DG Information Society **19.3.2004**
Future and Emerging Technologies

FET Proactive initiative 2004: Bio-inspired Intelligent Information Systems (Bio-i3)

POSITION PAPER

Aude Billard, Rodney Douglas, Beatrice de Gelder, Phil Husbands, Giorgio Innocenti, Marc de Kamps, Martin Kanje, Alois Knoll, Ricardo Sanz

INTRODUCTION

The "Extraction of meaning" is a fundamental operation of complex biological systems, geared to the evolutionary imperatives of survival and reproduction in a complex environment. For a given biological system "extraction of meaning" involves processing of sensory information in the context of a hierarchy of control targets or strategies at the top of which are biologically inscribed constants specifying conditions of survival and/or reproduction and at the bottom learned-adaptive variables, consisting of targets and/or strategies consistent with the biologically inscribed constants.

The targets and strategies at the lower levels of the "extraction of meaning" hierarchy can be viewed as perceptual, motor or cognitive "objects" to be separated from a noisy background. In the perceptual domain, the background consists of visual, auditory, tactile, olfactory, etc, stimuli which are not part of the object, in the motor domain it consists of activity in muscles not directly related to the action to be performed. In the cognitive domain the background consists in memories of object or strategies irrelevant to the actual behavioural situation. In all domains object and elements of the background can be interchanged. That is, the biological systems are endowed with flexibility in identifying the lower-level repertoire of learned targets and strategies. This flexibility has, in general, adaptive value. It increases with the increasing complexity of the neural networks involved, in particular those of the cerebral neocortex. Therefore, the study of the complex networks of neurons in the cerebral cortex and elsewhere can provide guidelines for the construction of man-made artefacts aimed at man-defined "extraction of meaning".

The "extraction of meaning" process is scaled-up through the hierarchy of neural systems up to the highest level cognitive functions in the brain, including the understanding of the own self and its relation with the surrounding world.

The networks involved in the "extraction of meaning" are the result of self-organizing processes which combine the implementation of predetermined templates with the selection of structures produced in excess, in both development and evolution. The latter processes endow the biological system with a high degree of flexibility in confronting changing environmental conditions. Further flexibility is provided in the adult brain by a redundancy of control systems including, possibly, awareness of the world and of the self, and by experience-dependent modulation of neuronal connections.

The extraction of technologically useful architectural, algorithmic and implementation principles from the biological models from low level neural structures to high level mental processes lies at the interface between the biological and technological fields and requires tight collaboration between neuroscientists, who show strong commitment to the systemic approach and interest in the relation between neuroscience, cognition and psychology, mathematicians, computer scientists and engineers. The expected benefits of this interaction are the acceleration of progress in the different fields by simultaneous cross-fertilization and the emergence of new interdisciplinary models. This will eventually contribute to the reshaping of the technologically oriented societies by the construction of better devices (user-friendly, performant, dependable etc) and, in medicine, to improved diagnostic criteria, and the definition of preventive or replacement therapies.

BRAIN ARCHITECTURES AS COMPUTATIONAL ARCHITECTURES

Background

The "extraction of meaning" i.e. the flexible structuring of the internal and external stimuli aimed at deriving biologically relevant information from a complex, dynamic and noisy background is performed by neural systems through two concurring strategies. The first is the implementation of neurons with selective activation properties, in the sensory, motor or cognitive domains. The second is the task-dependent implementation of large cooperative neuronal assemblies. For reasons of methodological feasibility, the first strategy has been more widely explored than the second, both experimentally, by recording the activity of single neurons, and in the field of brain theory. However, the activation properties of a large neuronal population cannot be predicted from the activity of the individual single neurons that compose the population. First, data harvesting in single neuron experiments is extremely laborious. This severely limits the number of neurons that can be analyzed in any experimental condition. Second, the activities of neurons are recorded sequentially. Such a procedure fails to render the overall population dynamics in real time. Third, the single neuron approach is mainly restricted to the analysis of neuronal firing. Neuronal firing represents the output stage of neuronal computation, while much of the computation occurs at the level of subthreshold activity in synapses and dendrites.

Progress toward reconciling these different levels of description is possible by combining modern techniques that together are able to bridge the broad scales of spatio-temporal organization encountered in the brain. These are, for instance: i) recording the electrical activity of neuronal populations (multi-unit, multi-microelectrode approaches, high resolution EEG or EMG); ii) imaging brain activity (intrinsically generated optical signals, voltage sensitive dies, fMRI); iii) detailed quantitative reconstruction of neural structures (analysis of neuronal morphologies and interconnections using neuroanatomical methods with different degrees of spatial resolution). Other techniques are available in the field of computer science, and engineering, in particular: iv) methods for the analysis of signals, including spatial and temporal de-blurring of population codes; v) methods for the analysis of non-linear dynamical systems; vi) simulation of the activation properties of large neuronal assemblies based on the information gathered with the biological methods and on the detailed knowledge of neuronal properties and of neuronal interactions derived from in-vitro studies; vii) implementation of biologically inspired control systems in physical platforms or in complex dynamic simulations, that allow one to study the complete sensor-motor loop.

Research challenges

The methodological advances mentioned above open a new era of brain studies in which the following goals, some of which have been long-standing far-to-reach targets, seem now to be finally attainable:

- To understand the relation between the different methodologies for studying neuronal assemblies with the goal of reducing data harvested with different methods to a common phenomenology.
- To explore the dynamics of neuronal assemblies in situation of increasing complexity in the perceptual, motor or cognitive domain. For example, the dynamics of a neuronal population can be studied in different environmental conditions, i.e. with various sensory input, or in a specific situation, in which animal experience and/or behaviour can lead to differential activation of the cortical circuits.
- To identify common principles and diversities in the principles of neuronal assembly formation and management in the different perceptual, motor and cognitive conditions.
- To model the activity of the neuronal populations on the basis of the underlying connectivity of cortical neurons in combined anatomical and recording experiments and of the dynamic interactions between cortical neurons derived from *in vitro* studies.
- To use information derived from the study of population dynamics in animal models in the analysis of signals recorded from the human brain.
- To implement principles gathered from the study of animal and human neural processes to the design of artificial systems that should be capable of extracting meaning from complex sets of stimuli and/or of controlling complex networks of operators.

GROWING INTELLIGENT INFORMATION SYSTEMS

Background

An understanding of the fully functional mature brain at the systems level is crucial for the design and construction of 'bio-inspired' intelligent information systems. However, across many scientific disciplines, there is an increasing awareness that the operations of the mature brain are better understood within the frame of evolution, development, and plasticity that give rise to them. Some important reasons for this point of view are:

The processes by which living organisms organise and repair themselves are still poorly understood – but they appear to play a crucial role in attaining and maintaining the robust performance of biological systems.

To understand why a biological organism acts the way it does, requires consideration of the interaction of its ontogenesis and the history of its environment. The principles of this interaction explain why biological functions are realized the way they are; what alternatives could possibly realise the same function; and why some aspects of biological function have lost their relevance in the organism's present environment.

Bio-inspired information systems necessarily include interfaces to the world such as artificial retinas and limbs. Development, plasticity and learning are important processes in establishing functional connections between these devices and the nervous system.

Finally, adult skills in biological organisms are based to a large extent on neural circuits, which are not yet operational at birth. These circuits generate new skills through learning. There is, however, also a set of basic innate skills, (in mammals e.g. sucking, crying). The fact that there are such innate skills also has to be considered when trying to build systems whose intelligence can scale up from simple to complex sensory-motor coupling (e.g. from reflexes and central pattern generators to walking, grasping and perception and cognition).

The above concerns are as relevant to the fields of robotics and multi-agent systems as they are to biological organisms. Learning and development in an embodied artificial system; robustness in performance; autonomous self-construction and growth of artefacts ('epigenetic robotics'); and adaptation to the environment (possibly over several 'generations') have deep functional and economic significance for the design, construction, and maintenance of intelligent artefacts.

Developments in materials science and nanotechnology may be highly relevant to this focus. They are essential for the construction of sensors, processors, and actuators that are able to *replicate and co-evolve physically* with their system.

Research challenges

Projects that would fit the scope of this focus include:

- Growable sensors and motors (muscles, supportive tissues, receptor surfaces)
- Self-assembling structures with different purposes and targets, e.g. for "vertebrates" and "invertebrates"
- Exploration and application concepts from "genetic engineering" for evolving morphology over a life-time of interaction with the environment
- Functional organisation of evolving and developing biological and artificial nervous systems, with respect to intelligent behaviour, e.g. :
    - Development of peripheral and central nervous systems; autonomic nervous systems; hierarchical structure; modularity vs. full connectivity
    - Self-Organizing Maps for "graphical" primitives from sensor image in robot head: Architectures for combination of orientation, colour, movement sketches
    - Development of sensori-motor abilities by operational architectures modelled on cortex: for example, the self-organization of orientation columns in visual cortex; and a possible blackboard architecture in sensory cortex which combine multimodal primitives into objects
- Systems that develop complex motor skills by learning to combine existing motor primitives in novel ways.
- Social agents/systems that learn from observing (imitation learning) or from interacting (kinesthetic learning, social learning) with other more advanced agents/systems. Skills developed by these agents include:
- Bimanual coordination
- Complex visuo-motor coupling for differentiated grasping of objects
- Multimodal (visuo, motor, auditory) coupling for perceiving and recognizing other agents
- Acquisition of a simple language

SELF-AWARE CONTROL SYSTEMS - INTEGRATED CONTROL ARCHITECTURES THAT GENERATE AND EXPLOIT WORLD- AND/OR SELF-AWARENESS[*]

Background

Computer-based control systems in all fields (process control, avionics, robotics, etc.) are facing an enormous challenge. Trends in embedded software show a steady increase in size and complexity but the capability of the development technologies hasn't followed the needs and these days there is a manifest shortage of engineering capability to build the required systems of tomorrow. This is a very serious problem because complex information systems are not only required to provide increased levels of performance (for example in plant-wide optimisation) or enhance the constructability and maintainability of systems (for example in X-by-Wire systems) but also to increase or at least maintain the levels of dependability of the smaller embedded systems of the past.

Machines are becoming much more skilful thanks to the incorporation of massive doses of IT but deteriorating global system quality is a major threat that is limiting the application of the technology in some fields where dependability and/or adaptability is essential.

Valuable contributions follow from the proof-driven engineering approach but this has a serious flaw concerning our capability of designing and analysing of complex systems in the presence of state combinatorial explosion. Other alternatives that are being explored are those approaches based on the implementation of architectural mechanisms for self-organisation and self-repair. The system is constructed in a base state and it adapts based on the circumstances of the environment that surrounds the computing system. Sample exemplary efforts are those of the autonomic computing field but they lack the capability of incorporating reasoning about the physical world in the process of adaptation.

In some sense, the problems of the different approaches can be traced back to a deep deficiency of control systems: they do not understand the perceptual flow in the terms and up to the level necessary to provide the robust performance required for those systems mentioned above.

Research challenges

The objective of this research is to explore new architectural approaches in the design of intelligent controllers that attribute meaning to complex patterns of sensory stimuli and generate sequences of elementary actions that satisfy high-level goals. This should be done not only in terms of perception/action over the external world but also in terms of perception/action over the internal world, i.e. the body of the agent itself. An important problem is the proper charac-

[*] with a contribution from, Andrei Kirilyuk, Owen Holland, Aaron Sloman, William Edmondson and Stephen Torrance

terisation of what constitutes the body of the agent and what constitutes the surrounding world (very different perspectives can be found today e.g. in mobile robotics and process control systems).

The ultimate aim is to build systems that exhibit flexible, autonomous, goal-directed behaviour in response to changes in internal and external conditions based on a deep understanding of the world and the self. They will have integrated control architectures that generate and exploit world- and/or self-awareness.

This research can achieve results of extreme importance in the construction of highly resilient, adaptable information systems that are urgently needed in many fields of computing and, in particular, in dependable autonomous systems. Some of the challenges for this research are:

- The analysis and design of the architectural foundations of integrated controllers with explicit representation and exploitation of "self". This is expected to undo the fragmentation in recent science and technology of mind (put together mental pieces into "complete" architectures).
- Design of robust mechanisms for self perception and representation.
- Implementation of tools for architecture exploration by rapid prototyping and definition of test scenarios for evaluation of bioinspired self-aware systems
- Adequate and coherent interaction in the complex-dynamical system of environment, controlled plant, machine controllers, and human controllers.
- Mechanisms for active and passive screening of world- and self-perception.
- Performance analysis of "self"-based autonomous controllers.
- Design-centric vs. ontogenic approaches to self construction and exploitation. Constructive interaction between design-centric and ontogenic (conceptual, mathematical) approaches.
- Multi-agent systems for controlling buildings and spaces, which would – over time – explore their sensory and actuator space and the relations between the agents
- Generation of theory and technology of synthetic phenomenology.
- Analysis of social, psychological, ethical and other philosophical aspects of the proposed engineering and scientific realisation of self-aware (conscious) machinery and the corresponding direction of development in the whole; necessary corrections and practical recommendations for the engineering solutions.

LINK TO OTHER FET INITIATIVES

The research topics of this initiative are linked to those addressed in the previous three FET neuro-IT proactive initiatives: The "Neuroinformatics for Living Artefacts" initiative, which views entirely grounded and situated artefacts, the "Life-Like Perception Systems" initiative that centres on perception-action integrations, and the "Beyond Robotics" initiative that has an augmented scope for learning, adaptation, integration and systems research.

With respect to earlier initiatives Bio-i3 will specifically focus on the computational architecture of the mature brain and on mechanisms of development, evolution, plasticity and self-awareness in biology and artefacts.

SCIENTIFIC COMMUNITIES ADDRESSED

The initiative addresses a broad range of scientific communities: to design and build intelligent information systems requires multi-disciplinary teams drawn from research fields, such as advanced robotics and control, machine learning, network learning, signal processing, perception and cognition, embedded systems, neuromorphic-, biomorphic- and biologically-inspired engineering. Given the broad nature of the call, also fields such as cognitive neuroscience, experimental neuroscience, cognitive and developmental psychology, biological cybernetics and neuroinformatics are addressed.

A primary aim of the initiative is to bring together a broad spectrum of scientists from the above scientific communities to work creatively and synergistically toward the realisation of the objectives and challenges of the initiative.

The research communities addressed are also those concerned with the construction of technical systems inspired by reflective and self-awareness control properties of the biological brain. Some examples are autonomous control systems, autonomic computing systems, and dependable information systems, complex and embedded systems.

ETHICAL ISSUES

This research is not focused on experimentation with humans and/or animals and hence it does not have ethical implications in this sense. The focus is on strictly technical matters and the only possible ethical issue is related to the remote possibility of achieving total success: i.e. being able to construct self-aware person-like synthetic systems. While this is not a concrete objective of the research proposed, it is obvious that if this research is totally successful the construction of synthetic persons should be considered as a possible application of the science and technology developed.

Powerful technological innovations tend to proliferate beyond control and if genuine artificially conscious systems became as widespread as PCs or mobile phones, questions of misattribution would be very serious. Also the consequences of a mass social misperception of artificial consciousness need to be considered. Such indirect social effects could be enormous, and difficult to predict. Like reproductive cloning, this involves ethical considerations that are quite novel, and difficult to focus on clearly. Ethical consequences of a novelty need not be negative, especially in this case, but any wide research programme should have a significant sub-project to investigate the ethical implications.