

What Can We Learn from the First Evolutionary Simulation Model?

Seth Bullock

Informatics Research Institute, School of Computer Studies, University of Leeds
seth@scs.leeds.ac.uk

Abstract

A simple computer program dating from the first half of the nineteenth century is presented as the earliest known example of an *evolutionary simulation model*. The model is described in detail and its status as an evolutionary simulation model is discussed. Three broad issues raised by the model are presented and their significance for modern evolutionary simulation modelling is explored: first, the utility of attending to the character of a system's entire dynamics rather than focusing on the equilibrium states that it admits of; second, the worth of adopting an evolutionary perspective on adaptive systems beyond those addressed by evolutionary biological research; third, the potential for the non-linear character of complex dynamical systems to be explored through an individual-based simulation modelling approach.

With the war-time and post-war development of the first modern computers came a surge of research into computational theory. Seminal work by mathematicians such as McCulloch and Pitts (1943) on the logic of neural circuitry, Turing (1952) on diffusion-reaction models of morphogenesis, Walter (1963) and Ashby (1956) on cybernetics, von Neumann and Burks (1966) on automata theory and self-replication, and later Holland (1975) on the formal properties of adaptation, involved the application of logic, mathematics, robotics, and control theory to essentially biological problems.

The above-cited pieces of research are now recognised as the intellectual precursors to the field that has come to be known as artificial life. Although, more proximally, artificial life can be considered to be the offspring of artificial intelligence (see Brooks, 1991, and Steels, 1994, for accounts of artificial life's relationship to artificial intelligence), it is becoming increasingly apparent that the work published under the artificial life rubric (e.g., models of morphogenesis, cellular automata models, behaviour based robotics, the simulation of adaptive behaviour, etc.) has inherited much of its method, and some would say madness, either directly, or circuitously, from these mid-century pioneers.

However, it will be claimed here that a particular kind of artificial life, the *evolutionary simulation model*, originated far earlier than even the first of these seminal works. Coincidentally, the first evolutionary simulation model holds many lessons that are pertinent today. After introducing the model and discussing its status as an evolutionary simulation model, a series of issues raised by the model will be presented and their implications for modern artificial life explored.

The Ninth Bridgewater Treatise

In 1837, twenty-two years before the publication of Darwin's *On the Origin of Species*, and over a century before the advent of the first modern computer, a piece of speculative work was published as an uninvited *Ninth Bridgewater Treatise*. The previous eight works in the series had been sponsored by the will of Francis Henry Egerton, Earl of Bridgewater, and a member of the English clergy. The will's instructions were to make money available to commission and publish an encyclopedia of natural theology concerning "the Power, Wisdom, and Goodness of God, as manifested in the Creation" (Brock, 1966; Robson, 1990).

The ninth publication in this series is noteworthy in that, unlike typical works of natural theology, it neither sought to draw attention to miraculous states of affairs deemed unlikely to have come about by chance, and thus thought to be the work of a divine hand (e.g., the length of the terrestrial day, which seems miraculously suited to the habits of man and other animals), nor did it seek to reconcile scientific findings with a literal reading of the Old Testament (e.g., disputing evidence that suggested an alarmingly ancient earth, accounting for the existence of dinosaur bones, or promoting evidence for the occurrence of the great flood, etc.). In contrast to these apologetic efforts, the author of the ninth Bridgewater treatise produced what is, to my knowledge, the first instance of an *evolutionary simulation model*.

The author of the *Ninth Bridgewater Treatise* was

Copyrighted Material

Charles Babbage, the designer of the difference engine and analytical engine (the first automatic calculating devices, and thus precursors to the modern computer). Indeed, in 1837 he was one of perhaps a handful of scientists capable of carrying out research involving automated computational modelling. His model stands as a usefully simple case study, and an elegant example of the use to which computers are being put in modern artificial life.

The aim of this paper will not be to draw attention to Babbage's model as an example of a particularly prescient piece of work, anticipating much of modern evolutionary simulation modelling. Rather, this antique evolutionary simulation model will be revived, not only to identify similarities between Victorian science and current artificial life, but in order to highlight important aspects of the contemporary practice of evolutionary simulation modelling. In pursuing this aim, I am well aware of the risks of Whiggism in interpreting historical material (see Hyman, 1990, for discussion of Whiggism in the study of Babbage and his work). Although a modern perspective on the past is unavoidable, the fact that this paper does not primarily seek to re-evaluate Babbage's work but to re-analyse modern evolutionary simulation modelling work in the light of Babbage's model ensures that the risk of mis-treating Babbage's work is slim.

The First Evolutionary Simulation Model

Babbage's (1837) model (see also Babbage, 1864, Chapter XXIX "Miracles" for a rather whimsical account of the model's development) was situated within what was then a controversial debate. It addressed the dispute between *catastrophists* and *uniformitarians*. Prima facie this debate was internal to geology, since it concerned the geological record's potential to show evidence of divine intervention (principally in the form of support for the Old Testament accounts of the Creation and the Deluge). Catastrophists argued for an interventionist interpretation of geological evidence, taking discontinuities in the record to be indicators of the occurrence of miracles (violations of laws of nature). In contrast, uniformitarians insisted that in order to carry out scientific enquiry, the entire geological record must be assumed to be the result of unchanging processes. Allowing a role for divine miracles, the uniformitarians claimed, would render competing explanations equally valid. No theory could be claimed to be more parsimonious or coherent than a competing theory that invoked *necessarily inexplicable* exogenous influences in its account of the phenomena at issue.

Although this dispute had already been dealt some-

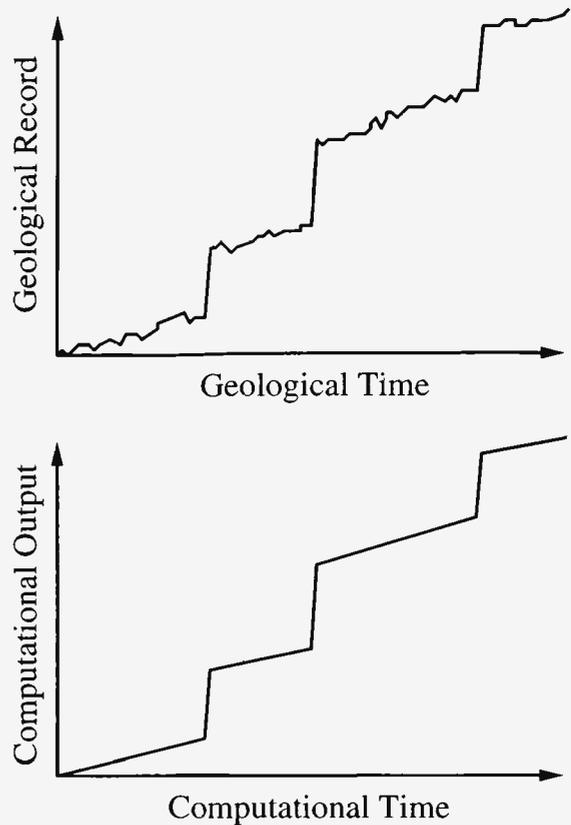


Figure 1: Babbage's (1836) evolutionary simulation model represented the empirically observed history of geological change as evidenced by the geological record (*upper panel*) as the output of a computing machine following a program (*lower panel*). A suitably programmed computing machine could generate sequences of output that exhibited discontinuities without requiring external influence. Hence discontinuities in the actual geological record did not require "catastrophic" divine intervention, but could be the result of "gradualist" processes.

thing of a death blow with Lyell's (1830) publication of his *Principles of Geology*, the publication of the Bridgewater treatises and works like them evidences its slow demise. Only subsequent to the *coup de grâce* provided by Darwin's work on evolution did natural theology texts finally cease to be published (Brock, 1966).

Babbage's response to the catastrophist position was to construct what can now be recognised as a simple evolutionary simulation model (see figure 1). He proposed that a suitably programmed difference en-

gine could be made to output a series of numbers according to some law (e.g., the integers, in order, from 0 onwards), but then at some pre-defined point (e.g., 100,000) begin to output a series of numbers according to some different law (e.g., the integers, in order, from 200,000 onwards). Although the *output* of such a difference engine (an analogue of the geological record) would feature a discontinuity (in our example the jump from 100,000 to 200,000), the *underlying process* responsible for this output would have remained constant (i.e., the general law, or program, that the machine was obeying would not have changed). The discontinuity would have been the result of the naturally unfolding mechanical (and computational) process. No external tinkering analogous to the assumed intervention of a providential deity would have taken place.

Babbage thus tried to show that what might appear to be discontinuities were not necessarily the result of meddling, but could be the natural result of unchanging processes. In doing this he cultivated the image of God as a programmer, engineer, or industrialist, capable of setting a process in motion that would accomplish His intentions without Him intervening repeatedly. In Victorian Britain, the notion of God as draughtsman of an 'automatic' universe, one that would run unassisted, without individual acts of creation, destruction, etc., proved attractive. This conception was subsequently reiterated by several other natural philosophers (e.g., Darwin, Lyell, and Chambers) who argued that it implied "a grander view of the Creator — One who operated by general laws" (Young, 1985, p.148).

For the purposes of this paper, what is of interest are not the theological implications of Babbage's work, nor the effect it had on the catastrophist/uniformitarian debate, but the manner in which Babbage mobilised his computational resources to attack a theoretical position. Babbage's computational system was a simple analogue of a natural system (the geology of the planet) implemented mechanically. Babbage did not seek to capture the complexity of real geology in his system. Indeed the analogy between the difference engine's program and geological processes is a crude one. However, the formal resemblance between the computing machine and the geological process is sufficient to enable a point about the latter system's dynamics to be made. Babbage's computing machine is thus clearly being employed as a *model*.

Evolutionary Simulation Models

What grounds do we have for claiming that Babbage's model is an example of an *evolutionary simulation model*? At first glance, the fact that the model involves

no mention of competition, heritable variation, limited resources, etc., would indicate that, even if it could be called a simulation model, Babbage's programmed calculator should not be awarded the status of *evolutionary simulation model*. However, the adjective *evolutionary* is being used here not to invoke the notion of biological evolutionary change, but to draw attention to the fact that Babbage's model was implemented as a dynamic, unfolding, process. For the moment, a few brief observations will serve to give a flavour of what is intended by the phrase¹.

First, the fact that Babbage's model is an unfolding computational process sets it apart from work in which models of dynamic change are constructed as mathematical proofs and are thus not evolutionary simulations. For example, Malthus' (1798) work on population dynamics, in which he demonstrated that population growth would outstrip that of agriculture, was constructed using paper and pencil.

However, it is not merely the computational nature of Babbage's model that ensures its status as an evolutionary simulation model. His reliance on the ongoing dynamic behaviour of his computational model, rather than on any end result it might produce, distinguishes it from modelling work which, although involving computational processes, uses computers as tools for solving what would otherwise prove to be intractable mathematical problems. For example, the use of computers to discover digits of pi, or to iteratively solve the differential equations that might comprise a model of population dynamics, do not count as *evolutionary simulation modelling* since the computational processes involved are merely the means of reaching a particular solution. In contrast, the substantive element of Babbage's model is the evolutionary aspect of the simulation (i.e., the manner in which it changes over time). In the scenario that Babbage considers, his suitably programmed difference engine will, in principle, run forever. Its calculation is not intended to produce some end product, but rather the ongoing calculation is itself the object of interest.

In the following sections particular aspects of Babbage's model will be expanded upon. First, the implications of considering the dynamic behaviour of a model to be central, rather than concentrating on the end-product of some calculation, will be discussed.

Dynamics and Stasis

Artificial life is perhaps exclusively concerned with systems that change over time, and, furthermore, the

¹For a more detailed treatment of the notion of evolutionary simulation modelling see Bullock (1997).

manner in which such systems change over time². Whether the system be a cellular automaton implementing the “game of life”, an autonomous robot designed to navigate an extra-terrestrial terrain, or a model of “complexity at the edge of chaos”, it is considered as a time-varying system with a certain dynamic character. It is this character that is of interest to the artificial life practitioner. Will the dynamic character of the cellular automaton admit of “universal computation”? Will the dynamic character of the autonomous robot result in robust walking behaviour? Will the dynamic character of the model exhibit “complexity at the edge of chaos”?

We can contrast this interest in dynamic change with the approach taken by game-theoretic models (von Neumann & Morgenstern, 1944). Despite being similarly concerned with systems that change over time (economies, individual economic agents, ecologies, populations of creatures, individual creatures, etc.), game-theoretic accounts of adaptive phenomena typically assume that the systems under consideration are at, or near, equilibria. Once this assumption is in place, the game theorist is faced with the task of specifying a model that admits of stable equilibria with a character that matches that of the observed economic or biological system.

For example, neo-classical economic theory asserts that since the economic agent, *homo economicus*, is an ideal maximiser of expected utility, such agents will clear a market at the equilibrium price. There is thus no sense in asking what behaviour would result from a system comprised of agents who cannot maximise expected utility. Such a system is far from equilibrium, and thus not likely to be found to reflect real economic situations in which markets are either at or near equilibria. Since, from this perspective, economic agents are assumed to be optimal players, one merely needs to identify the equilibrium price analytically in order to describe the behaviour of the market, as this is the price that will be settled upon. Attention to the global dynamics of the model is not necessary since the system will not spend time far from equilibrium.

Whilst these assumptions are, for the most part, eminently reasonable, there are scenarios in which the conditions that real economic agents find themselves in result in their inability to find the equilibrium price, e.g., sellers at an auction who are wary of the existence of cartels amongst their fellow bidders. In order to model these systems, approaches that take into account a richer pallet of dynamic behaviour are nec-

² Artificial life is clearly not the only field concerned with dynamical systems. Related fields such as cybernetics and control theory, for instance, share similar interests.

essary (see, e.g., Binmore, 1987, 1988, for critiques of the traditional axiomatic approach to game-theoretic economic modelling).

A similar perspective is evident within game-theoretic accounts of evolutionary systems. Maynard Smith (1982) identifies this problem at the outset of his book, *Evolution and the Theory of Games*,

“An obvious weakness of the game-theoretic approach to evolution is that it places great emphasis on equilibrium states, whereas evolution is a process of continuous, or at least periodic change. The same criticism can be levelled at the emphasis on equilibria in population genetics. It is of course mathematically easier to analyse equilibria than trajectories of change” (p. 8).

However, unlike economists, evolutionary game theorists have better grounds for pursuing a program of what Frank (1998) terms *comparative statics* than this appeal to the intractability of dynamic models. Identifying equilibria, and exploring their sensitivity to model parameters in order to make predictions about analogous real-world systems is a process with some chance of engaging with empirical biological observations, since, given that the natural systems around us are likely to be at or near equilibrium, we have some chance of collecting appropriate data. The likelihood of obtaining the observations necessary to decide between competing theories invoking trajectories of change is much smaller.

One area in which such data are routinely collected is in the construction of phylogenetic histories by systematists — the bioscience descendants of the geologists that Babbage’s model addressed. For our present purposes, these histories are important because differing evolutionary theories often make the same predictions concerning *current* states of affairs. This is because it is present-day phenomena that the theories attempt to account for. However, competing theories may make differing predictions concerning the *prior* states of affairs that have led to the current situation.

For instance, Ryan (1990) attempts to distinguish between theories that compete to account for the character of sensory systems and signalling behaviour extant in the natural world by constructing a phylogenetic tree for several species of frog. From this hypothetical history of speciation events Ryan attempts to discount certain theories whose predictions do not match the historical account he has constructed.

Evolutionary simulation modelling can contribute to this style of hypothesis testing in a way in which modelling methodologies that exclusively attend to equilibria cannot. An evolutionary simulation model provides an account of not only the behaviour of a system

at equilibrium, but also the behaviour which that system passed through before it reached this equilibrium. Such accounts of the trajectories followed by evolving populations prior to (potentially) achieving equilibria might be used to distinguish between competing theories.

This is not to say that data from simulations will simply augment data from phylogenetic reconstructions, but that the implications of evolutionary theories for the character of evolutionary trajectories might be clarified through evolutionary simulation modelling. The predictions resulting from such a clarification could then be compared to empirical data in the usual manner.

Although it is perhaps reasonable to expect the natural systems we see around us to be at stable equilibria given the evolutionary timescales involved, as with economic systems, there are situations in which evolving populations may consistently fail to reach equilibria, or in which the equilibria that evolutionary systems do reach are more complicated than point attractors. For example, Maynard Smith follows the passage quoted above with a prediction that *cyclic* attractors will be discovered to characterise much of the behaviour exhibited by players involved in asymmetric games. This prediction has been supported by the discovery of a species of lizard that occurs in three distinct morphs, each of which dominates one other morph, and is dominated by the remaining morph. Such a system is analogous to the parlour game scissors-paper-stone, in which playing one move consistently will never be a lasting strategy since any such strategy can be defeated (Sinervo & Lively, 1996; Maynard Smith, 1996).

In addition, theorists are coming to realise that many interesting games exhibit *multiple equilibria*. Unfortunately, naked game theory is unable to determine, given the existence of more than one equilibrium state, which equilibrium a population will arrive at. Additional criteria for deciding between equilibria (e.g., on grounds of parity, efficiency, etc.) have been offered (Harsanyi & Selten, 1988), but these often appear somewhat arbitrary. The natural solution to this *equilibrium selection problem* is to enquire which equilibria arise from which initial conditions (e.g., Binmore, Gale, & Samuelson, 1995a; Binmore, Samuelson, & Vaughan, 1995b) — a question addressing the dynamic character of the model.

Despite acknowledging that the behaviour of adaptive systems is inherently dynamic, and that attention to these dynamics might enable theorists to distinguish between competing theories, theorists often eschew the study of dynamic change. This accounts for the accent placed on the fixed points of models, whether they be

evolutionary stable strategies in biology, Nash equilibria in economics, or “exit points” in language change (Labov, 1994), rather than the general dynamic behaviour of such models. This is not to say that game theory and other formal approaches cannot tolerate limit sets of a higher order than constant trajectories, or have no consideration of initial conditions or transient behaviours. However, such matters are typically regarded as special cases that require additional analytic techniques if they are to be addressed at all (e.g., Maynard Smith, 1982, devotes an appendix to dealing with cyclic trajectories).

In contrast, evolutionary simulation modelling will principally concern itself with the *character* of a model's evolutionary dynamics rather than some “end product” of these dynamics, whether it be within a population of learning economic agents, a population of evolving creatures, or a population supporting a developing culture or language. From this inherently dynamic perspective, cyclic limit sets and start-up transients, drift and chaos, are on an equal footing with game theory's cardinal limit set, the point attractor.

Subject Matter

Perhaps the aspect of Babbage's model that most clearly distinguishes it from modern evolutionary simulation modelling is its subject matter. Several issues are related to this observation.

First, the overtly theological concerns of the debate Babbage engaged with are for the most part missing from contemporary science. Modern scientists now rarely struggle against matters of faith or religion in print (see Dawkins, 1998, however, for one contemporary example). Despite this, just as certain commentators claim that Babbage's model influenced the conception of God in the 19th century (Young, 1985), some researchers have taken modern artificial life to have implications for the relationship between religion and science (Helmreich, 1997).

Leaving religious matters to one side, it is also the case that whereas Babbage's system models an abstract *geological* process, the majority of current modelling work in this vein addresses *biological* subject matter. Despite the fact that modern evolutionary simulation modelling has tended to address biological questions, commentators have recognised its potential to address problems in other disciplines. However, it is probable that in 1836, Babbage and his contemporaries would not have recognised a difference between, for instance, geology and biology, since these fields and many others under the umbrella of natural philosophy had yet to part company and begin to specialise. Now that we understand that the mechanism by which or-

ganismic evolution proceeds (the differential transmission of genetic material) is not present in geological, economic, linguistic or psychological systems, should we be wary of the recent trend within artificial life of applying simulation models of adaptation to an ever broader class of topics and problems?³

Some of the non-biological disciplines that are beginning to adopt this approach have been exploring questions of change for some considerable time. For instance, both anthropology and linguistics involve a historical, diachronic or developmental element. The study of language change, for instance, predates Darwin's theory of evolution; indeed Darwin made use of research into the history and relatedness of languages in formulating his theory of natural selection.

In contrast, other disciplines have taken up the challenge of understanding the dynamics of their phenomena more recently. For example, economics (despite the prompting of Veblen, 1898, at the turn of the century) has only recently begun to consider the processes that might underly the evolution of economic systems (c.f., the inception of the *Journal of Evolutionary Economics* in 1991, that builds on the pioneering work of Joseph Schumpeter, e.g., 1934). Previously such matters were the preserve of historians of economics.

In the most extreme cases, novel scientific programs must be developed in order to pursue the implications of an adaptive systems perspective. Memetics, the study of the evolution of ideas, is one example of such a neonatal paradigm.

What links this rather disparate group of disciplines is their concern with the *dynamics of adaptation*. These dynamics are often studied by modellers using techniques developed specifically to deal with their indigenous problems, with no reference to evolution, or even with explicit rejection of evolutionary thinking. However, increasingly theorists are coming to see parallels between the dynamics underlying many different adaptive systems (c.f., the proliferation of journals, meetings and book titles involving *Evolutionary* as a leading adjective). Although, the systems that they study do not involve adaptation in the form of orthodox organismic evolution, nevertheless, the necessary ingredients for adaptation can be identified: competition for limited resources and heritable variation. For economic systems, the limiting resource is utility, the

³For example, Miller (1997) has recently appealed to the notion of self-organization in an attempt to resolve the apparent paradox presented by, on the one hand, the discovery of Hitler's indolence, and, on the other, the intensely structured order of the Third Reich. Miller quotes evidence from artificial life simulations that suggest that complex order may arise from simple local interactions, rather than requiring global co-ordination from some central executive.

variation exists at the level of economic strategy and inheritance occurs socially through some kind of learning mechanism. For linguistics, the limiting resource is language users, the variation exists at the level of language structure and inheritance again occurs socially through language transmission.

Evolutionary biologists enjoy an advantage over evolutionary simulation modellers dealing with non-biological systems in that they possess a detailed understanding of the mechanisms underlying biological evolution. As yet, comparable understanding of economic, linguistic, cultural, or psychological adaptation is relatively lacking. Non-biological adaptationists have made progress by appropriating mechanisms from biological evolution. However, the extent to which these adaptive mechanisms are merely being used metaphorically is often unclear. For example, the epidemiological notions of contagion used by Cavalli-Sforza and Feldman (1981) to model the spread of innovations may be considered to apply literally or metaphorically, depending, perhaps, on one's perspective. In contrast, it has been argued (Crick, 1989) that the concept of competitive exclusion that underlies "Neural Darwinism" (Edelman, 1987) does not constitute a literal example of natural selection in the brain due to crucial dissimilarities between Darwinism and "Edelmanism". Fundamental research into the unique character of these non-biological adaptive systems must eventually reify or replace these placeholders (e.g., Gatherer, 1998).

The prospect of multiple levels of adaptation interacting with one another further complicates the picture. The learning mechanisms invoked by economists, linguists and cultural anthropologists are not fixed entities, but are themselves the results of adaptive evolutionary processes. Some effort has been made to model the interaction between learning and evolution (e.g., Hinton & Nowlan, 1987) and to apply the insights thus gained to non-biological adaptationist research programs (e.g., Kirby & Hurford, 1997). But until theoretical approaches to parallel, interacting adaptive systems (e.g., Laland, Odling-Smee, & Feldman, in press) can be shown to be sound, these modelling enterprises will be less secure than their biological forebears.

In summary, any research paradigm that studies the behaviour of systems of entities which interact with each other and their environment over time such that they change in an adaptive fashion is amenable to the evolutionary simulation modelling approach. While Babbage's application of a simulation model to what is now considered to be a non-biological topic presaged modern simulations of non-biological adaptive systems, his model is lacking in a sophisticated notion of geo-

logical adaptation. Indeed it is unclear whether geological systems can be classed as adaptive in the sense described above. It is apparent that modern descendants of Babbage's model are more concerned with the details of *how* adaptation occurs in natural systems, rather than merely demonstrating that some phenomena can be replicated on a machine. This notion will be pursued further in the next section.

Emergence and Individuals

The phenomenon at the heart of Babbage's model — discontinuity — is emblematic of the concerns of present-day researchers employing evolutionary simulation models. However, it is in its approach to non-linearity that Babbage's model departs most significantly from modern models.

There is a superficial resemblance between the catastrophist debate of the 19th century and the more recent dispute over the theory of punctuated equilibria introduced by Eldredge and Gould (1973). Both arguments revolved around the significance of what appear to be abrupt changes at geological time scales. However, while Babbage's dispute centered on whether change could be explained by one continuously operating process or must involve two different mechanisms (the first being geological processes, the second Divine intervention), Gould and Eldredge take pains to point out that their theory does not supercede phylogenetic gradualism, but augments it. They wish to explain the two apparent modes of action evidenced by the fossil record (long periods of stasis, short bursts of change), not by invoking two processes, but by explaining the unevenness of evolutionary change. In this respect, the theory that Eldredge and Gould supply attempts to meet a modern challenge: that of explaining non-linearity, rather than merely accommodating it.

Whereas Babbage's aim was merely to demonstrate that a certain kind of non-linearity was logically possible in the absence of exogenous interference, modern researchers are probing questions of how and why non-linearities arise from the homogeneous action of low-level entities, and what implications these non-linearities may have for the systems under examination. In order to achieve this, modern simulation models have had to move beyond the elegant but simplistic form of Babbage's demonstration.

This increased sophistication stems, in part, from the use of an explanatory strategy that invokes a "constructive" relationship between at least two relevant levels of description: a level of explicitly modelled individual atomic entities and a higher level of aggregate phenomena. The notion that the possibly complex and non-linear behaviour of higher-level phenom-

ena emerges from the action of lower-level entities allows that an understanding of the constructive relationship that links them may be all that is required to account for the behaviour at the aggregate level of description. For example, a simulation model of traffic dynamics may involve routines that deal with the individual vehicles comprising the traffic without at any point explicitly invoking the higher-level phenomenon of "jams". Despite this, the simulation may be a useful way of modelling how jams behave if, through analysis of the simulation, an understanding of how the "emergent" phenomena derive from the atomic entities can be achieved.

It is important to explicate the differences between this kind of explanatory project and the task met by Babbage's model. Babbage did not need to model a system of entities at some atomic geological level of abstraction and then simulate the emergence of discontinuities in the geological record since his project was merely to demonstrate that a class of phenomena could exist in the absence of an element that had previously been considered necessary — external intervention. As such nothing hinged on the manner in which the natural phenomena actually did arise. This kind of explanation is a proof of concept of the type: "it is commonly thought that M is needed to generate P , but here is a model in which M is missing, but something that looks like P is exhibited". One of the challenges for modern evolutionary simulation models is to move beyond this kind of explanation, and reveal the constructive relationships that hold between atomic and emergent levels of description (see, e.g., Di Paolo, Noble & Bullock, this volume, for further discussion).

Babbage himself was not satisfied with merely demonstrating through simulation modelling that apparent discontinuities could be the result of unchanging mechanical processes. He also spent some time developing theories with which he sought to explain how specific examples of geological discontinuity could have arisen as the result of physical geological processes. One example of apparently rapid geological change that had figured prominently in geological debate since being depicted on the frontispiece of Lyell's *Principles of Geology* was the appearance of the Temple of Seraphis on the edge of the Bay of Baiae in Pozzuoli, Italy. The surface of the 42-foot pillars of the temple are characterised by three regimes. The lower portions of the pillars are smooth, their central portions have been attacked by marine creatures, while above this region the pillars are weathered but otherwise undamaged. These abrupt changes in the character of the surface of the pillars were taken by geologists

to be evidence that the temple had been partially submerged for a considerable period of time.

For Lyell (1830), an explanation could be found in the considerable seismic activity which, historically, had characterised the area. It was well known that eruptions could cover land in considerable amounts of volcanic material and that earthquakes could suddenly raise or lower tracts of land. Lyell reasoned, that a volcanic eruption could have buried the lower portion of the pillars before an earthquake lowered the land upon which the temple stood into the sea. Thus the lower portion would have been preserved from erosion, while a middle portion would have been subjected to marine perforations, and an upper section to the weathering associated with wind and rain.

Recent work by Dolan (1998) has uncovered the impact that Babbage's own thoughts on the puzzle of the pillars had on this debate. Babbage, while visiting the temple, noted an aspect of the pillars which had hitherto gone undetected: a patch of calciated stone located between the central, perforated, section, and the lower, smooth, portion. Babbage inferred that this calciation had been caused, over considerable time, by calcium bearing spring waters which had gradually flooded the temple, as the land upon which it stood sank lower and lower. Eventually this subsidence caused the temple pillars to sink below sea-level and resulted in the marine erosion evident on the middle portion of the columns.

Thus Babbage's explanation invoked gradual processes of cumulative change, rather than abrupt episodes of discontinuous change, despite the fact that the evidence presented by the pillars is that of sharply separated regimes. Babbage's account of this gradual change relied on the notion that a central, variable source of heat, below the earth's crust, caused expansion and contraction of the land masses above it. This expansion or contraction would lead to subsidence or elevation of the land masses involved. Babbage exploited the power of his new calculating machine in attempting to prove his theory, but not in the form of a simulation model. Instead, he used the engine to calculate tables of values that represented the expansion of granite under various temperature regimes. With these tables, Babbage could estimate the temperature changes that would have been necessary to cause the effects manifested by the Temple of Seraphis.

Here, Babbage is using a computer, and is moving beyond a gradualist account that merely tolerates discontinuities (i.e., his *Bridgewater Treatise*) to one that attempts to explain them. However, his engine is not being employed as an evolutionary simulation model, but as a prosthetic calculating device. The

complex, repetitive, computations involved in producing and compiling these tables of figures would normally have been carried out by "computers", people employed to make calculations manually. In replacing this error-prone, slow and costly manual calculation with his mechanical reckoning device, Babbage demonstrates the application of computing power to solving problems that are otherwise intractable. This use of computers has become widespread in modern science. Numerical and iterative techniques for calculating (or at least approximating) the results of what would be extremely taxing or tedious problems has become a mainstay of much academic practice.

In contrast, evolutionary simulation models of the kind discussed in this paper offer a new role for powerful computers. Where Babbage employed his machine to either (i) demonstrate that some natural phenomenon could be simulated in the absence of an element that had previously been deemed a necessary pre-requisite, or (ii) perform otherwise intractable calculations in order to support theory building/testing, modern simulation modellers attempt to move beyond these uses in pursuing an explanatory strategy in which simulations are used to directly explore explanatory theories proposed as ways of understanding how complex, aggregate behaviour might arise from the homogeneous action of lower-level entities.

A successful example of completing this explanatory task can be found in Di Paolo's (2000) model of the evolution of co-ordinated communication. Using an individual-based evolutionary simulation model featuring a spatially distributed population of agents and playing an action-response game, Di Paolo shows that even in situations for which game-theoretic considerations predict that co-ordination will be unstable, co-ordination may arise and persist. At this point the model has fulfilled the same explanatory role as that of Babbage's *Bridgewater Treatise*: an aggregate phenomena (co-ordinated communication) has been demonstrated in the absence of an element previously deemed necessary (i.e., an appropriate equilibrium in the underlying game). Di Paolo accounts for the presence and character of this co-ordinated communication by first drawing attention to the manner in which the spatial structure of the medium gives rise to clusters of individuals. Subsequent exploration reveals that individuals in the center of such clusters face a scenario that differs significantly from those at the periphery. Analysis demonstrates that the strategic asymmetry induced by this spatial organisation is sufficient to enable co-ordinated communication to persist as a stable strategy.

Notice that merely appealing to the existence of spa-

tial clustering, and invoking the idea that co-ordination “emerges” from the interactions of game-players that exist in a spatially-structured medium would be to fall short of such a successful explanation. Such an appeal would in fact be closer to the apologeticist use of “miracles” as explanations for phenomena that would otherwise be inexplicable. Indeed construing emergent phenomena to be those aggregate phenomena for which, as yet, we have no reductionist explanation (Ronald, Sipper, & Capcarrère, 1999) would seem to invite this comparison.

In summary Babbage’s model usefully demonstrates a simple explanatory strategy which modern evolutionary simulation modelling hopes to move beyond. In order to do so, evolutionary simulation models must do more than invoke the notion of emergence. The relationship between atomic and aggregate phenomena must be explicated successfully.

Conclusions

In many respects Babbage’s model second-guesses aspects of modern evolutionary simulation modelling. His use of an ongoing computational program to model the dynamics of a natural system, his attention to a debate that would now be regarded as external to evolutionary biology, and his concentration on high-level non-linear phenomena and the ability of low-level processes to give rise to them, are all prominent features of modern individual-based simulations of adaptive systems.

However, Babbage’s model does not address the manner in which exploring the dynamics of a simulation can explicate the constructive relationships that account for high-level aggregate phenomena in terms of the behaviour of systems of low-level atomic entities. The model does provide an example of a more simple explanatory strategy, that of demonstrating that some high-level phenomenon survives the removal of an element previously deemed to be a necessary prerequisite for it. Moving beyond this class of explanation has been identified here as an important challenge for modern evolutionary simulation modelling.

Babbage’s model may resemble current artificial life in one final respect that is perhaps worth noting. Upon publication of the *Ninth Bridgewater Treatise*, the work was treated with some disdain. His demonstration of the power of automatic computing was generally regarded as impressive. But although his machine was a remarkable feat of engineering it was perceived to be a tool ill-suited to the job of natural philosophy and his model gained little credibility as a result. It was generally agreed to have overstepped some boundary. In contrast, Dolan’s (1998) recent work has shown

that Babbage’s empirically-driven theories on the geological processes responsible for the appearance of the ‘Temple of Seraphis’ were readily taken on board by the eminent uniformitarian geologists of the time. Like Babbage’s *Bridgewater treatise*, contemporary evolutionary simulation models may also be “uninvited”. Perhaps the lessons we can learn from Babbage’s work will ensure that current modelling efforts have more chance of gaining a better reception.

Acknowledgments

Thanks to Ezequiel Di Paolo, Henrietta Wilson and two anonymous referees for helpful comments.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, London.
- Babbage, C. (1837). *Ninth Bridgewater Treatise: A Fragment* (2nd edition). John Murray.
- Babbage, C. (1864). *Passages from the Life of a Philosopher*. Longman, London.
- Binmore, K. (1987). Modelling rational players, part I. *Economics and Philosophy*, 3, 179–214.
- Binmore, K. (1988). Modelling rational players, part II. *Economics and Philosophy*, 4, 9–55.
- Binmore, K., Gale, J., & Samuelson, L. (1995a). Learning to be imperfect: The ultimatum game. *Games and Economic Behavior*, 8, 56–90.
- Binmore, K., Samuelson, L., & Vaughan, R. (1995b). Musical chairs: Modelling noisy evolution. *Games and Economic Behavior*, 11, 1–35.
- Brock, W. H. (1966). The selection of the authors of the *Bridgewater Treatises*. *Notes and Records of the Royal Society of London*, 21, 162–179.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.
- Bullock, S. (1997). *Evolutionary Simulation Models: On their Character, and Application to Problems Concerning the Evolution of Natural Signalling Systems*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.
- Cavalli-Sforza, L., & Feldman, M. (1981). *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton University Press, Princeton, NJ.
- Crick, F. (1989). Neural edelmanism. *Trends in Neurosciences*, 12(7), 240–248.
- Darwin, C. (1859). *The Origin of Species by Means of Natural Selection*. John Murray, London.
- Dawkins, R. (1998). *Unweaving the Rainbow: Science, Delusion and the Appetite for Wonder*. Allen Lane, London.

- Di Paolo, E. A. (2000). Ecological symmetry breaking can favour the evolution of altruism in an action-response game. *Journal of Theoretical Biology*, 203, 135–152.
- Di Paolo, E. A., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. This volume.
- Dolan, B. P. (1998). Representing novelty: Charles Babbage, Charles Lyell, and experiments in early Victorian geology. *History of Science*, 36, 299–327.
- Edelman, G. M. (Ed.). (1987). *Neural Darwinism: The Theory of Neuronal Group Selection*. Basic Books, New York.
- Eldredge, N., & Gould, S. J. (1973). Punctuated equilibria: An alternative to phyletic gradualism. In Schopf, T. J. M. (Ed.), *Models in Paleobiology*, pp. 82–115. Freeman, Cooper and Co, San Francisco.
- Frank, S. A. (1998). *Foundations of Social Evolution*. Princeton University Press, Princeton, NJ.
- Gatherer, D. (1998). Why the thought contagion metaphor is retarding the progress of memetics. *Journal of Memetics — Evolutionary Models of Information Transmission*, 2. <http://www.cpm.mmu.ac.uk/jomemit/1998/vol2/>. Last accessed 24.1.2000.
- Harsanyi, J., & Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- Helmreich, S. (1997). The spiritual in artificial life: Recombining science and religion in a computational culture medium. *Science as Culture*, 6(3), 363–395.
- Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1, 495–502.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor. Reprinted by MIT Press, 1992.
- Hyman, R. A. (1990). Whiggism in the history of science and the study of the life and work of Charles Babbage. *Annals of the History of Computing*, 12(1), 62–67.
- Kirby, S., & Hurford, J. (1997). Learning, culture and evolution in the origin of linguistic constraints. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life (ECAL'97)*, pp. 493–502. MIT Press / Bradford Books, Cambridge, MA.
- Labov, W. (1994). *Principles of Linguistic Change*, Vol. Volume 1: Internal Factors. Blackwell, Oxford.
- Laland, K. N., Odling-Smee, F. J., & Feldman, M. W. (in press). Niche construction, biological evolution and cultural change. To appear in *Behavioral and Brain Sciences*.
- Lyell, C. (1830/1970). *Principles of Geology*. John Murray, London.
- Malthus, T. (1798). *An Essay on the Principle of Population*. Patricia James (Ed.), 1989, Cambridge University Press, Cambridge.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.
- Maynard Smith, J. (1996). The games lizards play. *Nature*, 380, 198–199.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115–133.
- Miller, J. (1997). Start The Week. BBC Radio 4. Discussion of the television documentary ‘Nazis: A Warning from History’.
- Robson, J. M. (1990). The fiat and the finger of God: The Bridgewater Treatises. In Helmstadter, R. J., & Lightman, B. (Eds.), *Victorian Crisis in Faith: Essays on Continuity and Change in 19th Century Religious Belief*. Macmillan, Basingstoke.
- Ronald, E. M. A., Sipper, M., & Capcarrère, M. S. (1999). Testing for emergence in artificial life. In Floreano, D., Nicoud, J.-D., & Mondada, F. (Eds.), *Advances in Artificial Life: Fifth European Conference on Artificial Life (ECAL'99)*, pp. 13–20. Springer, Berlin.
- Ryan, M. J. (1990). Sexual selection, sensory systems, and sensory exploitation. *Oxford Survey of Evolutionary Biology*, 7, 157–195.
- Schumpeter, J. A. (1934). *The Theory of Economic Development*. Harvard University Press, Cambridge, MA.
- Sinervo, B., & Lively, C. M. (1996). The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380, 240.
- Steels, L. (1994). The artificial life roots of artificial intelligence. *Artificial Life*, 1(1/2), 75–110.
- Turing, A. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London, Series B*, 237, 37–72.
- Veblen, T. (1898). Why is economics not an evolutionary science? In Lerner, M. (Ed.), *The Portable Veblen*, pp. 215–240. Viking Press, New York. Collection published 1948.
- von Neumann, J., & Burks, A. W. (1966). *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, IL.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior* (Princeton, NJ edition). Princeton University Press.
- Walter, W. G. (1963). *The Living Brain*. W. W. Norton, New York.
- Young, R. M. (1985). *Darwin's Metaphor: Nature's Place in Victorian Culture*. Cambridge University Press, Cambridge.