

Artificial general intelligence: an organism and level based position statement

Leslie S. SMITH¹

Department of Computing Science and Mathematics, University of Stirling, UK

Abstract. Do simple animals, even single-celled animals, display general intelligence? If we take this to mean that they can cope with their environment, even when it is dynamic, the answer would have to be yes. What is it that gives them this easy capability, yet makes it hard for us to design systems which have it? We suggest that it is from the non-determinism which arises from the lowest levels of their operation, and from the way in which the different levels of their operation interact. We propose techniques whereby these might be incorporated into artificial systems.

Keywords. artificial general intelligence, brain model, paramecium, level interaction

Introduction

There are many views of what should be described as artificial general intelligence. General intelligence is not a well-defined concept: but then, neither is intelligence, nor are many other general concepts, such as a life. But this lack of definition does not mean that they have no meaning: rather that they occupy a diffuse area, rather than a single point: they are cluster concepts. Better then, to ask what some of the characteristics of these concepts are, and to attempt to approach them, rather than trying to define them into a corner then reconstruct them. In what follows, I have chosen two particular aspects of organisms that display what I consider general intelligence, and I have attempted to delineate a possible program of research based on these.

1. The low level organism view

Much of what has been researched in artificial intelligence comes from a view that humans have intelligence, and other animals do not. I suggest that this humano-centric view of the world has as much place in general intelligence research as an Earth-centred view of the Universe has in astronomy. General intelligence is about keeping an organism alive and thriving in a complex and ever-changing environment. It differs from purely reac-

¹Corresponding Author: Leslie S. Smith. Department of Computing Science and Mathematics, University of Stirling, Stirling FK9 4LA, Scotland, UK; E-mail: l.s.smith@cs.stir.ac.uk

tive behavior in that behaviors are not necessarily entirely contingent on the organisms' senses. Certainly, there are aspects of intelligence found in humans that are not found in other animals, most notably complex language; however, there are certainly aspects of general intelligence visible in the behavior of animals.

One is then led to ask what the simplest animal to display any form of general intelligence is, and whether we can learn from it, and build something based on it?

Consider the behaviour of a very simple single celled organism: a paramecium. This animal has no nervous system as such, yet has successfully survived in freshwater environments for many millions of years. The animal's single cell has a membrane surrounding it, and this is covered in small hair-like appendages (cilia). These allow it to move. It eats by endocytosis (incorporating food particles into its cytoplasm), and excretes by exocytosis. Endocytosis is mediated through protein sensors on the membrane surface. The membrane is not uniform: there is a gullet-like area for eating, and an anus-like area for excretion. The animal can move towards food. Reproduction can be either sexual or asexual. Does it display general intelligence? Or is it purely reactive? What aspects of its behavior display intelligence?

Consider a slightly more complex organism, one with a rudimentary nervous system. Actually, its visible behavior is likely to be very similar to that of the paramecium: the primary difference is that it has specialized cells for different processes, such as movement, digestion, and control. Yet in the end, eating and excretion are essentially orchestrated endo- and exocytosis. Control of the animal is achieved also through the action of the nerve cells on these other cells, and this is achieved through neurotransmitters sent through exocytosis, and received through receptors and ion channels, which then alter cell behaviour. Does it display intelligence or is it purely reactive? Or perhaps the animal displays general intelligence, but each cell is purely reactive?

Artificial intelligence designers often consider complexity to be part of the nature of intelligence. They build their systems out of logical entities (such as gates, or threshold logic elements), using some mixture of hardware and software to develop complex logical systems. Given appropriate input, and appropriate interpretation of their output, they can perform sophisticated tasks. Each element is utterly deterministic, and could be considered reactive. But does the overall system display general intelligence?

The examples above, two living and one not, may or may not be considered to display general intelligence depending on the standpoint of the onlooker. (I suggest that there is a range of "intelligences" between purely reactive systems and truly intelligent ones, but that attempting to delineate exactly where each example lies on a line connecting them is not a useful exercise.) However one thing that is clear is that the living entities are very different from the non-living one. Is this difference one purely of implementation (living cells as against logic gates), or are there deeper differences?

It has been suggested that life is ultimately based on logic (by which is usually meant that living entities are based on the same logical (and physical) rules that pervade the rest of the universe, rather than there being 'special sciences' that govern living entities). Yet this does not imply that life is necessarily based on the same two-valued logic that pervades digital computing. One of the great strengths of two-valued logic is also one of its problems: it is utterly deterministic: repeating the same two-valued logical steps on the same data will always give the same result. This goes against what we see in living organisms. They always show wide diversity of behaviour (and, indeed, this diversity

seems to exist almost no matter what level one examines the systems at). This leads to two different questions:

1. Where does this diversity originate from, and
2. Can we implement this diversity using the two-valued logic systems?

Even the simplest living organisms do not display two-valued logical behaviour. Non-determinism is built in from the very bottom level. Consider a piece of the cell membrane of a paramecium in water, which has touched a possible food particle. In order for the food particle to be ingested, a reaction needs to take place between the food particle and some sort of receptor on the surface of the cell. The molecules of both the food particle and the receptor are large (probably proteins, or fatty acids). The reaction occurs when the reactive parts of the molecules of the reactants become close enough to each other [1]. Both reactants are in motion because of their temperature, and because of Brownian motion induced by the water molecules. The likelihood of a reaction occurring is directly related to the likelihood of the reactive surfaces meeting each other before the same random movement causes the food particle to move away from the cell. This whole process is essentially nondeterministic.

Virtually all the processes that occur at the cell membrane are of a related form: for example, ion channels which alter shape due to variations in the charge patterns surrounding them are actually continuously in motion, and their shape is essentially best described as a stochastic system [2]. Thinking of ion channel behaviour as being like transistor behaviour (as Mead suggests in his 1989 book [3]) can therefore be slightly misleading. We conclude that there is diversity in response in biological systems even at the lowest level.

We suggest that the diversity seen at higher levels in biological systems arises from this lower level diversity, in the same way as the lack of diversity in digital electronic system behaviour has its roots in the absolute determinism of the lowest level operation of these electronic systems.

But how important is this diversity? We suggest that this diversity in response lies at the root of successful adaptivity: by responding differently to essentially identical environmental events, and then ‘rewarding’ responses which have the best outcomes by making them more likely to occur in future, an organism can adapt to its environment rapidly and effectively. This is not a new idea: Campbell [4] places blind adaptivity at the root of all learning, from this lowest level all the way up to scientific knowledge, and coined the term *evolutionary epistemology* to describe it. The adaptivity need not be entirely blind: more recent learning techniques such as reinforcement learning and TD learning [5] suggest that previous experience can be used to guide learning and hence adaptive evolution of behavior. We consider this evolutionary capability to be at the root of general intelligence, both in simple and higher organisms.

2. Interlinking levels

Brains operate at many different physiological levels simultaneously. There are numerous different species of ion channels, some voltage sensitive, some not, all opening and closing simultaneously; there are billions of synapses some excitatory, some inhibitory, some shunting, some transmitting, some not; there are billions of neurons, some spiking

some not; there are cortical columns processing information, and communicating with other columns; there are brain areas each with their (partial) specialisation: the whole lot comes together to produce the behaviour of that brain. All of these levels of viewing what is occurring in an animal brain are simultaneously valid. Clearly, they are not independent of each other. One might attempt a similar view of a modern computer system; there are billions of transistors, some on, and some off, billions of gates, some with logic 1 and some with logic 0 outputs; there are objects and functions all running (quasi-) simultaneously, all coming together to produce the behaviour of that computer system.

There is, however, a major difference between the two. Digital computer systems are designed and built hierarchically, whereas biological brains are not. This difference shows itself in the ways in which these different levels of operation interact. In a computer, the gates are built out of transistors, and, although the output of a gate is the output of a transistor (and the input to a gate is the input to a transistor) it is safe to ignore transistor operation entirely, as everything about the gate operation can be abstracted away from this lower level. Similarly, the entire operation of all the gates etc. in the hardware can be abstracted into the instruction set (and interrupt structure etc.), so that the software can be considered entirely independently of the hardware. The same is true for layers of software: each interface between levels allows the lower levels to be entirely abstracted away.

Often neural systems have been viewed in this way as well. Ionic channels and neuromodulators subserve entities like synapses, and synapses and the collection of charge on dendrites subserve the generation of spikes. Neurons in a cortical column are there in order to generate the behaviour of the column. In other words, the brain is often viewed as if it has been designed from an engineering perspective. But nature and evolution are not bound by engineering design laws! In fact there is considerable interaction between levels. For example, neurotransmitter released from one synapse will leak to another one. Further, small molecules which act as neuromodulators (such as Nitrous Oxide, NO) can diffuse over relatively wide areas, causing non-localised interactions which will cross levels. Research suggests that these non-local links can rapidly reconfigure neural systems [6].

Even with engineered systems, using genetic algorithms to optimize performance can result in level crossing being utilised in a problem solution. A good example of this is in Adrian Thompson's work [7] in which an FPGA was optimised to solve a problem, using a fitness function calculated directly from the operation of the chip. The result was a solution which used all sorts of cross-talk between entities which are supposed to be entirely independent directly in the system operation. Another example may be found in the work of the robotist Mark Tilden, who utilises the mechanical properties of the materials of his robots in the robot controller. Clearly, biological systems are quite free to develop such solutions to problems (and don't have to worry about making their systems comprehensible to engineers!).

How does this interlinking of levels contribute to general intelligence? It can provide direct (and therefore fast) interaction between different parts of the brain, permitting rapid reconfiguration. This could permit selection between different mechanisms that might otherwise be viewed as purely reactive. On the other hand it might also cause inappropriate crosstalk (which is why it is normally intentionally omitted from engineered systems).

3. Synthesizing general intelligence based on simple animals and linking levels: towards a program

We have stressed the importance of non-determinism at the lowest levels making itself visible at higher levels (up to eventual behavior). Could this non-determinism be replicated by noise in an otherwise deterministic system? And if so, would the form of the noise matter? It appears that the critical issue is the capability for selecting different behaviors, and rewarding those that are most appropriate. We have also stressed the way in which the different levels of operation are interlinked in biological systems. However, it is less clear how this contributes towards general intelligence. One possible answer is that this results in another form of variation in behavior because of the way in cross-level linkages alter behavior in (relatively) unexpected ways. It is not simply noise, because the effects will not be random, but (at least to some extent) repeatable. We note that it can also allow rapid changes to be made to the nature of processing.

Replicating this type of general intelligence is not straightforward. If we wish to take an Artificial Life approach, we clearly need to have a rich environment, as well as critters exhibiting multiple levels inhabiting this environment. The levels should not be precisely hierarchically arranged, but able to affect each other. These critters need to be non-deterministic, preferably with this non-determinism being implemented right at the lowest of levels. This should enable different behaviors in response to similar events, and adaptation towards the most appropriate behavior whether blindly or using (e.g.) reinforcement learning. An additional key may be to add a genetic algorithm system on top of this, enabling modification (and possibly optimization) of the internal structure, including the ways in which the levels are able to influence each other. This could permit fast reconfiguration of the system in the light of environmental cues.

Acknowledgements

Thanks to Catherine Breslin for suggesting the Paramecium as a very simple animal with interesting behavior, to Karla Parussel for interesting discussions, to the UK EPSRC Grand Challenges in Microelectronic Design Network for inspiration, and to the anonymous referees for historical context.

References

- [1] F.T. Hong, A multi-disciplinary survey of biocomputing: Part 1: molecular and cellular aspects, in *Information Processing and Living Systems* V. B. Bajić and T. W. Tan, Eds., Imperial College Press, London, 2005, 1–139.
- [2] T.F. Weiss, *Cellular Biophysics*, MIT Press, 1996.
- [3] C. Mead, *Analog VLSI and Neural Systems*, Addison Wesley, 1989.
- [4] D.T. Campbell, Evolutionary Epistemology, in *The Philosophy of Karl Popper*, P.A. Schlipp, Ed, Open Court, Illinois, 1974.
- [5] R.S. Sutton and A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 1998.
- [6] K. Parussel, A bottom-up approach to emulating emotions using neuromodulation in agents, University of Stirling Ph.D. thesis, <http://hdl.handle.net/1893/113>, November 2006
- [7] A. Thompson, I. Harvey and P. Husbands, Unconstrained evolution and hard consequences, in *Towards Evolvable Hardware: The evolutionary engineering approach*, E. Sanchez and M Tomassini, Eds, Springer Verlag, LNCS 1062, 1996, 136-165