Spielberg and Speculation

Included with this Quarterly is the first issue of AISB's new journal - AISBJ. This special issue on Agent Technology certainly does credit to the guest editors and production team. It is also an important addition to the value which AISB offers its members. As always, I'd like to hear what the members think.

On the subject of the journal - agent technology - I must also draw attention to the report of the Conference of Professors and Heads of Computing (CPHC) on agent-based systems. (Included in the News section). This is a call to action on this important area.

Elsewhere in this issue of the Quarterly, we have all the usual fare including plenty of news, views, and Father Hacker. There are some rather varied and sometimes outspoken opinions on AI: The Movie. What contribution this film makes to intelligent speculation about the future of our research remains to be seen, but I for one would welcome some speculation. (In the pages of this publication, naturally).

And on the subject of speculation, there's an interesting review of the C4 programme 'Battle of the Robots: The Search for AI', by John Jackson. This programme followed three prominent AI researchers: Steve Grand, Rod Brooks, and Hugo de Garis for about 18 months and attempted a fly-on-the-wall view of their successes (very few indeed) and failures (many). To the readership of this publication, this ratio of success to failure is no surprise and probably signals only that the problems are very hard. To a general lay audience, however, I would suggest that it is something very refreshingly different from the sort of hype that anything connected with AI normally gets from the media.

The reviewer, John Jackson, has much (probably well-justified) cynicism about the way in which some big name researchers are running into familiar problems. I'd welcome some emails on whether he's right. Kyran Dale in his response to Steve Grand's open letter (AISBQ 106) offers one view on why this might be. People have been too quick to speculate and build, he suggests. According to Dale, we should pursue understanding of simple mechanisms - not necessarily at the most basic level, but at a level well below that at which Brooks, Grand, and de Garis have chosen to tackle the problem.

Jackson has some similar observations. Genetic algorithms (GAs), for example, are too seductive, he says. Since we know that evolution *has* produced intelligence, we are too easily lead into trying to produce it by that route. In fact, says Jackson, evolution is constantly trying to avoid producing intelligence preferring short-cuts and kludges (like instinct) instead. A genetic algorithm is also unable to pursue any sort of long-range goal. Jackson allows that GAs may be part of the solution, but only in fleshing out the details of a well-understood problem.

Jackson is also opposed to building robots in the early stages of paradigm development. In this respect, his position is similar to Dale's - researchers should not be tackling "the biggest problem in the park" until they have achieved a reasonable level of success with some putative component problems. He offers suggestions, but many readers may consider them too difficult to tackle particularly the level of richness of the sensory apparatus of even the simplest animals.

All this paints a very different picture of the state of the art from the Spielberg movie with which I started the editorial. When journalists enquire (as they have done lately) "How far away are we from building a Gigolo Joe?" I am tempted to echo Richard Feynman's famous reply to a journalist: "I'm sorry but I don't know how to turn that into a sensible question".

Blay Whitby
Editor

Contents

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

News

CPHC Research Strategy Working Group Report on Agent- Based and Multi- Agent Systems

A working group was convened as part of the Conference of Professors and Heads of Computing (CPHC) Research Strategy Workshop organised by Chris Hankin and Ian Watson, held in Manchester in January 2000. This group comprised of leading UK researchers in the areas of agent-based and multi-agent systems. Its aims were to assess the current state of UK research in these areas, to identify the key challenges for their development and to make broad proposals concerning future funding priorities.

Agent-based systems are beginning to be used in a significant number of areas, and are suggested as providing appropriate solutions for an even wider range of problems. Although popularised relatively recently, the notion of an 'agent' can be seen as a natural evolution from work in distributed systems (for example, via coordination and mobility), object-based systems (for example, via concurrency and autonomy) and artificial intelligence (for example, via rationality and emotion); indeed, many researchers in those areas are tackling similar problems to those considered by the agent-based systems research community. In addition, it has significant industrial relevance.

The strength of the UK in agent-based systems research, for example exhibited through the large part played in the AgentLink EU Network of Excellence, together with advances in underlying software and hardware technology, mean that researchers in this area now have both the opportunity, and (viable) infrastructure in which, to develop significant agent-based systems.

Challenges
Key research questions that must be answered involve a number of areas, including the following:

Improved theories concerning massive numbers of interacting agents considering aspects relating to -
- communication
- management and organisation
- dynamic agent creation and open systems
- complexity and tractability

Refined techniques for organising multi-agent activity, for example -
- cooperation/coordination
- teams
- evolution/adaption/emergence of organisational structure

Improved theories of individual agents, capturing the range of possible components, for example learning, reflection, reactivity and perception.

Improved theories concerning how agents decide what to do, for example balancing deliberation and reactivity.

Techniques for interacting with and representing the 'real world'.

Design and implementation ontologies for communication/understanding between agents.

Action Plan
The team produced an action plan which included the following suggestions:

UK would benefit from a cooperative network of researchers tackling the interdisciplinary aspects across the whole of agent research.

Both understanding the techniques required in, and actually undertaking, the construction of large scale multi-agent systems is a difficult system engineering endeavour and the UK would benefit from some form of collaborative programme in this area with specific goals. For example, a programme tackling some identified challenges in agent research relevant to the UK.

There are many fundamental research questions that still require investigation. Targeted funding is needed for foundational research into, for example, models of cooperation and its genesis, methods of managing emergence, non-standard agent architectures, and new agent programming languages.

News

The UK needs specific (interdisciplinary) programmes to train researchers/practioners in agent development techniques.

There are a number of risks associated with research into agent-based systems that need to be avoided, namely -

- 'overhype', and a tendency to 're-invent wheels' within the agent community; lack of awareness of practice/methods used in the rest of the world
- a tendency to jump on the 'agent bandwagon' by the research community in general
- fragmentation of agent research community, partly due to the interdisciplinary nature of research area
- no standardisation or, even worse, premature standardisation
- pressure to make research 'applicable' may lead to 'short termism' and shallow research
- perceived failure, on the part of the agent research community, to explain what the research area is all about.

The full report of the CPHC Research Strategy Working Group on Agent-Based and Multi-Agent Systems is available at http://www.csc.liv.ac.uk/~michael/cphc00rep.html

Jim Doran, University of Essex
(doraj@essex.ac.uk)
Michael Fisher, University of Liverpool
(M.Fisher@csc.liv.ac.uk)
Nick Jennings, University of Southampton
(nrj@ecs.soton.ac.uk)
Chris Priest, HP Labs, Bristol
(Chris_Preist@hplb.hpl.hp.com)
Murray Shanahan, Imperial College,
London (m.shanahan@ic.ac.uk)
Aaron Sloman, University of Birmingham
(A.Sloman@cs.bham.ac.uk)
Mike Wooldridge, University of Liverpool
M.J.Wooldridge@csc.liv.ac.uk]

London-AI plans have changed

Those of you based an hour or less from London by train may be aware of London-AI: a listserver used to distribute a calendar of events in the London, Brighton, Oxford, Cambridge area in the general areas of robotics, cognitive science, neural networks, vision, and artificial intelligence. The idea was to help the huge community in the South East to take advantage of the numerous events going on in the area, network, and meet others working in their own field or locality.

The service is going to be reincarnated so that it is cheaper to run and so it can remain free for all those in the London +50mile area. Plans for the AISB to take it over fell through precisely because, as it was being run, the service was too expensive to maintain. Now, seminar/lecture/conference organizers will be asked to enter event details via the web, so that the information can be automatically compiled (after moderation) into a formatted calendar to go out via both the listserver and a web page.

The person who has saved London-AI is Tom Smith of the University of Sussex (Biols). He has set up a web page which, by the time you read this, should hopefully allow you to enter the relevant data. The URL follows. As the person who got London-AI started, I would like to thank Tom for rescuing the project. I hope the community will help to make it a success by using it to share the events they have organized.

The data-entry page and long term home of London-AI is now at:
http://www.cogs.susx.ac.uk/users/toms/London-AI/index.html

However, it may take a while before the listserver itself moves here. In the meantime, you can sign up to the service via:
http://www.jiscmail.ac.uk/lists/london-ai.html

If Tom's efforts are successful, I still hope, at some stage, to take London-AI onto the next stage: to have specifically-organized London-AI events. If all goes well, this programme would start in the 2002/2003 academic year. If you would be interested in being involved in or hosting such an event, please get in touch with me directly.

Sunny Bains
Imperial College
E-mail: sunny@sunnybains.com
http://www.sunnybains.com

AISB News

Chair's Message

AISBJ

The next major event on our agenda is the launch of AISBJ, the new interdisciplinary journal of Artificial Intelligence and the Simulation of Behaviour. As I write, the first issue of AISBJ is on schedule for circulation with this issue of AISBQ; having had sight of the pre-review papers, I am impressed by the quality of the publication, and I congratulate Simon Colton, his co-editors and the authors on a valuable contribution to the field.

No doubt many AISB members have read of the mass resignation of the editorial board of the Journal of Machine Learning in protest at the commercial requirements of their publisher. As publisher of AISBJ, AISB will not be distributing profits from the new journal: any revenue generated from its sales will be used to the benefit of the Society, its members and the field. We hope that this will be another good reason for members and others to submit their work for publication.

AISB 2002 Convention

Preparations for the AISB 2002 Convention are well under way, and the following symposia have been proposed and selected; further details are available via the AISB website at http://www.aisb.org.uk/:

Adaptive Agents in Multi-Agent Systems (Eduardo Alonso); AI and Creativity in Arts and Science (Amilcar Cardoso & Geraint Wiggins); Applications of AI in Authorship and Style (Peter Smith); Engineering Open Agent Societies (Jeremy Pitt); Intelligent Agents in Virtual Markets (Aspassia Daskalopulu); Grid Computing and e-Science (Omer Rana); Ontologies and Interaction (Luc Schneider & Jim Cunningham); Perception, Learning and Imitation in Lifelike Systems (Yiannis Demiris); Representing Time in Natural Language (Ian Pratt-Hartmann); Workshop on Automated Reasoning (Toby Walsh).

Warmest thanks to Jim Cunningham of Imperial College and his team for their on-going efforts, and to the symposium chairs for theirs to come.

Committee matters

We are, at long last, able to welcome two new members to the AISB committee, namely Eduardo Alonso (City University, London) and Louise Dennis (University of Nottingham). I look forward very much to working with them.

Finally, we have a change of roles on the committee. David Brée, who has now been treasurer and financial director for more than 5 years, will be moving on to a new post, created in response to the perceived need for a more focussed image of the AI world: that of publicity officer. As treasurer, David has seen AISB through a very difficult phase, and will be leaving the finances in a healthier state than they have been for some years. His aim in the new post will be to form a point of contact for journalists and others who need information on AI and Cognitive Science.

Thanks to David for all his hard work as treasurer, and for being prepared to work still further on AISB's behalf.

Taking over from David as treasurer, in January 2002, will be Paul Chung, who has contributed to the committee for a while now without a specific portfolio, and so has a good understanding of the running of the Society. I have no doubt he will take over the reins with a sure hand. Thanks to Paul for taking on this most important aspect of running AISB.

Geraint Wiggins
AISB Chair

Publications dates for the next issue of AISBQ:

All submissions for the March issue of AISBQ must be made no later than **31 January 2002.**

Publication of the AISBQ will be in March, June, September and December, with copies of the AISBJ being sent out with the June and December issues.

Letters

Kyran Dale responds to Steve Grand

Dear Editor,

I wish to respond to Steve Grand's (S.G) Open Letter (AISBQ 106).

First off, for various reasons I had no intention of making any comments on Mr Grand's talk; my response was provoked by a general (a few honourable exceptions notwithstanding) and surprising silence following a highly provocative talk.

One of the conventions of academia is that people then submit to a few questions. If S.G. "can do without this kind of outburst", then this difficulty is all the more surprising since he admits that his robot was "intentionally and provocatively made to resemble an orang-utan". In light of such provocation, I think S.G. did remarkably well and would have fared far worse had his talk been given to, for example, the Sussex Neuroscience Institute.

To place things in some sort of context, there was recently a small hoo-hah concerning the appointment of the cyborg previously known as Kevin Warwick to present the Royal Institution lectures on an alifey/aiey theme. The point was made by, among others, some members of COGS that Warwick's rampant self-publicity coupled with little achievement in the field and gross mis-representation, made him a poor heir to Faraday's title. The underlying premise was that it is a bad thing to have people out there spreading more disinformation in an already poorly understood, over-hyped corner of the media jungle.

Most of the positions attributed to me in the letter are not held by me nor could they reasonably have been derived from anything I said after the talk or have written in connected correspondence. To take a central allegation; I have actually thought for a long time now that the obsession with the minutiae of neuron physiology fatally distracts from more fundamental concerns.

To deal with the three main imputed elements of my argument:

1) I explicitly don't believe that "we can't study complex systems made of billions of parts until we understand the components from which they are made." Having spent five years doing just that, this allegation is odd. I would sign up to an alternative though:

"We shouldn't attempt to model complex systems made of billions of parts until we have successfully modelled complex systems made from hundreds of similar parts."

> '... as things currently stand replication of an insect is hopelessly ambitious. In the belief that ambition should always slightly outreach our grasp, a nemotode worm would be about right.'

2) I do not think that insects are necessarily the natural choice for study and replication at this current time. For the reasons given in 1), as things currently stand replication of an insect is hopelessly ambitious. In the belief that ambition should always slightly outreach our grasp, a nematode worm would be about right.

3) This is not an argument. I would not dream of asking anyone in this field, least of all S.G., "How dare you?". I don't question his right to take on such a hopelessly ambitious task, but by the same token, I reserve the right to point out its hopelessness and ask what, in scientific and engineering terms, are his justifications for embarking upon it; nothing amongst his current achievements suggests he has the slightest chance of success. He has also never provided a model which illuminates even a small mystery of animal, yet alone primate, behaviour.

Tackling the Biggest Problem in the Park
S.G. devotes much of his letter to demonstrating the rather obvious point that there are many interesting things that humans and vertebrates can do that insects

Letters

can't. Well yes of course, but there are also many interesting things that humans and vertebrates can do that insects can. And, it goes without saying, these things are not well understood at all. The belief that you can deal with those things that set us apart (culture, consciousness etc.) without first dealing with those things we have in common (contextually regulated behaviour, sensory stabilisation, walking, co-ordination, multi-modality etc.) has been, I would suggest, the bane of GOFAI; S.G., for all his statements to the contrary, actually sits quite well in this tradition. Although presenting himself throughout his letter as a rebel against some traditional academic orthodoxy, in this field S.G. is very traditional indeed, choosing to tackle the biggest problem in the park before showing any competence at simpler problems that would surely inform our progress. It strikes me as obvious that the higher-level cognitive abilities S.G. is interested in are grounded in those lower-level abilities common to all living animals and about which we currently have only the faintest of understandings. The point is not whether it would be more desirable to have a robot that could learn to fly a plane than one that could just maintain hexapedal navigation towards a specified goal in a foreign environment on uneven terrain, but that understanding how the latter is done is a pre-requisite to achieving the former.   Only if the two tasks were completely disconnected would S.G.'s project be justified and there is no reason to think this is so and many to think otherwise.

What constitutes a general ontogenetic (self-arising within the individual) learning ability seems fairly arbitrary; most of S.G.'s examples seem instead to be socio-genetic (arising within the group) learning abilities peculiar to man and not within the ambit of say orang-utans. If we accept a continuum of ontogenetic learning abilities then our current state of knowledge places us firmly at the lower reaches - far from being able to explain how an individual bee first locates an acceptable foraging site in an unknown environment, then learns enough about that site to be able to return there and communicate its whereabouts to other bees. Some progress is being made with respect to the bee's hardware but otherwise it's almost a complete mystery and a reasonably simple example of the ontogenetic learning of which S.G. is fond.   For someone contemplating human learning abilities solving it would surely be but the work of an afternoon and would lend enormous credibility to S.G.'s argument besides being worth oodles of cash and probably winning a Nobel Prize.

Insects are a Red Herring

Having brought them up it should be said, mea culpa, that with regard to S.G.'s work, insects are, if you'll excuse the phylogenetic scrambling, a bit of a red herring. The key themes at the table should be cognition and behaviour, the degree of the former being inferred from the sophistication of the latter. It is an open secret among neuroscientists, alifers, roboticists etc. that no behavioural model has been produced thus far that begins to approach, in terms of behavioural sophistication, the humble ant (approx. 500,000 neurons - take your pick of species). In fact the puny nematode worm C.Elegans (approx. 300 neurons) is well beyond us; the largest robotic networks being numbered in tens of greatly simplified artificial neurons. Insects may not be something that we would wish to emulate for practical purposes - having spent fifteen hours of the last week trapped in a small room with a bunch of obstreperous bumble-bees, the thought of larger robotic models is not enticing - but they do represent a benchmark, capable of many things such as visual gaze-stabilisation, sophisticated behavioural repertoires etc…, which we do not currently understand and are not yet able to realistically simulate.   And surely these 'lower level' abilities provide a necessary framework within

> 'No behavioural model has been produced thus far that begins to approach, in terms of behavioural sophistication, the humble ant (approx. 500,000 nuerons - take your pick of species).'

Letters

which higher cognitive abilities can be expressed; until we understand such things as sensory stabilisation, the reproduction of such hoary faves as consciousness is hopelessly premature.

Much of what we know about cell-cycle checkpoints, crucial to current cancer research, comes from research on the yeasts S. pombe and S. cerevisae. Developmental biology is the fruit-fly Drosophila (how's that for a category error). Currently thriving areas of biological research are discovering generally applicable mechanisms by studying simple organisms and yet, mysteriously, when it comes to the relatively moribund cognitive sciences, these rules don't apply. Even allowing that in the field of animal behaviour understanding simple organisms is not a pre-requisite for understanding more complex ones, S.G.'s approach seems methodologically untenable; the robot (Lucy) will never come close to human intelligence because it's just not built for it.

Lucy

Sitting at the front making occasional stereotyped movements, servo-motors grinding away, Lucy looked as pre-possessing and as cognitively enriched as those gum-ball mannequins used to frighten people away from news-agents; there really is an uncanny resemblance. It can't yet walk, crawl or visually fixate but my guess is that S.G. will soon decide to skip such trivial cognitive tasks and concentrate on getting it to independently formulate a theory of relativity. What specs we did get seemed to drive a stake through any pretensions to mammalian intelligence. From memory, it has monocular vision (stereopsis presumably being very difficult to implement) involving a grey-scale 128x128 pixellate array. Well no primatologist I, but that seems a wee bit under-powered for primate/human vision and the 120 million odd rods and cones that hook us up to the world, never mind the small issue of a visual cortex. What about the other senses? I don't recall mention of ears but a couple of microphones

won't hack it; these are hideously complex sensing devices in their own right, designed to intelligently process very noisy signals. And its proprioception is going to be hindered somewhat by a lack of skin - it has a few mechano-receptors but I'm guessing they don't begin to approach the sophistication of any one of the hundreds of hairs found on a single spider leg (see a recent J Comp Physiol (187/4) for very interesting analysis). One could go on but, in short, the sensory ability of Lucy would disgrace your average foraging honeybee, never mind a higher mammal. Now there are few generally accepted central tenets to cognitive science but the correlation between sensory sophistication and behavioural complexity seems a sure thing. In this sense Lucy is crippled from the start.

'Steve Grand's approach seems methodologically untenable; the robot (Lucy) will never come close to human intelligence because it's just not built for it.'

Creatures

S.G's other main claim to fame is, of course, authorship of the A-life software 'Creatures'. Great things are said of Creatures, mainly by S.G. himself, but they don't appear to stand up to much scrutiny. (I say 'appear' because the innards of Creatures Norns have never been detailed; S.G. would probably claim for reasons of copyright but closer to the truth, I suspect, is that this would reveal how much of a 'trick' the game is.) Biological metaphors abound with little attempt at justification, so we have Norns developing proto-language, cultural ties etc. but what we really have is a reasonably sophisticated game with a few neat twists. It doesn't look that bad, the Norns being a bit cutesy and disneyfied for my tastes (think ewok), with quite pretty animated backgrounds (for 1993 - 3d would be the norm now), but, hey, it's for the kids. Much of the engagement with the program comes from this innate tendency to project feelings onto anything with large brown eyes and a small nose; why kids can hold dialogues with a favourite teddy but few have a pet pencil. That's a great asset for a piece of entertainment but take away those big eyes that don't actually see anything (2d line projected magical awareness not

Letters

withstanding) and those big ears that don't actually hear anything and that big mouth that doesn't actually say anything and those legs that don't actually support a torso or maintain a bipedal gait on varying terrain (tough problem which has only just been cracked in the real-world by Hitachi, many millions of dollars and a few thousand man-hours later) and what you have left is less impressive. A more accurate representation, stripped of populist software obligations, might be some coloured circles moving on a squared 2-d grid, occasionally interacting, changing colour, texture etc.. Lets throw in a few variables, connected to such stuff as availability of 'food' sources derived from corresponding property vector describing each square etc, and call them a metabolism. Add an ability to form 'associations' between objects, maybe by manipulating gravitational properties etc. and hey-presto, the basis for an a-life MSc thesis. Except that such a project would pre-suppose at least some analysis of the behaviour totally lacking in S.G.'s discussion of Creatures.

> '...Creatures is pretty good but is not, a significant contribution to our understanding of animal behaviour or a suitable springboard to the kind of project Steve Grand is proposing...'

So if there isn't that much more to 'Creatures' - and, as mentioned, in lieu of further details one has to go by appearances - why the hype? Well one thing that distinguished Creatures from other reasonably sophisticated, highly constrained games is its use of an evolutionary algorithm (though little evidence that Ns. have evolved in any meaningful sense of the word), presumably the main reason for Dawkins interest. This seems primarily employed to mix characteristics in a trait-based fashion in the way a tall green x and a short blue x might produce a short green x. Well natural evolution has an attendant mystery that seems to demand respect - for a start many things, such as the origin and maintenance of sex, aren't fully understood - but evolutionary or genetic algorithms, while being hard to use well, are pretty easy to implement. Now S.G. mentions stuff like genotyping in regard to his GA but the analogy between natural genes and the artificial ones of which the Norns are an example is massively impoverished. The complex genetic interaction seen in living organisms, for example during morphogenesis, is nowhere to be found in S.G.'s model. It should be mentioned that the makers of Sims, a piece of software that seems to dwarf Creatures in terms of behavioural complexity, could probably graft on such an evolutionary component with ease, claim to have made a revolutionary breakthrough developing the skills necessary to tackle the simulation of higher life-forms and then seek funds to build a monkey in the basement, but they don't because they know, as does the gaming AI community, the massive constraints and sleights of hand necessary just to get the game up and running. In short Creatures is pretty good but is not, in any way, a significant contribution to our understanding of animal behaviour or a suitable springboard to the kind of project S.G. is proposing; an autonomous glider perhaps but certainly nothing as intelligent as a nematode worm.

Given that S.G.'s project is so obviously over-ambitious, why not just treat it as a bit of a joke and let be? Unfortunately those disciplines, both public and private, dealing with animal behaviour are, as a whole, so pitifully weak and fragmentary that I don't think a sense of humour is particularly appropriate just now. Given the lack of any agreed objectives, progressive canonical models or much to show for the last forty years or so I think sober reflection is the order of the day. Whereas a strong discipline like physics can condescend to occasional kookiness, in these fields the eccentrics hold centre field. Whereas, for example, the media is replete with reasonably informative, well-made popular science covering such stuff as physics, chemistry and cosmology, often managing to entice truly great scientists, an ALife programme wouldn't be complete without a contribution from Warwick, de Garis et. al. I know that S.G. often claims to be uninterested in the petty-fogging concerns of science, but that doesn't mean science

Letters

should be uninterested in him - and science should not, of course, presuppose academic science; the output from all camps is pretty modest. Hype, exaggeration, disinformation are bad per-se, whether you're applying for a grant to study the neural system of a sea-snail or trying to build a robust, off-line Mars-explorer. History shows that human enquiry, be it scientific or engineering - assuming it's possible to separate the two - is quite capable of disappearing elliptically up a certain orifice for hundreds of years. If there's one plot in the garden of knowledge where the fruit is not yet ripe for the picking it's here; this might not sit well with our egos but it's true nonetheless. Rather than contemplating artificial humanoid intelligence, we should all be indulging in a serving of humble pie; in the long run it could be good for us.

To conclude, I have no particular beef with S.G.; on the day of the talk he, at least, did his job. I find it hard to understand his surprise at my objections, but at least he presented his views in public. I had no intention of passing comment and would rather that others had said something and said it better.

Kyran Dale
Centre for Computational Neuroscience
and Robotics, University of Sussex

Emails to the Editor
AI: The Movie

Addressing two important questions

The film AI seems at first to be addressing two important questions for the discipline AI. One is: could a robot ever have real feelings? The other: could a person love a realistic child-like robot as she might her own son? Disappointingly, neither issue is explored in any interesting way: the answer to both questions is simply taken for granted as "yes". Equally, the film fails to acknowledge any of the practical challenges facing the discipline; more or less magical robotic technology is assumed. As a vehicle for saying something about real AI or its practitioners, the film has nothing going for

it. It is also easy to criticise on general grounds. There are some slow sections that needed editing. The human characters are shallow, one-dimensional, and stereotyped in the extreme. The gobbledegook about information storage in space-time (or whatever), used to justify the final scenes, would hardly be worthy of Star Trek on a bad day (and I think it matters - decent science fiction is not so lazy). The voice-over narration is heavy-handed, and the cars are feeble. Does all that make it a bad film? No: it's a good film. It has drama, pathos, and of course wonderful special effects and excellent photography. The plot hangs together, and has some real surprises. There are good characters amongst the robots: I particularly liked the gigolo-robot (Jude Law), and the teddy bear is pretty neat. The final section, sometimes criticised for being "tacked on" and weakening the structure, is in my view an important part of the film, with its evocation of a strange, protected, artificial world deliberately recalling the bizarre penultimate scene of Kubrick's 2001. Overall, I'd say: forget you work in AI, and see it.

David Young,
COGS

Printed with permission from Warner Brothers

Letters

Could have saved me five pounds

Blay Whitby wrote: My personal view is that it's unbelievably shallow, sentimental, and seems a poor competitor to various other films that have tackled the same issues over the years. Sorry if you went to see it expecting a half-decent film. If you disagree, I'd love to hear.

I have to say I totally agree, it would have saved me 5 pounds had I read your e-mail earlier :). If you find someone who disagrees, please let me know.

Helen Vassilakis
COGS

(I hope publishing these letters counts as informing you, ed.)

Avenues of Debate

I agree that AI:The movie tells us more about obsession than love. The android child behaved like a machine in its unfaltering and determined prayer, not like a human who would eventually have realised the absurdity of the situation or become doubtful of the veracity of his own emotions. This is part of what is interesting about the film.

There is an underlying tension about the idea of the mecha/orga divide. Should the robots have human rights? The whole film tries to get us to empathise with the android child and hence to endow him with human emotions and the ability to suffer. Yet the possibility that the android is only just a good simulator of emotions is also left wide open and there are several moments of ambiguity.

If the robot is just simulating emotions then a machine has been built that has an amazing ability to manipulate human sympathy. (I vaguely remember seeing something about MIT working on expressive interfaces for this purpose). Recall the scene at the circus/freak show when the apparently life-like emotion of the boy (did he cry?) causes the originally hostile audience to turn against the ringmaster. None of this is left rigid or clear cut, there are interesting themes raised, and enough room for interpretation there to inspire some nice avenues of debate.

Sam Woolf
COGS

Kubrick will be turning in his grave

As far as quality film making goes, Stanley Kubrick will be turning in his grave. I recommend the superior 80's classic, D.A.R.Y.L. (which can be also enjoyed in an ironic way). To me the film seems like a patchwork of tenuously related scenarios, that don't amount to much more than chewing gum for the eyes. No thanks, I have nachos. I have one positive thing to say for the film, 'robots fired out of cannons? Fantastic!'.

What amuses me more, is the fact that 'Joe Public' will now undoubtedly have some twisted view of what AI is all about. I was actually asked by someone "how close are you guys to creating robots that can't be told apart from a human". I felt that telling him about blocksworld and insect behaviour would be like dispelling the Santa Claus conspiracy (sorry, if I've upset anyone).

I would also like to thank Warner Bros. and all the Kevin Warwicks of this world for their outstanding marketing of the field, AI. Speaking of which, did anyone see the Channel 4 documentary about robots in Hollywood? How I laughed at Warwick's claim that moral issues in AI were 'addressed' by the film Terminator. Don't even get me started on the film 'Enigma' - starring an
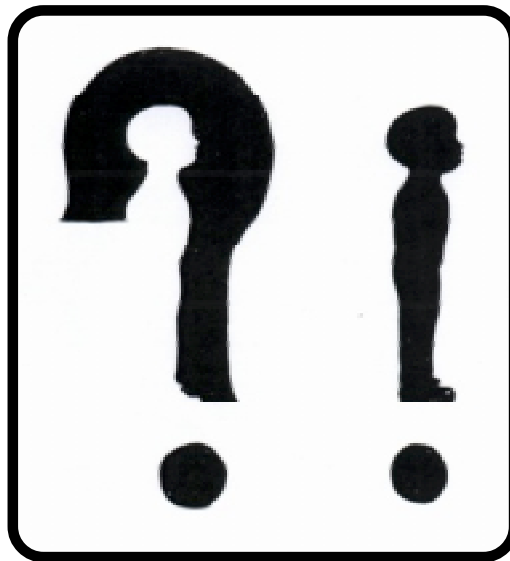
Letters

Alan Turing-type as a heterosexual, action-hero? Hmmm... computer science has never had so much kudos!

Alex Glover
COGS

What makes Spielberg great

I saw the film last night and from the stance you are taking, anyone could have thought you weren't talking about a Spielberg film. Have you seen any of his other films? I think this is definitely one of his better attempts. I think he went easy on the special effects (until the ridiculous ending of course - was that 2001 envy or what?) and in terms of sentimentality, I think he was biting his lip the whole way. IT'S A SPIELBERG MOVIE!!

What makes Spielberg great? For me it is his ability to take a very human and ageless perspective on everyday human life and society and portray how this might be affected by an extraordinary situation. AI began by setting a scene and building characters and a situation which was neglected in the rest of the film. I have probably seen most of Spielberg's films and on the whole, I think they are fantastic. The problem is that I really don't see AI as a Spielberg film, it was a project begun by Kubrick, and realised by Spielberg. Unfortunately I don't think that the two styles are complementary. Either Kubrick should have done it or Spielberg; it was a complete mismatch of creative drives. Since when does 'director = great film', there are always exceptions.

Have you read Super Toys? It's not exactly the most serious critique of the field I've ever read. It's more like a Philip K. Dick-style 'lets use a scientific topic as way to get a novel twist at the end' story telling tool.

It's a book that I feel didn't translate to film very well, it was apparent that the ending of AI was tacked on whilst in the book it was dealt with in more detail. I like the way Blade Runner (Dick) was dealt with because the elements that didn't work were changed or left out.

Kubrick doubtless wouldn't have added an ending specifically designed to loose you your lunch but he would have been directing a piece of entertainment and NOT "Computing Machinery and Intelligence - The Movie".

I've never had that effect from a Kubrick film; maybe I have a cast iron gut or something? When I go to see a film, I expect entertainment, not a documentary, and not a strange concoction of the two. And I agree he probably would have come up with a better title :)

As for the suggestion that this one film is going to give the public an inaccurate perception of the field? Well, it's called A.I. and it has a lot of A.I. hype surrounding it but given the contribution of certain 'experts' in the field and such <IRONY>classics<IRONY> as Virtuosity, I think the damage has already been done elsewhere and it's been done far more effectively, don't you?

John McFarlane
The Creative Assembly Ltd

If you want to write to the editor, or send an open letter for publication, please write to:

The Editor, AISBQ, School of Cognitive and Computing Sciences, University of Sussex, Brighton, East Sussex BN1 9QH

Or email editor@aisb.org.uk

Reviews

Thoughts on 'The Battle of the Robots. The Hunt for AI'.
Channel 4, 13.10.2001

We, of course, know any creeping cynicism for our subject is misplaced; the efforts of each one of us take it forward to ever higher glory. Perhaps though it might be instructive to slip occasionally into the mind frame of those of less pure faith.

The first of the three portrayed as contestants in the race to create a machine that was "really alive", Hugo de Garis, had pulled down massive Belgian funding for a genetic algorithm approach using a super-duper array processor.

After my initial terror that Hugo might succeed where I had failed subsided, I watched a Greek tragedy unfold as he followed my footsteps into an uncannily similar bog to the one I had foundered in 10 years before.

GA's are to AI'ers as candles are to moths. We know a GA can create intelligence - it's the only thing that ever has. Paradoxically though, intelligence is among the least suitable of tasks for a GA since each individual needs an evaluation function involving a complex learning regime, which is therefore time consuming. The individual's own operation will inevitably be computationally expensive anyway. (What a motto for the computer development company with whom Hugo was trying to work: "A place where 100 years means nothing"!) The ancient enmity between evolution and intelligence must also be faced: evolution tends to veer away from intelligence by trying to be clever and find short cuts based on instinct. The tests posed to the individuals must therefore vary, to avoid developing trivial solutions to standard features of the test; however each test must be roughly similar in difficulty to the last.

By far the worst hazzard in using GA's though is the failure to appreciate their inability to pursue a long range aim. Progressive optimisation of a task where the structure of the solution stays essentially the same throughout, can work well. However, when the task ascends through a cascade of levels different in character, or when the solution trail must progress through a sequence of structural developments that differ in some difficult-to-explain way, GA's can easily get stuck. They sometimes need to be guided in very small steps towards qualitatively different goals in turn, which unfortunately presumes knowledge of these stages on the part of the developer - and it was probably because he wasn't exactly sure what the stages were that he was using a GA to discover them in the first place! GA's are probably the most useful tool for the task but they need a lot of careful shepherding.

I could have told him all this, had he asked, but it was only through failing my M.Phil that I found it out - and who wants the advice of a failure! Few people take my advice at the best of times. Had he done so, I could also have warned him off using a huge machine. AI needs a powerful computer, but that's what a modern PC is. If you can't demonstrate your principle on a £1,000 machine of today, albeit slowly and a bit representatively, you haven't got a working principle at all. Hugo found that his monster muncher brought with it an army of beaurocrats brandishing deadlines and expectations, a financial impasse with the inventor disappearing off to the USA without telling anyone how to turn the machine on, meetings deciding your future you were barred from, and all the rest of it. It's hard enough writing the biggest program you've ever written, in which every line of code presumes to resolve some conundrum that's been puzzling philosophers for centuries, without having to fight off the infrastructure as well. This example of 'Less is More' is just one of the counter-intuitive secrets of cracking AI.

But at least Hugo didn't make the mistake of building robots, while only in the early stages of paradigm development. "Why are they still building those liddle robarts?" is my favourite Minsky quote. When they work they're a very concrete demonstration, but even then, are they more impressive than those of 40 years ago? Although Minsky didn't appear on the programme, his spirit hovers over the subject continually. His successor at MIT, Rod Brooks, had an collaborator, Brian Scassellati, who was the second contestant in the race. Of course

Reviews

Rod and Brian have a good understanding of principles essential to the grand project, as their conversations showed, but if anyone agreed to follow the approach of Brooks, whose attempts a few years ago to implement some admittedly sound high-level principles made my jaw drop, I would be surprised to see him breast the tape. Brooks' comment as Brian packed up to leave were even sadder than the sight of the robot's head falling as it was turned off. ('Cog' the robot had succeeded in mimicking a human's waving hand, but perhaps it was not waving but desperately signalling for more help.) At least Rod was honest: "There is no vision system which can, given these two things sitting on the table, say - That's a telephone, that's a pair of glasses." My own Ph.D. thesis was on that 20 years ago, and I didn't succeed, and no-one has done anything really successful, ever!

The third horse in the race was Steve Grand, who I first read about on the pages of this august organ. Luckily for my blood-pressure I convinced myself then that he wasn't on the final lap. His circumstances seemed to me far more promising though: no deadlines, a nice helpmate instead of a boss, and a sound style of philosophy. I could only fault him on his circumstances, for having a soldering iron in his workplace, and a real robot. They should both have been inside the computer. However, in his case the hardware was not for the benefit of anyone else, just himself, so perhaps it played a useful inspirational role.

I have toyed with the idea of mentioning some of my own thoughts to Steve, who I am sure would give me a better hearing than the EU or MIT would, (or that person at the OU who told me they only work on projects developed there). However, questors for the holy grail are doomed in many ways. If you even look like achieving it, everyone will hate you, as, frankly, I would if you beat me to it. I can't believe he would appreciate my advice if it were good. "If anyone does solve these problems, they will become immensely rich" according to the film's narrator. I doubt it. The solution will be a hunch-processing system, which will need to be embedded into something with high reliability if it is to be sold as anything other than a toy. There will be many "company men" between you and the reward, and as they openly admit: "We usually end up shooting the inventor."

C4's programme had the wrong kind of pessimism. Although it's been slow coming, we will get there. Unfortunately, I can't imagine the Prometheus responsible avoiding his classic fate.

John Jackson
Southampton

Review of 'Thinks...',
A novel by David Lodge, Secker & Warburg, London, 2001.

Reading David Lodge's novel *Thinks...*, I came to the words "It was the Second Sunday of Lent, I discovered." Looking at my calendar, I discovered it was the Second Sunday of Lent. The coincidence was appropriate, since Lodge has never shied away from placing coincidence at the heart of his plots. In this as in other respects, *Thinks...* is a characteristic work, exploring Lodge's preoccupations with sex, death, religion and the lives of academics. What makes it peculiarly interesting to readers of *AISBQ* is that here he also sets out to address questions about consciousness, as framed by AI and cognitive science.

The central characters are Helen Reed and Ralph Messenger. Helen, newly widowed and good-looking, is standing in at the University of Gloucester to teach prose narrative to a group of wonderfully intelligent and nice MA students. Ralph, the charismatic and womanising head of the University's Centre for Cognitive Science, is experimenting with tape-recording his own vocalised stream of consciousness. When they meet, Ralph decides that he would like to extend his study and explore Helen's consciousness, as well as her body. Magus-like, he gives her an introduction to some of the themes of cognitive science, aided by an extraordinary mural depicting crucial computational and thought experiments such as Searle's Chinese room and Schrödinger's cat. Intrigued by both Ralph and his studies, Helen tries to give some of these ideas a new perspective by pursuing them in her writing class, while Ralph pursues his own scientific and unscientific goals. Helen's equilibrium is

Reviews

rocked when she discovers a coincidental link between one of her students and her dead husband; Ralph's world too is shaken when the legacy of a youthful indulgence catches up with him. As we wait to discover how their relationship will develop, the realities of their environment come to overwhelm their explorations of the workings of their minds.

Whilst Helen is attractive as a writer and teacher, Ralph is a sorry specimen of a scientist. He fantasises about winning honours, and dreams of receiving a Nobel Prize, not of making the discovery that might lead to it; his fear is that a colleague might achieve recognition first. His dismissal of Turing is revealing: "… a truly great mind … but a totally screwed-up human being, a lonely, repressed, unhappy homosexual, eventually killed himself in a dreary flat in Manchester …". Ralph has no conception of what it must have been like to hold, for a while, the keys to our intellectual and technological development. There is no reason, of course, for Lodge to present us with an example of shining scientific virtue, but in Ralph Messenger he has given us someone who does science publicly whilst having no genuine scientific purpose. It is a pity that in a novel bold enough to take science seriously, the scientist-character should only be interested in acclaim, money, power and sex; that there is nothing to communicate the pleasure of making even the most modest discovery.

If scientists fare badly at Lodge's hands, what about science? Lodge has certainly done his homework, helped by Aaron Sloman, Professor of AI and Cognitive Science at Birmingham University. The bibliography at the end of the book covers a wide range of sources, which give an excellent introduction to studies of consciousness. Lodge has mined these selectively but fruitfully. The device of the mural (reminiscent of the carved doorway in Eco's *The Name of the Rose*) is an effective and graphic way to present a diverse, powerful set of ideas. The half-dozen pages where Ralph shows it to Helen are well worth considering as seminar material. Later Helen asks her students to explore two of the ideas illustrated, from Nagel's "What is it Like to be a Bat?" and from Jackson's "What Mary Didn't Know" (which is about a scientist who

has been kept in a monochrome prison, but has learnt all that can be learnt about colour without having experienced it). The students write pieces on these themes adopting the styles of various modern novelists, and the results, to which two chapters are devoted, amount to a stylistic *tour de force*. Alongside this overtly academic strand of the novel, Ralph's and Helen's perceptions of each other, of themselves, and of the events around them, recorded on Ralph's tapes and in Helen's diary, invite us to consider Ralph's statement of the problem of consciousness: "How to give an objective, third-person account of a subjective, first-person phenomenon."

Lodge thus presents cognitive science as science is rarely presented in novels: seriously, as something that can and should be thought about by non-scientists, and that affords real interest in its own right. Nevertheless, he is marking out intellectual territory: "I've always assumed, I suppose, that consciousness was the province of the arts, especially literature, and most especially the novel," writes Helen shortly after meeting Ralph, "I sort of resent the idea of science poking its nose into this business, *my* business." Later she addresses a scientific conference: "… literature is a written record of human consciousness, arguably the richest we have." Yet if Lodge shares Helen's feeling that scientists are reluctant to recognise this, he is surely wrong; a scientist with as limited a view of literature as Ralph is (one hopes) a rare animal.

In the end, rather than seeking to break down the science/arts divide, Lodge sustains it. Helen falls into the trap of opposing a literal-minded reductionism to a humanistic viewpoint: "We are told that … each of us is 'just a pack of neurons' … or just a parallel processing computer running by itself without an operator. As a human being and as a writer, I find that view of consciousness abhorrent - and intuitively unconvincing." She fails to realise that the word "just", in the sense she means it, is inserted not by the scientist, but by the non-scientist who does not see how a rich superstructure can be built on mechanistic foundations.

Nonetheless, *Thinks…* does have a message for cognitive science, even if it is not the one

Reviews

articulated in Helen's conference address. Earlier she notes that novels can, in a sense, be called thought experiments. Thought experiments are a vital tool of scientific investigation, yet, as the example of Maxwell's demon in physics shows, they give the wrong answers unless they reflect reality sufficiently accurately. (The demon seems to be able to break the Second Law of Thermodynamics by deflecting fast-moving molecules, but the thought experiment neglects the energy needed to see the molecules.) How far can we trust thought experiments like that conducted on Mary the colour scientist? Only to the extent that we can have confidence in every significant detail. It is debatable whether or not Lodge's rewritings of the experiment shed any light on the logic of the argument about qualia, but it seems to me that the exercise gives us a clear warning about missing out the equivalent of the demon's light source. Imagining Mary seeing colour for the first time tells you nothing, unless the Mary you imagine might be real.

It is easy to recommend this book. If you like Lodge's work you will enjoy it. As a picture of university life it is of course highly entertaining. If you are interested in a view of cognitive science and AI from an extraordinarily well-informed outsider, it is fascinating. And if you happen to be reading this on 24 February 2001, have a look at your calendar.

David Young
School of Cognitive and Computing
Sciences, University of Sussex

From the Archives – 24 Years Ago

Of Tortoises and Men: the development of electro-mechanical "organisms" and their implication for society.
Richard Rosenberg, University of British Colombia, Canada

The most general feature of news reporting in this field is the rampant animalization of machines. Consider the following extended extract from an article in TIME of March 27, 1950.

"But wandering around Dr. Walter's house, they act much like real live animals. …Elmer and Elsie are nocturnal. During the day repelled by too-strong light, they hide in a cozy "hutch" against the wainscotting. When night comes, they venture out in search of the mild artificial light they crave. Guided by their photoelectric eyes, they creep towards a lamp or the fireplace. When they hit an obstacle, they stop, growl faintly, back away and try again at a slightly different angle. Their wanderings take them all over the house. When they reach a light of the proper intensity, they bask under it blissfully in photoelectric euphoria.

But contentment does not last. As their batteries run down, Elmer and Elsie begin to feel uneasy. When hunger begins to dominate them, they lose interest in their gentle light. Now they want strong light: the bright glaring lamp that burns inside their hutch. They scuttle towards it eagerly. If all goes well, they pop into the hutch; where electrical contacts quiet their hunger by recharging their batteries. Not until their run down stomachs are full do they creep out again in search of gentle light."

It is almost superfluous to point out and catalogue the many words and phrases which are applied without qualification to these simple mechanisms as if they actually were living creatures. But a few outrageous examples must be given special attention: "growl faintly", "bask…blissfully", "feel uneasy", "lose interest".

…Thus, as of 1951, a small number of mobile devices had been created and in nearly all descriptions of their behaviour, words like thinking, intelligence and free will appeared most often without the use of quotation marks. To repeat: if such words can be applied to mechanisms of such low order of complexity, sometimes by the inventor, sometimes by the press, how much easier it will be to apply and accept them when the complexity increases significantly. There are obvious dangers in this process and we ignore them at our own peril.

AISB Quarterly Issue No. 28,   Oct 1977
Thanks to Rudi Lutz for selection of archive material.

Father Hacker's Guide for the Young AI Researcher

Cognitive Divinity Programme
Institute of Applied Epistemology

NEW

During a long distinguished career, I have accumulated profound wisdom and experience on the practice of AI research. Now, in the autumn of that career, it is time to pass on my collected knowledge to the next generation. Of course, my regular diary entries have carried important but implicit lessons for the aspiring researcher. But, in the fast-paced, modern world, the extraction of hidden messages from moral tales is too slow and uncertain a process. People now want the straight guff. So, inspired by the new look AISB Quarterly, I have remodelled my column as a guide for the young AI researcher. In each issue I will tackle a different skill that the ambitious AI researcher must perfect. I will identify techniques and pass on tips that will turn the aspiring novice into a successful master. For this inaugural entry in the guide, there is no more important skill than that of:

1) Creating Time for the Pursuit of Research

The modern research world contains numerous distractions, which must be avoided if a successful research career is to be pursued; administrative and clerical make-work, aimless committee meetings, training the obtuse, are duties all readily imposed that will fill up your life if active avoidance measures are not adopted. Here are some avoidance tips.

It is essential to cultivate a reputation for incompetence and irresponsibility. It would be counter-productive if this were seen to be a deliberate ploy: it must appear natural and artless.

The cultivation of a charming and harmless eccentricity will assist: wear outrageous clothing, ration your ablutions and use your memory selectively.

Above all, be an unreliable timekeeper. Arrive at meetings late, or not at all, or on the wrong day, or week. Respect deadlines - as the day to start the task in question.

If, despite all these precautions, you are allocated an onerous duty, it is not enough merely to make a mess of it. Your failure must be spectacular: the stuff of coffee room legend. This will require careful planning, especially if it is to appear unplanned.

When your boss (or your boss's boss) reprimands you, plead guilty, promise sincerely to try to improve, but don't.

Offer to take any time management or similar courses on offer. Turn up to the right venue, at the right time, but in the wrong week.

Plead an obscure, chronic, impossible to diagnose, disorganisational illness. Make regular visits to your psychiatrist to coincide with the most important departmental committee meetings.

Lastly, a tip about dealing with family pressures. If your family demands that you spend quality time with them, take them to the cinema or theatre. You can sleep while they watch. Later, you can work while they sleep.