

Models of Consciousness

ESF/PESC Exploratory Workshop

Scientific Report

Ricardo Sanz
Universidad Politécnica de Madrid

Ron Chrisley
University of Sussex

Aaron Sloman
University of Birmingham



October 15, 2003

Purpose of the document

This document contains the final scientific report of the ESF/PESC exploratory workshop *Models of Consciousness*.

Workshop WebSite

The workshop web site can be located at:

<http://aslab.disam.etsii.upm.es/public/events/moc/>

Downloading this Document

This document can be downloaded from:

<http://aslab.disam.etsii.upm.es/public/events/moc/report/ScientificReport.pdf>

Acknowledgements

The organisers of this workshop would like to acknowledge the help from Phillipa Rowe (European Science Foundation) and Donna Willmets (University of Birmingham).

We would like to specially recognise all the efforts that Ignacio López (Universidad Politécnica de Madrid) has done to have all those necessary things running smoothly.

Table of Contents

Executive summary.....	5
Scope and Objectives.....	5
The Context for the Workshop.....	5
Focus on Interdisciplinarity.....	6
Meeting place	6
Web Site.....	6
Continuation.....	7
Scientific Content of the event.....	8
François Anceau.....	8
Axel Cleermans.....	8
Jim Doran.....	9
William Edmonson.....	10
Petros Gelepithis.....	10
Pentti Haikonen.....	12
Germund Hesslow.....	13
Owen Holland.....	14
Jacques Lacombe.....	14
Riccardo Manzotti.....	14
Peter Redgrave.....	15
Geraint Rees.....	15
Antti Revonsuo.....	16
Miguel A. Salichs.....	17
Ricardo Sanz.....	18
Murray Shanahan.....	18
Aaron Sloman.....	19
John Taylor.....	21
Tom Ziemke.....	23
Assessment of the result.....	25
Expected outcome of the workshop	25
Real outcome of the workshop.....	25
Focus on implementation.....	26
Final Program.....	28
Day 1: Data and General Considerations.....	28
Day 2: Potential Unified Theories.....	29
Day 3: Research and Engineering.....	30
Final list of participants.....	31
Speakers.....	31
Non Speakers.....	35
Statistical information.....	37
Country.....	37
Sex.....	37
Position.....	37

This page has been intentionally left blank

Executive summary

Scope and Objectives

The **Models of Consciousness** workshop tried to advance the elaboration of a unified scientific theory of consciousness. A core topic of the workshop was *unified theories of natural and artificial consciousness* and this workshop focused on the particular aspects of *models of consciousness* that are also suitable for implementation, *i.e.* theories of consciousness that **can support the construction of conscious machines** and also serve as explanation of the experimental data about the natural kind of consciousness.

The objective of the workshop was the presentation, evaluation and discussion of available models of consciousness both from a *scientific point of view* (providing explanation of, *e.g.*, observed behavior) and *technological point of view* (serving as base design for building conscious machines). This would enable the evaluation of potential candidate architectures for various forms of consciousness that will serve as a focal points for future collaborative research projects.

The workshop gathered a multidisciplinary collection of European researchers into the fields of artificial intelligence, neuroscience, robotics, psychology, physics, automatic control, computer science, philosophy, *etc.*

This workshop will was very valuable to initiate the exchange of knowledge and experiences between researchers from across Europe in an emerging area of research; due to its necessary multidisciplinary it helped establish new collaborative links between different disciplines; and as a major activity it will serve to test innovative ideas and develop potential collaborations.

The Context for the Workshop

The objective of building *conscious machines* was already a research topic in the early years of artificial intelligence, but the extreme difficulties encountered at that time in developing implementable models of even the simplest features of human intelligence halted the research and put *machine consciousness* into the bin of *Utopian research topics* (more or less like *time-travel, immortality or hair-restoring*).

But the case for *consciousness* is a little bit different because consciousness *does exist now*. Consequently, we know *a priori* that the construction of a conscious entity is possible. Research in artificial consciousness is not any longer Utopian research for several reasons:

- Recent advances in *neuroscience* and *physiological psychology* have provided a deeper knowledge of how human and animal consciousness arises.
- Modern *robotic mental architectures* have refocused from abstract problem solving systems onto systems that continuously interact with the world, which some take to be an important step toward the explanation and creation of (a particular kind of) awareness (*i.e.* flea-type awareness).
- Developments in the field of *complex control systems* delved into theories of how

model-based control-loop integration can lead to an integrated perception of the self of the controller and the plant.

- *Computer technology* now provides computational power that is many orders of magnitude beyond what was available in the past. It has been estimated that in ten to twenty years, computers will reach the computational power of the human brain (from a certain perspective).
- A recent explosion of interest in consciousness on the part of philosophers has led to the most sophisticated conceptual understanding to date of the possibilities for a scientific theory of consciousness, and the potential obstacles to such.

We understand that the main problem for having a good scientific theory of consciousness lies in the wide scattering of knowledge among a wide collection of disciplines. It is necessary to employ coherent, interdisciplinary approaches to the problem to get a glimpse of a good solution. This is why this workshop is essential.

There have been previous attempts to advance the development of a theory of consciousness but, to our knowledge, this is the first attempt to filter theories on the basis of technological applicability (engineering filtering).

Focus on Interdisciplinarity

This work is interdisciplinary by nature and by need because there is no single discipline that can provide all the relevant knowledge, nor the necessary broadness, nor the essential tools to build such a machine.

- The people who have the data about natural consciousness can be found in the disciplines of **psychology**, **medicine**, **neurobiology**, **ethology**, **linguistics**, **anthropology**, and **philosophy**.
- The available scientific models of natural consciousness can be found in **neuroscience**, **cognitive science** and **biophysics**.
- Some **artificial intelligence** and **robotics** researchers now regard artificial consciousness as one of the basic research objectives of the field.
- If we are going to “engineer” conscious machines it is obvious that it is necessary to have engineering disciplines represented in the workshop. We have selected people from **computer engineering**, **software engineering** and **automatic control engineering**.

Meeting place

The meeting and accommodation place was be **Manor House** at the University of Birmingham, UK.

Web Site

The workshop web site can be located at:

<http://aslab.disam.etsii.upm.es/public/events/moc/>

Continuation

The MoC (Models of Consciousness) workshop was executed as planned with 28 people attending. The workshop was very welcomed by all researchers and, after the conclusion, many of them expressed their congratulations to the organisers.

The MoC workshop funded by the ESF was followed by another one funded by the European Network of Excellence Exystence. This was titled “Machine Consciousness. Complexity aspects and took place in Torino, Italy at the end of September.

The organisers of the Birmingham MoC workshop were requested to summarize results from the MoC workshop to the Torino attendants (many people repeated attendance) and this input was considered very valuable for the new workshop.

Scientific Content of the event

This section contains the abstracts of the talks during the workshop. The slides used by the authors are available at the workshop web site. A text version of the workshop (in the form of a book) is underway at the time of this writing.

François Anceau

A Model for Consciousness Based on the Sequential Behavior of Voluntary Actions

We have developed a model for consciousness in which the capability of focusing the attention seems to be the same mechanism as the one used for triggering the conscious mental functions. This mechanism makes conscious thinking an extension of the voluntary attention as defined by A.H.Luria.

We will call "Attention Point" the focusing point of the voluntary attention process. The moving of this point makes conscious thinking serial, as we can see every day. When a conscious process is operating, this point triggers automatic functions by visiting successively their triggering areas in the mental space.

From an evolution point of view, we will suppose that the emergence of the mechanism of consciousness has risen from the basic attention mechanism through an increase of its triggering possibilities. This evolution process had could give to the attention point the capacity to move outside the sensory areas, opening the possibility of paying attention to abstract notions.

We propose that the consciousness mechanism could be a framework providing facilities to make possible the very existence of the high-level non-conscious cerebral functions such as intelligence, long-term memory, reasoning,... One of the most important of these facilities would be to provide these mental functions with a timing coherence due to a kind of consistency mechanism. Thereby such a timing consistency could be the consequence of the seriality of the conscious activation of the automatic mental functions.

This model differs from the GW model defined by B.J. Baars where the sequentiality of the conscious processes appears to be a by-product of the shortage of neural resources needed to run them. It differs also from the models where consciousness consist of giving meaning to perceptions.

F. Anceau, *Vers une étude objective de la conscience*, Hermès Science, Paris, 1999

Axel Cleermans

The search for the computational correlates of consciousness

Over the past few years numerous proposals have appeared that attempt to characterize consciousness in terms of what could be called its computational correlates: Principles of information processing with which to characterize the differences between conscious and unconscious processing. Proposed computational correlates include architectural

specialization (such as the involvement of specific regions of the brain in conscious processing), properties of representations (such as their stability in time or their strength), and properties of specific processes (such as resonance, synchrony, interactivity or competition).

In exactly the same way that one can engage in a search for the neural correlates of consciousness, one can thus search for the computational correlates of consciousness. The most direct way of doing so consists of contrasting models of conscious vs. unconscious information processing.

In this talk I will review these developments and illustrate how computational modeling of specific cognitive processes can be useful in exploring and in formulating putative computational principles through which to capture the differences between implicit, explicit, and automatic cognition.

What can be gained from such approaches to the problem of consciousness is an understanding of the function it plays in information processing. Here, I suggest that the central function of consciousness is to make it possible for cognitive agents to exert flexible, adaptive control over behavior. Learning processes therefore play a central role in shaping conscious experience.

From this perspective, consciousness is best characterized as involving a continuum defined over quality of representation: Graded representational systems that can be adaptively modified by ongoing experience are thus viewed as a central feature of any successful model of the differences between conscious and unconscious cognition.

Jim Doran

Behaviours, Day-Dreams, Plans, and Consciousness

In this talk I shall describe an approach to the construction of conscious machines that initially seeks answers to three key questions:

- What minimum set of core properties is sufficient to support the emergence of high-level mental abilities?
- How may the presence of emergent high-level mental abilities be detected?
- What properties must a machine have if it is reasonably to be labelled "conscious"?

As regards the first and second questions, I shall first explain why "design for emergence" seems an interesting road to follow, and then focus on the emergent creation and use of behaviours, "day-dreams" and plans by a machine designed and implemented at the neural network level. I shall give some details of the design that I am currently elaborating and have begun to implement as a C program, and consider how high-level mental abilities might be detected in this specific context.

As regards the final question, I shall first recall the well-known distinction between "subjective consciousness" and "behavioural consciousness" and then put the former to

one side, and discuss how the latter may be achieved in a machine design of the foregoing type.

William Edmonson

General Cognitive Principles and the Structure of Behaviour or Prolegomena for a big ToE.

In order to consider the possibility of modelling consciousness it is necessary to have a functional specification of the brain – not just the human brain, but any animal brain. I argue that the brain is the organ for dealing with the sequential imperative and I introduce six general cognitive principles which cover important aspects of brain functionality:

GCP1 Sequentiality in behaviour is forced physiologically.

Corollary 1 Sequence penetrates the corporeal boundary.

Corollary 2 Sequence is semiotically free.

GCP2 Cognitive entities are i) inherently atemporal, & ii) dual in nature.

GCP3 Behaviour is sequencing; perception is de-sequencing.

GCP4 Learning serves the sequential imperative.

GCP5 Attention is the management of the processes of sequencing and de-sequencing.

GCP6 Thought is the production of cognitive entities.

The sequential imperative is expressed in the first principle, and the remaining principles cover aspects of functionality which are required for cognition.

A working model – PANTOME – has been built which gives expression to the first three principles. This will be described. Conjectures will be offered on the possibility that:

a) specification of functionality will lead to better/implementable models of cognition;

b) exploration of properties of such models as PANTOME will lead to identification of incompleteness of functionality, of unrecognized functionality, and of functional opportunities (spandrels);

c) consciousness occupies a spandrel.

Petros Gelepithis

Consciousness within Unified Theories of Mind

In consciousness studies there is a multitude of viewpoints but not much work in comparing and, hopefully, synthesizing at least some of the various conceptions. In

addition, theories of consciousness had not been developed, so far, within proposed unified theories of mind (UToM). This position paper is a brief outline proposal to that end. It consists of three parts. In the first one, we:

- a) Summarily, remind the reader of Newell's (1990, 1992) view of unified theories, his applied-AI approach to cognition and his proposed list of fundamental notions for Cognitive Science.
- b) Provide a bird's eye view of the three approaches to the study of mind and some key representative conceptions, theories, and hypotheses of the three approaches (Gelepithis 2002a).
- c) Provide a table of the currently available classes of tools for theory construction in cognitive science and organisation and social theory (Gelepithis 2004b).
- d) Present a juxtaposition of Newell's fundamental notions for Cognitive Science and our system of such notions (Gelepithis 2001b, 2004a).

Drawing upon the above, we suggest that the currently used mathematics are inadequate and propose that language (couched in axiomatic terms) is a better tool for capturing some of the complexity of mind and the nature of consciousness. The second and third parts of this paper take up these suggestions.

Specifically, in part two, we:

- a) Present a system of eight definitions that jointly constitute necessary and sufficient conditions for an intelligent system and the précis of our argument for the impossibility of human-robotic communication (Gelepithis 1991, 2001a, 2002b).
- b) Note a further consequence of our research, namely, that machine consciousness when truly developed will be fundamentally different from that of humans.

In parts one and two we outlined an alternative axiomatic approach to the development of unified theories of mind and made three points: (i) the necessity of placing a theory of consciousness within the wider framework of a UToM; (ii) the use of language as a tool for developing UToM; and (iii) the inadequacy of presently used mathematics.

In part three, we bring the above three points to bear upon the problem of the nature of consciousness. Specifically:

- a) We sketch a conception of consciousness, in both humans and machines, in terms of the process of understanding (within our overall approach outlined in part two).
- b) We suggest neighbourhood systems in topology as a mathematical tool able to capture some –but definitely not all- of the complexity of the human brain/mind.

The latter suggestion is partly expressed as a modeling correspondence between elements of a neighbourhood system and key semantical notions.

Gelepithis, P. A. M. (1991). The possibility of Machine Intelligence and the impossibility of Human-Machine Communication. *Cybernetica*, Vol. XXXIV, No 4, pp 255-268.

Gelepithis, P. A. M. (2001b). Complex Intelligent Systems: Juxtaposition of Foundational notions and a Research Agenda. *Entropy* 2001, Vol 3, 247-258.

Gelepithis, P. A. M. (2002b). An Axiomatic Approach to the Study of Mind. *Res-Systemica*, Vol. 2, Special Issue: Proceedings of the Fifth European Systems Science Congress, October 2002, Crete. <http://www.afscet.asso.fr/resSystemica/>.

Gelepithis, P. A. M. (2004a). Criteria and Evaluation of Cognitive Theories. *Behavioral and Brain Sciences*. Forthcoming.

Gelepithis, P. A. M. (2004b). Remarks on the Foundations of Cybernetics and Cognitive Science. *Kybernetes*. Forthcoming.

Pentti Haikonen

Conscious Machines and the Mind-body Problem

Our mind seems to be immaterial. This is our undeniable everyday experience. Our mind seems to refer to and be aware of external and internal, abstract and concrete entities effortlessly and without any perceived material machinery. We perceive the external world; tables, chairs, books, whatever, directly without the awareness of any intermediate symbols or representations. On the other hand our physical responses are definitely material, yet initiated and controlled by our immaterial mind. We get the mind-body interaction problem: How can something immaterial control material processes? This problem is carried over to robotics, too. A conscious robot should also have an immaterial robot mind that can control the material body. This would seem to be next to impossible and we are stuck with the mind-body problem. Any successful design proposal for conscious machines must also propose a solution to this problem. Immaterial minds cannot be created by material means; this is true by definition. If mind was indeed immaterial then obviously any material construction effort would fail and our work would be futile. We can go on only if we assume that in reality mind was based on carrying material symbols and processes, but how could this be given our everyday perception of the contrary?

What evidence do we have about the immaterial nature of mind? We have our subjective perception only; we are not aware of any material processes behind our thinking and consciousness. However, for the observer self the appearance of mind "without the awareness of any carrying material symbols or processes" and "without any carrying material symbols or processes" can be the same. Therefore, the missing perception of material carrying symbols or processes does not prove that these would not exist. Yet, for centuries common people and philosophers alike have made this logically unsound conclusion, which in turn has led to the idea of immaterial mind. If our mind is actually material then why can't we perceive the material machinery behind it? Or, what would it take for us to be able to perceive the machinery, the neural firings, etc.? It seems that whatever it takes we do not have and the material machinery remains beyond our introspective powers. The material carrier mechanism remains transparent to us and we perceive the carried information only; the percept of an object *is* the symbol for that object.

This view offers design guidance for conscious machines. Machine minds that appear as immaterial to the machine self can be created if the information representation and processing method supports carrier medium transparency.

An Implementable Architecture for Conscious Machines

An architecture for conscious machines is described here. This architecture is based on the author's consciousness studies, computer simulations and hardware development work. Actual microchips suitable for the implementation of this architecture are being developed.

The limitations of present-day AI seem to arise from the fact that computers do not really think or utilize meanings. To overcome these limitations it is argued that an advanced information processing machine should think in human fashion. It should have a flow of inner imagery and inner speech, which it could control in meaningful ways; it should be able to reason and imagine. It should utilize symbolic processing with meanings and significance in the human sense. Obviously this kind of machine could converse directly with humans and could also understand books, movies etc. without any artificial devices.

The proposed architecture is designed to comply with the requirements of perception, attention, match/mismatch/novelty detection, good-bad significance evaluation, system reactions, emotions, learning and introspection. Distributed signal representations are used with associative processing. Combinatorial explosion is avoided by the use of attention, controlled by match/mismatch situations and evaluated significance. The architecture supports the flow of inner speech and inner imagery. These and any sensory percepts form the instantaneous input to the machine.

The machine operates with grounded meanings and significance. The basic meanings of the signals are grounded to sensations, secondary meanings are attached by cross-associations. In this way the signal arrays can come to represent entities beyond their direct meanings.

This architecture supports an "immaterial mind". The meaning-carrying signal arrays are transparent to the machine. Therefore the machine "sees" only the carried meaning or modulation without any material carrier. The machine while able to report the mental content will not be able to report the existence of any carrying signals or material processes, hence the mental content will appear as immaterial mind.

Due to various cross-connections the machine is able to introspect its mental content and report it in various sensory terms, such as images and speech. This faculty is one of the hallmarks of consciousness, especially when taken together with the general operational way of the machine.

Germund Hesslow

The inner world as a consequence of behavioural and perceptual simulation

The 'simulation' theory of thinking rests on three assumptions about brain function. First, behaviour can be simulated by activation of motor mechanisms as during overt actions, but while suppressing its execution. Second, perception can be simulated by internal activation of sensory mechanisms as during normal perception of external stimuli. Third, both overt and covert actions can elicit perceptual simulation of their

probable consequences. A large body of evidence supports these assumptions. It is argued that the simulation mechanism automatically gives rise to (and thus explains) many of the phenomena that characterise our inner world and consciousness.

Owen Holland

Plans and the structure of consciousness

In certain classes of environment, an embodied autonomous agent that can plan will outperform an equivalent agent that cannot. This talk will examine what is required for effective planning, and will explore the effects on planning systems of the constraints acting on evolved agents. Many of the apparently anomalous characteristics of consciousness will be shown to be consistent with the operations of a planning system in which a structure, process, or entity forming an Internal Agent Model (IAM) interacts with an internal world model to generate predictions about possible, desirable, and probable future actions. In particular, the IAM will be identified as the experiencing self. The implications of this approach for machine consciousness will be examined, and an attempt will be made to relate the theory to Aleksander's axioms for consciousness, and to Sloman's virtual machine architectures.

Jacques Lacombe

The European Commission Perspective

Beginning in 1999 under the 5th Framework Programme for Research, the IST (Information Society Technologies) Programme has launched a series of FET (Future and Emerging Technologies, <http://www.cordis.lu/ist/fethome.htm>) initiatives involving collaboration of neurosciences and IT research communities for their mutual benefit. A synthesis of currently running projects will be presented (see <http://www.cordis.lu/ist/fetni-nt.htm>), followed by an overview of perspectives that could open for research on consciousness.

Riccardo Manzotti

A process based architecture for artificial consciousness: from ontogenesis to phenomenal experience

A conscious being is a system that experiences (feels) something. To build an artificial conscious being we must deal with what is to feel something. This something is the content of conscious experience. As a working hypothesis we propose that experiencing something is a kind of causal relation with that something. What kind of causal relation? We propose a particular kind of process (elsewhere called reciprocal causation, mutual causation, or co-causation). Ontologically speaking the process is prior both to the existence of the object and of its representation. Represented and representing entities are two different ways of looking at the process of becoming of what can be later described as a subject or as an object.

By applying the same kind of rationale, the same kind of indivisible process is responsible for all conscious representations by means of the ontogenetic development of the brain. We suggest that every external event, we are conscious of, has taken part in the developmental history of our brain. Be conscious of something means to be a

particular process. The ontogeny of the subject is the sum of these processes. A process that takes part in the constitution of the subject is called onphene (ontos + phenomenon) since it defines both what there is (the ontology of the world) and what is perceived (the phenomenology of the world). Ontogenesis and ontology are the two sides of the same coin. Representation and existence, mind and world, subject and object are different perspectives on the becoming of reality. There is no more need of a dualistic vision of reality. Reality is one.

An artificial conscious being is a system whose structure is totally built by these onphenes. Its development is driven by these onphenes triggered by physical reality. An onphene is a process in which the occurrence of an event creates the conditions for the occurrence of an event of the same kind (if I am impressed by a landscape, I will try to repeat that experience). Then the design of an artificial conscious being is based on an architecture capable of letting external events to provoke the repetition of events of the same kind. In a developing artificial conscious being, this attitude to repeat events can be named a motivation. An artificial conscious being is a system capable of developing new motivations on the basis of its experience.

Peter Redgrave

Biological solutions to the selection problem.

A selection problem arises whenever two or more competing systems seek simultaneous access to a restricted resource. Consideration of several selection architectures suggests there are significant advantages for systems that incorporate a central switching mechanism. We propose that the vertebrate basal ganglia have evolved as a centralised selection device, specialised to resolve conflicts over access to limited motor and cognitive resources. Analysis of basal ganglia functional architecture and its position within a wider anatomical framework suggests it can satisfy many of the requirements expected of an efficient selection mechanism. The implications of this view for the selection of actions and selective attention will be considered.

Geraint Rees

Imaging the NCC

The immediacy and directness of conscious experience belies the complexity of the underlying neural mechanisms, which remain incompletely understood. I will review some examples of our recent functional magnetic resonance imaging studies in normal subjects and patients with focal cortical lesions, which provide new insights into the mechanisms involved. These data complement behavioral, neuropsychological and electrophysiological findings by suggesting that activity in functionally specialized areas of ventral visual cortex is necessary for visual awareness. However, our recent work suggests that activity in ventral occipital and temporal cortex is not sufficient to support conscious vision without a contribution from parietal and prefrontal areas. Such a contribution may reflect processes such as selective attention and working memory. Reciprocal interactions between parietal and ventral visual cortex may thus serve to selectively integrate internal representations of visual events in the broader behavioral context in which they occur, leading to the richness of our conscious experience and providing a fundamental neural substrate for conscious visual experience.

Multilevel Mechanistic Models and the Explanation of Consciousness

When we search for explanatory unity in the study of consciousness, we should first consult philosophy of science to clarify the notions of “unity” and “explanation”. Theoretical unity could be achieved in the form of a coherent, progressive research program. Such a program should be committed to empirically plausible philosophical background assumptions about consciousness that could be shared by most empirical scientists working in the field (Revonsuo 2000). My proposal for such background assumptions would be to regard consciousness as a real biological phenomenon residing within the confines of the brain, or biological realism. This background assumption is in harmony with empirical cognitive neuroscience. Furthermore, biological realism implies that the framework of multilevel mechanistic explanation, widely applied in the biological sciences (Bechtel & Richardson 1993; Machamer et al. 2000), should also apply to the explanation of consciousness. In this framework, the explanation of consciousness is the task of constructing a multilevel mechanistic model with a detailed phenomenal-level description, and a causal-mechanical network surrounding the phenomenal level.

The multilevel model makes explicit the relationship of consciousness to the lower and higher levels of organization as well as to events causally modulating consciousness. This multilevel explanatory framework might bring significant unity to the empirical research on consciousness because it allows all the different sources of empirical evidence to be integrated under one framework. We could ask, concerning any particular empirical approach: (1) Does it tell us about the internal quality and organization and the temporal dynamics of consciousness, on phenomenological terms? If yes, then it contributes to the phenomenal-level description. (2) Does it describe, at an abstract computational level, what kind of input-output transformations the phenomenal level is capable of performing? If yes, then it contributes to the input-output description of consciousness. (3) Does it contribute to our understanding of where in the brain the underlying constitutive mechanisms of consciousness might be located, or what kind of neurophysiological activity might be involved (Revonsuo 2001)? If yes, then the data contributes to the downward-looking explanation of consciousness. (4) Does the data reveal how conscious information, in interaction with sensory and motor mechanisms, causally contributes to organism-environment interaction and behavior? If yes, then it contributes to the upward-looking, contextual explanation of consciousness. (5) Does it tell us about the preceding events (say, the early processing of sensory input) that causally modulate consciousness? Then it is a contribution to the etiological explanation of consciousness. Doesn't the empirical data tell us anything that could be fitted into the multilevel explanatory framework? Then that particular approach is not useful to the description and explanation of consciousness, and should perhaps be abandoned.

The framework of mechanistic biological explanation could show the contribution and the relative importance of different sources of empirical evidence to the overall explanation of consciousness. The overall goal would be to add new pieces to the multilevel and multidimensional causal-mechanical model of consciousness, which would be not unlike many other mechanistic models of complex biological phenomena. The model would also indicate the areas where we still lack empirical data, and which

parts of the description and explanation are already approaching completion. In this paper, I will first describe the mechanistic explanatory framework at a more general level, and then show how various sources of evidence concerning visual consciousness could be fitted into this framework.

Miguel A. Salichs

A systemic study of consciousness

From a point of view of an engineer who works in the design and construction of robots, that intend to be autonomous and intelligent, the approach to the problem of consciousness can be establish as a set of questions: What are the benefits that can be obtained from consciousness? Is consciousness an essential specification for an intelligent autonomous machine? In case of an affirmative answer, then the great challenges appear: Is it possible to make conscious machines? If so, which are the basic elements needed to produce consciousness? And how they must be assembled?

Consciousness is the result of very complex mental processes. The study of these processes using a systemic approach can help to answer the previous questions. Based on this analysis, we conclude that consciousness is an essential characteristic of human intelligence. In particular, it is associated to our capabilities of advanced learning and planning future actions. Consciousness should be, so on, a very valuable tool for an autonomous robot that might work in complex environments.

A key aspect to understand mental processes in humans, as well as in many animals, is the capacity to integrate the information obtained from the world as dynamic systems. A dynamic system is characterized by its components, the relations among them and its behaviours. When we are watching, we do not see a spatial distribution of light. We see objects. Objects that can be formed by related parts and that behaves in a specific way. That is, we perceive systems. The three characteristics of those systems: components, relations and behaviour are essential aspects to understand how we perceive the world, and none of them can be ignored. In some studies, the perception of the world has been classically interpreted only in terms of elements, obviating the importance of relations and behaviours. The perceived systems can be as simple as a stone and as complex as a person. Here the adjectives simple and complex are interpreted in terms of system complexity. That is, taking into account the complexity of the system elements, relations and behaviours. A primary level of consciousness, also called core consciousness, consist in the perception of the self, just as another system and using the same mechanisms used to perceive the rest of the world, but with the particularity of adding to the systems the extra information provided by introspection, like emotions or thoughts.

A secondary level of consciousness, also called extended consciousness, emerges from the human capacity to reason about the systems, either on the real world or in the virtual world created by our imagination. Those systems include the self, associated to the core consciousness. That gives us a formidable capacity. We are able to reason about ourselves. That permits us to analyse our past actions and learn of our errors and successes. But that is not all; using our capacity to reason about ourselves we can plan future actions. This is carried out by reasoning about the expected results of our possible actions confronted with the behaviour of other systems of the world.

Ricardo Sanz

Engineering Conscious Machines

There are several ongoing attempts to build conscious machines using different types of technologies and engineering methods: "conventional" symbolic AI, neural networks, non-linear dynamical systems, etc.

If a computer-based mind for a machine is going to be made as conscious as humans, it is quite clear that it will be a complex software/hardware application. There are several approaches to engineering complex software systems and they will be reviewed in this talk with a particular emphasis on constructive methods (obviously due to a bias of the speaker). Emergent methods will be also analysed and convergent technology introduced.

We will analyse the possibilities of the different methods and try to extrapolate from present software engineering technology the degree of effort needed, the time scope and the tradeoffs of different designs and implementation technologies.

Murray Shanahan

Consciousness and Cognitive Robotics: Two Research Agendas or One?

Many of the issues being addressed by researchers in cognitive robotics are also high on the agenda in the blossoming field of machine consciousness.

This talk will outline some ongoing work on visual perception with an upper-torso humanoid robot. The aim of this work is to give the robot equivalent sensory-motor skills to those of a young infant, in particular the ability to perceive, grasp, and manipulate arbitrary objects in a natural environment. The talk will offer some speculation about the relationship between consciousness and the kind of awareness of the world these skills necessitate. We also intend to endow the robot with a form of visual imagination, enabling it to internally simulate the effects of hypothetical actions and events. Accordingly, the talk will also offer some speculation about the role of the imagination in consciousness.

Meanings, Theories, and Models of Consciousness

Many people assume that the word "consciousness" refers to some kind of unique, well-defined state or process which is either present or absent, so that it makes sense to ask what IT is, how IT evolved, which animals have IT, which neural mechanisms bring IT about, whether machines can have IT etc.

This is actually a kind of deep linguistic self-deception, as indicated by the fact that people differ so much regarding what they mean by "consciousness" -- e.g. how they define IT. Even the same person can be inconsistent, e.g. sometimes claiming that consciousness is absent during sleep and later claiming that it is present during dreaming.

A popular alternative view that consciousness is just a matter of degree of something fails to account for the variety of types of phenomena. Differences between an ant and an ape are not merely differences of degree: there are many kinds of things an ape can do that an ant cannot. Searching for differences of degree is the wrong way to understand biological variety: e.g. because mutations and crossover produce differences of kind.

A third stance, our design-based stance, treats the noun "consciousness" as referring loosely to a large and ill-defined variety of states and processes in which organisms (and other machines) have or acquire information. From this viewpoint an eagle's consciousness will involve different kinds of phenomena from that of an insect. Likewise a newborn human infant's consciousness will lack many features of a typical normal, wide-awake adult's.

For instance: not all animals have adult human-like self-consciousness, i.e. consciousness of being conscious of anything, and probably human neonates don't have that sort of consciousness. Likewise different varieties of consciousness are to be found in people with various kinds of brain-damage or other pathologies.

In order to understand this huge variety of types of consciousness we need to consider the variety of information-processing architectures possible for organisms and machines, and then, for the different sorts of architectures, to analyse the kinds of consciousness they support, which will also involve analysing the varieties of perception, learning, decision making, affective states, learning, development, communications and actions they support.

To provide an ontological framework, we offer a (first draft) generic schema (CogAff) for thinking about and describing a wide range of types of architectures, and we conjecture that humans have a particularly rich instance of this schema, which we call H-Cogaff.

If normal adult humans conform to H-Cogaff, they will have at least reactive, deliberative and meta-management layers in their architectures, all concurrently active and not forming any simple control hierarchy. (These subsume the six layers in Minsky's 'The Emotion Machine' I think).

Different forms of consciousness will be supported by the different mechanisms in the different layers.

For instance, primitive kinds of self-consciousness involving proprioceptive feedback in reactive layers will be shared with many other animals. However the ability to monitor, categorise, compare, and evaluate or remember one's own deliberative processes and relate them to possible mental states of others will not be possible without both a deliberative and a meta-management layer.

A deliberative layer requires, and provides mechanisms for, consciousness of chunked aspects of the environment and the ability to learn and use associations between one's actions and their consequences. For instance advanced deliberative mechanisms support consideration of what might happen, what might be the case, what might explain something observed. Meta-management extends that to learning and thinking about one's own possible *internal* states and processes.

Insects very probably lack both: it's not just that they have tiny degrees of them. (This is an empirical claim and could turn out false!)

A fully developed version of this architectural theory will explain what it is about the human architecture that gives rise to particular popular beliefs and muddles about consciousness in humans.

For instance, robots with human-like architectures will fall into the same traps as humans when they philosophise about consciousness, as several science fiction writers have predicted.

Notice that this architecture-based theory implies that there can be many concurrent causally active processes in virtual machines within the same individual. It does NOT imply that mental states can be defined in terms of their externally observable input-output contingencies. On the contrary, many interesting ones cannot -- requiring us to replace the simple-minded forms of functionalism normally discussed by philosophers with engineering-inspired versions: Virtual Machine Functionalism. [Most states of a sophisticated operating system cannot be defined by input-output contingencies of the whole machine either.]

These ideas are developed in papers and presentations on the CogAff web-site:

<http://www.cs.bham.ac.uk/research/cogaff/>

<http://www.cs.bham.ac.uk/research/cogaff/talks/>

including a paper on consciousness jointly authored with Ron Chrisley published in a recent issue of the Journal of Consciousness Studies edited by Owen Holland.

The H-CogAff architecture

Some of the implications of the postulated H-Cogaff architecture will be analysed, for instance the implication that different perceptual and motor layers co-evolved with different central layers, and the implication that there are probably far more types of learning and development than have hitherto been studied (e.g. different kinds of learning in different parts of the architecture, and some kinds of learning linking different architectural layers). H-Cogaff also suggests a wide variety of types of affective states, making nonsense of simplistic taxonomies of types of emotion. If there's time, I'll discuss architectural and representational requirements for perception of affordances -- these inherently involve counterfactual conditionals.

However the theory still has many gaps, and the broad-brush divisions postulated in the CogAff schema are clearly inadequate in relation to the known diversity of biological phenomena. So we need a long-term research programme aimed at extending and refining our schema, along with its ontology both for architectures and for mental states and processes.

This, in turn, should lead to a deeper richer variant of the H-Cogaff architecture, which is better able to explain the enormous variety of commonplace and bizarre mental phenomena found both in everyday life and in neurological and psychiatric wards, especially in the many amazing kinds of development that occur in the first ten years of a child's life.

For instance we still lack a good theory of what it is to find something funny, or what aesthetic pleasure is. If our future robots cannot enjoy hearing a Bach two-part invention or be entertained by Asimov's story about robots becoming religious bigots, then we'll have missed something. But there's a long, long way to go, and part of our aim is to draw attention to the many phenomena that we cannot yet explain or model in order to produce requirements specifications for conscious machines of many kinds (biological and artificial).

This defines long term collaborative research goals and should help to defragment AI and cognitive science, and divert attention away from silly debates such as whether GOFAI has failed or whether connectionist models are best, or whether dynamical systems or embodiment are the key to anything.

John Taylor

The Data of Consciousness

There are two aspects to the data of consciousness: subjective and objective. For the former, both the areas of Western phenomenology and Western cognitive science must be considered, as well as the evidence from recent analyses of diseases of the mind (considered from the patients viewpoint). For the latter, much Western cognitive neuroscience data is relevant as to how conscious brain activity is created from that of which there is no consciousness.

This report starts by describing the analysis of the pre-reflective self or ipseity, and how it had developed from the earlier work of Husserl and the group of Sartre, Henry, Merleau-Ponty and others. The manner this is being used in understanding schizophrenia and related mental diseases will be described, following the work of early workers on schizophrenia and more recent analyses of Kimura Bin, Louis Sass and Josef Parnas. The relation of this to the 'null-point' of experience as arising from the body will be considered, as will the relation of *ipseity* to meditative states.

Cognitive scientific advances through brain imaging will be related to these inner experiences. Data from schizophrenics especially will be considered. Then the manner that attention functions as a gateway to consciousness will be described, with brief coverage of recent claims that this is not so. Various paradigms crucial to be explained by models of consciousness, such as the attentional blink and streaming, will also be described.

These, and related topics such as the motion after-effect, will be considered, and data indicating the brain regions involved described, and their related timings. Finally the nature of motor awareness will be considered, and the present experimental situation relating motor control to attention discussed. The basic problem: does there exist a motor awareness separate from sensory consciousness, will then be considered to complete this summary of some aspects of the data on consciousness.

Taylor JG (2002) Trends in Cognitive Sciences 6:206-210

Taylor JG (2002) Journal of Consciousness Studies 9:3-22

Taylor JG (2001) The Importance of the Parietal Lobes for Consciousness.

Consciousness and Cognition 10:379-417

The CODAM Model for Consciousness

The CODAM (Corollary Discharge of Attention Movement) Model was proposed in 2000 to develop a model of the creation of conscious experience based on an engineering control approach to attention. The model will be described, starting with a brief overview of control models and their highly successful development and application to motor control by the brain. These are extended by developing a basic control model for attention, and then its further extension to inclusion of an observer or predictor. The introduction of suitable sensory and efference copy buffers leads to a model of attention possessing the ultimate in efficiency in driving the movement of attention as rapidly as possible. This is CODAM.

The CODAM model will be analysed to explain how the sense of 'what it is like to be' or *ipseity* can be explained, being the early activation of the corollary discharge buffer, before activation of the sensory buffer by the attention-amplified input. This will be explored both for the siting of the various components in the brain as well as for its explanation of the range of experiences in schizophrenia (especially that of hyper-reflexivity).

The application of CODAM to the various pieces of objective data will then be

considered, using results from recent simulation of the Posner benefit paradigm and the attentional blink, as well as others. The evidence of the CODAM signal (activation of the corollary discharge buffer) as being the P2/N2 complex will be discussed. The nature of extension of CODAM to motor control will then be described. This leads to a full theory of the highest-order control circuitry of the brain, involving various predictors and monitors at various levels of conscious control. The question of separate motor awareness will be discussed in terms of CODAM.

Finally a program for future work to explore CODAM in various experimental areas, and the numerous open question raised by CODAM, will be described.

Taylor JG (2000) Attentional Movement: The Control Basis for Consciousness. Neuroscience Abstracts 26, 231 #839.3

Taylor JG (2003) Paying Attention to Consciousness (submitted to Progress in Neurobiology)

Taylor JG (2003) The CODAM Model and Deficits of Consciousness I: CODAM, KES (Oxford) to appear.

The Codam model and Deficits of Consciousness II: Schizophrenia and Neglect/Extinction. KES (Oxford) to appear.

Taylor JG & Fragopanagos N (2003) Simulations of Attention Control Models in Sensory and Motor Paradigms. IJCNN2003 to appear.

Tom Ziemke

What's life got to do with it? Why an artificial self might have to be autopoietic

Many researchers interested in the possibility of consciousness in computers, robots or other artefacts agree that several aspects of consciousness require an agent, i.e. a system that interacts with its environment by means of perception, action, etc. Much of the discussion focuses on the question which type(s) of information processing, representation, learning processes, etc. an artificial agent would have to be equipped with in order to achieve different types or levels of consciousness. The question exactly what it takes for an artificial (or natural) system to be an agent in the first place, and thus a candidate for an artificial self or consciousness, on the other hand, receives much less attention, although it has been pointed out by several authors (e.g. Franklin & Graesser) that current definitions of what constitutes an agent are somewhat vague. This talk discusses the theories of von Uexkuell, Maturana & Varela, Bickhard, Christensen & Hooker which, roughly speaking, argue that what it really takes for a system to constitute an agent capable of cognition, consciousness, etc. is that it is a living system. Implications for the possibility of consciousness in artefacts are discussed.

Towards neuro-robotic models of conscious thought as simulation of sensorimotor processes

Hesslow (2002) put forward the 'simulation hypothesis', i.e. the idea that conscious thought and the experience of an 'inner world' could be explained as an internal simulation of perception and behavior. This talk discusses experimental work from our own lab (e.g., Jirnhed, Hesslow & Ziemke, 2001; Ziemke, Jirnhed, Hesslow, subm.),

as well as related work by others, that aims to model the simulation hypothesis, in particular simulation of perception, in neural-net-controlled robots.

Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Science*, 6(6). 242-247.

Jirnhed, D.-A.; Hesslow, G. & Ziemke, T. (2001). Exploring Internal Simulation of Perception in Mobile Robots. In: Arras, Baerfeldt, Balkenius, Burgard, Siegwart (eds.) 2001 Fourth European Workshop on Advanced Mobile Robotics - Proceedings (pp. 107-113). Lund University Cognitive Studies, vol. 86. Lund, Sweden.

Ziemke, T., Jirnhed, D.-A., Hesslow, G. (subm.). Toward Internal Simulation of Perception in Mobile Robots. Submitted for journal publication.

Assessment of the result

Expected outcome of the workshop

The expected outcome of the workshop was divided into two parts.

The first part was a collection of scientific-technical results:

- A sound characterization of the field
- An engineering definition of consciousness
- A seed for a sound theory of consciousness
- Draft requirements for architectures for conscious machines of various sorts

The second part is a planning of future research activities in the field:

- Research roadmap
- Launch a network of excellence
- Elaborate concrete collaborative research projects

Let's evaluate the results based on this expected outcome.

Real outcome of the workshop

Scientific-technical results

The **sound characterization of the field** was, as expected, one the hardest problems in to deal. It was clear however, that the focus on machine implementation was contributing to the clarification of the research targets.

Quoting a private communication from an european commission officer regarding the emerging vision from these workshops:

In particular the report clearly shows that —far from being an esoteric endeavour— research in consciousness has clear roots in experimental science and (!!!) engineering.

It reminds me a bit of quantum computing: up to the time when experimentalists took quantum computing serious the whole area of 'fundamental quantum mechanics' was in the hands of philosophers (more precisely in this case pseudo-philosophers).

As soon as experimentalists took over many seemingly mysterious aspects became clear.

A **seed for a sound theory of consciousness** was envisioned as many of the speakers proposed theories based on control based on internal models. A basic design that is already being addressed by some of the attendants in projects funded by national bodies.

Draft requirements for architectures for conscious machines of various sorts were commented in relation with the collection of faculties that seem related to the conscious phenomena:

Started on-line discussion

An **engineering definition of consciousness** was proposed and discussed in a final wrap-up session of the workshop. A proposal centered around control mechanisms based on integrated mental models including the self were seen as a potential candidate both for natural and artificial systems. This agrees with Baars' vision about GW and supporting data

Planning of Research Activities

The workshop provided more results in the area of research planning and organization. All the participants agreed creating and launching a network of collaboration around the topic.

The results of the workshop were placed -with one exception.- in the workshop website, demonstrating an openness of the researchers to widely disseminate and create an open community.

A website (www.artificialconsciousness.org) is underway and a book with text version of speeches is being elaborated.

Some other actions were decided during the workshop:

- Solicitation of a IST Concerted Action in the area of FET Open¹
- Start activities toward a future Network of Excellence (including non Europeans)
- Contacting Neuro-IT people
- Define Potential Research Projects
- Preparation of an ESF Scientific Programme to be submitted to ESF PESC (already done in the outline phase).

Focus on implementation

In the original proposal of the workshop and in the call for participation, it was said that the workshop focused on artificial consciousness and hence rejecting non-implementable theories:

Non-implementable theories (*i.e.* theories that claim that consciousness is *indefinable or unknowable*, theories that say that consciousness is *epiphenomenal* and hence has no causal powers, theories that say that consciousness is *just a myth invented* by philosophers, theories that say that *no machine could have it* or theories that say that *machines that are indistinguishable from us could lack consciousness*) are not useful for the engineering work. It is necessary to re-consider their suitability as scientific

¹ Also participation in the NEST initiative “what it means to be human” has been proposed after the workshop.

explanations.

This was considered a very valuable approach by the workshop participants because it helped limit -although not completely- the amount of metaphysical discussion.

Final Program

The program was three days long. Each day had a similar structure but was somewhat focused on a particular aspect of the domain (data, theories, research). This program (with active links) can be found at:

<http://aslab.disam.etsii.upm.es/public/events/moc/program.html>

Day 1: Data and General Considerations

09:00	Welcome and introduction
09:10	Opening speech: Ron Chrisley <i>Data to be explained by a Theory of Consciousness</i>
10:10	Coffee Break
10:30	Morning session
	Germund Hesslow <i>The inner world as a consequence of behavioural and perceptual simulation</i>
	John Taylor <i>The Data of Consciousness</i>
12:30	Lunch
13:30	Afternoon Session
	Pentti Haikonen <i>Conscious Machines and the Mind-body Problem</i>
	Jim Doran <i>Behaviours, Day-Dreams, Plans, and Consciousness</i>
	Geraint Rees <i>Imaging the NCC</i>
15:30	Coffee Break
16:00	Miguel Angel Salichs <i>A Systemic Study of Consciousness</i>
17:00	Discussion: <i>The relevant data and the targets of potential explanations of consciousness</i>
19:00	Dinner

Day 2: Potential Unified Theories

09:00	Opening speech: Aaron Sloman <u>Meanings, Theories and Models of Consciousness</u>
10:05	Coffee Break
10:30	Morning session
	John Taylor <u>The CODAM Model for Consciousness</u>
	Peter Redgrave <u>Biological Solutions to the Selection Problem</u>
	Antti Revonsuo <u>Multilevel Mechanistic Models and the Explanation of Consciousness</u>
12:30	Lunch
13:30	Afternoon Session
	Tom Ziemke <u>What's life got to do with it? Why an artificial self might have to be autopoietic</u>
	François Anceau <u>A Model for Consciousness Based on the Sequential Behavior of Voluntary Actions</u>
	Axel Cleermans <u>The Search for the Computational Correlates of Consciousness</u>
15:30	Coffee Break
16:00	Owen Holland <u>Plans and the structure of consciousness</u>
	Petros Gelepithis <u>Consciousness within Unified Theories of Mind</u>
17:30	Discussion: <i>The (multiple) nature of consciousness and the feasibility of a ultimate theory</i>
19:00	Dinner

Day 3: Research and Engineering

09:00	Opening speech: Ricardo Sanz <u>Engineering Conscious Machines</u>
10:05	Coffee Break
10:30	Morning session
	Jacques Lacombe <u>The European Commission Perspective</u>
	Tom Ziemke <u>Towards neuro-robotic models of conscious thought as simulation of sensorimotor processes</u>
	Aaron Sloman <u>The H-CogAff architecture</u>
12:30	Lunch
13:30	Afternoon Session:
	Riccardo Manzotti <u>A Process-based Architecture for Artificial Consciousness: from Ontogenesis to Phenomenal Experience</u>
	Pentti Haikonen <u>An Implementable Architecture for Conscious Machines</u>
	Murray Shanahan <u>Machine Consciousness and Cognitive Robotics: Two Research Agendas or One?</u>
15:30	Coffee Break
16:30:00	William Edmonson <u>General Cognitive Principles and the Structure of Behavior</u>
17:00	Closing Session: <i>Roadmapping for consciousness research and engineering</i>
19:00	Dinner

Final list of participants

Speakers

Aaron Sloman

School of Computer Science
The University of Birmingham
Birmingham, B15 2TT
England, UK

A.Sloman@cs.bham.ac.uk

Antti Revonsuo

Research Fellow of the Academy of Finland
Department of Philosophy
University of Turku
FIN-20014 FINLAND

antti.revonsuo@utu.fi

Axel Cleermans

Cognitive Science Research Unit
Université Libre de Bruxelles CP 122
Av. F.-D. Roosevelt, 50
1050 Bruxelles
BELGIUM

axcleer@ulb.ac.be

François Anceau

Chaire des Techniques Fondamentales de l'Informatique
CNAM
292 rue S^t Martin
75141 PARIS Cedex 03
France

anceau@cnam.fr

Geraint Rees

Institute of Cognitive Neuroscience,
University College London,
Alexandra House, 17
Queen Square London WC1N 3AR
UK

g.rees@fil.ion.ucl.ac.uk

Germund Hesslow

Department of Physiological Sciences
Neurophysiology
Lund University
BMC F10
Tornavägen 10
SE-221 84 LUND
Sweden

Germund.Hesslow@mphy.lu.se

Jacques Lacombe

European Commission
DG INFSO - F1 (BU33 3/65)
B - 1049 Brussels
Belgium

Jacques.Lacombe@cec.eu.int

Jim E. Doran

Department of Computer Science
University of Essex
Wivenhoe Park
Colchester CO4 3SQ
UK

doraj@essex.ac.uk

John Taylor

Director of the CNN
Department of Mathematics
King's College,
UK

john.g.taylor@kcl.ac.uk

Miguel A. Salichs

Catedrático de Ingeniería de Sistemas y Automática
Universidad Carlos III de Madrid
Escuela Politécnica Superior
Butarque, 15. Leganés
28911 Madrid
Spain

salichs@ing.uc3m.es

Murray Shanahan

Dept. Electrical and Electronic Engineering
Imperial College London
Exhibition Road
London SW7 2BT
UK

m.shanahan@imperial.ac.uk

Owen Holland

University of Essex
Department of Computer Science
Wivenhoe Park
Colchester Essex CO4 3SQ
UK

owen@essex.ac.uk

Pentti Haikonen

Principal Scientist
Cognitive Technology
Nokia Research Center
P.O. Box 407
FIN-00045 NOKIA GROUP
FINLAND

pentti.haikonen@nokia.com

Petros Gelepithis

Cognitive Science Laboratory
School of Computing and Information Systems
Kingston University, Penrhyn Road,
Kingston, Surrey KT1 2EE,
UK

Petros@kingston.ac.uk

Peter Redgrave

Department of Psychology
The University of Sheffield
Sheffield S10 2TP
UK

P.Redgrave@sheffield.ac.uk

Ricardo Sanz

Department of Automatic Control
Universidad politécnica de Madrid
José Gutiérrez Abascal 2
28006 Madrid
SPAIN

Ricardo.Sanz@etsii.upm.es

Riccardo Manzotti

Research Assistant in Robotics
LIRA-Lab,
DIST, Department of Computer Science, Robotics, and Bioengineering
University of Genoa
Via all'Opera Pia 13
16145 Genova
Italy

manzotti@dist.unige.it

Ron Chrisley

Senior Lecturer in Philosophy
School of Computer Science
University of Birmingham
Birmingham, B15 2TT,
United Kingdom

rlc@cs.bham.ac.uk

Tom Ziemke

Professor of Cognitive Science
Department of Computer Science
University of Skövde
Sweden
P.O. Box 408
S-541 28 Skövde
SWEDEN

tom@ida.his.se

William Edmonson

School of Computer Science,
University of Birmingham,
Edgbaston,
Birmingham,
B15 2TT,

UK

w.h.edmondson@bham.ac.uk

Non Speakers

Mercedes Lahnstein

Neural Systems Researcher
Interactive and Intelligent Systems Group
Dept. of Electrical and Electronic Engineering
Imperial College of Science, Technology and Medicine
Exhibition Rd., London SW3 5LX
UK

mlahnstein@hotmail.com

Rafael Chinchilla

ETS Ingenieros Industriales
Universidad Politécnica de Madrid
28006 Madrid
Spain

rchinchilla@etsii.upm.es

Dean Peters

School of Computer Science
University of Birmingham
Edgbaston, Birmingham B15 2TT
U.K.

D.A.Peters@cs.bham.ac.uk

Ignacio Lopez

ETS Ingenieros Industriales
Universidad Politécnica de Madrid
28006 Madrid
Spain

ilopez@etsii.upm.es

Marcos I. Quintana

School of Computer Science
University of Birmingham
Edgbaston, Birmingham B15 2TT
U.K.

M.I.Quintana@cs.bham.ac.uk

Catriona Kennedy

School of Computer Science
University of Birmingham
Edgbaston, Birmingham B15 2TT
U.K.

C.M.Kennedy@cs.bham.ac.uk

David Booth

School of Psychology
University of Birmingham
Edgbaston, Birmingham B15 2TT
U.K.

D.A.Booth@bham.ac.uk

Andrew Cook

CM Software
Oxford
UK

support@cmsoftware.co.uk

Statistical information

Country

UK: 16
Spain: 4
Belgium: 2
Sweden: 2
Finland: 2
Italy: 1
France: 1

Sex

Women: 2
Men: 26

Position

Academics: 19
Students: 6
Industry: 3