EXCERPTED FROM

STEPHEN WOLFRAM
A NEW KIND OF SCIENCE

SECTION 8.5

Fundamental Issues in Biology

Fundamental Issues in Biology

Biological systems are often cited as supreme examples of complexity in nature, and it is not uncommon for it to be assumed that their complexity must be somehow of a fundamentally higher order than other systems.

And typically it is thought that this must be a consequence of the rather unique processes of adaptation and natural selection that operate in biological systems. But despite all sorts of discussion over the years, no clear understanding has ever emerged of just why such processes should in the end actually lead to much complexity at all.

And in fact what I have come to believe is that many of the most obvious examples of complexity in biological systems actually have very little to do with adaptation or natural selection. And instead what I suspect is that they are mainly just another consequence of the very basic phenomenon that I have discovered in this book in the context of simple programs: that in almost any kind of system many choices of underlying rules inevitably lead to behavior of great complexity.

The general idea of thinking in terms of programs is, if anything, even more obvious for biological systems than for physical ones. For in a physical system the rules of a program must normally be deduced indirectly from the laws of physics. But in a biological organism there is genetic material which can be thought of quite directly as providing a program for the development of the organism.

Most of the programs that I have discussed in this book, however, have been very simple. Yet the genetic program for every biological organism known today is long and complicated: in humans, for example, it presumably involves millions of separate rules—making it by most measures as complex as large practical software systems like *Mathematica*.

So from this one might think that the complexity we see in biological organisms must all just be a reflection of complexity in their underlying rules—making discoveries about simple programs not really relevant. And certainly the presence of many different types of organs and other elements in a typical complete organism seems likely to be related to the presence of many separate sets of rules in the underlying program. But what if one looks not at a complete organism but instead just at some part of an organism?

Particularly on a microscopic scale, the forms one sees are often highly regular and quite simple, as in the pictures on the facing page. And when one looks at these, it seems perfectly reasonable to suppose that they are in effect produced by fairly simple programs.

But what about the much more complicated forms that one sees in biological systems? On the basis of traditional intuition one might assume that such forms could never be produced by simple programs. But from the discoveries in this book we now know that in fact it is possible to get remarkable complexity even from very simple programs.

So is this what actually happens in biological systems?

There is certainly no dramatic difference between the underlying types of cells or other elements that occur in complex biological forms and in the forms on the facing page. And from this one might begin to suspect that in the end the kinds of programs which generate all these forms are quite similar—and all potentially rather simple.

For even though the complete genetic program for an organism is long and complicated, the subprograms which govern individual aspects of an organism can still be simple—and there are now plenty of specific simple examples where this is known to be the case. But still one might assume that to get significant complexity would require something more. And indeed at first one might think that it would never really be possible to say much at all about complexity just by looking at parts of organisms.

But in fact, as it turns out, a rather large fraction of the most obvious examples of biological complexity seem to involve only surprisingly limited parts of the organisms. Elaborate pigmentation patterns, for instance, typically exist just on an outer skin, and are made up of only a few types of cells. And the vast majority of complicated

Examples of highly regular forms occurring in biological systems. Most of these forms are simple enough that it seems immediately plausible that they could in effect be generated by simple programs. The majority show either simple geometrical shapes, or repetition of identical elements. A few, however, show various types of nesting. Note that there seems to be no obvious correlation between the sophistication of a form and when in geological time it first appeared. ▸

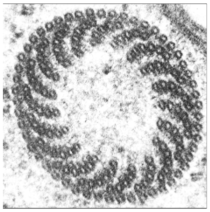
protist microspine section

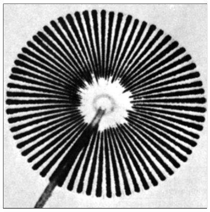
alga

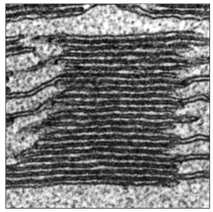
chloroplast

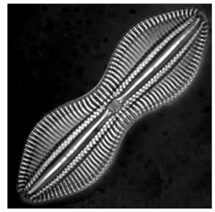
diatom


coccolithophorid

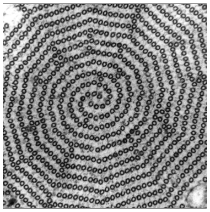
heliozoan axopod

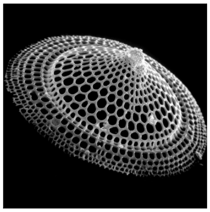
radiolarian


thallus

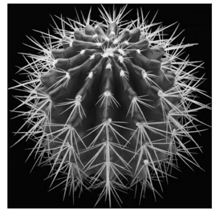
cactus

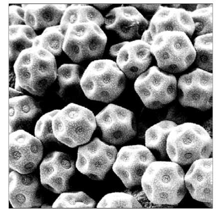
chickweed pollen

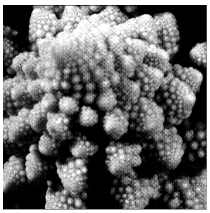
romanesco broccoli

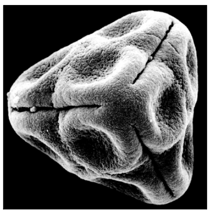
heather pollen

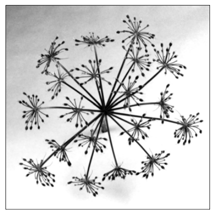
cow parsley


carrot leaf

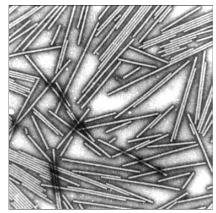
tobacco mosaic viruses

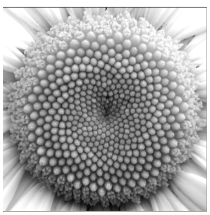
daisy

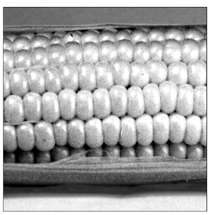
corn

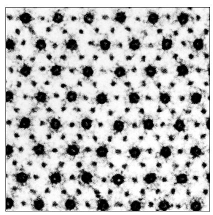
insect muscle section

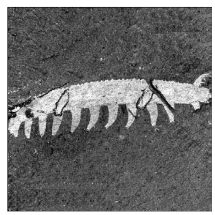
Burgess Shale fossil

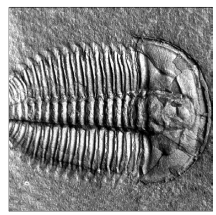
trilobite


barnacle

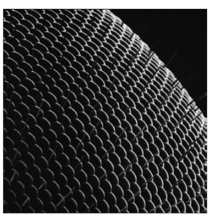
fly eye

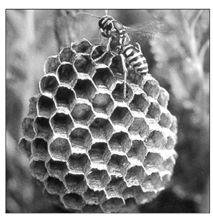
wasp nest

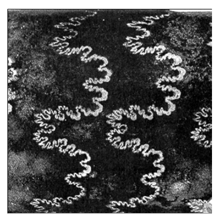
fossil ammonite septa


nautilus shell section


octopus tentacle


opened brachiopod

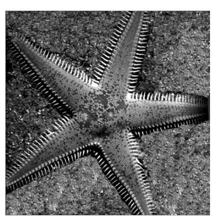
starfish

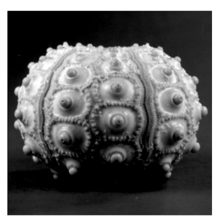
sea urchin

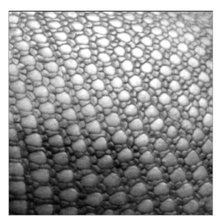
armadillo skin

morphological structures get their forms from arrangements of very limited numbers of types of cells or other elements.

But just how are the programs for these and other features of organisms actually determined? Over the past century or so it has become almost universally believed that at some level these programs must end up being the ones that maximize the fitness of the organism, and the number of viable offspring it produces.

The notion is that if a line of organisms with a particular program typically produce more offspring, then after a few generations there will inevitably be vastly more organisms with this program than with other programs. And if one assumes that the program for each new offspring involves small random mutations then this means that over the course of many generations biological evolution will in effect carry out a random search for programs that maximize the fitness of an organism.

But how successful can one expect such a search to be?

The problem of maximizing fitness is essentially the same as the problem of satisfying constraints that we discussed at the end of Chapter 7. And what we found there is that for sufficiently simple constraints—particularly continuous ones—iterative random searches can converge fairly quickly to an optimal solution. But as soon as the constraints are more complicated this is no longer the case. And indeed even when the optimal solution is comparatively simple it can require an astronomically large number of steps to get even anywhere close to it.

Biological systems do appear to have some tricks for speeding up the search process. Sexual reproduction, for example, allows large-scale mixing of similar programs, rather than just small-scale mutation. And differentiation into organs in effect allows different parts of a program to be updated separately. But even with a whole array of such tricks, it is still completely implausible that the trillion or so generations of organisms since the beginning of life on Earth would be sufficient to allow optimal solutions to be found to constraints of any significant complexity.

And indeed one suspects that in fact the vast majority of features of biological organisms do not correspond to anything close to optimal solutions: rather, they represent solutions that were fairly easy to find, but are good enough not to cause fatal problems for the organism.

The basic notion that organisms tend to evolve to achieve a maximum fitness has certainly in the past been very useful in providing a general framework for understanding the historical progression of species, and in yielding specific explanations for various fairly simple properties of particular species.

But in present-day thinking about biology the notion has tended to be taken to an extreme, so that especially among those not in daily contact with detailed data on biological systems it has come to be assumed that essentially every feature of every organism can be explained on the basis of it somehow maximizing the fitness of the organism.

It is certainly recognized that some aspects of current organisms are in effect holdovers from earlier stages in biological evolution. And there is also increasing awareness that the actual process of growth and development within an individual organism can make it easier or more difficult for particular kinds of structures to occur.

But beyond this there is a surprisingly universal conviction that any significant property that one sees in any organism must be there because it in essence serves a purpose in maximizing the fitness of the organism.

Often it is at first quite unclear what this purpose might be, but at least in fairly simple cases, some kind of hypothesis can usually be constructed. And having settled on a supposed purpose it often seems quite marvellous how ingenious biology has been in finding a solution that achieves that purpose.

Thus, for example, the golden ratio spiral of branches on a plant stem can be viewed as a marvellous way to minimize the shading of leaves, while the elaborate patterns on certain mollusc shells can be viewed as marvellous ways to confuse the visual systems of supposed predators.

But it is my strong suspicion that such purposes in fact have very little to do with the real reasons that these particular features exist. For instead, as I will discuss in the next couple of sections, what I believe is that these features actually arise in essence just because they are easy to produce with fairly simple programs. And indeed as one looks at more and more complex features of biological organisms—notably texture and pigmentation patterns—it becomes increasingly difficult to find any credible purpose at all that would be served by the details of what one sees.

In the past, the idea of optimization for some sophisticated purpose seemed to be the only conceivable explanation for the level of complexity that is seen in many biological systems. But with the discovery in this book that it takes only a simple program to produce behavior of great complexity, a quite different—and ultimately much more predictive—kind of explanation immediately becomes possible.

In the course of biological evolution random mutations will in effect cause a whole sequence of programs to be tried. And the point is that from what we have discovered in this book, we now know that it is almost inevitable that a fair fraction of these programs will yield complex behavior.

Some programs will presumably lead to organisms that are more successful than others, and natural selection will cause these programs eventually to dominate. But in most cases I strongly suspect that it is comparatively coarse features that tend to determine the success of an organism—not all the details of any complex behavior that may occur.

Thus in a very simple case it is easy to imagine for example that an organism might be more likely to go unnoticed by its predators, and thus survive and be more successful, if its skin was a mixture of brown and white, rather than, say, uniformly bright orange. But it could then be that most programs which yield any mixture of colors also happen to be such that they make the colors occur in a highly complex pattern.

And if this is so, then in the course of random mutation, the chances are that the first program encountered that is successful enough to survive will also, quite coincidentally, exhibit complex behavior.

On the basis of traditional biological thinking one would tend to assume that whatever complexity one saw must in the end be carefully crafted to satisfy some elaborate set of constraints. But what I believe instead is that the vast majority of the complexity we see in biological systems actually has its origin in the purely abstract fact that among randomly chosen programs many give rise to complex behavior.

In the past it tends to have been implicitly assumed that to get substantial complexity in a biological system must somehow be fundamentally very difficult. But from the discoveries in this book I have come to the conclusion that instead it is actually rather easy.

So how can one tell if this is really the case?

One circumstantial piece of evidence is that one already sees considerable complexity even in very early fossil organisms. Over the course of the past billion or so years, more and more organs and other devices have appeared. But the most obvious outward signs of complexity, manifest for example in textures and other morphological features, seem to have already been present even from very early times.

And indeed there is every indication that the level of complexity of individual parts of organisms has not changed much in at least several hundred million years. So this suggests that somehow the complexity we see must arise from some straightforward and general mechanism—and not, for example, from a mechanism that relies on elaborate refinement through a long process of biological evolution.

Another circumstantial piece of evidence that complexity is in a sense easy to get in biological systems comes from the observation that among otherwise very similar present-day organisms features such as pigmentation patterns often vary from quite simple to highly complex.

Whether one looks at fishes, butterflies, molluscs or practically any other kind of organism, it is common to find that across species or even within species organisms that live in the same environment and have essentially the same internal structure can nevertheless exhibit radically different pigmentation patterns. In some cases the patterns may be simple, but in other cases they are highly complex.

And the point is that no elaborate structural changes and no sophisticated processes of adaptation seem to be needed in order to get these more complex patterns. And in the end it is, I suspect, just that some of the possible underlying genetic programs happen to produce complex patterns, while others do not.

Two sections from now I will discuss a rather striking potential example of this: if one looks at molluscs of various types, then it turns out that the range of pigmentation patterns on their shells corresponds remarkably closely with the range of patterns that are produced by simple randomly chosen programs based on cellular automata.

And examples like this—together with many others in the next couple of sections—provide evidence that the kind of complexity we see in biological organisms can indeed successfully be reproduced by short and simple underlying programs. But there still remains the question of whether actual biological organisms really use such programs, or whether somehow they instead use much more complicated programs.

Modern molecular biology should soon be able to isolate the specific programs responsible, say, for the patterns on mollusc shells, and see explicitly how long they are. But there are already indications that these programs are quite short.

For one of the consequences of a program being short is that it has little room for inessential elements. And this means that almost any mutation or change in the program—however small—will tend to have a significant effect on at least the details of patterns it produces.

Sometimes it is hard to tell whether changes in patterns between organisms within a species are truly of genetic origin. But in cases where they appear to be it is common to find that different organisms show a considerable variety of different patterns—supporting the idea that the programs responsible for these patterns are indeed short.

So what about the actual process of biological evolution? How does it pick out which programs to use? As a very simple idealization of biological evolution, one can consider a sequence of cellular automaton programs in which each successive program is obtained from the previous one by a random mutation that adds or modifies a single element.

The pictures on the facing page then show a typical example of what happens with such a setup. If one starts from extremely short programs, the behavior one gets is at first quite simple. But as soon as the underlying programs become even slightly longer, one immediately sees highly complex behavior.

Traditional intuition would suggest that if the programs were to become still longer, the behavior would get ever richer and more complex. But from the discoveries in this book we know that this will not in general be the case: above a fairly low threshold, adding complexity to an underlying program does not fundamentally change the kind of behavior that it can produce.

And from this one concludes that biological systems should in a sense be capable of generating essentially arbitrary complexity by using short programs formed by just a few mutations.

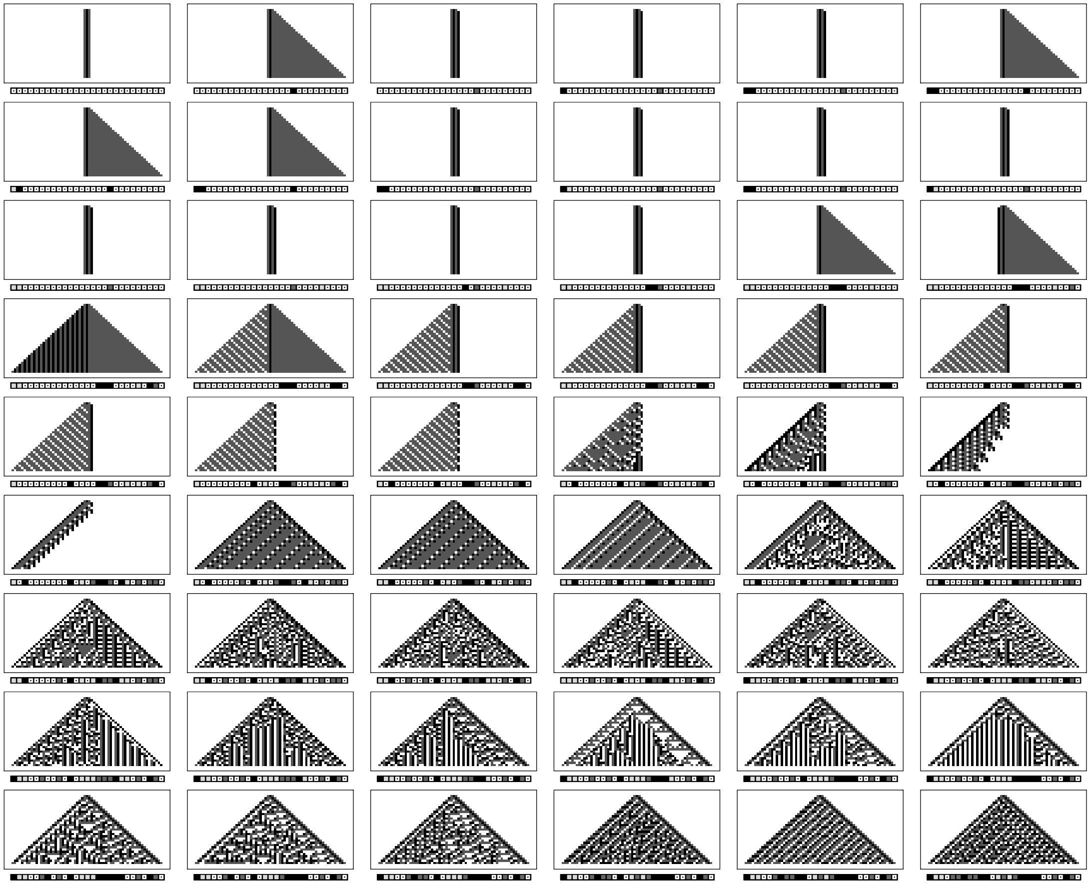

The behavior of a sequence of cellular automaton programs obtained by successive random mutations. The first program contains no rules for changing the color of a cell with any neighborhood. Mutations in successive programs add rules for changing the colors of cells with specific neighborhoods, or modify these rules. Each program in the sequence differs from the previous one by a single mutation, made completely at random. The sequence provides a very simple idealization of biological evolution without explicit natural selection. The cellular automata shown here all have 3 possible colors and nearest-neighbor rules. The label for each picture gives a representation of the rules for each of the 27 possible 3-cell neighborhoods. A dot signifies that the rule does not change the color of the center cell in the neighborhood.

But if complexity is this easy to get, why is it not even more widespread in biology? For while there are certainly many examples of elaborate forms and patterns in biological systems, the overall shapes and many of the most obvious features of typical organisms are usually quite simple.

So why should this be? My guess is that in essence it reflects limitations associated with the process of natural selection. For while natural selection is often touted as a force of almost arbitrary power, I have increasingly come to believe that in fact its power is remarkably limited. And indeed, what I suspect is that in the end natural selection can only operate in a meaningful way on systems or parts of systems whose behavior is in some sense quite simple.

If a particular part of an organism always grows, say, in a simple straight line, then it is fairly easy to imagine that natural selection could succeed in picking out the optimal length for any given environment. But what if an organism can grow in a more complex way, say like in the pictures on the previous page? My strong suspicion is that in such a case natural selection will normally be able to achieve very little.

There are several reasons for this, all somewhat related.

First, with more complex behavior, there are typically a huge number of possible variations, and in a realistic population of organisms it becomes infeasible for any significant fraction of these variations to be explored.

Second, complex behavior inevitably involves many elaborate details, and since different ones of these details may happen to be the deciding factors in the fates of individual organisms, it becomes very difficult for natural selection to act in a consistent and definitive way.

Third, whenever the overall behavior of a system is more complex than its underlying program, almost any mutation in the program will lead to a whole collection of detailed changes in the behavior, so that natural selection has no opportunity to pick out changes which are beneficial from those which are not.

Fourth, if random mutations can only, say, increase or decrease a length, then even if one mutation goes in the wrong direction, it is easy for another mutation to recover by going in the opposite direction. But if there are in effect many possible directions, it becomes much more difficult to recover from missteps, and to exhibit any form of systematic convergence.

And finally, as the results in Chapter 7 suggest, for anything beyond the very simplest forms of behavior, iterative random searches rapidly tend to get stuck, and make at best excruciatingly slow progress towards any kind of global optimum.

In a sense it is not surprising that natural selection can achieve little when confronted with complex behavior. For in effect it is being asked to predict what changes would need to be made in an underlying program in order to produce or enhance a certain form of overall behavior. Yet one of the main conclusions of this book is that even given a particular program, it can be very difficult to see what the behavior of the program will be. And to go backwards from behavior to programs is a still much more difficult task.

In writing this book it would certainly have been convenient to have had a systematic way to be able to find examples of programs that exhibit specified forms of complex behavior. And indeed I have tried hard to develop iterative search procedures that would do this. But even using a whole range of tricks suggested by biology—as well as quite a number that are not—I have never been successful. And in fact in every single case I have in the end reverted either to exhaustive or to purely random searches, with no attempt at iterative improvement.

So what does this mean for biological organisms? It suggests that if a particular feature of an organism is successfully going to be optimized for different environments by natural selection, then this feature must somehow be quite simple.

And no doubt that is a large part of the reason that biological organisms always tend to consist of separate organs or other parts, each of which has at least some attributes that are fairly simple. For in this way there end up being components that are simple enough to be adjusted in a meaningful fashion by natural selection.

It has often been claimed that natural selection is what makes systems in biology able to exhibit so much more complexity than systems that we explicitly construct in engineering. But my strong suspicion is that in fact the main effect of natural selection is almost exactly the opposite: it tends to make biological systems avoid complexity, and be more like systems in engineering.

When one does engineering, one normally operates under the constraint that the systems one builds must behave in a way that is readily predictable and understandable. And in order to achieve this one typically limits oneself to constructing systems out of fairly small numbers of components whose behavior and interactions are somehow simple.

But systems in nature need not in general operate under the constraint that their behavior should be predictable or understandable. And what this means is that in a sense they can use any number of components of any kind—with the result, as we have seen in this book, that the behavior they produce can often be highly complex.

However, if natural selection is to be successful at systematically molding the properties of a system then once again there are limitations on the kinds of components that the system can have. And indeed, it seems that what is needed are components that behave in simple and somewhat independent ways—much as in traditional engineering.

At some level it is not surprising that there should be an analogy between engineering and natural selection. For both cases can be viewed as trying to create systems that will achieve or optimize some goal.

Indeed, the main difference is just that in engineering explicit human effort is expended to find an appropriate form for the system, whereas in natural selection an iterative random search process is used instead. But the point is that the conditions under which these two approaches work turn out to be not so different.

In fact, there are even, I suspect, similarities in quite detailed issues such as the kinds of adjustments that can be made to individual components. In engineering it is common to work with components whose properties can somehow be varied smoothly, and which can therefore be analyzed using the methods of calculus and traditional continuous mathematics.

And as it turns out, much as we saw in Chapter 7, this same kind of smooth variation is also what tends to make iterative search methods such as natural selection be successful.

In biological systems based on discrete genetic programs, it is far from clear how smooth variation can emerge. Presumably in some cases it can be approximated by the presence of varying numbers of repeats in the underlying program. And more often it is probably the result of combinations of large numbers of elements that each produce fairly random behavior.

But the possibility of smooth variation seems to be important enough to the effectiveness of natural selection that it is extremely common in actual biological systems. And indeed, while there are some traits—such as eye color and blood type in humans—that are more or less discrete, the vast majority of traits seen, say, in the breeding of plants and animals, show quite smooth variation.

So to what extent does the actual history of biological evolution reflect the kinds of simple characteristics that I have argued one should expect from natural selection?

If one looks at species that exist today, and at the fossil record of past species, then one of the most striking features is just how much is in common across vast ranges of different organisms. The basic body plans for animals, for example, have been almost the same for hundreds of millions of years, and many organs and developmental pathways are probably even still older.

In fact, the vast majority of structurally important features seem to have changed only quite slowly and gradually in the course of evolution—just as one would expect from a process of natural selection that is based on smooth variations in fairly simple properties.

But despite this it is still clear that there is considerable diversity, at least at the level of visual appearance, in the actual forms of biological organisms that occur. So how then does such diversity arise?

One effect, to be discussed at greater length in the next section, is essentially just a matter of geometry. If the relative rates of growth of different parts of an organism change even slightly, then it turns out that this can sometimes have dramatic consequences for the overall shape of the organism, as well as for its mechanical operation.

And what this means is that just by making gradual changes in quantities such as relative rates of growth, natural selection can succeed in producing organisms that at least in some respects look very different.

But what about other differences between organisms? To what extent are all of them systematically determined by natural selection?

Following the discussion earlier in this section, it is my strong suspicion that at least many of the visually most striking differences—

associated for example with texture and pigmentation patterns—in the end have almost nothing to do with natural selection.

And instead what I believe is that such differences are in essence just reflections of completely random changes in underlying genetic programs, with no systematic effects from natural selection.

Particularly among closely related species of organisms there is certainly quite a contrast between the dramatic differences often seen in features such as pigmentation patterns and the amazing constancy of other features. And most likely those features in which a great degree of constancy is seen are precisely the ones that have successfully been molded by natural selection.

But as I mentioned earlier, it is almost always those features which change most rapidly between species that show the most obvious signs of complexity. And this observation fits precisely with the idea that complexity is easy to get by randomly sampling simple programs, but is hard for natural selection to handle in any kind of systematic way.

So in the end, therefore, what I conclude is that many of the most obvious features of complexity in biological organisms arise in a sense not because of natural selection, but rather in spite of it.

No doubt it will for many people be difficult to abandon the idea that natural selection is somehow crucial to the presence of complexity in biological organisms. For traditional intuition makes one think that to get the level of complexity that one sees in biological systems must require great effort—and the long and ponderous course of evolution revealed in the fossil record seems like just the kind of process that should be involved.

But the point is that what I have discovered in this book shows that in fact if one just chooses programs at random, then it is easy to get behavior of great complexity. And it is this that I believe lies at the heart of most of the complexity that we see in nature, both in biological and non-biological systems.

Whenever natural selection is an important determining factor, I suspect that one will inevitably see many of the same simplifying features as in systems created through engineering. And only when natural selection is not crucial, therefore, will biological systems be able to exhibit the same level of complexity that one observes for example in many systems in physics.

In biology the presence of long programs with many separate parts can lead to a certain rather straightforward complexity analogous to having many physical objects of different kinds collected together. But the most dramatic examples of complexity in biology tend to occur in individual parts of systems—and often involve patterns or structures that look remarkably like those in physics.

Yet if biology samples underlying genetic programs essentially at random, why should these programs behave anything like programs that are derived from specific laws of physics?

The answer, as we have seen many times in this book, is that across a very wide range of programs there is great universality in the behavior that occurs. The details depend on the exact rules for each program, but the overall characteristics remain very much the same.

And one of the important consequences of this is that it suggests that it might be possible to develop a rather general predictive theory of biology that would tell one, for example, what basic forms are and are not likely to occur in biological systems.

One might have thought that the traditional idea that organisms are selected to be optimal for their environment would already long ago have led to some kind of predictive theory. And indeed it has for example allowed some simple numerical ratios associated with populations of organisms to be successfully derived. But about a question such as what forms of organisms are likely to occur it has much less to say.

There are a number of situations where fairly complicated structures appear to have arisen independently in several very different types of organisms. And it is sometimes claimed that this kind of convergent evolution occurs because these structures are in some ultimate sense optimal, making it inevitable that they will eventually be produced.

But I would be very surprised if this explanation were correct. And instead what I strongly suspect is that the reason certain structures appear repeatedly is just that they are somehow common among programs of certain kinds—just as, for example, we have seen that the

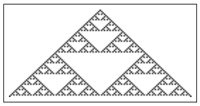

An example of a basic pattern that is produced in several variants by a wide range of simple programs.

intricate nested pattern shown on the left arises from many different simple programs.

Ever since the original development of the theory of evolution, there has been a widespread belief that the general trend seen in the fossil record towards the formation of progressively more complicated types of organisms must somehow be related to an overall increase in optimality.

Needless to say, we do not know what a truly optimal organism would be like. But if optimality is associated with having as many offspring as possible, then very simple organisms such as viruses and protozoa already seem to do very well.

So why then do higher organisms exist at all? My guess is that it has almost nothing to do with optimality, and that instead it is essentially just a consequence of strings of random mutations that happened to add more and more features without introducing fatal flaws.

It is certainly not the case—as is often assumed—that natural selection somehow inevitably leads to organisms with progressively more elaborate structures and progressively larger numbers of parts.

For a start, some kinds of organisms have been subject to natural selection for more than a billion years, but have never ended up becoming much more complicated. And although there are situations where organisms do end up becoming more complicated, they also often become simpler.

A typical pattern—remarkably similar, as it happens, to what occurs in the history of technology—is that at some point in the fossil record some major new capability or feature is suddenly seen. At first there is then rapid expansion, with many new species trying out all sorts of possibilities that have been opened up. And usually some of these possibilities get quite ornate and elaborate. But after a while it becomes clear what makes sense and what does not. And typically things then get simpler again.

So what is the role of natural selection in all of this? My guess is that as in other situations, its main systematic contribution is to make things simpler, and that insofar as things do end up getting more complicated, this is almost always the result of essentially random sampling of underlying programs—without any systematic effect of natural selection.

For the more superficial aspects of organisms—such as pigmentation patterns—it seems likely that among programs sampled at random a fair fraction will produce results that are not disastrous for the organism. But when one is dealing with the basic structure of organisms, the vast majority of programs sampled at random will no doubt have immediate disastrous consequences. And in a sense it is natural selection that is responsible for the fact that such programs do not survive.

But the point is that in such a case its effect is not systematic or cumulative. And indeed it is my strong suspicion that for essentially all purposes the only reasonable model for important new features of organisms is that they come from programs selected purely at random.

So does this then mean that there can never be any kind of general theory for all the features of higher organisms? Presumably the pattern of exactly which new features were added when in the history of biological evolution is no more amenable to general theory than the specific course of events in human history. But I strongly suspect that the vast majority of significant new features that appear in organisms are at least at first associated with fairly short underlying programs. And insofar as this is the case the results of this book should allow one to develop some fairly general characterizations of what can happen.

So what all this means is that much of what we see in biology should correspond quite closely to the typical behavior of simple programs as we have studied them in this book—with the main caveat being just that certain aspects will be smoothed and simplified by the effects of natural selection. Seeing in earlier chapters of this book all the diverse things that simple programs can do, it is easy to be struck by analogies to books of biological flora and fauna. Yet what we now see is that in fact such analogies may be quite direct—and that many of the most obvious features of actual biological organisms may in effect be direct reflections of typical behavior that one sees in simple programs.