

Cognitive Systems: Insights, Examples, Systems — Report from the EUCog 3 x 3 Questionnaire

Nathan Labhart
Artificial Intelligence Lab, University of Zurich
labhart@ifi.uzh.ch
2014-05-29

Table of Contents

Introduction.....	1
Embodiment.....	2
Human-machine/human-computer interaction.....	2
Robots.....	3
Emotions.....	3
Medical, rehabilitation, care.....	4
Ethics and societal issues.....	4
Vision and perception.....	5
Autonomy and navigation.....	5
Planning and control.....	5
Manipulation.....	6
Cognition.....	6
Language.....	7
Bio-inspiration, understanding by building.....	7
Understanding and modeling the human brain, Neural Networks.....	8
Knowledge representation.....	8
Learning.....	8
CSR/EUCognition/Interdisciplinary research.....	9
Issues.....	10

Introduction

Since early November 2013 a questionnaire had been available on the EUCog website, asking the members for three (or more) key insights, characteristic examples, and working systems or system components enabled by Cognitive Systems Research (CSR).

By mid-May 2014 a total of 153 response sets were submitted (EUCogIII has roughly 950 members), resulting in about 360 insights, 330 examples, and 260 systems mentioned.

As a general observation it should be noted that the responses were completely unstructured, they varied from a single word or URL (to a project website or scientific article) to entire paragraphs with multiple references – and often unspecific. Even in the “systems” category there were few concrete “outcomes” mentioned. In addition, the focus was on CSR in general, not so much Europe-specific.

This report provides a thematically clustered overview of the responses, interspersed with “quotable” statements printed in italics.

Embodiment

One of the key insights of CSR is clearly the importance of “embodiment” as a concept to better understand cognitive processes, i.e., the view that cognition (or intelligence¹) emerges from the interaction of brain, body, and environment. Flexibility and intelligence are seen as properties of embodied, situated systems, and embodiment is needed for the development of cognitive abilities as cognition is strongly contingent upon our bodily immersion in an environment to which we are coupled through sensory and motor processes. In short: Embodiment and cognition are inseparable – the body shapes the mind and the mind shapes the body.

From this follows that complex behavior does not necessarily need complex internal mechanisms but can be the result from “simple” interactions with the environment. For example, many higher cognitive abilities rely on basic sensorimotor capacities. Embodied systems are never purely reactive, as sensory input may stimulate motor output (a “proactive” change of the environment), which in turn leads to new sensory input, etc.

The deep questions that arose in the computational form of the Turing Test appear again with embodied systems: when does an imitation become the real McCoy?²

On a more practical side, the body of an agent itself can compute complex tasks (“morphological computation”), i.e., the computational expense can be offloaded into the body and the environment.

A fundamental aspect of CSR is to provide a deeper understanding of the interaction of brain, body, and environment as a whole.

Cognition cannot be understood without at the same time understanding how cognition is linked to the sensory and motor surfaces and how cognition takes place while an organism/agent is situated in a structured environment to which it is adapted.³

Human-machine/human-computer interaction

An area where CSR has an immediate impact is human-machine/human-computer interaction, i.e., how people relate to machines and how this informs the design of these machines. Example applications include assistive robots for caregiving (see separate section below), artefacts that support communication or cooperation, devices to empower the cognitive capacities of human-machine teams, and the design of consumer electronics informed by psychological models.

In natural human-human communication, a multitude of verbal and non-verbal behaviour is used, e.g. speech, communicative gestures, nods, gazes etc. The verbal and the non-verbal acts together comprise the information necessary for a learner to understand. If robots are to interact with humans in natural ways in the future, mechanisms accounting for the multi-modal complexity of oral human communication need to be developed.⁴

CSR is needed for robots to interact with humans, especially in unstructured, dynamically changing

1 The two concepts are sometimes used interchangeably, which – depending on their definition – may be justified; it is one of the issues of CSR that there are no clear definitions of what “intelligence” or “cognition” exactly mean (see section “Issues”)!

2 NN

3 Gregor Schöner

4 Stephanie Schreitter

environments which are usually not a problem for humans but still pose formidable problems to machines. Semiautonomous operation, active or proactive perception strategies are ways to deal with such environments, which exist for example in rescue scenarios, as addressed e.g. by the projects NIFTi, TRADR, and WALK-MAN. Robot navigation and interaction in human environments is also addressed by STRANDS.

Verbal and behavioral communication plays a major role in human-machine interaction too, where projects like HUMAVIPS, ALIZ-E, and JAMES address the issue of learning from human behavior and recognizing/understanding/interacting with people. Issues include recognizing the dynamics of social interaction, predicting humans' actions, evaluating and generating communicative gestures – in short, finding the conditions for appropriate cooperation between humans and robots.

A service robot in human habitable environment should be able to respond efficiently to human actions and emotion state. The robot should have the ability to express its state in an appropriate for a human way and to find the best action in respect to its task and the observed human's gestures, emoting state etc.⁵

Further examples of human-machine interactions that benefit from CSR include table-tennis robots that predict the moves of their human opponent, robotic tour guides, and light-weight assembly line robots that are capable of safely operating next to human workers. Safety can be increased by using compliant or soft bodies (Probo, COMAN).

Robots

One of the most visible outcomes of CSR are robots and mechatronic systems, which come in all sizes and shapes and address various issues. Integrating cognition, perception and action in robots brings them closer and closer to human-like machines; among humanoid robots, the iCub is an excellent example of a usable and useful system that could not have been developed without the interplay of various fields in European CSR. As an open platform deployed in many different projects, it features open-source software repositories also useful for other platforms. Another popular robotic system is the commercial “Nao” humanoid by Aldebaran.

iCub, which at least is a reasonably common platform in Europe (which can be counted as a key result)⁶

Aldebaran Nao: the best selling humanoid relies heavily on cogsys research for its autonomy⁷

Other humanoid robots stemming from European CSR include the ARMAR family of robots, ECCE, Kaspar (which helps autistic children), the open-source Poppy platform, and Roboy. Non-EU humanoids include Actroid, ASIMO, Cog, PR2, Leonardo, and Zeno. Animal-like robots include Paro, Pleo, Big Dog, and Wildcat; KUKA is a good example of an industrial robot.

Emotions

The perception, understanding, and expression of emotions is still a hard problem. CSR has led to

5 Kiril Kiryazov

6 NN

7 Tony Belpaeme/Sue O'Donnell

systems that analyze facial expressions and recognize emotions from these, and ventured into giving an emotional behavior to robots. Applications for emotional monitoring include marketing (automatically measuring consumer response) and human-machine interaction such as caregiving robotic companions.

Medical, rehabilitation, care

CSR has shown that cognitively different people (for example, patients with ASD) can perform quite well in society when helped by appropriately designed cognitive systems. Mathematical or computational models of the brain and its cognitive processes combined with psychological treatment and robotics approaches led to several programs that use humanoid or animal-like robots to treat autism in children by helping them communicate with other people and learn certain social skills. Developments in clinical cognitive science are essential for new psychological treatments.

Database support systems, knowledge representation models, and expert systems may supplement human expertise and play a significant role in decision making, especially in the medical domain (where IBM's Watson is being deployed, too). Such systems aid in difficult medical diagnosis problems by analyzing clinical data, (semantic) searching in drug databases and/or medical documents for treatment suggestions, and similarity search in medical image databases. As an example for such frameworks, the DebugIT project allows monitoring infectious diseases by extracting microbial case data from clinical information systems.

After surgery, which in some cases is supported by robotic devices such as the “da Vinci” robot, electrical and mechanical engineering approaches lead to applications in physical rehabilitation, e.g., brain-computer interfaced wheelchairs for severely disabled humans, exoskeletons (such as ATLAS), and sports assistants (e.g, playing table tennis). Control learning strategies are important in human motion rehabilitation, and motor control learning is needed in the neurological rehabilitation of people with stroke, Parkinson's, or cerebral palsy.

In general, robots find applications as companions and caregivers by assisting patients who require long-term or even permanent care, such as elder or physically disabled people. Of course, as in other human-machine interaction scenarios, safety needs to be ensured, and ethical issues may arise.

Ethics and societal issues

A European framework for an ethical agenda in robotics is needed, as “robot ethics” is an important issue especially as robotic applications in health care (e.g., Paro robots used in care environments) and assisted living are increasing: Safety, privacy, and security need to be addressed and liability issues may pose a major challenge to the distribution of cognitive systems. Therefore, ethical implications in connection with the use of cognitive systems need to be discussed as early as possible. In a future where humans and robots “live” together, machines should be “socially acceptable”, even if they are only assisting in certain tasks (care for the elderly, rehabilitation scenarios, “intelligent avatars” in human-computer interfaces). In the engineering domain, research on failure detection and handling should get more weight.

Autonomous cars are likely to be part of traffic in the near future; the problems to be solved are now more of legislative and societal nature than technical. The use of autonomous drones in military also raises ethical questions.

Vision and perception

Vision and perception are hard problems. Incorporating social and cultural aspects in visual cognition may improve robotic vision systems, thus, CSR is needed for vision to become more useful in domains such as navigation (driver assistance), surgery, image/object recognition and classification, scene understanding in mobile robotics, face detection, or object tracking.

A concrete example is the HERMES project which employs cognitive computer vision for surveillance systems in factory automation, security, and many other applications.

Perception is a multimodal concept that includes vision. It links to psychology, where for example the phenomena of change blindness and choice blindness may help in the creation of cognitive models of the human mind. By “borrowing ideas” from how human perception is thought to work, i.e., by coming up with computational models that seem to mimic human perception (such as Bayesian reasoning), massive advances have been made in computer vision, for example in recognizing people in pictures.

Applications of CSR in perception are found in human-machine interaction scenarios (facial expression, aggression detection, intention recognition).

Autonomy and navigation

Robotic systems that autonomously navigate in largely unknown, changing environments are gaining importance, a prominent example are “self-driving” vehicles. Some of these systems are not fully autonomous but complement humans in their navigation tasks. Traditionally, one of the biggest problems to be tackled was the real-time integration of large amounts of sensory data. An important insight from CSR is that

Autonomous behavior may be generated based on much simpler interfaces with sensory and motor systems than previously postulated in the sense-plan-act approach to autonomous robotics⁸

In search and rescue scenarios, long-term autonomy is necessary for artificial cognitive systems to “survive” in real-world environments.

Systems that currently employ autonomous navigation skills to a certain extent include vacuum robots, warehouse robots, driver assistance/navigation devices that keep the car on the track and avoid collisions with other cars or pedestrians, and unmanned aerial vehicles.

Planning and control

A great deal of robotics research is concerned with control mechanisms and with methods of how to optimally plan and execute the robots’ motions depending on their tasks and environments. CSR led to biomimetic approaches to control, which will be of great importance for example in bipedal locomotion: understanding the role of (human) body morphology aids in the robot’s control optimization of whole-body motion. As many mechatronic systems have physical (technical) constraints, accuracy and speed can be exploited to the fullest by learning optimal control parameters.

8 Gregor Schöner

Cognitive systems can efficiently control complex systems⁹

Novel robot architectures are based on the concept of sensorimotor contingencies (eSMCs), which – unlike traditional “perception-thinking-action” control strategies – take into account the relations between actions and associated changes in sensory input.

Manipulation

Traditionally, (industrial) robots were used for precisely controlled manipulation tasks. Advances in CSR have led to more powerful, non-trivial grasping and manipulation capabilities, i.e., the handling of “soft” materials such as fabrics, textiles, and garments (with two arms, CloPeMa project) and paper. Grasping consists of sub-problems such as affordance, intention, hand-eye coordination, and compliance, which can be helped by improved vision systems (as used in PACO-PLUS). Another important manipulation task, especially in production lines, is assembly (addressed e.g. by DARWIN). Depending on the operating environment, anthropomorphic solutions to movement generation are more desirable than the traditional approach of programming exactly every move without any “perception” of the environment. This perception may be aided by specialized computer vision for grasping.

Cognition

While many important functions of a cognitive system are not yet completely understood, and the distinction between perception and cognition as separate psychological processes is sometimes considered to be arbitrary, CSR has provided a number of important insights, for example that cognition and development are closely related (cognition emerges during development, taking sensorimotor coordination into account), that both top-down and bottom-up approaches in cognitive modeling are necessary for efficient implementations of these models, and that

Cognition is social: cognition of the human species is unique in being rooted in social interaction, which sets us apart from other species and resulted in uniquely human cognitive capacities, such as language¹⁰

Many cognitive abilities cannot be fully understood in isolation, but rather require context of an actual complete, working system¹¹

Better understanding the mechanism of cognition is of crucial importance in order to build any kind of “cyborg” or HRI systems¹²

In turn, (evolutionary) robotics is suggested as a new tool to study cognition (the “understanding by building” approach). CSR has implications in many different areas, for example to understand human cognition and build better machines:

[...] cognition in general can (and indeed should) be explored in a principled manner, enabling results to be related and applied across domains¹³

Within CSR a number of cognitive architectures have been developed (e.g., ACT-R, CERA-CRANIUM, CLARION, CRAM, ISAC, LIDA, etc.) with various applications, such as

9 Dimitris Iakovidis

10 Tony Belpaeme/Sue O'Donnell

11 Ron Chrisley

12 Kiril Kiryazov

13 Paul Baxter

implementing controllers for autonomous robots by integrating different cognitive components, or simulating tasks in cognitive and social psychology. As such, cognitive systems can enhance human cognitive capabilities.

Language

Understanding the neural basis of language or the acquisition thereof can lead to better natural language processing systems, for example machine translation or speech recognition and generation. CSR has contributed to neurophonetic models for speech processing, and “embodied” approaches provided insights into the relation of fine motor control to language acquisition.

Cognitive functions previously assumed to be innate (for example language) can be learned from embodied experience¹⁴

CSR is also concerned with how speech learning, production, and perception are linked. Further research is needed to transform sensor data into symbolic representation in order to “ground” language (“Symbol Grounding Problem”). This may also lead to improved machine translation algorithms.

The origins of speech sounds shared by individuals of linguistic communities can be understood as a self-organized macro-structure resulting from local vocal imitation among peers. This new view of the origins and evolution of speech has been made possible thanks to a whole series of computational models.¹⁵

Apple’s Siri is a prominent application of CSR, as it comprises of speech recognition, context awareness, information collection, question answering, and “artificial consciousness”.

Bio-inspiration, understanding by building

Biological systems provide rich inspiration for the development of artificial cognitive systems such as robots, and (technically) re-engineering or modeling natural mechanisms and processes can identify important features of biological systems that may not have been obvious or thought possible previously. This constructive approach is also known as “understanding by building”. A central insight from CSR is that

To some extent, biological and artificial cognitive systems can be conceptualized based on similar principles¹⁶

Biologically inspired solutions can simplify the mechanical and control design of robotic applications, for example by building soft robots that exploit their material properties in grasping tasks, or by using “morphological computation” in passive walking instead of relying on a central controller. This has led to

biomimetic robotic systems that mimic animals such as bats, dolphins, snakes, birds, ants etc., artificial models of biological systems or biological organs, [and] software/hardware related to the above two items¹⁷

14 Anthony Morse

15 Pierre-Yves Oudeyer

16 Billur Barshan

17 Billur Barshan

Understanding and modeling the human brain, Neural Networks

Partly following the “understanding by building” approach, the development of artificial neural networks is, on the one hand, implementing discoveries from human brain research, and on the other, contributing to new models of the human brain. Apart from practical applications such as pattern recognition systems, mathematical or computational models of the brain contribute to the understanding of cognitive processes, e.g., associative memory (using spiking neuron models), episodic memory, cortical-hippocampal relations, or the emergence of autobiographical memories.

In neuroscience the discovery of mirror neurons led to a new understanding of imitation learning, empathy and emotionality. This area of research is helped by brain scanning and near-real time decoding of images (fMRI pattern recognition), of which one outcome is an accurate mapping of the human brain. The interdisciplinarity of CSR (combining e.g. neuroscience, psychology, and electrical engineering) leads to brain implants for the visually impaired (see also “Medical, rehabilitation, care” above).

Knowledge representation

CSR is shedding light on the acquisition, management, representation and exploitation of knowledge of the (human) mind and is applying the findings in various fields. A particular insight is on

Modeling the knowledge an artificial system possesses or may possess in [the] future¹⁸

Cognitive Systems research led to integration of quantitative low-level and qualitative high-level knowledge representation [and] showed that robot knowledge can be learned and need not be manually designed¹⁹

Search engine technology, particularly more recent iterations, that rely on semantic network-like functionality, or the more human-like representation of information on a massive scale to facilitate more accurate search results²⁰

Examples include WordNet, Deep Blue, IBM’s Watson (with applications e.g. in medical database support systems), search engines like Google and KartOO, and recommender systems such as the one used by Amazon.

Learning

All forms of learning help to improve cognitive capabilities and to obtain flexibility and adaptability in real-world scenarios. Several learning methods exist (active learning, reward learning, learning by demonstration, etc.), many of which are based on natural systems. Since (natural) learning is a multimodal process, the interdisciplinarity of CSR approaches can explain how learning actually occurs and at the same time employ or adapt learning methods for applications. Insights include the role of prediction and simulation in learning and cognition, the fact that neuro-structural changes are not necessary for Hebbian learning, and the importance of mechanisms for autonomous learning for improving the behavior of artificial agents that are integrated in human environments.

18 Giorgio Metta

19 Bernd Neumann

20 Paul Baxter

Simultaneous Localisation and Mapping (SLAM) is a hugely successful machine learning approach²¹

CSR has contributed to teaching robots by reinforcement and machine learning (learning controllers from scratch) and imitation learning of motor skills (e.g., CogX, ERC project EXPLORERS). Understanding curiosity-driven learning made robots possible that develop novel skills during their entire “life”, and that can self-calibrate/adapt and recover after failures.

Further applications of learning methods include the adaptive powering of antennae (e.g., mobile phone masts) and handling huge amounts of data (e.g., search machines with spell checkers and translation engines).

CSR/EUCognition/Interdisciplinary research

Many problems that are easy for humans to solve turn out extremely hard for machines, e.g., vision, social interaction, language processing, walking, and prediction. Advances in CSR are necessary to tackle these problems and can lead to the discovery of distinct “categories” of intelligence, such as swarm intelligence or morphological intelligence, which in turn may inspire new technologies. Even though the path is long and complex, increasingly sophisticated solutions for specific application areas are being developed, for example expert systems like IBM’s Watson.

Pretty much any problem could be improved by more advanced cognitive systems. These could help us to understand the problem, formulate a solution or actually physically solve it.²²

Combining traditional approaches with modern CSR can be a successful path:

Robots acting in everyday situations - borrows ideas from classical AI about how to plan and represent knowledge, and uses more modern ideas to control the sensorimotor robotics side²³

Holistic approaches for artificial systems are necessary for them to function in the real world, as too specialized or narrow solutions will fail in uncontrolled environments. The EUCog network enables such an interdisciplinary culture between humanities and technological disciplines and brings together people from theory and practice to shape the future of the field. CSR is a combination of interdisciplinary advances, where empirical studies take place alongside engineering, and provides critical analysis of fundamental aspects of cognitive systems and robotics.

Most current electronic technologies originated in their embryonic forms as AI research projects, or involve key components that originated in that way.²⁴

For example, in the field of computer vision, face detection/recognition or activity recognition (e.g. Kinect), ... all originated from decades of cognitive vision research²⁵

In general, there are a multitude of systems based on cognitive systems research developed in working research labs (both academic and more industry-focussed) that

21 Alan Winfield

22 David Gamez

23 Frank Guerin

24 Steve Torrance

25 Tony Belpaeme/Sue O'Donnell

have been applied to (admittedly limited) domains - the media attention that these systems gain, and the potential influence this has on the public consciousness should not, in my view, be ignored²⁶

Concrete “products” of CSR include Aldebaran’s Nao robot, first implementations of robots that learn by demonstration and can then work alongside human workers, real-time handwriting recognition, and systems that understand natural language (a field whose problems haven’t been completely solved, but CSR is essential for its advancement). Especially in human-machine interaction scenarios (in education, rescuing, companionship, etc.) insights from CSR need to be employed to achieve higher levels of functionality, considering multi-modal sensory feedback (not just speech, but also gestures, facial expressions, etc.).

There are also a number of open-source software and hardware developments originating from CSR, such as the iCub humanoid robot, the RocoComp robotic framework, the EmbedIT educational toolkit, SpikeStream (neural simulator), and iSpike (interface between the iCub and spiking neural simulators).

Issues

Needless to say, CSR cannot solve every problem instantly and thoroughly – there is much about human cognition and action that probably cannot be explained through AI and related sciences and technologies –, and it has its share of issues that come in various forms. For example, there is no clear definition of terms like “cognition” or “intelligence”, and the term “cognitive systems” is very broad. One reason is that the involved disciplines (psychology, artificial intelligence, computer vision, robotics, etc.) could not agree on what were the important issues at stake.

Some people argue that if we do not know what intelligence or cognition really is, we cannot build robots that truly embody these concepts – and we do not have the methodologies to benchmark systems that are said to be based on CSR. Even without such measuring instruments, it can be seen that dealing with unknown environments and situations is still an unresolved issue in many CSR projects.

There is a large gap between the theorists of cognitive systems that limit their applications to abstract domains and the more practical researchers in embodied cognition that need to make ungrounded abstractions and loose generalizations in order to make these approaches applicable in real robots.²⁷

Due to the many interpretations of "cog. systems" I have a hard time identifying clear questions that we (and nobody else) adress, and therefore solutions are also provided by other fields²⁸

I have been involved in projects to develop a talking teddy bear and a more sophisticated smartphone assistant. Neither is working yet!²⁹

More research is needed in the areas of Symbol Grounding (i.e., transforming sensory data into symbolic representations), adaptive and flexible robot control in industry (already partly solved),

26 Paul Baxter

27 Georgios Pierris

28 Erik Billing

29 David Gamez

natural language processing, and mathematical models of self-organization.

Lots of pieces, but--to my knowledge--no clearly convincing direction as yet!³⁰

Very sadly actually, I can't think of an example where the major effort in solving an issue did not in fact originate from another field....³¹

However:

This is a bit tricky, because the moment a lot of problems are solved (algorithmically), they tend to leave the domain of cognitive systems/AI and become "common" computer science³²

One big problem with many CSR projects is the overly enthusiastic tone in the proposals, which may lead the audience – politicians, companies, the public – to expect a perfect solution, when in effect only a slight improvement of a small part of a much larger theoretical or practical issue is eventually addressed by the project.

Needed: a reduction in the levels of hyperbole and over-expectation that surround cognitive systems/AI research³³

30 NN

31 NN

32 Joachim de Greeff

33 Alan Winfield