Seth Bullock

Institute for Complex Systems Simulation, and
School of Electronics and Computer Science
University of Southampton
Southampton, SO16 1BJ, UK

sgb@ecs.soton.ac.uk

1. In what sense do you find it meaningful to talk about "living technology?"

It strikes me that this question has two related interpretations: what is the significance of the research activity organising around the term 'living technology', and what issues are raised by the notion of living technology as a concept or category. First, I think that there are many different valid responses to these questions, and that mine are largely coloured by the fact that my work is far removed from the actual attempt to synthesise living technologies. For me, the significance of that enterprise lies in the challenge that it presents to the design, engineering, control and management communities – a radically new perspective on what technology can be, on how technologies can be built, and on how we are expected to interact with them. It seems to me that the perspective on systems engineering presented by living technology research represents a very timely new paradigm that is sorely needed.

As for deciding what things should and should not count as exemplars or instances of living technology: Since we are dealing with a notion that has yet to develop into a consensually defined mature category, we will not always know whether particular systems should be deemed "in" or "out" (for two examples, see question 4, below). However, this is nothing to be ashamed of or worried about. I started out in cognitive psychology and artificial intelligence, and got used to working on a set of problems that are also hard to pin down: perception, memory, reasoning, intelligence, mind, and consciousness. Consensually agreed mature definitions that operationalise these terms rather than simply delineating a set of phenomena to be explained just don't exist yet. Achieving them is the end of the story, not something to be established before the story can begin.

So my short answer to the opening question would be "in the widest sense possible." I tend to subscribe to a generous reading of "living technology" because I am interested in the boundary cases, whether they are man-made or naturally occurring, organic or inorganic, or even socio-technological. My guess is that most contributors to this volume are involved in the attempt to *bring about living technology*, which might understandably mean that they are more interested than me in tightening up the definition ahead of time. Possibly they feel that some synthetic biotic machines either already fall into the category of living technology or will soon. Most may feel that simple, tiny, lifelike systems built from non-biological materials will eventually join them. After all, to feel otherwise is tantamount to the assumption that "life's core properties" will forever be possessed only by systems made from already living materials. The same practitioners may concede that, technically, there are also naturally occurring, *non-synthetic* examples of living technology, where whole animals or plants are employed as machines either by us (e.g., bees, yeast) or even each other (e.g., parasitic flukes that directly influence the brain of their ant host in order to control its movement); but it is likely that these are not their primary concern. More controversial are socio-technological systems that may include living entities as parts, but are also claimed to exhibit "life's core properties" at the level of the whole system, e.g., hospitals, governments, universities, online communities, etc. For some, my guess is that this may be a step too far, but personally I would welcome them in too.

I don't feel the need to strongly defend the claim that any of these kinds of systems are *really* examples of living technology, or that they *really* are not, since my interest in the field stems precisely from the way in which it will shed light on the nature of the "core properties of living systems" that lie at the heart of both the living technology enterprise and my own research interests.

2. How does your research relate to living technology, and why were you initially drawn to do this work?

Well, I see two core reasons for working on living technology: first, to synthesize useful new technologies that solve important problems for society, and second, to better understand what the core properties of living systems are by generating new examples of technologies that purport to exhibit these properties. The former is an increasingly significant activity, but my primary interest is in the latter, and I see the living technology field as an important new way of stimulating this project.

As I mentioned, my background is in psychology and artificial intelligence (AI). Like others involved in what became known as "nouveau AI," my personal trajectory has been ever downwards. Within cognitive science, I climbed down from the heady heights of my undergraduate degree studying human logic, language, rationality and reason, and writing AI algorithms designed in the image of clinical, Kasparov-like ratiocination, in order to pursue a PhD in Sussex's Evolutionary Robotics lab. There I was confronted by the crude, clumsy effectiveness of robots inspired by tiny insects, driven by toy neurons, automatically wired together by witless evolutionary algorithms: AI from the bottom up. In parallel, I was working with experimental economists. There, I dutifully worked through the prescriptive Bayesian analyses of rational choice theory, but then plumbed the depths of real people's inability to string two rational decisions together in a risky choice experiment. As my PhD took a biological turn, I coded up Dawkins' eerily simple algorithm for abstracting the raw power of natural selection freed from the messy particularities of terrestrial biochemistry, and then struggled to understand the manifest gap between the resultant insipid simulated evolution and the full-blown complexity of the real thing.

In each case, I was left feeling that the realities of life, mind and society were far removed from the clean algorithms, theorems and calculations of classical AI, neo-classical economics and neo-classical Darwinian biology. Their abstracted, formal accounts were attractively clean, elegant and economical, but this "neat" approach was fundamentally misleading. What was missing seemed obvious to me at the time. I was embedded (and embodied and situated) in a research environment suffused with a Heideggerian conviction that all behaviour happens *in the world*, that it is "embedded, embodied, situated," that Simon's emphasis on "satisficing" rather than "optimising" was the appropriate attitude, and that Kauffman's "order for free" could be the bedrock for adaptive behaviour.

There is not space here to do much justice to the ideas that I've just name-checked, but I will indicate roughly what brings them together for me, which is a thoroughgoing naturalisation of life and mind. Martin Heidegger's philosophy challenges a still pervasive Cartesian mindset that takes thinking and reasoning to be essentially general, rational, logical, computational and (tacitly) magical, in that they are assumed to take place in an idealised mental realm of symbols, logico-syntactic structure-sensitive functions, etc. For Heidegger, cognition, thought, language, logic and ultimately *being*, must all be recognised as situated in the world no matter how angelic and transcendental they might appear. In taking this turn he opens the door for an entirely new enactivist cognitive science that prioritises a concern with the coupling between creatures and their environments as much as their sensing and acting or beliefs and desires, and does this for creatures irrespective of whether they can be imagined to possess an internal "language of thought". (See Wheeler, 1995, 2005, for more on the link between Heidegger, cognition and artificial life.)

Thirty years after the publication of Heidegger's (1927) *Being and Time* (but a decade before it's first translation into English), Herb Simon (1957) articulated his idea of *bounded rationality*. He operationalized a notion of grounded cognition, pointing out that it is not in the interest of an organism to reach optimal solutions to the problems posed by its environment if the time or resources consumed in doing so were prohibitive. Decision making is always *situated* in an evolving environment that is specific to the decision maker. The dynamics of this *Umwelt* are critical to understanding its behaviour, cognitive or otherwise. (See Goodie et al., 1999, for more on the threads of bounded rationality that run through cognitive science, biology and economics.)

If the Heidegerrian attitude of Sussex's evolutionary robotics group was appropriate for cognitive behaviour, it was also surely true of life more generally. Just as the locus of intelligent behaviour was not an algorithm running inside a head, the locus of 'vitality' was not to be found in a free-floating evolutionary algorithm comprising heritable variation plus differential reproduction. Stuart Kauffman's (1993) *Origins of Order* set out a series of results that demonstrated how useful, complex biological organisation might arise in the absence of natural selection and that this ordered organisation was more than just the backdrop or raw materials for evolution. Rather, facts about terrestrial physics and chemistry had a substantive part to play in how adaptive processes had played out on Earth and would continue to do so.

The challenge, then, was to understand what aspects of this dirty, frustrating *embeddedness* were involved critically in underpinning the phenomena of life, mind and, ultimately, society. Modelling was the route that I took to exploring these questions, building simulations in which to explore the role of "logistics" (who does what to whom, where and when) in effective adaptive biological organisation. How does the spatial embedding of a population influence natural selection (Clark and Bullock, 2007)? How does a population's environment structure influence the costs of irrationality (Bullock and Todd, 1999)? How do logistics influence the evolvability of signalling (Noble et al., 2001)?

Along with the construction of real-world living technologies, simulation models of this kind are an example of the *synthetic methodology*. They are attempts to understand systems not by taking them apart and exploring the properties of the pieces in the reductionist hope that systems' secrets are located in the properties of their atoms, but by synthesising systems bottom-up. By assembling a system's parts together in a computer or a petri dish or on a workbench, the aim is to explore the relationship between the organisation of these parts and the properties of the whole that they form.

It seems to me incontrovertible that brute facts about the nature of the implementation level will be key to understanding the behaviour of complex adaptive systems and that we cannot therefore continue to abstract these facts away in neat, tidy theorem-friendly formalisations. Self-organisation, thermodynamics, spatial embedding, etc. are key to unlocking the secrets of how brains, cells, organisms, and communities work when they do. Attempts to synthesise and understand living technologies directly and necessarily confront these implementation issues. My hope is that in doing so they will shed new light on the problem of what underpins living behaviour in much the same way that examining patients with brain damage sheds light on regular cognitive behaviour. Note that even in cases where such brain damage results in abnormal cognition, or in behaviour which is not "cognitive" at all, they are valuable to the study of cognition, since exploring where and how a system breaks down is a good way of finding out how it was working in the first place. Analogously, whether or not we end up deciding that a particular new technology is "truly living" is to a large extent irrelevant to how much it might teach us about what it is about the world that brings systems to life.

3. How is living technology related to overlapping or nearby research areas, such as nanotechnology, molecular biology, cloning and stem cell research, genetic engineering and synthetic biology? How is it related to social and technological systems such as social networks or information networks, such as the World Wide Web, cell phone networks and electronic banking networks?

From my perspective, the most interesting boundaries are between the synthesis and study of living technologies and our understanding of what makes certain biological, social and socio-technological organisations *vital*. What makes one hospital *click,* while another consumes massive amounts of energy and people in delivering a mediocre service? What makes one community resilient to insults and shocks, while another is fragile or moribund? How could buildings or transport networks become systems that organise themselves creatively in a symbiotic relationship with their users?

In some sense, artificial synthetic cells are just a set of exciting new analogies for existing complex adaptive systems – just a new set of vocabulary with which to gloss the same ideas and questions. In my experience it is easy to undervalue such new ways of speaking, since they lead to new ways of thinking. But in this case there is a second and potentially more significant cross-disciplinary contribution to be made by living technologists: the struggles, failures and successes in living technology labs will provide real experimental data on the nature of autopoiesis.

Unlike my simulation models, where the abundant degrees of freedom offered by a modern computer ensure that almost anything can be made to hold within a digital world, synthesis in the real world is massively constrained by physics and chemistry. It is clear that what works must work despite these constraints. From my perspective, the fascinating possibility is that living technologies will work not despite these constraints, but *because* of them – i.e., the constraints of physics and chemistry will be *enabling* for life (Bullock and Buckley, 2009).

4. What do you think are the most important open research questions about living technology, and how you think they should be pursued?

There are obviously many complicated problems that remain to be solved in realising examples of living technology in the lab, and other contributors to this volume are far better qualified to assess them than myself. So I will limit myself to raising one issue that is perhaps further down the line but will need to be addressed before autonomous, adaptive, free-living technologies are employed in earnest. How will these technologies adapt during their use? How will we ensure that this adaptation is benign? How could we in fact exploit the adaptive power of living technologies rather than seek to attenuate it?

Here, I would suggest that there may be value in studying naturally occurring living technologies – living technologies that were not deliberately engineered by people, but arose spontaneously in nature and have persisted and adapted ever since.
Two (candidate) examples of non-anthropogenic living technology (that have perhaps not been recognised as such) have occupied me in some of my work: termite mounds (Ladley and Bullock, 2005) and biological signalling systems (Bullock and Cliff, 1997).

> Termite mounds are clear examples of technology in that they are carefully constructed homes, they are just not our homes, not homes that are planned or engineered, but instead self-organise. If we consider the mound and the termite colony that built it together as a single system, then there are grounds for considering it as an example of *living* technology (in much the same way that the Internet is sometimes taken as such). The structure itself is sophisticated and multi-functional, serving as shelter and protection from weather and predators, but also as a functionally segregated environment with specialised areas in which to care for offspring, bury the dead, raise crops, etc. The mound is adaptive and homeostatic, maintaining critical parameters such as chamber temperature and humidity via self-regulatory air flows, and integrity via self-repair and ultimately self-reproduction. However, to grant this system the status of living technology is complicated by the role of the living beings that produce it – aren't they the only living aspect of the system, with the mound simply being a product of their activity? While this is a legitimate perspective, it may be that we will not fully understand the nature of termite mounds, how they adapt and their stigmergic relationship with their inhabitants, until we take a perspective that recognises them as autopoietic living systems in their own right.

> Across the natural world signalling systems are rife: from the messenger molecules employed by the simplest cells to the syntactically structured utterances of the most complex primates. Signalling occurs for a multitude of different reasons and takes a huge variety of different forms: from songs and calls to smells, gestures, postures, and patterns. All this sharing of information takes place despite the competition to survive and reproduce. How do natural signalling systems arise, persist, and adapt? Why do signals take the form that they do? To some extent these questions can be answered by treating signalling systems as living technology. Signals are tools for transmitting information and their form reflects this function. The temporal structure of some bird song, for example, is adapted in a way that allows it to resist degradation by reverberation in forested habitats and thereby travel further. Simon Kirby's group has shown that properties of human language may arise as a consequence of different language variants competing to successfully be transmitted through the bottleneck of language learning in a human infant (Kirby, 2002). While signalling systems are not *engineered* technology, and they are not even *substantial* in the sense of a piece of physical hardware, nevertheless it may be the case that a full understanding of them requires us to consider them as evolving adaptive living technologies.

In both cases, we can view these techno-social systems as comprising an organismal population that produces, supports and maintains a technological superstructure. However, it is also possible to see causal processes that run in the reverse direction. Technology clearly impacts directly on the organisms that use it, sheltering them or informing them, but the selective processes that these organisms are subject to also slowly shape them, fitting them to an environment in which the technology is a dominant feature. In some sense then, a natural language comes to direct the way that its speakers think and see the world, and a termite mound *builds and maintains itself* by steering the short-sighted behaviours of its termite slaves.

If we are to enter a design space populated by living technologies of the kind described above, or other kinds, we will need to understand how these co-adaptive relationships unfold, and be able to steer them when necessary.

5. What do you consider to be the most interesting and important <u>human or societal implications</u> of research and development in living technology?

In a world that cannot continue to sustain us in our current resource-hungry mode of existence, we must quickly learn how to build and use systems that *self-organise* to deliver the quality of life that we need. Living technology is one route to such a future. If civilisation's job is, very broadly, to lower entropy in our local environment at the expense of increased entropy out in space somewhere, then technologies driven by self-organisation, if designed correctly, will offer us solutions that require minimum energy in order to deliver the ordered world that we need. By contrast, our traditional approach to technology is analogous to Maxwell's demon, carefully creating order by deliberately, manually shifting stuff from one place to another, but at the expense of consuming the massive amounts of energy and time required to measure and control the world, and resist and recover from the natural processes that are going on around us. Living systems demonstrate that we can do better. A few genes steer a process of self-organisation to create a creature, perhaps a person. A sub-set of these genes steer processes of self-organisation to create a learning brain. A further sub-set steer processes of self-organisation in order to create a living cell operating at efficiencies unheard of in human engineering, in environments that are massively challenging, and at collaborative scales that are currently unthinkable.

References

Bullock, S., & Buckley, C. L. (2009). Embracing the tyranny of distance: Space as an enabling constraint. *Technoetic Arts*, 7(2), 141-152.

Bullock, S., & Cliff, D. (1997). The role of 'hidden preferences' in the artificial co-evolution of symmetrical signals. *Proceedings of the Royal Society of London, Series B*, 264, 505-511.

Bullock, S., & Todd, P. M. (1999). Made to measure: Ecological rationality in structured environments. *Minds and Machines*, 9(4), 497-541.

Clark, B., & Bullock, S. (2007). Shedding light on plant competition: Modelling the influence of plant morphology on light capture (and vice versa). *Journal of Theoretical Biology*, 244(2), 208-217.

Goodie, A. S., Ortmann, A., Davis, J., Bullock, S. & Werner, G. M. (1999). Demons versus heuristics in artificial intelligence, behavioral ecology, and economics. In G. Gigerenzer and P. M. Todd (Eds.), *Simple heuristics that make us smart* (pp. 327-355). Oxford: Oxford University Press.

Heidegger, M. (1927/1962). *Being and Time*. Trans. by John Macquarrie & Edward Robinson, London: SCM Press.

Kauffman, S. (1993). Origins of order: Self-organization and selection in evolution. Oxford: Oxford University Press.

Kirby, S. (2002). Natural language from artificial life. *Artificial Life*, 8(2), 185-215.

Ladley, D., & Bullock, S. (2005). The role of logistic constraints on termite construction of chambers and tunnels. *Journal of Theoretical Biology*, 234, 551-564.

Noble, J., Di Paolo, E. A., & Bullock, S. (2001). Adaptive factors in the evolution of signalling systems. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (pp. 53-78). Heidelberg: Springer.

Simon, H. (1957). A behavioral model of rational choice. In *Models of man, social and rational: Mathematical essays on rational human behavior in a social setting*. New York: Wiley.

Wheeler, M. (1995). Escaping from the Cartesian mind-set: Heidegger and artificial life. In F. Moran, A. Moreno, J. J. Merelo, & P. Chacon (Eds.), *Advances in artificial life: Proceedings of the third European conference on artificial life* (pp. 65-76). Heidelberg: Springer.

Wheeler, M. (2005). *Reconstructing the Cognitive World: The Next Step*. Cambridge, MA: MIT Press.