ARTIFICIAL COGNITIVE SYSTEMS IN FP7

A report on expert consultations for the EU seventh Framework Programme 2007-2013 for research and technology development*.

1. Motivation

The future prospect of **artificial cognitive systems** (ACS) which can perceive, reason and interact robustly in open-ended environments is an ambitious challenge which goes beyond today's systems engineering paradigm. Present day systems engineering relies on specifying every eventuality a system will have to cope with in the execution of its task(s) and programming the appropriate response in each case. With the abundance of ever cheaper, smaller sensors, actuators and wireless links that bring systems in contact with the real world and with other systems, this approach faces serious limitations:

- The **real world** and **rich virtual environments** (such as the Internet) are generally too nuanced, too complicated and too unpredictable to be summarised within a limited set of specifications; there will inevitably be novel situations and the system will always have gaps, conflicts or ambiguities in its own knowledge and capabilities.
- Even in situations where unpredictable events are less likely, the problem of **extracting meaning and purpose** from bursts of sensor data or strings of computer code arises, because we don't have a formalisation of information processing that embodies semantics. Coping with novelty would require that systems be capable of constructing novel meanings; presently however it is the system designer who encodes the consequences of semantic content for system behaviour.

The extensive programming and customisation which would be required to cope with all foreseeable events is computationally tractable only **for isolated problems or in one particular context**.

Artificial cognitive systems, as opposed to traditional machine or computer systems, can be characterised as systems which cope with *novel or indeterminate* situations, which aim to achieve *general goals* as opposed to solving specific problems, and which integrate *capabilities* normally associated with people or animals such as perception, learning, reasoning, communication, etc. Specifically they should be:

- *more robust* - performance should not degrade when they are presented with unexpected data;
- *more adaptive* – performance should be open (within reasonable constraints) to changing service requirements, without the need for extensive human intervention;
- *more effective* - performance should improve because they can predict or anticipate what might happen at some point in the future, near or far;
- *more natural* - performance should be tolerant to the ambiguity and uncertainty that is a consequence of dealing with humans and performance should improve with time.

This will require rethinking the way we engineer systems and – because we aim to get machines to exhibit performance capacities that resemble those of humans or animals - will borrow inspiration and insights from the bio sciences, social sciences and humanities. Engineering progress will depend on advancing scientific understanding of what both natural and artificial systems can and cannot do, and how and why.

2. Questions

With the seventh EU Framework Programme in mind, the Commission's Cognition Unit in DG Information Society and Media consulted a number of leading researchers in different disciplines to explore the potential for making progress toward the creation of a **scientific foundation** for

* drafted jointly by members of the "Cognition" unit of Directorate "Content" of Directorate General "Information Society and Media". The consultations took place in the period December 2005 – March 2006.

engineering artificial cognitive systems, and the positioning of artificial cognition as an **enabling technology** in many areas of applied systems engineering.

Five workshops[1] were held on selected topics with a view to soliciting recommendations on cutting edge research, longer and medium-term R&D goals, scope for interdisciplinary co-operation and impact.

The topics chosen do not define R&D modules; rather, they are supposed to highlight partly complementary, partly overlapping, aspects and directions of current ACS research. They cover a large ground but are by no means exhaustive.

Cognition and Machine Learning

What is the potential of Machine Learning research for advancing the engineering of artificial cognitive systems?

Cognitive Robotics

What is the potential of robotic devices for ACS research and application?

Cognition and Control

What impact can systems control research make on ACS research and vice versa?

Cognition and Language

What is the role of language-based communication for cognition and of cognition for language-based communication?

Cognitive Models and Paradigms

Can ACS R&D benefit from new developments in mathematical modelling and non-standard computing? To what extent would ACS require the emulation of the physical substrate of natural cognition?

3. Answers

The consultations strongly supported research into cognitive systems, pointing to areas where synergies between different scientific and engineering disciplines could be sought at European level. Salient points:

(i) **Machine Learning** or, more generally, **learning in artificial systems**, comprises a set of methods and techniques that are extremely relevant for ACS. It extends the remit of the latter from immediate natural environments (as for example of robots, robotic devices and other appliances and machines) to more generalised notions of environment, including all sorts of digital spaces.

(ii) **Robotics** is considered an important platform for ACS research, with clear mutual benefits to both domains, allying the embodiment and situatedness of a robot machine in the real world to advanced perceptual and higher-level cognitive capabilities. ACS research is badly needed to advance the scientific theories and models of robot behaviour.

(iii) Cognition and **control** are linked at several levels. Control theory can bring formal foundations to ACS research, and cognition can add higher levels of abstraction or sophistication to control systems. In industrial applications of control systems, cognition helps to develop more adaptive technology, more complex and at the same time more easily manageable engineering products.

(iv) **Natural language** is a "key product" of natural cognition. Hence, technology dealing with natural language would greatly profit from pertinent ACS research. This concerns equally all forms and modalities of purposeful communication across human-machine interfaces. On the other hand, ACS R&D benefits for example from models of language development in infants.

(v) Appropriate **models and paradigms** and their theoretical / mathematical underpinnings are needed to progress beyond the constraints of today's artificial systems, for example linked to the notions of emergence, self-organisation, evolution, and development, and the possibility of

[1] http://cordis.europa.eu.int/ist/cognition/presentations.htm

basing cognitive systems on future new substrates which would directly support functions related to self-modification, for example memory and learning.

(vi) **Bio-inspired** methods are considered crucial to bridging the gap between natural and artificial cognitive performance capacity, whether drawn from neuroscience, or broader life sciences and humanities. This holds for the lower (e.g., sensorimotor functions) as well as for the higher (e.g., impact of emotion on reasoning and decision-making) levels of cognition.

(vii) **Applications** which have an intrinsic need of cognitive methods, such as exploration and navigation, object recognition and manipulation in all sorts of environments, dealing with unexpected requests from human users, or developing better ways to manage complex, networked systems (e.g. intelligent traffic control) are likely to be in the first order.

(viii) The **potential impact** of research into artificial cognitive systems can be expected in several domains (see below), including **scientific research** itself, by promoting highly inter-disciplinary work on new combinations of topics not previously explored at a European level[2].

Detailed reports on the workshop discussions can be found at the following Web addresses:

Cognition and Machine Learning
ftp://ftp.cordis.europa.eu/pub/ist/docs/cognition/workshop_report20051219.pdf
Cognitive Robotics
ftp://ftp.cordis.europa.eu/pub/ist/docs/cognition/workshop_report20051220.pdf
Cognition and Language
ftp://ftp.cordis.europa.eu/pub/ist/docs/cognition/workshop-report-270206_en.pdf
Cognition and Control
ftp://ftp.cordis.europa.eu/pub/ist/docs/cognition/workshop-report-090306_en.pdf
Cognitive Models and Paradigms
ftp://ftp.cordis.europa.eu/pub/ist/docs/cognition/workshop-report-march06_en.pdf

4. Conclusions

ACS R&D provides a technology that enables significant progress in many key application domains, for example: robotics (for industrial manufacturing; field and service) and other types of assistive devices; human-machine interaction; vehicle control and traffic safety; management and control of transport, energy and communication networks; remote and on-site (environmental) sensing and monitoring; agency in content and service networks; medical diagnostics and therapeutics.

ACS R&D generated / inspired technology will be instrumental in making systems in these domains more robust and effective, more adaptive and efficient, safer and easier to use.

ACS R&D typically contributes to:

- the theory and application of learning in artificial systems;
- solving fundamental problems related to the extraction of meaning from all sorts of environments (through vision, audition, and other senses, and – more generally - through data analysis and interpretation);
- the theory, design and implementation of cognitive systems architectures;
- the development of criteria for benchmarking desirable system properties such as robustness, scalability and adaptivity, depending on the application environment at issue;
- the exploration and validation within the context of creating artificial cognitive systems, of the use of
  - existing or emerging (possibly bio-inspired) information-processing paradigms, and models of self-organisation and natural cognition (including human mental and linguistic development);
  - advanced sensor, actuator, memory and control elements and platforms, based on new materials and hardware designs.

Last but not least, ACS R&D provides insights into the limits of artificial and natural cognition.

[2] For work already undertaken see the EU Cognitive Systems initiative in FP6, http://www.cognitivesystems.eu

ANNEX – Expert participants list.

Elisabeth André, University of Augsburg;

Peter Auer, University of Leoben;

José Balcazar, Universidad Politècnica de Catalunya;

Rens Bod, University of St Andrews;

Ernst Buchberger, Medical University of Vienna;

Nicola Cancedda, Xerox Research Centre Europe;

Angelo Cangelosi, University of Plymouth;

Nicolò Cesa-Bianchi, Università degli Studi di Milano;

Anthony Cohn, University of Leeds;

Erik de Schutter, University of Antwerp;

Renato De Mori, Université d'Avignon et des Pays de Vaucluse;

Sue Denham, University of Plymouth;

Ezequiel Di Paolo, University of Sussex;

Jan-Olof. Eklundh, Royal Institute of Technology, Stockholm;

Nicolas Franceschini, Université de la Méditerranée, Marseilles;

Owen Holland, University of Essex;

Timo Honkela, Helsinki University of Technology;

Auke Ijspeert, EPFL, Lausanne;

Frederic Kaplan, Sony, France;

Daniel Kayser, Université Paris Nord;

Giuseppe Longo, Ecole Normale Supérieure, Paris;

John McCaskill, Ruhr University Bochum;

Barry McMullin, Dublin City University;

Martin Middendorf, Universität Leipzig;

Dunja Mladenic, Josef Stefan Institute, Ljubljana;

Alvaro Moreno, University of the Basque Country (UPV/EHU);

Klaus Obermayer, Technische Universität Berlin;

Paavo Pylkkänen, University of Skövde;

Sajit Rao, MIT - CSAIL;

Eytan Ruppin, Tel Aviv University;

Antonio Sala-Piqueras, Universidad Politécnica de Valencia;

Miguel Salichs, Universidad Carlos III, Madrid;

Ricardo Sanz, Universidad Politécnica de Madrid;

Brian Scassellati, Yale University;

Hartmut Schmeck, Universität Karlsruhe;

Juergen Schmidhuber, IDSIA, TUM Munich;

Gregor Schoener, Ruhr-Universität Bochum;

Michele Sebag, Université Paris-Sud Orsay;

Murray Shanahan, Imperial College, London;

John Shawe-Taylor, University of Southampton;

Mark Steedman, University of Edinburgh;

Albert Tarantola, Institut de Physique du Globe de Paris;

Gianluca Tempesti, EPFL, Lausanne;

Louis Ten Bosch, Radboud University of Nijmegen;

Volker Tresp, Siemens AG Central Research;

Theo Ungerer, University of Augsburg;

David Vernon, University of Genova;

Sethu Vijayakumar, University of Edinburgh;

Carl Vogel, Trinity College Dublin;

Jiří Wiedermann, Academy of Sciences of the Czech Republic;

Chris Williams, University of Edinburgh;

Klaus-Peter Zauner, University of Southampton;

Jianwei Zhang, University of Hamburg;

Andrew Zisserman, University of Oxford.