

Welcome to the Cogric book of abstracts

We hope that the abstracts give a flavour of the exciting, diverse and interesting talks throughout the workshop. The talks were selected not only to summarise the different fields and state-of-the-art work, but also to stimulate interesting discussions. After the plenary and paper abstracts, which are in order of presentation, there are volunteered poster abstracts and stimulating articles. We would encourage you to peruse the following as soon as possible:

Wednesday: Plenary one:

The sensorimotor approach to phenomenal consciousness, with empirical tests and possible applications to robotics.

J. Kevin O'Regan
Laboratoire Psychologie de la Perception
Centre National de Recherche Scientifique
Université Paris 5 - René Descartes

Phenomenal consciousness is the "raw feel" part of consciousness. It is the "what it is like" of experience: what makes sensory stimulation feel the way it does. It is considered by philosophers to be the "hard" part of the problem of consciousness: How could a system that functions using known physico-chemical mechanisms actually feel anything at all?

Most approaches to this question assume that some new brain mechanism will one day be discovered that generates feel. Perhaps feel arises from special oscillations, reverberations, quantum mechanisms or phase transitions in brain dynamics...

But I argue that there is a different way of thinking about feel that dissipates the problem. This new way of thinking takes feel not to be something that is generated by the brain. Instead feel is taken to be a *sensorimotor skill* that the organism is engaged in. Thinking about feel in this new way evades the logical and philosophical problems generally associated with phenomenal consciousness. Furthermore, the approach has the advantage that it makes empirical predictions that can be tested experimentally. I will describe such experiments on change blindness, sensory substitution, and adaptation to color. I shall also show how the perceptual structure of color space is more accurately predicted from the sensorimotor approach than from the normal neurophysiological approaches to color vision.

The sensorimotor approach can also be applied to problems that may be relevant in robotics: how can one design a system which must successfully adapt to new and unknown environments, all the time making use of unreliable sensors and effectors? I shall illustrate how the sensorimotor approach suggests a method for an agent to deduce the geometry of space using unknown sensors and effectors. I shall also show a method of dimensionality reduction which makes use of the agent's actions to provide a more natural metric for organizing sensory inputs than is usually used. Finally I shall show some results on the perceptual structure of color space that may be useful in artificial vision.

Plenary two:

Cognitive Robotics and Global Workspace Theory

Murray Shanahan

Baars's global workspace theory (GWT) has been highly influential within the scientific study of consciousness. It is also potentially important to cognitive robotics - partly because of its basis in a parallel computational architecture, but chiefly because it advances the idea that consciously processed information, construed according to the theory, is cognitively efficacious in a way that non-consciously processed information is not. After a short overview of the rudiments of GWT, this talk

will suggest possible fundamental links between consciousness and cognition relating to the so-called frame problem. Finally, a cognitive architecture will be outlined that combines a global workspace with an internally closed sensorimotor loop.

Papers

Virtual agency, embodiment and analgesia in phantom limb pain.

J Cole, G Austwick, S. Crowle, C Dawson, Z Zhang and R Wynne, University of Bournemouth.

Phantom limb pain (PLP) had been considered due to loss of sensory input and central plasticity, though Ramachandran, following use of a mirror box (Ramachandran and Rogers-Ramachandran, Proc Roy Soc B, 1996, 263, 377-386), suggested that it may result from mismatch between sensory input and motor intention.

We have developed a virtual arm seen on a screen, or in VR-spectacles, which moves in real time relation to movement of the amputee's stump recorded using a magnetic motion sensor. The arm moves to a table and grasps an apple as the subject guides his stump or shoulder forward and medially.

Six patients with severe PLP (aged 32 - 87 and with forequarter to mid humerus amputations 2 months to 15 years previously) have tried this for several hours over 1-2 days. Four learnt to move the virtual arm and felt their phantom arms move and grasp. With virtual agency and re-embodiment their pain reduced, in three from 7 – 8 to 2 - 4 on a visual analogue scale and in one from 4 to 0). No effect was seen in two; one had poor motor control of the stump following root avulsion, whilst both had had paralysis of their phantoms (and arms) for years before the trial.

Since leg amputation is more common we are now developing a lower limb task for trial.

This work was supported by The Wellcome Trust.

Moving towards collaboration: Using computational cognitive models to enable better human-robot interaction

Alan C. Schultz

Director, Navy Center for Applied Research in Artificial Intelligence Chair, Human Subject Institutional Review Board Naval Research Laboratory Code 5510 Washington DC 20375
(202) 767-2684 fax: (202) 767-3172

As we move along the scale from teleoperation towards collaboration, human-robot interactions become more complex and require that the human and the robot share more common knowledge about the world and how things within the environment are related.

At the collaborative level of interaction, the robot and human must exercise mixed initiative in solving a problem, each taking advantage of their unique skills, location, and perspective of the current situation. We believe that at this level and beyond, the robot will need representations and procedures that are similar to those used by humans, in order to collaborate successfully.

Our working hypothesis is that a system that uses representations and processes similar to a person's will be able to collaborate with a person better than a computational system that does not. I suggest three reasons for the representational hypothesis and then describe empirical and computational evidence in several domains.

Modelling the classic Attentional Blink and its emotional variant

Nikos Fragopanagos

A brain-based neural model of attention is used to simulate results for the 'attentional blink', observed when a subject is exposed to a rapid stream of stimuli and required to monitor for two successive targets in the stream. The 'blink' occurs when the time between the first and second targets is 200-500 ms, when there is reduced accuracy for report of the second target. The model gives a qualitative explanation of the phenomenon, especially of how attention is bolstered, during the processing to report of a given stimulus, in order to defend reportable information from attack by distracters. Finally, an extension of the attention model through the addition of an emotional component is presented in order to account for recent results of an emotional/attentional blink experiment showing that when the second target is emotional charged, subjects can detect it much more easily.

Thursday: Plenary three:

Neuroprosthetics and useful signals from motor cortex

Andrew B. Schwartz
University of Pittsburgh

We have shown that detailed predictive information of the arm's trajectory can be extracted from populations of single unit recordings from motor cortex. Using drawing movements as a behavioral paradigm, these signals have been shown to contain instantaneous velocity information and many of the invariants describing animate movement. Furthermore, this technique can be used to study visuo-perceptual processes taking place as objects are drawn. By developing techniques to record these populations and process the signal in real-time, we have been successful in demonstrating the efficacy of these recordings as a control signal for intended movements in 3D space. Having shown that closed-loop control of a cortical prosthesis can produce very good brain-controlled movements in virtual reality, we have been extending this work to robot control. We are using an anthropomorphic robot arm with our closed-loop system to show how monkeys can control the robot's movement with direct brain-control in a self-feeding task. The animals control the arm continuously in 3D space to reach out to the food and retrieve it to their mouths. Since the recorded signals are a high fidelity representation of the intended behavior and contain features of animate movement, neural prosthetic devices derived from this technology will be capable of producing agile, natural movement.

Plenary four:

Learning and Remapping Sensory Motor Control

Ferdinando A. Mussa-Ivaldi
Northwestern University
Rehabilitation Institute of Chicago

Studies of adaptation to patterns of deterministic forces have revealed the ability of the motor control system to form and use predictive representations of the environment. These studies have also pointed out that adaptation to novel dynamics is aimed at preserving the trajectories of a controlled endpoint, either the hand of a subject or a transported object. I will review some of these experiments and I will present new studies aimed at understanding how the motor system forms representations of space and time. An extensive line of investigations in visual information processing has dealt with the issue of how the Euclidean properties of space are recovered from visual signals that do not appear to possess these properties. The same question is addressed here in the context of motor behavior and motor learning and motor learning by observing how people remap hand gestures and body motions that control the state of an external device. I will discuss some theoretical considerations and experimental evidence about the ability of the nervous system to create novel patterns of coordination that are consistent with the representation of extrapersonal space. I will also discuss the relevance of this to rehabilitation and body-machine interfaces.

Papers:

Adaptive Working Memory: From Computational Neuroscience Model To Robot Control Module

David C. Noelle

Assistant Professor of Computer Science Assistant Professor of Psychology Vanderbilt University

Cognitive robotics has often been inspired by research in human and animal cognition. Recently, progress uncovering the computational basis of working memory – a memory system that focuses attention on task-essential information to aid the learning and execution of tasks – has advanced substantially in the field of cognitive science. Since humans and animals make extensive use of working memory for cognitive processing, it is reasonable to assume that robots could benefit from having working memory capabilities, as well. Such potential benefits include: focusing attention on the most task-relevant features, support of learning that transfers across various tasks, limiting the perceptual search space, providing a means to avoid the out-of-sight/out-of-mind problem, and more robust behavior in the face of distracting or irrelevant events.

In the work reported here, we have focused on one particular account of working memory function based on computational neuroscience models of the human prefrontal cortex. The prefrontal cortex has been widely implicated in human working memory function, and our computational account explains how interactions between the prefrontal cortex and the mid-brain dopamine system allow working memory to be adaptive — efficiently learning, from experience, what needs to be remembered and what can be safely forgotten. In this presentation, we review several important properties of human working memory and our adopted modeling approach. We then introduce a recently developed open source software library called the Working Memory Toolkit, which provides an abstraction of our computational neuroscience models. Finally, through a robot simulation study, we show that the toolkit offers sufficient working memory function to learn a delayed saccade task – a task commonly used to assess the working memory abilities of humans and non-human primates.

Internal models, adaptation, and uncertainty

Reza Shadmehr

Professor of Biomedical Engineering and Neuroscience

Johns Hopkins University

When the brain generates a motor command, it appears to also predict the sensory consequences of that command via an "internal model". The reliance on a model appears to make the brain able to sense the world better than is possible from the sensors alone. However, this happens only when the models are accurate. To keep the models accurate, the brain must constantly learn from prediction errors. Here I suggest that rules that govern learning are a reflection of the natural changes that can occur to the motor system:

- muscles can fatigue, and our body can be injured. According to a Bayesian formulation of this idea, movement error results in a credit assignment
- problem: what timescale is responsible for this disturbance? Our model predicts that the training schedule influences the behavior of the learner, changing estimates at different timescales as well as their uncertainty. A theoretical system that adapts in this way appears to account for many properties that have been observed in how the nervous system learns motor control.

Consciousness, Cognition, and Internal Models

Owen Holland

Department of Computer Science

University of Essex

In this talk I will describe some of the work done in connection with our attempt to develop a robot with a form of machine consciousness. The hypothesis behind the project is very similar to that proposed by the philosopher Thomas Metzinger: consciousness involves an internal model of the agent (or robot) that interacts with an internal model of the world, and the internal model of the agent is 'transparent' – that is, it has no access to information that it is only a model, and behaves in all respects as if it is the real agent. To investigate this hypothesis we have built CRONOS, an anthropomorphic robot – one that is qualitatively similar to the human body, with a multi-degree of freedom articulated skeleton, and with biomimetically-arranged elastic actuators. The properties of the robot are interesting in their own right, and will be described as such, but the key advance is that we have also developed an accurate physics-based real-time simulation of the robot (SIMNOS) and of the environment, and the simulated interaction of these two entities is expected to have predictive powers for the real robot's interactions with the world. I will describe our progress towards the next stage of work, which involves two aspects: the population of the internal world with appropriate representations of the real world using information from visual and other sensors (the perceptual process); and the development of a motor control system that will yield the same results on both CRONOS and SIMNOS.

Plenary five:

Socially Intelligent Robots – The Human Perspective

Prof. Kerstin Dautenhahn, University of Hertfordshire, UK

Social intelligence in robots has a quite recent history in Artificial Intelligence and robotics. However, it has become increasingly apparent that social and interactive skills are necessary requirements in many application areas and contexts where robots need to interact and collaborate with other robots or humans. In my talk I will focus on research in human-robot interaction (HRI) which poses many challenges regarding the nature of interactivity and 'social behaviour' in robot and humans. HRI has become a growing research field at the intersection of psychology, ethology, robotics and engineering. My talk will discuss examples of 'social intelligence' in robots, including two projects that I have been involved in: Firstly, robots as 'playmates' for children, research carried out as part of the Aurora project which investigates the potential use of robot as educational or therapeutic 'toys' for children with autism (www.aurora-project.com). Here I will highlight the concept of interactive emergence whereby turn-taking games emerge in play between children and a simple robot that only possesses very basic behavioural 'rules' but appears 'social' when situated in a play context. I will present examples where children demonstrate interactive competencies in interactions with the robot and/or the experimenter who is part of the experimental scenario. Secondly, I will survey research conducted as part of the European project Cogniron (www.cogniron.org) which investigates the development of a cognitive robot companion. I will provide examples of results from a series of HRI experiments with human-sized robots, including experiments in the University of Hertfordshire 'Robot House' dedicated to HRI experiments in a domestic living room environment (<http://www.exn.ca/dailyplanet/view.asp?date=4/11/2006>). The aim of this work is to develop social rules for robot behaviour (a 'robotiquette') that is comfortable and acceptable to humans (<http://news.bbc.co.uk/1/hi/technology/3962699.stm>).

Plenary six:

Information and Embodiment

Olaf Sporns
Department of Psychological and Brain Sciences
Indiana University Bloomington, USA

Biological organisms continuously select and sample information used by their neural structures for perception and action, and for creating coherent cognitive states guiding their autonomous behavior. However, information processing is not solely an internal function of the nervous system. Sensorimotor interaction and body morphology can induce statistical regularities and create additional information in sensory inputs and within the neural control architecture. This effect of embodied interaction on information in the nervous system can be explored by quantitative analysis of sensory, neural and motor data sets collected from different robotic platforms. The analysis reveals the presence of information structure and (directed or “causal”) information flow induced by dynamically coupled sensorimotor activity. Information structure and information flow in sensorimotor networks has spatial and temporal specificity, depends on aspects of body morphology and changes with learning. Our approach establishes a quantitative link between information and embodiment.

References:

- Sporns, O., Karnowski, J., and Lungarella, M. (2006) Mapping causal relations in sensorimotor networks. Proc. ICDL 2006.
- Sporns, O. and Lungarella, M. 2006. Evolving coordinated behavior by maximizing information structure. In Rocha, L.M., Yager, L.S., Bedeau, M.A., Floreano, D., Goldstone, R.L., and Vespigniani, A. (eds.) Artificial Life X. Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems, pp. 323-329, Cambridge, MA: MIT Press.
- Lungarella, M., Pegors, T., Bulwinkle, D., and Sporns, O. (2005) Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics* 3, 243-262.

Papers:

Designing a complex mind - is it possible?

Troy Kelley

Army Research Laboratory

Human Research and Engineering Directorate, AMSRD-ARL-HR-SE, APG, MD, 21005-5425

voice: 410-278-5869, fax: 410-278-9523, tkelley@arl.army.mil

How does one design a mind? This is obviously a difficult question. This question has also turned out to be much more difficult than was first thought during the early days of Artificial Intelligence (AI). If we understand how the mind was created, by evolution and through natural selection, does that necessarily mean that we can create a mind in computerized format? Or, could it be, that in order to create a complex mind, we would have to re-create evolution, and such an undertaking would be impossible. Or, could we speed up the evolutionary process for mind creation? If we did speed up the evolutionary process, how could we be sure that the constraints we have imposed for our own evolutionary process, are the correct constraints to create a mind? What are the possible approaches, constraints and considerations for the creation of a truly complex mind? These questions will be addressed in the following discussion.

Principles Underlying the Construction of Brain-Based Devices

Jeff Krichmar

Without a doubt organisms whose behavior is guided by a nervous system demonstrate the most sophisticated behavior seen in biological agents. 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 15 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:(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 physical instantiation, which allows for active sensing and autonomous movement in the environment, (iv) it should engage in a task that is initially constrained by minimal set of innate behaviors or reflexes, (v) it should have a means to adapt the device's behavior, called value systems, when an important environmental event occurs, 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.

Friday: Plenary seven:

Morphological computation – connecting brain, body, and environment

Rolf Pfeifer, Artificial Intelligence Laboratory, University of Zurich, Switzerland

Traditionally, in robotics, artificial intelligence, and neuroscience, there has been a focus on the study of the control or the neural system itself. Recently there has been an increasing interest into the notion of embodiment in all disciplines dealing with intelligent behavior, including psychology, philosophy, and linguistics. In this talk, we explore the far-reaching and often surprising implications of this concept. While embodiment has often been used in its trivial meaning, i.e. „intelligence requires a body“, there are deeper and more important consequences, concerned with connecting brain, body, and environment, or more generally with the relation between physical and information (neural, control) processes. Often, morphology and materials can take over some of the functions normally attributed to control, a phenomenon called “morphological computation”. It can be shown that through the embodied interaction with the environment, in particular through sensory-motor coordination, information structure is induced in the sensory data, thus facilitating perception and learning. A number of case studies are presented to illustrate the concepts introduced. I conclude with some speculations about potential lessons for robotics.

Papers:

Social cognition and social robots

Shaun Gallagher

The theoretical underpinnings of social cognition have been shifting away from overly-intellectualistic conceptions based on theory of mind approaches, toward more embodied versions of simulation and interaction theories. In theory of mind approaches social cognition is framed in terms of gaining access to the other person's mind; in simulation theory (especially implicit versions) the activation of neural resonance systems (mirror neurons, shared neural representations) puts conspecifics into the same or similar sensory-motor states, and this is the basis for social understanding that is informed by action schemas or emotion based empathy (e.g., Decety 2004, 2005; Gallese 2003).

Interaction theory appeals to the same neuroscience of resonance systems, and builds on research in developmental psychology, to show that the basis of social cognition is both perceptual and contextual (Gallagher 2001, 2004, 2005).

Conspecifics who are communicating or interacting in pragmatic or social contexts depend to a high degree on the perception of the other's movements, postures, gestures, facial expressions (e.g., Trevarthan 1979 -- ‘primary intersubjectivity’). Conspecifics also make use of pragmatic contexts (environmental features, specific objects of shared attention) to understand the actions of others (Trevarthan and Hubley 1978 -- ‘secondary intersubjectivity’). These capacities for understanding embodied, non-verbalized meanings are important mechanisms for understanding others even prior to

(and preparatory to) language acquisition, and they are integrated with linguistic communication and narrative competency in human development. They deliver very basic elements that often are sufficient (without verbal communication) for delivering meaning.


Here I want to explore some questions about just one aspect of human-robot interaction. What are the specific elements that robots and intelligent agents need to manifest in second-person interactions with humans to facilitate human understanding of the robot? How dependent on, or independent of human embodiment are these elements, and is it nonetheless possible to create homologous substitutes in robotic bodies for those elements that are highly dependent?

From Intelligent Control to Cognitive Control: A mechanism for robust robotic sensorimotor intelligence

Kazuhiko Kawamura
 Center for Intelligent Systems, Vanderbilt University
 Nashville, TN, 37235-0131, USA


Engineers have long used control systems utilizing models and feedback loops to control real-world systems. Limitations of model-based control led to a generation of intelligent control techniques such as adaptive and fuzzy control. Human brain, on the other hand, is known to process a variety of inputs in parallel, ignore noises and distractions to focus on the task in hand. This process, known as executive or cognitive control on cognitive psychology, is unique to humans and some animals. We are interested in implementing cognitive control in robots. This presentation tries to answer the following two questions: (1) How does one define cognitive control to fit into the field of robotics or intelligent machines? (2) How is cognitive control implemented in robots?

In 2004, we proposed a multi-agent based cognitive architecture [1] for our humanoid robot ISAC. In this presentation, I will give brief overview of our current implementation status.



Cognitive Control Function

ISAC's cognitive control function is modeled partially based on Baddeley and Hitch's psychological human working memory model [2] and on O'Reilly's computational model of working memory [3]. In our cognitive architecture, an IMA agent called the Central Executive Agent (CEA) is responsible for providing cognitive control function to the rest of the system as shown below. It interfaces to the Working Memory System (WMS) to maintain task-related information (or "chunks") during task execution.



Our WMS is based on a computational neuroscience model and its toolkit developed by Phillips and Noelle [4]. Under the current design, CEA will have the four key functions: 1) situation-based action selection, 2) episode-based action selection, 3) control of task execution, and 4) learning sensory-motor actions. Each function will be realized through interaction between CEA, other IMA agents, and various memory systems.

REFERENCES

- [1] Kawamura, K., R.A. Peters II, R. Bodenheimer, N. Sarkar, J. Park, A. Spratley, and K. A. Hambuchen, "A Parallel Distributed Cognitive Control System for a Humanoid Robot, International Journal of Humanoid Robotics, vol. 1, no. 1, pp. 65-93, 2004
- [2] Baddeley, A.D. and G.J. Hitch, "Working Memory", In G.A. Bower (Ed.), Recent Advances in Learning and Motivation, (Vol. 8, pp. 47-90), New York: Academic Press, 1974
- [3] O'Reilly, R.C., T.S. Braver and J.D. Cohen, "A biologically based computational model of working memory," in Models of Working Memory: Mechanisms of active maintenance and executive control, A. Miyake and P. Shah, eds., Cambridge University Press, 1999, pp. 375-411.
- [4] Phillips, J. L. and D.C. Noelle, "Working Memory for Robots", In Proceedings of the 5th International Conference on Development and Learning, Bloomington, Indiana, 2006

Supplementary Abstracts (including Poster Abstracts):

Why Robots Don't Feel Pain

Mark Bishop, Goldsmiths College, UK.

In Science and Science Fiction the hope is periodically reignited that a robot may one day achieve consciousness in virtue of its execution of an appropriate program; indeed in the UK the EPSRC recently funded a large project with a goal of instantiating 'machine consciousness' through appropriate computational 'internal modelling'. In contrast, in this poster I outline a simple reductio style argument based on [1] that either suggests such optimism is misplaced or that panpsychism, (the belief that 'the physical universe is composed of elements each of which is conscious'), must be true.

1. Bishop J. M., 2002, 'Counterfactuals cannot count: a rejoinder to David Chalmers', Consciousness and Cognition, v11, pp 642-652.

Spatial Reference Language for Human Robot Interaction

Samuel N. Blisard, Robert H. Luke III, Marjorie Skubic, and James Keller
Computational Intelligence Research Lab, University of Missouri-Columbia

One of the key components for natural interaction between humans and robots is the ability to understand the spatial relationships that exist in the natural world. We, as humans, do not tend to use precise language in our day-to-day interactions with the environment (e.g. “Turn 85 degrees and walk forward 2.2 meters” as opposed to “Stand in front of the camera.”); instead, we use a qualitative model of the space around us and spatial referencing language to communicate to others about our environment. Previous research has shown that modeling the 2D spatial relationships of FRONT, BEHIND, LEFT, RIGHT, and BETWEEN can be accomplished with results consistent with that of a human being. Upcoming research will involve a human subject study to investigate the use of spatial relationships in 3D space. This will be the first step in extending previous research of the 2D spatial relations into a 3D representation through the use of 3D object point clouds generated by the SIFT algorithm and stereo vision.

REFERENCES

- [1] M. Skubic, D. Perzanowski, S. Blisard, A. Schultz, W. Adams, M. Bugajska, and D. Brock, Spatial language for human-robot dialogs, *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, vol. 34, no. 2, Special Issue on Human-Robot Interaction, pp. 154-167, 2004.
- [2] M. Skubic, P. Matsakis, G. Chronis, and J. Keller, Generating multilevel linguistic spatial descriptions from range sensor readings using the histogram of forces, *Autonomous Robots*, vol. 14, no. 1, 2003.
- [3] S. N. Blisard, Modeling spatial references for unoccupied spaces for human-robot interaction, Master's thesis, University of Missouri-Columbia, 2004.
- [4] D. G. Lowe, Distinctive Image Features from Scale-Invariant Keypoints, *International Journal of Computer Vision*, vol. 60 no. 2 pp 91-110, 2004
- [5] R.H. Luke, S.N. Blisard, J.M. Keller, M. Skubic, Modeling Spatial Referencing Language for Human Robot Interaction, in *Proc. of the IEEE Intl. Workshop on Robot and Human Interactive Communication*, Nashville, TN, Aug. 2005, pp. 704 - 709

Working Memory Experiment: A Robot in a Water Maze

Marjorie Skubic, Mark Busch, Kevin Stone, and James Keller
Computational Intelligence Research Lab, University of Missouri-Columbia

Using the Working Memory toolkit developed by David Noelle and Josh Phillips at Vanderbilt University, we have designed an experiment that is inspired by Jeff Krichmar's work at the Neurosciences Institute, which is modeled on experiments with a rat in a water maze. A robotic “rat” is placed in a rectangular room where its perceptual information is in the form of colored panels on the walls. A circular platform is placed somewhere in the room; the robot can sense the platform only when directly on top of it. The robot must learn to locate the platform using only the perceptual cues from the colored panels. We have implemented a learning system using the Working Memory toolkit. First, the perceptual space is discretized using a self organizing feature map (SOFM). The node of the SOFM is used as a state vector, and the robot must learn which action to take, given the current state. Five possible actions are used – move forward, turn right, turn left, turn hard right, turn hard left. Preliminary results indicate that the robot can learn how to navigate to the platform from a consistent starting location in about 60 iterations.

REFERENCES

- [1] Phillips, J. L. and Noelle, D. C. “A Biologically Inspired Working Memory Framework for Robots”, Proceedings of the 27th Annual Conference of the Cognitive Science Society, Stresa, Italy, pp. 1750-1755, 2005.

[2] J.L. Krichmar, D. A. Nitz, J.A. Gally, and G. M. Edelman, “Characterizing functional hippocampal pathways in a brain-based device as it solves a spatial memory task,” *Proc Natl Acad Sci USA*, vol. 102, pp. 2111-2116, 2005.

Shared Control for Motor Skill Training

Volkan Patoglu, Yanfang Li and Marcia K. O’Malley

Department of Mechanical Engineering, Rice University, Houston, TX – 77005

The neural mechanisms for acquisition of sophisticated motor behavior in humans are complex. There exists evidence in the literature that human nervous system may be utilizing internal dynamic models of the external systems with which they interact [1]. We are investigating existence and formation of internal models during sensorimotor learning through human subject experiments. To this end, we conduct catch trial based experiments in which subjects interact with an underactuated dynamic task. We use optimal control paradigms to explain our experimental results and to model the internal model formation in humans. Our preliminary results (see Figure 1) indicate that humans tune their excitation through an optimization scheme that is subject to the dynamics of the external objects they manipulate.




Fig. 1. This figure represents the frequency power spectrum of excitation for a typical human subject interacting with an external dynamic system with two different set of parameters. Subjects are overtrained to perform an underactuated target hitting task with one set of dynamic system parameters. Excitation frequency spectrum plot of a typical trial for this overtrained task is shown in Subfigure 1 (a). Subfigure 1 (b) depicts the adaptation of human excitation frequency when the system parameters are changed unexpectedly.

We are interested in accelerating the adaptation of the human internal model to achieve desired performance for a given task. We aim to design novel interaction paradigms for virtual environments in order to facilitate motor learning of these dynamic tasks. Shared control has been proposed by the authors as a means to achieve this goal. Shared control is an active assistance based training paradigm where dynamic intervention through an automatic feedback controller, acting upon the system, is employed to modify the system dynamics adaptively during training. A schematic representation of shared control is given in Figure 2. Prior studies using shared control with error reduction, implemented as assistance in order to improve task performance, indicate that the rate of learning of dynamics tasks can be expedited by such paradigms [2]. Further work by the authors revealed that the benefit from an active shared controller with error reduction is not through demonstration of the preferred strategy [3]. These results suggest that faster learning may be promoted through a reduction in number of control parameters in the subjects’ adaptation mechanism.

Aforementioned results are in good agreement with the Bernstein hypothesis [4], which proposes that simplification of system dynamics through reduction in degrees of freedom can facilitate motor learning of dynamic tasks. The hypothesis is analogous to well-developed multi-phase optimization

techniques, where a coarse global search phase is followed by a fine local search that is initialized with the parameters suggested by the coarse global phase. Utilizing the knowledge of simplified dynamics as a useful foundation for learning the more complex task, developmental progression has been shown to be the most effective training mechanism for neural networks [5].




Figure 2 (b) is analogous of the bicycle riding task but is implemented in a virtual setting, where dynamic intervention through an automatic feedback controller is utilized as an active assistance based training paradigm.

In addition to number of control parameters, human perceptual issues play an important role in the design of training methodologies. A thorough understanding of human dynamic response to changes in external system parameters is required for effective design of adaptive training schemes. Towards this end, we are carrying out human subject studies to understand and quantify human interaction with external systems in terms of perception related parameters, such as just noticeable differences. The results of these experiments will be utilized to optimize shared control based training. This research will contribute too many areas with high societal impact, e.g. rehabilitation of stroke patients and training of surgeons.

REFERENCES

- [1] J. B. Dingwell, C. D. Mah, and F. A. Mussa-Ivaldi, "Manipulating objects with internal degrees of freedom: Evidence for model-based control," *The Journal of Neurophysiology*, vol. 88, no. 1, pp. 222–235, 2002.
- [2] M. K. O'Malley, A. Gupta, M. Gen, and Y. Li, "Shared control in haptic systems for performance enhancement and training," *ASME Journal of Dynamic Systems, Measurement and Control (JDSMC) Special Issue on Novel Robotics and Control*, vol. 128, no. 1, pp. 75–85, 2006.
- [3] Y. Li, V. Patoglu, and M. K. O'Malley, "Shared control for training in virtual environments: Learning through demonstration?" in *EuroHaptics'06*, 2006.
- [4] N. Bernstein, *The Coordination and Regulation of Movements*. Oxford: Pergamon, 1967.
- [5] V. Ivanchenko and R. A. Jacobs, "A developmental approach aids motor learning," *Neural Computation*, vol. 15, pp. 2051–2065, 2003.

The Role of an Adaptive Working Memory System in a Cognitive Robot Architecture for Task Execution

Stephen Gordon and Kazuhiko Kawamura
Center for Intelligent Systems, Vanderbilt University

Humans use a short-term memory system called working memory to manage memory attention and retrieval during task execution [1] [2]. Under the sponsorship of DARPA and NSF, we in the CIS have

developed a cognitive robotic architecture with distinctive memory systems [3]. Most recently, a software implementation of a biologically inspired adaptive working memory system has been developed [4]. The work described in this presentation demonstrates our attempts to implement such a system onto our cognitive robot architecture. Through training, our humanoid robot, ISAC, will adapt its own working memory system (WMS) and learn over time what information from its long-term memory (LTM) and short-term memory (STM) is most useful for particular tasks. This information will represent three-to-five “chunks” of information necessary for task execution or learning. Appropriate reward signals and close monitoring of both internal and external states will guide this learning process. Over time, the system is expected to acquire capability through the training and implementation of a wide variety of WMSs applicable to certain situations that ISAC may encounter. The goal of this work is to show that ISAC will be capable of recognizing particular (non-routine) situations in order to retrieve the most relevant chunks of trained WMSs from LTM and therefore be able to utilize WMS to guide ISAC’s response to the current situation.

- [1] Baddeley, A. D. and Hitch, G. J., “Working Memory”, In G. A. Bower (Ed.) *Recent Advances in Learning and Motivation*, New York: Academic Press, pp 47 – 90, 1974
- [2] Braver, T. S. and Cohen, J. D., “On the Control of Control: The role of dopamine in regulating prefrontal function and working memory”, In Monsell, S. and Driver, J., eds., *Control of Cognitive Processes*, 18, Attention and Performance, Chapter 31, MIT Press, pp. 713 – 737, 2000
- [3] Kawamura, K., Peters, R. A., Bodenheimer, R. E., Sarkar, N., Park, J., Clifton, C. A., Spratley, A. W., and Hambuchen, K., “A Parallel Distributed Cognitive Control System for a Humanoid Robot”, *International Journal of Humanoid Robotics*, 1, pp. 65-93, 2004.
- [4] Phillips, J. L. and Noelle, D. C. “A Biologically Inspired Working Memory Framework for Robots”, *Proceedings of the 27th Annual Conference of the Cognitive Science Society*, Stresa, Italy, pp. 1750-1755, 2005.

Bonus abstracts:

When a Robot Thinks 'I'

Igor Aleksander – Imperial College, London

Does a robot need an explicit representation of 'self' and, if so, what might this be? I refer to some models we have built in which there are several differing interpretations of self: some refer to the focus of a sensory world ‘in the head’ and others to whatever phenomenal mechanisms may be active in the system.

But stepping back to get a perspective on Cognitive Robotics in general, one finds several philosophies. At the one extreme, classical AI techniques make use of concepts familiar in cognitive psychology where a representation of self is largely couched in terms of various memory modes (episodic, working, etc.). In 'situated' robots (e.g. Brooks) and 'sensory-motor contingency' paradigms (e.g. O'Regan) the world is a place-holder for the phenomenal needs of a robot to drive its engagement with the environment. Here the self comes largely from a representation of action selection in the 'I (will) do' mode. At the other extreme (e.g. Kritchmer) there are attempts at very close modelling of neurophysiology, and here the same question as in living systems arises: what are the neural correlates of a sense of self? The answer to this is anything but clear.

Highly interesting are models such as Baars’ Global Workspace where a distinction is made between those computational elements in which any consciousness can take place (the GW) and those where it does not (the supporting memory processes). However there are difficulties here in doing more than just saying that a conscious self is represented in GW. How is it encoded and how does this encoding differ between conscious and unconscious representation? These remain open questions.

As stated, in our own work, we distinguish between different uses of the word ‘I’ as in: (a) ‘I am in a field with cows and daisies in it’; (b) ‘I remember Paris when it drizzles’; (c) ‘Today I want penne

arrabiata rather than pizza napoletana; (d) I am scared of the dark. Briefly, these modes point to different computational mechanisms.

Now for the mechanisms in our model. (a) is the result of depictive representations in fine-grain neural fields which create a virtual point of view that feels to be somewhere in the middle of one's head. (b) is the result of the existence of attractors in a dynamic neural field that is iconically coupled to the field in (a). Here the 'I' comes from the emergence of a meaningful state in this 'imagination' field. (c) relates to action selection mechanisms which depend both on the following of planning trajectories in the imagination neural mechanism and receiving an emotional input from unconscious 'emotion' nets in our kernel architecture (this is work done by Rabinder Lee in my group at Imperial College – a demonstration using webots will be available). Finally, a raw emotion such as (d), we argue occurs by association with innate muscular reactions largely as a mode of operation of the imaginative network.

The upshot of all this is that when talking of a 'self' as a phenomenon in a robot, it should be considered that this does not point to a cohesive mechanism but the activity of one or more of a variety of mechanisms. Then, when asking the question, 'does a robot need a 'self'?' the answer may be couched in terms of the contributing mechanisms. I submit that the answer is indeed positive. Given mechanisms of depictive perception, depictive imagination, planning emotional evaluation, would indeed produce a robot capable of improving its behaviour in a complex world and, given language, express its feelings about itself and its actions using the word 'I' in an appropriate way. I would not recommend trying to build 'self' boxes. Indeed when a person goes to the doctor saying "I have an acute pain here in my stomach", the doctor looks up 'appendicitis' in his manual and not 'the self' despite the appearance of the word 'I'.

[* There is no space here for full definitions of some of our idiosyncratically used terms, but they are defined in our recent papers or, indeed in I. Aleksander, 'The World in my Mind ...' Imprint Academic, 2005]

Discussion Forum

This workshop is about having ideas, some of which don't take off and others will fly. The discussion forum associated with our website was one idea that didn't take off. We've had constructive comment that researchers are simply too busy to search such forums. However, it would be a shame to miss the post below from Aaron Sloman:

...I thought I'd invite people to look at and criticise three related online documents which are informal discussion papers arising out of the (EC-funded) CoSy robotic project.

1. A critique of common ideas about sensorimotor contingencies whose importance is generally inflated beyond their worth in relation to animals or robots able to perceive learn about, think about, and manipulate complex 3-D objects.

See this HTML file:

[url]<http://www.cs.bham.ac.uk/research/projects/cosy/papers/index.php#dp0603>[/url]

2. A set of conjectures about kinds of learning required to produce the sorts of competence, learning and creativity found in human children and some other altricial species, based on acquiring a large number of (nearly) orthogonal recombinable competences (discussed recently on the psyche-D list):

[url]<http://www.cs.bham.ac.uk/research/projects/cosy/papers/index.php#dp0601>[/url]

3. A related document attempts to make explicit the requirements for 'fully deliberative' architectures and mechanisms, which future robots will need if they are to be used as congenial helpers, e.g. in domestic contexts:

[url]<http://www.cs.bham.ac.uk/research/projects/cosy/papers/index.php#dp0604>[/url]

As far as I know most people who use the word 'deliberative' think it merely refers to the ability to make a selection from a set of options, which is about the simplest possible case (I presume even insects do that). Human deliberative competences are far richer and more complex.

...., please email me if you wish to discuss these things: A.Sloman@cs.bham.ac.uk
<mailto:A.Sloman@cs.bham.ac.uk>

Or use this new jiscmail list set up by John Jackson
[url]<http://www.jiscmail.ac.uk/lists/mindmechanisms.html>
[/url]

Or use the psyche-D list:
[url]<http://listserv.uh.edu/archives/psyche-d.html>[/url]

Aaron Sloman

oooOOooo

Thank you to all the participants, both those who kindly took time to prepare a talk or poster and all those who will discuss the topics raised.