Towards a Scientific Foundation for Engineering Cognitive Systems
- A European Research Agenda, its Rationale and Perspectives -[1]

Hans-Georg Stork[2]

Abstract For more than ten years, beginning in the early 2000s, the European Commission has been supporting targeted research in the fields of Cognitive Systems and Robotics. In this note we discuss the rationale of the underlying research agenda (including its relevance to the BICA challenge), structure the large set of funded projects, and outline perspectives for future research.

Keywords: Cognitive Systems; Bio-Inspiration; Robotics; Cognitive Robotics; Research agendas; public funding.

1 Introduction: facts and figures

Topics broadly related to Artificial Intelligence (AI) have been part and parcel of European research programs ever since research funding at the European level began in the 80s of the previous century. *Cognitive Systems*, however, became prominent as a specific item on the research agenda only in the late 90s when, under the heading *Cognitive Vision*, a cluster of eight projects was launched in response to a growing demand for more powerful computer vision systems that were able to interpret what they saw and sensibly to act upon it.

From 2002 onwards, this line of funding has been extended to cover both, *Cognitive Systems* in general and *Robotics* (CSR for short). It has been firmly established as a key chapter of the 6th and 7th multiannual Framework Programmes for Research and Technology Development (FP6 from 2002-2006 and FP7 from 2007-2013 respectively), and codified in a series of usually biannual Work Programmes that underly the regularly published Calls for Proposals (EU Commission, n.d.).

So far, under this chapter, the Commission of the European Union (EU) (represented by its Directorate-General Information Society and Media) has spent more than half a billion Euros on nearly 140 projects and ancillary activities in the areas at issue.

[1] Preprint; to appear in Biologically Inspired Cognitive Architectures, Elsevier, 2012.
[2] From 1992 to 1/2012 the author served as Research Programme Officer in Directorate General Information Society and Media of the European Commission. From 2004 to 2012 he was involved in preparing and running the Cognitive Systems and Robotics chapter of the Information and Communication Technology research programme. In this article he presents his personal views which in no way engage his former employer.

As required by the rules of the European Framework Programmes, projects take the form of collaborations between at least three partners, academic and/or industry labs, that are institutionally independent of each other and reside in at least three different EU Member States or associated countries ("EU+"). In justified cases non-EU+ partners (e.g., from the US) can also form part of consortia and receive monies. These collaborations, based on a contractually agreed Description of Work, usually are scheduled to last between 2 and 5 years. Measurable progress (e.g., in terms of milestones and deliverables) towards proposed concrete goals is required. Proposals as well as progress and results are peer reviewed. The amount of funding typically ranges from 2 to 3 M€ (million Euro) for smaller projects and 6 to 10 M€ for larger ones.

2 Objectives, rationale and research questions

The overarching objective of European research funding (and presumably of public research funding anywhere in the world) is to boost economic growth through science-based innovation. Area-specific Work Programmes, including those soliciting basic research, have to take account of this, one way or another.

Accordingly, the CSR programme, while covering a fair number of basic research topics, ultimately aims at developing technical systems that are more robust, versatile, reliable, adaptable, and less demanding of human intervention than is possible given the current state-of-the-art; systems that are, in other words, capable of responding intelligently and largely autonomously to gaps in "their knowledge" and to situations or contexts that have not been specified in their design. Such systems should also be able to effectively improve their performance, for instance in delivering specific services; and they should be more natural in dealing with people where dealing with people is a requirement. (EU Commission, n.d.)

Progress towards endowing systems with these characteristics can only be achieved by finding, studying and adopting new engineering principles and approaches, based on largely common but as yet not fully explored scientific grounds.

Hence we can rephrase the main focus of the CSR programme as follows:

... to strengthen the scientific foundation for engineering artificial cognitive systems - i.e., artificial systems that perceive and (inter-)act, in accordance with a suitable understanding of their environment;

and, in doing so ...

... to foster technologies that enable a variety of applications involving interaction within "real world" environments pertaining to, for instance, robotics, assistive technologies, and multimodal man-machine interaction.

A key term in this description is real-world environment. The notion of environment encompasses more than we or our animal companions can cope with, only using eyes, ears and muscles, but no technical support. More generally, an "inner world" arising from an environment (an "outer world") is determined by the (types of) signals and data an entity operating in it is fit to process.

From a human perspective one may distinguish various classes of environments, for instance:

1. *Common Sense* environments, where people perceive and act, determined by the familiar spatiotemporal dimensions, limited parts of the electromagnetic and acoustic spectra, touchable objects, etc.; they may be *natural* (e.g., a forest, a mountain range) or *civilized* (e.g., a town, a house, a shopfloor, etc.);

2. natural environments at various scales, not directly or fully accessible through our own (bodily) senses and actuators; these may include the deep sea, outer space, etc., and even our own bodies;

3. virtual environments consisting of external representations of Man's (and machines') perceptions and reflections; e.g., *Digital Content spaces*, the Web;

4. technical systems embedded in environments of the first and second types.

While ultimately stemming from the *real world*, data pertaining to type 3 environments need not necessarily originate directly from sensing. Indeed, as far as the underlying data spaces are concerned, the primary focus of the CSR programme, is on data directly obtained through sensors within environments that belong to one of the first two classes. In turn, this focus entails an emphasis on systems with spatially delimited bodies (*embodied systems*), in particular robots; machines, that is, which are able to interact with physical entities and, in many instances, to move around in a given environment. Artificial sensors and actuators are of course not bound to the limitations of their biological counterparts. Hence machines equipped with them may generate individual *inner worlds* that from the very outset, differ greatly from those of living entities. A common trait though, of the environments these *worlds* must relate to is their open-endedness; they are dynamic and usually full of surprises.

What should a scientific foundation for engineering such systems support? There are *generic* and *specific* answers. The generic ones should be obvious, given the historical record of science and its relationship to technology and engineering. Almost none of our modern technologies would exist without the basic knowledge generated through scientific research: knowledge expressed in terms of models, principles and laws of nature, which set the boundaries within which engineering can happen and which inform the practice of engineering itself: the interwined activities usually referred to as specification, design, implementation, testing and deployment of systems and machines.

The specific answers depend on the specific issues that arise as a particular technology evolves. These issues become apparent through growing sets of research problems. A scientific foundation provides the ground for tackling such problems methodically and with some likelihood of success.

The most general (big) question, presumably, in relation to the CSR programme, is the following:

What (if anything) do we need to understand about cognition as a biological phenomenon in order to specify, design and build artificial cognitive systems?

The fact that cognition is first and foremost a biological phenomenon may not be a reason sufficient for drawing inspiration from biology when trying to build artificial systems with the above mentioned properties. Studies of birds' flight, for example, did not greatly contribute to advancing the art of building aeroplanes (apart from the beginnings of human flight when bird morphology actually did have some impact on engineering planes, but avian mechanics and physiology did not (Tennekes, 1996)). Likewise, although *mainstream* Artificial Intelligence (AI) research was more impressed with man's unique reasoning and planning capabilities than (for instance) with his *gut feelings* it managed to yield many interesting and useful results. But little did it contribute to creating the kind of systems aimed at under the CSR programme, or for that matter, to achieving what has come to be termed Artificial General Intelligence (Goertzel & Pennachin, 2006). (By the same token, modern aeroplanes and even drones do lack some of the most outstanding avian faculties.)

One must admit though, that contemporaneously with largely *meta-mathematics inspired* AI research (that appeared to dominate the agenda for a considerable period of time in the previous century) there have always been *engineering oriented* and *biology motivated* strands of activities. While picking up on issues pertaining to cognitive systems these strands were often not explicitly associated with *mainstream* AI. Both had in fact once (in the 40s and 50s of the last century) been intertwined in Cybernetics, the science of *communication and control in animal and machine* as Norbert Wiener put it (Wiener, 1948). It had certainly been an early contender in the race for artificial intelligence, yet fell behind for several reasons, one of them presumably being that its hardware requirements could not be met when (micro-)electronics and other relevant technologies were not sufficiently mature.

But the flame has been kept burning. And in light of the fact that natural cognitive agents (as individuals or species) are (up until now) practically the only entities that are capable of learning through acting on or interacting with complex dynamic environments, it now seems evident that the engineering of artificial cognitive systems can be informed by studying natural processes related to cognition and control, including the role of the physical substrates of these processes. Hence, to start from the biological origins of cognition in order to gain proficiency in building such systems, appears to be widely accepted ever since the emergence of a "new AI" (Brooks, 1999).

Understanding natural cognition should contribute to a general *Theory of Cognition* (or Theories ...), which in turn should feed into a Discipline of Engineering Artificial Cognitive Systems.The big question requires a methodical and conceptual framework within which answers can be situated and formulated; it should give operational (rather than denotational) meaning to the very terminology that ought to underly a scientific discourse on cognition, permitting a shared understanding of key concepts (e.g., information, knowledge, organisation, organism, body, machine, environment, and many others, not to mention cognition itself). A sort of *physics of cognition* would be needed, with (ideally) a set of precisely defined basic and derived notions, and which allows for falsifiable hypotheses and replicable experiments and observation.

Among the aspects, characteristics and elements of natural cognition one may wish to understand better, are the following[3]:

- The emergence, evolution and development of cognitive capabilities in living organisms; the way they relate to the needs and activities of living entities (one of the most challenging research problems may indeed be the *How* and *Why* of the emergence of human intelligence and creativity);
- the physical structures and functions underlying cognitive capabilities and processes (e.g., "*What is the role of the physical substrate?*", "*What are minimal requirements on the physical substrate?*", "*What is the role of embodiment?*");
- the mechanisms of recognising objects, actions and situations, and of generating and adapting behaviour within non-deterministic environments;
- the types and levels of internal representations (of external phenomena in the space and time of the relevant environment) and the ways they come about, change and interrelate through cognitive processes;
- the role and instantiations of memory and learning in cognitive systems (e.g., learning through system-environment interaction / communication, or based on schemas originating from evolution, culture or previous individual experience);
- goal-setting mechanisms and the development of strategies for achieving goals (sub-goal discovery and purpose-driven learning, e.g., learning preda-

[3] Some of these bullet points reflect the author's interpretation of discussions and exchanges he had while preparing the CSR funding agendas. The author is particularly grateful to Aaron Sloman for pointing out several of the issues mentioned in this note.

tion / predator-avoidance; how can goals (desires, preferences, values, moods, intentions, etc.) be identified by (or to) a natural cognitive system?);

– the nature and role of emotion and affect;

– self-awareness, consciousness (for example in the sense of anticipation / simulation of bodily activity), intentionality and *Theory of Mind*, and how these relate to higher-level human cognition;

– the role of language in cognition and of cognition in language.

Based on this understanding an engineer of artificial cognitive systems might wish to address questions such as the following:

– Which artificial cognitive systems need what form of embodiment and why?

– Can (some of the) cognitive toil inherent in perception-action cycles, be offloaded onto physical processes that are peculiar to *the shape of things*, the material things are made of, and the way they are put together?

– Which sorts of memory (mechanisms) are required, and what are the modes and mechanisms of learning needed in an artificial cognitive system?

– What form and degree of autonomy is desirable and achievable?

– To what extent can natural cognitive traits such as affect, consciousness or *Theory of Mind*, be modeled and used in artificial systems? (One may of course turn this around: to what extent can research on artificial versions of cognition shed light on the natural counterparts?)

– For an artificial system to be (or to become) cognitive, does it necessarily have to be self-X (X $\in$ {monitoring, maintaining, modifying, debugging, healing, configuring, controlling, adapting, understanding, aware, generating, organising . . . })?

– Where does design end and (semi-)autonomous evolution, development, self-organisation and learning begin?

– How does all this (representations, concurrent processes, memory, autonomy, self-X . . .) boil down to integrating architectures (*anatomy & physiology*) for artificial cognition?

Questions like these mark out the domain CSR projects are expected to explore. They provide the rationale for the structure of the CSR research agenda depicted below (cf., Figure 1).

This structure takes account of a dictum ascribed to Max Planck: "*Knowledge has to precede application*" (Gruss, 2002). However, it also takes account of the

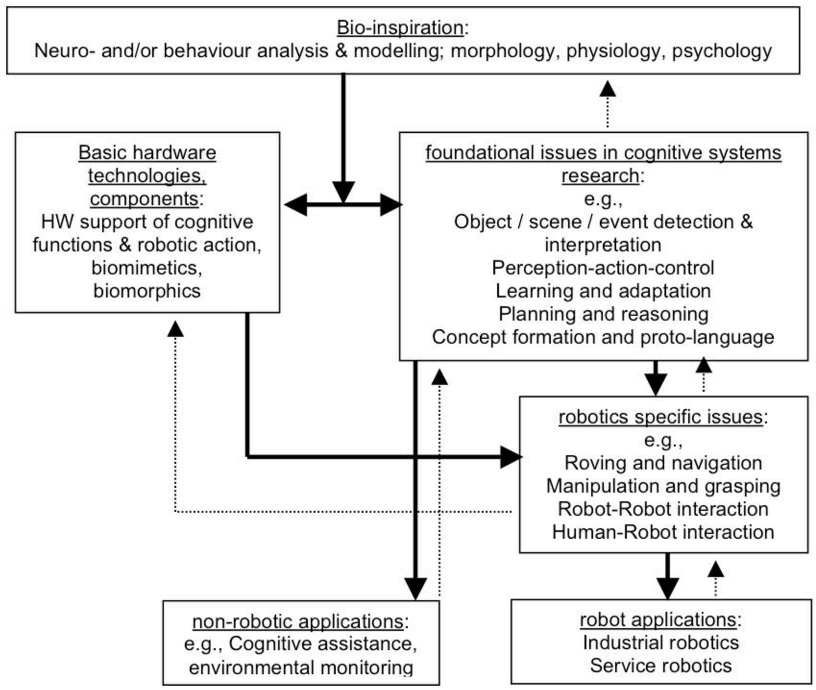

Figure 1: Components of the CSR research agenda

fact that research with a view to supporting engineering must not be confined to an ivory tower. Rather, it should be motivated by and cater to real needs, in line with the strategic goals of (public) European research funding. The two boxes at the bottom of the above diagram represent a wide range of possible scenarios that may provide this motivation, and also guide and validate the research effort proper. In fact, most CSR projects are committed, one way or another, to some concrete application scenario pertaining to one of these boxes.

On the other hand, bio-inspiration in particular, is not a one-way street. Technical experimentation platforms may serve to falsify or back up hypotheses concerning the cognitive capabilities of biological entities; work along these lines may even lead to new clinical approaches to dealing with certain diseases and/or disabilities. While this would be a welcome side-effect if and when it happens in projects, it is not a requirement under the CSR agenda.

3 Structuring the set of CSR projects

The set of CSR projects funded so far by the European Commission can be structured in many ways. The above diagram indicates some. In this note we shall discuss the following facets and cluster projects accordingly.

- environments and platforms
- cognitive competencies
- methods, models and paradigms

Most projects (including those addressing specific robot functionalities) are multi-disciplinary, involving computer science, engineering, neuroscience, behavioural sciences and other pertinent disciplines; in many cases companies with a vested interest in taking up research results also participate..

Apart from RobotCub, ICEA, and SPARK1 all projects mentioned in this section have been or, as of this writing, are being funded under FP7. Listing a project under a specific heading or allocating it to a specific category at best indicates its main thrust. It does not preclude its relevance within other contexts or to another theme. The same holds, *mutatis mutandis*, for projects in the CSR portfolio that are not being cited in this note.

Environments and platforms

Robots, as explained above, are suitable development and experimentation platforms, given the environments the CSR programme mainly focuses on (type 1 and 2 and possibly, 4; as described in section 2).

The history of industrialisation could partly be written in terms of the progress achieved in building such time-and-motion saving machines. Refining and improving their mechanics (grippers, artificial hands, exoskeletons and locomotion gear, etc.) and sensorial capacities (including ones that living organisms do not possess) has always been a major concern.

To reduce the extent of human intervention in the operation of these machines, has been another persistent trend. Ultimately, this means more than merely automating the completion of a task according to some preset rules (as for instance so called numerically controlled (NC) machine-tools do). It means that machines are able to take decisions autonomously on how to proceed with a given task, should conditions arise that had not been foreseen when that task was defined. A (conceptually) simple example would be a roving robot that is supposed to retrieve some object from a distant place but on its way encounters obstacles which had not been anticipated by its human commander (or programmer). Another scenario would involve a machine-tool on the shop floor that need not be laboriously re-programmed upon changing the objects it has to deal with, but that can learn what it can and should do with the new objects (Kroes, 2011).

The industrially and economically motivated desire to make machines in that sense cognitive so that they can for instance take over manual labour in a reliable way, turns them into rewarding platforms for Cognitive Systems research and indeed justifies the combination of Cognitive Systems and Robotics. Endowed with rich acting and sensing capabilities robots are, for example, suitable for testing theories of cognition that propose the interplay between perceiving and acting to be the very basis for developing cognitive faculties. Much of the work under the above mentioned *new AI* label (or *embodied A*I) is in fact being done on such platforms. It has become a research area in its own right, known as Cognitive Robotics (a robotic creature named "*Cog*" being its - perhaps - most famous early representative, see (Brooks et al, 1999)).

Several large CSR projects are addressing a broad range of general aspects pertaining to this area; a number of smaller ones are focussing on particular issues and on creating specific robotic systems, operating in specific (type 1 or 2) environments, to demonstrate the viability of their approach.

The following is a batch of projects that aim to endow robots with advanced grasping and manipulation skills:

- DEXMART ("Dexterous and autonomous dual-arm/hand robotic manipulation with smart sensory-motor skills: A bridge from natural to artificial cognition", `http://www.dexmart.eu/`)
- FIRST-MM ("Flexible skill acquisition and intuitive robot tasking for mobile manipulation in the real world", `http://www.first-mm.eu/`)
- GeRT ("Generalizing robot manipulation tasks", `http://www.gert-project.eu/`)
- GRASP ("Emergence of cognitive grasping through emulation, Introspection and surprise", `http://www.csc.kth.se/grasp/`)
- HANDLE ("Developmental pathway towards autonomy and dexterity in robot in-hand manipulation", `http://www.handle-project.eu/`)
- STIFF ("Enhancing biomorphic agility through variable stiffness", `http://stiff-project.eu/`)
- THE ("The hand embodied", `http://www.thehandembodied.eu/`)
- TOMSY ("Topology based motion synthesis for dexterous manipulation"`http://www.tomsy.eu`)
- DARWIN ("Dextrous assembler robot working with embodied intelligence", `http://www.darwin-project.eu/`)

– Goal-Leaders ("Goal-directed, adaptive builder robots", http://www.goal-leaders.eu/)

The first eight projects on this list are more directly concerned with improving through cognitive enhancement of the underlying control mechanisms, the precision and versatility of a robot's grasping and manipulation operations, whereas the latter two focus on issues related to cognitive architectures prerequisite to carrying out assembly tasks.

By contrast, project ROBLOG ("Cognitive robot for unloading containers", http://www.roblog.eu/) deals with large scale handling of heavy deformable objects. It has a clear and immediate industrial motivation.

Project ROBOSKIN ("Skin-based technologies and capabilities for safe, autonomous and interactive robots", http://www.roboskin.eu/) develops capabilities based on tactile feedback from a suitably sensitized robotic skin, with a view to making human-robot interaction safer.

Specific environments and pertinent functionalities provide the defining challenges in a number of projects:

Roadscapes and driver assistance:

– DIPLECS ("Dynamic interactive perception-action learning in cognitive systems", http://www.diplecs.eu/)

– RADHAR ("Robotic adaptation to humans adapting to robots", http://www.radhar.eu/)

(Deep) sea and underwater operations:

– Co3AUVs ("Cognitive control for autonomous underwater vehicles", http://robotics.jacobs-university.de/projects/Co3-AUVs/)

– CoCoRo ("Collective cognitive robots", http://cocoro.uni-graz.at/)

– FILOSE ("Robotic fish locomotion and sensing", http://www.filose.eu/)

– NOPTILUS ("Autonomous, self-learning, optimal and complete underwater systems", http://www.noptilus-fp7.eu/)

– SHOAL ("Search and monitoring of harmful contaminants, other pollutants and leaks in vessels in port using a swarm of robotic fish", http://www.roboshoal.com/)

– TRIDENT ("Marine robots and dexterous manipulation for enabling autonomous underwater multipurpose intervention missions", http://www.irs.uji.es/trident/)

The sky and aerial inspection and monitoring:

- AIRobots (“Innovative aerial service robots for remote inspections by contact”, http://www.airobots.eu)
- sFLY (“Swarm of micro flying robots”, http://www.sfly.org/)

Urban environments and services:

- EUROPA (“European robotic pedestrian assistant”, http://europa.informatik.uni-freiburg.de/)
- IURO (“Interactive urban robot”, http://www.iuro-project.eu/)
- V-CHARGE (“V-Charge - Autonomous valet parking and charging for e-mobility”, http://www.v-charge.eu/)

In a way, project COGNITO (“Cognitive workflow capturing and rendering with on-body sensor networks”, http://www.ict-cognito.org/) takes as its defining environment the entire human body as it appears to an outside observer.

The human body is also at the centre of attention of three projects dealing with surgical robotics, albeit not on a micro-scale:

- ACTIVE (“Active constraints technologies for ill-defined or volatile environments”, http://www.active-fp7.eu/)
- I-SUR (“Intelligent surgical robotics”, http://www.isur.eu/)
- ROBOCAST (“Robot and sensors integration as guidance for enhanced computer assisted surgery and therapy”, http://www.robocast.eu/)

The industrial orientation of robotics research becomes most prominent in projects BRICS (“Best practice in Robotics”, http://www.best-of-robotics.org) and ECHORD (“European clearing house for open robotics development”, http://www.echord.info/). Both pursue the goal of forging strong links between academic research and industrial development. The latter is an umbrella grant through which many small-scale industry-academia collaborations are funded.

While many of the above projects make use of commercially available mobile or stationary robotic platforms for experimentation and development, others include the design and construction of speciific robotic devices, often on biomimetic principles. Most noteworthy in this regard is FP6 project RobotCub (“An Open Framework for Research in Embodied Cognition”, http://www.robotcub.org/, 9/2004-2/2010). One of the objectives of this project was the provision of a humanoid robot of the approximate shape and size of a three year old child. The resulting machine, known as the iCub (http://www.icub.org/), has 53 degrees of freedom and is operated under a variant of YARP (“Yet Another Robotic

Platform", `http://eris.liralab.it/yarp/`). Both, hardware and software are open-source and as of this writing, used by more than 20 research labs worldwide, including at least ten sites of current CSR projects.

Cognitive competencies

It seems reasonable to assume some sort of layering of the totality of cognitive functions, where functions on a lower layer are prerequisite to (implementing) those above. But what is low, and what is high? A proposal made some 2350 years ago by Aristotle, the towering figure of ancient Greek philosophy, still seems to be popular and widely accepted. He put it in terms of faculties of the soul, broadly classified as *nutrition* (peculiar to plants), *movement* (peculiar to animals), and *reason* (peculiar to humans) (Aristotle, 350 BC). It can be understood as an early attempt to describe what we now call a cognitive architecture.

In modern terms, Aristotle's layers can be interpreted as pertaining to *materials and energy supply*, *sensory-motor co-ordination*, and *reasoning, planning and communication* functions, respectively (Brooks, 1986). Consequently, we may identify a hierarchy of lower layers that comprise competencies covering the basic needs of a living entity, and a hierarchy of upper layers of functions that enables the entity to use the lower layer functions more effectively in order to ensure its survival, for instance by triggering anticipatory action. For artificial embodied systems these hierarchies of layers could be succinctly described in terms of sets of functions that

1. maintain the artificial system in its environment (low-level, sub-symbolic);
2. make the artificial system act deliberately and purposefully, and explicitly understand its world (high-level, symbolic).

The first set covers materials and energy supply and includes mechanisms of autonomic control of *service-sustaining* processes, important features in light of requirements such as robustness and independence of external control, for instance in hostile environments (note that in the artificial domain *service* substitutes *life*). Here, capabilities for autonomic exteroception, proprioception and sensory-motor co-ordination are crucial not only for locomotion and obstacle avoidance in 3D environments but also for setting and reaching survival-related goals. Whether such autonomic mechanisms are within the cognitive domain is of course a matter of debate. It is justified at least in so far as we know of no natural cognitive system (in any agreed sense of the term) where such mechanisms are not in place. In addition one can argue that they implicitly represent effective interpretations (albeit not actively and willingly obtained) of an environment (the outer world) and their host system's situation in it.

The second set comprises functions considered by many to be the essence of cognition. It starts with (elementary) cognitive capabilities to establish and recognise patterns in sensor-generated data; they are prerequisites for higher level operations like conceptualisation, scene interpretation, reasoning, planning, intelligent control, complex goal-oriented behaviour, and communication.

CSR projects can be classified according to their focus on one of these two *thick* layers of cognitive competencies, or (some of) their respective sublayers.

The SPARK projects ("Spatial-temporal patterns for action-oriented perception in roving robots", `http://www.spark.diees.unict.it/` and `http://www.spark2.diees.unict.it/`) as well as their successor EMICAB ("Embodied Motion Intelligence for Cognitive, Autonomous Robots", `http://www.emicab.eu/`) exemplify the implementation in analog circuitry, of low-level competencies that do not normally hinge on formal and symbolic reasoning. A collaboration of biologists, cyberneticists, physicists and engineers, they built seeing, hearing and moving artefacts (multi-legged robots), linking (auditory and visual) perception and action *at the lowest level*, based on principles gleaned from comparatively *primitive* animals, such as crickets and other arthropoda (Frasca et al, 2004).

A fair number of projects deal with sensory-motor coordination in complex multi-degree of freedom machines. ECCEROBOT ("Embodied cognition in a compliantly engineered robot", `http://eccerobot.org/`) for instance, tackles the problem of controlling a robot whose compliant structure mimics the anatomy of the human musculo-skeletal system. AMARSi ("Adaptive Modular Architecture for Rich Motor Skills", `http://www.amarsi-project.eu/`) sets out to develop *low-level* control architectures, based on *Reservoir Computing* (Schrauwen et al, 2007), inter alia for a compliant version of the iCub robot.

Material and energy supply (*nutrition*) issues are less frequently addressed under the general theme of artificial cognitive systems. ICEA, an FP6 project, (`http://www.aslab.upm.es/public/projects/ICEA/`) is the only one in the CSR portfolio which considers the energy dimension of autonomy and the need for self-sustenance as potential prerequisites for and drivers of cognition-based behaviours like foraging and self-defence and ultimately, high-level processes like planning and reasoning. It comprises experimentation with robots (called ecobots by their designers, see (Ieropoulos et al, 2004)) that are equipped with microbial fuel cells to "digest"'food (e.g., sugar, insects or fruits) and generate electricity.

Like ICEA, the project SF (Synthetic Forager, `http://specs.upf.edu/sf/`) undertakes to study the behaviour and neurophysiology of rodents. SF does so with a view "*to identify the neuronal, cognitive and behavioural principles underlying optimal foraging in rodents and to implement these principles in a real-world foraging artefact*" (SF project, n.d.); principles belonging to the *low-level* layer, identified above.

Several of the projects listed in the preceding section may also with good justification be subsumed under the heading *low-level* competencies, for instance those that deal with locomotion (like FILOSE) and/or biomorphic design (like DEXMART, STIFF, and THE).

The *high-level* functions take us into the realm of human-likeness. Indeed, a human brain/body system can actively make its *inner world* explicit to itself (e.g., for acting through planning and deliberate reasoning) and to fellow systems (for joint planning, reasoning, and acting; for deliberate communication, etc.) in an intersubjective consensual domain, spanned by signs and symbols and, in particular, language. A variety of *mental mechanisms*, collectively known as *consciousness* and supported by the brain's anatomy and physiology, are likely to make these feats possible.

How can artificial systems (compact machines or distributed systems, a robot or a network of sensors and actuators) be made to behave - in this sense - human-like? How should a machine be designed so that it can tell us in our terms and at to-be-determined levels of detail, what is happening around it, to it, or inside it? What it is doing and why? How can we make a machine understand what we want it to do? How can a system improve its performance (in rendering a service, for instance) by observing what it is doing and how; by taking account of the effects of its actions, and then drawing the right conclusions, based on some sort of symbolic reasoning? How can it make *its mind* explicit to itself and fellow systems?

Answers to these questions would not only enable a wide range of new services but also make the delivery of many common services more natural and human-like (in traditional terminology: more *user-friendly*), more robust and less dependent on human intervention (i.e., more *autonomous*).

Answers of sorts have been given - and implemented – in the past. However, scene analysis and interpretation, speech recognition and understanding, object classification and understanding affordances, in real-world contexts and under real-time constraints, do pose problems that, on a practical level, we can approach only now, with suitable and affordable sensor, memory and processing hardware at hand. Powerful hardware apart, solving these problems still requires - as in the past - deep research into the principles and methods of producing symbolic descriptions of multimodal real-world scenarios, and - conversely - of understanding their symbolic content.

Understanding a complex dynamic environment with a view to acting in and/or communicating about it is a key issue in many CSR projects. We list but a few, grouped in clusters that can be labeled *scene analysis and interpretation*, *man-machine co-operation and communication*, and *language acquisition* respectively.

Scene analysis and interpretation

- GARNICS ("Gardening with a cognitive system", `http://www.garnics.eu/`)
- IntellAct ("Intelligent observation and execution of actions", `http://www.intellact.eu/`)
- SCOVIS ("Self-configurable cognitive video supervision", `http://www.scovis.eu/`)
- SEARISE ("Smart eyes: attending and recognizing instances of salient events", `http://www.searise.eu/`)
- SCANDLE ("Acoustic scene analysis for detecting living entities", `http://scandle.eu/`)
- TACO ("Three-dimensional adaptive camera with object detection and foveation", `http://www.taco-project.eu/`)
- VANAHEIM ("Video/audio networked surveillance system enhancement through human-centered adaptive monitoring", `http://www.vanaheim-project.eu/`)

Man-machine co-operation and communication

- ALIZ-E ("Adaptive strategies for sustainable long-term social interaction", `http://www.aliz-e.org/`)
- CHRIS ("Cooperative human robot interaction systems", `http://www.chrisfp7.eu/`),
- CORBYS ("Cognitive control framework for robotic systems", `http://corbys.eu/`)
- EFAA {"Experimental functional android assistant", `http://efaa.upf.edu/`)
- HUMAVIPS ("Humanoids with auditory and visual abilities in populated spaces", `http://humavips.inrialpes.fr/`)
- HUMANOBS ("Humanoids that learn socio-communicative skills by observation", `http://www.humanobs.org/`)
- JAMES ("Joint action for multimodal embodied social systems", `http://james-project.eu`)
- NIFTi ("Natural human-robot cooperation in dynamic environments", `http://www.nifti.eu/`)

– ROSETTA ("Robot control for skilled execution of tasks in natural interaction with humans; based on Autonomy, cumulative knowledge and learning", `http://www.fp7rosetta.org/`)

– SPACEBOOK ("Spatial & personal adaptive communication environment: behaviors & objects & operations & knowledge", `http://www.spacebook-project.eu/`)

Language acquisition

– ALEAR ("Artificial language evolution on autonomous robots", `http://www.alear.eu/`)

– ITALK ("Integration and transfer of action and language knowledge in robots", `http://www.italkproject.org/`)

– ROSSI ("Emergence of communication in robots through sensorimotor and social interaction", `http://www.rossiproject.net`)

Projects CogX ("Cognitive systems that self-understand and self-extend", `http://cogx.eu/`) and XPERIENCE ("Robots bootstrapped through learning from experience", `http://www.xperience.org/`) are examples of tackling the more general problem of creating systems capable of adapting their behaviour to changing or entirely new features of their environment. CogX does so by endowing the system with the ability of modelling its own knowledge of its world, including gaps and uncertainty. While CogX does not draw (at least not directly and explicitly) on insights gained from studying natural cognitive systems, XPERIENCE does so extensively. It takes an approach called structural bootstrapping which generalises models of child language acquisition to "understanding actions and dynamic situations" (Steedman, 2002).

Methods, models and paradigms

Learning in all modes and variants plays an eminent role in virtually every project, whether focused on low-level or on high-level cognitive competencies, or on specific robot skills in whatever environment. This is certainly no accident as learning has long since been "*perceived as the gateway to understanding the problem of intelligence*" (Poggio & Shelton, 2000). And the New York Times quotes the co-founder of Microsoft as saying that "*If you invent a breakthrough in artificial intelligence, so machines can learn, that is worth 10 Microsofts.*" (New York Times, 2004)

PinView ("Personal Information Navigator Adapting Through Viewing", `http://www.pinview.eu/`), ComPlacs ("Composing Learning for Artificial Cognitive Systems", `http://www.complacs.org/`) and MASH ("Massive Sets of Heuristics for Machine Learning", http://mash-project.eu/) are projects that explicitly demonstrate the power of *classical* machine learning in cognitive systems research.

The CSR programme also acknowledges the relevance of machine learning in the broadest sense and the need to advance the mathematical underpinnings of engineering cognitive systems that machine learning provides. It therefore gives additional financial support to a large community of European researchers working in pertinent fields, through the umbrella grant PASCAL ("Pattern analysis, statistical modelling and computational learning", http://www.pascal-network.org/).

The basic structure of many projects that focus on general cognitive competencies (rather than on specific robot skills) follows the familiar analysis - modelling - implementation triad where these components correspond almost one-to-one to the disciplines represented in project consortia:

– analysis of natural cognition ($>$ neuroscience, biology, psychology, ...),
– abstract modelling of cognitive processes and architectures ($>$ mathematics, "algorithmics"),
– implementation ("synthesis") of cognitive machines or of cognitive processes in machines and other systems, based on abstract models ($>$ engineering - hardware and software).

However, the CSR programme is agnostic as far as paradigms and modelling are concerned. It leaves room for all schools of thought: cognitivist, connectionist, enactivist, dynamicist, or any variant or hybrid thereof. Inviting competition, it is reminiscent of the late Chinese leader Deng Xiaoping's famous dictum: "*I don't care if it is a white cat or a black cat. It is a good cat as long as it catches mice.*" (Unsourced quote)

Several projects take an *action precedes perception* approach to constructing cognitive architectures and thus either implicitly or explicitly cater to an enactivist worldview. This holds for instance for eSMCs ("Extending sensorimotor contingencies to cognition", http://esmcs.eu/) which, building on the concept of sensorimotor contingencies (O'Regan & Noe, 2001)), attempts to bridge the gap between *low-level* and *high-level* cognitive competencies (Maye & Engel, 2011). Another current project that stands out in this regard is NeuralDynamics ("A neuro-dynamic framework for cognitive robotics: scene representations, behavioural sequences, and learning", http://cordis.europa.eu/projects/97477_en.html) which applies the conceptual apparatus of Dynamic Field Theory to endowing robots with *high-level* faculties (Reimann et al, 2011).

Apart from creating the iCub, it has been one of the main goals of the aforementioned RobotCub project to conceive and implement a cognitive architecture, informed by insights into the *phylogeny* and *ontogeny* of living entities.

It lends itself to emulating mental development as it occurs in infants and toddlers. To this end RobotCub participants have made substantial contributions some of which are documented and referenced in (Vernon et al, 2010).

Lastly, it is worth mentioning that a number of projects listed under *Platforms and environment*, aim to set up multi-robot systems. sFly and, apart from FILOSE, all underwater systems, belong to this category. They represent different multi-robot paradigms. CoCoRo comes closest to emulating a self-organising swarm of underwater devices (Schickl et al, 2011) whereas the other projects largely deal with robot-to-robot communication under adverse conditions. Asynchronous robot-to-robot communication is paramount in project RoboEarth ("Robots sharing a knowledge base for world modelling and learning of actions", `http://www.roboearth.org/`). Here the idea is to create a distributed database (*cloud*), accessible via the World Wide Web, of pertinent *robot knowledge* to be provided and used by robots.

4 Perspectives

It is probably no understatement to claim that CSR research, in spite of its impressive record to date, is still in its beginnings. Its long term goal is indeed elusive. It still has a long way to go from the syntactic systems of early AI, with their externally defined semantics, to enabling systems with intrinsic semantics, grounded in evolution, growth and action; from machines we have to understand, to machines that understand (at least parts of) the world (including us); to a technical solution, one might say, of the age-old mind-body problem (Taylor, 2009). That understanding would become manifest through the system's ability to *do the right things* of its own accord (i.e., act autonomously, yet in compliance with human ontologies). This may include deliberation and communication on human terms, as well as graceful adaptation to novel situations in its environment.

The ongoing increase in processor speed and memory capacities, self-modifying and biomorphic hardware (e.g., (Boahen, 2005)), advances in *intelligent materials*, and much more may be required; but these are most certainly no panacea. Above all, "*Understanding Understanding*" (von Foerster, 2003) provides the key to unlocking the potential of these technologies. Understanding for instance, how we understand the world, and how animals understand the world, in terms of what we and they can and should do in it. (It is useful to remember the remarkable feats that bees, birds and beavers, to name but a few, are capable of (Wynne, 2001)). Arriving at this understanding and making it actionable will likely take more time than might be optimistically expected.

The question of embodiment also appears far from settled (Ziemke, 2003; Sloman, 2008). More research along the lines of the above mentioned ICEA project may be called for. One of the hypotheses this project set out to corroborate is that "*the emotional and bioregulatory mechanisms that come with the organismic embodiment of living cognitive systems also play a crucial role in the constitution of their high-level cognitive processes*" (Ziemke, 2008). More generally, one may ask *to what extent can function be divorced from structure?* or: *does mind-matter matter to what a mind can think (and do)?*. (Amusingly, the 19th century philosopher Friedrich Nietzsche maintained that "*The body is a big sagacity, a plurality with one sense, a war and a peace, a flock and a shepherd. An instrument of thy body is also thy little sagacity, my brother, which thou callest 'spirit' - a little instrument and plaything of thy big sagacity*" (Nietzsche, 1896).)

Cognitivist (or *computationalist*) paradigms (e.g., reducing cognition to symbol manipulation, simply put) have been extensively challenged in the past. But we may still need to find new and hitherto unexplored ways for making the physico-chemical dynamics of suitable material structures *map (or: reflect) an outer world* into themselves. Whether these mappings will be *Turing-realisable* is an open question. (Note that the mappings effectuated by the members of a certain class of artifical neural networks are provably *super-Turing* (Siegelmann, 1999).)

We may have to reinvent life, presumably merging *Artificial Life* and CSR research, with a view to marrying the precision, capacity and reach of digital systems to the robustness, adaptivity and effectiveness of (certain) biological systems (and – of course – our little human ingenuity, greatly amplified).

5 Epilogue

It should be noted that the current European CSR programme includes arguably the largest non-military (or non-military sponsored) robotics programme of all time. But it should also be clear that the CSR agenda and its underlying research questions directly address the challenge of creating computational equivalents of human-like cognitive traits and capabilities in embodied systems. In particular, this holds for projects that focus on high-level competencies, on learning, or on linking low-level to high-level competencies (e.g., through developmental robotics, or extending sensorimotor contingencies), as described in Section 3.

Whether a line of public research funding (as the one presented in this note) will be continued, one way or another, often not only depends on motivations intrinsic to the field at issue, but also, and perhaps more so, on priorities set within a given political and economic context. There is, however, good news. Neelie Kroes, the Commissioner in charge of European ICT (Information and Communication Technologies) funding, in (Kroes, 2011), quotes from Bertolt Brecht's famous play *The Life of Galileo* (Brecht, 1952) the great scientist's adage: "*I maintain that the only goal of science is to alleviate the drudgery* *of human life.*" And she adds: "*Sound advice indeed! We will continue to fund research whose results help create better living conditions for everyone on this planet and research that helps us to better understand ourselves and the world we live in. Both go hand in hand—and robots should take their fair share in this ICT landscape.*"

Acknowledgement

These notes draw on a number of workshops I had the pleasure to (co-)organise while preparing the Cognitive Systems and Robotics part of the European Commission's IST/ICT programmes, on projects I came in contact with, on many discussions (on various occasions) with workers in pertinent fields and, last but not least, with my colleagues in the Cognition Unit of the European Commission's Directorate General Information Society and Media.

References

Aristotle (350 BC). De Anima (On the Soul). Translated by J. A. Smith, retrieved from `http://classics.mit.edu/Aristotle/soul.html`.

Boahen, K. (2005). Neuromorphic Microchips. Scientific American, vol 292, no 5, pp 56-63

Brecht, B. (1952). Galileo. Trans. Charles Laughton. Ed. Eric Bentley. Works of Bertolt Brecht Ser. New York: Grove Press

Brooks, R. A. (1986). A Robust Layered Control System for a Mobile Robot. IEEE Journal of Robotics and Automation, Vol. 2, No. 1, pp. 14-23.

Brooks, R.A. (1999).Cambrian Intelligence: The Early History of the New AI. Cambridge (Mass.): The MIT Press.

Brooks, R. A., Breazeal C., Marjanovic M., Scassellati B., & Williamson M. (1999). The Cog project: building a humanoid robot. in: Computation for Metaphors, Analogy, and Agents. C. Nehaniv (ed), Lecture Notes in Artificial Intelligence 1562. New York: Springer, pp. 52-87.

EU Commission (n.d.). ICT Work Programmes 2007-2008, 2009-2010, 2011-12. Retrieved from `http://cordis.europa.eu/fp7/ict/newsroom/library_en.html`.

Frasca, M, Arena P. & Fortunato L. (2004). Bio-inspired Emergent Control of Locomotion Systems. Singapur: World Scientific Publishing Co. Plc. Ltd.

Goertzel, B. & Pennachin C. (Eds). (2006). Artificial General Intelligence (Cognitive Technologies). Berlin, Heidelberg: Springer.

Gruss, P. (2002). Scientific Autonomy and Financial Support - Prerequisites for Creative Research. President's address, 53rd General Meeting of the Max Planck Society for the Advancement of Science Halle, Germany. Retrieved from `http://www.mpg.de/727610/020614hv2002Gruss_en.pdf`.

Ieropoulos, I., Melhuish C. & Greenman J. (2004). Energetically Autonomous Robots. Proceedings of the 8th Intelligent Autonomous Systems Conference (IAS-8), Amsterdam, pp. 128-35.

Kroes, N. (2011). Robots and Other Cognitive Systems: Challenges and European Responses. Philosophy & Technology, Volume 24, Number 3, pp 355-357.

Maye, A. & Engel A.K. (2011). A computational model of sensorimotor contingencies for object perception and control of behavior, IEEE International Conference on Robotics and Automation (ICRA 2011), 9-13 May 2011, Shanghai, China, pp 3810-3815. doi: 10.1109/ICRA.2011.5979919.

New York Times. (2004 March 1). Microsoft, Amid Dwindling Interest, Talks Up Computing as a Career. Retrieved from http://www.nytimes.com/2004/03/01/ business/microsoft-amid-dwindling-interest-talks-up-computing -as-a-career. html?pagewanted=print&src=pm.

Nietzsche, F. (1896). Thus Spake Zarathustra (translated by Alexander Tille). New York and London: Macmillan.

O'Regan, K. & Noe A. (2001). A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24, pp 939–1031.

Poggio, T. & Shelton C.R. (2000). Learning in brains and machines. Spatial Vision, 13(2,3), pp 287-296.

Reimann, H., Iossifidis, I. & Schöner G. (2011). Autonomous movement generation for manipulators with multiple simultaneous constraints using the attractor dynamics approach. ICRA 2011. pp 5470-5477.

Schmickl, T., Thenius R., Timmis J., et al. (2011). CoCoRo: The self-aware swarm of underwater robots. IROS 2011 - IEEE/RSJ International Conference on Intelligent Robots and Systems.

Schrauwen, B., Verstraeten D. & Van Campenhout J. (2007). An overview of reservoir computing: theory, applications and implementations. Proceedings of the 15th European Symposium on Artificial Neural Networks. pp. 471-482.

SF project. (n.d.). Retrieved from `http://specs.upf.edu/sf/`

Siegelmann, H. T. (1999). Neural Networks and Analog Computation: Beyond the Turing Limit. Boston: Birkhäuser.

Sloman, A. (2008). Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress. In Sendhoff B. et al (Eds), Creating Brain-like Intelligence. Berlin: Springer.

Steedman, M. (2002). Plans, affordances, and combinatory grammar, Linguistics and Philosophy 25 (5-6), pp 723- 753.

Taylor, J.G. (2009). Mind-body problem: New approaches. Scholarpedia, 5(10) : 1580. Retrieved from `http://www.scholarpedia.org/article/Mind-body_problem:_New_approaches`

Tennekes, H. (1996). The Simple Science of Flight. Cambridge (Mass.): The MIT Press

Vernon, D., von Hofsten C., & Fadiga L. (2010) A Roadmap for Cognitive Development in Humanoid Robots. Berlin, Heidelberg: Springer.

von Foerster, H. (2003).Understanding Understanding: Essays on Cybernetics and Cognition. New York: Springer.

Wiener, N. (1948). Cybernetics or Control and Communication in the Animal and the Machine. New York: Wiley & Sons.

Wynne, C.D.L. (2001). Animal Cognition: The Mental Lives of Animals. Houndmills: Palgrave.

Ziemke, T. (2003). What's that Thing Called Embodiment? In: Proceedings of the 25th Annual Meeting of the Cognitive Science Society. pp. 1305-1310,

Ziemke, T. (2008). Emotion in Cognitive Systems Architectures. In SETN '08 Proceedings of the 5th Hellenic conference on Artificial Intelligence: Theories, Models and Applications. Berlin. Heidelberg: Springer.