write genetic algorithms, gives an introduction to the various theoretical tools that explain why they work, and provides an overview of a number of interesting application areas. A particularly strong feature of this book is the mixture

of exercises at the end of each chapter, including a mixture of thought and practical computing exercises.

In recent years, a number of additional topics and techniques have been introduced into the evolutionary-computing field. Of particular importance have been techniques that hybridise evolutionary ideas with other techniques: e.g. in 'memetic algorithms' that combine genetic algorithms with local search techniques. A book which gives a solid basic introduction to evolutionary computing, whilst also covering many such recent developments, is **Introduction to Evolutionary Computing** by Eiben and Smith. Each chapter of this book also contains a nice annotated summary of further reading from the research literature.

An important sub-area is the application of evolutionary ideas to the evolution of programs, rather than just fixed solutions to problems. The most significant technique so far in this area is *genetic programming*, which takes a population of programs described as parse trees and applies variants of the traditional genetic operators to evolve programs that perform well on a problem. A good tutorial on this is **Genetic Programming: An Introduction** by Wolfgang Banzhaf *et al.*. This covers both the core material on this topic and a number of variants and application areas. Another technique for evolving programs that is of increasing importance is *grammatical evolution*. This uses a traditional genetic algorithm as its problem-solving engine, but converts the members of the population into programs by treating them as encodings of sentences in a grammar that defines a programming language. O'Neill and Ryan's little book **Grammatical Evolution** provides a summary of these ideas.

Another area of biology that has proven a particularly rich source of inspiration in the last few years has been immunology. Our immune systems protect us on a day-to-day basis from a huge range of infections. Importantly for computing, the immune system learns: once we have been exposed to a disease, or immunised against it, we are protected from re-infection for many years. The immune

system has proven to be a particularly powerful metaphor for computer security: systems have been developed to recognise both patterns of attempted attacks on networks and computer viruses. Linked to this are many other application areas based on extracting unusual behaviour. Artificial immune systems have been used by financial organizations to detect fraudulent behaviour, and by others to detect spam. However there are a number of applications of these techniques beyond security and anomaly detection, e.g. in data classification, robot control, and scheduling problems. All of these ideas are discussed in **Artificial Immune Systems** by de Castro and Timmis, which also gives a good introduction to the immunological background necessary for understanding.

There are many other biologically-inspired paradigms that have received attention in recent years. One of the most important is 'swarm intelligence', which mimics the behaviour of large groups—swarms of social insects, human crowds, or particles in a fluid—through local interactions. A good recent book on a particular swarm intelligence concept, *viz.* taking inspiration from ants, is **Ant Colony Optimization** by Dorigo and Stützle. Some articles on other swarm techniques are available in the edited collection **New Ideas in Optimisation**. This book also contains some interesting articles on immune system methods and memetic algorithms.

Finally, we can look beyond simply taking inspiration from nature for computing and look instead at using natural systems to build new kinds of computers. An example of this is using chemical operations on problems encoded in long-chain molecules as a way of doing massively-parallel computing. A good overview of this area is Calude and Paun's **Computing with Cells and Atoms**, which gives a good overview of many different techniques. Martyn Amos's edited collection **Cellular Computing** looks like it will be a good reference source for recent developments of this kind. (*It was not available for review*)

Continued on p. 10

BOOK REVIEW

Mechanical Bodies, Computational Minds: Artificial Intelligence from Automata to Cyborgs

Stefano Franchi and Güven Güzeldere (eds) **Publisher:** MIT Press/Bradford Books
Hardcover/Paperback: January 2005, 544pp, £61.95/£29.95 **ISBN:** 0262562065

The aim of this book is to, "provide a forum for intellectual exchange between the artificial intelligence community and scholars in the traditional humanities and social sciences" which—the editors claim—has been "noticably rare" hitherto. Although there are substantive existing interactions between AI and psychology, philosophy and linguistics, management, music, art, and so on, it is true that most AI researchers do not much concern themselves with, "anthropology, history, literary criticism, cultural studies, religious studies, art history, theology and aesthetics". These—the editors argue—might well be thought relevant to AI issues such as the nature of human reason and creativity. So does this book make a convincing case that more attention should be paid to these subjects within AI?

The first section is a 140 page overview, by the editors, of automata, AI, Cybernetics, and cyborgs. Although necessarily this is highly selective, there were some odd compressions, for example identifying 'cybernetics' with 'perceptrons'. I found the organisation of topics in this essay unclear, as it was in the book overall. However, individually, many of the subsequent contributions make interesting reading.

The second section, *Automata, cybernetics and AI*, contains essays by Agre, Mazlish, Keller, and Pickering. Agre makes the useful point that algorithmic instantiation inevitably biases AI's view of human activity towards what can be most easily formalised: "formalisation becomes a highly organised form of social forgetting". What has been left behind, he says, often resurfaces as an impasse in technical research. Mazlish rehearses the history of automata and then provides a summary of several literary accounts of man-machines, but provides little analysis. Keller focuses on cellular automata, artificial life, and genetic algorithms, discussing how this research encroaches on the boundaries between machines and organisms. Pickering also explores these boundaries by describing three examples of 'weird' systems that lie between the animate and the inanimate: the anti-aircraft predictor of Wiener; Ashby's homeostat; and Prigogine's discussion of the Belousov-Zhabotinsky reaction.

The third section is titled *Controversies*

of *Artificial Intelligence* that includes an essay by a Hofstadter about his studies of analogy, and an exchange of letters between Dreyfus and Dennett about Deep Blue and Cog. It also contains an article by Dretske in which he sets out to defend the claim that computers, "don't solve problems, prove theorems, recognise patterns, let alone think, see and remember...": which seems to boil down to the need for symbol grounding.

After a diversion consisting of some dialogues from well-known AI programs, the next section finally begins to present viewpoints that may be more novel to an AI researcher. Latour and Teil's article on the 'Hume machine' suggests the natural strength of computers could be exploited in social science fields by using simple co-occurrences of words to build networks of associations, thus deriving concepts from raw text (they do not reference the similar research of Landauer & Dumais). In *Knowing subjects: AI from feminist philosophy*, Adam argues that AI has a limited view of intelligence as mere problem-solving logic and accumulation of knowledge. Other than characterising this view as typically masculine (itself a disputable claim) this criticism is hardly new, and her targets—CYC and SOAR—poorly reflect the current diversity of approaches in AI.

The following article by Harry Collins distinguishes 'regular actions'—in which the agent's intention can be executed in a variety of different behaviours (e.g. different words can be chosen to express the same meaning)—and 'behaviour-specific acts', where the intent is to perform a specific behaviour (e.g. to recite a poem correctly). He argues that the latter can be easily mechanised, whereas the former can be obtained only through socialisation. In the final article in this section, Michael Johnson proposes a semiotic approach to the frame problem, but I must admit I found this obscured the issues rather than clarifying them.

My copy of the book was missing 35 pages of section 6, so I cannot comment on the articles by Rieu and Sharoff. Wilson, an artist, describes several of his installations, and insights gained from these for AI: that human interactions with computers involves aesthetic as well as technical

considerations; that we should be wary of adjusting ourselves to fit technology; and that reflective intelligence requires self- and (cultural) world-models. Again, none of these points seem particularly novel. Matteuzzi present a definition of 'theory', and argues that AI as a science lacks its own universe; he relates this to several standard criticisms of AI.

Burke, following Mead and Dewey, discusses how perception and thought are two different modes of agent-environment interaction, the latter characterised by its symbolic character, the possibility of disengagement, and having evolved to support social interaction. Finally, Foerst provides a theological view of AI, noting that some of its questions are ones usually discussed within religious frameworks, for example, our experience of internal conflicts, and the possibility of immortality. However I was not convinced by her claim that these issues form part of the 'mythos' of AI, and are hence best analysed by theologians.

In summary I found this book engaging rather than enlightening. The more valuable articles seemed to come from fields (e.g. philosophy) that already have a strong interaction with AI, and the more remote fields (e.g. literary theory, art, feminism, theology) seemed to provide fewer novel insights. Nevertheless, it was useful to consider this broader perspective, and I would recommend the collection overall.

Barbara Webb

Barbara Webb is based at the School of Informatics at the University of Edinburgh. Her main research area is biorobotics, and she has an active interest in methodological issues of AI, particularly modelling.

AISB Convention '06

Details of the Bristol convention are now available at:
<http://www.aisb.org.uk>

Also, this years **Treasurer's report** can be downloaded by going to [.../treasurer](#)

Biologically-inspired computing books

Continued from p. 8

at the time this was written—Ed.)

Colin G. Johnson

Colin Johnson is a senior lecturer in computer science at the University of Kent at Canterbury. His research interests are mainly in the areas of computing and mathematical methods in biomedical research and in the use of natural science as an inspiration for novel computing methods. He is also interested in media and music technologies.

An Introduction to Genetic Algorithms, Melanie Mitchell, The MIT Press, 1998, 218 pp.

Ant Colony Optimization, Marco Dorigo and Thomas Stützle, The MIT Press, 2004, 328 pp.

Artificial Immune Systems, Leandro N. de Castro and Jonathan Timmis, Springer-Verlag, 2002, 375 pp.

Cellular Computing, Martin Amos, ed., Oxford University Press USA, 2004, 238 pp.

Computing with Cells and Atoms, Christah Calude and Gheorghe Paun, Taylor and Francis, 2000, 320 pp.

Digital Biology, Peter Bentley, Simon & Schuster Inc, 2002, 272 pp.

Genetic Programming: An Introduction by Wolfgang Banzhaf et al., Morgan Kaufmann, 1998, 512 pp.

Grammatical Evolution Michael O'Neill and Conor Ryan, eds., Kluwer Academic Publishers, 2003, 160 pp.

Introduction to Evolutionary Computing, A. E. Eiben and J. E. Smith, eds., Springer, 2003, 314 pp.

New Ideas in Optimisation, David Corne, Marco Dorigo, and Fred Glover, eds., McGraw-Hill Education, 1999, 450 pp.

Chair's Report

First I'd like to restate my great pleasure in the fact that Tony Cohn of the University of Leeds has become a fellow of the society. I am also delighted that Fiona McNeill of the University of Edinburgh has joined the society's committee as a result of the last election, and that Louise Dennis and Eduardo Alonso were re-elected.

I expect all those who attended AISB'05 will agree that it was a great success, both in terms of the interestingness of the papers and the unusually high number of delegates. I would like especially to thank the chair, Kerstin Dautenhahn, for her energy, enthusiasm and leadership, and for generating a higher-than-usual degree of international involvement. AISB'06 will be at the University of Bristol, and the society is indebted to the chair, Tim Kovacs, for this. The theme will be *Adaptation in Artificial and Biological Systems*.

This year the society and the British Computer Society (BCS) Specialist Group on AI are jointly organizing a public-understanding event at the Royal Society of Edinburgh, co-located with this year's International Joint Conference on Artificial Intelligence (IJCAI-05). It is a new type of endeavour for the society and a further step in a fruitful cooperation with the BCS group.

Your committee worked hard on making a variety of nominations to Research Assessment Exercise panels, both to the relevant main panel and to two sub-panels (so as to cover both computer

science and electronic engineering), and for various different types of member (academic, research-user, international, etc.). In particular, it sought to serve the 'SB' end of 'AISB' by making a nomination for a psychology observer on the CS sub-panel. We coordinated with the UK Computing Research Committee on some of the nominations.

Some of you will have seen an article and editorial in the New Scientist magazine (23 April 2005) concerning what has happened to AI and the need for more of a public debate about its ethical dimensions. I submitted a short letter in response, after consultation with the committee and fellows, basically saying that the society would welcome such a debate—and pointing out that for the time being any problems arise from deficiencies in people, not the evil machinations (word chosen advisedly) of AI programs! I encourage you to take a leading part in any debate that unfolds.

Finally, I am impelled to repeat my call for the more senior or the more unusually generous members of the society to become *patron* or *benefactor* members. This involves a higher membership fee. It also involves a warm sense of satisfaction. A good deal, I feel.

John Barnden

School of Computer Science
University of Birmingham
E-mail: J.A.Barnden@cs.bham.ac.uk

Perspex Machine IV

Continued from p. 8

words, in theory, global reasoning can be delivered by a compiler that compiles any existing program.

And there are deeper properties too. The Walnut Cake Theorem⁴ shows that, in general, when a discrete system approximates a finer discrete or continuous system it does so non-monotonically. Thus, non-monotonic reasoning is a general property of discrete machines operating in spacetime. Of course, monotonic reasoning can be had in certain special cases, but these are unrepresentative of the spectrum of computing machines

that can exist in spacetime.

There is a great deal more that could be said about the perspex machine, but this must suffice. Unifying the Turing machine with geometry has produced a new class of machines, perspex machines, that describe the shape and motion of objects in the world in a natural way, one that combines symbolic and non-symbolic computation in a single machine, and one which offers geometrical methods of computation that are, theoretically, more powerful than the Turing machine. Even Turing computable simulations of the perspex machine have surprising properties that make it a very powerful virtual machine with many potential applications in AI.

James Anderson, Matthew Spanner, and Christopher Kershaw

Computer Science, Univ. of Reading, UK
E-mail: author@bookofparagon.com
<http://www.bookofparagon.com>

References

1. J. A. D. W. Anderson, *Perspex Machine in Vision Geometry X1*, **Proc. SPIE 4794**, pp 10-21, 2002.
2. J. A. D. W. Anderson, *Exact Numerical Computation of the Rational General Linear Transformations in Vision Geometry X1*, **Proc. SPIE 4794**, pp 22-28, 2002.
3. J. A. D. W. Anderson, *Perspex Machine III: Continuity Over the Turing Operations in Vision Geometry XIII*, **Proc. SPIE 5675**, pp 112-123, 2005.
4. J. A. D. W. Anderson, *Perspex Machine II: Visualisation in Vision Geometry XIII*, **Proc. SPIE 5675** (entire volume).

OBITUARY

Robert William Milne

A key figure in pioneering artificial intelligence applications, Robert William Milne died while climbing Mount Everest early on 5 June 2005. He was 48.

Rob was born in Libby, Montana, and held dual US and UK citizenship. Brought up in Colorado, he was educated at MIT, receiving a B.Sc. in Electrical Engineering and Computer Science in 1978. He then moved to Edinburgh where he met and married his wife, Valerie, in 1981. Following the award of his Ph.D. in artificial intelligence from the University of Edinburgh in 1983, he began to seek increasingly innovative applications of AI in the real world, becoming Chief AI Scientist for the Pentagon in 1985.

Returning to Scotland in 1986, he founded Intelligent Applications Ltd. in Livingston, one of the first UK companies to market expert systems technology. Under his astute direction, the company became an industry leader in developing intelligent software solutions: a fact recognised by many accolades including the Queen's Award for Technology.

Rob has been a leader in the information technology field in Scotland, having for a time been Director of ScotlandIS, the industry body for IT and software. He was a mentor to a number of start-up companies and guided other entrepreneurs in their efforts to establish successful businesses.

Despite these demands on his time, Rob also engaged enthusiastically with academia and the wider AI and software-engineering communities. He was one of those rare individuals able to maintain a link between his academic and industry work. Through a variety of visiting and honorary posts at a number of universities, he assisted in maintaining relevance to industry and still found time to publish the results of his own work in traditional academic journals. He chaired many of the major conferences in AI and played a leading role in the field in Europe, becoming the president of the European Coordinating Committee for Artificial Intelligence in 2000. Most recently, he led the successful bid to bring the world's major AI conference, the International Joint Conference in Artificial Intelligence, to Scotland in 2005, only the second time that the meeting has been held in the UK (the last was in 1971). In recognition of his research work and leadership, Rob was elected to Fellowship of the Royal Society of Edinburgh in 2003.

Rob was already a keen mountaineer when he arrived in Scotland to begin his Ph.D.. Indeed, in his first meeting with his prospective supervisor he demonstrated how to climb a vertical brick wall: the supervisor declined to try. Milne was Munroist number 1860—a Munroist is a mountaineer who has successfully ascended the complete list of 284 mountains in Scotland—who 'bagged' his final Munro



Top: Robert Milne, who died on Mount Everest in June, as he will be remembered by his colleagues in artificial intelligence and by the mountaineering community (inset).

Bottom: Milne with Rick Hayes-Roth (far left), Edward Feigenbaum (centre right), and Austin Tate (far right) in 1986.

in 1997. He went on to become a senior figure in the Scottish Mountaineering Club and the author of a book on the Corbett hills.

Rob's life was characterised by setting very ambitious goals and single-mindedly pursuing them until he succeeded. His prominence in AI and software engineering and the achievements and accolades that followed are testament to his vision and tenacity. He led, inspired, and befriended many of the people he met and will be sorely missed.

He is survived by his wife, Val, and his children, Alex and Rosemary.

Alan Bundy and Austin Tate
University of Edinburgh

AISB COMMITTEE MEMBERS

Chair

John Barnden
University of Birmingham
chair@aisb.org.uk

Vice Chair

Eduardo Alonso
City University, London
vicechair@aisb.org.uk

Secretary/Webmaster

Louise Dennis
University of Nottingham
louise.dennis@aisb.org.uk

Treasurer/Travel

Patrick Olivier
U. Newcastle-upon-Tyne
treasurer@aisb.org.uk

Awards

travel@aisb.org.uk

Membership

Gary Jones
University of Derby
membership@aisb.org.uk

Publications

Natalio Krasnogor
University of Nottingham
publications@aisb.org.uk

Managing Editors

AISB Journal
Eduardo Alonso
City University, London

Geraint Wiggins
Goldsmiths, U. London
aisbj@aisb.org.uk

Editor, AISB Quarterly

Sunny Bains
Imperial College London
aisbq@aisb.org.uk

Public Understanding of AI

Simon Colton
Imperial College London
sgc@imperial.ac.uk

Public Relations

David Brée
publicrelations@aisb.org.uk
University of Manchester

AISB Convention 2006

Tim Kovacs
University of Bristol
aisb06@aisb.org.uk

AISB Convention 2005

Kerstin Dautenhahn
University of Hertfordshire
aisb05@aisb.org.uk

Ordinary members

Dimitar Kazakov
University of York
kazakov@cs.york.ac.uk
Fiona McNeill
University of Edinburgh
f.j.mcneill@ed.ac.uk



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>

AISBQ EDITORIAL BOARD

Sunny Bains
Imperial College London
Joanna Bryson
University of Bath
Lola Canamero
University of Hertfordshire
Simon Colton
Imperial College London
Yiannis Demiris
Imperial College London
Louise Dennis
University of Nottingham
Colin Johnson
U. Kent at Canterbury
Gary Jones
University of Derby
Natalio Krasnogor
University of Nottingham
Kia Ng
University of Leeds
Frank Ritter
Pennsylvania State U.
Barbara Webb
University of Edinburgh
Geraint Wiggins
Goldsmiths, U. London

An appointment as the editor of a major journal is one of the highest accolades that your peers can bestow. Thoughtlessly handled, however, it can also be a major burden. So, it is vital to learn....

15. How to edit a journal.

Editorial work can be a back-breaking and thankless task: attracting and processing submissions, chasing dilatory referees, reading and checking papers, dealing with publishers and printers, consulting colleagues. In the past, entrepreneurial academics would press their research students and assistants into service to relieve them of these chores, but nowadays, wily researchers demand payment. We will see how to avoid these problems by avoidance or automation.

1. With so many rival journals and conferences available, it can be difficult to attract authors to submit to your journal, especially if it is a new one. Fortunately, several programs now exist for automatically generating research papers. Don't hesitate to use these programs to boost your journal's submission rates and give it a head start over the opposition.
2. Authors prefer a journal that gives them both a rapid and a positive decision. Unfortunately, AI referees are notoriously slow and critical. The Hacker Way to solve this mismatch is to cut out the middleman. By return of post, send every author an acceptance letter with a short, commendatory review. Instantiate these reviews from a set of templates. Don't worry if the review is not a perfect fit to the paper: no one will complain about an inappropriate review if it is also complimentary.

**Rapidly and routinely write
referee reports for all submitted
papers with CRAP™ (Customised
Reviews Automatically Produced)**

*CRAP™ uses the latest
computational linguistics techniques to
analyse submitted papers and provide
plausible paeans of praise.*

3. A journal's reputation depends on its impact factor, which is calculated from the average number of citations of the papers it publishes. Luckily, you have control over this statistic. In feedback to authors, you should insist that they cite lots of other papers from your journal as a condition of acceptance. Authors have become acclimatised to demands from referees to cite their papers, however distant the connection, so no one will find this requirement unusual.
4. Making your journal electronic avoids the need

to deal with publishers and printers, and opens up opportunities for automatic paper processing that will reduce your workload to a minimum. But how can you generate income without hardcopy sales? Advertisements are the answer. You have an audience of professional, cosmopolitan technophiles. This, combined with your glowing impact factors, will make advertisers drool, especially those selling 'adult' products.

**Advertisement targeting
matches the advertisement to the
reader**

*Hacker's SPOT-ON™ (Specialised
Promotion Offered to Target One's
Needs) identifies your quarries' needs
from an analysis of their reading
habits, then offers only products that
address these.*

5. Stephen Hawking was famously told that every formula in his book would half his audience. The same is true for the papers in your journal. Instead of page charges, adopt formula charges to authors who insist on using mathematics to describe their research. One way or another, the formula charge will increase your income. Authors will either pay the charge or decrease the number of formulae, leading to an increased circulation and higher advertiser's fees.

This is the last in the Hacker's Guide series. Look out for fresh directions for Fr. Hacker in the next edition of the AISB Quarterly.

Next AISBQ deadline:

5 August 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>