AISB Quarterly No.126

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

"I rarely use it myself, sir, it promotes rust": Doing Conversation with machines

Robby the Robot's comment about the oxygen content of the atmosphere in "Forbidden Planet" [MGM, 1956] captures three points about conversations with a machine that I believe will shape future research in the area. In his opening conversation with Commander Adams, Robby is explicitly introducing the topic of his non humanness. Robby is not doing the Turing Test, and as such is making no attempt to deceive his conversational partner. Robby's strategy has interesting consequences for those of us designing conversational agents. As de Angeli et al [2001] point out, ALICE also acknowledges the fact it is a machine and as such has an out: ALICE does not need to be able to talk about food and what it tastes like, about fashion and clothes, or about the weather. This does, however, mean that all topics are up for negotiation.

The second point is that the difference between the embodiment of man and machine makes a difference to the metaphorical meaning of what is said. To Robby, oxygen is like low level background radiation or carcinogens in food - unavoidable and unpleasant. If, as Lakoff and Johnson [1980] argued, we either have direct experience of something allowing us to ground meaning in *doing*, or meaning is built on metaphors of this experience, then robots are going to ground language in different ways to us humans. What would language grounded in action look like for a machine, and what would such a model tell us about the workings of human language?

Thirdly, the comment is funny. Commander Adams was attempting to make polite conversation upon meeting a stranger. "Nice [place] you have here..." He was behaving normatively and expected the second part of the normal sequence along the lines of "Thank you, we like it." Robby could have ignored the reference to oxygen, but Robby performs a discourse act that emphasises his otherness. Indeed the script spends considerable time introducing Robby, his abilities (strong, obedient etc) and inabilities (he does not feel pain).

These three are key points about a situated approach to building conversational agents. Contrast these with conventional approaches to human language technology which have tended to focus on either generic machine learning over larger and larger collections of recorded human behaviour, or focused on information flow. The tendency is to ignore the fact that people are social animals and that, in human-human conversation, the primary role of language is to manage social relations [Dunbar, 1996]. Such behaviour is mostly invisible to us humans as it is just common-sense. The challenge for those working in this area of artificial intelligence is to come up with some means of capturing this common-sense in a form that is amenable to programming.

One might agree that humans need to manage social relations but surely the usual case when talking to a machine is that the human wants to be informed, or to command, or some other useful speech act. When things go well this is true and information transfer is what conversation with a machine is all about. When things don't work out however even minor discourse repair strategies require social intelligence. Consider this extract from the DARPA Communicator data in which the caller thinks she is booking a flight from Duluth, Minnesota to Chicago, Illinois - a route flown by four airlines with more than 20 flights a week.

```
...
System: let's see then. I have you
going from Toulouse Minnesota to
Chicago Illinois on November twenty
ninth is that correct?
Caller: yes
System: please let me check
availability and see what I can find
```

In this issue:

Features

Reviews

Society

Father Hacker

Doing Conversation with Machines (contd.)

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

ISSN 0268-4179

Editor
Colin G. Johnson.
Computing Laboratory,
University of Kent at Canterbury
aisbq@aisb.org.uk
http://www.aisb.org.uk/aisbq/index.shtml

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


```
System: no records satisfy your
request. would you like to depart on
a different date?
Caller: no!
System: would you like to start over?
...
```

How would you react? Verbal abuse is common with these systems and linguists have explanations for why, but as a socialised human being, common-sense tells you that the system's response in this case is just wrong. A comprehensive theory would be good, but at this stage a systematic approach to "noticing" would help us collect the data that needs explaining.

The problem here is of course caused by the ASR system confusing "Duluth" with "Toulouse" and one might think more resources should be spent improving word recognition rates. However, humans make mistakes all the time and, even when the "error rate" of a human operator is considerably larger than one might expect of a machine, the people using the service still don't get annoyed [Wallis et al, 2001]. Why? Choosing what to say next is non trivial and, it seems, can compensate for less than perfect automatic speech recognition.

The Companions Project [http://www.companions-project.org] is a large FP6 integrated project that will be touching on some of these issues. In addition, the computer science department at Sheffield has a series of proposals in the pipe-line that look at meaning and embodiment, language and emotional content, and interdisciplinary work on noticing. In particular we would like to meet people interested in "scripting" conversational art work, or characters in virtual worlds. Please do feel free to contact us!

References

1. de Angeli, Antonella, Graham I. Johnson and Lynne Coventry The Unfriendly User: Exploring social reaction to chatterbots in Proceedings of the International Conference on Affective Human Factors Design, Asean Academic Press, London, 2001

2. Lakoff, George and Mark Johnson. Metaphors We Live By, Cambridge University Press, 1980

3. Dunbar, Robin. Grooming, Gossip, and the Evolution of Language, Harvard University Press, 1996

4. Wallis, Peter, Helen Mitchard, Damian O'Dea and Jyotsna Das. Dialogue Modelling for a Conversational Agent, in AI2001: Advance in Artificial Intelligence, 14th Australasian Joint Conference on Artificial Intelligence Springer (LNAI 2256), 2001.

Peter Wallis
Department of Computer Science
The University of Sheffield
Email: p.c.wallis@gmail.com
http://www.dcs.shef.ac.uk/~peter/

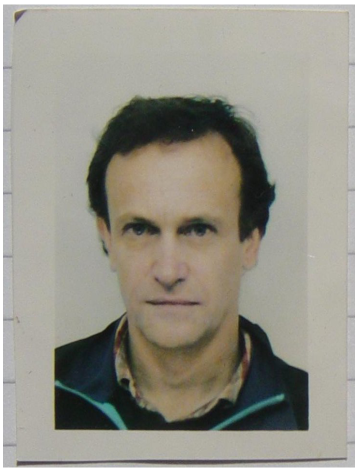

Programming and Verifying Complex Systems

Research in the area of autonomous agents and multi-agent systems has produced an enormous variety of techniques that deal with particular aspects of either modelling and simulating complex social systems (such as human societies) or indeed developing software that can operate in highly complex, dynamic, unpredictable environment. In fact, just in the last issue of AISB Quarterly we saw an example of such techniques, and indeed the motivation for an approach—often called agent-based computing—for designing software that can deal with such ambitious, typically safety critical, areas of application. There is little point denying though that many AI and Agents techniques do note scale to real-world problems. Interestingly, this is not, in my opinion, the greatest impediment for the uptake of agent technology, but rather that those techniques do not fit immediately with existing programming languages and design practice as normally taught in Universities and used in Industry. Unfortunately, for many years, little attention was given by multi-agent systems researchers to the software engineering aspects of such systems which would allow real-world applications make use of the sophisticated techniques emerging from theoretical research in the area (for an overview see [21]).

Recent years, however, have seen an incredibly fast increase in the volume and quality of the research on agent-oriented computing as a viable, practical approach to designing and programming sophisticated (distributed) systems. In [4, 12], a number of agent-oriented programming languages for multi-agent systems were introduced ,Äì and many others have appeared in the literature, drawing inspiration from varying sources (such as logic programming, functional programming, process algebra, etc.). Most of these language either directly use or were influenced by the BDI (Belief-Desire-Intention) agent architecture that originated in the AI literature. In any case, BDI-based agent programming represents a significant compromise on some of the original AI expectations of BDI agents, but very promising as a programming technique for multi-agent systems. For example, programmers write libraries of plans that agents use at run time, but they (typically) do no planning as such. On the other hand, the interpreter of a BDI agent language provides sophisticated means of executing plans that agents use to achieve long-term goals as well as to react to changes perceived in the environment where they are situated; they also help an agent to handle concurrent foci of attention when various changes in a very dynamic environment require the agent to take action. Clearly, agent programming draws heavily on ideas originated in AI, but also this programming approach has a lot to give back to AI as it provides an appropriate system design for using and experimenting with AI techniques.

Frame systems provided some of the fundamental ideas leading to the development of object orientation. It seems to me that history is repeating itself (as in the Hegelian dialectical spiral of historical progression) in that AI has again provided the basis—this time around perhaps much more unashamedly—for a major shift in the established programming paradigm. It also seems to me that logic-based programming was attempted at the wrong level, essentially contraposing the long hierarchy of programming abstractions that had been built on top of each other since the early days of computer science. If one could claim that logic programming does not work well as a general-purpose programming language, it is certainly more difficult to deny that it sits very comfortably at the level of practical reasoning (i.e. reasoning about courses of action as opposed to beliefs) done by an autonomous agent controlling computational resources and services that are themselves often better programmed using object orientation, and taking advantage of the massive number of libraries, tools, methodologies, and legacy systems available at that level. The use of logic-based approaches at the agent abstraction level also makes this paradigm lend itself more easily to combining with other current technology trends such as (semantic) web services and the semantic web.

Our endeavour has been in selecting particular theories that worked well together, and creating or choosing existing simplifications of those to develop a practical approach for putting together the various "pieces of the puzzle" of available agent techniques. I believe it gives very powerful abstractions for programmers, that work well on top of all the existing abstractions used in object orientation. Contrary to what some had expected, learning how to program does not get any easier just because of the folk psychology element of the BDI architecture. Rather, it requires experience with existing abstractions and it does take time to learn all the new ones. However, with the ever increasing sophistication of computational systems for current and future social demands, it is hard to deny that we will need more powerful abstractions to cope with this trend. Starting from the definition of an abstract programming language called AgentSpeak(L) [18], we have worked on various extensions of the language, for example to allow agents programmed in that language to communication using a speech-act based agent communication language [20], plan (i.e., know-how) exchange among cooperating agents [2], the use of ontologies and ontology reasoners [16], belief revision for AgentSpeak agents [1], plan patterns for advanced BDI (declarative, goal-oriented) programming [13], and reasoning about goals [19]. We also did work to formally relate the AgentSpeak language constructs and the interpreter data structures to the modalities of BDI logic, using the operational semantics of the language [9]. This is important for ongoing work on formal verification which I briefly refer to later in this article.

While that work was mainly theoretical, it served as a basis for the development of a very elaborate, highly customisable platform for developing multi-agent systems called , which was done mostly in joint work with Jomi Hübner, and made available open source at http://jason.sf.net. That work culminated in the recent publication of a textbook to put together all those contributions in a practical and accessible way [8]. Various features of the *Jason* platform make it useful for modelling and simulation of social phenomena. Most of the available tools for social simulation allow only very simple (i.e., reactive as opposed to cognitive) agents to be used. Prominent researchers in social simulation have advocated the need for cognitive agents in certain advanced types of social simulation [10]. We are, therefore, in ongoing work [7, 3], incorporating features in *Jason* which will make it an even more powerful platform for developing software based on multi-agent systems techniques, but also facilitate its use as a tool for social simulation, in particular for modelling human goal-directed behaviour. Two examples of such work are discussed next. In *Jason*, environment models have to be programmed in Java. For some applications, a more declarative, high-level language could be very useful. This was the motivation that led to the development of an environment description language which has recently been extended [17] to allow environment descriptions to have objects containing social norms that are to be observed only within the confines of an environment location, possibly where an institution or organisation is situated (similar to "please refrain from smoking" or "keep silent" signs).

An important part of agent-oriented software engineering is related to agent organisations, which has received much research attention in the last few years. There is ongoing work on allowing specifications of agent organisations (with the related notions of roles, groups, relationships between groups, social norms, etc.) to be used in combination with *Jason* for the programming of individual agents. The particular organisational model being integrated with our approach is called Moise+; an initial integration with *Jason* is discussed in [14], and available from http://moise.sf.net. One of the advantages of the approach is that the organisation specification is available for the agents to access and possibly modify at run time.

As mentioned earlier, there is also ongoing work on reasoning about goals [19] which in the future we aim to incorporate into *Jason*. This type of reasoning allows agents to consider the interactions between the various plans they have in their plan library and the consequences of plan choices for the goals an agent has to achieve or maintain. Other recent work is looking at how to implement in practice the use of ontological information as part of an agent program [15]. By making use of existing ontological knowledge available on the web, we can make agent programming simpler, besides facilitating the integration of agent programming for semantic web

Programming and verifying complex systems (cont.)

applications and applications that make use of (semantic) web services.

An important use of logic-based techniques in the context of software development in multi-agent systems (in particular with so many of its areas of application requiring "dependable systems") is for formal verification. In previous work, we devised model checking techniques for multi-agent systems programmed in AgentSpeak [6], and also state-space reduction techniques to be used in model checking AgentSpeak programs [5]. While in that work we were specifically interested in model checking AgentSpeak multi-agent systems, in a recent ongoing project, joint with Michael Fisher, we are interested in developing techniques that would allow model checking for a variety of agent-oriented programming languages [11]. In this work, we use Java PathFinder (JPF)—see http://javapathfinder.sf.net/)—as the target model checker, and we formally specified and developed (in Java) a library of general agent-related concepts, called AIL. The idea is to optimise JPF to model check systems that use AIL, and to automatically translate different agent programming languages to run on top of AIL, facilitating the use of JPF for model checking agent programs written in various modern agent programming languages.

References

[1] N. Alechina, R. H. Bordini, J. F. Hübner, M. Jago, and B. Logan. Automating be- lief revision for AgentSpeak. In Proc. of DALT 2006, LNCS 4327, pp 61-77. Springer, 2007.

[2] D. Ancona, V. Mascardi, J. F. Hübner, and R. H. Bordini. Coo-AgentSpeak: Cooperation in AgentSpeak through plan exchange. In Proc. of AAMAS-2004, pp 698-705. ACM Press, 2004.

[3] R. H. Bordini, A. C. da Rocha Costa, J. F. Hübner, Á. F. Moreira, F. Y. Okuyama, and R. Vieira. MAS-SOC: a social simulation platform based on agent-oriented programming. Journal of Artificial Societies and Social Simulation, 8(3) - JASSS Forum,

2005. <http://jasss.soc.surrey.ac.uk/8/3/7.html>

[4] R. H. Bordini, M. Dastani, J. Dix, and A. El Fallah Seghrouchni, editors. Multi-Agent Programming: Languages, Platforms and Applications. Springer, 2005.

[5] R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. State-space reduction techniques in agent verification. In Proc. of AAMAS-2004, pp 896-903. ACM Press, 2004.

[6] R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. Verifying multi-agent programs by model checking. Journal of Autonomous Agents and Multi-Agent Systems, 12(2):239-256, 2006.

[7] R. H. Bordini and J. F. Hübner. Agent-based simulation using BDI programming in Jason . In A. M. Uhrmacher and D. Weyns (edts), Agents, Simulation and Applications. Taylor and Francis, 2008. To appear.

[8] R. H. Bordini, J. F. Hübner, and M. Wooldridge. Programming Multi-Agent Systems in AgentSpeak Using Jason. Wiley Series in Agent Technology. John Wiley & Sons, 2007.

[9] R. H. Bordini and Á. F. Moreira. Proving BDI properties of agent-oriented programming languages: The asymmetry thesis principles in AgentSpeak(L). Annals of Mathematics and Artificial Intelligence, 42(1,Äì3):197-226, 2004.

[10] C. Castelfranchi. The theory of social functions: Challenges for computational social science and multi-agent learning. Cognitive Systems Research, 2(1):5-38, 2001.

[11] L. A. Dennis, B. Farwer, R. H. Bordini, M. Fisher, and M. Wooldridge. A common semantic basis for BDI languages. In Proc. of ProMAS-2007. To appear in Springer, LNCS series, 2008.

[12] M. Fisher, R. H. Bordini, B. Hirsch, and P. Torroni. Computational logics and agents: a road map of current technologies and future trends. Computational Intelligence, 23(1):61-91, Feb. 2007.

[13] J. F. Hübner, R. H. Bordini, and M. Wooldridge. Programming declarative goals using plan patterns. In Proc. DALT 2006, LNCS 4327, pp 123-140. Springer, 2007.

[14] J. F. Hübner, J. S. Sichman, and O. Boissier. Developing organised multi-agent systems using the MOISE+ model: Programming issues at the system and agent levels. International Journal of Agent-Oriented Software Engineering, 1(3/4):370-395, 2007.

[15] T. Klapiscak and R. H. Bordini. JASDL: A practical programming approach combining agent and semantic web technologies. In Proc. DALT-2008, to held as part of AAMAS-2008, Estoril, Portugal,

2008.

[16] F. Moreira, R. Vieira, R. H. Bordini, and J. F. Hübner. Agent-oriented programming with underlying ontological reasoning. In Proc. DALT-05, LNCS 3904, pp 155-170. Springer, 2006.

[17] F. Y. Okuyama, R. H. Bordini, and A. C. da Rocha Costa. Spatially distributed normative objects. In Proc. of COIN-2006, LNCS 4386, pp 133-146. Springer, 2007.

[

18] A. S. Rao. AgentSpeak(L): BDI agents speak out in a logical computable language. In Proc MAAMAW, LNCS 1038, pp 42-55. Springer, 1996.

[19] P. H. Shaw, B. Farwer, and R. H. Bordini. Theoretical and experimental results on the goal-plan tree problem (short paper). In Proc. AAMAS-2008. IFAAMAS, 2008. To appear.

[20] R. Vieira, Moreira, M. Wooldridge, and R. H. Bordini. On the formal semantics of speech-act based communication in an agent-oriented programming language. Journal of Artificial Intelligence Research, 29:221-267, 2007.

[21] M. Wooldridge. An Introduction to MultiAgent Systems. John Wiley & Sons, 2002.

Note

The work on model checking agent programming languages is supported by EPSRC grant numbers EP/D054788 (Durham) and EP/D052548 (Liverpool).

Rafael H. Bordini
University of Durham
R.Bordini@durham.ac.uk

Computational Creativity

The Science Museum in London once exhibited some interesting machines made from Meccano which were able to perform complex mathematical calculations. As these machines were built in the 1930s, the Meccano magazine from June 1934 speculated about the future in an article entitled: "Are Thinking Machines Possible?" They couldn't have possibly known the impact the computing age would have, but they were already certain about one thing: to my horror, at the end of the article, the author said: "Truly creative thinking of course will always remain beyond the power of any machine".

A full ten years before programmable computers were born, and decades before they started to show any signs of intelligence, people were already condemning them to an uncreative future. This reaction is entirely understandable. We can't run as fast as tigers, swim like dolphins or climb like goats. But we have our smarts—that's our thing. It's been hard enough to admit that computers might rival us at tasks which require intelligence. So we won't easily relinquish our position as the only creative beings on the planet.

Under this weight of prejudice, research into computational creativity has lagged behind other areas of AI research a little. However, there is a small band of us who pursue the goal of getting computer programs to creatively produce poems, sonatas,paintings, theorems, jokes, and much more. We've recently reached the stage where there is a sufficiently large number of such programs for us to be able to compare and contrast them in a meaningful way. This has enabled us to begin to come to consensus about the main issues in computational creativity.

Firstly, we've realised that we're working in a different paradigm to the majority of AI researchers. When faced with an intelligent task to perform, AI people generally think of the task in terms of solving a problem. That problem might be planning a route from London to Liverpool, proving a mathematical theorem, or recognising hand-written words. But in each case, AI techniques are developed that can (hopefully) solve the problem as well as or better than humans.

Computational creativity researchers, on the other hand, work in an artefact generating paradigm. Here, the task is to generate artefacts of real value to someone. Those artefacts may be plot lines for a play, a mathematical theorem, or a harmonisation for a Bach chorale. There is a lot of overlap between the two paradigms, but they each have their peculiarities. For instance, problem solving AI programs know when to stop: when they have solved the problem to a satisfactory degree. With AI artefact generation programs, however, it's often not clear when to stop them (humans have this problem, of course: it's never easy to know when to stop painting a picture).

Another difference is aesthetics. Usually, the value of solutions generated by problem solving programs is measured only by the success of the solution to the problem. Nice and simple. With artefact generation, however, there are many competing ways of assessing the artefacts and different aesthetic considerations have to be taken into account. Many of these will be specific to the musical, artistic, literary or scientific application at hand. Others will be more general, such as whether the artefact is novel, surprising, or evokes an emotional response.

So far, to engineer our creative programs, we've stolen anything we can from Artificial Intelligence and elsewhere. There are attempts underway to look at the methods being used and characterise some of them in terms of the kind of search they perform: do they just look very hard through thousands of similar artefacts for a good one, or do they somehow transform the way in which artefacts are generated, which might be considered more creative. Often, if the methods involve some randomness, or are so complicated that we can't explain their actions, this may increase our perception of a program's creativity. If we can completely describe how a program produced an artefact, the chances are that no matter how pleasing the artefact is, we would not think of the program as being particularly creative.

Usually, creative artefact generation programs need three types of methods: those which mimic a human skill; those which mimic human appreciation of artefacts; and those which mimic our imagination. Imagine an artist missing one of skill, appreciation or imagination. Without skill, they would never produce anything. Without appreciation, they would produce things which looked awful. Without imagination, everything they produced would look the same. It is usually the imagination part that we have the most difficulty simulating, and often we have to approximate this by getting our programs to search through millions of possibilities.

A very important issue is the assessment of the creativity of programs. In the problem solving paradigm, if a new program solves a previously unsolvable problem, or solves a bunch of problems faster than all other programs, then clear progress has been made. As creativity is such a subjective notion (is your child really as creative as you say?), it's much more difficult for us to compare the creative abilities of different programs. However, much progress has been made towards telling whether we should use the word creative to describe a program and telling whether one artefact generation program is performing more creatively than another.

It is wrong to think that when computers start acting creatively,artists, musicians, and poets will be out of a job. Why on earth would Lucien Freud stop painting just because a computer can paint as well? Moreover, people will always appreciate the blood, sweat, and tears expended by creative people in producing their works. With this in mind, much research in computational creativity has looked at how to enhance or supplement the creativity of people undertaking creative endeavours. In the same way that a composer expects creativity in a performer, we can begin to expect that computers will act as creative collaborators in our projects.

As a society, we leave behind our creations, so surely it is worthwhile having more creativity in the world. And think of the great gadgets we can have if computers become truly creative: a website to generate topical jokes for the speech you're giving tonight, IPods which can generate entirely new tracks to suit your mood, and fridges which can concoct a delicious recipe to fit their meagre contents. As the phrase goes, we are limited only by our imagination. But there is the crux of the matter: if we have such a limitation, can we really afford to ignore computational creativity?

Simon Colton
Imperial College, London
sgc@doc.ic.ac.uk
http://www.doc.ic.ac.uk/~sgc/

Conference Report: ICANN 2007

ICANN is one of the main European conferences on neural networks and it took place this year from the 9-13th September 2007 in Porto, Portugal. The conference program included one day of tutorials and half a day of workshops, which gave ample scope for discussions. The main emphasis of the conference was on machine learning, and there were also sessions on other topics, such as spiking neural networks and cognitive systems.

One of the most interesting presentations was given by Felix Schürmann on the Blue Brain Project. This is taking place at the Ecole Polytechnique Fédérale de Lausanne in Switzerland and its overall goal is to reverse engineer the mammalian brain. The current focus of the Blue Brain Project is on accurately reproducing the behaviour of in-vitro neural tissue from the rat and clarifying what kinds of data need to be recorded during in-vitro experiments (to enable in-vitro results from different groups to be reproduced and more systematically compared).

The Blue Brain project is starting with a single cortical column, which is the basic functional unit of the cortex. Cortical columns occupy a cylindrical volume 0.5 mm wide by 2 mm high and contain around 10,000 neurons interconnected with 30 million synapses. This microcircuit is repeated millions of times across the cortex and it is similar between species - the main interspecies difference is that human brains have many more cortical columns than smaller mammals. The Blue Brain Project has been simulating single cortical columns using biologically accurate neurons with realistic connectivity. The simulations are being carried out on an IBM Blue Gene supercomputer, which contains 8192 processors and 2 TB of RAM —a total of 22 x 1012 teraflops processing power. The simulation generates 160 GB/s of data and the team has had to develop strategies to store and process this information and they have also created some impressive 3D visualizations. The software used for the simulation is a combination of the large scale Neocortical Simulator [7] and NEURON [8]. The first simulation of the rat cortical column was carried out in 2006 and it is currently running at about two orders of magnitude slower than real time.

Another highlight of the conference was the presentation on 'Modelling Consciousness with Neural Architectures' by Igor Aleksander, which gave an overview of his work on machine consciousness. Aleksander aims to understand consciousness by building systems that are not living and might be attributed some form of consciousness, and his theoretical approach is based around five axioms that he claims are minimally necessary for consciousness[1] These axioms are depiction, imagination, attention, planning and emotion, with depiction being the most important. According to Aleksander, these axioms are a preliminary list of mechanisms that could make a system conscious, which should be revised as our knowledge of consciousness develops—a useful starting point that can be used to test ideas and develop the field. These axioms were deduced by Aleksander using introspection, and he also identifies neural mechanisms that could implement them in the brain.

Aleksander has developed a kernel ar-

ICANN07 (cont.)

chitecture that covers all five of his axioms. This includes a perceptual module that depicts sensory input, a memory module that implements non-perceptual thought for planning and recall of experience, an emotion module that evaluates the 'thoughts' in the memory module, and an action module that causes the best plan to be carried out. Aleksander has built a number of brain-inspired implementations of this kernel architecture with the Neural Representation Modeller (NRM) [9], which uses weightless neurons containing lookup tables that match input patterns to an output response. During training, these neurons store the link between each input pattern and the specified output; during testing, the neurons produce the output of the closest match to a known input pattern or a random sequence of 1s and 0s when there is more than one match. These neurons are assembled into large recurrent networks and trained using the graphical and scripting abilities of NRM. These brain-inspired simulations of the kernel architecture are minimal implementations of Aleksander's five axioms, and so they have the potential for phenomenal consciousness according to the axiomatic theory. Full details about how the kernel architecture implements the axioms can be found in a recent paper by Aleksander and Morton [2]. More information about other work in machine consciousness is available [5].

A third interesting topic was presented by Sophie Deneve, whose work bridges the gap between more abstract interpretations of the nervous system, such as Bayesian approaches, and the behaviour of real neurons. She is particularly interested in how probabilistic knowledge and uncertainties are represented by cortical neurons and how biological neural networks perform or approximate Bayesian inference. For example, in a recent paper with Pouget,4 she showed how basis functions can be combined with a Bayesian framework to provide optimal multisensory integration, and she has also demonstrated how iterative basis function networks can learn forward and inverse models simultaneously, which could be applied in sensorimotor integration and motor control. Her current work is focused on testing the predictions made by her Bayesian and iterative basis function models using psychophysical experiments and neurophysiological recordings.

References

1. Aleksander, The World in My Mind, My Mind in the World: Key Mechanisms of Consciousness in People, Animals and Machines, Exeter: Imprint Academic, 2005.
2. Aleksander and H. Morton, Why Axiomatic Models of Being Conscious?, in Ron Chrisley, Rob Clowes, and Steve Torrance (eds.), Journal of Consciousness Studies special issue on Machine Consciousness, forthcoming in 2007.
3. Blue Brain Project website: http://bluebrain.epfl.ch.
4. S. Deneve and A. Pouget, Bayesian multisensory integration and cross-modal spatial links, Journal of Physiology (Paris) 98(1-3): 249–258.
5. D. Gamez, Progress in Machine Consciousness, Consciousness and Cognition doi:10.1016/j.concog.2007.04.005, forthcoming.
6. ICANN conference website: http://www.icann2007.org.
7. Neocortical Simulator: http://www.cse.unr.edu/brain/FILES_HTML/welcome_frame.html.
8. NEURON simulator: http://www.neuron.yale.edu/neuron.
9. NRM: http://www.iis.ee.ic.ac.uk/eagle/barry_dunmall.htm.

David Gamez
University of Essex
daogam@essex.ac.uk

Conference Report: Thirteenth Portuguese Conference on AI

The 13th Portuguese Conference on Artificial Intelligence had a number of commendable papers to offer. In this article, three papers are selected for presentation, two of which are from the 2nd Intelligent Robotics (IROBOT) workshop of the conference, and one from the 3rd workshop on Artificial Life and Evolutionary Algorithms (ALEA).

The paper "An Omnidirectional Vision System for Soccer Robots" by A.J.R. Neves, G.A. Corrente, and A.J. Pinho presented a complete and efficient vision system developed by the University of Aveiro in Portugal. Their robotic soccer team, CAMBADA, earned the first prize in their national robotic soccer competition, ROBOTICA 2007. The importance of this paper lies in the proven effectiveness of their system and algorithms, especially regarding object detection and the robots' self-localization. Their system is an omnidirectional vision system that can find all artifacts in the robot's world. It consists of a firewire camera mounted vertically on top of the robots and a hyperbolic mirror above this camera, which reflects the entire 360 degree field around the robot. Additionally, the system relies on a real-time database, which is shared by all players. The system speed is due to its distributed architecture that allows color extraction, object detection and image visualization to run in parallel. Color extraction is using radial search lines, as opposed to the more frequent range scan, and this makes processing time independent of the world and more fitting for Real-Time Systems, like this one. Object segmentation is simplified by the fact that a pixel's color is a strong hint for what the object is, since they're all color coded: black robots play on a green field with white lines, passing around an orange ball, trying to get it in to the yellow and blue goals. The system detects efficiently all these objects via a set of simple algorithms, which also calculate the position of these objects within certain limits in an angular representation. The authors claim that future work will involve improving the current object detection algorithm to also include shape, instead of only color, information, and also the creation of novel algorithms for camera and color calibration.

Portuguese Conference on AI (cont.)

The paper "Generalization and Transfer Learning in Noise-Affected Robot Navigation Tasks" by Lutz Frommberger from Universitat Bremen examines the transferring of knowledge learned. The ability to learn from one task, generalize and apply to other problems is fundamental in human intelligence, and the reason of the importance of this particular paper. Specifically, the author looks into the transfer of the knowledge of an acquired, by reinforcement learning, optimal policy for a goal-directed navigation task, to other tasks with varying conditions (generalization) or to different, but structurally similar tasks. In this work, the author has shown that although Reinforcement Learning is the obvious technique to use in such tasks, the agents lack an intelligent understanding of the general structure of the world. The author has chosen to represent the structural knowledge of the robot via a recent technique specifically designed for indoor robot navigation, the Relative Line Position Representation (RLPR). RLPR divides the system state into the "goal-directed behavior" and the "generally sensible behavior." The technique assumes that the walls are the only structural elements used for navigation and that they are of a specific color. The advantages of this method are the small and discrete state space, the fact that there is no need for function approximation, the fast and robust learning and the readiness of collaborating with an arbitrary learning algorithm. Moreover, it is shown that RLPR is very successful in noisy environments. Due to its abstraction it avoids severe misclassifications, such as false line or landmark detection. Finally, the paper presents a new algorithm for transferring knowledge learned by such navigational tasks into arbitrary environments with particularly good results, even though they are all simulation results. The author suggests that the next step will be testing this strategy with real robots.

N. Pillay and W. Banzhaf were the authors of "A Genetic Programming Approach to the Generation of Hyper-Heuristics for the Uncapacitated Examination Timetabling Problem." As the authors also state in their paper, the importance of this paper lies within the fact that most evolutionary techniques to solving the examination timetabling problem go about it by evolving solutions to the problem. In this paper, the authors present a new genetic programming technique with which the compilation of heuristics used to solve the specific problem is instead being evolved. The results were encouraging, since the Genetic Programming-based system outperformed the other systems used for comparison on most of the problems the systems were tested on.

The papers presented above were all selected for a different reason, and in no way does their selection for inclusion in this particular review mean that they were in any way better or more interesting than other work presented in Guimaraes, Portugal during the conference, which was held on December 3-7, 2007.

Konstantinos Bousmalis
University of Edinburgh

Conference Report: World Congress of Engineering and Computer Science

My personal work has centred on using data for making intelligent decisions—using artificial intelligence and neural network techniques where possible. I recently presented a paper discussing the use of Kohonen Self-Organising Maps and Neural Networks to target vehicles for inspection upon entering the United Kingdom through the Channel Ports of Dover and Ramsgate.

Thanks in part to a generous grant from the AISB I was able to present this paper the World Congress of Engineering and Computer Science (WCECS 2007) at the University of California, Berkeley, San Francisco, California. Of the many interesting and varied papers discussed at the conference, three stood out in particular. The most fascinating, in my opinion was one of the keynote speeches, which discussed a combination of utilising sensor data and advanced simulation in a host of applications the speaker referred to as Dynamic Data Driven Application Systems (DDDAS). Professor Craig Douglas's latest work had centred on predicting the behaviour of wild fires - prescient as at the time there were many out of control blazes in South California.

DDDAS differed from conventional modelling approaches in that the system uses real-time weather data images and sensor streams. The system changes its forecasts as new data is received, using a long term running simulation that self corrects using out of order and imperfect sensor data. The DDDAS version replaces the existing simulation which was previously run using data only in initial conditions.

The speaker suggested that much more accurate models could be generated by DDDAS as the paradigm is able to both dynamically incorporate additional data into an executing application and use this new data to reactively steer the measurement process.

Much of the speakers work had focused around managing disaster recovery situations utilising this methodology, and as a consequence his work was well funded by the US Government. It is hoped that a suite of technologies developed through a number of DDDAS projects can be put together to enable end-to-end, prediction, impact, response and mitigation for critical events, like natural and man-made disasters.

World Congress of Engineering and Computer Science (cont).

Professor Douglas's project developed new methods and technological infrastructure, including a prototype disaster response test bed which combined an actual evolving crisis event in-tandem with a simulation framework where the on-going event data was continually and dynamically integrated with the on-line simulations. The computational aspects of the framework developed under the project are built using peer-reviewed published models, driven and validated by real world data. The framework encompasses data and models from the public domain, such as GIS data, fire-models, models for blast and impact dynamics, models of building structures under stress, and models of wind and other environmental influences. Real-time sensor data, video streams, and human inputs from the actual or exercise scenarios can be bridged to the virtual environment. Agent-based emergency response simulations can dynamically invoke models, to represent a given scenario and they also encompass agent modelling of model human behaviour. All this information and data—real-time, historical, and computed—may be transmitted to users to enable coordination of a comprehensive plan for responding to the crises, minimising the event's impact, and assisting in the recovery from a crisis event.

I felt these embryonic DDDAS systems had much potential, certainly in fields such as traffic and railway management, as well as the disaster management applications suggested by the speaker.

Although not directly related to my field, I found some research completed by the Jisan Research Insititute regarding Swarm Engineering and Multi Agent Systems fascinating and of no doubt of interest to the AISB. Dr. Kazadi described swarms as groups of bi-directionally communicating autonomous agents which exhibit emergence, which allows them to undertake actions that are not explicitly part of their control algorithm. The problem that the research was attempting to address was the lack of a generally principled approach to swarm design; whilst many researchers have built preliminary systems for monitoring or understanding these emergent behaviours, these studies do not yet generalize a methodology that works for a large number of swarm systems. As a result no particular method exists for generating swarms of particular design.

The research has its roots in work which advocated the creation of simple robots with simple behaviours in the place of more complex abstract models for intelligent behaviour. Swarm engineering consists of two related steps. The first step is to propose an expression of the problem which leads to a set of conditions on the individual agents which, when satisfied, will complete the task. The second step is to produce a behaviour or set of behaviours for one or many robots which accomplishes these requirements. Much of Dr Kazadi's work centred on accurate and robust definition of these initial conditions. One of his group's research pieces attempted to actually model market economics utilising a swarm—real economic systems being systems of autonomous agents with bi-directional communication and therefore swarm like. A number of studies have been made which used agent based simulations in which interactions between agents define what the economy will do, this simulation appeared to be unique as it applies a method of generating the global behaviours and then designing the system around that behaviour.

The authors showed some success in their work, and many interesting observations; for example the "Invisible hand of the Market" could be explained by initial conditions, particularly an agent's sensitivity to the profit margin a vendor was receiving. This factor had a more powerful affect then vendor competition, and the author postulated that such initial conditions could be responsible for the dot com bust and the American housing slow down.

The final papers which interested me were the large number of authors discussing efficiency, accuracy and reliability tweaks to a number of relatively mature supervised data analysis methods, such as decision tree optimisation and a well presented paper regarding an improved Fuzzy ART neural network algorithm for predictive analysis. I think that it is clear that with so much data in organisations, more people are trying to do something useful with this data, and a number of experts can see opportunities in this growing field.

Grant Brown

Society News

AISB Members Only Area

A few of you may be aware that the AISB website includes a small members only area (http://www.aisb.org.uk/membersonly/) from which it is possible to access PDF copies of back issues of AISBJ, details of some of the journal discounts AISB members are entitled to, and a couple of items of miscellaneous information and resources. Up until recently this area was accessible using a catch all username and password.

We have now extended this system so that the members only area consults a limited database of member details. This means that individual members will need to create their own password for access to the area in future, but it also allows individual members to check (and alter) some of the information we hold on them. This information falls into two categories: contact information is held centrally and cannot be altered directly from the web, however a member can use the page to send a message to the AISB administrator requesting a change of details. The second category is the information we use to create a publicly accessible page of our members' web pages and research interests (http://www.aisb.org.uk/membership/members_interests.php). It is now possible to directly update this information.

To log on to the members only area you will need your AISB membership number and the email address we use to contact you. This information will allow you to set up a password, once this is done you may log on with just your email address and the password. Logging on takes you to the main members only page, but there is a link at the top of the page "Edit your Details" which will allow you to check and alter our information.

As this system is very new please don't hesitate to report any problems or bugs.

We hope in future to link this system into an online membership payment and renewal system (using PayPal). We would also be interested in any input from the members about content they would like to see on the web which would be inappropriate for publicly viewable areas (e.g., sample teaching and open day materials).

Louise Dennis, AISB Webmaster

Recent Committee Changes

Following a call in September 2007, four nominations were received for committee membership. Since four committee places were available these were all elected unopposed. The new (and returning) committee members are Aladdin Ayesh (De Montfort), Louise Dennis (Liverpool), Berndt Farwer (Durham) and Manfred Kerber (Birmingham).

Louise Dennis, AISB Secretary

Olympics 2012

The SSAISB Committee has at various times raised the possibility of the society seeking to use the 2012 Olympics as an opportunity to inspire new research and public understanding on aspects of AI related to sport, games, security, infrastructure, etc. We have asked in particular for volunteers to form a task force to push the matter forward. However, in the absence of any volunteers we are now proposing to drop the matter, as the Committee is not itself in a position to act as a task force. We must leave any society members who would like to pursue Olympics-related angles of AI research or public understanding to do so by the means of individual efforts. We are also aware of a feeling on the part of some researchers that enough government resources are going into the Olympics already.

John Barnden, AISB Chair

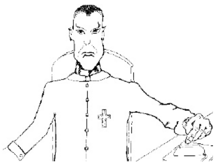

The Life of A. Hacker
by Fr. Aloysius Hacker

Cognitive Divinity Programme
Institute of Applied Epistemology

Episode 5: Four Seasons of AI

UK weather is famous for its variability – so is AI weather. From the mid-60s to the mid-80s, AI went through four seasons, but not in the usual order.

The AI Spring started in the mid-60s, as the diaspora from the demise of CATHOLIC™ (Church of Aloysius Theobald Hacker for Ordinations, Liturgy, Inquisitions and Christenings) founded new research groups in Edinburgh, Sussex, Essex and elsewhere. Unfortunately, relations with Mickey MacDonald's group were soon soured by his insistence on misnaming the field as Machine Intelligence and by his failure to acknowledge me for the loan of all those matchboxes. Naming his learning system MENACE™ (Mickey's Engine owing Nought to Aloysius's Creative Endeavours) made me particularly cross. MacDonald also resented my founding and chairing of the AISB™ (Aloysius IS Boss) society, whose remit was the promotion of AI in the UK.

But by the early 70s, the AI Winter was created by the infamous Lighthill Report. I must take this opportunity to refute the persistent rumours that I steered Lighthill to his conclusions in order to damage MacDonald's research funding. Nothing could be further from the truth – and in any case, he deserved it. Lighthill made a strong intellectual case against MacDonald's research programme. He accused AI roboticists of suppressed maternal yearnings in trying to build artificial people. They accused Lighthill of fear of premature ejaculation in studying hydrodynamics.

As the echoes of the Lighthill Report faded, AI Winter turned to AI Autumn. We took advantage of the more relaxed atmosphere to launch WHIM™ (the Werner Heisenberg Intelligence Module) a hardware plug-in that uses an embedded radiation source to import quantum uncertainty into all your AI systems, thus giving them genuine freewill. Unfortunately, we were forced to recall and withdraw WHIM™ due to its potential for generating undesired legal activity. I realised this potential after a junior Institute clerk plugged a WHIM™ into our payroll system - and I read my subsequent pay cheque. Before I could sack the clerk, she resigned and bought a luxury yacht to sail to the Caribbean.

The AI Summer was launched by the advent of expert systems, the Japanese Fifth Generation Programme and the UK's reaction to it in the ALVEY™ (Alvey, unlike Lighthill, Vigorously Enriches You) Programme. Wet Sock Ltd, was created by the Institute to take advantage of the new economic opportunities. Its main product was PINCH™ (Probabilistic Inference in Nought and Crosses by Hacker). This was selling well, but had to be withdrawn after legal action from MacDonald over infringement of copyright. Thus, the four seasons ended much as they had started.

Want more say?

If you have lots of ideas about what we should have in the Quarterly, contact the Editor about becoming an *Editorial Board Member.*

Want to advertise?

The AISBQ reaches hundreds of Artificial Intelligence and Cognitive Science researchers in the UK, Europe, and beyond. Advertising and general information is available via:
http://www.aisb.org.uk/aisbq/index.shtml

Want to write something?

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

The deadline for the next issue is:

Friday 16th May 2008