Grand Challenge 5:
Architecture of Brain and Mind

(Mindware implemented in wetware-inspired hardware ???)

Aaron Sloman
(With much help from others)

Introduction

What is the most powerful and most complicated computer on the planet? Wrong! it's not a machine you can buy for millions of dollars, it's the amazing system that we all own, the few kilos of grey and white mush in our heads.....

Biological information processing systems produced by evolution still far outstrip both our understanding and our practical achievements.

In order to reduce both the theoretical and the practical gap, we need to look closely at two of the most impressive products of evolution: human *brains* and human *minds* – and attempt to construct a combined vision of how they work:

Brains, the contents of our skulls, are composed of extraordinarily intricate physical structures, performing many tasks in parallel at many scales, from individual molecules to large collections of cooperating neurones or chemical transport systems.

Minds are more abstract and contain ideas, percepts, thoughts, feelings, memories, mathematical knowledge, motor control skills and other things that cannot be seen by opening up skulls. Yet their existence and their power to do things depends on all the 'wetware' components that make up brains.

Can we ever replicate all that functionality? At least we can improve our understanding, by developing new deep theories, tested in working models.

The challenge

We aim to combine insights from several disciplines: from *neuroscientists* studying brain mechanisms, from *psychologists, linguists, social scientists, ethologists and philosophers* studying what minds can and cannot do, and from *researchers in computer science and AI* developing techniques for *implementing* many kinds of abstract mechanisms and processes in concrete physical machines.

We should be able to implement many kinds of abstract mental processes and mechanisms in physical systems, just as abstract machines like spelling checkers, email systems, databases and AI planners and problem-solvers are already implemented in electronic computers and networks.

This research will address two age-old quests: the attempt to understand what we are, and the attempt to make artificial human-like systems, whether entertaining toys, surrogate humans to work in inhospitable environments or intelligent robot helpers for the aged and the infirm.

How will it be done?

Several mutually-informing tasks will be pursued in parallel:

Task 1 Build computational models of brain function, at various levels of abstraction, designed to support as many as possible of the higher level functions identified in other tasks.

Task 2 Codify typical human capabilities, for instance those shared by young children, including perceptual, motor, communicative, emotional and learning capabilities and use them to specify a succession of increasingly ambitious design goals for a fully functioning (partially) human-like system.

Task 3 Develop a new theory of the kinds of *architectures* capable of combining all the many information-processing mechanisms operating at different levels of abstraction, and test the theory by designing and implementing a succession of increasingly sophisticated working models, each version adding more detail.

Analysing what those models can and cannot do and why, will feed back information to the other two tasks.

A possible 15 to 20 year target providing an extremely demanding test of the scientific advances might be demonstration of a robot with some of the general intelligence of a young child, able to learn to navigate a typical home and perform a subset of domestic tasks, including some collaborative and communicative tasks. Unlike current robots it should know what it is doing and why, and be able to discuss alternatives. Linguistic skills should include understanding and discussing simple narratives about things that can happen in its world, and their implications.

Achieving all this will require major scientific advances in the aforementioned disciplines, and could provide the foundation for a variety of practical applications in many industries, in unmanned space exploration, in education, and in the ever-growing problem of caring for disabled or blind persons wishing to lead an active life without being dependent on human helpers. Perhaps many people reading this will welcome such a helper one day.

International collaboration

Several international research programmes, including “Cognitive Systems” initiatives in Europe and in the USA are now supporting related research: international collaboration is essential for success in such a demanding project.