

Some Representational and Ecological Aspects of Evolvability*

Tim Taylor

School of Computing, University of Abertay Dundee, Bell Street, Dundee DD1 1HG, Scotland
tim.taylor@abertay.ac.uk

Introduction

In this brief position paper I highlight a number of issues relevant to the evolvability of artificial life systems. These issues became apparent over the course of my doctoral research which involved studying the evolutionary dynamics of a Tierra-like system (Taylor 1999a). The first two issues (*multifunctional phenotypic components* and *semantic closure*) relate to the representation and function of individuals. The other two issues (*embeddedness* and *ecological interactions*) concern how individuals are related to the shared environment and to other individuals. While much of the recent work on evolvability has focussed on issues concerning individuals (e.g. genotype-phenotype mapping), which may give a system the *capacity* for high evolvability, I believe that the role of the ecological aspects of the system (e.g. the final two issues addressed in this paper) in providing a *drive* for evolvability has been neglected; even if a system has the capacity for high evolvability, it will not realise this capacity if the appropriate selection pressures are absent.

Multifunctional, Multimodal Phenotypes

John Maynard Smith, echoing observations by Charles Darwin (Darwin 1859) (p.220), has remarked that “it seems to be a general feature of evolution that new functions are performed by organs which arise, not *de novo*, but as modifications of pre-existing organs” (Maynard Smith 1986) (p.46). This principle could partially solve¹ a problem raised by Howard Pattee, of how new measuring devices (or other novel phenotypic interactions) arise during evolution (Pattee 1988): a structure with multiple properties (perhaps in a num-

ber of different modalities, e.g. computational, structural, chemical, electrical) might originally be selected for one of these properties (the “focal” property), but it might later turn out (quite accidentally) that some of its other properties also confer (unrelated) adaptive advantages upon the bearer of that structure. In such a scenario, an organism which duplicated this structure might have an adaptive advantage over those possessing a single copy, because each structure could be optimised for a single property. In this way, the organism can acquire new (possibly radically new) phenotypic properties based upon existing structures. This process is related to the concept of neutral networks, in that evolution will be able to “experiment” with variations of non-focal properties of a structure as long as such variations do not harm the focal property; if a variation of a non-focal property is found that is particularly beneficial for the organism, then this may then become the new focal property of the structure. This perspective may bring some light to bear upon the evolution of radical innovations, but it also opens up a whole range of new problems relating to the modelling of multiple, and mostly (initially at least) irrelevant, properties of objects; generally when devising a model we wish to strip away all irrelevant details of the system we are modelling, but the current perspective suggests that we should model structures with many different properties, in various modalities, many of which will, initially at least, be “irrelevant details”.

When there is a change in the focal property of a structure being acted upon by natural selection, we might regard the genetic information encoding the structure as now representing a different function (i.e. the function of the new focal property). However, in reality the semantics of the genetic system itself have not changed—all that has changed is our interpretation of the system as external observers. The issue of how completely new semantic information may arise in an evolutionary system is discussed in the following section.

*Some sections of this paper are revised versions of work previously published in (Taylor 1999a) and (Taylor 2000).

¹A full solution would also entail proper consideration of semantic closure, discussed in the next section.

Semantic Closure

From an epistemological point of view, Pattee points out that symbolic information (such as that contained in an organism's genes) has "no intrinsic meaning outside the context of an entire symbol system as well as the material organization that constructs (writes) and interprets (reads) the symbol for a specific function, such a classification, control, construction, communication ..." (Pattee 1995b). He argues that a necessary condition for an organism to be capable of open-ended evolution is that it encapsulates this entire self-referent organisation (Pattee refers to this condition as *semantic closure*). From this it follows that organisms should be constructed "with the parts and the laws of an artificial physical world" (Pattee 1995a) (p.36). In other words, for fully open-ended evolution, the whole organism, including the genome, the machinery for interpreting and executing the genome, and all phenotypic structures, should be explicitly represented in the (artificial) physical environment.²

Most existing artificial evolution systems, including genetic algorithms and Tierra (Ray 1991), do not fully satisfy these requirements. For example, in Tierra the machinery for interpreting and executing the instructions of individual programs resides in the Tierran operating system rather than being explicitly encoded by the programs themselves; see (Taylor 1999b). These systems therefore lack the ability to generate new symbolic representations; to take Tierra as an example again, the instruction set in which programs are written could not evolve *new* instructions to, say, sense aspects of the environment not monitored in some way by existing instructions. One example of an artificial system which *may* satisfy the requirements of semantic closure is von Neumann's cellular automata model of self-reproduction (von Neumann 1966).³ A fundamental aspect of semantic closure is *genetic relativism*—the ability of an evolutionary system to switch between different genetic languages; this topic is discussed in relation to von Neumann's work in another paper at this workshop (McMullin 2000).

²By extrapolation, this argument seems to imply that explicit interpretation machinery was required even at the earliest stages of prebiotic evolution. While this is clearly not the case for biological evolution, I believe that the argument still holds in that prebiotic self-replicators must at least have had the *potential* for explicit interpretation machinery to evolve. See (Taylor 1999a) (p.212) or (Taylor 2000) for further discussion.

³Although I think that an implementation of von Neumann's model would not exhibit high evolvability for other reasons, not least because the organisms have no capacity for self-maintenance in the face of environmental perturbations.

Embeddedness in the Arena of Competition

An essential requirement for an evolutionary process is that some form of selection mechanism exists, so that some variations of the reproducing entities are favoured over others. The selection mechanism therefore introduces a form of competition between the individual reproducers; they become engaged in a struggle for existence. The presence of such a mechanism implies that, in some form, the individuals coexist in an arena of limited capacity, and that they are competing with their neighbours (either globally or locally) for the right to be there.

An evolutionary system must therefore have a shared environment—an arena of competition—of some description, although there are few restrictions on the particular form it should take. All that is required is that it introduces the concept of a resource that is: (a) a vital commodity to individuals in the population; (b) of limited availability; and (c) that individuals can compete for (at either a global or local level). This resource can usually be interpreted as energy, space, matter, or a combination of these.

An issue that arises when considering different evolutionary systems is the extent to which individuals are embedded in this arena of competition. In von Neumann's cellular automata design, individuals are fully embedded—there is no 'hidden' state information (i.e. information which is not embedded in the cellular space itself). The same can be said of the biosphere, at least according to materialism. At the other extreme, individuals in a genetic algorithm (GA) have minimal embeddedness—the arena of competition merely contains place holders for the chromosomes, and the restriction is generally on the number of individuals, regardless of their size (although most GAs have constant-size chromosomes anyway). These two extremes, together with intermediate situations arising in Tierra and Avida,⁴ are depicted in Figure 1. Note that individuals in Avida are not really embedded in the arena of competition at all; the two-dimensional environment only holds pointers to the cells, in much the same way as in a GA.⁵ In Tierra, a program's in-

⁴Avida is a platform developed by Chris Adami and colleagues (see <http://www.krl.caltech.edu/avida/>). It is based upon Tierra, but significant differences include the fact that individual programs in Avida occupy positions in a two-dimensional arena.

⁵That is, the two-dimensional environment in which all of the programs coexist is distinct from the one-dimensional memory in which each individual program is stored. Furthermore, in the default settings of Avida, programs cannot read instructions of neighbouring programs, so no parasitism of this nature can emerge.

structions are embedded in the arena, although each program still has some additional state information (its ‘virtual CPU’ state). In Avida the fundamental space limitation applies to the number of *programs* that can fit in the arena of competition, whereas in Tierra it applies to the total number of *instructions* contained in all of the programs in the population.

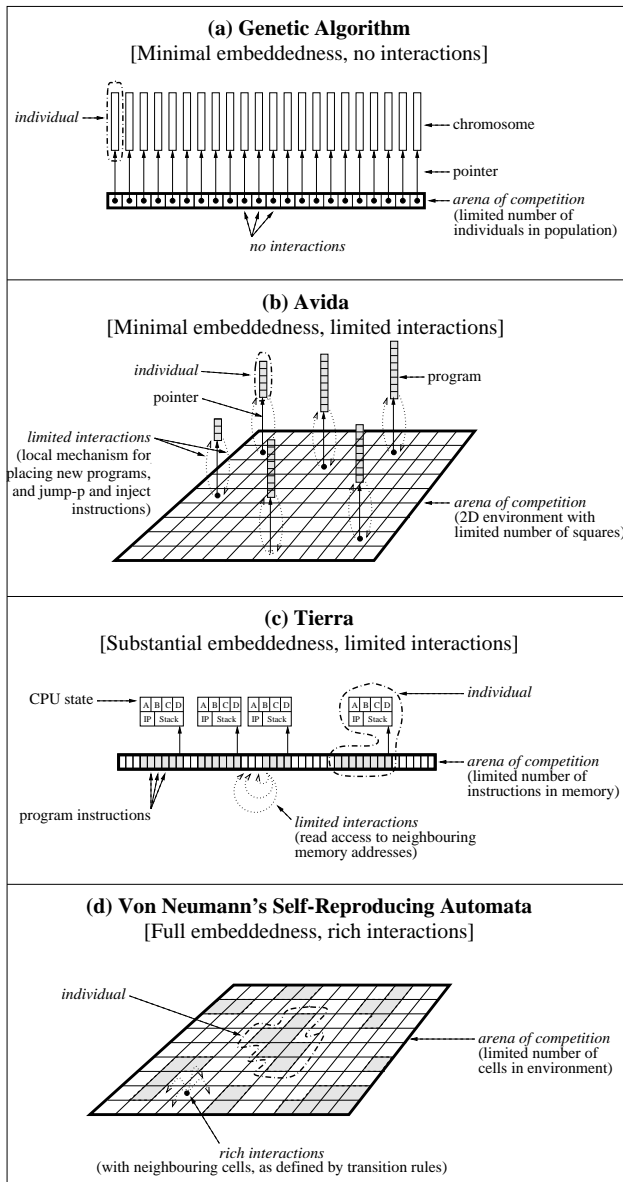


Figure 1: Embeddedness of Individuals and Richness of Interactions in Various Evolutionary Platforms.

Rich Ecological Interactions

Related to the issue of physical embeddedness is that of how restricted is the range of interactions that are

allowed between objects within the arena. In a standard GA, no direct interactions are allowed between chromosomes at all; the continued existence of an individual is decided by the extrinsically-defined selection mechanism. As already mentioned, in the default configuration of Avida programs cannot read the instructions of their neighbours. Although programs in Tierra are embedded in the arena of competition to a much greater extent than they are in Avida, the range of interactions allowed with neighbouring programs is still fairly restricted; programs can read the code of their neighbours, but they cannot directly write to neighbouring memory addresses.

In contrast, von Neumann's cellular automata implementation is far less restrictive; the transition rules of the cellular automata define neighbourhood interactions which occur at the level of individual cells and which therefore do not respect boundaries between individual organisms. This is of course similar to the situation of biological organisms, which have the freedom to interact with their environment in a variety of ways only limited by the laws of physics (although organisms themselves generally evolve mechanisms to restrict such free interaction).

From the point of view of the evolvability of individuals, the more embedded they are, and the less restricted the interactions are, then the more potential there is for the very *structure* of the individual to be modified. Sections of the individual which are not embedded in the arena of competition are ‘hard-wired’ and likely to remain unchanged unless specific mechanisms are included to allow them to change (and the very fact that specific mechanisms are required suggests that they would still only be able to change in certain restricted ways). Additionally, recall that Pattee has argued that open-ended evolution fundamentally requires the evolution of new meaning in the system, and this can only be achieved in the context of a semantically closed organisation which is completely embedded within the physical world.

One of the tenets of Darwinism is that organisms are engaged in a struggle for existence. However, it is difficult to identify the precise nature of this struggle, as Darwin himself observed. In *The Origin of Species*, he wrote “What checks the natural tendency of each species to increase in number is most obscure ... The amount of food for each species of course gives the extreme limit to which each can increase; but very frequently it is not the obtaining food, but the serving as prey to other animals, which determines the average numbers of a species” (Darwin 1859) (pp.119-120). Thus, an important aspect of the struggle for existence is the obtaining of food not from passive, abi-

otic sources, but through predator-prey relationships. In the biological realm, the struggle for existence involves organisms killing other organisms, because the very matter from which they are constructed is a valuable resource of matter and energy. This competition is therefore very much a matter of life or death.

It may be difficult to identify the precise nature of the struggle for existence, but it may be useful to categorise the numerous forms of competition in terms of a small number of fundamental resources. In the biosphere, a (speculative) list might be: matter, energy, space and information.

Tierra-like systems generally do not have any notion of competition for matter. Indeed, they cannot really be said to have a notion of matter at all, in terms of fundamental units from which all structures are built, and which are conserved during reactions. Instead, when a program is writing a copy of itself, it can produce the copied instructions spontaneously rather than first having to collect a copy of the individual instruction from somewhere else in memory. In other words, the individual instructions are represented as states of specific memory locations, rather than as units of matter. The only fundamental competition that exists in Tierra is for space (memory) into which to divide. This is allocated at a global level by the Tierran operating system's memory allocation services. The programs are not even really competing for energy (CPU-time), because any number of programs are allowed to execute instructions at each time slice; the limiting factor is how many programs can fit into the available memory. Programs in Tierra can act as resources for other programs in another way, by acting as 'library code' which can be read by their neighbours (as happens in the evolution of parasites). In other words, they can act as information resources. However, this is not as strong an ecological interaction as when one organism acts as a resource of matter or energy, in the sense that acting as an information resource is not a direct matter of life or death for the host. These stronger interactions may introduce selection pressure for the evolution of mechanisms for organismic self-maintenance and other processes associated with biological life.

Summary

I have highlighted a number of issues which have some bearing on the evolvability of a system. One representational issue which could improve the capacity of a system for high evolvability is the modeling of multifunctional, and possibly multimodal, phenotypic structures. While this may provide new transitional pathways to possibly more complex phenotypes, new symbols can only be introduced into the genetic language

of the organisms if their full semantically-closed organisation (including genotype, interpretation machinery, and phenotype) is explicitly represented in the artificial physical environment. Furthermore, only those organismic structures which are representationally embedded in the shared environment of the system (the arena of competition) will generally be able to become adapted to the environmental conditions presented by other organisms in the arena as they compete for limited resources. However, even if a system has the capacity for high evolvability, this capacity will not be realised in the absence of suitable selection pressures. Such drives will strongly depend upon the nature of the allowed interactions between organisms, and of the types of resources for which they are competing.

References

- Darwin, C. 1859. *The Origin of Species*. London: John Murray. (Reprinted in Penguin Classics, 1985. Any page numbers given in the text refer to this reprint.)
- Maynard Smith, J. 1986. *The Problems of Biology*. Oxford University Press.
- McMullin, B. 2000. The von Neumann self-reproducing architecture, genetic relativism and evolvability. ALIFE7 Evolvability Workshop.
- Pattee, H. 1988. Simulations, realizations, and theories of life. In Langton, C., ed., *Artificial Life*, volume VI of *Santa Fe Institute Studies in the Sciences of Complexity*, 63-77. Reading, MA: Addison-Wesley. Reprinted in "The Philosophy of Artificial Life", M.A. Boden (ed.), Oxford University Press, 1996.
- Pattee, H. 1995a. Artificial life needs a real epistemology. In Morán, F.; Moreno, A.; Merelo, J.; and Chacón, P., eds., *Advances in Artificial Life: Third European Conference on Artificial Life*, Lecture Notes in Artificial Intelligence, 23-38. Berlin: Springer.
- Pattee, H. 1995b. Evolving self-reference: Matter, symbols, and semantic closure. *Communication and Cognition—Artificial Intelligence* 12(1-2):9-28.
- Ray, T. S. 1991. An approach to the synthesis of life. In Langton, C.; Taylor, C.; Farmer, J.; and Rasmussen, S., eds., *Artificial Life II*. Redwood City, CA: Addison-Wesley. 371-408.
- Taylor, T. 1999a. *From Artificial Evolution to Artificial Life*. Ph.D. Dissertation, Division of Informatics, University of Edinburgh.
- Taylor, T. 1999b. On self-reproduction and evolvability. In Floreano, D.; Nicoud, J.-D.; and Mondada, F., eds., *Proceedings of the Fifth European Conference on Artificial Life (ECAL'99)*, number 1674 in Lecture Notes in Artificial Intelligence, 94-103. Springer.
- Taylor, T. 2000. Creativity in evolution: Individuals, interactions and environments. In Bentley, P. J., and Corne, D. W., eds., *Creative Evolutionary Systems*. Morgan Kaufman. (to appear late 2000).
- von Neumann, J. 1966. *The Theory of Self-Reproducing Automata*. Urbana, Ill.: University of Illinois Press.