Symposium GC5: Architecture of Brain and Mind Integrating high level cognitive processes with brain mechanisms and functions in a working robot

Monday and Tuesday 3-4th April 2006
University of Bristol

Part of the AISB'06 Symposium:

Adaptation in Artificial and Biological Systems

Edited by Aaron Sloman[⋆†]

[⋆]School of Computer Science
The University of Birmingham UK
http://www.cs.bham.ac.uk/˜axs/
A.Sloman@cs.bham.ac.uk

Abstract

Overview of (provisional) symposium programme, and list of invited speakers, papers and abstracts

1 Invited speakers

Jackie Chappell

The University of Birmingham, UK
http://users.ox.ac.uk/.kgroup/jackie.html

How do animals gather useful information about their environment and act on it?
Slides online (PDF)

Abstract
Animals are much more successful than current robots in their ability to gather information from the environment, detect affordances, attribute causes to affects, and sometimes generate individually novel behaviour. What kinds of mechanisms might make this possible? I will discuss different mechanisms for acquiring information in animals, and their strengths and weaknesses given different life histories and niches. I will discuss experiments which have attempted to uncover the extent of animals' abilities to use information from their environment, and the mechanisms that might be used to accomplish this. The development of these kinds of competences (in evolutionary time and over the course of an individual's lifetime) is another interesting problem. Exploration and play seem to be very important for some kinds of behaviour, particularly flexible responses to novel problems, but there is also the possibility that animals come equipped with certain kinds of 'core knowledge', which might help to direct and structure the acquisition of more complex competences.

[1] R. C. Barnet, R. P. Cole, and R. R. Miller. Temporal integration in second-order conditioning and sensory pre- conditioning. Animal Learning and Behavior, 25:221–233, 1997.
[2] J. Chappell and A. Kacelnik. New Caledonian crows manufacture tools with a suitable diameter for a novel task. Animal Cognition, 7:121–127, 2004.
[3] S. E. Cummins-Sebree and D. M. Fragaszy. Choosing and using tools: Capuchins (*Cebus apella*) use a different metric than tamarins (*Sanguinus oedipus*). Journal of Comparative Psychology, 119(2):210–219, 2005.
[4] M. Domjan and N. E. Wilson. Specificity of cue to consequence in aversion learning in the rat. Psychonomic Science, 26:143–145, 1972.
[5] M. Hayashi and T. Matsuzawa. Cognitive development in object manipulation in infant chimpanzees. Animal Cognition, 6:225–233, 2003.
[6] A. A. S. Weir, J. Chappell, and A. Kacelnik. Shaping of hooks in New Caledonian crows. Science, 297(9 August 2002):981, 2002.

Mike Denham

Plymouth University
http://www.plymouth.ac.uk/pages/dynamic.asp
?page=staffdetails.id=mdenham.size=l

The role of the neocortical laminar microcircuitry in perception, cognition, and consciousness
Paper online (PDF)

Abstract
The talk will focus on the objectives of the EPSRC- and EU-funded projects I am involved in, including the role of the neocortical laminar microcircuitry in perception, cognition, and (dare I say it) consciousness. Of particular interest is the question of how the brain uses context to modify perceptual awareness, as illustrated for example by visual illusions.

Steve Furber

Manchester University
http://www.cs.manchester.ac.uk/apt/people/sfurber/

High-Performance Computing for Systems of Spiking Neurons
Paper online (PDF)

Abstract
We propose a bottom-up computer engineering approach to the Grand Challenge of understanding the Architecture of Brain and Mind as a viable complement to top-down modelling and alternative approaches informed by the skills and philosophies of other disciplines. Our approach starts from the observation that brains are built from spiking neurons and then progresses by looking for a systematic way to deploy spiking neurons as components from which useful information processing functions can be constructed, at all stages being informed (but not constrained) by the neural structures and microarchitectures observed by neuroscientists as playing a role in biological systems. In order to explore the behaviours of large-scale complex systems of spiking neuron components we require high-performance computing equipment, and we propose the construction of a machine specifically for this task - a massively parallel computer designed to be a universal spiking neural network simulation engine.

Jeffrey L. Krichmar

The Neurosciences Institute, San Diego, USA
http://www.nsi.edu/users/krichmar/

Brain-based devices for the study of nervous systems and the development of intelligent machines.
Paper online (PDF)

Abstract
Without a doubt the most sophisticated behavior seen in either biological or artificial agents is demonstrated by organisms whose behavior is guided by a nervous system. Thus, the construction of behaving devices based on principles of nervous systems may have much to offer. Our group has built series of brain-based devices (BBDs) over the last 14 years to provide a heuristic for studying brain function by embedding neurobiological principles on a physical platform capable of interacting with the real world. These BBDs have been used to study perception, operant conditioning, episodic and spatial memory, and motor control through the simulation of brain regions such as the visual cortex, the dopaminergic reward system, the hippocampus, and the cerebellum. Following the brain-based model, we argue that an intelligent machine should be constrained by the following design principles[1, 2]: (i) it should incorporate a simulated brain with detailed neuroanatomy and neural dynamics that controls behavior and shapes memory, (ii) it should organize the unlabeled signals it receives from the environment into categories without a priori knowledge or instruction, (iii) it should have a means to adapt the device's behavior, called value systems, when an important environmental event occurs, (iv) it should have a physical instantiation, which allows for active sensing and autonomous movement in the environment, (v) it needs to engage in a task that is initially constrained by minimal set of innate behaviors or reflexes, and (vi) it should allow comparisons with experimental data acquired from animal nervous systems. Like the brain, these devices operate according to selectional principles through which they form categorical memory, associate categories with innate value, and adapt to the environment. Moreover, this approach may provide the groundwork for the development of intelligent machines that follow neurobiological rather than computational principles in their construction.

1. Krichmar, J.L. and G.M. Edelman, Brain-based devices for the study of nervous systems and the development of intelligent machines. *Artif Life*, 2005. 11(1-2): p. 63-77.

2. Krichmar, J.L. and G.N. Reeke, The Darwin Brain-Based Automata: Synthetic Neural Models and Real-World Devices, in *Modeling in the Neurosciences: From Biological Systems to Neuromimetic Robotics*, G.N. Reeke, et al., Editors. 2005, Taylor & Francis: Boca Raton. p. 613-638.

Mark Lee

University of Wales, Aberystwyth
http://users.aber.ac.uk/mhl/

Developmental Robotics: an emerging paradigm for intelligent agents.

Abstract
Much research has explored the issues involved in creating truly autonomous embodied learning agents but only recently has the idea of a developmental approach been investigated as a serious strategy for robot learning. This is now emerging as a vibrant new research area. We examine the goals and methods of Developmental Robotics, and assess the current state of play. We give some requirements for a developmental system and relate these to the UK Grand Challenge 5 (The architecture of mind and brain) in terms of design issues for future robotic systems.

Peter Redgrave

Dept of Psychology, Sheffield University
http://www.shef.ac.uk/psychology/staff/academic/
peter-redgrave.html

Is it just a question of priority? Inspiration from the vertebrate basal ganglia.

Abstract
As soon as an agent, biological or physical, is provided with two or more parallel processing sensory or motivational systems that can guide movement there is a problem. Indeed, the same problem arises when a single system has the capacity to represent two or more features that can guide movement. If competing systems/features seek to guide incompatible movements (e.g. approach/avoidance) which one should be given priority ? Our supposition is that one of the vertebrate brain's fundamental processing units, the basal ganglia, has evolved to deal with such issues. Throughout the brains of vertebrates, parallel processing sensory, motivational and cognitive systems that can direct movement all provide phasic excitatory inputs to the basal ganglia. In turn, the basal ganglia output nuclei provide returning tonically active inhibitory connections to all input structures. Thus, the architectural principle describing basal ganglia connections with both cortical and sub-cortical systems is one of largely segregated parallel projecting loops. Winner-take-all selection is achieved by selective disinhibition of behavioural

systems targeted by basal ganglia output. Within this general framework, the implications of having of sub-cortical motivational systems (basic urges?) competing directly for behavioural expression with cortical (intellectual models of the world) will be considered. Additionally, an important quality of adaptive action selection systems is the capacity to adjust response probabilities (selections) based on reinforcement outcome. What is being reinforced by phasic dopaminergic neurotransmission within the basal ganglia is currently a topic of some dispute. Evidence will be considered suggesting dopaminergic teaching signals play a central role in identifying components of context and behaviour that are critical for causing unexpected biologically significant outcomes; in other words, learning those events for which the agent is responsible.

Murray Shanahan

Imperial College, London

http://www.doc.ic.ac.uk/˜mpsha/

Cognition and Consciousness: Is There a Fundamental Link?

Abstract

Some contemporary theories posit an intimate link between cognition and consciousness. For example, according to Baars's global workspace theory, the hallmark of consciously processed information is that it involves competition between and broadcast to widespread, multiple brain regions, while non-conscious information processing is localised. On this account, consciously processed information - because it integrates the activity of massively parallel processing resources, sifting out the relevant contributions given the ongoing situation for the organism - is cognitively efficacious in a way that non-conscious information processing is not. From the perspective of understanding the "architecture of mind and brain", this suggests that the issue of consciousness cannot be ignored, but should be a central element of the research programme of our Grand Challenge.

Aaron Sloman

University of Birmingham

http://www.cs.bham.ac.uk/˜axs/

Requirements for a robot with human child-like or crow-like visual and learning capabilities.

Abstract

I regard explaining vision as the hardest unsolved problem in AI and psychology. In part that's because identifying the functions of vision is so difficult. What are the functions of vision? There are many AI and robot systems that include a small set of visual abilities, e.g. the ability to analyse static or changing video images, usually in a very limited way, e.g. identifying instances of a few types of objects (e.g. vehicles), or tracking moving objects treated merely as blobs, or locating a robot relative to a previously stored map. A vast amount of research is driven by benchmark-based competitions which use arbitrarily selected collections of tests, based on human performances that are not understood at all, and which do not relate vision to its animal functions in enabling and controlling actions in a 3-D world.

In contrast, in humans and many animals, vision involves a rich and deep variety of functions, including perceiving static and changing 3-D structures, perceiving many kinds of positive and negative affordances, controlling actions both ballistically and online, and (at least in humans) interpreting the intentions of others, reading text and music, interpreting gestures, understanding how some mechanism works, solving mathematical problems with the aid of diagrams, and many more.

I previously[1] thought (like many others) that most of these functions could be explained in terms of perception of *structure* at different levels of abstraction processed concurrently, from which perception of affordances (information about what is and is not possible) could arise. Recently, while working on 3-D manipulation tasks for the CoSy robot project[2], I realised that most normal perception is not of structures but of complex *processes* (represented at different levels of abstraction concurrently). For instance, as two objects move in relation to each other each typically has parts that move in relation to other parts and to parts of the other object. Thus we are surrounded by "multi-strand" processes.

This simple obsrvation has profound implications regarding requirements for explanatory models, which I shall attempt to explain. This is closely related to the Emulation Theory of Representation pre-

[1] http://www.cs.bham.ac.uk/research/cogaff/crp/chap9.html

[2] http://www.cs.bham.ac.uk/research/projects/cosy/PlayMate-start.html

sented by Rick Grush in BBS 2004.[3] In particular, detailed analysis of requirements for such capabilities sheds light on the variety of types of learning that need to occur, e.g. as a result of active and playful exploration of the environment, and also points to some deep requirements for cognition that are ignored by many researchers who emphasise the importance embodiment, for example the requirement to perceive "vicarious affordances" (affordances for others, or for oneself in the past or future). The mechanisms that provide this have been misleadingly labelled "mirror neurons" rather than "abstraction neurons". Moreover analysis of the implications reveals that as far as humans and some other altricial species are concerned, the role of embodiment is reduced compared with, for instance insects, and sensorimotor contingencies are replaced by more "objective" condition-consequence contingencies.

An incomplete overview of background ideas for my talk can be found in a PDF presentation on vision as perception of processes[4], a PDF presentation on Two views of child as scientist: Humean and Kantian[5] and this web page on Orthogonal recombinable competences acquired by altricial species.[6]

Mark Steedman

University of Edinburgh

http://www.cogsci.ed.ac.uk/˜steedman/

Plans and the Structure of Mind and Language

Abstract

For both neuro-anatomical and theoretical reasons, it has been argued for many years that language and planned action are related. I will discuss this relation using a formalization related to those used in AI planning, drawing on linear and combinatory logic. This formalism gives a direct logical representation for the Gibsonian notion of "affordance" in its relation to action representation. Its relation to universal syntactic combinatory primitives implicated in language is so direct that it raises an obvious question: since higher animals make certain kinds of plans, and planning seems to require a symbolic representation closely akin to language, why don't those animals possess language in the human sense of the term? I will argue that the lexicalization of recursive propositional attitude concepts concerning the mental state of others provides almost all that is needed to generalize planning to fully lexicalized natural language grammar. The conclusion will be that the evolutionary development of language from planning may have been a relatively simple and inevitable process. A much harder question is how the capacity for symbolic planning evolved from neurally embedded sensory-motor systems in the first place.

Tom Ziemke

University of Skövde, Sweden

http://www.ida.his.se/˜tom/

Integrating Cognition, Emotion and Autonomy: Embodied Cognition in Organisms and Robots

Abstract

Much research in embodied AI and cognitive science emphasizes the fact that robots, supposedly unlike purely computational models of cognition, are "embodied". However, in this talk it is argued that the physical embodiment that robots share with animals provides only one aspect of the "organismic embodiment" that is underlying natural cognition, emotion and consciousness. The talk discusses the living body's relevance to embodied cognition and agency, and outlines a European research project that aims to model the integration of cognition, emotion and bioregulation (self-maintenance) in robots.

[3] http://journals.cambridge.org/action/displayIssue?jid=BBS.volumeId=27.issueId=03

[4] http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505

[5] http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0506

[6] http://www.cs.bham.ac.uk/research/cogaff/misc/orthogonal-competences.html

2 Poster summaries

Ruth Aylett

MACS, Heriot-Watt University, Riccarton, Edinburgh EH10 4ET

Emotion as an integrative process between non-symbolic and symbolic systems in intelligent agents

Abstract

This paper briefly considers the story so far in AI on agent control architectures and the later equivalent debate between symbolic and situated cognition in cognitive science. It argues against the adoption of a reductionist position on symbolically-represented cognition but in favour of an ac- count consistent with embodiment. Emotion is put forward as a possible integrative mechanism via its role in the management of interaction between processes and a number of views of emotion are considered. A sketch of how this interaction might be modelled is discussed.

Joanna Bryson

Artificial models of natural Intelligence, University of Bath, United Kingdom

Embodiment vs. Memetics: Is Building a Human getting Easier?

Abstract

This heretical article suggests that while embodiment was key to evolving human culture, and clearly affects our thinking and word choice now (as do many things in our environment), our culture may have evolved to such a point that a purely memetic AI beast could pass the Turing test. Though making something just like a human would clearly require both embodiment and memetics, if we were forced to choose one or the other, memetics might actually be easier. This short paper argues this point, and discusses what it would take to move beyond current semantic priming results to a human-like agent.

Nick Hawes, Aaron Sloman and Jeremy Wyatt

School of Computer Science, University of Birmingham

Requirements & Designs: Asking Scientific Questions About Architectures

Abstract

This paper discusses our views on the future of the field of cognitive architectures, and how the scientific questions that define it should be addressed. We also report on a set of requirements, and a related architecture design, that we are currently investigating as part of the CoSy project.

John Knapman

School of Computer Science

The University of Birmingham (Formerly IBM)

Integration and Decomposition in Cognitive Architecture

Abstract

Given the limitations of human researchers' minds, it is necessary to decompose systems and then address the problem of how to integrate at some level of abstraction. Connectionism and numerical methods need to be combined with symbolic processing, with the emphasis on scaling to large numbers of competencies and knowledge sources and to large state spaces. A proposal is briefly outlined that uses overlapping oscillations in a 3-D grid to address disparate problems. Two selected problems are the use of analogy in commercial software evolution and the analysis of medical im- ages.

Hagen Lehmann

University of Bath

How to build a brain - An Evolutionary Approach

(Poster only, no paper)

Abstract

To simulate processes in the human brain it will be very useful to create a model of how individuals plan and how the decision making in the process of action selection works. It would be important to understand what consciousness is, how it evolved and why humans have the ability to become aware of them selves. Since these phenomena are the results of evolutionary processes, the key to create something like a human brain is to replicate and model these processes. The question is not what material we would need to build a brain, the question is how and can we describe and model the complex operations the human brain does. In order to answer this question we have to gain consolidated knowledge about the natural processes, which made our brain evolve to the state it is in now. How are humans able to learn in a way, that they can adapt their behaviour very quickly in dynamic environments in order to respond to new situations. The concept of intentionality and the concept of consciousness seem to play an important role in this action selection process.To be able to make complex plans about what to do in the future and then overthrow them quickly if the situation makes it necessary an individual has to understand itself and other individuals in its group as intentional agents. The evolution of language is another key problem to be addressed thinking about how human intelligence works. In order to communicate ideas and solutions for certain problems it was necessary to develop a system of abstract vocal symbols.

Maria Petrou and Roberta Piroddi

Department of Electrical and Electronic Engineering, Imperial College London

On the structure of the mind

Abstract

The focus of any attempt to create an artificial brain and mind should reside in the dynamic model of the network of information. The study of biological networks has progressed enormously in recent years. It is an intriguing possibility that the architecture of representation and exchange of information at high level closely resembles that of neurons. Taking this hypothesis into account, we designed an experiment, concerning the way ideas are organised according to human perception. The experiment is divided into two parts: a visual task and a verbal task. A network of ideas was constructed using the results of the experiment. Statistical analysis showed that the verbally invoked network has the same topological structure as the visually invoked one, but the two networks are distinct.

Mark A. Wood

University of Bath, Department of Computer Science

Artificial models of natural Intelligence

Abstract

Social learning is an important source of human knowledge, and the degree to which we do it sets us apart from other animals. In this short paper, I examine the role of social learning as part of a complete agent, identify what makes it possible and what additional functionality is needed. I do this with reference to COIL, a working model of imitation learning.

Acknowledgements

Thanks to the euCognition network for funding this symposium, and for members of the EU-funded CoSy project for help with ideas and with arrangements for the symposium.