Research plan

General

My research is quite unique in using a computer science approach to understanding the brain. It is distinct from neural modeling which only uses neural representations, and distinct from cognitive modeling in using a computer system design approach. As such this approach is capable of connecting psychological phenomena with brain activity, and in synergizing with cognitive modeling and neural modeling.

I would like to establish a research group, to collaborate with professors in Computer Science, Psychology and Neuroscience, and to teach courses that would train students in this multidisciplinary area. Since I was tenured faculty at London University in the United Kingdom for eleven years, it may only be necessary to commit to a three year tenure clock.

Overview of my work

There is a comprehensive description of me and my research on my website at: http://www.exso.edu/alan_us.html which includes overview papers and an overview set of slides. I also have a rough draft of a book that I have been writing, which gives a more complete idea of my approach, research and directions in which it is headed.

Computer science and the brain

The subject of my research will be separately described by computer scientists and by biologists each in their own terms and from their own perspectives. From a computer science perspective, I seek to describe a computer based on the brain, which will have some of the legendary properties of the brain, its parallelism, intelligence, flexibility and resilience. I am developing an intelligent parallel architecture which can be applied to computation problems of all kinds, including multiagent systems. My work is unique in developing an approach to intelligent agents which is based on the brain. The architecture should have a straightforward implementation in hardware.

From a biological perspective, I seek to bring computer science concepts and methods to bear on the problem of the scientific understanding the brain. The development of a brain model based on computer science will also break through many current intellectual barriers, allowing many issues that are now treated informally to be grounded in a properly defined computational model. It will allow neurological models of disease as well as normal functioning. Cognitive psychology will be now grounded in a neurally plausible architecture. There will be wholesale clarification of concepts concerning the brain, including memory, mental state, stream of consciousness and self. This is the intellectual revolution that computer science has been promising for the last several decades.

Creative research

This type of research is very creative and original, and requires a wide knowledge of several subjects, including neuroanatomy, neurology, cognitive psychology, linguistics, sociology and of course computer science, including artificial intelligence and multiagent systems. It also requires an extended effort over many years since each subject has its own concepts and culture which have to be understood, and then there is the process of creative synthesis which cannot be hurried.

Overall research goals

I am continually revising my computational model. In general, the overall problems I am trying to solve include:

1. Behaviors: (i) how to model problem solving behavior, (ii) how to recognize natural language sentences, (iii) how to model social interaction, and (iv) how to model motivation of the system.

2. Learning: (i) how to learn problem-solving knowledge for "simple" problems using in psychological experiments and clinical diagnosis, and (ii) how to model the development of social relationships.
3. Relationship to neural nets: (i) how to find an equivalent neural representation of the model, (ii) how to model neural development, and (iii) how to model neural pathologies.
4. How to implement the model as a computer using present-day technologies such as FPGA or asynchronous VLSI.
5. How to understand the model theoretically as a computational system.

More detailed research plan
The following detailed plan is intended to estimate what research is achievable in the next three years. Now that I have a basic brain model designed and implemented, good delimited research problems can be defined which can realistically be solved in periods of the order of one year or less.

1. Development of the computational model.
The most important goals I have right now concern getting the model to do problem solving and learning of problem solving strategies.
(i) I am currently focusing on the Tower of Hanoi, since this is the only problem for which very detailed data and psychological theories are available. I've represented in the brain model the different strategies used by humans. I also have been working with an experimental set up for studying human subjects doing this problem and gathering verbal think-aloud protocols and eye saccade data during long experimental sessions of the order of an hour.
(ii) I am using a separate learning module with mechanisms for episodic memory, that is, memory for contexts, and also for learning new rules and plans which are then relocated to appropriate performance modules. Another module will provide routinization, that is, learning of stimulus-response connections, based on frequency of occurrence.
(iii) the resulting performances will be able to be compared with clinical data on solution of Tower of Hanoi by amnesics and by parkinson patients, since these ailments result from disability of episodic memory and routinization respectively.
(iv) I'd like to also do learning to perform a simple collaborative task, such as learning how to greet someone, or to groom if nonhuman. This is for me logically required since the model is naturally suited to joint action of two or more collaborating agents. Another case is scaffolded learning where a caretaker teaches an infant by flexible sharing of roles. I am expecting the joint case to be a straightforward generalization of the single agent learning case.

2. Collaborations in computer science.
I am interested to link up with projects and people where my work can contribute. For example, applications to virtual agents and intelligent robots.
(i) The robot as person. My approach is to emphasize believability through social interaction, mediated by joint plans. I'd like to explore different types of human-computer interaction scenario.
(ii) Interacting with virtual groups. The idea is to achieve human-computer relationships based on social interaction, to investigate the experience of an interacting user in becoming a member of the virtual group.
(iii) with connectionist modelers to understand how system level models can be implemented as neural nets, and how learning mechanisms in neural nets can be abstracted to give learning mechanisms at the system level.
(iv) with AI researchers, this is an approach to parallel AI and to multiagent systems.
(v) with computer theoreticians, can we characterize these systems and understand their properties.
(vi) with computer architecture experts, the implementation of the model in hardware.

(vii) with people interested in modeling, brain architecture can be viewed as a perception-action hiearchy of models.

3. Collaboration in life sciences.
(i) with cognitive psychologists, modeling different cognitive functions for problem solving, episodic memory, visual search, and so on.
(ii) with neurologists. There is a host of other possible applications of my model in neurology and neuroscience, for example, frontal lobe pathologies, action slips, amnesias, dyslexias, obsessive-compulsive disorder, etc.
(iii) with developmental psychologists, for example, modeling attachment, anxiety, aggression, etc.
(iv) with primatologists, for example, modeling a set of monkey troops.

4. My programming language.
I will continue developing, and will distribute my brain programming language, BAD. A key structural principle for my research is the development of a network of collaborators and the provision of a common representation and programming language by means of which scientific knowledge of brain function can be expressed and tested. The current situation is that:
(i) I have developed the first version of a programming language, which I call BAD, for Brain Architecture Description language. BAD basically allows one to construct abstract systems which consist of agents made of sets of interconnected parallel modules.
(ii) I have collaborators who also have graduate students and postdocs who will also work with us.
(iii) I have developed BAD programs - (a) programmed and tested natural language processing mechanisms, (b) programmed motivational mechanisms, and (c) am currently programming problem solving mechanisms.
(iv) I can create stand-alone BAD systems for Mac OS X, Windows 98, Solaris 7 and Linux. This is straightforward using Sicstus Prolog, which also has versions for other platforms such as IRIX, AIX and FreeBSD.
Thus any scientific theory of brain function will be expressed precisely in BAD, and will also be tested against experimental data obtained by my collaborators.

5. Industrial collaboration.
I plan to continue my collaboration with the Fujitsu Network Agents Research Group in San Jose, California, headed by Frank McCabe, who is active in FIPA. Our research concerns language design for network agents, and also involves Keith Clark at Imperial College, London University.

Alan Bond
2006-1-10