Post-publication version of "Aaron Sloman Travels Forward to – Virtual Machinery and Evolution of Mind" (Part 3) pp 849-856
In *Alan Turing - His Work and Impact*, eds S. B. Cooper & J. van Leeuwen, Elsevier 2013.

Aaron Sloman Travels Forward to
VIRTUAL MACHINERY AND EVOLUTION OF MIND
Meta-Morphogenesis: Evolution of Information-Processing Machinery
(PART 3)*

September 8, 2013

1 Introduction: Types of emergence

In 1936, Turing made major contributions to our understanding of certain types of emergence, by showing how a Turing machine can be set up so as to generate large numbers of very simple processes that cumulatively produce qualitatively new large scale results, e.g. TM operations producing results related to results of human mathematical reasoning. Later work by Turing and others led to electronic computing machinery enabling a small set of very simple operations to produce very many kinds of novel, useful, complex, and qualitatively varied results – now impacting on many aspects of everyday life. Universal TMs showed that both the construction of *mechanisms* and the construction of *things on which mechanisms operate* can in some cases be handled in a uniform way, by having mechanisms that can construct and manipulate mechanisms (e.g. self-modifying computer programs). A similar theme was implicit in his 1950 paper. As far as I know, Turing's last work on micro-macro emergence was the 1952 paper on morphogenesis, explaining how micro-interactions in physicochemical structures could account for global transformations from a fertilised egg to an animal or plant, within a single organism.

All those ways in which complex configurations of simple structures and processes can have qualitatively new features are examples of micro-macro relationships that can be labelled as 'emergent' Cohen and Stewart (1994)).

It is now clear that physical and chemical mechanisms involved in biological reproduction can, like computational machinery, include specifications not only for (partially) controlled construction of new physical mechanisms (where some of the control comes from the environment) but also for production of new construction specifications, and new mechanisms for using such specifications, as well as development and learning mechanisms for growing and modifying already functioning machinery, and mechanisms for detecting damage and producing repairs. The combined products of all these mechanisms, ecosystems and

*This is Part 3 of a sequence of linked papers in this Volume. *Alan Turing – His Work and Impact*. Part 1 is in Part I of the volume and Part 2 is in Part III of the volume. This part, concerned with Meta-Morphogenesis, is in Part IV. The ideas presented here are developed further in `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html` socio-economic-political systems, together constitute the most complex known example of emergence.

Much research on evolution and development has focused on production of new physical forms and new physical behaviours. We also need to understand micro-macro relationships involving creation and use of *new forms of information-processing*, without which much of the complexity could not have arisen. There is much knowledge and expertise about information processing in computer science, software engineering and more generally computer systems engineering, but relatively little understanding of biological 'meta-morphogenesis' (MM), the information processing mechanisms involved in producing biological novelty, including new forms of information-processing.[1].

2 Layered computational emergence

Computing systems developers create new micro-macro relationships, using a set of micro components: types of hardware or software structure, a small collection of possible processes, and ways of ways of combining processes and structures using syntactic composition methods. The resulting new macro components (e.g. electronic circuits, or computer programs) have complex and varied structures, and can support yet more new types of complex and varied processes, some of which provide 'platforms' for constructing further layers of complexity. As argued in Part 1 (this volume), the functions, states and processes in the new layers often cannot be defined in the language of physics and chemistry, or digital circuits. In that sense, although the new layers may be fully implemented in the old ones, they are not reduced to them. E.g. the concepts 'win' and 'lose', required for describing a running chess program, are not definable in the language of physics. So the chess machine is implemented in, but not reducible to physical machinery.

Achieving such micro-macro bridges requires understanding the deep and unobvious generative potential of the initial fragments and their possible relationships. The full potential was unobvious in the early days of computing, but new programming languages, new development environments, new operating systems, new re-usable packages and, above all, new problems, have continually revealed new, more complex, achievable targets. The complexity we now take for granted was achievable only through *layered* development of tools and techniques, often depending on use of earlier layers. Similar constraints apply to biological evolutionary and developmental trajectories. Many biological mechanisms, structures and functions that developed recently could not have occurred in earlier times, despite availability of all the *physical* materials, because many small intermediate changes were required in order to produce the *infrastructure* for newer more complex mechanisms.

New layers of computing machinery were in part a response to *external* pressures from application domains, with which new computing systems had to interact, e.g. using sensors (e.g. cameras, pressure sensors, etc.), effectors (e.g. grippers, wheels, paint sprayers, etc.), or network connections. Similar, still unidentified, environmental pressures led to new emergent mechanisms and processes in biological evolution. Other pressures can come from *internal* requirements to improve speed, reliability, energy efficiency, easy of monitoring, ease of debugging and ease of extension.

[1]For an answer to 'What is information?' see Sloman (2011).

The physical universe produces objects of varying complexity, from subatomic particles through molecules, planets, galaxies and beyond. Large lumps of solid or liquid matter can result from materials being brought together concurrently. But many intermediate-sized structures of great complexity, including organic molecules and organisms of many kinds, require special construction mechanisms, or intermediate scale components, that are not always directly available even when the physical materials are available. Such complex systems need to be assembled over time using precisely controlled selections from among physically and chemically possible alternatives. For example, there was no way the matter on this planet several billion years ago could have immediately reorganised itself into an oak tree or an orangutan.

A Tornado could not assemble a 747 airliner from a junkyard full of the required parts. Assembling an airliner requires not only prior assembly of smaller parts, but also machinery for producing the intermediate structures, and maintaining them in relationships required for subsequent operations. Biological evolution also requires intermediate stages including intermediate mechanisms of reproduction and development. Intermediate stages in evolution require increasingly complex forms of information processing, so biological information processing mechanisms, like computer systems engineering, must have involved many intermediate forms of information processing. Compare how later stages of a mathematical proof depend on earlier stages, preventing simultaneous discovery of all parts of the proof.

Successive information-processing mechanisms must have had successively more complex physical components, forms of representation, ontologies, algorithms, architectures, and functions, especially information processing functions relating to the environment. We need to understand the mechanisms of meta-morphogenesis.

Some thinkers assume that there must be a single master designer controlling such processes of assembly of complex living structures from inanimate matter. But development of software engineering sophistication over the last six decades did not require some super-engineer controlling the whole process. There was only a large collection of successively discovered or created bootstrapping processes in a multitude of forms of competition and co-operation partly driven by a plethora of new more complex goals that became visible as horizons receded. Humans mostly stumbled across more and more complex ways in which previous achievements could be extended. Natural selection had much in common with this, except that there were no designers detecting new targets – until species emerged with sufficient intelligence to engage in mate selection and other selective breeding activities.

3 Meta-morphogenesis and biological complexity

A feature of growth of complexity is that as new mechanisms are developed some of them transform, and simplify, opportunities for subsequent developments, as illustrated for individual cognitive development in Chappell and Sloman (2007). Related points were made in Cohen and Stewart (1994). New mechanisms, new forms of representation and new architectures, can sometimes be combined to provide new 'platforms' bringing entire new spaces within (relatively) easy reach. Examples in the history of computing include new operating systems, new programming languages, new interfacing protocols, new networking technologies, new constraints and requirements from users, including requirements for reliability, modifiability, security, ease of learning, ease of use, etc. We don't know all the new pressures that influenced developments of biological information-processing mechanisms, both in evolution and in individual development, though we can guess some of them.

Evolution of biological information processing was much slower, and did not need goal-direction, only random 'implicit' search (implicit because there were no explicit goals, only opportunities that allowed some changes to be relatively advantageous). Identifying those opportunities and the evolutionary changes they influenced is hard. An example is the difference between organisms in an amorphous chemical soup and organisms whose environment has distinct enduring parts with different properties (e.g. providing different, persistent, nutrients and dangers, in different locations). Only the second type could benefit from mechanisms for acquiring and storing information about those enduring structures. Such information would necessarily have to be built up piecemeal over time. If the organism had visual mechanisms, it could rapidly take in information about complex structures at different distances. If it only had tactile/haptic sensors, the information would have to be acquired in much smaller doses with more movements required. .

Some computing developments, such as a new notation, or a new ontology (e.g. for types of communication, or types of event handler, or types of data-structure), or creation of a new type of operating system, can provide a 'platform' supporting a very wide range of further developments. There were probably also many different kinds of platform-producing transitions in biological evolution, e.g. development of new means of locomotion, new sensors, new manipulators, new forms of learning. Some of these were changes in physical form or structure or forms of motion, or types of connectivity, whereas others were changes concerned with information processing. Smith and Szathmáry (1995) discussed changes in forms of *communication*, but there must have been many more transitions in information processing capabilities and mechanisms, some discussed in Sloman (2008).

When a new multi-function platform is developed, searches starting from the new platform can (relatively) quickly reach results that would previously have involved intractable search spaces. After learning a powerful language like Prolog, a programmer can often quickly produce programs that would have been very difficult to express using earlier languages. New high level languages add new opportunities for rapid advances. Likewise, as Dawkins and others have pointed out, some biological developments, including new forms of information processing, could, in principle, dramatically shorten time spans required for subsequent developments, even though there is no goal-directed design going on. Even random search (though not a tornado?) can benefit from a billion-fold reduction in size of a search space.

4 Less blind evolutionary transitions

Some animals can formulate explicit goals and preferences, and select actions accordingly. The evolution of that capability can provide a basis for selecting actions that influence reproductive processes, for example selecting mates, or favouring some offspring over others, e.g. bigger, stronger or more creative offspring. When animals acquire such cognitive capabilities, such choices can be used, explicitly or as a side-effect of other choices, to influence selective breeding, in ways that may be as effective as explicit selective breeding of another species, e.g. domestic cattle or hunting dogs. Which types of selective breeding a species is capable of will depend on which features they are capable of recognising. If all they can distinguish among prospective mates or their offspring is size or patterns of motion, that could speed up evolution of physical strength and prowess. If they can distinguish differences in information processing capabilities that could lead to kinds of selective breeding of kinds of intelligence. (N.B. I am not endorsing eugenics.)

These are examples of ways in which production of a new platform can transform something impossible into something possible, overcoming limitations of pre-existing mechanisms of composition. That includes bringing yet more new platforms within reach, as has happened repeatedly in computer systems engineering when new tools allowed the construction of even more powerful tools, e.g. using each new generation of processor design to help with production of subsequent designs.

A major research task is to identify evolutionary and developmental transitions that facilitate new subsequent evolutionary and developmental transitions. Innate learning capabilities produced at a late stage in evolution may include important pre-compiled partial information about the environment that facilitates specific kinds of learning about that sort of environment. (Compare Chomsky's claims about human language learning, and Karmiloff-Smith (1992).) Special-purpose evolved learning systems may, on this planet, outstrip all *totally general*, domain-neutral, mechanisms of learning or evolution sought investigated by some researchers. In his 1950 paper, Turing suggested that 'blank slate' learning would be possible, which I find surprising. In contrast, McCarthy (2008) argues convincingly that evolution produced new, specialised, learning capabilities, required for human learning in a human lifetime, in certain sorts of changing 3-D environments.

5 From morphogenesis to meta-morphogenesis

In the same general spirit as Turing's paper on morphogenesis, I have tried to sketch a rudimentary theory of 'meta-morphogenesis' showing how kinds of development that are possible in a complex system can change dramatically after new 'platforms' (for evolution or development) are produced by pre-existing mechanisms.

Biological evolution is constantly confronted with environmental changes some of which reduce or remove, or in some cases enhance, the usefulness of previously developed systems, while blocking some opportunities for change and opening up new opportunities. In that sense the environment (our planet) is something like a very capricious teacher guiding a pupil. Initially the 'teacher' could change only physical aspects of the environment, through climate changes, earthquakes, volcanic eruptions, asteroid collisions, solar changes, and a host of local changes in chemical soups and terrain features. Later, the teacher itself was transformed by products of biological evolution, including global changes in the composition of the atmosphere, seas, lakes, and the land-water distribution influenced by evolution of microbes that transformed the matter with which they interacted.

As more complex organisms evolved, they formed increasingly significant parts of the environment for other organisms, of the same or different types, providing passive or active food (e.g. prey trying to escape being caught), new materials for use in various forms of construction (e.g. building shelters, protective clothing, or tools) active predators, mates, and competitors for food, territory, or even mates. As a species evolved new physical forms and new information-processing mechanisms, those new developments could make possible new developments that were previously out of reach, e.g. modification of a control mechanism might allow legs that had originally evolved for locomotion to be used for digging, fighting or manipulation. As new control subsystems evolved, they could have produced new opportunities for system architectures containing those subsystems to develop, allowing old competences to be combined in novel ways. So developments in the 'learner' can be seen as developments in the 'teacher', the environment. Two concepts used in educational theory, Vygotsky's concept of *Zone of proximal development* (ZPD) and Bruner's notion of 'scaffolding' can therefore be generalised to evolution. Evolutionary and other changes can modify the ZPD of an existing species and provide scaffolding that encourages or supports new evolutionary developments. Further details would contribute to a theory of meta-morphogenesis.

6 Evolved information processing: beyond Gibson

Almost all organisms are control systems, using stored energy (sometimes externally supplemented, e.g. when birds use up-draughts) to produce internal and external changes that serve their needs. The control details depend on information acquired through sensors of various kinds, at various times. So organisms are 'informed control systems'. Information available, and also the control possibilities, vary enormously: from the simplest micro-organisms, mostly responding passively in chemical soups, to animals with articulated bodies and multiple sensors, who were capable of performing many different sorts of action, and requiring increasingly complex information processing to notice opportunities, to select goals, to select ways of achieving goals, to carry out those selected actions, to deal with unexpected details of the environment detected during execution, and to learn both from experiences of performing successful and unsuccessful actions and from observation of other things occurring in the environment. A full account of these transitions requires several generalisations of James Gibson's notion of 'affordance' Sloman (2009).

We need to extend not only Turing's work but also the work of Maynard Smith and Szathmáry, on transitions in evolution, to include detailed investigation of transitions in types of *information-processing*. Transitions in forms of communication are often noted, for instance the development in humans of communication using syntactic structures, but there are far more biological processes involving information than communication (internal or external). The need for them will be obvious to experienced designers of intelligent, autonomous robots. The information processing requirements include interpreting sensory information, controlling sensors, learning, forming plans, dealing with conflicts, evaluating options, and many more (Sloman, 2006). Many of the requirements are not obvious; so researchers often notice only a tiny subset and therefore underestimate the problems to be solved – as has happened repeatedly in the history of AI. An extreme example is assuming that the function of animal vision is to provide geometric information about the surfaces in view (Marr, 1982), ignoring the functions concerned with detecting affordances, interpreting communications, and continuous control of actions (Gibson, 1979).

A particularly pernicious type of myopia occurs in research in robotics, biology, psychology, neuroscience and philosophy that focuses entirely on the continuous or discrete online interactions between an organism (or robot) and the immediate environment, ignoring requirements for planning, explaining and reasoning about things going on in other locations, and past and possible future events discussed in Sloman (2006, 2009). Overcoming this myopia can be very difficult, but progress is possible if instead of focusing attention on single organisms or particular designs, we examine *spaces of possibilities*: possible sets of requirements for organisms and robots, and possible sets of design features capable of meeting those requirements. For example, noticing an organism or individual failing to do something may draw attention to the problem of explaining how others succeed – a requirement that may previously have gone unnoticed. A special case of this is the work of Jean Piaget on the many partial or missing competences of young children, which help to draw attention to the hidden complexities in the competences of (normal) adults. Likewise events following brain damage can expose unobvious aspects of normal cognition.

Simply observing or dissecting organisms will not reveal their information-processing: we also need to engage in detailed analysis of differences between environments and morphologies, showing how, as environments change, a succession of increasingly complex demands and opportunities can make possible cumulative changes not only in physical structure, size, strength, and behaviours, but also in the kinds of information available, the kinds of information processing mechanisms, and the uses of such information.

We also need to identify different requirements for belief-like and desire-like states that inform behaviours as discussed (incompletely) in Sloman et al. (2005). Changes in the environment can affect the goals that are essential or useful for an organism to pursue. In some cases goals remain the same, but the information processing and behaviours required to achieve them change: for example if drought or competition makes a certain kind of fruit more scarce, requiring the animals to travel further, climb higher up trees, or physically engage with competitors after the same food. In other cases, changes in the environment may produce new constraints or new opportunities, making it useful to acquire new types of goal. For example, a new kind of food may become available, and if food is scarce the species that acquire desires to find and consume the new food will benefit. However, the physical actions required to obtain and consume that food (e.g. breaking open a shell) may benefit from new forms of control, thereby allowing yet another genetic change to be useful.

Even when environment and sensorimotor morphology remain the same, changes in the *mode of processing* of the information available may provide benefits, e.g.,

- acquiring new ways of learning correlations between sensorimotor signals,
- acquiring new actions that provide or refine information about the environment, e.g. approaching objects, viewing them from new locations, rotating them, acting on them by prodding, pushing, squeezing, twisting, pulling apart, etc. (Gibson, 1979),
- developing a new ontology extending old semantic contents (e.g. developing an exosomatic ontology of 3-D structures and processes that exist independently of being sensed, or developing an ontology that allows information about the past or the future or states of affairs out of sight to be represented),
- developing new explanatory theories about the materials, structures, processes, and causal interactions in the environment,

- developing ways of exploring future possible actions to find good plans before initiating behaviours (Craik, 1943; Sloman, 2006),
- developing new meta-semantic competences that allow the information processing of other organisms to be taken into account (e.g. prey, predators, conspecifics, offspring, mates).

7 Monitoring and controlling virtual machinery

Some changes produce new opportunities for informed control of monitoring and other processes, including operations on the intermediate virtual machine structures in perceptual sub-systems. Parts 1 and 2 of this chapter point out that such biological developments involving virtual machinery can explain philosophically puzzling features of animal (including human) minds, such as the existence of 'qualia'. This can enhance our understanding of requirements for future machines rivalling biological intelligence. We need to explore the space of *possible* minds, and the different requirements different sorts of minds need to satisfy – a very difficult task, since many of the requirements are unobvious. In particular, I hope it is now clear that not all the requirements for embodied organisms (and future robots) are concerned with real-time, continuous, online interactions with the immediate environment, except for very simple organisms with very simple sensory-motor capabilities.

Turing was interested in evolution and epigenesis and made pioneering suggestions regarding morphogenesis – differentiation of cells to form diverse body parts during development. As far as I know, he did not do any work on how a genome can produce *behavioural competences* of the complete organism, including behaviours with complex conditional structures so that what is done depends on internal and external sensory information, nor internal behaviours that extend or modify previously developed information processing architectures, as discussed in Karmiloff-Smith (1992).

Even if we can understand in the abstract that evolution produces behavioural competences by selecting brain mechanisms that provide those competences, explaining how it actually works raises many deep problems, especially where the competences are not themselves behavioural.

The human-produced mechanisms for constructing more and more complex computing systems from a relatively small set of relatively simple types of components are all examples of 'emergence' of qualitatively new large-scale structures and processes from combinations of much simpler building blocks.[2] Perhaps a deeper study of the evolution of tools, techniques, concepts and theories for designing complex systems in the last half century will stimulate new conjectures about the evolution of natural information processing systems, including those that build themselves only partly on the basis of an inherited specification. I suspect that people who predict imminent singularities underestimate the extent of our ignorance about what evolution has achieved, and some of the difficulties of replicating it using known mechanisms. Most biological meta-morphogenesis remains undetected.

[2]Part 1 introduced a distinction between implementation and reduction, where a Running Virtual Machine (RVM) can be fully implemented in physical machinery (PM) even though the concepts required to describe the processes in the RVM cannot be defined in terms of concepts of physics. In that case the RVM is implemented in but not reduced to physical machinery. Part 2 showed how this might account for some of the special properties of mental phenomena such as qualia.

References

Chappell, J. M., & Sloman, A. (2007). Natural and artificial meta-configured altricial information-processing systems. *International Journal of Unconventional Computing*, *3*(3), 211–239.

Cohen, J., & Stewart, I. (1994). *The collapse of chaos.* New York: Penguin Books.

Craik, K. (1943). *The nature of explanation.* London, New York: Cambridge University Press.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin.

Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science.* MIT Press.

Marr, D. (1982). *Vision.* San Francisco: W.H.Freeman.

McCarthy, J. (2008). The well-designed child. *Artificial Intelligence*, *172*(18), 2003-2014.

Sloman, A. (2006, May). *Requirements for a Fully Deliberative Architecture (Or component of an architecture)* (Research Note No. COSY-DP-0604). Birmingham, UK: School of Computer Science, University of Birmingham.

Sloman, A. (2008). *Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)?* (Research Note No. COSY-PR-0702). Birmingham, UK.

Sloman, A. (2009). Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress. In B. Sendhoff, E. Koerner, O. Sporns, H. Ritter, & K. Doya (Eds.), *Creating Brain-like Intelligence* (pp. 248–277). Berlin: Springer-Verlag.

Sloman, A. (2011). What's information, for an organism or intelligent machine? How can a machine or organism mean? In G. Dodig-Crnkovic & M. Burgin (Eds.), *Information and Computation* (pp. 393–438). New Jersey: World Scientific.

Sloman, A., Chrisley, R., & Scheutz, M. (2005). The architectural basis of affective states and processes. In M. Arbib & J.-M. Fellous (Eds.), *Who Needs Emotions?: The Brain Meets the Robot* (pp. 203–244). New York: Oxford University Press.

Smith, J. M., & Szathmáry, E. (1995). *The Major Transitions in Evolution.* Oxford, England:: Oxford University Press.