Nichiversal Computation: Specifying the Limits of Bio-Inspiration

Seth Bullock
School of Computing, University of Leeds, Leeds, LS2 9JT, UK
+44 (0)113 343 5322, seth@comp.leeds.ac.uk

Nature has served as the inspiration for some of the most generic and widely used tools in our possession. For instance, the development of clothing, habitation, weapons, etc., by early man was probably influenced by his experience with furs, hides, nests, lairs, horns and claws. Within the much more recent context of computing, we routinely employ metaphors derived from nature: bugs, trees, inheritance, killing, spawning, sleeping, etc. In terms of novel computational paradigms, biology has also been a rich source of ideas. Artificial neural networks, evolutionary algorithms, swarm intelligence, artificial immune systems, etc., have all arisen as a result of our growing appreciation of the sophisticated computational abilities of naturally occurring biological systems.

There is an undeniable attraction to be found in mechanisms that, despite comprising simple elements governed by simple rules, can manifest powerful, organised, problem-solving behaviour. This attraction is only compounded by the possibility that these mechanisms are, in some powerful sense, *general purpose*. For instance, a common class of continuous-time recurrent artificial neural network can be shown to approximate any dynamical system to an arbitrary degree of accuracy (Funahashi & Nakamura, 1993); a swarm of artificial insects can implement a general-purpose optimisation algorithm (Bonabeau, Dorigo, & Théraulaz, 2000a); a cellular automaton is capable of universal computation (Cook, 2004); certain (hotly debated) classes of evolutionary search algorithms may be suited to solving the poorly defined but extremely general class of "real-world" problems (e.g., Harvey, 2001); wasp and termite behaviour might be idealised to deliver general-purpose construction algorithms for self-organising architectures (Bonabeau, Guérin, Sners, Kuntz, & Théraulaz, 2000b; Howsman, O'Neil, & Craft, 2004); artificial immune systems could offer the ability to efficiently classify arbitrary classes of patterns (e.g., Tarakanov, Skormin, & Sokolova, 2003).

The generality of these idealised, bio-inspired systems suggests that they might enjoy very wide applicability, and holds out the possibility that, give or take some simple parameter setting, we might not need to gain a deep understanding of a particular problem before we are able to generate a solution to it. However, things have not turned out to be this straightforward. Successfully applying a bio-inspired approach to a real, non-trivial problem is often surprisingly difficult. Tailoring bio-inspired algorithms to achieve what we require of them often involves a painstaking and opaque process of tinkering, reworking, and removal/addition of algorithmic components, in addition to mere parameterisation. Moreover, when success is achieved in a particular instance it is often difficult to see how it can be generalised to a wider class of problem. This process is often seriously impeded by the difficulty that we face in understanding how small alterations to a decentralised, complex adaptive system impact on the global behaviour of the whole system. If bio-inspired approaches have given rise to powerful, general-purpose algorithms, why are they not more generally applicable and successful?[1]

One approach to addressing this question commences with the observation that real biological systems are only ever general-purpose *accidentally*. No biological species, organism, organ, trait or mechanism has ever evolved to be general-purpose, i.e., to apply to a class of problems that is wider than the set of problems actually encountered by its ancestors so far. It is true that some are more or less specialised than others, but natural selection is not in the business of fashioning devices that solve future problems or potential problems, only actual historical ones.

For instance, the behavioural mechanisms that termites use to construct their amazing mounds (Bruinsma, 1979) were not evolved for construction, *per se*, but for constructing termite mounds, specifically (Ladley & Bullock, 2004, in press). Our immune system has not

[1] Perhaps the only current exception to this argument is to be found in the mainstream success of some kinds of artificial neural network (ANN) for pattern recognition tasks—algorithms that by now bear little relation to the biological neural systems that ultimately inspired them, being more closely related to forms of multiple non-linear regression. Furthermore, despite having become a somewhat established tool, there remains a degree of black magic in getting an ANN approach to work efficiently.

evolved to classify arbitrary patterns, but to deal with the particular kinds of pathogen to which we have historically been exposed. Even the human brain, indisputably the most awesome problem-solving mechanism that we know of, is not a general purpose cognitive machine. It is specialised to undertake particular cognitive tasks (language learning, face recognition, social cognition, etc.). It is not organised to solve any problem or deal with every cognitive challenge (witness the large literature on our cognitive shortcomings, e.g., Kahnemann, Slovic, & Tversky, 1982). Rather our brain exhibits properties that allow it to successfully tackle the reproductively significant cognitive problems that faced our evolutionary ancestors on the African savannah.

Stated more generally, biological devices are shaped by natural selection such that they tend to be well suited to the challenges posed by their Environment of Evolutionary Adaptedness (or EEA, see Foley, 1997). This "environment" is actually the sum total of the selection pressures that have been brought to bear on a device's lineage (weighted by recency). It is the finite set of reproductive problems that a particular contemporary biological device's ancestors solved in order that this device (rather than competing forms) currently exists. The EEA is thus similar to the notion of a biological *niche*, in that the "design" of a biological device can be understood as an attempt to satisfy the demands, pressures, and challenges that characterise this niche. From an alternative perspective, one can expect biological devices to function successfully only under Normal conditions: the conditions that the device's ancestors tended to find themselves in historically (Millikan, 1984, 1993). Outwith such conditions, the performance of an evolved device may be suboptimal, or even pathological (e.g., some forms of human obesity may result from some of our evolved devices operating in a modern environment featuring many abNormal foodstuffs).

What this means is that the biological systems that inspire novel computational paradigms are likely to be suited to particular tasks.[2] Even when (idealised abstractions of) these mechanisms are *capable* of exhibiting a very general class of behaviour, they will not do so *uniformly*—they will be more suited to some tasks than others. While continuous-time recurrent neural networks are capable of exhibiting *arbitrary* dynamics (given enough nodes), it is still true that a certain kind of dynamic is *characteristic* of such networks, i.e., this class of device does exhibit a *generic* behaviour (Beer, 1995). Attempting to find or construct networks that exhibit dynamics very different from this generic behaviour is difficult. Similarly, even if termite construction behaviours can be idealised such that they are, in theory, capable of generating arbitrary structures (Howsman et al., 2004), it will remain the case that some classes of structure are more readily buildable by such systems. In order to configure such a system to construct architectures that are *uncharacteristic*, one faces a very difficult reverse engineering task that cannot typically be solved by hand and is often difficult to solve using some kind of search algorithm (for difficulties in evolving and hand-designing artificial termite systems for arbitrary construction tasks, see Ladley, Bullock, & Prangnell, in preparation) .

[2]It might be suggested that this argument does not apply to evolutionary algorithms, since these were inspired by the evolutionary *process* rather than any *product* of evolution. The evolutionary process is already an abstraction and is, by definition, not tailored to particular tasks. However, even so, the evolutionary process cannot be claimed to be a general purpose search algorithm—heritable variation coupled with competition for some scare resource(s) is not guaranteed to optimise. Moreover, evolutionary algorithms are typically influenced not only by the evolutionary process in the abstract, but also by the particular mechanisms by which it is instantiated in nature (genetic encoding, sexual reproduction, coevolution, etc). These mechanisms *are* the product of natural selection, and are thus subject to the argument outlined here.

Guidelines

In one sense this argument boils down to a well-known fact: every tool is good for some things and not so good for others. However, the ramifications of this fact imply alterations to the typical working methodology of bio-inspired computing researchers. The rest of this paper attempts to spell out these methodological implications via four guidelines and, subsequently, a brief example.

1. Embrace the "nichiversal" (literally niche-facing) nature of bio-inspired computation.

First and foremost, we should adopt the working assumption that the limited, task-specific, generic behaviour of a bio-inspired system is what is important, rather than its potential for universality. After all, demonstrating that a class of mechanism is capable of universal computation tells us more about the nature of computation than the class of mechanism. Characterising the "niche" of a class of bio-inspired system is a challenging, but critically important goal.

2. Accept that multiple idealisations of a biological mechanism/organisation/process can coexist.

Within a particular domain of bio-inspired computing, there often appears to be competition between different flavours of system. The evolutionary algorithms literature offers many clear examples. At one level, genetic programming, genetic algorithms, evolutionary strategies, etc., "compete" (with each other and with alternative search and optimisation algorithms) to demonstrate their ability to solve hard optimisation problems. At a lower level, different genetic encodings, genetic operators, selection schemes, multi-population set-ups, etc., also compete to outperform each other. There is rarely an attempt to specify the ways in which these different flavours of algorithm relate to one another, or to specify how one might decide between them when attempting to solve a particular problem.

3. Take note of negative results, carefully examined.

In the context of (2), above, one can see why negative results are unpopular: "I've invented a new type of a genetic algorithm—here's a number of ways in which it is outperformed by existing genetic algorithms". Of course, if we understand that no evolutionary algorithm will outperform *all* others on *all* classes of problem, then it is precisely this type of negative result that can be valuable, when carefully analysed. Zaera, Cliff, and Bruten (1996), for example, present a failed attempt to evolve realistic flocking behaviour as an indicator of what makes constructing a fitness function hard or easy. Unfortunately, this type of research is rarely undertaken and remains difficult to get published when it is.

4. Attend to the limits of natural biological mechanisms/organisations/processes *in situ*.

Biologists cannot completely and accurately characterise a biological mechanism's EEA or its Normal conditions for functioning. However, biologists often know something about the character of a mechanism's niche. In particular, where a mechanism varies across different populations, there is scope for explaining these differences as resulting from the different selection pressures that these populations have been subjected to. This information can be useful in determining what one might expect a bio-inspired approach to be good for. However, gathering it involves serious engagement with the relevant biological community and their literature, which is time-consuming and difficult work.

A Brief Example

Within neuroscience there is an increasing realisation that the traditional abstraction of neural systems as essentially networks of units interacting via neurotransmission is unsatisfactory since it neglects the role of the chemical substrate within which this electrical activity is embedded (Katz, 1999). The chemicals involved are implicated in numerous kinds of adaptive behaviour, from triggering plasticity and learning, reconfiguring neural circuits, and balancing gross levels of activity, to switching between multiple modes of behaviour. From this research is emerging a new "liquid brain" perspective on real neural networks (Changeux, 1993).

GasNets are a class of recurrent artificial neural network inspired by this line of neuroscience research (Husbands, Smith, Jakobi, & O'Shea, 1998). In addition to a relatively standard explicit network of neurons communicating via idealised neurotransmission, these ANNs employ an idealised type of chemical signalling in the form of simulated neuromodulators. GasNets have been artificially evolved successfully for a range of tasks including the control of autonomous mobile robots. In fact, they appear to be particularly suited to this kind of application (*op. cit.*).

What is important for the purposes of this paper is that this particular flavour of bio-inspired robot control architecture should not be regarded as in a flat competition with alternative schemes, e.g., continuous-time recurrent neural networks or CTRNNs (Beer & Gallagher, 1992; Beer, 1995). Instead, the pertinent question should be: in what circumstances are GasNets, or CTRNNs, the appropriate architecture to employ?

Answering this question requires more than collecting a large number of examples of one paradigm outperforming another. Rather, a combination of carefully analysed successes (Smith, Husbands, Philippides, & O.Shea, 2002), basic research into the original biological mechanisms (Philippides, Husbands, & O'Shea, 2000), new conceptual frameworks (Philippides, Husbands, Smith, & O'Shea, 2002), and fundamental modelling work (e.g., Buckley, Bullock, & Cohen, 2004) is necessary in order to reveal why, for example, ANN schemes that involve analogues of neuromodulatory chemicals are able to exhibit robust, evolvable, adaptive behaviour over multiple timescales (temporal adaptivity). It is only through this interdisciplinary activity that the GasNet niche can be characterised.

Summary

Only once we accept that in general biological devices, processes and organisations are properly viewed as specific to their particular niches, and develop theoretical accounts of what it is that individual biological devices, processes or organisations are good at—what it is that they have been "designed" to achieve—are we in a position to exploit idealisations of them efficiently.

Acknowledgements Thanks to Chris Buckley, Dan Ladley, and the Leeds Biosystems group for wide-ranging discussions of these issues. Responsibility for the opinions stated here rests solely with myself.

References

Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, *3*, 469–509.

Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, *1*, 91–122.

Bonabeau, E. W., Dorigo, M., & Théraulaz, G. (2000a). Inspiration for optimization from social insect behaviour. *Nature*, *406*, 39–42.

Bonabeau, E. W., Guérin, S., Sners, D., Kuntz, P., & Théraulaz, G. (2000b). Three-dimensional architectures grown by simple 'stigmergic' agents. *BioSystems*, *56*, 13–32.

Bruinsma, O. H. (1979). *An analysis of building behaviour of the termite* Macrotermes subhyalinas *(Rambur)*. Ph.D. thesis, Landbouwhogeschool, Wageningen, Netherlands.

Buckley, C., Bullock, S., & Cohen, N. (2004). Toward a dynamical systems analysis of neuromodulation. In Schaal, S., Ijspeert, A. J., Vijayakumar, A. B. S., Hallam, J., & Meyer, J.-A. (Eds.), *Eighth International Conference on Simulation of Adaptive Behavior*, pp. 334–343. MIT Press, Cambridge, MA.

Changeux, J. (1993). Chemical signaling in the brain. *Scientific American*, *8*, 58.

Cook, M. (2004). Universality in elementary cellular automata. *Complex Systems*, *15*, 1–40.

Foley, R. A. (1997). The adaptive legacy of human evolution: A search for the environment of evolutionary adaptedness. *Evolutionary Anthropology*, *4*, 194–203.

Funahashi, K., & Nakamura, Y. (1993). Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks*, *6*, 801–806.

Harvey, I. (2001). Artificial evolution: A continuing SAGA. In Gomi, T. (Ed.), *Evolutionary Robotics: From Intelligent Robots to Artificial Life*, pp. 94–109. Springer-Verlag, Heidelberg.

Howsman, T. G., O'Neil, D., & Craft, M. A. (2004). Stigmergic cooperative multi-robot control architecture. In Pollock, J., Bedau, M., Husbands, P., Ikegami, T., & Watson, R. A. (Eds.), *Ninth International Conference on Artificial Life*, pp. 88–93. MIT Press, Cambridge, MA.

Husbands, P., Smith, T. M. C., Jakobi, N., & O'Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, *10*, 185–210.

Kahnemann, D., Slovic, P., & Tversky, A. (1982). *Judgement Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge.

Katz, P. S. (1999). *Beyond Neurotransmission: Neuromodulation and its Importance for Information Processing*. Oxford University Press, Oxford.

Ladley, D., & Bullock, S. (2004). Logistic constraints on 3D termite construction. In Dorigo, M., Birattari, M., Blum, C., Gambardella, L. M., Mondada, F., & Stützle, T. (Eds.), *Fourth International Workshop on Ant Colony Optimization and Swarm Intelligence*, pp. 178–189. Springer-Verlag, Berlin.

Ladley, D., & Bullock, S. (in press). The role of logistic constraints in termite construction of chambers and tunnels. *Journal of Theoretical Biology*.

Ladley, D., Bullock, S., & Prangnell, D. (in preparation). Constraints and limitations on stigmergic structural evolution. Unpublished manuscript.

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. MIT Press, Cambridge, MA.

Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. MIT Press, Cambridge, MA.

Philippides, A. O., Husbands, P., & O'Shea, M. (2000). Four-dimensional neuronal signaling by nitric oxide: A computational analysis. *Journal of Neuroscience*, *20*, 1199–1207.

Philippides, A. O., Husbands, P., Smith, T. M. C., & O'Shea, M. (2002). Fast and loose: Biologically inspired couplings. In Standish, R. K., Bedau, M. A., & Abbass, H. A. (Eds.), *Eighth International Conference on Artificial Life*, pp. 292–301. MIT Press, Cambridge, MA.

Smith, T. M. C., Husbands, P., Philippides, A. O., & O.Shea, M. (2002). Neuronal plasticity and temporal adaptivity: GasNet robot control. *Adaptive Behavior*, *10*, 161–184.

Tarakanov, A. O., Skormin, V. A., & Sokolova, S. P. (2003). *Immunocomputing: Principles and Applications*. Springer-Verlag, Heidelberg.

Zaera, N., Cliff, D., & Bruten, J. (1996). (Not) evolving collective behaviors in synthetic fish. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., & Wilson, S. (Eds.), *Fourth International Conference on Simulation of Adaptive Behavior*, pp. 635–644. MIT Press, Cambridge, MA.