A Roadmap for NeuroIT

Challenges for the Next Decade

BY MARC DE KAMPS AND ALOIS C. KNOLL

What is NeuroIT? And why does it need a roadmap? After all, there is something called "neuroinformatics," which has been around for a while and which is growing rapidly. The term neuroinformatics is often used to refer to the application of information technology (IT) to the "brain sciences." Almost all of the brain sciences have become considerably more complex, and recording and managing the results from experiments entails the use of ever-more complex and larger databases and analysis tools. Examples are the large heterogeneous data sets produced by fMRI machines, the complex electro-chemical mechanisms and genetic factors that determine neuron function, and so on. To understand these data, increasingly complex and time-consuming models are necessary, which run on ever-larger computers, which sometimes need novel architectures to run efficiently. An interesting overview of the activity in this area can be found in two reports published by the Organization for Economic Cooperation and Development (OECD) [1], [2].

Where our knowledge of the brain and brain function is expanding rapidly, our ability to make use of this information has somehow not increased at the same rate. Living creatures still outperform computers in a large range of skills, many of which are considered to be "simple." Computer scientists can only dream of the object-recognition skills of humans, and roboticists would love to create service robots with the same degree of autonomy as an ant. The possibility for the development of artifacts that are able to learn over their lifetime and are able to adapt their behavior in the face of changing circumstances seems even more remote. In general, every complex artifact has to be programmed carefully, by hand, and for a new range of applications this has to be done again, from scratch.

The relatively slow progress in the creation of bio-inspired artifacts and IT applications is a source of frustration for policy makers, scientists, and engineers alike. Scientists and engineers that try to emulate methods used by nature find that their bio-inspired approaches work very well on some problems, while failing on other, seemingly related ones. Or they find that approaches that are promising on toy problems do not scale well with the problem size. Behind these problems is a lack of systematic understanding of how nature accomplishes things, which, as we will see, is one of the central issues in the roadmap. This makes bio-inspired engineering difficult and solutions to many problems can only be found by trial-and-error. The haphazard development of bio-inspired engineering is also undesirable from a political point of view: clearly there is great economic potential in some fields of research and in order to stay competitive it is important to know which research to fund. Moreover, the impact of new technology on society can be considerable, as we have seen recently with the Internet.

These considerations have led to the creation of nEUro-IT.net (http://www.neuro-it.net), a thematic network in the area of NeuroIT. NeuroIT has here been defined loosely as "neuroscience for IT," whereas in "neuroinformatics" the emphasis has more been on "IT for neuroscience." The distinction, therefore, is not so much in the field of study, or in the techniques used, but rather in the long-term objective of the research. nEUro-IT.net is funded by the Future and Emerging Technology (FET) arm of the Information Society Technology (IST). One of its most important activities is the creation of a roadmap.

The reasons for a roadmap have been stated implicitly already: it tries to develop a vision for where the field will be in the next decade. It tries to identify problems that affect the field as a whole and how they can be solved. It serves as a reference for the state-of-the-art research in various fields for researchers and for decision makers. This is perhaps even more important for NeuroIT than for other fields, since it is highly interdisciplinary: a computer scientist cannot be expected to know about the latest developments in primate vision, and yet these developments may provide crucial inspiration for new approaches in computer vision.

The Creation of a Roadmap

The NeuroIT community is very heterogeneous and the creation of a roadmap entails a number of practical problems. First of all, a constituency must be formed in a field that does not yet exist. The FET's systematic funding of NeuroIT proved instrumental here: although mostly unaware of each other's existence, quite a large group of people had been working in multidisciplinary projects for some years. They were first brought together in a kick-off meeting in 2002 December in Leuven, Belgium. Second, people must be convinced that

 0739-5175/07/$25.00©2007IEEE

Since nature has solved this problem very efficiently in the course of evolution, it is necessary to look to nature for guidance on how to design complex artificial systems.

roadmapping is worth their time. At first, this was a difficult process. After a "start-up" meeting it was simply decided that each of the attendees write a "grand challenge." The "challenges" together constituted the first version, which was published on the nEUro-IT.net Web site. After a Web consultation, which was supported by FET, the document had generated considerable attention and it became easier to ask for contributions, especially since the document had since been used in the formulation of several funding call texts.

The Roadmap: Topics in NeuroIT

We briefly discuss the topics of the roadmap here, the expected benefits, the resources required, and the most important obstacles that stand in the way of their realization. The discussion is necessarily brief, but the interested reader can always obtain the roadmap at the nEUro-IT.net Web site. References can be found there.

"Brainship": Human–Machine Interaction

Controlling machines by mere thought is an old engineering dream. It has obvious applications; for example, controlling prosthetics in the form of artificial limbs. Another possibility would be teleoperation of remote exploratory vehicles, equipped with artificial sensors, ranging from microendoscopes to deep-sea vehicles. Yet another application would be a direct interface with information systems. Although this sounds like science fiction, recent progress in the use of multielectrodes implanted in the brain suggests that this is a real possibility. A major breakthrough is the discovery that the brain has enough plasticity to adapt its signals to communication over a limited number of channels, and it has therefore become possible to predict limb movements from the activity of multiple single-neuron recordings in the motor cortex [3]. This was first demonstrated on rats and later in monkeys. An impressive demonstration was given in two experiments [4], [5], where brain signals directly control the position of a cursor, using visual feedback from the screen.

While important as proof of principle, a major problem in controlling a prosthetic device is the lack of somato-sensory feedback in current experimental designs. There is still a long way to go in this area: a better understanding of principles of neural coding is needed, in particular how the brain integrates sensory and motor systems, both in fast motor control and in decision. To name a few problem areas: one needs stimulation multielectrode arrays to allow for a direct input of sensory information. It is important to find out if there are alternatives to implanting electrodes in the brain, and, as long as this is not the case, to improve the durability of electrodes implanted and to reduce the impact that they have on brain tissue. Important ethical issues are the possibility of brain damage caused by invasive techniques: under what circumstances is this acceptable, if at all? And for invasive and noninvasive techniques alike, there is the possibility that the induced brain plasticity interferes with normal brain function.

Factor X: A Machine that Grows by a Factor of 10 in Size, Strength, and Cognitive Abilities

Inspired by the possibilities that new materials and nanotechnology offer, some roboticists dream about machines that coevolve their brains and bodies in continuous interaction with the environment over a limited period of time. This vision is largely inspired by the development of living organisms and the theory of action-centered cognition [6]. One may think of a self-assembling robot, based on "genetic" information, under the influence of the environment. Such robots would be autopoietic (built from the inside out) as opposed to the current generation of robots, which is allopoietic (built from the outside in). A possible starting point for such robots would be a (modified) biological substrate. This has several disadvantages: the creation of new organisms, which is the implication, would be a major ethical problem. The size of such artifacts and the capabilities of its sensors and actuators would be limited to biological ranges. It seems more realistic to look at the possibilities that recent developments in material science offer. Currently there are a few developments that could be taken as a starting point for this challenge: Modular robots are built from a certain number of identical motor modules and can be combined into different shapes and macro structures, evolutionary and epigenetic robotics, and nanoscale self-assembling structures.

From today's perspective there are four lines of research that could be considered to bring a project like this underway. First is molecular robotics, which is the exploration and design of materials and substrates that lend themselves to build "cells" that can be made to meet the different requirements for various body areas. Second are distributed growable sensors for distributed areas of sensor cells. Here it will be necessary to investigate how they can be coordinated and produce sensible results when they are distributed over a large surface of the outer body ("skin") and are physically connected by a medium that has a high degree of flexibility ("body"). Third is growable distributed information processing, which is a demanding research area because the information processing has to control the artifacts from the moments of "inception." Hence, this system not only has to control its actuators and sensors but it must interact with the environment to control the growth of the artifact and co-evolve with the increasing capabilities of its sensors and actuators. Fourth are growable motor entities and spatially distributed actuators. The actuators must be controllable as they develop their actuator part and their support structure. The development must be in sync with the growth of the size of the artifact.

The benefits of this kind of construction method are quite clear, but it will take considerable research to see if basic building blocks can be designed that are cheap and have the required properties.

Conscious Machines

From the point of view of NeuroIT, three motivations to study consciousness in relation to IT artifacts appear repeatedly in the literature. First of all, it has proved to be far more difficult than expected to deliver artifacts that are truly autonomous and that are able to learn over their lifetime. Artifacts still have to be programmed very carefully for every new task that they have to perform. Such programming is obviously limited, since new and unforeseen situations can not be preprogrammed. This is turn sets severe limitations on the autonomy of the artifact. The second motivation has to do with the way that humans interact with artifacts and how they perceive them. The third motivation has to do with the fact that the neurosciences start to present the outline of a systems-level understanding of the brain. Still, there is much to be understood on this level, and given the experimental difficulties that are involved (see the "Brainprobe Project" section), it is widely believed that building artifacts is actually an interesting way to study aspects of higher-level cognition, including consciousness.

It is generally assumed that complex autonomous artifacts need a sense of "self" [6], [7] to monitor the effect of their actions, in past and present, on the world, and to be able to "learn" from this in a way that is not preprogrammed. Many people also believe that emotions, awareness, and attention play an important role in such processes, and so there is a considerable interest in a more rigorous definition of these concepts and in the way humans and animals use them. An interesting problem that relates to consciousness is how it emerges: a living creature is specified by a genetic code, which is rather limited and cannot prescribe the development of an organism in detail. This raises the question of how consciousness emerges in a growing organism, and if lessons can be learned concerning the construction of artifacts from this. Other questions are how the embodiment of an organism or artifact influences its mental representations.

The interest in several fields of engineering and IT for concepts, which used to be studied exclusively by psychologists (e.g., [8]) is relatively new; e.g., [9]–[11]. To make progress, other challenges in the roadmap are instrumental. Ethical issues involved first of all relate to humans. Truly autonomous service robots may replace humans, for instance. If progress in this field is really successful, however, we might have to extend such considerations to machines.

Successful in the Physical World

This challenge analyzes the reasons for why artificial systems perform so poorly compared to living ones from a slightly different perspective. For the creation of an intelligent system, a designer has options that basically can be classified as the choice for computation and control strategies, the choice for morphology, the choice of materials, and using the environment. The hypothesis behind this challenge is that living creatures are optimized with respect to all of these options, as opposed to conscious machines, where the emphasis was on the assumption that living creatures have computational capabilities for reasoning, planning, etc., that are vastly superior to man-made algorithms devised to reproduce these skills.

Significant simplifications in the control loop of an agent can be achieved if its computational capabilities are distributed over the central nervous system, the peripheral system, the materials of its body, and its interaction with the environment. A well-integrated periphery may be the key in lowering the difficulty of a task to a point where it becomes feasible as well as take the load off a central processing module. Prime research objectives for this challenge are, therefore: intelligent periphery, system integration, morphology and materials, and "environment models" used to codify task/world knowledge. Despite the fact that neural processing will remain an integral part of agent construction, the focus of this challenge is on making the periphery smarter and integrating it better with central computations. An important aspect of this is the development of universal standards ("bus standards") for smart reusable peripherals, which would facilitate cooperation between different projects and would facilitate the introduction of new technology into product design. This would then create building blocks for individual systems but also allow capabilities for robot interaction and collaborative behaviors.

Automated Design of Artificial Cognitive Systems

In this chapter of the roadmap, it is argued that the reason that artificial intelligence in its various guises has been unsuccessful in modeling and explaining sophisticated cognitive functions is the fact that sophisticated cognitive algorithms are complex in the technical sense that they cannot be compressed into a compact piece of code [12]. The question of how to design artificial cognitive systems then arises. Since nature has solved this problem very efficiently in the course of evolution, it is necessary to look to nature for guidance on how to design complex artificial systems. The aim is therefore to create a theory of the evolution of complex systems in nature and to apply the theory to generate biologically inspired techniques for the automated design of artificial cognitive systems.

A key reason that is cited for lack of progress so far is the fact that the biological processes that artificial evolution has to emulate are incredibly hard to model. As a result, the main thrust of recent biological research is toward the investigation of specific organisms and systems rather than broad theory. Traditional evolutionary theory lent itself naturally to mathematical modeling, whereas more recent research has generated a vast wealth of disjointed information that has yet to be adequately organized.

To solve this problem, one needs a theory of the evolution of complex systems. In the challenge, various lines of research are proposed that include evolutionary techniques together with a realistic modeling of the physics and the chemistry that is involved. Importantly, a set of benchmark problems must be included: projects that are too difficult to solve by contemporary software techniques but sufficiently simple to give developers a real hope of success.

Constructed Brain

In this challenge the possibilities for the "complete" simulation of a brain are investigated. There are at least three major problems that hamper the understanding of the brain. The first is the sheer complexity of the brain, in terms of number of components and in terms of the physical and chemical processes that control its function. The second is that the brain is hard to divide in modules with a well-defined function and that many aspects of the brain are hard to study in isolation. The third is that a large number of disciplines are involved in the study of the brain, each with its own methodology, terminology, and traditions.

To overcome these problems it is suggested to create a framework that allows a large-scale, coarse simulation of the brain, with sufficient flexibility to create more detailed simulations locally, where needed, or to increase overall sophistication when computer power increases.

Theoretical methods are reviewed that could be important for the creation of such a framework. It is clear that the simulation of the 100 billion neurons of the human brain is not currently possible, and even if this were possible, one would still need to find the correlates of cognitive states and behavior. Special consideration is therefore given to techniques that yield a "mesoscopic" or "thermodynamic" description of groups of neurons; e.g., [14], [15].

Another area that is important is the identification of computational architectures; e.g., [16], [17]. The large-scale organization of the cortical networks is starting to emerge from a combination of imaging data, psychophysical experiments, and theoretical modelling (see Figure 1). At the same time, multi-electrode arrays provide new insights in the function of the local cortical circuits, which are repeated over the entire brain (see Figure 2). To find organizational principles that are repeated over and over again, and which explain how a massively parallel network of relatively slow elements can perform complex computations, is extremely important for NeuroIT.

There are many issues involved in the start of such a project, too many to mention here, but one of the most important is software engineering. It is obvious that in such a project many software libraries have to be integrated in a flexible way. Taking this one step further, it will be necessary, given the complexity of the project, to invent new ways of publishing models in addition to journal articles. The publication of models in itself would be an interesting issue in software engineering as well.

It is hard to overestimate the potential benefits of such a project. Modeling and software development in brain science is primitive compared to disciplines such as high-energy physics, which have established traditions of software engineering. This is odd, given the heterogeneous data and models in "brain science" and the inherent complexity of models of the brain. It is probably necessary to establish an agency to start up this project, since it is beyond the capabilities of individual research groups.

The Brainprobe Project

The Brainprobe Project starts with a few observations on the importance of an understanding beyond the single neuron level for NeuroIT, it attempts to identify the most important obstacles for such an understanding from the point of experimental neuroscience, and it suggests several lines of research to overcome them. It starts with an extensive review of the current experimental techniques to measure brain activity: PET, fMRI, (multiple) single electrode measurements, multielectrode measurements, and optical techniques. It discusses the (combinations of) relative strengths and limitations of these techniques and how they relate to unresolved issues concerning the interpretation of neuroscientific data and, as such, contains a valuable comprehensive review of the experimental state-of-the-art in neuroscience.

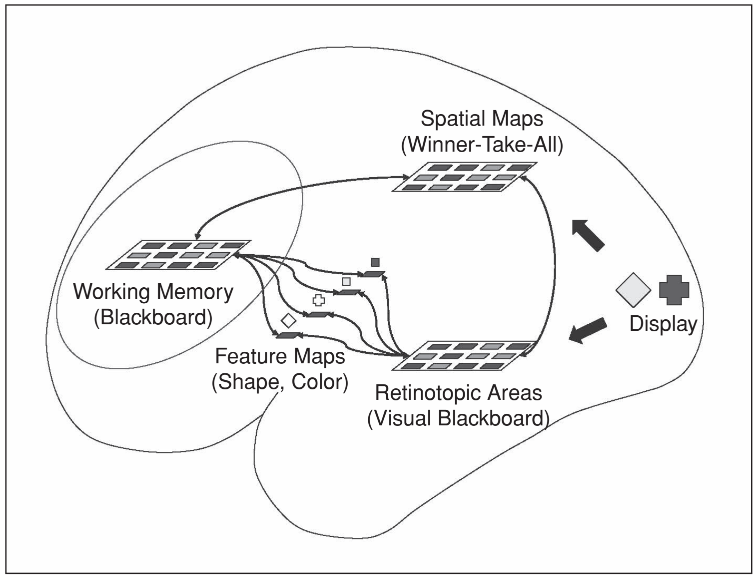

Fig. 1. We slowly start to understand how local networks are organized into large-scale cortical networks. Can we organize massively parallel computers in a similar way?

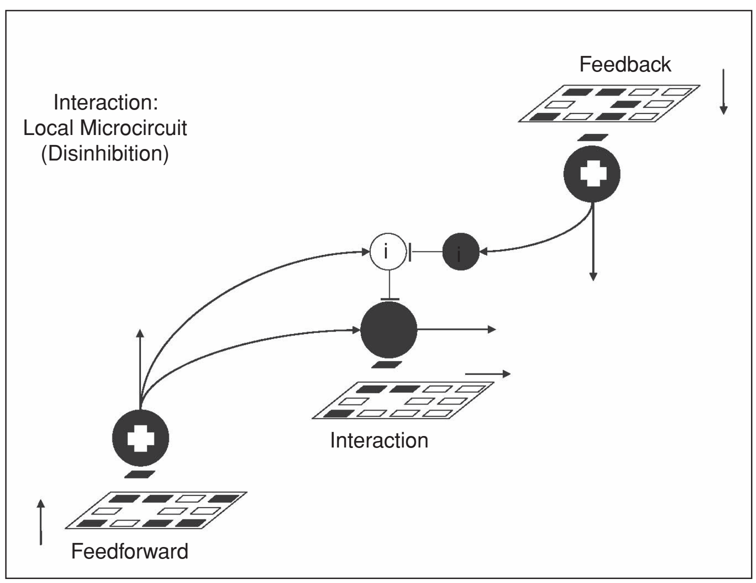

Fig. 2. Detailed simulations of local microcircuits are sometimes necessary to understand the big picture. The interaction between local and global structures is a characteristic of most natural systems and this seems to apply to successful bio-inspired artifacts as well.

There are several experimental issues that hold the promise for progress. The combination of different techniques in one experiment (fMRI with EEG, for instance) would enable the use of the good spatial resolution of the former and the good temporal resolution of the latter. Multi-electrode measurements consisting of several hundreds of electrodes, implanted in several brain regions, would offer the possibility to observe synchrony between different brain areas, to study functional architecture, and to input signals into the areas. Finally, the development of new mathematical tools to interpret the new data is essential.

Conclusions

It is clear that the challenges presented here are very ambitious, and that some of them can only be started in the future. It is also clear that there are strong interrelations. "Conscious Machines," for instance, will need some of the research that was described in the "Constructed Brain" and the "Brainprobe Project" sections. Very interesting is the fact that in at least three challenges the need was recognized for collaborations that extend beyond small research groups. This is true for the "Acting" challenge, which calls for standardization of peripherals; for "Constructed Brain," which claims that models created by individual research groups cannot capture the complexity of the brain and calls for a framework to connect such models; and also for the "Brainprobe Project." This shows the need for a NeuroIT community and also for more permanent funding initiatives in these areas that would allow the creation of communities and institutes that would tackle these standardization and collaboration issues in a more systematic way than is currently the case.

Recently, a new version of the roadmap was published (version 2.0). It includes a new chapter on bio-inspired hardware, which describes VLSI design of neuromorphic chips and evolvable hardware, among others. Nearly all chapters have been updated to the state-of-the-art of summer 2006.

Acknowledgments

Contributors to the roadmap are Marc de Kamps (Constructed Brain), Alois Knoll (Factor X, Conscious Machines), Rolf Müller (Successful in the Physical World, together with John Hallam and Herbert Peremans), Guy Orban (The Brainprobe Project, with contributions from Giacomo Rizzolatti), Giulio Sandini (Conscious Machines, with contributions from Ricardo Manzotti and Vincenzo Tagliasco), Erik de Schutter (Brain Ship), and Richard Walker (Automated Design of Artificial Cognitive Systems). The Conscious Machines chapter of the roadmap is currently being reviewed and extended by Igor Alkesander, John Taylor, and Ricardo Sanz. New contributors to version 2.0 were: Anders Lansner (Constructed Brain), Eduardo Ross, Fransisco Pelayo, Giacomo Indiveri, and Tobi Delbrück (Bio-Inspired Hardware).

Marc de Kamps studied theoretical physics at the University of Amsterdam. He obtained a Ph.D. in 1996 in experimental high-energy physics. After finishing his thesis, he worked at the Psychology Department of Leiden University, modelling the neural substrate of attention. He then worked at the Technical University of Munich, modelling visual attention and combinatorial productivity in neural networks and assisting Alois Knoll in the coordination of nEUro-IT.net. He is now a lecturer in the Biosystems Group of the School of Computing at the University of Leeds.

Alois C. Knoll received the diploma (M.Sc.) degree in electrical engineering from the University of Stuttgart, Germany, in 1985 and his Ph.D. (summa cum laude) in computer science from the Technical University of Berlin, Germany, in 1988. Since 2001 he has been a professor of computer science at the Computer Science Department of the Technical University of Munich. His research interests include sensor-based robotics, multiagent systems, data fusion, adaptive systems, and multimedia information retrieval. He is currently coordinator of nEUro-IT.net.

Address for Correspondence: Marc de Kamps, Biosystems Group, School of Computing, University of Leeds, LS2 9JT Leeds, United Kingdom. E-mail: dekamps@comp.leeds.ac.uk.

References

[1] Biological Informatics Group, "Neuroinformatics: A developing field, uniting the neurosciences and interfacing it to the medical and information technology domains (version 6)," in *Proc. OECD Megascience Forum, Neuroinformatics Subgroup*, 1998 [Online]. Available: http://www.oecd.org/dataoecd/24/32/2105199.pdf

[2] Report on Neuroinformatics from the Global Science Forum Neuroinformatics Working Group of the OECD, 2002 [Online]. Available: http://www.oecd.org/dataoecd/58/34/1946728.pdf

[3] J.K. Chapin, K.A. Moxin, R.S. Markowitz, and M. Nicolelis, "Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex," *Nature Neurosci.*, vol. 2, pp. 664–670, 1999.

[4] M.D. Serruya, N.G. Hatsopolous, L. Paninski, M.R. Fellows, and J.P. Donoghue, "Instant neural control of a movement signal," *Nature*, vol. 416, pp. 141–142, 2002.

[5] M.D. Taylor, S.I. Helms-Tillery, and A.B. Schwartz, "Direct cortical control of 3D neuroprostethic devices," *Science*, vol. 296, pp. 1829–1832, 2002.

[6] I. Aleksander, "The self 'out there'," *Nature*, vol. 413, pp. 23–23, 2001.

[7] I. Aleksander, *How to Build a Mind*. London: Weidenfeld & Nicolson, 2000.

[8] B.J. Baars, *A Cognitive Theory of Consciousness*. Cambridge, U.K.: Cambridge Univ. Press, 1988.

[9] S. Grosberg, "The link between brain learning, attention, and consciousness," *Consciousness Cognition*, vol. 8, pp. 1–44, 1999.

[10] S. Harnard, "Can a machine be conscious? How?," *J. Consciousness Studies*, vol. 10, pp. 67–75, 2003.

[11] J.G. Taylor, "Paying attention to consciousness," *Trends Cognitive Sci.*," vol. 6, pp. 206–210, 2002.

[12] G.J. Chaitin, "Randomness and mathematical proof," *Sci. Amer.*, vol. 232, pp. 47–52, May 1975.

[13] J. Hertz, A. Krogh, and R.G. Palmer, *Introduction to the Theory of Neural Computation*. Reading, MA: Addison-Wesley, 1991.

[14] J. Eggert and J.L. van Hemmen, "Modeling neuronal assemblies: Theory and implementation," *Neural Computat.*, vol. 13, pp. 1923–1974, Sept. 2001.

[15] B.W. Knight, D. Manin, and L. Sirovich, "Dynamical models of interacting neuron populations in visual cortex," in *Symp. Robotics and Cybernetics: Computational Engineering in Systems Applications*, E.C. Gerf, Ed. Lille, France: Cite Scientific, 1996.

[16] F. van der Velde, "On the use of computation in modelling behavior," *Netw. Computation Neural Syst.*, vol. 8, pp. R1–R32, 1997.

[17] F. van der Velde and M. de Kamps, "Neural blackboard architectures of combinatorial structures in cognition," *Behav. Brain Sci.*, vol 29, no. 1, pp. 37–70, 2006.