Biological inspiration: short-circuiting evolution?

Dr. M. Neal, Intelligent Systems Lab., University of Wales, Aberystwyth
mjn@aber.ac.uk

Endowing engineered systems with the ability to maintain their mode of operation in dramatically changing circumstances and during component failure and degradation is a key goal of A.I. research and robotics in particular. The best examples of complicated systems that have these properties are biological organisms of dramatically varying scales and types. Attempting to imitate the evolutionary approaches that have lead to biological success have provided very limited success in artificial systems, and there is some evidence that "biological-style" evolution requires very large numbers of organisms and long periods of time to succeed in the initial production of complicated organisms.

A quicker route to the emulation of biological (organism-level) homeostasis might be through the virtual dissection of organisms in order to identify key information processing systems, mechanisms and representations. Fortunately our colleagues in biology began this process several thousand years ago and have identified a relatively small number (arguably three) of systems that are key to the maintenance of homeostasis in higher organisms.

These systems are the neural system, the immune system and the endocrine system. Due to the nature of evolution these three systems are highly interdependent and cannot readily be disentangled. This has not stopped biological theories (some very successful) from considering them as independent units. The three systems have quite different functions that act at quite different time-scales and in quite different ways. The neural system provides almost instantaneous control of effectors both within and without the body, the endocrine system relies on usually slower mechanisms such as diffusion and operates mostly at the scale of several seconds to several months, and the immune system maintains its defence over periods of hours to years.

Robotics: the obvious test-bed

The three systems provide a set of perception, action and memory mechanisms which result in the homeostatic properties of biological organisms. Given that biologists have a relatively good understanding of the three systems and their operation it is tempting to hope that some of the processing systems, mechanisms and representations might be useful for building systems that are rather good at surviving the trials and tribulations of the world. Robotics provides the obvious test-bed for such systems and is the focus of much work in the intelligent systems laboratory at Aberystwyth. The use of artificial neuro-endocrine controllers and the emulation of some simple endocrine functions is ongoing and shows promise especially for arbitration between behaviours in neural networks. The integration of artificial immune systems is still in its infancy, but their potential both as growth moderators and maintenance subsystems is now being explored.

Ubiquitous and mobile systems: more appropriate?

Ubiquitous and mobile computing represents a change in emphasis from desktop and other static computing solutions to systems which operate in harmony with the potentially unpredictable and dynamic environments that their users inhabit. The ability to tolerate and exploit the fluid nature of their use and abuse is fundamental to successful development, deployment and robustness of such systems. These requirements have much in common with autonomous robotic systems and with living organisms. The vogue for biologically inspired approaches to developing and controlling robots has lead to some significant changes in the way that robots are viewed, and many of the lessons learnt in robotics are transferable to the mobile and ubiquitous domains. There is strong evidence that the properties of ubiquitous and mobile systems are far more likely to benefit from biologically inspired approaches than robotic and other applications.

There are a number of properties of ubiquitous and mobile systems that distinguish them from more traditional computer/communication systems. These are fairly widely accepted to include: openness, reconfigurability, scalability, heterogeneity and variability of connectivity (in pattern, bandwidth and quality).

Openness in this context implies that the system is open to additions and deletions of components and connections without any real limit. This has a number of implications and consequences, probably the most important of which is that systems cannot rely upon algorithms and technologies that will not scale indefinitely. For example transmission of data around an (in principle) unbounded network of agents is a very different proposition to transmission of data around networks with a fixed maximum size. Even the desire to uniquely identify agents in the network must be questioned, as must the use of communication strategies other than those which are essentially independent of context.

Reconfigurability implies that the agents in the system are liable to be moved physically and/or logically within the network, as well as assigned new functions, priorities and statuses. These types of change are likely to occur frequently and probably without warning or notification. This has a number of effects upon the types of algorithm and technology that are appropriate and in the extreme case may dictate the use of effectively "stateless" operation.

Scalability of such systems must be essentially unbounded as mentioned above. This is essential for truly "open" systems and sets a number of constraints upon both efficiency and potential optimality of the systems produced. For example it is not possible in potentially unbounded networks to achieve complete connectivity. It is worth noting that the strict sense of unbounded need not apply: systems are in effect unbounded as far as the agents are concerned when the local environment and resources available to an agent are insufficient to generate or store "addresses" for a large portion of the network with which it can communicate.

Heterogeneity of the components within such systems is often mooted as essential to the successful growth and operation of truly ubiquitous systems. This can be tackled in a number of ways which includes relying on a base level of standard capability and negotiation to make good use of resources at each agent. It is also possible to use a small number of standard mechanisms without explicit negotiation. It can be argued that the latter is preferable, as it requires less (or no) bi-directional exchange of information between agents.

Variability of connectivity Clearly as agents move or are moved around the environment their pattern of connectivity will vary dramatically. Any communication protocols and algorithms dependent upon them must be capable of dealing intelligently with a wide range of connectivity. Isolated groups of agents may wish to become "aware" of their isolation and must also avoid frequently repeated requests for unobtainable resources around such disconnected parts of the network of agents.

The mechanisms offered by the immune, neural and endocrine systems seem to offer a number of features that match up well with some of these properties. The growth of organisms and proliferation of cells demonstrates a degree of openness, and the ability to differentiate and move cells within the body demonstrates reconfigurability. Vertebrates from the scale of the blue whale to the mouse use very similar immune, neural and endocrine systems and seem to function very well. The heterogeneity of cells within the organism is both tolerated and essential for its continuing survival. Finally the movement of cells around the body is exploited by the use of receptors and localised hormone release to perform specific functions in specific areas.

Stupid not to try?

It could be argued that *not* to try to copy a number of the promising looking functions and mechanisms used by the higher organisms would be a gross dereliction of duty by computer scientists. The long timescales required for successful artificial evolution may ultimately be the only answer, but the potential for bootstrapping with simple artificial immune, neural and endocrine systems might provide an enormous head start: it took about 3 billion years to evolve a multi-cellular organism from scratch, but only about 450 million years to get from the first vertebrates to humans and the other complex creatures that exist today. It seems that providing well designed building blocks for developmental processes to run with (whether evolutionary or otherwise) yields great returns.

Designing a set of components that are based around biological principles and using them both "as is" and as building blocks for developmental systems looks like an interesting strategy.