

Real-time camera-based localisation and mapping

In the field of mobile robotics, there has recently been significant progress in such real-time simultaneous localisation and mapping:¹ a field so popular that the acronym SLAM has come to be commonly used to describe it. If a robot is to be dropped into an unknown environment (on a remote planet, for instance) it needs a SLAM capability to navigate on the fly without the need for prior maps. Although we do not explicitly consider a specific robot platform, our work is very much in this spirit.

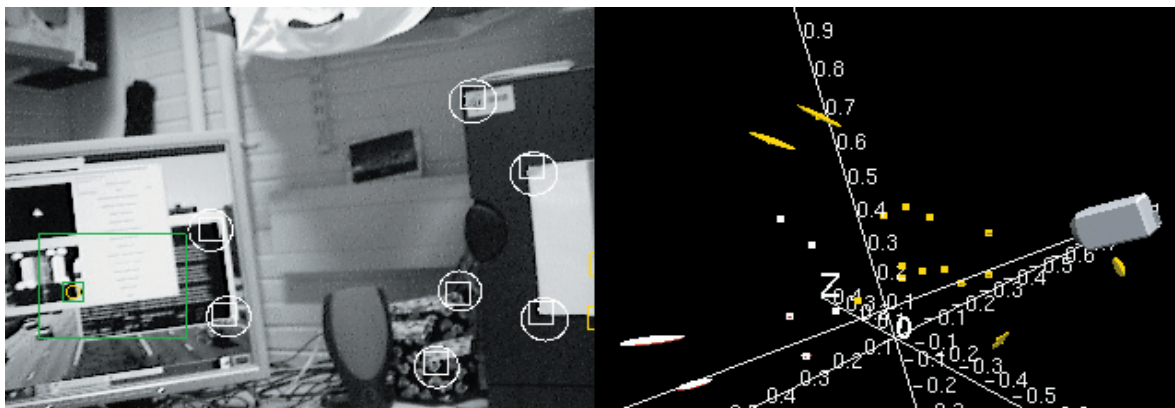
In 'structure from motion' research in computer vision, it has long been known that it is possible to recover the trajectory of a moving camera and the 3D geometry of the arbitrary scene it views purely from image information, on the assumption that that scene is rigid. However, successful implementation have, until now, used off-line processing to produce an after-the-event reconstruction of the camera trajectory and scene geometry. They have also used as input a complete image sequence captured, for instance, with a camcorder.²

In research published earlier this year³ we have demonstrated that it is in fact possible to recover such camera motion and scene-structure information with standard modern processors in real time (30 frames per second). Thus, a standard 'web-cam' module attached to a desktop or laptop PC becomes a flexible real-time position sensor. It can then be attached to a robot, human or other device whose motion we would like to estimate.

The demand for real-time operation imposes a strict constraint on the number of processing operations that can be used to digest the information available in each new image received from the camera. This is because the subsequent frame will be arriving a fixed time later—the speed of the CPU used and its frame rate pre-determine a constant upper bound. This necessitates not only an efficient approach to programming in general but, in fact, a quite different method to those used by those working in off-line structure from motion. The problem

for us is essentially one of sequential state-based estimation. Real-time constraints force the amount of information that can be extracted from each image to be sparse, and the key is understanding and propagating the uncertainty inherent in this data. Specifically, a sparse 3D model or 'map' of features of the static environment around the camera is constructed and maintained on the fly in the computer's memory. Also stored is information on the uncertainty in the map: information that is updated over time in a single extended Kalman filter. As certain image features (assumed to correspond to static entities in the 3D world) are repeatedly located in the images received over time, their image motions allow both the camera motions to be estimated and the feature-location estimates themselves to be improved.

Figure 1 shows a snapshot of our system in action, highlighting the issue of uncertainty propagation.



The 'features' described are square image templates selected automatically from images as points of high image contrast, these being generally associated with point geometrical features of the 3D world that are easy to locate repeatedly in subsequent images with a correlation operator. At any point in time, estimates of both the camera and feature locations are accompanied by uncertainty measures: the ellipsoid shapes around the estimated feature locations in the right-hand external 3D view above are uncertainty-bounded regions within which the features are believed to lie. The key to the efficient real-time operation of the system is that this uncertainty information can actually be used to predict, within bounds, the values of measurements before they are carried out. The elliptical search regions

Davison, continued on p.9

In this issue:

Features

- **Manfred Kerber** 2
Living with paradoxes
- **Kerstin Dautenhahn** 3
Robots and autistic children
- **John Zeleznikow** 4
Modelling legal knowledge
- **Anthony Pipe** 5
The Whiskerbot project
- **Hu and Calderon** 6
Robot learning through imitation

Reviews

- **Book** 9
Gary Jones on Shultz' 'Computational Developmental Psychology'
- **Conference** 10
Penny Noy on 'Engineering Societies in the Agents World'

News

- **Membership report**

Father Hacker

- **How to write papers** 12

Figure 1. Graphical output from our prototype system for single-camera SLAM. Left is the in-camera view, displaying square image-feature patches, elliptical feature-search regions, and the feature initialisation search box. Right is the external 3D view showing the recovered camera and feature locations with uncertainty ellipsoids.

Living with paradoxes

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

ISSN 0268-4179

Copyright © AISB 2004

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

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

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

Editorial Assistant/ Layout
Stuart Barr

Advertising and Administration:
Therrie Hendrey-Seabrook
aishq@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

Paradoxes have, from the days of the ancient Greeks, been a considerable source of exhilaration. They can be very funny, but also quite confusing and lead to inappropriate conclusions. In the most concise form, a logical paradox can be formulated by the sentence: "This sentence is false." Let's assume it is true, then it must be false, since it says that it is false. Hence we have a contradiction since it is true and false at the same time. So our initial assumption that it is true can't have been correct, hence it must be false. This leads to a contradiction as well, since if it is false, its content can't be true, hence the sentence can't be false either.

The problems with paradoxical sentences have been known for a very long time, but were thought to be of no relevance for the development of set theory and logic until around the end of the 19th, and start of the 20th, century. At that time, paradoxes in set theory and logic seemed to endanger progress in the field and, for a while, it wasn't clear whether "the paradise of sets" built by Cantor (as Hilbert called it) and "the paradise of logic" built by Frege weren't inhabited by terrible lions that could attack and kill at any time. Frege had formalised predicate logic and achieved a new level of clarity and rigour. Naive set theory was built by Cantor and first published around the same time. The first modern paradoxes were found in 1897 by Burali-Forti, and have been refined to those involving the cardinality of the set of all sets. Russell's notion of the set of all sets that do not contain themselves presented a major problem for set theory, and also meant that Frege's original system was reflective and paradoxical.

Zermelo¹ presented in 1908 the first axiomatisation of set theory, and Russell,² in the same year, the theory of types. Thus it was possible to fence the lions out. In the case of set theory, this was achieved using the foundation axiom that excludes

sets that contain (or, worse, do not contain) themselves. In the theory of types, it forbids a predicate to refer to itself. This way it was possible to build a safe area, free of paradoxes. More precisely, there might still be some wild beasts out there in the bush, but for a hundred years there were no further attacks and so we are pretty confident that there aren't any inside the fenced-off area. But, of course, you never know, that may just be a false kind of security.

While the constructs of axiomatic sets and types may

be considered as adequate for mathematical reasoning, there are quite a number of examples in the areas of natural language understanding and

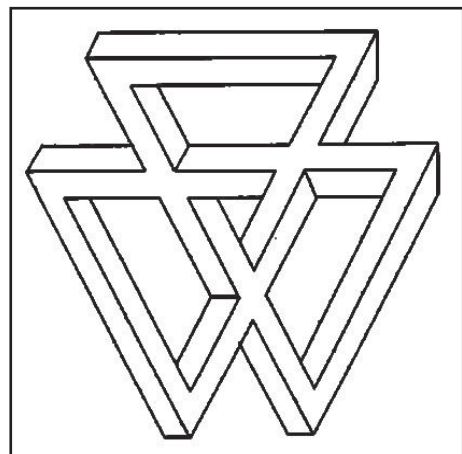
knowledge representation in a more general sense, where we need a more powerful language: where we can't exclude the self-referential *a priori* on mere syntactic grounds. If we did, we would end up with a system that is difficult to use and/or is not close to the everyday usage of language, which does allow for paradoxes. Actually quite a number of complicated formalisms have been developed and are currently in the mainstream of investigations in AI, although some simpler system may do the job better. Perlis³ pleaded in 1985 that we "can have everything in first-order logic." He investigated first-order logic plus strings, and redefined Tarski's definition of truth. I advocate a three-valued logic here.

First-order logic plus strings is an alternative knowledge representation formalism built on logic, and can be compared to modal logic. As Davis⁴ points out, the difference between the two approaches is pretty much the same as the difference between direct quotation—John knows "The evening star is the morning star," formally represented in syntactic theory by $\text{Knows}(\text{John}, \text{"EveningStar"} = \text{MorningStar})$ —and indirect quotation: John knows that the evening star is the morning star, formally represented in modal logic by $[\text{John}] \text{EveningStar} = \text{MorningStar}$. In both approaches, the truth values of the formulae are not extensional, i.e. the truth value of composed formulae cannot be calculated as functions of the sub-formulae. In modal logic the truth value in all possible worlds reachable from an initial world has to be known. In syntactic theory, expressions like "EveningStar = MorningStar" stand for strings of symbols and not for the objects they denote.

The approach is problematic when we adopt Tarski's definition of truth $\text{True}("A") = A$. Since the language is self-referential we are able to express the sentence "This sentence is false," formally as $L = \neg \text{True}("L")$. Together with Tarski's definition of truth, we get $L = \neg L$. This is contradictory in a two-valued setting. We want to be able to express such sentences not because they are particularly useful in their own right, but because it is difficult to draw a line between those self-referential statements that are useful, and those that aren't.

There are different ways out: one is to abandon classical logic and to live with contradictions (adapt a para-consistent logic), another is to forbid self-referential statements (Russell's approach), a third is to go for a three-valued logic in which the third truth value stands for paradoxicality. Just adding a third truth value doesn't solve the problem, since "higher-order" paradoxes, which involve reference to paradoxicality, are possible. When we disallow this, we can speak about the truth of sentences, for instance, we can say, "This sentence is false," but we must not say something like, "This sentence is paradoxical or false." This way, it seems possible to use efficient reasoners and come close to the treatment used in

Kerber, continued on p.7



Robots as assistive tools in the therapy and education of children with autism

The Aurora project¹ was set up to investigate the potential use of robots to help educate and provide therapy for children with autism who show impairments in communication, social interaction, imagination and fantasy, as well as a tendency towards repetitive and other stereotypical behaviours.² In the area of assistive technology, developing computer or robotic systems for children with autism is a challenging task.³ Literature suggests that people with autism enjoy interacting with computers,⁴ while interaction with people is usually a difficult and often frightening experience. Interactions with physical robots are fundamentally different from interactions with, for example, a computer game. Physical interactions are embodied and situated in the real world, and require the child to involve his/her body in a more extensive way than just operating a mouse or keyboard. Thus, the starting point of our work is the assumption that autistic children enjoy playing with robots, similar to an approach put forward by Ferrara and Hill.⁵

So far we have been conducting trials involving mobile robots and a small humanoid robotic doll. We use playful, enjoyable scenarios as a context in which children can engage in therapeutically- or educationally-useful behaviour, such as turn-taking, imitation, or joint attention. Importantly, our motivation is not to develop the robot as a replacement for carers. Rather, the robots should mediate

between what—from an autistic child's perspective—is the widely-unpredictable world of people, and the much more predictable world of machines. The envisaged purpose of our robots is to help autistic children to better understand and interact with other people.

We use two types of robots (see Figures 1 and 2): mobile robots with a vehicle-like appearance, and a stationary robot with a human-doll shape. The mobile robots are programmed so that the children can play simple interaction games with them, such as chasing, following, and other simple turn-taking games. The small humanoid robot, Robota, can move its arms, legs, and head, and has facilities for vision, speech, and producing music⁶ (initial trials are described by Dautenhahn and Billard⁷). In a recent longitudinal study, we allowed the children to move freely in a large room where Robota had been placed. In this setup, the robot was controlled remotely (a fact not known to the children) so that it could respond to subtle movements of the child playing with it. Results indicate that interactions with

this social scenario in a very constructive manner, demonstrating some communicative competence: i.e. they might use the robot as a focus of attention in order to interact and/or communicate with other people in the room.⁹

An important part of our work is the development of appropriate scenarios and



Figure 1. The mobile robot used in the trials.



Figure 2. The humanoid robot called Robota.

the robot in a social context involving the children's carer and the experimenter could encourage imitative behaviour.⁸

A key issue is that the children proactively initiate interactions rather than merely respond repetitively to particular stimuli. Additionally, the chosen set up is social, i.e. it involves not only the robot and the autistic child, but can include other children, the teacher, or other adults. Some children use

techniques in order to evaluate details of robot-child interactions.¹⁰ A range of different qualitative and quantitative evaluation techniques are likely to be needed in order to reveal not only statistical regularities and patterns, but also meaningful events of behaviour in context.¹¹

The Aurora project is highly multidisciplinary. The main investigators are primarily pursuing research in artificial intelligence, robot-human interaction, and assistive technology, but work in collaboration with psychologists and educators. The research questions involved touch upon a number of hard and challenging issues in robot-human interaction: e.g. evaluation methods and methodologies, design guidelines for interactive systems in therapy and education, and architectures and algorithms for robots that can adapt to individual people.

Our current and future work with the mobile robots and Robota includes both further detailed analysis of robot-child inter-

Dautenhahn, continued on p. 7

Modelling legal knowledge

At the Donald Berman Information Technology and Law Research Unit, Victoria University, Australia, we have been building legal decision-support systems for a dozen years. This work has been conducted with associated universities such as La Trobe, Ballarat, Bar Ilan (Israel), and the Joseph Bell Centre for Forensic Statistics and Legal reasoning at the University of Edinburgh. Industry partners have included Software Engineering Australia, Victoria Legal Aid, Phillips and Wilkins, Allan Moore and Company and JustSys.

Our goal has been to use artificial-intelligence techniques to develop legal decision-support systems. Recently, in cooperation with Victoria Legal Aid, we have been developing web-based legal decision support systems to increase access to justice. Our initial foray into this area started with an examination of the British Nationality Act of 1981 as a logic program. This led us to realise that ideally, such systems must deal with the issue of open texture—the notion that empirical concepts are necessarily indeterminate.

To better understand the nature of legal decision support, open texture, and discretion, we have used a number of inferencing techniques: association rules, case-based reasoning, machine learning, neural networks, and rule induction. Domains investigated include: workers' compensation, credit law, both property distribution and mediation within family law, refugee law, eligibility for legal aid, copyright law, eye-witness identification, examining the causes of death (natural causes, suicide, or homicide), sentencing, and the building industry.

Our research uses the jurisprudence of legal realists. For legal realists, rules and principles may be invoked after a decision has been reached in order to ensure that a decision is just, moral, and legally correct. Rules and principles are invoked to explain a decision but there is no need to assume they are used to reach the decision.

Given our desire to move beyond rule-based systems when modelling law, we commenced the IKBALS (intelligent knowledge-based legal systems) project. IKBALS used the object-oriented approach to build a hybrid rule-based/case-based system to advise upon open texture in the domain of workers compensation. A later version, IKBALSIII, included induction as the basis for its case-based retrieval function, and relied on distributed

artificial-intelligence techniques and the object-oriented paradigm, rather than a blackboard architecture. Induction was used in IKBALSIII to generate the indices for the cases. Thus, the developer could specify a number of cases, including the relevant factors and the outcome, and the induction algorithm would generate the indices automatically.

Our research concluded that legal reasoning in discretionary domains can be modelled using machine-learning techniques. The domain we chose was property distribution in Australian family law. Split-Up is a programme that provides advice on property distribution following divorce. The aim of the approach used in developing it was to identify, with domain experts, relevant factors in the distribution of property. They wanted to assemble a dataset of values on these factors from past cases that could be fed to machine learning programs such as neural networks. In this way, the manner in which judges weighed factors in past cases could be learned without the need to advance rules. The legal realist jurisprudence movement inspired this approach.

In the Split-Up system, the relevant variables were structured as data and claim items in 35 separate arguments. The claim items of some arguments were the data items of others, resulting in a tree that culminated in the ultimate claim that indicated the percentage split of assets a judge would likely to award the husband. In 15 of the 35 arguments, claim values were inferred from data items with the use of heuristics, whereas neural networks were used to infer claim values in the remaining 20 arguments. The neural networks were trained with data from only 103 cases. This was feasible because each argument involved a small number of data items.

In consultation with experts, 94 variables were identified as relevant for a determination. The way the factors combine was not elicited from experts as rules or complex formulae. Rather, values on the variables were to be extracted from cases previously decided, so that a neural network could learn to mimic the way in which judges had combined them.

However, according to neural network rules of thumb, the number of cases needed to identify useful patterns given 94 relevant variables is in the many tens of thousands. Data from this number of cases is rarely available in any legal do-

main. Furthermore, few cases involve all 94 variables. For example, childless marriages have no values for all variables associated with children so a training set would be replete with missing values. In addition to this, it became obvious that the variables were in no way independent.

In the Split-Up system, the relevant variables were structured as separate arguments following the argument structure advanced by Toulmin.¹ It performed favourably on evaluation, despite the small number of samples.

Because the law is constantly changing, it is important to update legal decision-support systems. The original hybrid rule-based/neural network version of Split-Up was constructed in 1996. Currently, the tree of arguments is being modified in conjunction with domain experts from Victoria Legal Aid to accommodate recent changes in legislation.

The argument-based representation facilitates the localization of changes and makes maintenance feasible. The use of the argument-based representation of knowledge enables machine learning techniques to be applied to model a field of law widely regarded as discretionary.

John Zeleznikow

School of Information Systems
Victoria University
Melbourne City, Australia
E-mail: john.zeleznikow@vu.edu.au

Reference

1. S. E. Toulmin, *The uses of argument*, Cambridge University Press, 1958.

Get free books!

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

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

Robot learning through imitation

The robots of tomorrow are expected to be able to co-habit with humans and perform services and tasks in a highly dynamic and unstructured environment. Such service robots should be able to understand human intention and emotion in order to help humans in personal tasks and even become friends with them. Hu-

issues: building a map of proprioception and the end-states of body parts; and both identifying body parts and perceiving the relationship between those of others and those of the robot itself. We propose an integrated approach to address these two issues, and a matching process is implemented in a real robot: the Pioneer

II mobile robot carrying an on-board arm with five degrees of freedom.

Firstly, we develop an algorithm that permits our robot to learn how to control its actuators as they map their movements to the end state. In other words, our robot should be able to describe its body configuration and calculate its joint motion trajectories by its internal model in order to achieve the motion trajectory of its end effector. In this way, the imitator (which has an on-board vision system) need only pay attention to the final effector state of the robot or human it is trying to imitate. Figure 1 shows that the imitator uses its onboard camera to observe the end-effector motion of the imitator (model), and treat it as the reference input. Since the imitator is able to calculate its body

configuration based on its physical internal model automatically, the joint commands will then be output to individual actuators to rotate in order to achieve the final target position of the end effector.

Our next aim is to reach the state of imitation of action on objects, which will introduce more complexity into the robot tasks and in turn increase the level of the interaction model for the imitator.

Hu and Calderon, continued on p.7

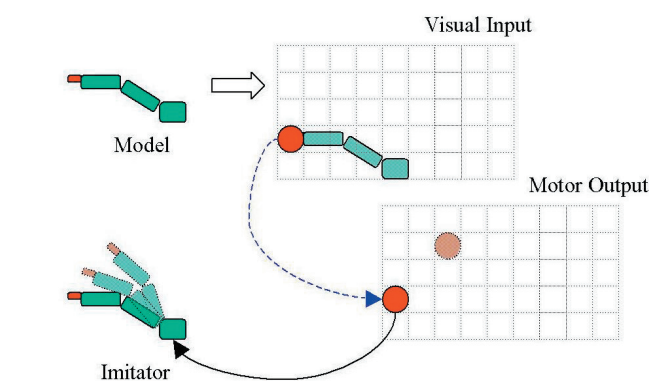


Figure 1. The focus of robot attention in on the end effector.

man-robot interaction becomes a central issue in these applications. However, many current robots are equipped with only a predetermined set of behaviours or actions, and have very limited learning and adaptation capability: in short, they are inadequate for such applications. As a result, robotics researchers have sought to engineer new robotics abilities through increased observation and communication. At Essex, we are currently developing intelligent service robots that are able to adapt their behaviours through learning by imitation as our human does.

Imitation is a very useful tool for human beings to acquire new behaviours and will have a key role to play in the development of intelligent service robots. Physiologists have addressed imitation of behaviour, and a four-stage progression of imitative abilities was proposed by Meltzoff: body babbling, imitation of body movements, imitation of actions on objects, and imitation based on inferring intentions of actions.¹ At the current stage of our research, we have identified two fundamental

trajectories by its internal model in order to achieve the motion trajectory of its end effector. In this way, the imitator (which has an on-board vision system) need only pay attention to the final effector state of the robot or human it is trying to imitate. Figure 1 shows that the imitator uses its onboard camera to observe the end-effector motion of the imitator (model), and treat it as the reference input. Since the imitator is able to calculate its body

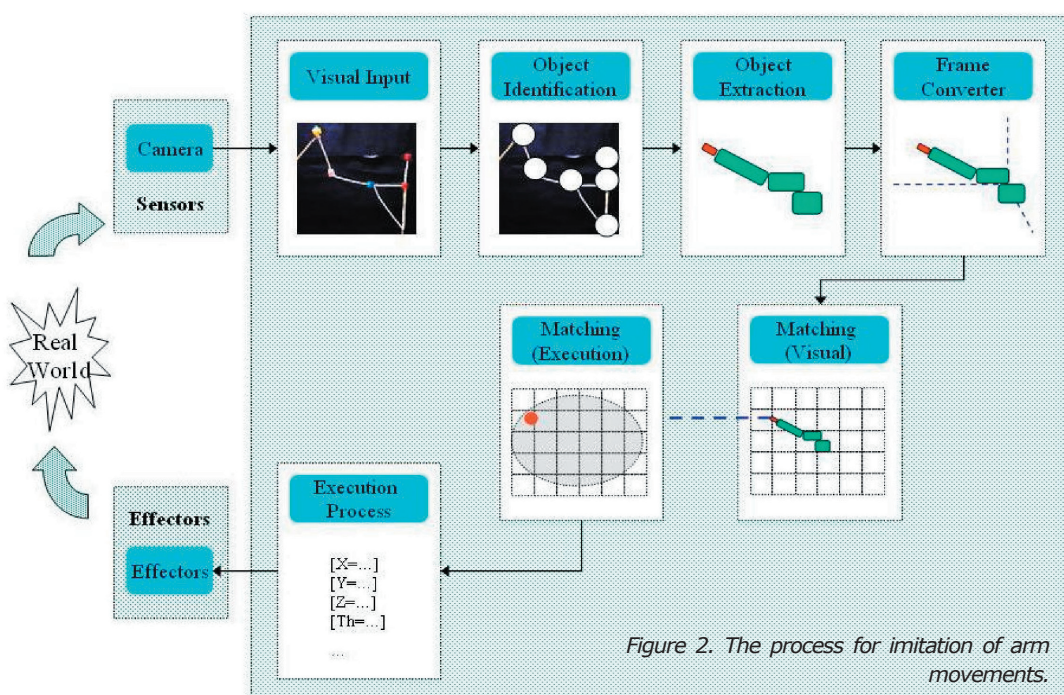


Figure 2. The process for imitation of arm movements.

Robot learning through imitation

Continued from p.6

This will have many potential applications for service robots to assist elderly individuals with cognitive impairment and other disabilities.

Huosheng Hu and Carlos A. Acosta Calderon

Department of Computer Science
University of Essex

E-mail: {hhu, cacaos}@essex.ac.uk

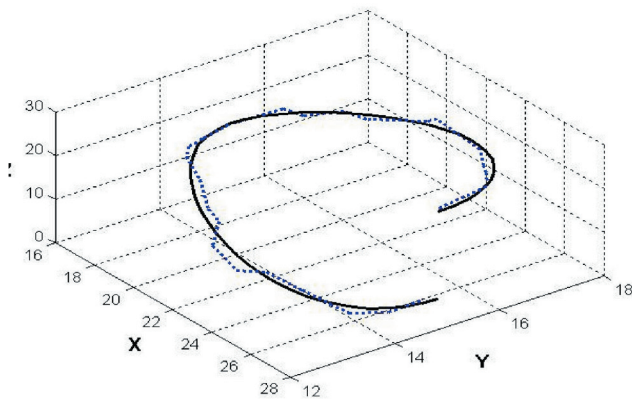


Figure 3. Motion trajectories by the imitator (solid line) and imitate (dotted line).

References

1. A. N. Meltzoff and M. K. Moore, *Earlier imitation within a functional framework: the important person identity, movement and development*, **J. Infant Behaviour and Development**, **15**, pp. 479-505, 1992.
2. C. A. Acosta-Calderon and H. Hu, *Goals and Actions: Learning by Imitation*, **Proc. of Int'l Symp. on Imitation in Animals and Artefacts**, Aberystwyth, Wales, pp. 179-182, 2003.

Living with paradoxes

Continued from p.2

everyday language. However, to deal with the full phenomenon seems—as usual—to be more difficult. For more details see my recent paper.⁵

Manfred Kerber

School of Computer Science
The University of Birmingham
Birmingham, England

e-mail: M.Kerber@cs.bham.ac.uk

WWW: <http://www.cs.bham.ac.uk/~mmk/>

References

1. Ernst Zermelo, *Untersuchungen über die Grundlagen der Mengenlehre I*, **Mathematische Annalen**, pp. 261-281, 1908.
2. Bertrand Russell, *Mathematical logic as based on the theory of types*, **American Journal of Mathematics**, **30**, pp. 222-262, 1908.
3. Donald Perlis, *Languages with self-reference I: Foundations (or: We can have everything in first-order logic!)*, **AI**, **25**, pp. 301-322, 1985.
4. Ernest Davis, *Representations of Commonsense Knowledge*, Morgan Kaufmann, San Mateo, California, USA, p. 77, 1990.
5. Manfred Kerber, *On Truth, Strings, and Paradoxes*, **Cognitive Science Research Papers CSRP-03-01**, The University of Birmingham, School of Computer Science, January 2003. <ftp://ftp.cs.bham.ac.uk/pub/authors/M.Kerber/TR/CSRP-03-01.pdf> or [/CSRP-03-01.ps.gz](http://www.cs.bham.ac.uk/~mmk/CSRP-03-01.ps.gz).

Robots as assistive tools in the therapy and education of children with autism

Continued from p.3

actions and the development of controllers that will allow the robots to identify the specific play and interaction patterns of particular children: information that can be used by the robot to adapt to individuals (see first results with non-autistic children by Salter et al.¹²). In the long term, robots that possess a behaviour repertoire that 'grows' alongside the child could guide children through different therapeutically-relevant interactions.

Prof. Kerstin Dautenhahn

Department of Computer Science
University of Hertfordshire, UK
E-mail: K.Dautenhahn@herts.ac.uk
<http://homepages.feis.herts.ac.uk/~comqkd/>

References

1. <http://www.aurora-project.com>
2. <http://www.nas.org.uk>
3. K. Dautenhahn, I. Werry, *Towards Interactive Robots in Autism Therapy: Background, Motivation and Challenges*, to appear in **Pragmatics and Cognition**, **12** (1), 2004.
4. S. Powell, *The use of computers in teaching people with autism*, **Autism on the agenda: papers from a National Autistic Society Conf.**, London, 1996.
5. C. Ferrara, S. D., Hill, *The responsiveness of*

autistic children to the predictability of social and non-social toys, **J. Autism and Developmental Disorders**, **10** (1), pp. 51-57, 1980.

6. A. Billard, *Robota: Clever Toy and Educational Too*, **Robotics & Autonomous Systems**, **42**, pp. 259-269, 2003.

7. K. Dautenhahn, A. Billard, *Games Children with Autism Can Play With Robota, a Humanoid Robotic Doll*, **Universal Access and Assistive Technology**, S. Keates et al. (eds.), Springer-Verlag, London, pp. 179-190, 2002.

8. B. Robins, K. Dautenhahn, R. te Boekhorst, A. Billard, *Effects of repeated exposure to a humanoid robot on children with autism*, **Proc. Universal Access and Assistive Technology (CWUAAT)**, Springer Verlag, London, 2004.

9. I. Werry, K. Dautenhahn, B. Ogden, W. Harwin, *Can Social Interaction Skills Be Taught by a Social Agent? The Role of a Robotic Mediator in Autism Therapy*, **Proc. CT2001 LNAI 2117**, M. Beynon et al. (eds.), Springer Verlag, pp. 57-74, 2001.

10. I. Werry, *Development and evaluation of a mobile robotic platform as a therapy device for children with autism*, Ph.D. thesis, University of Reading, UK.

11. K. Dautenhahn, I. Werry, J. Rae, P. Dickerson, S. Stribling, B. Ogden, *Robotic Playmates: Analysing Interactive Competencies of Children with Autism Playing with a Mobile Robot*, **Socially Intelligent Agents - Creating Relationships with Computers and Robots**, K. Dautenhahn et al. (eds.), Kluwer Academic Publishers, pp. 117-124, 2002.

12. T. Salter, R. te Boekhorst, K. Dautenhahn, *Detecting and analysing children's play styles with autonomous mobile robots: A case study comparing observational data with sensor readings*, **Proc. IAS-8**, IOS Press, 2004.

Help make a difference in the AISB

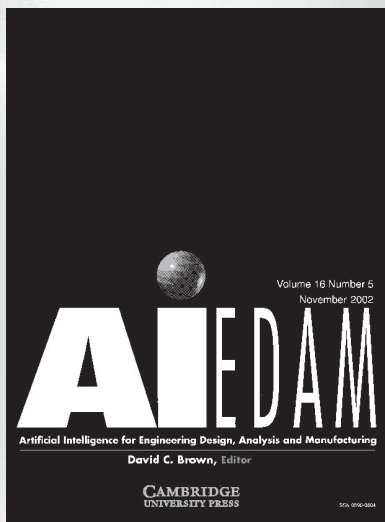
We are always looking for new talent to make the AISB a better organization: offering more services better-suited to the needs of its members.

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

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

chair@aisb.org.uk

The latest research in Artificial Intelligence ... from Cambridge



AIEDAM (Artificial Intelligence for Engineering, Design, Analysis and Manufacturing)

AI EDAM is a journal intended to reach two audiences: engineers and designers who see AI technologies as powerful means for solving difficult engineering problems; and researchers in AI and computer science who are interested in applications of AI and in the theoretical issues that arise from such applications. The journal publishes original articles about significant theory and applications based on the most up-to-date research in all branches and phases of engineering.

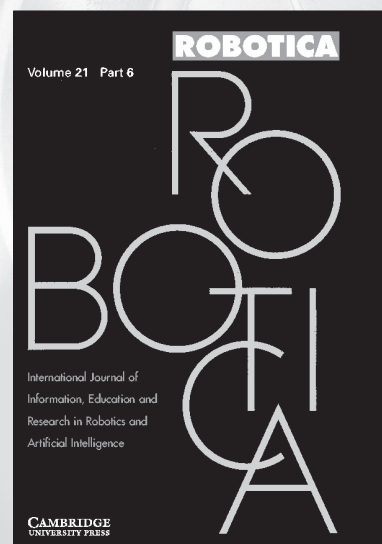
For a FREE online sample visit:
www.journals.cambridge.org/jid_AIE

Robotica

An Official Journal of the International Federation of Robotics

Robotica is a forum for the multidisciplinary subject of robotics and encourages developments, applications and research in this important field of automation and robotics with regard to industry, health, education and economic and social aspects of relevance. Interest is focused on various Artificial Intelligence topics of theoretical and practical interest.

For a FREE online sample visit:
www.journals.cambridge.org/jid_ROB



**To order, please contact Cambridge University Press at tel: +44 (0) 1223 326070
or email journals@cambridge.org**

Book Review

Tom Shultz is a Professor at McGill University in Montreal and has been working on computational models of development for some time, primarily using cascade-correlation neural networks. I've always been quite a fan of his method of modelling children's development and, though this book is mainly a summary of that work, it also includes some general modelling work.

There are seven chapters, the first being a concise introduction to the use of computational models in development, together with the reasons behind the use of these models. One primary use is to force the precision of verbal theories, which obviously applies to a host of domains. Another advantage, particularly for development, is that modelling allows confounding variables to be controlled for. This is particularly useful in children's development because their knowledge is being enriched all the time, so a method that allows knowledge to be controlled for has an obvious advantage. Shultz concentrates on neural networks for the main part of this book (in particular, cascade-correlation networks which are the focus of his work), which is why such an advantage is omitted. Neural networks usually begin with no knowledge of a particular domain apart from the design of the architecture, so although things like knowledge are being controlled for, it is a by-product of the system starting with a blank slate. In fact, the majority of the remainder of the book concentrates on neural networks, so much so that a more accurate title might be 'Connectionist developmental

Computational developmental psychology Thomas. R. Shultz

Publisher: *Bradford Book.* **Hardback:** *published 31 July 2003, 336 pp, £24.95.*

psychology'.

Chapter 2 covers how neural networks work, both in terms of the functions used to compute things (like how activation spreads), and how different learning algorithms work. Chapter 3 explains two general computational methods: production systems and connectionism. The aim here, to an extent, seems to be to dismiss production systems and promote connectionism, which seems unnecessary. Chapter 4 continues on this theme: the first part is a very good summary of proposed theoretical transition mechanisms, the second a comparison of a rule-based model versus a cascade-correlation connectionist model. There are distinct benefits for each type of modelling architecture, so it seems slightly harsh to dismiss one (albeit not altogether) in favour of another.

The real modelling of development occurs in Chapter 5 where the cascade-correlation algorithm is used in a variety of connectionist models to simulate behaviour in various developmental domains. This chapter begins with a thorough discussion of the opposing ideas of continuous versus discontinuous development which developmental psychologists have argued over for years. As Shultz points out, the recent thinking is that longitudinal views of development see it as discontinuous but a more micro view sees it as continuous. Shultz himself argues that, even at a micro level, you can still see discontinuous development. In terms of cascade-correlation networks, such discontinuities can be modelled by the addition of hidden layer units that cause sudden

spurts in development. A series of models of development are presented (e.g. on the balance-scale problem and seriation) which illustrate how cascade-correlation networks achieve this.

The final two chapters cover criticisms and future directions. In terms of criticisms, both critiques of modelling itself and of using connectionism are addressed. In terms of future directions, Shultz suggests various ideas that could be implemented using modelling, many of which are linked to how we know human development occurs. For example, we know that the human brain both grows and prunes neurons, and so this may be a significant factor in development that models need to capture.

In summary, this book is a very useful text for anyone who is doing research in developmental psychology, particularly if they have a computational slant to their research. It is possibly a little highbrow for undergraduates but certainly postgraduates and researchers should consider purchasing it: it provides a very good account of the use of cascade-correlation models of development, together with reasons for selecting this type of model. The book is marred only by its critique of production-system models that seems unnecessary, especially given the title of the book. Cut out the two half chapters that form the basis of the critique and change the title of the book to 'Connectionist developmental psychology', and you've got a winner.

Gary Jones

Psychology, University of Derby

Real-time camera-based localisation and mapping

Continued from p.1

shown in the left side of Figure 1 represent image regions within which the centres of the features of interest are predicted to lie with high probability. The sizes of these regions are determined by the uncertainty in the map and camera location at the previous frame, plus the extra modelled uncertainty introduced during one time-step of motion (which is assumed to be smooth to some parameterised degree). Computationally-expensive image search is therefore minimised by only searching within these regions.

This is the first system to be demonstrated that uses commodity hardware as simple as a single web-cam and standard PC, and that operates in full 3D in the information-rich visual domain. This is the type of technology that could at last open up applications such as domestic robotics, where the implementation of expensive dedicated sensors is not feasible. We have also received interest in the system from companies in such diverse fields as television, computer-user interfaces, industrial inspection and automobiles.

Andrew Davison

Robotics Research Group
Department of Engineering Science
University of Oxford
E-mail: ajd@robots.ox.ac.uk

References

1. P. M. Newman, J. J. Leonard, J. Neira and J. Tardós, *Explore and Return: Experimental Validation of Real Time Concurrent Mapping and Localization*, **Proc. Int'l Conf. on Robotics and Automation**, p. 1802-1809, 2002.
2. A. W. Fitzgibbon and A. Zisserman, *Automatic Camera Recovery for Closed or Open Image Sequences*, **Proc. 5th European Conf. on Computer Vision**, pp. 311-326, 1998.
3. A. J. Davison, *Real-Time Simultaneous Localisation and Mapping with a Single Camera*, **Proc. 9th Int'l Conf. on Computer Vision**, 2003.

Conference Report

Engineering Societies in the Agents World

29-31 October 2003, Imperial College London

The range of papers presented at this recent international workshop illustrates the multidisciplinary efforts being made to support the successful deployment of software systems involving very large numbers of autonomous components. The new computing environments will mean ubiquitous, persistent, pervasive devices, while at the same time allowing individual machines to access the resources and information from across the world, as enabled by the realization of the semantic web. The vision here is of agent societies embedded within human ones. The ESAW workshops have been held over the last four years to consider the different aspects of this problem.

At the most recent event, several presenters emphasized the importance of the environment, or infrastructure, in the coordination of general collective action by agents. Jean-Pierre Müller addressed the question of how we define and design emergence in his invited talk, *Emergence of collective behaviour from problem solving to social engineering and back*. He characterized weak emergence as where the interpretation is made by an external observer and strong emergence as where the interpretation is made by the entities involved. The role of the environment was emphasized on the micro level—for instance, to provide resources to the entities, and as an interaction medium—and on the macro level: e.g. to provide resources for emergent phenomena such as collective memory. It was emphasized that strong and weak emergence co-exist in social systems: there are multiple observers and implications, and thus a multiplicity of points of view and organizational levels.

Anthony Karageorgos and Nikolay Mehandjiev presented work on evaluation in their paper, *A design complexity evaluation framework for agent-based system engineering methodologies*. Agents are sophisticated software artefacts, associated with a large number of features, and therefore agent-based-system (ABS) engineering methodologies involve considerable design complexity. This paper explores different concepts of such complexity and proposes a framework to evaluate ABS methodologies. Since software engineering complexity relates to how difficult a system is to implement, anything that can reduce such complexity is important. The framework was used to evaluate a number of methodologies and indicate gaps in provision of support of those

already in use.

One element in the engineering of electronic societies is the examination of structures and mechanisms in human societies in order to assess their suitability for use with agents. Paul Feltovich and his colleagues, in their paper, *Social order and adaptability in animal and human cultures as an analogue for agent communities: towards a policy-based approach*, discuss some of the ways social order is maintained in animal and human realms in order to enrich our thinking about communities of agents. Specific work in agent-human interactions, using signalling of intended behaviour, is described, as is technical and social policy.

Federico Bergenti, in *Formalizing the reusability of software agents*, compares the reusability of agent-oriented and component-based technologies. His work suggests that—since agents are abstractions belonging to a new knowledge system level, and are described in terms of high-level concepts—the agent-oriented approach enables some concepts to be formalized more easily than traditional component-based methods. Illustrating this is his formalization of two problems, studied for some time in the literature: semantic interoperability and semantic composability. Another conclusion he reaches is that agents are potentially more reusable than components, though this may come at a cost.

The notion of competent agents is discussed in a paper written by Ulle Endriss and his colleagues: *Competent agents and customising protocols*. It argues that the ability to merely conform to a protocol, in the sense of not uttering any illegal dialogue moves, is not sufficient for an agent to be a competent user of that protocol. The notion of competence as the ability of an agent to reach a particular state of the interaction allows the authors to demonstrate automated competence checking of logic-based agents. Customization of

protocols is suggested for use by agents that are not fully competent.

Robert Tolksdorf and Ronaldo Menzes's paper—*Using swarm intelligence in Linda systems*—addresses the problem of scalability of tuple-space systems. Here the environment is explicitly considered as the context in which the individuals work and observe. The environment has a state that can be observed and changed by individuals.

In general discussion, aside from particular references in the presentations, the importance of the coordination medium was stressed: including the use of coordination artifacts used for embedding instructions etc.. Also, the interactions between agents and ordinary objects, as well as with the infrastructure, were considered. The overall message was not to be so focussed on the agents that the environment is forgotten. At the very least, one might want a blackboard or tuple space in the project to collect results and information. It may sometimes be better to have smarter environments than smarter individuals, or a combination of both.

Papers and presentations from the workshop are available online,¹ and selected papers are to be published in the *Lecture Notes in Artificial Intelligence* series (Springer-Verlag).

Penny Noy

Department of Computing
City University, London, UK
E-mail: p.a.noy@soi.city.ac.uk
<http://www.soi.city.ac.uk/~dk707>

Reference

1. <http://www.ai.univie.ac.at/~paolo/conf/ESAW03/esaw03accpapers.html>

Penny Noy is a research student working in the visualization of complex systems, and is part of agents@city within the Distributed and Intelligent Systems Group at City University.

Let us fund your travel

The AISB offers two student travel grants per year, each of £300. In 2004/5 two of the four grants are being tied to specific conferences: ECAI '04 and IJCAI '05. If you'd like to apply for the cash, please see the details at:

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

Membership Report

Our current membership is well over 400 people with interests in a variety of disciplines within the fields of artificial intelligence and simulation of behaviour. This has been reflected in the variety of articles that we have seen in the Quarterly, which is now packed with summary articles on the latest findings in AI. In addition to this, we have a twice-yearly journal that covers in-depth articles on AI research (and for which we have found a number of new Editorial Board members, see below). We feel this serves our members interests well and hope that you share this view.

Part of our membership fees go toward funding travel awards so that students are able to attend national and international conferences. We now give two £300 awards each year, which we believe is more realistic in terms of today's financial realities. If you are a student member and wish to apply, then you should email travel@aisb.org.uk with details of the conference you wish to attend. Those interested in attending the ECAI 2004 conference are particularly welcome to apply, as we have set aside one of our grants for this meeting both to strengthen our links

with ECAI and encourage student attendance.

We are also currently developing a database of membership interests. This will be useful for several things, such as putting members in contact with others located nearby or who do research similar areas, and for putting names forward as potential reviewers of articles. If you are interested in this and have not yet given your details to AISB, then print off the relevant page from the membership form on the AISB web pages¹ and send it to our offices.

Last, but not least, we always like to encourage current members to contact friends/colleagues who they feel might like to join AISB. One approach is to send them the web address¹ so they can see what the Society offers. For example, they can have a look at previous Quarterly newsletters on the site, plus the contents of past issues of the Journal...

Gary Jones

Membership Secretary

E-mail: membership@aisb.org.uk

1. <http://www.aisb.org.uk/>

Submit your paper to the Journal

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

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

John Barnden, The University of Birmingham, UK

David Bree, University of Manchester, UK

Paul Davidsson, Blekinge Institute of Technology, Sweden

Ramon López de Mántaras, Artificial Intelligence Research Institute, Spain

Nick Jennings, University of Southampton, UK

Rob Milne, Intelligent Applications, UK

Ulrich Nehmzow, University of Essex, UK

Tony Pridmore, University of Nottingham, UK

Frank Ritter, Penn State University, USA

Peter Ross, Napier University, UK

Gerhard Weiß, Technische Universität München, Germany

Geraint Wiggins, City University, UK

The editorial board is pleased to call for papers for Volume 1, Number 6, to be published at the end of 2004. Papers must be e-mailed as PDF files to the eduardo@soi.city.ac.uk. Please use the LaTeX template for the submissions available at AISBJ web-page (<http://www.aisb.org.uk/aisbj/index.html>). Papers are limited to 25 pages and should include only text and black-and-white line diagrams (authors requiring more complicated diagrams, such as half-tones and photographs, should contact the editors).

AISB COMMITTEE MEMBERS

Chair

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

Vice Chair/Co-Managing Editor, Journal

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

Secretary/Webmaster/Schools Liaison

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

Treasurer/Travel

Awards
Paul Chung
Loughborough University
of Technology
treasurer@aisb.org.uk
travel@aisb.org.uk

Membership

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

Public Relations

David Bree
University of Manchester
bree@cs.man.ac.uk

Public Understanding of AI

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

Industrial Liaison

Vacant

Co-Managing Editor, Journal/Publications

Geraint Wiggins
City University, London

AISBQ

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

AISBQ EDITORIAL BOARD

Eduardo Alonso

City University

Sunny Bains

Imperial College London

John Barnden

University of Birmingham

David Bree

University of Manchester

Lola Canamero

University of Hertfordshire

Paul Chung

U. of Loughborough

Simon Colton

Imperial College London

Kerstin Dautenhahn

University of Hertfordshire

Yiannis Demiris

Imperial College London

Louise Dennis

University of Nottingham

Colin Johnson

University of Kent at
Canterbury

Gary Jones

University of Derby

Frank Ritter

Pennsylvania State Uni-
versity

Barbara Webb

University of Edinburgh

Geraint Wiggins

City University



Father Hacker's Guide for the Young AI Researcher

Cognitive Divinity
Programme
Institute of Applied
Epistemology

About the Society

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

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

AISB membership includes the following benefits:

- Quarterly newsletter
- Biannual Journal
- 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>

The key to a successful research career is to publish the maximum number of papers with the minimum amount of effort. To achieve this it is essential to know....

10. How to write papers

1. To succeed in getting your work accepted for publication, you must pick referees who will like your papers. But surely the journal editor or conference programme chair chooses the referees, not the author? That kind of defeatist thinking is for losers. Editors are busy people—and programme chairs even busier. They use simple heuristics to choose referees; master these heuristics and you can choose your own. Cite several papers by your chosen referees. Explain clearly in the introduction, abstract and even title (the editor will not read further than this) that your work builds on that of your preferred referees.

2. Choose a topic that will appeal to your referees. Something accessible, i.e. not too technical, with lots of fireworks, will leave them with a warm glow. A polemic usually goes down well. Choose a popular target and you will carry both your referees, and eventual readers, with you. Attacking very technical and mathematical approaches to the field can work very well. You will be supporting the prejudices of many of your readers and reassuring them that they can safely avoid the hard work involved in mastering difficult material. You don't need to say anything original: reinforcing common misconceptions and restating previous criticisms are sufficient. Set your students the task of trawling the literature and summarising previous critiques—especially obscure or forgotten ones that you can safely claim as your original contribution. A polemic also has the major advantage of not requiring prior experimental or theoretical research, minimising the effort required to write it, as desired.

Ensure you get the referees of your choice with Hacker's DECK- STACK™...

*...Determine Editorial Choice of
Kindly Souls who on your Treatise
Are Commendably Keen.* It applies
a large collection of heuristic rules
to persuade the editor to be kind to
your paper.

3. Just as you have carefully identified your potential supporters, you should also identify those who will hate your paper. Get your retaliation in first by attacking them fearlessly—and by name. Such attacks will provide the necessary fireworks. Programme chairs, especially, love papers that create controversy at their conference. But isn't this a dangerous strategy, your wounded opponents are likely to be picked as referees? This was likely anyway, and they were never going to recommend

your paper. At least this way their counter-attack will be weakened: they may be stampeded into a rash over-reaction, undermining their criticism, and you can always accuse them of bias. If you are really lucky, they will be stung into publishing a rejoinder. Your polemic will now be promoted into a controversial issue. This is your meal ticket to a sequence of papers as the issue is batted back and forwards.

Turn on the heat with Hacker's FLAME™...

*...Fault-find, Libel, Argue-against
and Mock your Enemies.* Given a few
carefully chosen quotations, it will au-
tomatically construct your polemic.

4. Polemics sidestep the standard scientific conventions when it comes to organisation: background, specification, implementation, results, further work, are all irrelevant in the polemic. Start your paper with (your version of) your opponents' position, followed by a bold denunciation. Continue by illustrating each of your criticisms with selective quotation from their papers: choose obscure and over-technical passages, which you are then free to interpret for the reader so as to most effectively make your point. State your views forcefully and colourfully. This will not only provide the necessary fireworks and stir up a controversy, but it will allow plenty of openings for your opponent to come back with a rejoinder and plenty of scope for your second paper, clarifying the first.

5. If you are really stuck for a topic, you can always write a polemic against polemics.

Next AISBQ deadline 21 May 2004

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

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

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