AISB Quarterly

No.118

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

In this issue:

Features

Reviews

Society News

Father Hacker

Looking for logic in perceptual observations

In his seminal work[1] Bertrand Russell suggests that the *discovery* of what is really given in perceptual data is a process full of difficulty. We have taken this issue seriously in AI by applying a knowledge-discovery tool to build theories from perceptual observation. In particular, this article introduces some of the research we are conducting on autonomous learning rules of behaviour from audio-visual observation.

In our framework, classification models of perceptual objects are learned (in an unsupervised way) using statistical methods that enable a symbolic description of observed scenes to be created. Sequences of such descriptions are fed into an off-the-shelf inductive logic programming (ILP) system, called Progol,[2] whose task is, in this context, to construct theories about the perceptual observations. The theories thus constructed are further used by an automatic agent to interpret its environment and to act according to the protocols learned.

In its current implementation, our system is capable of learning the rules of simple table-top games (and how to use those rules) from the observation of people playing. Our basic motivating hypothesis for assuming game-playing scenarios is that these can provide both rich domains—allowing multiple concepts can be learned—and domains with gradually-increasing complexity, so that concepts can be learned incrementally. Moreover, it has been largely argued that games provide interesting ways of modelling social interaction.[3]

The experimental setting used in this research is composed of two video cameras: one observing a table-top where a game is taking place, and another pointed at one of the players. This player will also have a microphone recording utterances that he produces when playing the game. The purpose of the second camera is to capture the facial movements of the player, whose voice is being recorded, so that a synthetic agent can reproduce them in similar situations. A schematic of the system is shown in Figure 1.

The vision system consists of a spatio-temporal attention mechanism and an object classifier. Classes obtained from vision data are used to provide a symbolic description of states of the objects observed on the table top. This is used as input data for Progol.

In particular, our interests in this research are twofold. The first aim of this project was the autonomous learning of perceptual categories from continuous data, grounding them into meaningful (symbolic) theories. The second (but no less important) aim is the autonomous discovery of simple mathematical rules from the perceptual observation of games. We shall consider these goals in turn.

Grounding symbols to the world

In the current guise of this project, symbolic data provided by the audio and vision systems are input to the knowledge discovery engine as atomic formulae. Within these, symbols for utterances are arguments of predicates representing *actions* in the world whereas symbols for visual objects compose atomic statements representing the *state* of the world. Both statements are time-stamped with the time point relative to when they were recorded. A relation *successor* connects two subsequent time-points. The task here is to use Progol to *discover* the relationship between the utterances produced by one of the players, and the objects played on the table. Therefore, we are interested in the autonomous learning of the connection of audio and visual

Paulo Santos, University of Leeds

Continued on p. 8

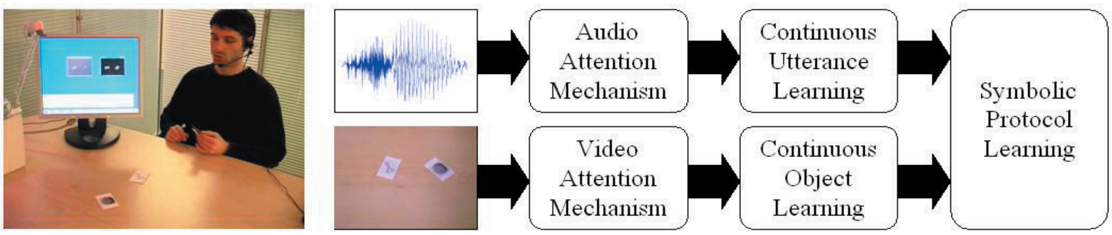

Figure 1. A schematic of the experimental set-up that allows machine learning of game rules through observation.

The rise of the meme machines

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

ISSN 0268-4179

Memes have had a bad press. Since Dawkins coined the term in 1976,[1] they have been called "a meaningless metaphor," an "empty analogy," and a "superstitious notion". Critics claim there is no evidence that memes even exist. Yet these objections may be based more on caricature than Dawkins' original idea.

The idea is this: whenever there is information that is copied with variation and selection, the information is a replicator and the principles of universal Darwinism apply. So, claimed Dawkins, the gene is not the only replicator on this planet. All around us, in its primeval soup of culture, information is being copied by imitation, teaching, and reading: it is rich in variation and selected by the limits on human memory, time, and resources. He called this information memes.

From this you can immediately see that the basis of memetics is not "let's divide culture up into bits and make an analogy with genes," but to treat the information we copy as a replicator. With this view, analogies between memes and genes may be worth exploring, but they are unlikely to be close because genes and memes are copied by such very different processes. To object that memes have never been proved to exist is to miss the point because, if people can imitate, then whatever they imitate is, by definition, a meme. So the very words you are reading now are memes. The question is not whether memes exist, but whether meme theory can do any useful scientific work.

By highlighting the significance of imitation, memetics provides a new way of understanding human evolution. Imitation is often taken for granted because it comes so naturally to us. Yet imitation, as anyone who has tried to build an imitating robot knows, is actually computationally extremely challenging. This is probably why our species is more or less alone in having the capacity. Some songbirds, whales and dolphins can imitate sounds, and chimpanzees and orangutans can, arguably, copy some actions, but only humans spontaneously and enthusiastically imitate from an early age and take pleasure in doing so.[2]

A recent surge of research on imitation has revealed some interesting comparisons.[2] For example, when observing others carrying out actions, chimpanzees are more likely to copy the goal or outcome, such as opening a box, while children are more likely to copy the precise movements made, even if this makes them less efficient at getting the reward. This may seem odd from a biological point of view, but it makes sense for a system that is evolving for the benefit of selfish memes rather than genes.

Here we get to the crux. From a memes'-eye view, human beings are meme machines that have been designed by the interaction of two competing replicators: genes and memes. The enormous human brain was not designed primarily for genetic benefit (it's a terrible liability in childbirth for example) but was forced on the genes by the demands of memetic evolution. And language, that peculiarly human trait, was similarly created by and for the memes, not the genes. I have called this process memetic drive[3]–the idea being that once early hominids were capable of imitation, memetic evolution would have taken off, changing the selection pressure on genes by favouring people capable of imitating the currently successful memes. In this way, the direction taken by memetic evolution would have driven genes to create the machinery necessary for copying those memes. The resulting copying machine is us.

There are many implications of this for artificial systems. One is that no humanoid robot, whether it is Cog, Kismet, or a household servant, is going to be remotely human unless it is a first-class imitator. Another is that it should be possible to simulate the evolution of language by building imitating robots. When I first suggested this[3] I thought it was extremely rash, especially given the long disputes over the origins of language. Yet people were already beginning to build robots that could imitate sounds, and they found that reference—and indeed entirely novel languages—can emerge this way.[4] The shift needed here is from thinking of language as a form of communication that serves the genes, to a process of copying that lets loose a new replicator.

Finally, memetics makes it clear that the photocopiers, phones, fax machines, and computers that we think we created for ourselves, are really meme machines created by and for the memes. They are the inevitable products of a co-evolutionary process. If we don't understand this we are likely to under-estimate the power of this accelerating process, and misjudge our own and our artefacts' role in its evolution.

Susan Blackmore
E-mail: susan.blackmore@blueyonder.co.uk
http://www.susanblackmore.co.uk and
http://www.memetics.com

References

1. R. Dawkins, **The Selfish Gene,** Oxford University Press, 1976.
2. S. Hurley and N. Chater (Eds), **Perspectives on Imitation: From Mirror Neurons to Memes,** MIT Press, 2004.
3. S. Blackmore, **The Meme Machine,** Oxford University Press, 1999.
4. L. Steels et al., Crucial factors in the origins of word-meaning, in Wray, A., et. al. (eds.), **The Transition to Language,** Oxford University Press, 2002.

Reifying building blocks in evolutionary algorithms

In recent years we have been witness to the application of bio-inspired algorithms for the solution of a plethora of hard problems in computer science, engineering, biology, and so forth.[1] Arguably the most popular of these bio-inspired methodologies is the evolutionary algorithm. Based on—and inspired by—the workings of evolution by natural selection, the basic meta-algorithm is seductively (and, I might add, deceptively) simple, and can be expressed in a mere eight lines of pseudocode:

1. produce an initial **population** of individuals, these latter being candidate solutions to the problem at hand
2. **evaluate** the fitness of each individual in accordance with the problem whose solution is sought
3. *while* termination condition not met *do*
4. **select** fitter individuals for reproduction
5. **recombine (crossover)** individuals
6. **mutate** individuals
7. **evaluate** fitness of modified individuals
8. *end while*

Over the past two decades, evolutionary algorithms have proven their worth beyond a doubt: at times not only solving hard problems but indeed competing with their human designers.[2] A major reason for the success of this seemingly blind search is its not being blind at all: though randomness and probability do play a part, selection is a function of fitness, and the recombination of good sub-pieces—or building-blocks—from two or more individuals enables the creation of offspring that outperform their parents (see Figure 1).

Building blocks—those good (or less good) sub-parts of individuals from which increasingly good solutions may be constructed—are part and parcel of evolutionary algorithm theory, explaining why these algorithms function and emphasizing the need for good genomic representations. Yet building blocks are usually relegated to playing an implicit role.

Borrowing another idea from nature—that of co-evolution—my colleague Assaf Zaritsky and I recently brought the building blocks to the fore, thus enhancing the performance of evolutionary algorithms in general. Co-evolution is the simultaneous evolution of two or more species with coupled fitness. Such coupled evolution favours the discovery of complex solutions when they are required. Simplistically speaking, one can say that co-evolving species either compete (e.g., predator-prey) or cooperate (e.g., symbiosis).

Applying cooperative co-evolution, we created two populations: the first containing candidate solutions (SO), and the second containing building blocks (BB). Fitness evaluation, selection, and mutation of individuals in the SO population are done as in the standard evolutionary algorithm. Crossover, however, is performed quite differently, in a more 'intelligent' manner: rather than choosing at random sub-pieces to exchange, these latter are chosen by referring to the BB population. Essentially, recombination will aim to conserve the—now known—good building blocks, and destroy only the bad ones (see Figure 2).

Starting from short chunks (two consecutive bits in the example of Figure 1) the BB population assembles better and longer blocks through evolution. This is done by assigning fitness to building blocks through referral to the SO population: blocks that appear in good solutions are assigned a higher fitness value than those that appear in bad solutions. The assembly of better and better building blocks is reminiscent of the way one assembles a puzzle, hence we dubbed our approach the *puzzle algorithm*.[3]

We showed that the co-evolution of building blocks and solutions is a powerful idea by applying our algorithm to a hard problem known as the *shortest common*

Moshe Sipper, Ben-Gurion University

Continued on p. 6

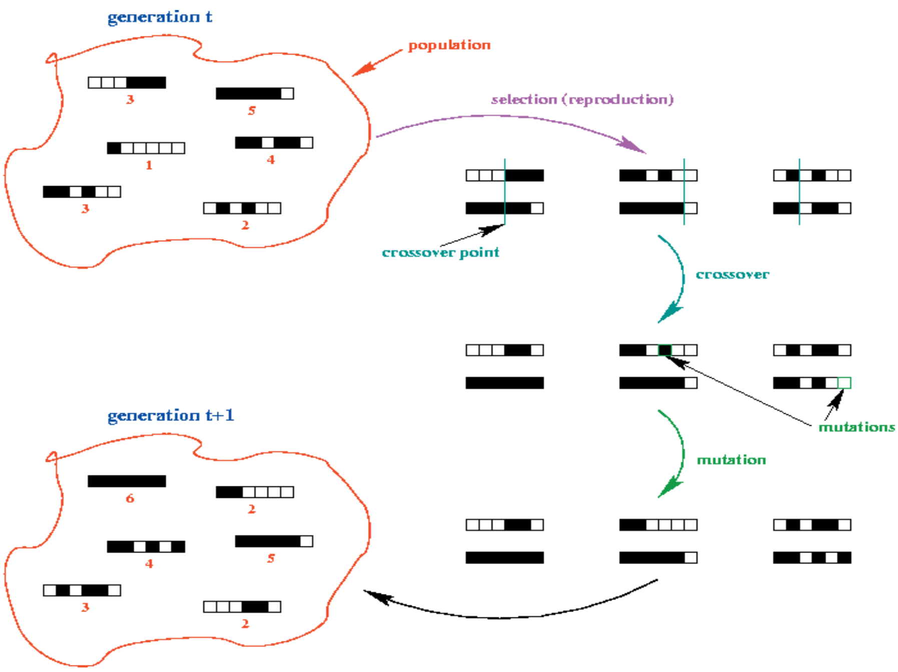

Figure 1: One generation (evaluation-selection-crossover-mutation cycle) of a simple evolutionary algorithm. Individuals are represented as bit strings (the so-called genomes). Fitness in this toy example is the number of bits with value 1 (shown in black). Selection is performed proportionately to fitness, so that high-fitness individuals are more likely to get selected. Recombination—or crossover—is performed on two individuals by selecting a crossover point at random and exchanging the chunks beyond this point. Mutation is performed by flipping a small number of bits (with low probability). As can be seen, this simple evolutionary scenario produced a perfect solution by the next generation: in this case due to the crossover operation, which has glued a good sub-piece from one individual (three 1s on the right) with another good sub-piece from the second individual (three 1s on the left). These sub-pieces are known as building blocks.

Connectionist modelling of the development of analogical completion

Connectionist modelling has become the foremost computational method for investigating mechanisms of learning and development in early childhood and infancy.[1] These are cognitive models loosely based on neural information processing. Learning consists of the gradual adjustment of connection weights in a network made of simplified neurone-like processing units. Although each unit is very simple, the network as a whole can mimic very complex patterns of behaviours. Connectionist networks are ideal for modelling development because they gradually develop task-appropriate internal representations through interactions with the outside world. In addition, they show graded levels of response—as do infants and children—on many different tasks.

Analogical reasoning is one area where a connectionist approach can provide insights into possible mechanisms of development. An example of an analogy might be: puppy is to dog as kitten is to cat. The interesting similarity is not between the attributes of the objects composing the analogy but between how puppy relates to dog and how kitten relates to cat. Consequently, analogy is often defined as involving seeing the relational similarity between domains.

Traditional approaches to modelling analogical reasoning have tended to take complex analogies as their starting point. Such analogies include comparing heat and water flow or comparing the structures of an atom and a solar system. These high-level, complex analogies have naturally emphasised structured representations (e.g. predicate logic) and structure mapping between a base and a target domain.[2]

Our alternative developmental approach highlights fundamentally different aspects of analogical reasoning.[3] Explicit, structured representations are seen as inflexible and very difficult to learn. Instead, of central importance is how analogy relates to the acquisition of knowledge and how mechanisms of analogical reasoning merge with other cognitive processes. Unsurprisingly, the earliest examples of analogical reasoning demonstrated by children are quite simple (e.g. pictorial analogies such as `apple' is to `cut apple' as `bread' is to `cut bread'). Consequently, our work takes these simple analogies as its starting point.

We propose that these simple examples of analogical reasoning can be explained as a by-product of normal memory processes: in particular, priming. Initial exposure to a situation primes a relation which can then be applied to a novel situation to make an analogy. Relations are represented as transformations from one state to another (e.g. the relation cutting can be understood as involving a transformation in an object from `apple' to `cut apple'.) This has the advantage that relations do not have to be represented explicitly, avoiding the attendant difficulties of that approach. In this account, more complex analogical reasoning (e.g. explicit structure mapping) is viewed as a meta-cognitive skill. For example, analogies involving systems of objects may be built up by repeatedly applying a much simpler process in a controlled way.

This framework has been implemented in a recurrent network architecture trained with contrastive Hebbian learning (see Figure 1). The latter is a supervised training algorithm involving two phases of activation: a minus phase where only the input layers are clamped on, and a plus phase where all external units are clamped on. The change in connection weights is calculated locally as the difference between a Hebbian term for the plus phase and an anti-Hebbian term for the minus phase. Through training, the network develops appropriate attractor states corresponding to desired states of activation. Priming is achieved through the recurrent connections maintaining prior activation. Analogical completion occurs through a combination of priming and the pattern-completion abilities of this type of network.

A wide range of behaviours consistent with the developmental literature can be demonstrated the network. There is a shift in responding—from initially being based on object features to relational features—consistent with observations on children.[4] There is no general analogical reasoning ability that suddenly comes online at one point in training. Instead the network can complete individual analogies when it has sufficient knowledge of the underlying relations. Again this is consistent with developmental studies.[5] Furthermore, the network is never trained on analogy: instead analogical completion emerges from the way relational information is represented and tested. This is con-

Leech & Mareschal, Birkbeck

Continued on p. 8

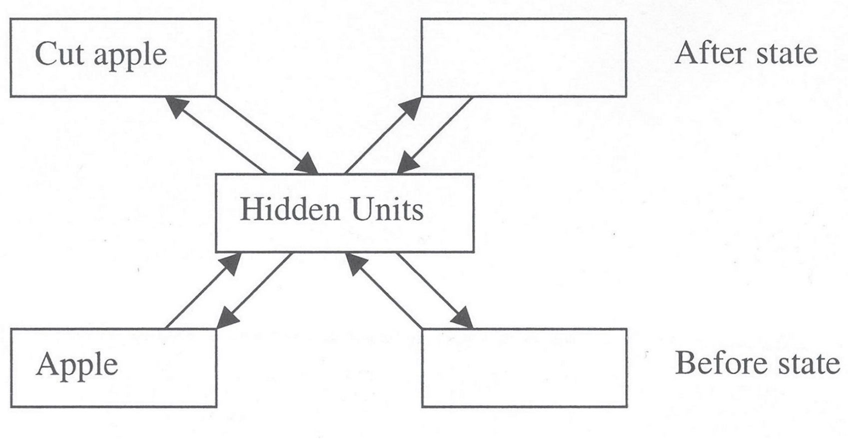

Figure 1: An illustration of the network architecture and how relations are implemented as transformations. The input layer corresponds to the `before' state of a relation (e.g. apple) and the output layer corresponds to the `after' state (e.g. cut apple).

Unweaving the analogical rainbow with lightweight lexical ontologies

Analogical reasoning is a decidedly knowledge-hungry faculty. It is, after all, one of the foremost cognitive tools we possess for shedding light on a poorly-understood domain by importing the structure of one more-clearly understood. The approach to analogical reasoning most familiar to AI researchers will undoubtedly be the structure-mapping approach, first suggested by Patrick Winston[1] and Dedre Gentner[2] and given an algorithmic realization in SME, the *structure-mapping engine*.[3] These and other models that operate on similar principles (such as my own *Sapper* model)[4] presuppose that analogy operates by systematically projecting the causal propositional structure of one domain onto another. In effect, structure-mapping can be viewed in mathematical terms as a variant of the well-understood—but NP-complete—problem of finding the largest isomorphic sub-graph of two representations. Variants of structure-mapping mainly differ in the sub-optimal choices they make to achieve tractable performance.

For researchers like myself, whose first exposure to analogy was via structure-mapping, the 1990s were heady years in which competing models of analogical mapping were pitted against each other on specially-crafted domain descriptions of Aesopian fables and Shakespearean plots. Indeed, so spirited was this competition that some have referred to the whole enterprise as the 'analogy wars'. However, as the name suggests, structure-mapping is vexingly dependent on the availability of explicitly-structured domain descriptions, and this dependency means that analogical research in this period relied for the most part on hand-coded representations.

More recently, my group has been attempting to implement robust and scalable models of analogy, both interpretative and generative, that rely instead on large-scale representations from third-party sources. This has lead us to consider a number of possible knowledge-sources, from Cyc[5] to WordNet.[6] Indeed, the quest for large-scale structured resources that were independent of their analogical uses lead me in 1999 to Cycorp Inc. of Austin, Texas, where I spent a year applying structure-matching ideas to the propositions and axioms stored in the Cyc knowledge-base. Cyc is a heavy-duty ontology with extensive cross-linking between concepts, yet I was forced to conclude that there is far too much structural variation between the descriptions of different domains—usually entered by different engineers—to make it a viable knowledge-base for structure-mapping purposes. This realization has lead my group to look instead to freely available, if flawed, light-weight ontologies like Princeton WordNet[6] (PWN), to develop more scalable and less structure-dependent approaches.

Our work reveals PWN to be a sufficiently-rich basis for processing lexical analogies, such as those found on scholastic aptitude tests[7] (or SATs). For example, "Doubloon is to coin as what is to ship?" The answer is galleon, since a doubloon is a *Spanish* coin and a galleon is a *Spanish* ship. PWN can be used to understand and generate analogies like these by unlocking the implicit references contained in the textual glosses that annotate each WordNet sense entry. Certain gloss terms will be shared in common between lexical analogues (e.g., as *Spanish* is shared by galleon and coin) while others will be domain-shifted (e.g., *coin/ship* in the above analogy, or *spacecraft/airplane* in the glosses of astronaut and pilot). Algorithmically recognizing which gloss terms serve which role is the essence of lexical analogy in PWN.

The best lexical analogies thus involve a combination of overt similarity tempered by constrained difference, the latter juxtaposing terms that reside within the same semantic field. For example, Table 1 summarizes the analogical mappings that can be resolved using PWN in the taxonomic domain of deities.

Our work shows that analogy is a powerful retrieval tool that allows users of PWN to locate concepts not just via synonymy, but through complex allusions. For example, "Muslim bible" can be used to retrieve Qu'ran while "Jewish alpha" retrieves the letter Aleph. Our goal is the re-invention of the humble thesaurus as an actively creative resource, essentially an *analogical thesaurus* capable of understanding a user's allusions and even generating creative allusions of its own. This expressive power also fuels our current interest in generating analogical puzzles and riddles for use in both computer games and scholastic tests.[7]

Difference / *Commonality*	*Greek*	*Roman*	*Hindu*	*Norse*	*Celtic*
Supreme	Zeus	Jove	Varuna	Odin	N/A
Wisdom	Athena	Minerva	Ganesh	N/A	Brigit
Beauty, love	Aphrodite	Venus	Kama	Freyja	Arianrhod
Sea	Poseidon	Neptune	N/A	N/A	Ler
Fertility	Dionysus	Ops	N/A	Freyr	Brigit
Queen	Hera	Juno	Aditi	Hela	Ana
War	Ares	Mars	Skanda	Tyr	Morrigan
Hearth	Hestia	Vesta	Agni	N/A	Brigit
Moon	Artemis	Diana	Aditi	N/A	N/A
Sun	Apollo	Apollo	Rahu	N/A	Lug

Table 1: Summary of the analogical mappings that can be resolved using Princeton WordNet in the taxonomic domain of deities.

Tony Veale
Dept. of Computer Science
University College Dublin, Ireland
http://www.cs.ucd.ie

References

1. P. H. Winston, *Learning and Reasoning by Analogy,* **Communications of the Association for Computing Machinery 23** (12), 1980.
2. D. Gentner, *Structure-mapping: A theoretical framework for analogy,* **Cognitive Science 7,** pp. 155-170, 1983.
3. B. Falkenhainer, K. Forbus and D. Gentner, *Structure-Mapping Engine: Algorithm and Examples,* **Artificial Intelligence 41,** pp. 1-63, 1989.
4. T. Veale, and M. T. Keane, *The Competence of Sub-Optimal Structure Mapping on Hard Analogies,* **Proc. IJCAI'97,** 1997.
5. D. Lenat and R. V. Guha, **Building Large Knowledge-based systems,** Addison Wesley, 1990.
6. G. A. Miller, *WordNet: A Lexical Database for English,* **Communications of the ACM, 38**-11, 1995.
7. T. Veale, *WordNet sits the S.A.T.: A Knowledge-based approach to Lexical Analogy,* **Proc. ECAI'2004,** 2004.

The Intelligent Media Institute: A chance to broaden our horizons

The convergence of computing, networks, and media—fuelled by the 'digital revolution'—is starting to collide with the convergence of technology and biology. Out of this is emerging something entirely new and of profound importance: intelligent media. Early in 2003, University College London's (UCL's) Professor Philip Treleaven and Microsoft's Dr Stephen Emmott came together with one aim: to establish a major new institute focused on research, teaching, and enterprise called the Intelligent Media Institute (IMI). Their motivation comes from the belief that intelligent media will almost certainly create an enormous impact on our relationships with technology, with each other, and with the things around us that define our culture and society: from television to toys, film to fashion, product design to packaging, architecture to advertising, cities to commerce. It will also create equally-significant new innovation and market opportunities in the sectors that are at the centre of all these areas: the creative industries.

Treleaven and his colleagues define intelligent media as any physical or digital medium (e.g. software, fabrics, music, video, television, print, film, building materials, paper, paint, content) that contains or exhibits some computationally-derived or biologically-inspired intelligent behaviour. These could include the artefact's ability to learn about, adapt to, or communicate or interact with, its environment. An important additional constraint is that this ability should not currently exist in that medium.

The geographical hub for the IMI is the capital, the creative industries being absolutely vital to the economic prosperity of London, adding £28 billion annually to the city's output, and employing more than 500,000 people.[1] Consequently, Treleaven and Emmott have gathered top academics from UCL, Imperial College London, University of the Arts London, King's College London, Goldsmiths College, Queen Mary College, City University, and The Royal College of Music, as well as key industry figures, to gauge the extent and level of support for the concept. In return, the Institute has received around 13 project proposals from research clusters composing of researchers from computer-science, AI, mathematics, materials science, chemistry, physics, product design, art, neuroscience, and electrical engineering. These projects come under the following themes: music; entertainment (television, film, games); intelligent cultural heritage; interaction and design; and retail and advertising. Project examples span from musical-performance-enhancement tools to virtual clones, digital museums to 3D goggles, smart fabrics and fashion to intelligent and aware buildings.

The establishment of the IMI is now in phase two. After receiving potential proposals, Treleaven and his colleagues are in the process of producing the business case and research agenda and are visiting key industries and government bodies to secure initial financing. Thus far, they have been successful in obtaining £1.5 million in initial funding from the LDA (the London Development Agency) and will pre-launch with a workshop in September.

London now has a major opportunity to benefit from developing and exploiting these opportunities, and to lead the world in pioneering research and innovation in intelligent media through the creation of an important new inter-disciplinary research institute that uniquely brings together and unites science, technology, creativity, and enterprise. Likewise, those of us working in narrowly-defined research programmes can get the chance to work on more diverse and interesting projects: broadening our horizons and, potentially, initiating a fundamental paradigm shift that actually matters.

For further information please contact Prof Philip Treleaven via p.treleaven@cs.ucl.ac.uk.

Dr Julie McCann
Department of Computing
Imperial College London
http://www.doc.ic.ac.uk/~jamm

References

1. Greater London Authority, *Creativity: London's Core Business,* **Greater London Authority Report,** 2003

Reifying building blocks in evolutionary algorithms

Continued from p. 3

superstring. This involves finding the shortest string that covers all substrings in a given set, and has important applications in DNA sequencing. The puzzle algorithm blithely outperformed its competitors.

Moshe Sipper
Department of Computer Science, Ben-Gurion University, Israel
E-mail: sipper@cs.bgu.ac.il, http://www.moshesipper.com

References

1. M. Sipper, **Machine Nature: The Coming Age of Bio-Inspired Computing,** McGraw-Hill, New York, 2002.
2. J. Koza, *Genetic Programming: Routine, Human-Competitive, High Return Machine Intelligence,* **AISB Quarterly 114,** Autumn 2003.
3. A. Zaritsky and M. Sipper, *The Preservation of Favoured Building Blocks in the Struggle for Life: The Puzzle Algorithm,* **IEEE Transactions on Evolutionary Computation,** 2004, (in press).

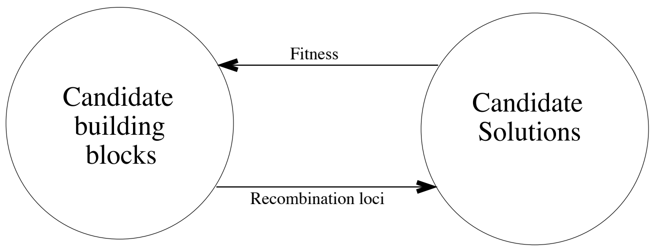

Figure 2: The puzzle algorithm's general architecture involves two co-evolving species (populations): solutions and building blocks. The fitness of an individual in the population of building-blocks depends on individuals from the solution population. The choice of recombination (crossover) loci in the solution species is driven by individuals from the building-block population.

Art and artificial intelligence

There is a new spirit of inter-disciplinarity at Goldsmiths College. Among its manifestations is the advent of two interlocking research centres in which the Department of Computing is playing a major role. The first of these is the centre for *Cognition, Computation, and Culture*, which includes work on both brain and mind modelling. The second is called *Goldsmiths Digital Studios* and is dedicated to interactions between arts practice and digital technology. This centre has recently spread and it now encompasses a second site at Martlesham, where we are working closely with British Telecom to develop research on the boundary of arts and technology.

The Department of Computing is deeply committed to this new vision and we are rapidly developing research that spans cognitive sciences, cultural studies and artistic practices, including music, though in this article I will focus on visual arts and interactive film. A feature of all of the arts-based research is that it includes serious practicing artists working with computer scientists, and, consequently, the projects all have integral deliverable results in both art and computer science. For example, we are working with the artist Warren Neidich on a project, sponsored by the Arts Council and Arts and Humanities Research Board, to make both dynamic web-based and three-dimensional art works out of brain-process-inspired computation. The project takes as its starting point time-lapsed photographs of the hands of deaf people signing a conversation. Neidich calls this pictures these pictures *conversation maps*, an example of which is shown in Figure 1.

Figure 1: A Warren Neidich conversation map: time-lapse photograph of a deaf person

The maps will evolve with interaction from users, following various algorithms that simulate theories of brain function. As they evolve, they change as art works on the screen. Periodically, the trace of the evolving work will, through the use of a three-dimensional printer, generate a solid art work. As well as the art works, new techniques and insights in both information visualisation and neural computing will be developed through this project

A second line of research that similarly explores relationships between art, information representation, and human intelligence is centred on the use painters make of abstraction as a way of compactly imparting information. We are studying this as a way of understanding human perception and visual representation. See Reference 1 for a brief description of the framework. We will then use this understanding to help us to design computer-based systems. One of the goals is to apply the results of this study to questions concerning abstraction and computer representations for AI tasks.[2]

We have discussed projects that involve web-based art, three-dimensional sculpture, and painting. For the rest of this paper, we turn our attention to three related projects that are based on digital film. A common technical theme throughout the three projects is that films are built out of segments of media. In every one of the projects, larger objects—whole films, interactive television, or tours through an archive—are derived at viewing time. The largest of these projects NM2 (*New Media for the New Millennium*), which is about to start as an European-Union-funded Sixth Framework project, has twin missions. The first is the development of a large tool set that supports sophisticated production and viewing of non-linear multi-media productions, and the second the development of several prototype non-linear productions. NM2 covers all facets of the technical and practice questions in the system design: from image understanding to software architectures to questions of non-linear narrative. The part for which we are primarily responsible involves defining a language, reasoning techniques and tools that will enable the system to express both the visual narratives and expertise associated with authoring them. Descriptions in the language will then be used, for example, to reason about what non-linear routes through the material have what affects.

In a related project, we are working with the Tate Museums to build a system with which the Tate's archive of long-streaming media objects can be taken apart by the curators and put together again by curators or end-users to create individual compositions that follow different themes, ideas, people, and levels of engagement. And, finally, we are working with Andrew Shoben, an artist, to produce generative films. These are films that will reconfigure themselves by randomised algorithms that are dependent on external factors. In the film projects, interesting questions of representation and constraint reasoning arise and the computer science and productions will be developed hand-in-hand.

In all of these projects, artists and technologists are working closely and new kinds of art and computing are being developed. This is the multi-disciplinary vision that is behind the work of *Goldsmiths Digital Studios*.

Robert Zimmer
Goldsmiths College
University of London
E-mail: R.Zimmer@gold.ac.uk

References

1. R. Zimmer, *Abstraction in art with implications for perception*, **Philos. Trans. R. Soc. Lond. B. Biol. Sci. 358** (1435), pp. 1285-91, 2003.
2. R. C Holte, T. Mkadmi, R. M. Zimmer, A. J. MacDonald, *Speeding up problem solving by abstraction: a graph oriented approach,* **Artificial Intelligence 85,** pp. 321-261, 1996.

Looking for Logic in Perceptual Observations

Continued from p.1

objects within a particular context.

Some preliminary results of this research, presented in Reference 4, show that the system could learn accurate descriptions of simple table top games.

Learning simple axioms from observation

We next submitted the system to the challenge of finding the transitivity, reflexivity, and symmetry axioms from the observation of games without assuming any preconceived notion of number or any pseudo definition of ordering. This research is described in References 5 and 6, where these axioms are obtained from noisy data with the help of a new ensemble algorithm that combines the results of multiple Progol processes by means of a ranking method.

It is worth pointing out that the discovery of simple mathematical axioms from observation, without assuming any explicit background knowledge, was not set for pure intellectual pleasure. In fact, I believe that the capacity of abstracting general truth from simple data is essential if any system is to solve new problems and to overcome obstacles not seen before.

Conclusion

The final aim of the project summarized above is to build a machine capable of learning (and further executing) human tasks by observing people accomplishing them. The preliminary results of this research suggest that the combination of statistical methods with inductive logic programming, applied to the task of learning behavior by observation, is a promising route towards our final goal.

Chris Needham, Derek Magee, Anthony Cohn, and David Hogg are all active participants in the development of the project described above. Thanks to Brandon Bennett and Aphrodite Galata for discussions and suggestions. (The Editor would also like to thank SPIE, The International Society for Optical Engineering, for permission to print this article, commissioned first for the newsletter of their Technical Group on Robotics and Machine Perception: find them at www.spie.org).

Paulo Santos
CogVis Group
School of Computing
University of Leeds, UK
http://www.computingleeds.ac.uk/
E-mail: psantos@comp.leeds.ac.uk

References

1. B. Russell. **Our Knowledge of the External World,** The Open Court Publishing Company, 1914. Reprinted in 2000 by Routledge.
2. S. Muggleton. *Inverse entailment and Progol,* **New Generation Computing, Special issue on Inductive Logic Programming, 13** (3-4), pp. 245-286, 1995.
3. S. P. H. Heap and Y. Varoufakis. **Game Theory: a critical introduction,** Routledge, 1995.
4. D. Magee, C. J. Needham, P. Santos, A. G. Cohn, and D. C. Hogg. *Autonomous learning for a cognitive agent using continuous models and inductive logic programming from audio-visual input,* **Proc. AAAI Workshop on Anchoring Symbols to Sensor Data,** 2004.
5. P. Santos, D. Magee, and A. Cohn. *Looking for logic in vision,* **Proc. 11th Workshop on Automated Reasoning,** pp. 61-62, Leeds, 2004.
6. P. Santos, D. Magee, A. Cohn, and D. Hogg. *Combining multiple answers for learning mathematical structures from visual observation,* **Proc. European Conf. on Artificial Intelligence (ECAI-04),** Valencia, Spain, 2004.

Connectionist modelling of the development of analogical completion

Continued from p.4

sistent with observations of spontaneous analogical reasoning in young children.[6]

The impressive range of developmental phenomena that this network can simulate suggests the appropriateness of this approach for investigating a domain (analogical reasoning) that has normally been considered the preserve of high-level explicitly-structured models. Furthermore, our approach also allows us to consider the key cognitive mechanisms that underlie the developmental transitions that occur in children's emerging ability to reason by analogy. The framework suggests how learning interacts with the development of representations and how this can affect behaviour.

Robert Leech and Denis Mareschal
School of Psychology
Birkbeck, University of London
http://www.psyc.bbk.ac.uk/

References

1. Y. Munakata and J. C. McClelland, *Connectionist models of development*, **Developmental Science 6,** pp. 413-429, 2003.
2. D. Gentner, *Structure-mapping: A theoretical framework for analogy.* **Cognitive Science 7,** pp. 155-170, 1983.
3. R. Leech, D. Mareschal, and R. P. Cooper, *A connectionist account of analogical development,* in R. Alterman and D. Kirsh (eds.), **Proc. 25th Int'l Conference of the Cognitive Science Society,** pp. 710-715, 2003.
4. D. Gentner, *Metaphor as structure-mapping: The relational shift,* **Child Development 59,** pp. 147-59, 1988.
5. U. Goswami and A.L. Brown, *Melting chocolate and melting snowmen: Analogical reasoning and causal relations,* **Cognition 35,** pp. 69-95, 1989.
6. S. Pauen and F. Wilkening, *Children's analogical reasoning about natural phenomena,* **J. of Experimental Child Psychology 67,** pp. 90-114, 1996.

Let us fund your travel

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

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

Note: those given travel awards are required to write a review of the conference they attend for the AISBQ.

CONFERENCE REVIEW

16th European Conference on Artificial Intelligence
Valencia, Spain, 22 - 27 August 2004

The Sixteenth European Conference on Artificial Intelligence (ECAI) attracted a large and varied audience this year, indicating a growth of interest in AI worldwide, and the promising emergence of high-quality AI research in the traditionally somewhat-less-involved continents. The distribution of accepted papers over the various subfields of AI indicates that knowledge representation and reasoning, constraint satisfaction, multi-agent systems, and machine learning are currently of most interest to this AI community, while less attention is being devoted to case-based reasoning, cognitive modeling, neural networks, philosophical foundations, planning, and robotics.

A topic that received considerable attention is that of dealing with a multitude of possibly-heterogeneous loosely-related knowledge fragments instead of one single homogeneous knowledge base. To my mind, this is of great importance for the development of AI in its most general sense. Very much related to this is the issue of focusing the attention of AI programs on that aspect of their knowledge base which is relevant for the task they are currently performing. The necessity to do this has long been argued for in the philosophical literature, but now becomes more and more evident from an application-oriented perspective as well.

For instance, the first invited talk, given by Christian Freksa, sketched an AI perspective on certain issues in the field of spatial reasoning. He clarified the notion of focusing one's (spatial) attention, by means of an orienteering example: imagine you are trying to locate a friend's house in a city you are not familiar with using an ordinary street map. You never consider all the information provided by the map, but rather focus on the relevant part of it only. You may adapt your focal viewpoint now and then in one of three ways: by looking at another part of the map (panning), looking at the map from a different angle (turning), or looking at the map at a different level of granularity (zooming). These three seem to correspond to fundamental dimensions along which one could adapt one's focus: in a more general setting we could refer to these dimensions as partiality, perspective, and approximation, respectively. An excellent paper has been published recently on this.[1]

Invited speaker Gloriana Davenport highlighted results and problems encountered in automating story generation from personal (multi-)media collections. To perform this task, computer programs need to have a sufficient understanding of both the content of multi-media fragments, and the nature of story-construction mechanisms. Story scenes and the relations between them can often be understood from different perspectives and, although people mostly seem to automatically address a suitable one, this is not a trivial task for an artificial story maker.

Take a digital image of a father giving a full-size guitar to a young boy, who might be his son. This scene can be seen as a representation of an event in which possession is transferred. In the light of such an interpretation, subsequent scenes could elaborate on what the boy may do with the guitar now that he owns it (play on it, take it on a trip, give it to someone else). On the other hand, the same scene may be regarded as a purely physical act: a big and heavy object is put into the hands of a young boy. Subsequent scenes may tell how the boy cannot carry the object and drops it on the floor. In fact, to understand how his father may react to this, a story maker needs to keep in mind both the possession transfer and the physical perspective. These ideas go back to Marvin Minsky's seminal and ever inspiring work on multiple selves and realms of thought.[2]

The last invited talk was given by Carole Goble, who pointed out that the enormous research efforts currently devoted to establishing the grid on the one hand and the semantic web on the other, could both highly benefit from greater interaction between their associated communities. Again, both the organization of online ontologies—which are to serve as the basis of the semantic web—and the establishment of large-scale data grids call for a sophisticated understanding of how to deal with distributed, heterogeneous information.

This point was actually taken somewhat further in some of the technical sessions of the conference. Maybe most significantly, Chiara Ghidini and Fausto Giunchiglia[3] argued that abstraction (cf. Freksa's zooming) may well be the most important relationship between two different representations in this setting. They proposed a formal semantics for abstraction, which boils down to a slightly restricted first-order version of the local-model semantics, which was originally devised for a contextual knowledge-representation framework. In line with Goble's plea, Ghidini and Giunchiglia call for further work on the characterization of the notion of abstraction and of other fundamental kinds of mappings that may exist between different representations.

Floris Roelofsen

References

1. M. Benerecetti, P. Bouquet, and C. Ghidini, *Contextual Reasoning Distilled,* **J. of Experimental and Theoretical Artificial Intelligence (JETAI), 12** (3), 2000.

2. M. Minsky, **The Society of Mind,** Basic Books, 1985.

3. C. Ghidini and F. Giunchiglia, *A Semantics for Abstraction,* **Sixteenth European Conf. on Artificial Intelligence,** pp. 343-347, 2004.

Floris Roelofsen is based at the Institute for Logic, Language, and Computation in Amsterdam. His is interested in philosophical, logical, and computational perspectives on contextual knowledge representation and reasoning processes.

BOOK REVIEW

Models of a man: Essays in memory of Herbert A. Simon

Mie Augier and James G. March (eds.)

Publisher: *MIT Press.*
Hardback: *Published May 2004, 592pp, £29.95. ISBN: 0262012081*

This book—like Simon's work—ranges broadly across many topics, including economics, management, political science, cognitive science, and AI. The collection of 42 contributions also vary widely in style, from personal reminiscence and discussion of how Simon influenced colleagues and students, to more rigorous research papers. The picture that emerges from this collage clearly shows the wide influence of Simon's work, and his unique contribution to issues of decision-making as manifested in many different fields. Central to this was the belief that mathematics could be applied to seemingly vague issues in the study of human behaviour, and the development of the concept of 'satisficing' as an alternative to assumptions of optimal rationality.

Two biographical chapters introduce the book. The first of these, by Augier and March, provides an excellent summary of the development of Simon's ideas and the links between the apparently disparate topics he studied. Landmarks from his contribution to AI (mostly in association with Alan Newell) include the invention of the first AI program, the *logic theorist,* in 1955; soon followed by the *general problem solver;* publication of the *Sciences of the Artificial* in 1968; the articulation of the 'symbol-processing' view of intelligence in 1972; his work on models of scientific discovery; and the EPAM (elementary perceiver and memorizer) model of human verbal learning. He is also (in)famous for his highly-optimistic predictions for the field: e.g. in 1958 he said, "There are now in the world machines that think, that learn, that create…in a visible future—the range of problems they can handle will be co-extensive with the range to which the human mind has been applied."

The remainder of the book is in four sections which cover ideas in economics, public administration, systems modelling, and 'minds and machines'. Many of these chapters are personal views and anecdotes about Simon's influence. As such, they are largely of historical rather than scientific interest, although they provide some useful reflections on Simon's approach to problems. For example, Langley provides a set of 'heuristics for scientific discovery' as reflected in Simon's career, that include: be audacious, ignore discipline boundaries, use a secret weapon (that is, apply methods and metaphors that you have mastered but others in your field have not), balance theory and data, satisfice, and persevere. Consequently the volume contains much useful advice for any ambitious young scientist.

However, the more valuable articles in my view were those that took the opportunity to evaluate current understanding of some of the key issues introduced by Simon. For example, Gigerenzer provides a discussion of how 'bounded rationality' should be distinguished from 'optimization under constraints', and also from the study of systematic errors in human reasoning. Kahneman and Frederick discuss current views of intuitive judgement. Pitt reflects on Simon's view on models and what can be learnt from them despite their essential simplifications.

It is also interesting to have this overview of Simon's legacy in the light of recent opposition between 'symbol-processing' and 'behaviour-based' approaches to AI. While most of his work is clearly associated with the former camp, the latter often cite his analogy of the ant, which produces complex behaviour as a result of interaction with a complex environment rather than any inherent internal complexity. In fact, the issue of environmental interaction as a critical determining factor in cognition was woven throughout his approach, albeit that he studied this in the problem space of reasoning rather than sensorimotor control.

In summary, despite the disparity of style and topics of these essays, I found much of interest in this book. In particular, stepping back to see how a leader of our field tackled the central issues provided a useful perspective that is often hard to obtain in these times of highly specialised study.

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.

Secretary's Report

The results of the current AISB committee election were just in as we went to press. We are pleased to announce that John Barnden and Dimitar Kazakov won the election. John Barnden, our current Chair, is in the process of contacting the winners and we hope that they will accept their places on the committee.

In total, 37 people voted in the election which represents a turnout of roughly 10%. Of these, approximately one quarter failed to include the voter's name on the ballot paper. We realise this is an unusual request in an election but it is the simplest mechanism available to us to ensure that only society members vote and that no one can vote twice. After considerable deliberation the committee decided not to include these votes in the count. In future we intend to make this requirement clearer on ballot papers. Full details of the voting process and the affect of the spoilt papers is available on request from secretary@aisb.org.uk.

You will no doubt have noticed yet another call for nominations with this issue of the Quarterly. The society has failed to hold any elections since March 2002, as a result of which many of the current committee's terms of office are coming to an end. We hope from now on to issue a call for nominations for 3-4 committee places each Autumn, with an election around the new year, allowing new committee members to be introduced at the convention. Hopefully, therefore, this is the last call for nominations you will see until next year.

Louise Dennis
SSAISB Secretary

From the Chair

As this is my first statement in the Quarterly as SSAISB Chair, perhaps the first thing for me to do is emphasize that our committee is constantly seeking to amplify the Society's value to its members. If you have any ideas about how to do this I encourage you to contact me or someone else on the committee. In these days of viewing cognition as extending beyond the individual's head and into society, we should expect the SSAISB to bring its emergent, distributed intelligence to bear! It has also struck me that it would be beneficial if I, or other committee members, were to visit departments doing work in the technical areas we cover in order to discuss how the Society could develop. If you'd be interested in such a visit during 04/05, please contact me or another committee member.

One thing on my wish list is for there to be more enhanced (Patron and Supporting) members. The current number is small, to put it politely. I invite senior colleagues in the Society to consider enhancing their membership level. A single extra Patron membership could, for instance, go a long way to funding the attendance of an extra foreign researcher or a couple of extra students at the Convention each year.

Since starting as Chair in September 2003 I have been delighted by the commitment and energy of committee members: I think the committee is now at its most effective and efficient since I got involved about six years ago. This is in no small measure due to the previous Chairs' own energy, commitment, and forward thinking, most recently that of Geraint Wiggins. I am delighted also at the number of people who have come forward at and since our 2004 Convention as candidates for open positions.

I hope you will agree that this year's Convention at Leeds, organized by Kia Ng, was another success. I believe it shows the worth—in terms of innovation and distinctiveness—of our somewhat unusual Convention style, based on symposia and invited speakers within an overall loose theme. However, if you have views on what could be done in the future then, again, please contact us. We are also very open to new possibilities for connections between the Conventions and other conferences, workshops, and so forth.

As I write, we are in the run-up to the September 15 deadline for nominating members of the RAE 2008 discipline 'sub-panels' and lumped-discipline 'main panels'. The Committee is actively working on preparing nominations of various sorts, especially as the SSAISB is one of the learned societies that have been invited to make nominations. In our deliberations we will be liaising with certain other bodies. After nomination time we will be trying to do what we can to help the RAE 2008 process serve the interests of SSAISB members.

I am pleased that we have started a fruitful dialogue with the BCS special interest group on AI (SGAI), and collaborated with them on a careers day at City University. I look forward to stimulating and mutually-beneficial relationship, given the overlapping but healthily different remits and constituencies of the two organizations.

Finally, I'd like to mention here the goals for my chairmanship that I put into to my statement for the recent committee-member election process, except for the ones that are already implicit above. The goals are: to foster increased interdisciplinary links between AI and other fields, to strengthen the cognitive science flavour of SSAISB, to make the Society more visible to national policy makers, to expand the membership, to help some current moves of the Society into more concern with education, careers and public understanding, and to develop further links with other societies, publishers, and so forth. I would welcome any help or advice on achieving these goals.

Wishing you a happy and intelligently productive 2004/5!

John Barnden
Chair, SSAISB

Contacting the AISB

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

Areas of responsibility and e-mail addresses are available on the right column of this page (top), and via the web site. Find us at:

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

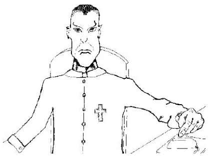

Father Hacker's Guide
for the Young AI Researcher

Cognitive Divinity Programme
Institute of Applied Epistemology

Organising a major international conference both demonstrates the esteem in which your peers hold you and illustrates your selfless dedication to serving your field. You must seek out this double opportunity, so it is essential to know...

12. How to Organise a Conference

1. If you wait to be *asked* to organise a major conference, it may never happen. Cut out the middleman by founding your own, new conference series and organising the first one. Not only will this enhance your international reputation, but it will also create a new publication outlet, with a high success rate, for you and your close colleagues. Recruit them all to help you. Reward them by accepting their papers for your conference.
2. There are three major organising roles in running a conference.
• Local Arrangements Chair: This involves a lot of low-level administrative work with little kudos. Avoid it at all costs.
• Programme Chair: This *does* carry a lot of kudos, but involves a lot of high-level organisational work and research judgement. This one is not for you either.
• Conference Chair: This is a prestigious, overseeing role, with a low workload, but high kudos. This is the job for you.
3. To ensure that your conference is well attended, you need to focus on three factors: location, location, and location. Poll your colleagues to identify their ideal holiday location and then book a top hotel there with stunning views and lots of local attractions. Note that you will need at least one pre-conference visit to check out the facilities, investigate the site-seeing opportunities, and sample the local cuisine.

4. To guarantee a coherent theme to your conference, make sure that the programme chair, the programme committee, and the referees share your view of the field: your ex-students, research fellows and close collaborators are ideal candidates. They will have no difficulty in recognising the high quality of each other's work, and rejecting any inharmonious intrusion from interlopers. In fact, so clear-cut are the assessments likely to be, you may decide that the refereeing process is redundant.

5. To stand out from the crowd and leave an abiding impression, a conference must be the battleground for a major academic fight. Don't leave this to chance. As protagonists, pick two self-opinionated egotists with diametrically-opposing views. Make sure they receive each other's papers to review. If their reports are insufficiently vitriolic, then sharpen them up yourself. Send the reports to the authors and accidentally forget to delete the referees' names. Tip off both protagonists that the other one is planning a hatchet job in their presentation. Schedule them to have adjacent presentation slots and leave plenty of time for discussion. Then sit back and watch the sparks fly. Don't forget to alert the press. Your conference will remain a fond memory for all but two attendees.

Next AISBQ deadline:
19 Nov 2004

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/qhome.html