A real-time computer simulation of the human brain

Steve Furber, 13 October 2002

A draft proposal for the UKCRC Grand Challenge workshop

Understanding the human brain is, arguably, the grandest challenge of them all. It is the most powerful information processing system known to mankind, with capabilities we are nowhere near to emulating with our artificial computing machines. Using fairly low performance unreliable component technology, humans can understand the visual world and control a complex of muscle systems in ways that surpass our engineering knowledge. Then there is the phenomenon of natural language…

[In what follows I propose a largely bottom-up connectionist approach to the problem, but others would argue that it is better to start higher up – after all, human brains are not the largest in the animal kingdom, so cognitive functions are not simply a property of size. A system level approach to the problem might be more promising, and would at least be a valuable counterbalance to the approach proposed below.]

The human brain consists of 10^{11} neurons each with, on average, around 10^4 inputs. Neurons typically fire at a few 10s of Hz, so the computing power required to model the brain at this level must be able to handle around 10^{16} connections per second. If we knew exactly how each neuron behaved, how it is connected and how to model it, this would represent a formidable computing challenge. But we don't even know that much. There are major challenges in understanding the abstract computational function of a neuron (that is, what aspects of its natural behaviour are essential to its computational function as opposed to keeping it alive, providing energy, or gratuitous complexity resulting from its evolutionary heritage), the coding of information in neural connections, the high-level architecture of brain structures, and so on.

Progress is, however, being made. A recently-launched Foresight Project into Cognitive Systems...
http://www.foresight.gov.uk/servlet/Controller/ver=575/userid=2/
...is looking into the prospects for establishing inter-disciplinary activity that brings together neuroscientists and computer scientists to see if there is potential for mutual benefit from sharing their respective understandings of natural and machine intelligence. Analytical tools employed by neuroscientists are providing insights into the functions of neural systems in more detail than ever before, and computer models of natural systems are increasingly successful in modelling their behaviours.

Even partial progress towards this goal offers great benefits in the understanding (and therefore potentially the treatment of) brain pathologies and the development of computer systems that can emulate human capabilities more readily, easing human-computer interactions. Achieving the Grand Challenge will yield machines with capabilities which up to now have existed only in the imaginations of writers of science fiction.

Feasibility

The feasibility of this Grand Challenge depends on computing resources (processing speed, memory, communications) and the availability of biological data. It also depends on less tangible factors such as a breakthrough in our understanding of neural codes, new mathematical models, neurons not being a whole lot more complex than they appear to be, and so on. Here I just look at the modelling issues and do some rough calculations on the machine resources that might be required:

Memory: If we assume that each of the 10^{15} synaptic connections requires a byte to indicate its strength the memory needed is 10^{15} bytes (regularity in the structure may reduce this significantly), which with memory today costing around $1 for 10Mbytes would cost $100M, reducing by a factor 1,000 over the next 20 years. This is clearly within the scope of a research budget, but not a high-volume consumer product.

Processing: The processing power required can only be guessed, but let's assume that each connection requires 10 instructions. We therefore require 10^{17} IPS or 10^{11} MIPS. Current general-purpose processors deliver 10^{3} MIPS, rising to 10^{4} MIPS over the next 10 years. 10^{7} processors will therefore be required, costing perhaps $100 each, so $1,000M in all. As another reference point, high-end supercomputers are aiming for a petaflop (10^{15} FLOPS) over the next 5 years and will cost around $100M. A lot depends here on the complexity of the neural algorithms. If they can be cast into hardware the cost may be reduced by several orders of magnitude; the estimates above suggest this is a requirement for meeting this Grand Challenge.

Communication: A hardware-based system might require around 10^{5} nodes, each node modelling 10^{6} neurons with 10^{9} bytes of RAM and 10^{6} MIPS. Each neuron would generate 10 events/s, so the communication requirement would be around 1Gbit/s per node, which is quite manageable today. The total communication bandwidth is 10^{5} Gbit/s, so the communication architecture would have to exploit locality and be carefully designed to avoid bottlenecks.

Phases

Phase 1: understand neural coding and neural models. Current wet science delivers coarse activity measures ("this area of the brain is active at this point"), individual neural activity measures ("this neuron produces this series of spikes") and very detailed characterisation of the operation of an individual neuron ("this chemical has a role in long term potentiation through this mechanism"). What is missing is any insight into the fundamental computational abstraction ("this group of neurons encodes this information in this way"). Computer modelling of neurons works the other way, postulating neural encoding schemes and then demonstrating that these schemes can emulate natural behaviours. What is required is for these two activities to join in the middle so that the computer models generate hypotheses that can be verified experimentally. A key question here relates to how the brain achieves the immense dimensionality reduction apparent, for

example, in the visual system (where all visual information is compressed to pass down a million fibres in the optic nerve).
Timescale: years 1-10

Phase 2: Model natural systems of increasing complexity. The sea slug (7 neurons) has already been modelled. The nematode worm C. elegans has a hundred neurons and a completely known morphogenesis. The fruit fly has 400 neurons, the functions of all of which are understood. Perhaps a bee brain (10^6 neurons) would be another good intermediate challenge. There has been quite a lot of work on synthetic neural systems that will produce the cockroach walking system (for example), including its response to external stimuli. Constructing a fully working house fly would be an interesting challenge for the micromachine folk too!
Timescale: years 5-15

Phase 3: Build major sub-functions of the human brain, such as the visual, auditory, motor control or language systems. There is considerable information available on high-level pathways and structures, but without the information encoding knowledge from phase 1 it is difficult to progress here.
Timescale: years 10-20

Phase 4: Complete the brain. This involves scaling up the computing resource by the final order of magnitude.
Timescale: years 15-20

These "plans" could easily turn out to be very cautious; all that is required is a major breakthrough in understanding neural encoding and appropriate abstractions and the whole lot could fall into place in half the time I suggest here! Where are the theories that describe the representation and processing of information encoded in partially coherent waves of spikes?

Grand Challenge criteria

Significance.

Is it driven by curiosity about the foundations, applications or limits of basic Science?
In some sense it is driven by curiosity about one of the great mysteries of the human condition: How does our brain work? What is consciousness? Can machines ever be conscious?

Is there a clear criterion for the success or failure of the project after fifteen years?
Total success will be clearly recognisable. Partial success will be less clear. Fifteen years may be a bit aggressive.

Does it promise a revolutionary shift in the accepted paradigm of thinking or practice?
True machine intelligence will be revolutionary; intermediate results much less so.

Does it avoid duplicating evolutionary development of commercial products?
The scale and scope of this challenge puts it outside commercial development roadmaps.

Impact.

Will its promotion as a Grand Challenge contribute to the progress of Science?
In-line with the Cognitive Systems initiative, this activity will promote inter-disciplinary work between computer scientists and neuroscientists to the benefit of both and, even if the work does not achieve total success, its potential to contribute to the treatment of mental disorders and the improvement of human-computer interfaces is highly significant.

Does it have the enthusiastic support of the general scientific community?
It runs a significant risk of being rejected as having an objective perceived as outrageous. Earlier claims of AI, going back to Turing, have also rather tainted the territory that this work proposes to explore. This would need handling with great care.

Does it appeal to the imagination of the general public?
Clearly the general public will see the appeal, though pandering too overtly to this appeal will have a negative impact on the support of the scientific community.

What kind of benefits to science, industry, or society may be expected?
Direct benefits may include more intelligent, more usable machines and better treatments for mental disorders.

Scale.

Does it have international scope?
Clearly. I am in an early stage of trying to work my way into the neural network community, so I can't yet identify the key players in the computing side of this field let alone in the neuroscience space. However, the literature shows that they clearly exist!

How does the project split into sub-tasks or sub-phases, with identifiable goals and criteria, say at five-year intervals?
See outline plans presented above.

What calls does it make for collaboration of research teams with diverse skills?
The neuroscience – computer science collaboration that is required here must be very strong. I hope it is clear that there is a great deal of scope for all sorts of CS folk – mathematical types (information modelling, coding…), engineers (digital and analogue, high-speed comms…), programmers, system designers – the task is very broad.

How can it be promoted by competition between teams with diverse approaches?
The problems are very open and there is still enormous scope for diversity in the approaches taken by different, competing teams.

Timeliness.

When was it first proposed as a challenge? Why has it been so difficult so far?
The general objective was proposed at least as long ago as Turing in the 1950s. It is difficult because many fundamentals are not understood and (in my opinion) because computers have not been, and still are not, powerful enough for the task. This is set to change over the next 20 years (again, in my opinion).

Why is it now expected to be feasible in a ten to fifteen year timescale?
Faster machines and better "wet" science are providing converging knowledge bases and capabilities that, when they overlap sufficiently, should support rapid progress.

What are the first steps? What are the most likely reasons for failure?
The first steps are outlined in the plan above.
The risk factors mainly relate to the unknowns:

- Are neurons really much more subtle than they appear?
- Are quantum effects in microtubules really important, as claimed by Penrose? (If Penrose is right the project is doomed, but in my view his 'proof' is flawed!)
- Can we build machines with sufficient power and connectivity?
- Is the way a brain grows through life vital to model and, if it is, is this tractable?

Relevance to Computer Science

As a final comment, this is clearly as much a grand challenge for brain scientists as it is for computer scientists. The machine I have outlined earlier in this document could be used for a variety of purposes and not just neural modelling. However, it is clear that computer science has a unique opportunity to contribute to this work:

- The issues of information coding, communication and processing are central, and these are likely to require new theories to be developed.
- If you accept the view of the brain as the ultimate information processing device, understanding it *is* computer science (whereas working with human-built artefacts is arguably better described as computer *engineering*).
- The machine requirements are clearly extreme, if my estimates are anywhere near correct. The challenge can therefore drive computer architecture into new territory.

Initially this work should, perhaps, be led by brain scientists. However, there is a great deal in it for computer scientists, and CS as a discipline should throw its resources enthusiastically into the work!