

Complexity by Subtraction

Daniel W. McShea · Wim Hordijk

Received: 9 January 2013 / Accepted: 20 March 2013
© Springer Science+Business Media New York 2013

Abstract The eye and brain: standard thinking is that these devices are both complex and functional. They are complex in the sense of having many different types of parts, and functional in the sense of having capacities that promote survival and reproduction. Standard thinking says that the evolution of complex functionality proceeds by the addition of new parts, and that this build-up of complexity is driven by selection, by the functional advantages of complex design. The standard thinking could be right, even in general. But alternatives have not been much discussed or investigated, and the possibility remains open that other routes may not only exist but may be the norm. Our purpose here is to introduce a new route to functional complexity, a route in which complexity starts high, rising perhaps on account of the spontaneous tendency for parts to differentiate. Then, driven by selection for effective and efficient function, complexity *decreases* over time. Eventually, the result is a system that is highly functional and retains considerable residual complexity, enough to impress us. We try to raise this alternative route to the level of plausibility as a general mechanism in evolution by describing two cases, one from a computational model and one from the history of life.

Keywords Evolution · Complexity · Constructive neutral evolution · Irreducible complexity · ZFEL

D. W. McShea
Duke University, Durham, NC, USA
e-mail: dmc Shea@duke.edu

W. Hordijk (✉)
SmartAnalytiX.com, Lausanne, Switzerland
e-mail: wim@WorldWideWanderings.net

Introduction

In a famous passage of *On the Origin of Species*, Darwin answers the skeptic's charge that natural selection could not possibly explain the evolution of complex structures such as the eye, what he called "organs of extreme perfection."

To suppose that the eye, with all its inimitable contrivances for adjusting the focus to different distances, for admitting different amounts of light, and for the correction of spherical and chromatic aberration, could have been formed by natural selection, seems, I freely confess, absurd in the highest possible degree. Yet reason tells me, that if numerous gradations from a perfect and complex eye to one very imperfect and simple, each grade being useful to its possessor, can be shown to exist; if further, the eye does vary ever so slightly, and the variations be inherited, which is certainly the case; and if any variation or modification in the organ be ever useful to an animal under changing conditions of life, then the difficulty of believing that a perfect and complex eye could be formed by natural selection, though insuperable by our imagination, can hardly be considered real. (Darwin 1859, pp. 186–187)

The argument is straightforward and sensible. Starting with a simple functional structure—for the eye, he proposed, a nerve ending with a light-sensitive pigment—selection builds up to complex structure by incremental addition. Intermediate stages, in this argument, are not only functional and thus preserved, but are increasingly functional. They are improvements.

Darwin's answer has been a model for evolutionists answering modern challenges to evolution, from the

argument from design to the problem of “irreducible complexity.” For example, skeptics have charged that the bacterial flagellum is so complex—consisting as it does of so many interdependent parts—that it could not possibly have arisen by incremental addition. Intermediates would not have been adaptive, the complaint goes. They would not have been preserved by natural selection. Half a flagellum does not propel a bacterium. The Darwinian reply has been to argue that intermediates could indeed have been adaptive. And the modern argument adds to Darwin’s tactic the possibility of exaptation, or change of function. The intermediates in the incremental build-up to the bacterial flagellum could have functioned for something other than propulsion.

The Darwinian route to complexity works, and indeed it could be right—right in the sense that it could correctly describe the most common route by which complex structures arise in evolution. But there are other possible routes. Here we show how complexity could arise, not by incremental addition but by incremental subtraction. We offer an evolutionary logic in which function arises in structures that are already complex, sometimes more complex than they need to be. Natural selection then favors a reduction in the complexity of these structures. They lose parts, to produce structures that are still functional, sometimes improvements, and often still sporting considerable residual complexity. There is nothing undarwinian about this route. It relies heavily on the principle of natural selection. But as will be seen, it also invokes a second principle to account for the initial complexity, what has been called the zero-force evolutionary law (ZFEL, McShea and Brandon 2010), the spontaneous tendency for parts to differentiate.

We begin with a discussion of complexity, of how we are using the term in this paper and how it can arise via the ZFEL. We then offer an abstract demonstration of our alternative route, showing how complexity by subtraction works in a now-standard computational model, evolving cellular automata. We then examine what we take to be a parallel case in a biological system, the evolution of the vertebrate skull. Finally, we discuss some implications of this alternative route.

Complexity

In his classic treatment of orchids, Darwin refers to “differentiation of parts and consequent complexity of structure” (Darwin 1862, p. 333), making what must have been for him a meaningful connection between the two concepts, differentiation and complexity. A century and a half later, this connection still has strong intuitive appeal in biology. For example, Buchholtz and Wolkovich (2005) adopted it

in their study of whale vertebral columns, measuring complexity using a number of metrics, all of them variance analogues, in other words, functions of the degree of phenotypic differentiation among the vertebral bones (see also McShea 1992, 1993).

Differentiation is complexity in its continuous sense. But it also has a discrete sense, number of part types. For example, in multicellular organisms, variation among cells is often discontinuous, so that the complexity of a multicellular organism at the cell level can be measured as the number of cell types (Bonner 1988; Valentine et al. 1994). McShea (2002) used this measure to assess complexity at the subcellular level, measuring complexity of cells at roughly the organelle level with counts of number of organelle-sized part types. Marcus (2005) counted part types to investigate complexity of microbes. Cisne (1974) and Adamowicz et al. (2008) counted part types in studies of complexity of arthropod limb series. Finally, complexity in the sense of number of part types is becoming the industry standard in molecular biology. For example, Finnigan et al. (2012) used the term “complex” to describe “molecular machines” with more part types (see also Doolittle 2012).

Notice that complexity as number of part types includes no notion of function. This is complexity in what might be called its “pure” sense (McShea and Brandon 2010), uncontaminated with any consideration of the degree of adaptedness, sophistication, or function. It is not that functionality is unimportant. On the contrary, in studies of the evolution of complexity, a central question has to do with a possible connection between complexity and functionality. (Indeed, it is central in this one.) Rather, it is that in order to investigate that connection, it is essential to keep the concepts separate.

Parts and Levels

A technical definition of a part is offered elsewhere (McShea and Venit 2001). For present purposes, it is enough to say that parts are entities that are isolated to some degree from adjacent entities, with a boundary or change in composition serving as an indicator of isolation. Importantly, complexity in the sense of part types is level relative. A fish has about 120 part types at the level of cells (i.e., 120 cell types), but at the tissue/organ level it has only about 90. At the molecular level it has some very large number (the huge number of molecular species present in a fish), while at the atomic level, the number of part-types (i.e., atom types) is probably about a dozen, i.e., the number of elements present in any appreciable number. There is no contradiction here. And there is no privileged level, no level at which a “true” complexity can be measured. In the examples that follow, we have chosen to

measure complexity at certain levels because those levels reveal certain patterns of change that make our point. But it is important to recognize that different patterns of change might occur at other levels.

The Zero-Force Evolutionary Law

In any system with reproduction and heritable variation, the expectation in the absence of opposing forces or limiting constraints is increasing complexity. The reason is that parts in a system will tend, on average, to accumulate variation and therefore tend to become more different from each other. And to the extent that variation is or becomes discontinuous, the number of part types is also expected to increase. The underlying principle is what McShea and Brandon (2010) call the *zero-force evolutionary law* (ZFEL). Essentially, the ZFEL is a general statement of a simple principle that underlies a number of widely acknowledged evolutionary mechanisms, including the tendency for parts to duplicate and differentiate, much discussed in early twentieth century paleontology (Gregory 1934, 1935, and see below), the tendency for left-right asymmetry to rise in bilaterians in the absence of developmental buffering, i.e., fluctuating asymmetry (Van Valen 1962), in the recently much-discussed tendency for duplicate genes to spontaneously differentiate (Taylor and Raes 2004; Lynch 2007), and what is being called constructive neutral evolution (Stoltzfus 1999, and see below; Gray et al. 2010, and see below). The same principle also applies to differentiation among individuals and among taxa, in other words, to diversity, as well as to complexity (McShea and Brandon 2010).

In the ZFEL view, there are two routes by which parts in an organism can come to be different from one another. One is obviously drift, the absence of constraint and selection. The other is selection acting differently on each part. For example, selection on a bipedal primate's hip to improve walking ability and simultaneous selection on the shoulder to improve throwing ability will tend to make hip and shoulder even more different from each other. Thus, whenever parts vary to some degree independently, whether due to drift or to selection acting differently on each, the ZFEL predicts increasing complexity. Two clarifications are needed here. First, the ZFEL is called a "zero-force" principle because it describes what is expected in the absence of selection. But seemingly oddly, the shoulder-hip example invoked selection. The explanation is that the requirement for zero selection applies only to selection acting *on the differences among parts*, on complexity itself, favoring the advantages of differentness. This includes selection favoring differentiation on account of the advantages of division of labor, as doubtless occurred in hominid evolution in the differentiation the thumb from the other

fingers. And it also includes selection opposing differentiation, as when it favors the similarity of the left and right legs. In contrast, in the shoulder-hip example, the assumption is that they become more different not because there is any particular advantage to their being different but simply because they change differently. In that case, complexity rises as the passive result of independent change, not as a result of selection favoring complexity. And it is the presence of any force directly favoring or opposing complexity that the zero-force clause in the ZFEL prohibits. The difference is analogous to the ZFEL-like increase in diversity of locations that occurs when a group of initially clustered individuals each goes his or her own way. They spread out, each under the influence of his or her own will, each person's will an independent force. And yet the zero-force condition is met if there is no force driving them apart, if they are not fleeing each other.

Second, in the history of life there have obviously been many instances when complexity did not increase and many in which complexity decreased. In the ZFEL view, what these instances point to is the existence of either constraints on variation or selective opposition to differentiation. In other words, the ZFEL reverses standard intuitions on complexity: increase is easy. It is the expectation. While stasis and decrease demand opposing constraints and forces (see McShea and Brandon 2010).

The principle is quite general, applicable not just to biological systems but to computer simulations of evolution with reproduction and heredity, to replicating crystals with memory for errors, and even to certain systems without replication but that nevertheless have memory and therefore retain variation (such as the surface of the moon, which retains and accumulates complexity from meteorite impacts).

Complexity is easy. It is spontaneous. No special mechanism beyond the simple tendency for parts to become different from each other is needed to account for it. No selective advantage to complexity needs to be invoked. Of course, selection *can* favor complexity, in which case one would expect differentiation to occur even more quickly, but a rapid rise in complexity is not, all by itself, evidence for a selective advantage. Of course, selection is necessary to explain *adaptive* complexity, to explain why a system with many part types functions, why it *does* something. But not to explain pure complexity, not to explain the existence of many part types.

Early Complexity, Later Reduction

The examples discussed in the next two sections are both systems that begin with many part types. They are unquestionably adaptive, so clearly selection is not entirely absent. But there is no evidence that they are highly adaptive, at least not initially. More precisely, given the

ZFEL tendency for parts to differentiate spontaneously, there is no reason to think that their initial state—characterized by high levels of differentiation—is more than minimally functional.

What there is evidence for, and what we draw attention to, is the reduction in complexity that followed, apparently from selection for improved function, which in turn seems to have required simplification. The resulting structure still has considerable residual complexity. But that complexity was arrived at not by accumulation, not by a build-up from a simple starting condition, not by addition. Rather it was produced by reduction, by building down from an even-more-complex starting condition, by subtraction.

A Computational Example

The following examples are drawn from the *Evolving Cellular Automata* (EvCA) project (Hordijk 2013). In this project, a genetic algorithm was used to evolve cellular automata to perform a non-trivial computational task, with the aim of answering the general question: “How does evolution produce sophisticated emergent computation in systems composed of simple components limited to local interactions?”

Cellular automata are a class of mathematical models of complex systems that consist of a large number of relatively simple components which are limited to local interactions only. Yet they are able to produce a wide variety of intricate patterns in their dynamical behavior which are often considered to be *emergent*, i.e., they are not “programmed” into the simple rules and limited interactions of the underlying individual components, but arise at a higher, global level. A more detailed overview, mostly by example, is given below. A genetic algorithm is a stochastic search and optimization method that is modeled after natural evolution, and is therefore also often used as an actual computer simulation of an evolutionary process. A more detailed description follows below. Combining these two methods provides a useful and versatile computational framework to study the evolution of complexity.

Cellular Automata

Cellular automata were introduced by John von Neumann (after a suggestion by his colleague Stanislaw Ulam) to study the logic of self-reproduction (von Neumann 1966; Burks 1970). They were popularized with John Conway’s “Game of Life” (Gardner 1970). Because of their interesting, sometimes even surprising dynamical behaviors, cellular automata are often used as simple computer models to study pattern formation, self-organization, and emergence.

In its simplest form, a *cellular automaton* (CA) consists of a linear array (or *lattice*) of identical “cells”, each of which can be in one of two states, say zero or one. At each time step (or iteration), all cells simultaneously update their state according to a fixed update rule, depending on their current local “neighborhood configuration” which consists of the cell itself, its left neighbor, and its right neighbor. This update rule simply states for each possible neighborhood configuration what the new state of the center cell will be. Given two possible states and a three-cell neighborhood, there are $2^3 = 8$ possible neighborhood configurations. An example of such an update rule is given in Table 1.

This simplest form of a CA is known as an *elementary cellular automaton* (ECA), and the particular update rule shown in Table 1 is known as ECA 18. Note that the bottom row (new state values) could have any configuration of zeros and ones, giving rise to $2^8 = 256$ possible ECA update rules. Depending on which particular rule is used, the system as a whole (the array of cells), can show very different types of dynamical behaviors, from fixed point or simple periodic, to very complex or even (seemingly) random behavior. Figure 1 shows the dynamical behavior of ECA 18 in a so-called *space-time diagram*, with white representing the state zero and black representing the state one. In this diagram, space is horizontal and time is vertical. The top row (100 cells wide) is the *initial configuration*, which in this case was generated at random. Each next row is the CA configuration after applying the update rule to all cells, for a total of 100 iterations. Note that *periodic boundary conditions* are used, i.e., the array of cells is considered to be circular so that the left-most cell and the right-most cell are each others neighbors.

Of course there are many possible variations on the basic definition of elementary CAs. For example, more than two states can be used, or a larger neighborhood size, or a higher-dimensional array of cells (each of which increases the size of the update rule as well as the total number of possible rules). Other variations such as asynchronous updates, non-uniform update rules, or randomness in the update rule can be included. The possibilities are endless, and so are the range and complexity of corresponding dynamical behaviors.

As Fig. 1 shows, ECA 18 generates intricate patterns in its global dynamical behavior which are at a larger scale than just the local neighborhood configurations. These patterns (and those in many other CAs) are sometimes very

Table 1 An example of an update rule known as elementary CA 18

Neighborhood	000	001	010	011	100	101	110	111
New state	0	1	0	0	1	0	0	0




Fig. 1 A space-time diagram of ECA 18. Space is horizontal (100 cells) and time is vertical (100 iterations), starting from a random initial configuration in the top row. Periodic boundary conditions are used

similar to patterns one observes in natural systems, such as spiral waves, synchronous oscillations, or patterns on sea shells, insect wings, or in animal fur. For this reason, CAs are used frequently as models of (emergent) pattern formation in natural systems (see, e.g., Vichniac 1984; Margolus et al. 1986; Tamayo and Hartman 1988; Manneville et al. 1990; Boerlijst and Hogeweg 1991; Ermentrout and Edelstein-Keshet 1993, and countless more recent publications). In fact, they have even been used to model patterns in road traffic (Simon and Nagel 1998) or social systems (Brown and McBurnett 1996).

CAs are also capable of performing computations. For example, they can perform simple arithmetic, generate pseudo random numbers, and have been shown (in several specific cases) to be capable of universal computation, i.e., they are (theoretically) equivalent in computing power to a Universal Turing Machine (see, e.g., Mitchell 1998 and the many references therein).

Imagine the following computational task for a cellular automaton. Given an initial configuration (IC) of white and black cells, the CA has to decide whether there are more white cells or more black cells in the IC. If there are more white cells, then the CA has to settle down, within a given maximum number of iterations, to an all-white configuration (i.e., all cells becoming white and staying in that configuration in subsequent iterations). Otherwise it has to settle down to an all-black configuration. This task is known as *density classification* (i.e., the CA has to classify the density of black cells in the IC as either below or above 0.5). Note that this is a non-trivial task for a CA, as it requires global information processing, even though each

individual cell can only communicate locally (with its direct neighbors). So, it is not just a matter of “reducing” any IC to an all-white or all-black configuration, but the CA has to “choose” the correct answer state according to a non-trivial property (from the perspective of an individual cell) of the entire IC. This can be compared to, for example, requiring the people in some village to come to a (unanimous) agreement even if they can only communicate with their direct neighbors.

Since there are only 256 elementary CAs, it is easy to check whether any of them can perform this density classification task. However, none of them is capable of doing this. Therefore, a variant of the elementary CA definition is considered here, with the local neighborhood consisting of a cell itself and its nearest *three* neighbors on either side (called a *radius* of three), i.e., a neighborhood of seven cells in total. This gives rise to an update rule with $2^7 = 128$ entries (possible neighborhood configurations), which means there are $2^{128} \approx 3.4 \times 10^{38}$ possible update rules. The question is then: Is there a two-state, radius-three CA (update rule) that can perform the density classification task, and if so, how does it perform the necessary global information processing?

Genetic Algorithms

Given the large number of possible two-state, radius-three CA update rules (too large to do an exhaustive search as with the ECAs), it is useful to apply an automated method that searches through this large space of possible CA rules in an intelligent way to try and find reasonably good solutions. One such method is a genetic algorithm.

A *genetic algorithm* (GA) (Holland 1975; Goldberg 1989; Mitchell 1996) tries to *evolve* better and better solutions to a given (optimization) problem. The idea is to maintain a population of candidate solutions, and create subsequent generations by applying selection and recombination on the individuals in the current population, thus mimicking real evolution. Individuals in the population are assigned a *fitness value* which indicates how well they solve the given problem, and based on which they are then selected to “mate” and create offspring. This simulated evolutionary process generally leads to more and more fit (i.e., better and better) solutions to the given problem. GAs have been used widely and successfully to find good approximate (near-optimal) solutions to problems for which there is no analytical or efficient algorithmic way to find the best possible solution [so-called *NP-complete problems* (Garey and Johnson 1979)].

To apply a genetic algorithm, first the candidate solutions need to be represented by a suitable *genetic encoding*. In the case of two-state cellular automata, there is a natural and straightforward encoding in the form of *bit strings*, i.e.,

strings of zeros and ones. Consider the update rule of ECA 18 as given in Table 1. This rule can simply be represented by the bit string $b = 01001000$, i.e., the bit values in the lower row in the table (which determine the new state of a cell), given a lexicographical ordering of the possible neighborhood configurations. In case of radius-three CAs, this bit string will actually be of length 128 (as calculated above). An initial GA population of candidate solutions is now created by picking a certain number (say 100) of bit strings of length 128 at random (out of the more than 3.4×10^{38} possible ones).

Next, a *fitness function* is needed. This function takes as input an individual from the GA population (in this case a bit string of length 128), translates it into the candidate solution it represents (a CA update rule), and returns a number indicating its *fitness*, i.e., how well it solves the given problem. For example, in the density classification task, the given CA update rule is iterated on, say, 100 random initial configurations (ICs), and the fraction of ICs on which it gives the correct answer (i.e., settles down correctly to all-white or all-black depending on the density of black cells in the IC) is then taken as its fitness value. This way, individuals in the GA population can be directly compared to each other in terms of their fitness.

Once an initial population of candidate solutions with assigned fitness values is created, new generations of individuals are produced by selecting individuals from the current population based on their fitness values, and allowing them to create “offspring”. Here, a very strong form of selection is used, called *elitism*. The individuals in the current population are ranked according to their fitness values (from high to low), and the best 20 individuals are copied to the next generation without modification. Next, pairs of individuals are selected at random, regardless of fitness, but with replacement, from these 20 elite individuals to act as “parents” and create “offspring”.

Offspring individuals are created by randomly recombining the genetic material of a selected pair of parents through *one-point crossover*. For each pair of parents, a random *crossover point* is chosen somewhere between the first and the last bit. The substrings behind this crossover point are then swapped between the two parents. For example, if we have the following pair of parents

$$b_1 = 00000000 \quad \text{and} \quad b_2 = 11111111$$

and the crossover point was randomly chosen, say, between the third and fourth bit, the two offspring individuals will look like

$$b'_1 = 00011111 \quad \text{and} \quad b'_2 = 11100000.$$

Note again that for the CA density classification problem the bit strings are actually of length 128, but the crossover mechanism is the same.

Finally, newly created offspring individuals are subjected to random *mutation*, where a randomly chosen bit is flipped, i.e., a zero is changed to a one or vice versa. For example, given the first offspring individual b'_1 above, and assuming the sixth bit was randomly chosen, it will be mutated to

$$b'_1 = 00011011.$$

Again, for the CA density classification task there are actually 128 bits, and in each mutation event *two* randomly chosen bits are flipped.

This process of parent selection, crossover, and mutation is repeated until a new population of individuals is produced, which will then replace the current population. This whole process is then repeated for a given number of generations. Here we have described the genetic algorithm as used in the specific case of the CA density classification task, but of course there are many ways in which a genetic encoding, fitness function, selection, crossover, and mutation can be implemented, depending on the given optimization problem, and the search performance of a GA can depend strongly on these choices. But the main idea is to search through (or sample) the (generally very large) space of candidate solutions in an intelligent way to try and find good (near-optimal) solutions to the given problem. The GA performs such a search by simulating an evolutionary process with the aim of evolving better and better solutions over time. In short, a GA can generally be described as follows.

1. **Initialization:** Create an initial population of candidate solutions at random, using an appropriate genetic encoding.
2. **Fitness:** Calculate the fitness of each individual in the current population, using an appropriate fitness function.
3. **Selection:** Select individuals, based on their fitness values, to act as parents (possibly using elitism to preserve the current best individuals).
4. **Crossover:** From each next pair of selected parents, create two offspring individuals through crossover.
5. **Mutation:** Apply mutation to the newly created offspring individuals and place them in the new population.
6. **Replacement:** Once the new population is filled up with offspring individuals, replace the current population with the new one and go back to step 2, until a given number of generations is reached.

Note that in the evolving CA case, a bit string in the GA population, representing a CA update rule, can be considered the “genotype”, and the actual dynamical behavior of the corresponding CA can be considered the “phenotype”. So, as in real evolution (to a large extent, at least), the

evolution (random changes) happen at the level of the genotype, but the fitness determination and selection happen at the level of the phenotype. In the CA case, the genotype and phenotype can also be said to be linked through a (possibly complex) “developmental and behavioral process”.

Evolving Cellular Automata with Genetic Algorithms

The first experiments on evolving cellular automata with a genetic algorithm to perform the density classification task were described in Packard (1988), and were later repeated and examined in more detail in Mitchell et al. (1993, 1994a, b). The first real high-performance evolved CAs, using sophisticated emergent computation, were found in a series of subsequent experiments, as described in Das et al. (1994), Crutchfield and Mitchell (1995). We repeated these experiments, and obtained similar results. Here, we use results from the best CA that was found in the experiments reported in Das et al. (1994).

Figure 2 shows space-time diagrams of three CA rules that occurred during the GA run that produced this overall

best CA. Each space-time diagram shows 149 cells across (with periodic boundary conditions) and 149 iterations down the page, starting from a random initial configuration. These three CAs (simply called ϕ_1 , ϕ_2 , and ϕ_3 here) were the best individuals in their respective generations (16, 18, and 63) during the same GA run, and thus are part of a single evolutionary sequence. All three CAs use a similar “strategy” to solve the density classification task, but the actual “implementation” of this strategy improves significantly during their evolution, causing each next CA to have a somewhat higher fitness (i.e., the correct answer is given on a larger fraction of random ICs) than its predecessor.

As Fig. 2 shows, all three CAs quickly settle down into local regions of all-white (W), all-black (B), or a checkerboard ($\#$) pattern. The boundaries between these regions interact with each other, which leads to the annihilation or creation of new stable regions and boundaries, until eventually only one pattern (either W or B) is left as the answer state. Following Hanson and Crutchfield (1992); Crutchfield and Hanson (1993), these local, stable regions are called *regular domains*, and the boundaries between

Fig. 2 Space-time diagrams of ϕ_1 (top-left), ϕ_2 (top-right), and ϕ_3 (bottom)





Fig. 3 Filtered space-time diagrams of ϕ_1 (top-left), ϕ_2 (top-right), and ϕ_3 (bottom)



them *particles* with their *particle interactions*. In Das et al. (1994); Crutchfield and Mitchell (1995) it was argued that it is these particles and their interactions which perform the necessary global information processing to solve the density classification task. This was shown more formally and convincingly in Hordijk et al. (1996, 1998), Hordijk (1999) by modeling the dynamical behavior of evolved CAs at the level of these emergent particles and their interactions.

Figure 3 shows the same three space-time diagrams as in Fig. 2, but with the regular domains (the W , B , and $\#$ patterns) filtered out. This filtering method, formally described in Hanson and Crutchfield (1992), Crutchfield and Hanson (1993), leaves an image that shows the particles and their interactions more clearly and explicitly.

Reduction in Complexity in Evolved CAs

A detailed description and analysis of the particle strategy implemented by these evolved CAs can be found in Das et al. (1994), Crutchfield and Mitchell (1995), Hordijk et al. (1996, 1998), Hordijk (1999), but what is of interest here is that there is reduction in complexity.

Consider the particle types in Figs. 2 and 3 labeled with the letters a , b , and c , respectively. In this labeling, a particle that forms a boundary between a black (B) domain on the left and a white (W) domain on the right, is always considered a particle of type a , even if it might look somewhat different in the three (evolutionary related) CAs. Similarly for particle type b always being a $B\#$ boundary and particle type c always being a $\#W$ boundary.

The top row of Fig. 4 shows one of these particles (type a) enlarged, as it occurs in the three evolved CAs: ϕ_1 (left), ϕ_2 (center), and ϕ_3 (right), respectively. The bottom row shows the corresponding filtered particles.

Recall that complexity is a function of number of part types. At a small scale, each particle is made up entirely of black and white squares and so at the scale of single squares, all have the same complexity, namely two. However, at a larger scale, each particle consists of subsequences of black and white squares. Look closely at the particle's initial structure in the bottom left of Fig. 4. At the time points (rows) where its width is maximal, the particle consists of two 3-square-length subsequences, a BWB subsequence followed by a BBW subsequence. In




Fig. 4 Evolution of particle *a*. *Top row* original particle. *Bottom row* filtered particle

other words, at its maximum width, it has two part types. Compare this with the particle in the bottom middle figure. At any time point (any row, because particle width is the same everywhere), if we resolve the pattern into blocks three squares long, it consists of a WBW subsequence followed by another WBW subsequence. The two are identical, so the number of part types is one. Sampling at different resolutions, i.e., different subsequence lengths, would yield different absolute part-type counts, but on average the complexity of the particle at the bottom left is greater than the bottom middle. In other words, maximum particle width is an increasing function of number of part types, and therefore a good proxy for complexity.

Complexity can also be understood in a temporal sense, as the number of part types a particle contains over time. At issue here is the temporal periodicity of a particle, the number of unique sequences it passes through before repeating, with each unique sequence understood as a part type. The greater the temporal periodicity, the greater the complexity. In biology, an example of an increase in temporal complexity would be an increase in number of stages, or morphs perhaps, in a life cycle.

In the space-time diagrams in Fig. 4, it is clear that particle *a* becomes less complex as it evolves. From ϕ_1 to ϕ_2 , it goes from a maximum width of five cells to a maximum width of four cells, and from a temporal periodicity of three time steps to a temporal periodicity of one (There is no further reduction, as the particle remains the same in ϕ_3 , 43 generations later).

Figure 5 shows similar images for the evolution of particle type *b*. This particle type reduces its temporal periodicity of two and maximum width of six to both a temporal periodicity and maximum width of one (i.e., minimal complexity) in the two generations from ϕ_1 to ϕ_2 .

As with particle type *a*, particle type *b* then also remains the same during the remainder of the evolutionary process.

Finally, Fig. 6 shows the evolution of particle type *c*. Here, the reduction in complexity happens in two stages. First, its temporal periodicity reduces from four to two and its maximum width from 11 to seven, going from ϕ_1 to ϕ_2 . Then there is a further reduction in ϕ_3 to both a temporal periodicity and maximum width of one, where it reaches the minimum possible complexity.

The reason we can consider each of these particle types to be the same across the three different (but similar and evolutionarily related) CAs ϕ_1 , ϕ_2 , and ϕ_3 , is that they form the same domain boundaries in each CA, and they interact with each other in similar ways. For example, particle type *a* forms a boundary between black and white domains in each of the three CAs, and an interaction between particle types *b* and *c* always produces a particle of type *a*. So, even though the exact spatial and temporal structure of a particular particle type can be different between the three CAs (as shown in Figs. 4, 5, and 6), “functionally”, i.e., in terms of the emergent particle strategy for solving the computational task, they are the same.




Fig. 5 Evolution of particle *b*. *Top row* original particle. *Bottom row* filtered particle




Fig. 6 Evolution of particle *c*. *Top row* original particle. *Bottom row* filtered particle

Selective Pressure for Reduced Complexity

This reduction in complexity of the particles is not arbitrary. In fact, there is *selective pressure* during the evolutionary process for them to become less complex. This is a consequence of the way the particles interact with each other. Particles carry information about local densities of black and white cells in different parts of the array of cells, and particle interactions are the loci of exchange and processing of this local information, either by annihilation or by the creation of other particles. This way, an emergent “particle strategy” is implemented to perform the necessary global information processing to successfully perform the density classification task (Das et al. 1994; Crutchfield and Mitchell 1995; Hordijk et al. 1996, 1998; Hordijk 1999).

Obviously, such a particle strategy will be more accurate if the local information is transferred and processed efficiently and without ambiguities. However, more complex particles (i.e., a larger temporal periodicity and larger maximum width) have a higher chance of being less efficient and more ambiguous than less complex particles, simply because they require more time and space to interact with each other. Furthermore, there is a higher chance that two interacting particles get “interrupted” by a third, nearby particle if the particles are more complex than what is minimally required.

Figure 7 shows an explicit example of this. The figure on the left shows a detail of a space-time diagram for ϕ_3 where two particles (a $W\#$ and a $\#W$ boundary) collide and annihilate each other, which results in the disappearance of the $\#$ domain. Just to the left of this particle interaction is a particle of type a , but as long as the distance between the two-particle interaction on the right and the a particle on the left is at least four cells, there is no interference.

The figure on the right shows a similar situation for ϕ_1 . However, here the different particles do interfere with each other due to their higher complexity. If the distance between the two-particle interaction on the right and the a particle on the left is less than seven cells, this




Fig. 7 The two-particle interaction for ϕ_3 (left; no interference) and ϕ_1 (right; interference)

interference will happen, resulting in the a particle to be shifted over to the right, causing a bias towards black domains in the implementation of the particle strategy. So, instead of needing a minimum distance of only four cells, a distance of at least seven cells is required to avoid interference, due to the higher complexity of both the a and the c particle.

The more complex particles, and the consequently less efficient implementation of the particle strategy in ϕ_1 , is one of the causes for its lower fitness for solving the density classification task as compared to ϕ_3 , where the particles have evolved towards lower complexity resulting in a more efficient particle strategy implementation. There are, of course, also other causes for this difference in fitness, but the above example illustrates why and how there does indeed exist selective pressure for particles to become less complex. In Hordijk (1999) several examples are provided of how the difference between more complex and minimally complex (“ideal”) particles can make the difference between a correct and an incorrect answer to the density classification task on a given initial configuration.

A Different Computational Task

To show that the reduction in complexity in these CAs evolved for the density classification task is not, somehow, an artifact of the given task, we show a similar result with evolved CAs for a different task, known as *global synchronization*. In this task, the CA has to settle down, from any initial configuration, to a synchronized oscillation between an all-white configuration and an all-black configuration. This task was first described and analyzed in Das et al. (1995). We also repeated the experiments for this task, and found similar results again. Here, we use results from one of the best evolved CAs from our own experiments.

Figure 8 shows space-time diagrams of two CAs that occurred during one of the runs of the GA on the global synchronization task (again using periodic boundary conditions). The diagram on the left shows one of the best CAs in generation 15 (ϕ_4) while the diagram on the right shows one of the best CAs in the final generation (ϕ_5). As with the density classification CAs, these CAs quickly settle down into local regular domains and particles. In this case the regular domains are the locally synchronized regions (although possibly out of phase with each other) and the regions with the repeating “L-shaped” pattern.

Figure 9 shows the same space-time diagrams, but with the two regular domains filtered out. This again clearly reveals the particles and their interactions. Two particle types (d and e) are labeled in these diagrams. Particle type d exists in ϕ_4 , but not in ϕ_5 . In fact, this particle is an evolutionary “relic” from one of ϕ_4 ’s ancestors, but is not

Fig. 8 Space-time diagrams of ϕ_4 (left) and ϕ_5 (right)





Fig. 9 Filtered space-time diagrams of ϕ_4 (left) and ϕ_5 (right)



playing a useful role in the emergent particle strategy as it evolved in ϕ_4 . During the subsequent course of the evolution, this particle disappears altogether. This shows a reduction in complexity in terms of the number of particle types. Furthermore, particle type e undergoes a drastic reduction in complexity similar to the examples shown above for the density classification CAs. These additional results on a very different task show that reduction in complexity under selective pressure indeed seems to be a more general phenomenon.

A Biological Example

Certain biological systems also show early high levels of complexity, with a subsequent loss of part types. The sequence in Fig. 10 appeared in a 1935 paper by the American paleontologist William Gregory. It purports to show a trend toward reduction in number of skull bones in the evolutionary transitions from fish to amphibian to reptile to mammal. While it is not quite an evolutionary trajectory—since it represents transitions among evolutionary grades of skull organization, rather than an ancestor-descendant sequence—the trend has long been acknowledged to be real. Indeed, a general pattern of

reduction in parts in evolution was recognized in Gregory's time and even before, notably by the paleontologist Samuel Williston. In the early pages of his major work, *Water Reptiles of the Past and Present* (quoted by Gregory in his 1935 paper), Williston wrote:

And it is also a law of evolution that the parts in an organism tend toward reduction in number, with the fewer parts greatly specialized in function, just as the most perfect human machine is that which has the fewest parts and each part most highly adapted to the special function it has to subserve (Williston 1914, p. 3).

In the case of the skull, the reduction was understood to occur by the loss of bones or by the fusion of adjacent ones. And as can be seen in Fig. 10, most of the skull bones are different from each other, and thus represent different part types, even in the earlier part of the sequence (*Eusthenopteron* and *Ichthyostega*). Thus number of bones is well correlated with number of bone types. (Bones that are bilaterally symmetrical, or paired, are obviously quite similar to each other, but they are typically tightly connected in development and share the same evolutionary fate, leaving the correlation intact.) Thus the loss of bones amounts to a reduction in the complexity of the skull.

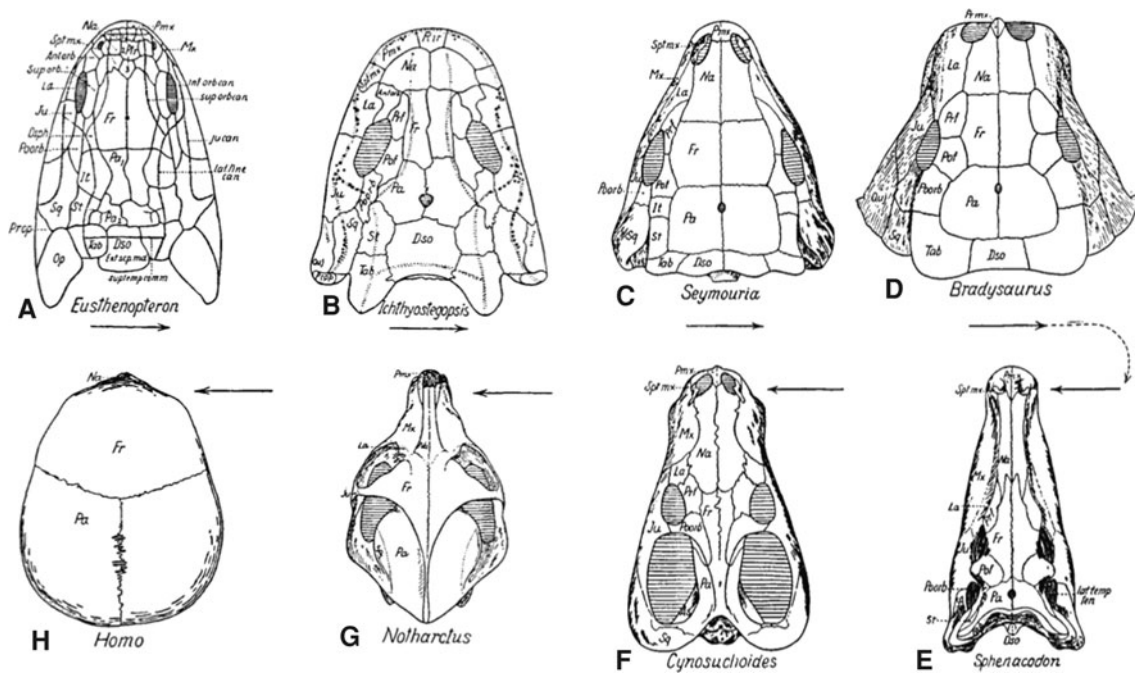


Fig. 10 *Eusthenopteron* is a Devonian lobe-finned fish. *Ichthyostegopsis* (now *Ichthyostega*) and *Seymouria* are Permian (labyrinthodont) amphibians, *Bradysaurus*, *Sphenacodon*, and *Cynosuchooides* (now

Cynosaurus) are Permian reptiles, and *Notharctus* is an Eocene primate. From Gregory (1935)

A Trend in Complexity, Driven by Selection

Almost a century after Williston framed his law of evolution, Christian Sidor took a new look at the trend using current phylogenetic methods (Sidor 2001). He devised what he called a skull simplification metric (SSM), a function of the number of bone types present primitively, and of new bones arising, reduced by the number of losses and fusions. Sidor's study was limited to a subgroup of the tetrapods, the Synapsida, over 150 million years, from the Upper Carboniferous through the Lower Jurassic. Figure 11a shows the trend in SSM for the entire group, plus three modern mammals (top right). Figures 11b–f show the trend for synapsid subgroups. The X axis is SSM, and the Y axis is time—measured discretely as age rank—moving upward toward the present. Notice that SSM is really a complexity measure, increasing with rises in number of bone types and decreasing with losses, but that Sidor has reversed the usual protocol for plotting variables on the X axis. Movement to the right is a decrease in complexity.

The figures reveal not just a trend—a decrease in mean skull complexity (that is, a movement to the right)—but a special kind of trend. The pattern of loss appears to be “driven,” that is, the product of a strong bias in evolution of a sort that is commonly interpreted as a strong selection pressure, one that acted on all or most lineages over the group's history (McShea 1994). The main evidence for a drive is the decrease of the complexity maximum, in other

words the movement of the maximum away from the Y axis and the opening up of an empty region on the left side of the graph at high complexity values. It is as though a strong selective wind were blowing to the right throughout the SSM space, carrying the group like a plume of smoke toward lower complexity.

This is not the only way a long-term trend in the mean could have been produced. In a different world, we might still have observed a trend but it might have been “passive” instead of driven. That is, selection might have favored increases and decreases in complexity equally often over most of the group's history while opposing increases in complexity above the primitive maximum. In that case, mean complexity would still have decreased, but the maximum complexity would not have increased. In other words, the upper left portions of the graphs in Fig. 11a would have been filled in with taxa, rather than empty as they actually are. It is the emptiness of that region that suggests a driven trend, rather than a passive one (McShea 1994). Interestingly, as Sidor observed, the same pattern is repeated in the subgroups (Fig. 11b–f), suggesting that the same drive toward reduced complexity was present over time and over the whole complexity range.

Alternative Interpretations

A passive trend resulting from equal increase and decrease is ruled out, but the data do not by themselves point decisively to a strong selection pressure. Other mechanisms can

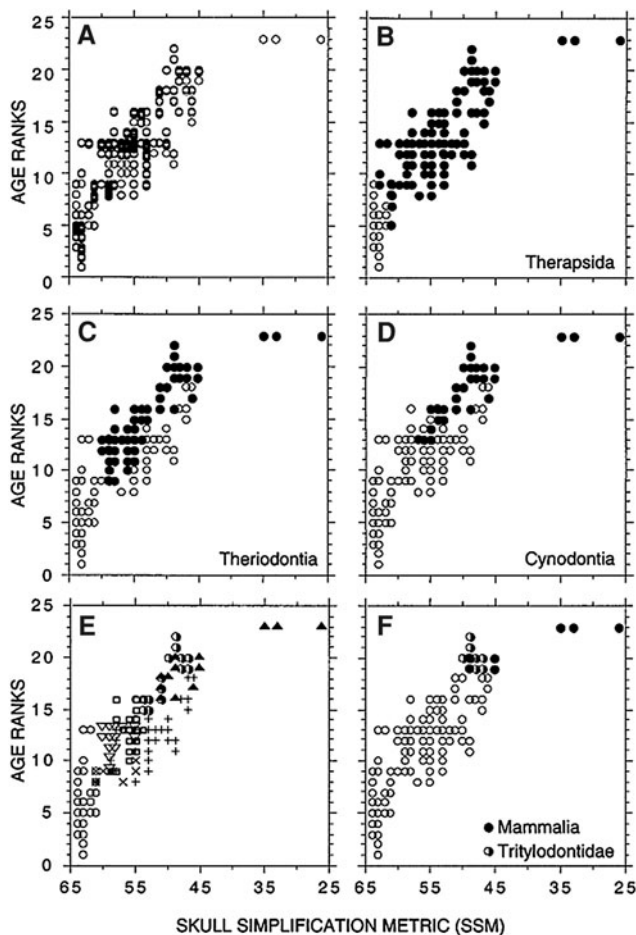


Fig. 11 Data from Sidor (2001) showing the decline in complexity of skulls in synapsids. The *horizontal axis* is his skull simplification metric (SSM), which is directly related to complexity, and the *vertical axis* is “age rank,” a proxy for time, with time moving upward toward the present. **a** Data for all synapsids that Sidor examined; **b–f** filled and half-filled circles correspond to data points for various synapsid subgroups (see Sidor (2001) for the key to e)

produce a similar pattern (Wagner 1996; Alroy 2001), for example, a lower taxon extinction rate, or rising origination rate, at lower complexity values. Also, intrinsic biases—perhaps ease of bone loss in development relative to bone gain—could produce a similar pattern (Sidor 2001). Still, a selection pressure favoring reduced complexity is consistent with expectations based on biomechanical considerations. Sidor writes: “Because sutural joints permit some degree of interbone mobility, reducing skull bone number . . . is a way to solidify the skull against the forces induced by mastication. Continued selection pressure for stronger and more rigid skulls could therefore produce the observed pattern” (Sidor 2001, p. 1430).

Does a decrease in SSM necessarily point to a decrease in skull complexity? Recall that complexity is understood here as number of part types, or where variation is continuous, degree of differentiation. And further recall that the discrete

measure, part types, is a special case of the continuous, differentiation. Both capture differentiation, but they are also to some degree independent. In principle, part types can increase while differentiation decreases, and vice versa. Further, the continuous measure is the preferred one in that it captures differentiation with higher resolution (McShea and Brandon 2010). Thus part types might decrease while differentiation increases or stays the same. And so there is one reason to doubt Sidor’s finding. Adding to the doubt is Gregory’s view that the loss of skull bones is accompanied by an increase in specialization of those that remain, a pattern that Gregory interpreted as an increase in skull complexity. Human skulls, for example, have many fewer bones than fish skulls but those few bones seemed to Gregory to be more different from each other, on average, than fish skull bones are from each other. However, this assessment is entirely impressionistic. To demonstrate it, one would need a morphometric study in which the bones of a skull are plotted in an appropriate morphospace and the complexity of each skull measured as a function of degree of disparity among the bones, perhaps as the variance of their locations in morphospace. In the meantime, Sidor’s finding based on part types and the SSM is the state of the art.

Interestingly, Williston’s Law has been revisited recently from an entirely different perspective. In a fascinating paper, Esteve-Altava et al. (2012) examine the relationship in both living and fossil tetrapods between skull-bone number and degree of connectivity among bones. Connectivity was understood as adjacency—bones in a skull are connected if they contact each other—and degree of connectivity was assessed using three different graph-theoretical metrics. Esteve-Altava et al. found that for all three metrics, connectivity increased as bone number decreased. They describe connectivity as a superior measure of complexity, overcoming the “limitations” of a simple count of bone types. And they concluded that the real trend in skull evolution was an increase in what they call complexity of organization, rather a simplification trend. We have no argument with the use of these metrics or with the results. Indeed they would seem to offer a new and promising perspective on this very old problem. But there is no contradiction between it and the Sidor findings. The reason is that complexity in its non-technical sense is a compound concept, consisting of many logically independent aspects. There is complexity in the sense of number of levels of organization (McShea 2001). There is complexity in the sense of number of different functions or capabilities (McShea 2000). And in the sense of part types. And in the sense of number of different interactions or connections (Esteve-Altava et al. 2012). And many more (McShea 1996). What the Esteve-Altava finding means is that complexity in one technical sense (connections) is negatively correlated with complexity in another technical sense (part types). And while there is something intriguing

about that, something worth investigating further, there is no contradiction. Regardless of what happened to complexity in the sense of connectivity, it is still true, as Sidor demonstrated, that skull complexity in the sense of the part types decreased.

Conclusions and Discussion

The evolutionary trajectories in the computational and skull examples are similar. The system starts with many part types, at high complexity, and driven by selection it loses complexity over time.

In the computational case, it is clear that the original complexity of the system was excessive. Not only can we point to ways in which the complexity of particles hindered their ability to perform the tasks, density classification and global synchronization, but we know that selection was present and that the task did not change. Thus, the ancestral more-complex particles were less fit than the less-complex derived ones. The computational approach has the obvious virtue that evolutionary trajectories and selective forces are perfectly known. The approach also has some obvious limitations. In particular, from a biological perspective, a concern is that the mechanisms of change are imperfect analogs of biological mechanisms. For example, “sexual reproduction” in a genetic algorithm is very different from sexual reproduction in real (biological) organisms and populations. On the other hand, at a higher level of abstraction, the dissimilarity among mechanisms may not matter much. We have a system evolving to perform a particular task and in doing so its trajectory spontaneously took it from high complexity to low, which at least raises the possibility that such a trajectory might be available to any selection-driven evolutionary system, including biological systems. At worst, the computational case offers a proof of concept, a demonstration that early excessive complexity followed by adaptive reduction is a possible route to adaptation. More optimistically, it suggests a route that may be generally available to evolving systems.

It is also worth noting that the computational case provides a versatile testbed to perform additional experiments and investigate the relevant issues in more detail. Experiments can be repeated under different circumstances (parameter settings) to study the influence of various factors on the occurrences of reduction in complexity. For example, the selective pressure can be varied by changing the selection operator in the genetic algorithm, or by incorporating additional measures in the fitness function, such as the number of iterations it takes the CA to get to an answer state. Also, statistics can be collected on how often a reduction in complexity is observed (over a number of simulation runs), or how many generations it takes on

average for particles to reach a minimal level of complexity while maintaining the same functionality. Such experiments and statistics are difficult, if not impossible, to perform and obtain in biological systems.

The skull simplification case cannot be interpreted with the same confidence of course. Part type counts may not reflect actual phenotypic disparity among skull elements. The pattern of change suggests a selective driving force toward simplification but there are other possibilities. And if selection is in fact involved, we do not know for certain that it was a single selection pressure acting more or less continuously over hundreds of millions of years. The task, so to speak, may have changed. It could be that the late-Paleozoic tetrapod skull was optimal for tetrapods needs at the time and those needs happened to require more part types than skulls in later contexts. In that case, we cannot say that the original complexity was excessive. On the other hand, part-type counts do offer the best estimate currently available of skull complexity, and they clearly indicate a decrease in skull complexity. And the selective story is buttressed by the known biomechanical advantages of a less complex, more rigid skull, arguably advantageous for chewing across a wide range of diets and ecological contexts. This combined with the fact that the trend appears to have been driven not only in the group as a whole but in subgroups distributed over the group’s history, suggests a common cause. At worst, the case for early excessive complexity followed by adaptive reduction is plausible and consistent with what is known.

Complexity by Subtraction

In the passage from the *Origin* quoted earlier, Darwin was addressing the famous argument from design, as well known in his time as in our own (Dembski and Ruse 2004). The challenge was, and remains: how to explain the evolution of structures with many parts from simpler ancestral ones, especially in those cases in which the gradual addition of novel parts seems improbable, because it is hard to imagine how the intermediates could have been functional. The standard answer from evolutionary biology—from Darwin to the present—has been that complexity did in fact arise by stepwise addition of parts from a primitively simple ancestral condition, and that the intermediates were in fact functional, although importantly, the function of these increasingly complex structures often changed in the course of the process.

This answer could be right, and not just to explain the occasional complex structure, but in general. We do not know. Our point here is that there is an alternative, initial high complexity followed by loss, complexity by subtraction. Its advantage as an explanation for complexity is that the problem of nonfunctional intermediates does not arise.

Fig. 12 Complexity by subtraction. From *left to right*: 1 Some process—perhaps the ZFEL—produces part-type excess. 2 Functional structures are present among a subset of the parts. 3 & 4 Selection pares away the excess





Fig. 13 Two views of complexity over time in evolution: **a** standard (*left*), monotonic increase toward adaptive peak; **b** complexity by subtraction (*right*), rapid increase to adaptive excess, followed by decrease toward optimum

Consider a cartoon version of the problem, the construction of a stone arch. Arches are stable and weight supporting only when completed, when the keystone is finally lowered into place. And so in the standard method of building an arch, intermediates require scaffolding for support. But there is another route, shown in Fig. 12. Start with a large pile of stones of various shapes. (Suppose that stone types are cheap and easy, available for free through some common physical process, just as part types are available via the ZFEL.) Within a sufficiently large pile, weight supporting structures are likely to be present. The engineer's job, then, is not to build an arch out of stones but to remove the excess, the stones that do not participate in the already existing arch (and perhaps to reshape the remaining stones). The resulting structure is still complex, although obviously reduced from what might be called the "excessive complexity" of the structure it arose from.

The Trajectory of Complexity

Figure 13 contrasts the standard view of the trajectory of complexity with our proposal here. In the standard view (Fig. 13a), a structure starts simple and becomes more complex with the gradual addition of new part types, rising to some maximum that presumably marks optimal efficiency or efficacy. Two features distinguish the alternative (Fig. 13b). First, the process begins with a simple structure, as in the standard view, but the rise in complexity is relatively rapid. This is plausible, we argue, because complexity in the sense of part types does not require an

extended process of variation and selection. Indeed, a rapid rise in complexity is the expectation in any system in which selection is absent or relaxed, or in which selection acts to some degree independently on each part (the conditions for the ZFEL). Second, as the graph shows, complexity declines from its initial high, leveling off at a point where complexity is lower than its initial maximum but still higher than the ancestral starting point, again presumably at a point corresponding to some efficiency or efficacy optimum.

The route to complex functionality that we propose is not entirely new. It may be present implicitly in the literature in biology on self-organization (e.g. Kauffman 1996). Consider Depew and Weber's reply to the intelligent-design argument, where they write (in a footnote) of "the emergence of machines that involve some self-assembly, combined with paring away of less-than-fully-efficient parts." (Weber and Depew 2004, p. 186). Also, a critical component of the route we propose is implicit in the current molecular evolution literature, in the notion that gene duplication commonly leads to neutral differentiation and therefore to a spontaneous rise in number of gene part types, on which selection acts to produce function (what is now called either neofunctionalization or subfunctionalization) (Lynch 2007). Finally, complexity by subtraction shares an intuition that is present in the recent literature on constructive neutral evolution (Stoltzfus 1999), the idea that, as Gray et al. write, "Many of the cell's macromolecular machines appear gratuitously complex, comprising more components than their basic functions seem to demand" (Gray et al. 2010, p. 920).

We conclude by noting a perhaps-obvious consequence of the complexity-by-subtraction view and by posing some questions, rhetorically, in the hope of inspiring others to pursue answers. The consequence is that, if complexity-by-subtraction is the rule in evolution, then the complexity of functional biological devices is merely residual. And that residual complexity is to some degree a secondary effect of the route taken, not necessarily favored in its own right. Indeed, what is favored is streamlined simplicity. If functional structures are complex, it may be in part because they start that way, because initial complexity is easy.

The questions we raise have to do with the generality of the route we propose. First, there is the question of how

commonly it occurs in evolution, and then, whether and how commonly it occurs in non-biological systems. Do machines typically evolve along similar lines, starting with excess complexity and becoming simpler as they are improved? Do languages start complex and become more streamlined over time, under pressure perhaps for fast and efficient communication? How about the complexity of human institutions, such as businesses, that experience pressures at least analogous to natural selection for improved functionality? Do they follow a similar trajectory? As in the evolutionary case, a case-study approach of the sort we attempt here would be useful as a first step. What is needed then—for the evolutionary case as well—is a broader and more systematic study, to discover whether or the extent to which complexity-by-subtraction occurs generally.

Acknowledgments The main ideas described in this paper originated at a catalysis meeting at the National Evolutionary Synthesis Center (NESCent) in Durham, NC, USA. They were developed further and finalized into the current paper during a subsequent short-term research visit of WH at, and supported by, NESCent. We thank Robert Brandon for suggesting the apt and evocative phrase “complexity by subtraction.” Finally, one of us (DM) would like to thank Benedikt Hallgrímsson for discussions decades ago, discussions that turned out to be foundational in the development of the ZFEL.

References

- Adamowicz, S. J., Purvis, A., & Wills, M. A. (2008). Increasing morphological complexity in multiple parallel lineages of the Crustacea. *Proceedings of the National Academy of Sciences*, *105*, 4786–4791.
- Alroy, J. (2001). Understanding the dynamics of trends within evolving lineages. *Paleobiology*, *26*, 319–329.
- Boerlijst, M., & Hogeweg, P. (1991). Self-structuring and selection: Spiral waves as a substrate for prebiotic evolution. In C. G. Langton, C. Taylor, J. D. Farmer & S. Rasmussen (Eds.), *Artificial life II* (pp. 55–276). Reading: Addison-Wesley.
- Bonner, J. T. (1988). *The evolution of complexity by means of natural selection*. Princeton: Princeton University Press.
- Brown, T. A., & McBurnett, M. D. (1996). The emergence of political elites. In M. Coombs & M. Sulcoski (Eds.), *Proceedings of the International Workshop on Control Mechanisms for Complex Systems* (pp. 143–161).
- Buchholtz, E. A., & Wolkovich, E. H. (2005). Vertebral osteology and complexity in *Lagenorhynchus acutus*. *Marine Mammal Science*, *21*, 411–428.
- Burks, A. W. (Ed) (1970). *Essays on cellular automata*. Urbana: University of Illinois Press.
- Cisne, J. L. (1974). Evolution of the world fauna of aquatic free-living arthropods. *Evolution*, *28*, 337–366.
- Crutchfield, J. P., & Hanson, J. E. (1993). Turbulent pattern bases for cellular automata. *Physica D*, *69*, 279–301.
- Crutchfield, J. P., & Mitchell, M. (1995). The evolution of emergent computation. *Proceedings of the National Academy of Sciences*, *92*(23), 10742–10746.
- Darwin, C. (1859). *On the origin of species*. London: J. Murray.
- Darwin, C. (1862). *On the various contrivances by which British and foreign orchids are fertilised by insects, and on the food effects of intercrossing*. London: J. Murray.
- Das, R., Mitchell, M., Crutchfield, J. P. (1994). A genetic algorithm discovers particle-based computation in cellular automata. In Y. Davidor, H. P. Schwefel & R. Manner (Eds.), *Parallel problem solving from nature—PPSN III* (pp. 344–353). Berlin: Springer.
- Das, R., Crutchfield, J. P., Mitchell, M., & Hanson, J. E. (1995). Evolving globally synchronized cellular automata. In L. J. Eshelman (Ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms* (pp. 336–343). Los Altos: Morgan Kaufmann.
- Dembski, W. A., & Ruse, M. (Eds.) (2004). *Debating design*. Cambridge: Cambridge University Press.
- Doolittle, W. F. (2012). A ratchet for protein complexity. *Nature*, *481*, 270–271.
- Ermentrout, G. B., & Edelstein-Keshet, L. (1993). Cellular automata approaches to biological modeling. *Journal of Theoretical Biology*, *160*, 97–133.
- Esteve-Altava, B., Marugán-Lobón, J., Botella, H., & Rasskin-Gutman, D. (2012). Structural constraints in the evolution of the tetrapod skull complexity: Williston's Law revisited using network models. *Evolutionary Biology*. doi:10.1007/s11692-012-9200-9.
- Finnigan, G. C., Hanson-Smith, V., Stevens, T. H., & Thornton, J. W. (2012). Evolution of increased complexity in a molecular machine. *Nature*, *481*, 360–364.
- Gardner, M. (1970). The fantastic combinations of John Conway's new solitaire game “life”. *Scientific American*, *223*(120), 123.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. New York: W. H. Freeman.
- Goldberg, D. E. (1989) Genetic algorithms in search, optimization, and machine learning. Reading: Addison-Wesley.
- Gray, M. W., Lukeš, J., Archibald, J. M., Keeling, P. J., & Doolittle, W. F. (2010). Irremediable complexity? *Science*, *330*, 920–921.
- Gregory, W. K. (1934). Polyisomerism and anisomerism in cranial and dental evolution among vertebrates. *Proceedings of the National Academy of Sciences*, *20*, 1–9.
- Gregory, W. K. (1935). Reduplication in evolution. *Quarterly Review of Biology*, *10*, 272–290.
- Hanson, J. E., & Crutchfield, J. P. (1992). The attractor-basin portrait of a cellular automaton. *Journal of Statistical Physics*, *66*(5/6), 1415–1462.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor: University of Michigan Press (2nd edn., MIT Press, 1992).
- Hordijk, W. (1999). *Dynamics, emergent computation, and evolution in cellular automata*. PhD thesis, Albuquerque, NM, USA: University of New Mexico.
- Hordijk, W. (2013). *The EvCA project: A brief history*. Complexity (To appear).
- Hordijk, W., Crutchfield, J. P., & Mitchell, M. (1996). Embedded particle computation in evolved cellular automata. In T. Toffoli, M. Biafore & J. Leão (Eds.), *Proceedings of the Conference on Physics and Computation* (pp. 153–158). Cambridge: New England Complex Systems Institute.
- Hordijk, W., Crutchfield, J. P., & Mitchell, M. (1998). Mechanisms of emergent computation in cellular automata. In A. E. Eiben, T. Bäck, M. Schoenauer & H. P. Schwefel (Eds.), *Parallel Problem Solving from Nature-V* (pp. 613–622). New York: Springer.
- Kauffman, S. A. (1996). *At home in the universe*. Oxford: Oxford University Press.
- Lynch, M. (2007). The frailty of adaptive hypotheses for the origins of organismal complexity. *Proceedings of the National Academy of Sciences*, *104*, 8597–8604.
- Manneville, P., Boccaro, N., Vichniac, G. Y., & Bidaux, R. (1990). *Cellular automata and modeling of complex physical systems*, volume 46 of Springer Proceedings in Physics. New York: Springer.

- Marcus, J. M. (2005). A partial solution to the C-value paradox. Lecture Notes in Computer Science, p 3678.
- Margolus, N., Toffoli, T., & Vichniac, G. (1986). Cellular-automata supercomputers for fluid-dynamics modeling. *Physical Review Letters*, *56*(16), 1694–1696.
- McShea, D. W. (1992). A metric for the study of evolutionary trends in the complexity of serial structures. *Biological Journal of the Linnean Society*, *45*, 39–55.
- McShea, D. W. (1993). Evolutionary changes in the morphological complexity of the mammalian vertebral column. *Evolution*, *47*, 730–740.
- McShea, D. W. (1994). Mechanisms of large-scale evolutionary trends. *Evolution*, *48*, 1747–1763.
- McShea, D. W. (1996). Metazoan complexity and evolution: Is there a trend? *Evolution*, *50*, 477–492.
- McShea, D. W. (2000). Functional complexity in organisms: Parts as proxies. *Biology and Philosophy*, *15*, 641–668.
- McShea, D. W. (2001). The hierarchical structure of organisms: A scale and documentation of a trend in the maximum. *Paleobiology*, *27*, 405–423.
- McShea, D. W. (2002). A complexity drain on cells in the evolution of multicellularity. *Evolution*, *56*, 441–452.
- McShea, D. W., & Brandon, R. N. (2010). *Biologys first law*. Chicago: University of Chicago Press.
- McShea, D. W., & Venit, E. P. (2001). What is a part? In G. P. Wagner (Ed.), *The character concept in evolutionary biology* (pp. 259–284). New York: Academic Press.
- Mitchell, M. (1996). *An introduction to genetic algorithms*. Cambridge: MIT Press.
- Mitchell, M. (1998). Computation in cellular automata: A selected review. In T. Gramss, S. Bornholdt, M. Gross, M. Mitchell & T. Pellizzari (Eds.), *Nonstandard computation*. Weinheim: VCH Verlagsgesellschaft.
- Mitchell, M., Hraber, P. T., & Crutchfield, J. P. (1993). Revisiting the edge of chaos: Evolving cellular automata to perform computations. *Complex Systems*, *7*, 89–130.
- Mitchell, M., Crutchfield, J. P., & Hraber, P. T. (1994a). Dynamics, computation, and the “edge of chaos”: A re-examination. In G. A. Cowan, D. Pines & D. Melzner (Eds.), *Complexity: Metaphors, Models, and Reality* (pp. 497–513). Reading: Addison-Wesley. Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume 19.
- Mitchell, M., Crutchfield, J. P., & Hraber, P. T. (1994b). Evolving cellular automata to perform computations: Mechanisms and impediments. *Physica D*, *75*, 361–391.
- Packard, N. H. (1988). Adaptation toward the edge of chaos. In J. A. S. Kelso, A. J. Mandell & M. F. Shlesinger (Eds.), *Dynamic patterns in complex systems* (pp. 293–301). Singapore: World Scientific.
- Sidor, C. A. (2001). Simplification as a trend in synapsid cranial evolution. *Evolution*, *55*, 1142–1419.
- Simon, P. M., & Nagel, K. (1998). Simplified cellular automaton model for city traffic. *Physical Review E*, *58*(2), 1286–1295.
- Stoltzfus, A. (1999). On the possibility of constructive neutral evolution. *Journal of Molecular Evolution*, *49*, 169–181.
- Tamayo, P., & Hartman, H. (1988). Cellular automata, reaction-diffusion systems and the origin of life. In C. G. Langon (Ed.), *Artificial life* (pp. 105–124). Reading: Addison-Wesley.
- Taylor, J. S., & Raes, J. (2004). Duplication and divergence: The evolution of new genes and old ideas. *Annual Review of Genetics*, *38*, 615–643.
- Valentine, J. W., Collins, A. G., & Meyer, C. P. (1994). Morphological complexity increase in metazoans. *Paleobiology*, *20*, 131–142.
- Van Valen, L. (1962). A study of fluctuating asymmetry. *Evolution*, *16*, 125–142.
- von Neumann, J. (1966). *Theory of self-reproducing automata*. In A. W. Burks (Ed.), Urbana: University of Illinois Press.
- Vichniac, G. Y. (1984). Simulating physics with cellular automata. *Physica D*, *10*, 96–116.
- Wagner, P. J. (1996). Testing the underlying patterns of active trends. *Evolution*, *50*, 990–1017.
- Weber, B. H., & Depew, D. J. (2004). Darwinism, design, and complex systems dynamics. In W. A. Dembski & M. Ruse (Eds.), *Debating design* (pp. 173–190). Cambridge: Cambridge University Press.
- Williston, S. W. (1914). *Water reptiles of the past and present*. Chicago: University of Chicago Press.