Available online at www.sciencedirect.com

ScienceDirect

Cognitive Systems Research 14 (2012) 1–9

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Towards truly human-level intelligence in artificial applications

Action editors: Minho Lee and Wlodzislaw Duch

Marc de Kamps

Institute for Artificial Intelligence and Biological Systems, School of Computing, University of Leeds, Leeds LS2 9JT, UK

Received 1 April 2010; accepted 12 December 2010
Available online 31 January 2011

Abstract

Despite the fact that there are now a large number of successful bio-inspired applications in use in science and technology, we are still quite far removed from creating applications that display human-like intelligence. Putting together successful bio-inspired applications remains something of a black art; this is due to a lack of fundamental understanding of brain function. The causes for these problems were analysed in a 'Roadmap for Neuro-IT' and were deemed to be sufficiently pressing to motivate one of five 'Grand Challenges' in Neuro-IT: the 'Constructed Brain'. The challenge argued that one of the main bottlenecks to progress is that data taking and modelling in the neurosciences are being fractured across many research groups and communities; it makes proposals for addressing the issue. Similar observations, raised in two OECD workgroup papers have led to the formation of the International Neuroinformatics Coordinating Facility. As a consequence we can conclude that there is now a much higher awareness of the problems and that in the neurosciences the situation has improved dramatically. I will review recent initiatives to facilitate data management, modelling and simulation in the neurosciences. One problem remains unaddressed, however. The project-based funding of the brain sciences sets an upper limit to the complexity of brain models. Since the brain is truly complex, any individual project will fall short of capturing the brain's complexity. The creation of a central infrastructure for the brain sciences is inescapable, but is unlikely to be realised soon. I will outline suggestions to handle the current situation.

Keywords: Neuroinformatics; Computational neuroscience; Cognitive neuroscience; Framework; Middleware

1. Introduction

There is a considerable optimism that our understanding of the brain and the computational principles on which it operates will improve substantially in the near future. This optimism is fuelled in part by the incredible development of experimental techniques in the neurosciences, such as fMRI, EEG and PET, but also by the explosive development of multi-electrode arrays (MEAs) and the data analysis methods to make optimal use of these techniques. Computing power to analyse the data that these techniques deliver has become cheaper rapidly. There is hope that we will learn some fundamental computational principles, which can then be transferred into technology. There are at least four major challenges, which, when taken on, could lead to significant technological advances:

- The interface between the Central Nerve System (CNS) and machines, sometimes called Brain-Computer Interfacing (BCI).
- The creation of 'intelligent' machines: machines that demonstrate flexible behaviour and that are able to adapt to unforeseen circumstances. The creation of a brand of machines which can be customised easily to tasks outside the original design specifications, without requiring a re-design of the machine.
- New and better ways to study the brain. If one were able to do *in silico* simulation of psychopharmica effects, the pharmaceutical industry would be able to design them in a more systematic way, and, hopefully, in the long run be able to do without animal (and human) testing.

E-mail address: M.deKamps@leeds.ac.uk

1389-0417/$ - see front matter
doi:10.1016/j.cogsys.2011.01.003

- An understanding of the brain itself. A better understanding of brain function will beyond a doubt have a profound impact on the treatment of psychological disorders, which are a source of distress to many and also a source of substantial economic damage. It will also shed light on long standing questions in the philosophy of mind.

The neurosciences have received substantial investments during the last years: the amount spent on research-dedicated functional Magnetic Resonance Imaging (fMRI) scanners alone is astronomical. The USA, Germany and Japan have dedicated programs for the neurosciences (e.g. http://www.riken.go.jp/engn/index.html, http://crcns.org/about, http://portal.uni-freiburg.de/nncn) that run in the hundreds of millions of dollars. Although there is a fundamental interest in the brain, most investment will be motivated with these four areas in mind. If substantial returns will not be forthcoming in the next decade it is likely that future investments will be reallocated to other fields of science so the current interest in neurosciences represent something of a window of opportunity. It is worth to examine the obstacles.

Many funding organisations try to predict funding requirements to achieve their strategic objectives. The Future and Emerging Technology arm of Information Society Technology commissioned a so-called Roadmap for Neuro-IT (de Kamps & Knoll, 2007; Knoll & de Kamps, 2003) in early 2003 to help decide its funding objectives for the 7th Framework Programme. Its scope was: ‘neuroscience for IT, not IT for neuroscience’. It intended to promote a general understanding of neural processing, so that engineering could be more ‘intelligent’ (in the human sense of the word, whatever that may mean), flexible and adaptive. It was less concerned with IT applications to construct databases or visualisation tools to handle neuroscience data better, except where it could be argued that this was inevitable to achieve the main objective: an understanding of the engineering principles of the brain to the point were they could be applied routinely. In the process of writing the Roadmap, experts from different areas of ‘the brain sciences’ identified several obstacles to progress. The identification of these obstacles is crucial for the field as a whole because it should identify where funding should go.

In this paper, I will discuss the obstacles that were identified in the Roadmap. Since the first publication of that Roadmap some of the concerns raised have been addressed. In the neurosciences some form of coordination has been established in the form of the International Neuroinformatics Coordination facility (INCF), but I will argue that many of the problems discussed in the Roadmap are still on the table. In the absence of a shared common infrastructure for the brain sciences individual researchers can address some of the obstacles in their own work. In the last section of the paper I will give examples of how to do this. Much of the discussion will relate to current developments in computational neuroscience and neuroinformatics. This is predicated on the assumption that large-scale integrated models of brain function will be necessary to thoroughly understand brain function. If this is true, the issues raised in this paper relate to psychology, artificial intelligence and clinical research as well. My prediction is that in the next decade the issues raised in this paper will emerge in these areas.

2. Obstacles to brain-inspired engineering applications

The basic operation of neurons has been understood since the 1950s (Hodgkin & Huxley, 1952). The realisation that such neurons must be studied in networks came quickly and because of limitations in computers at the time such networks were studied with simplified neural models during the three decades thereafter. Important fundamental principles of neural computation were established relatively quickly: the importance of winner-take-all circuits, the realisation that feedforward networks can essentially implement any input–output transformation (Hornik, Stinchcombe, & White, 1989; Rumelhart & McClelland, 1986). These initial successes, achieved with very simple network models, shed light on some fundamental ways in which neural circuits can do computation. Although this has led to a large number of engineering applications, a general theory of neural computation has never been established. There are some candidates for a proto-theory, but at present they all still fail to the establish generic principles of neural computation: if we were given a brain-sized massively parallel piece of hardware today we still would not be in a position to recreate object-recognition, object-handling, navigation, language recognition or production on par with their biological counterparts. We would not know how to program it. This means that the function of many engineering applications can not be guaranteed: it is sometimes not quite clear why they work, or there are insufficient guarantees that they will perform to a high enough standard in all situations.

In response, two main approaches have been taken. The field of machine learning has studied theoretical aspects of processes such as pattern recognition, categorisation, and object-recognition. As a consequence it has focused on algorithms and methods for which good theoretical rationales can be found. Interestingly, this field seems to have moved away from neural networks as a consequence. In modern textbooks on Machine Learning (Bishop, 2007) neural networks are only mentioned in passing. There are counter examples: Olshausen and Field’s work (Olshausen & Field, 1996) give a statistical explanation for receptive field characteristics in primary visual cortex, suggesting that cells in primary visual cortex are ideally suited to decompose natural images and represent them by a minimal number of neurons (sparse representations). Such explanations are very important because they suggest a statistical explanation for cortical function which is not directly tied up with biological substrate: other implemen-