STEPHEN WOLFRAM A NEW KIND OF SCIENCE

EXCERPTED FROM

SECTION 8.6

Growth of Plants and Animals

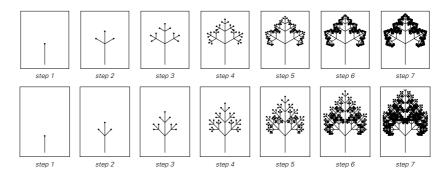
Growth of Plants and Animals

Looking at all the elaborate forms of plants and animals one might at first assume that the underlying rules for their growth must be highly complex. But in this book we have discovered that even by following very simple rules it is possible to obtain forms of great complexity. And what I have come to believe is that in fact most aspects of the growth of plants and animals are in the end governed by remarkably simple rules.

As a first example of biological growth, consider the stem of a plant. It is usually only at the tip of a stem that growth can occur, and much of the time all that ever happens is that the stem just gets progressively longer. But the crucial phenomenon that ultimately leads to much of the structure we see in many kinds of plants is that at the tip of a stem it is possible for new stems to form and branch off. And in the simplest cases these new stems are in essence just smaller copies of the original stem, with the same basic rules for growth and branching.

With this setup the succession of branchings can then be represented by steps in the evolution of a neighbor-independent substitution system in which the tip of each stem is at each step replaced by a collection of smaller stems in some fixed configuration.

Two examples of such substitution systems are shown in the pictures below. In both cases the rules are set up so that every stem in effect just branches into exactly three new stems at each step. And this



Steps in the evolution of substitution systems that provide simple models for the growth of plants. At each step every growing stem is replaced by a collection of three new stems according to the rules shown. For individual stems this type of branching is known in botany as monopodial.

means that the network of connections between stems necessarily has a very simple nested form. But if one looks at the actual geometrical arrangement of stems there is no longer such simplicity; indeed, despite the great simplicity of the underlying rules, considerable complexity is immediately evident even in the pictures at the bottom of the facing page.

The pictures on the next page show patterns obtained with various sequences of choices for the lengths and angles of new stems. In a few cases the patterns are quite simple; but in most cases they turn out to be highly complex—and remarkably diverse.

The pictures immediately remind one of the overall branching patterns of all sorts of plants—from algae to ferns to trees to many kinds of flowering plants. And no doubt it is from such simple rules of growth that most such overall branching patterns come.

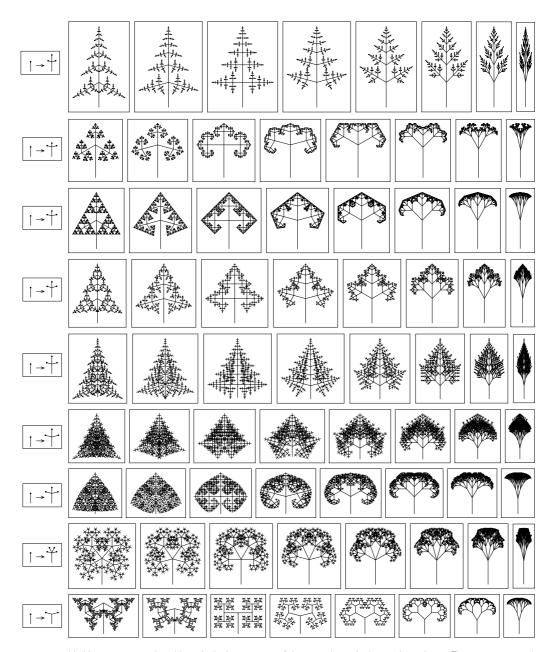
But what about more detailed features of plants? Can they also be thought of as consequences of simple underlying rules of growth?

For many years I wondered in particular about the shapes of leaves. For among different plants there is tremendous diversity in such shapes—as illustrated in the pictures on page 403. Some plants have leaves with simple smooth boundaries that one might imagine could be described by traditional mathematical functions. Others have leaves with various configurations of sharp points. And still others have leaves with complex and seemingly somewhat random boundaries.

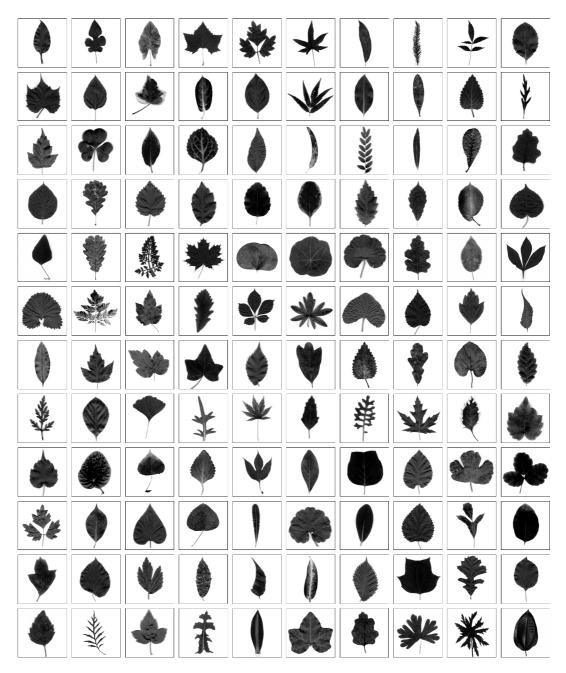
So given this diversity one might at first suppose that no single kind of underlying rule could be responsible for what is seen. But looking at arrays of pictures like the ones on the next page one makes a remarkable discovery: among the patterns that can be generated by simple substitution systems are ones whose outlines look extremely similar to those of a wide variety of types of leaves.

There are patterns with smooth edges that look like lily pads. There are patterns with sharp points that look like prickly leaves of various kinds. And there are patterns with intricate and seemingly somewhat random shapes that look like sycamore or grape leaves.

It has never in the past been at all clear how leaves get the shapes they do. Presumably most of the processes that are important take place while leaves are still folded up inside buds, and are not yet very solid.



Limiting patterns produced by substitution systems of the type shown in the previous picture. The patterns on each row are obtained from rules that are set up to give branches with particular relative lengths. The angles between the branches are taken to increase by 15° in successive pictures across the row. Note that pictures shown on different rows are scaled differently—so that the initial vertical stem does not always appear with the same height. The similarity between pictures on this page and overall branching patterns and shapes of leaves in many kinds of plants is striking.



Examples of different kinds of leaves, mostly from common flowering plants. The diversity of shapes is remarkable, as is the similarity to the forms shown on the facing page. The leaves range in size from under an inch to many feet.

For although leaves typically expand significantly after they come out, the basic features of their shapes almost never seem to change.

There is some evidence that at least some aspects of the pattern of veins in a leaf are laid down before the main surface of the leaf is filled in, and perhaps the stems in the branching process I describe here correspond to precursors of structures related to veins. Indeed, the criss-crossing of veins in the leaves of higher plants may be not unrelated to the fact that stems in the pictures two pages ago often cross over—although certainly many of the veins in actual full-grown leaves are probably added long after the shapes of the leaves are determined.

One might at the outset have thought that leaves would get their shapes through some mechanism quite unrelated to other aspects of plant growth. But I strongly suspect that in fact the very same simple process of branching is ultimately responsible both for the overall forms of plants, and for the shapes of their leaves.

Quite possibly there will sometimes be at least some correspondence between the lengths and angles that appear in the rules for overall growth and for the growth of leaves. But in general the details of all these rules will no doubt depend on very specific characteristics of individual plants.

The distance before a new stem appears is, for example, probably determined by the rates of production and diffusion of plant hormones and related substances, and these rates will inevitably depend both on the thickness and mechanical structure of the stem, as well as on all kinds of biochemical properties of the plant. And when it comes to the angles between old and new stems I would not be surprised if these were governed by such microscopic details as individual shapes of cells and individual sequences of cell divisions.

The traditional intuition of biology would suggest that whenever one sees complexity—say in the shape of a leaf—it must have been generated for some particular purpose by some sophisticated process of natural selection. But what the pictures on the previous pages demonstrate is that in fact a high degree of complexity can arise in a sense quite effortlessly just as a consequence of following certain simple rules of growth.

No doubt some of the underlying properties of plants are indeed guided by natural selection. But what I strongly suspect is that in the vast majority of cases the occurrence of complexity—say in the shapes of leaves—is in essence just a side effect of the particular rules of growth that happen to result from the underlying properties of the plant.

The pictures on the next page show the array of possible forms that can be produced by rules in which each stem splits into exactly two new stems at each step. The vertical black line on the left-hand side of the page represents in effect the original stem at each step, and the pictures are arranged so that the one which appears at a given position on the page shows the pattern that is generated when the tip of the right-hand new stem goes to that position relative to the original stem shown on the left.

In some cases the patterns obtained are fairly simple. But even in these cases the pictures show that comparatively small changes in underlying rules can lead to much more complex patterns. And so if in the course of biological evolution gradual changes occur in the rules, it is almost inevitable that complex patterns will sometimes be seen.

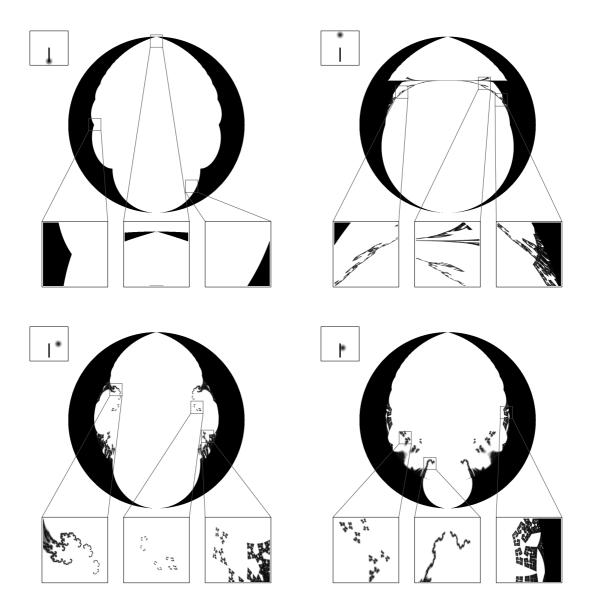
But just how suddenly can the patterns change? To get some idea of this one can construct a kind of limit of the array on the next page in which the total number of pictures is in effect infinite, but only a specific infinitesimal region of each picture is shown. Page 407 gives results for four choices of the position of this region relative to the original stem. And instead of just displaying black or white depending on whether any part of the pattern lies in the region, the picture uses gray levels to indicate how close it comes.

The areas of solid black thus correspond to ranges of parameters in the underlying rule for which the patterns obtained always reach a particular position. But what we see is that at the edges of these areas there are often intricate structures with an essentially nested form. And the presence of such structures implies that at least with some ranges of parameters, even very small changes in underlying rules can lead to large changes in certain aspects of the patterns that are produced.

So what this suggests is that it is almost inevitable that features such as the shapes of leaves can sometimes change greatly even when the underlying properties of plants change only slightly. And I suspect

-	,	\bigcirc					fr. A	A BA
-	Ţ	\bigcirc	Y Y	Y				
Ţ,	Ŷ	(Cr)					
Ĩ	Ŷ	۲ ۲	C					
γ	Ŷ	Ŷ	ϵ_{γ}	C ^r 3				
Ŷ	Ŷ	Ŷ	$\mathcal{C}_{igert}^{\infty}$	etta K	Cores			
Ŷ	Ŷ	Ŷ	ကိုာ	\mathfrak{C}				/ #*************** / #*#################
Ŷ	Ŷ	נ ר ייזי	ເຕັ	ଟ୍ୟୁତ				
Ŷ	Ŷ		; ; ;	€3 €3 €3¥€3	63 63 63 63		/ 274000000000000000000000000000000000000	
Ť	Ŷ	*	€/¥ ₽/₩	**	00 00 02 60			900×000 80 8010
† T	Ţ	:- <u>:</u>	**					· · · · · · · · · · · · · · · · · · ·
Î	Tr.	~	2	36 36				a #*# # ∉⊺ ∰ #
^	*	*	*** ***	*** **∕** **				
Å	**	**						
↑	*	78.55	N N					
♠	*	~**	we see			· · · · · · · · · · · · · · · · · · ·	· · · · · · · · · · · · · · · · · · ·	
♠						· / 🍂 🕅	. "Agaan vaa kastiin Nietoorikaan kastiin ka	
₼	+							
+	-			- 96				NIN . MARAN
+								
	/M	<u>s </u>	A.S.					

The full array of patterns that can be produced by simple substitution systems in which each stem branches into exactly two symmetrical stems at each step. The patterns are arranged on the page so that the pattern shown at a particular position corresponds to what is obtained with a rule in which the tip of the right-hand stem goes to that position (corrected for the aspect ratio of the array) relative to the original stem shown as a vertical line on the left-hand side of the page. In each case the result of 10 steps of evolution is shown, and the pictures are scaled so that all points above the bottom of the original stem can be included. Note that for rules outside of a distorted semicircle centered on the dot at the left-hand side of the page, and touching the three other sides of the page, the patterns generated grow at each step, rather than tending to a limit of fixed size.



Maps of where in the space of parameters for the substitution systems on the facing page the patterns obtained overlap the region indicated in the icon at the top left of each picture. Black corresponds to complete overlap, while white corresponds to no overlap. The maps shown can be thought of as being made by taking an infinitely dense limit of the array of pictures on the facing page, but keeping only what one sees in each picture by looking through a peephole at a particular position relative to the original stem.

that this is precisely why such diverse shapes of leaves are occasionally seen even in plants that otherwise appear very similar.

But while features such as the shapes of leaves typically differ greatly between different plants, there are also some seemingly quite sophisticated aspects of plants that typically remain almost exactly the same across a huge range of species.

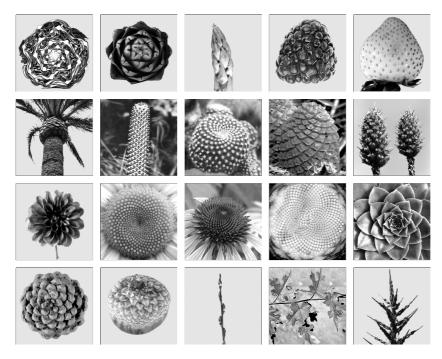
One example is the arrangement of sequences of plant organs or other elements around a stem. In some cases successive leaves, say, will always come out on opposite sides of a stem—180° apart. But considerably more common is for leaves to come out less than 180° apart, and in most plants the angle turns out to be essentially the same, and equal to almost exactly 137.5°.

It is already remarkable that such a definite angle arises in the arrangement of leaves—or so-called phyllotaxis—of so many plants. But it turns out that this very same angle also shows up in all sorts of other features of plants, as shown in the pictures at the top of the facing page. And although the geometry is different in different cases, the presence of a fixed angle close to 137.5° always leads to remarkably regular spiral patterns.

Over the years, much has been written about such patterns, and about their mathematical properties. For it turns out that an angle between successive elements of about 137.5° is equivalent to a rotation by a number of turns equal to the so-called golden ratio $(1 + \sqrt{5})/2 \approx 1.618$ which arises in a wide variety of mathematical contexts—notably as the limiting ratio of Fibonacci numbers.

And no doubt in large part because of this elegant mathematical connection, it has usually come to be assumed that the 137.5° angle and the spiral patterns to which it leads must correspond to some kind of sophisticated optimization found by an elaborate process of natural selection.

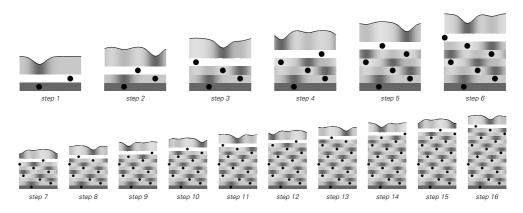
But I do not believe that this is in fact the case. And instead what I strongly suspect is that the patterns are just inevitable consequences of a rather simple process of growth not unlike one that was already discussed, at least in general terms, nearly a century ago.



Examples of spiral arrangements of elements in various plant systems. The details of the final geometry are different in different cases. But in all cases it turns out that the original angle between successive elements is almost exactly 137.5°. The first row shows red cabbage (cut open), artichoke, asparagus, raspberry and strawberry. The first two objects on the last row are a pinecone and an acorn.

The positions of new plant organs or other elements around a stem are presumably determined by what happens in a small ring of material near the tip of the growing stem. And what I suspect is that a new element will typically form at a particular position around the ring if at that position the concentration of some chemical has reached a certain critical level.

But as soon as an element is formed, one can expect that it will deplete the concentration of the chemical in its local neighborhood, and thus inhibit further elements from forming nearby. Nevertheless, general processes in the growing stem will presumably make the concentration steadily rise throughout the ring of active material, and eventually this concentration will again get high enough at some position that it will cause another element to be formed.

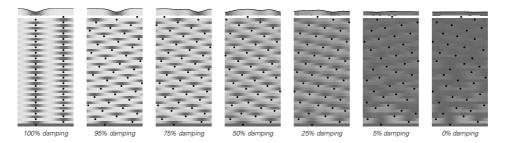


A simple model for the arrangement of leaves or other elements produced at the growing tip of a plant stem. The stem is taken to grow up the page, and for purposes of display it is unrolled into a line. The positions of leaves or other elements are indicated by black dots. The concentration of a chemical is indicated by gray level, and for the top line at each step, it is also plotted. The rule for the system places a new black dot at whatever position this concentration is largest. The black dot is then assumed to deplete the concentration around it, but the overall concentration is uniformly increased before the next step. It turns out that successive black dots rapidly become spaced at almost exactly 137.5°.

The pictures above show an example of this type of process. For purposes of display the ring of active material is unrolled into a line, and successive states of this line are shown one on top of each other going up the page. At each step a new element, indicated by a black dot, is taken to be generated at whatever position the concentration is maximal. And around this position the new element is then taken to produce a dip in concentration that is gradually washed out over the course of several steps.

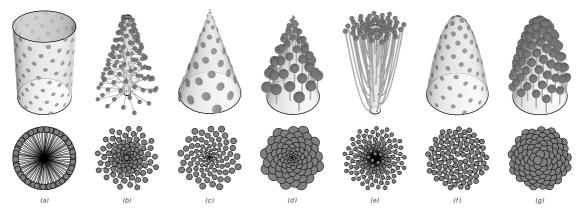
The way the pictures are drawn, the angles between successive elements correspond to the horizontal distances between them. And although these distances vary somewhat for the first few steps, what we see in general is remarkably rapid convergence to a fixed distance which turns out to correspond to an angle of almost exactly 137.5°.

So what happens if one changes the details of the model? In the extreme case where all memory of previous behavior is immediately damped out the first picture at the top of the facing page shows that successive elements form at 180° angles. And in the case where there is very little damping the last two pictures show that at least for a while elements can form at fairly random angles. But in the majority of cases one sees rather rapid convergence to almost precisely 137.5°.



Examples of changing the amount of damping used in the model on the facing page. 100% damping corresponds to increasing the overall concentration at each step so much that no memory of previous steps remains. 0% corresponds to no increase in overall concentration at each step. Away from these extreme cases, rapid convergence is seen to a spacing between black dots of almost exactly 137.5°.

So just how does this angle show up in actual plant systems? As the top pictures below demonstrate, the details depend on the geometry and relative growth rates of new elements and of the original stem. But in all cases very characteristic patterns are produced.



Examples of structures formed in various geometries by successively adding elements at a golden ratio angle 137.5°. Each of these structures is seen in one type of plant growth or another, as illustrated on page 409.



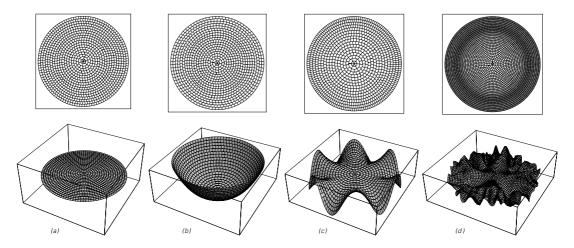
Overall patterns formed by successively adding elements at a variety of different angles. In each case the n^{th} element appears at coordinates $\sqrt{n} \{Cos[n \theta], Sin[n \theta]\}$. Stripes are seen if θ/π (with θ in radians) is easy to approximate by a rational number. (The size of the region before stripes appear depends on Length[ContinuedFraction[θ/π]].)

And as the bottom pictures on the previous page demonstrate, the forms of these patterns are very sensitive to the precise angle of successive elements: indeed, even a small deviation leads to patterns that are visually quite different. At first one might have assumed that to get a precise angle like 137.5° would require some kind of elaborate and highly detailed process. But just as in so many other situations that we have seen in this book, what we have seen is that in fact a very simple rule is all that is in the end needed.

One of the general features of plants is that most of their cells tend to develop fairly rigid cellulose walls which make it essentially impossible for new material to be added inside the volume of the plant, and so typically force new growth to occur only on the outside of the plant—most importantly at the tips of stems.

But when plants form sheets of material as in leaves or petals there is usually some flexibility for growth to occur within the sheet. And the pictures below show examples of what can happen if one starts with a flat disk and then adds different amounts of material in different places.

If more material is added near the center than near the edge, as in case (b), then the disk is forced to take on a cup shape similar to many



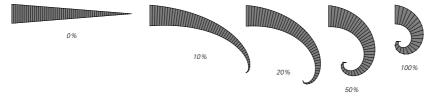
Disks with varying amounts of material at different distances from their centers. In the top row the disks are always flat, forcing the cells of material to vary in size and shape. In the bottom row, the disks form shapes in three dimensions in which all cells are the same size and shape. Relative to case (a), the amount of material going out from the center decreases linearly in case (b), increases linearly in case (c), and increases exponentially in case (d).

flowers. But if more material is added near the edge than near the center, as in case (c), then the sheet will become wavy at the edge, much like some leaves. And if the amount of material increases sufficiently rapidly from the center to the edge, as in case (d), then the disk will be forced to become highly corrugated, somewhat like a lettuce leaf.

So what about animals? To what extent are their mechanisms of growth the same as plants? If one looks at air passages or small blood vessels in higher animals then the patterns of branching one sees look similar to those in plants. But in most of their obvious structural features animals do not typically look much like plants at all. And in fact their mechanisms of growth mostly turn out to be rather different.

As a first example, consider a horn. One might have thought that, like a stem in a plant, a horn would grow by adding material at its tip. But in fact, like nails and hair, a horn instead grows by adding material at its base. And an immediate consequence of this is that the kind of branching that one sees in plants does not normally occur in horns.

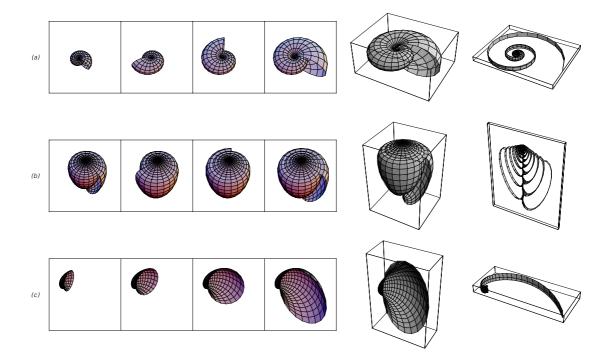
But on the other hand coiling is common. For in order to get a structure that is perfectly straight, the rate at which material is added must be exactly the same on each side of the base. And if there is any difference, one edge of the structure that is produced will always end up being longer than the other, so that coiling will inevitably result, as in the pictures below.



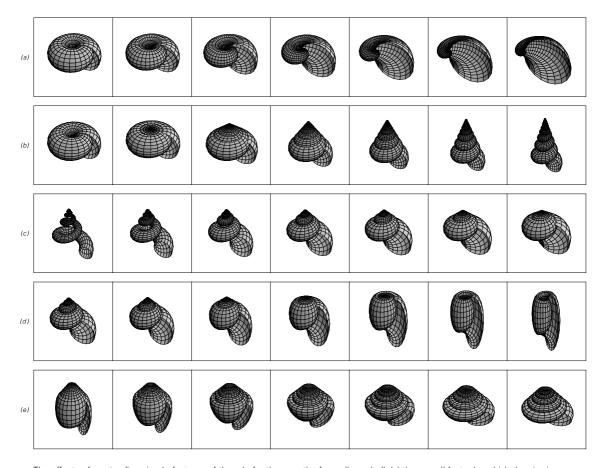
Idealized horns generated by progressively adding new material, with the amount of material on the upper edge of the base always being the specified percentage larger than the amount on the lower edge. These pictures can be viewed as one-dimensional analogs of those on the facing page.

And as has been thought for several centuries, it turns out that a three-dimensional version of this phenomenon is essentially what leads to the elaborate coiled structures that one sees in mollusc shells. For in a typical case, the animal which lives at the open end of the shell secretes new shell material faster on one side than the other, causing the shell to grow in a spiral. The rates at which shell material is secreted at different points around the opening are presumably determined by details of the anatomy of the animal. And it turns out that—much as we saw in the case of branching structures earlier in this section—even fairly small changes in such rates can have quite dramatic effects on the overall shape of the shell.

The pictures below show three examples of what can happen, while the facing page shows the effects of systematically varying certain growth rates. And what one sees is that even though the same very simple underlying model is used, there are all sorts of visually very different geometrical forms that can nevertheless be produced.



A simple model for the growth of mollusc shells. In each case new shell material is progressively added at the open end of the shell. The rule on the left shows the amount of material added at each stage at different points around the opening; the line from the center indicates the progressive lateral displacement of the opening. Case (a) is typical of a nautilus shell, (b) of a cone shell and (c) of one-half of a clam shell. All shells produced by adding material according to fixed rules of the kind shown here have the property that throughout their growth they maintain the same overall shape.



The effects of varying five simple features of the rule for the growth of a mollusc shell: (a) the overall factor by which the size increases in the course of each revolution; (b) the relative amount by which the opening is displaced downward at each revolution; (c) the size of the opening relative to the overall size of the shell; (d) the elongation of the opening; (e) the orientation of elongation in the opening. The pictures at the beginning and end of each row correspond roughly to the following: (a) pond snail shell, cockle shell; (b) pond snail shell, horn shell; (c) worm shell, bonnet shell; (d) periwinkle shell, cowrie shell; (e) olive shell, sundial shell.

So out of all the possible forms, which ones actually occur in real molluscs? The remarkable fact illustrated on the next page is that essentially all of them are found in some kind of mollusc or another.

If one just saw a single mollusc shell, one might well think that its elaborate form must have been carefully crafted by some long process of natural selection. But what we now see is that in fact all the different forms that are observed are in effect just consequences of the Shell shapes generated by the simple model and found in nature. The array shows systematic variation of the first two parameters from the previous page. Similar arrays could be made for the other parameters.

0		
		E.

application of three-dimensional geometry to very simple underlying rules of growth. And so once again therefore natural selection cannot reasonably be considered the source of the elaborate forms we see.

Away from mollusc shells, coiled structures—like branched ones—are not especially common in animals. Indeed, the vast majority of animals do not tend to have overall forms that are dominated by any single kind of structure. Rather, they are usually made up of a collection of separate identifiable parts, like heads, tails, legs, eyes and so on, all with their own specific structure.

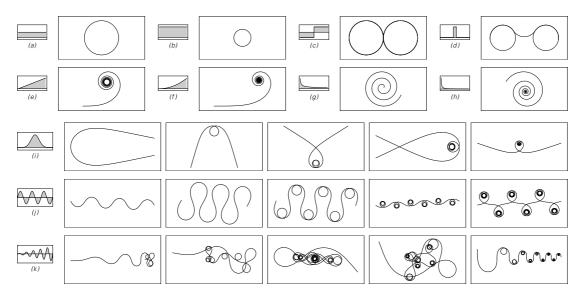
Sometimes some of these parts are repeated, perhaps in a sequence of segments, or perhaps in some kind of two-dimensional array. And very often the whole animal is covered by a fairly uniform outer skin. But the presence of many different kinds of parts is in the end one of the most obvious features of many animals.

So how do all these parts get produced? The basic mechanism seems to be that at different places and different times inside a developing animal different sections of its genetic program end up getting used—causing different kinds of growth to occur, and different structures to be produced. And part of what makes this possible is that particularly at the stage of the embryo most cells in an animal are not extremely rigid—so that even when different pieces of the animal grow quite differently they can still deform so as to fit together.

Usually there are some elements—such as bones—that eventually do become rigid. But the crucial point is that at the stage when the basic form of an animal is determined most of these elements are not yet rigid. And this allows various processes to occur that would otherwise be impossible.

Probably the most important of these is folding. For folding is not only involved in producing shapes such as teeth surfaces and human ear lobes, but is also critical in allowing flat sheets of tissue to form the kinds of pockets and tubes that are so common inside animals.

Folding seems to occur for a variety of reasons. Sometimes it is most likely the direct result of tugging by microscopic fibers. And in other cases it is probably a consequence of growth occurring at different rates in different places, as in the pictures on page 412.



Curves obtained by varying the local curvature according to definite rules as one goes from one end to the other. Each sequence of curves shows what happens when the local curvature is multiplied by a progressively larger factor. The local curvature at any particular point is defined to be the reciprocal of the radius of a circle that approximates the curve at that point. The formulas for local curvature as a function of arc length for each set of pictures are as follows: 1 (circle); *s* (Cornu spiral or clothoid); s^2 ; $1/\sqrt{s}$ (involute of circle); 1/s (logarithmic or equiangular spiral); $1/s^2$; e^{-s^2} ; Sin[s]; sSin[s]. The curvature functions f[s] can be thought of as specifying how much to turn a vehicle at every moment in order to keep it driving along the curve. The curves have been rotated so as to fit into the frames provided.

But what kinds of shapes can folding produce? The pictures above show what happens when the local curvature—which is essentially the local rate of folding—is taken to vary according to several simple rules as one goes along a curve. In a few cases the shapes produced are rather simple. But in most cases they are fairly complicated. And it takes only very simple rules to generate shapes that look like the villi and other corrugated structures one often sees in animals.

In addition to folding, there are other kinds of processes that are made possible by the lack of rigidity in a developing animal. One is furrowing or tearing of tissue through a loss of adhesion between cells. And another is explicit migration of individual cells based on chemical or immunological affinities.

But how do all these various processes get organized to produce an actual animal? If one looks at the sequence of events that take place in a

typical animal embryo they at first seem remarkably haphazard. But presumably the main thing that is going on—as mentioned above—is that at different places and different times different sections of the underlying genetic program are being used, and these different sections can lead to very different kinds of behavior. Some may produce just uniform growth. Others may lead to various kinds of local folding. And still others may cause regions of tissue to die—thereby for example allowing separate fingers and toes to emerge from a single sheet of tissue.

But just how is it determined what section of the underlying genetic program should be used at what point in the development of the animal? At first, one might think that each individual cell that comes into existence might use a different section of the underlying genetic program. And in very simple animals with just a few hundred cells this is most likely what in effect happens.

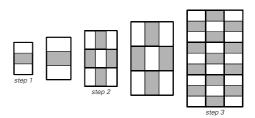
But in general it seems to be not so much individual cells as regions of the developing animal that end up using different sections of the underlying program. Indeed, the typical pattern seems to be that whenever a part of an animal has grown to be a few tenths of a millimeter across, that part can break up into a handful of smaller regions which each use a different section of the underlying genetic program.

So how does this work? What appears to be the case is that there are cells which produce chemicals whose concentrations decrease over distances of a few tenths of a millimeter. And what has been discovered in the past decade or so is that in all animals—as well as plants—there are a handful of so-called homeobox genes which seem to become active or inactive at particular concentration levels and which control what section of the underlying genetic program will be used.

The existence of a fixed length scale at which such processes occur then almost inevitably implies that an embryo must develop in a somewhat hierarchical fashion. For at a sufficiently early stage, the whole embryo will be so small that it can contain only a handful of regions that use different sections of the genetic program. And at this stage there may, for example, be a leg region, but there will not yet be a distinct foot region. As the embryo grows, however, the leg region will eventually become large enough that it can differentiate into several separate regions. And at this point, a distinct foot region can appear. Then, when the foot region becomes large enough, it too can break into separate regions that will, say, turn into bone or soft tissue. And when a region that will turn into bone becomes large enough, it can break into further regions that will, say, yield separate individual bones.

If at every stage the tissue in each region produced grows at the same rate, and all that differs is what final type of cells will exist in each region, then inevitably a simple and highly regular overall structure will emerge, as in the idealized picture below. With different substitution rules for each type of cell, the structure will in general be nested. And in fact there are, for example, some parts of the skeletons of animals that do seem to exhibit, at least roughly, a few levels of nesting of this kind.

A schematic illustration of the successive subdivisions which presumably occur in the growth of animals. Here the subdivisions are taken to occur in two directions, always giving three simple rectangles which all grow at the same rate. In practice, the geometry will usually be much more complex.



But in most cases there is no such obvious nesting of this kind. One reason for this is that a region may break not into a simple line of smaller regions, but into concentric circles or into some collection of regions in a much more complicated arrangement—say of the kind that I discuss in the next section. And perhaps even more important, a region may break into smaller regions that grow at different rates, and that potentially fold over or deform in other ways. And when this happens, the geometry that develops will in turn affect the way that subsequent regions break up.

The idea that the basic mechanism for producing different parts of animals is that regions a few tenths of a millimeter across break into separate smaller regions turns out in the end to be strangely similar to the idea that stems of plants whose tips are perhaps a millimeter across grow by splitting off smaller stems. And indeed it is even known that some of the genetic phenomena involved are extremely similar.

But the point is that because of the comparative rigidity of plants during their most important period of growth, only structures that involve fairly explicit branching can be produced. In animals, however, the lack of rigidity allows a vastly wider range of structures to appear, since now tissue in different regions need not just grow uniformly, but can change shape in a whole variety of ways.

By the time an animal hatches or is born, its basic form is usually determined, and there are bones or other rigid elements in place to maintain this form. But in most animals there is still a significant further increase in size. So how does this work?

Some bones in effect just expand by adding material to their outer surface. But in many cases, bones are in effect divided into sections, and growth occurs between these sections. Thus, for example, the long bones in the arms and legs have regions of growth at each end of their main shafts. And the skull is divided into a collection of pieces that each grow around their edges.

Typically there are somewhat different rates of growth for different parts of an animal—leading, for example, to the decrease in relative head size usually seen from birth to adulthood. And this inevitably means that there will be at least some changes in the shapes of animals as they mature.

But what if one compares different breeds or species of animals? At first, their shapes may seem quite different. But it turns out that among animals of a particular family or even order, it is very common to find that their overall shapes are in fact related by fairly simple and smooth geometrical transformations.

And indeed it seems likely that—much like the leaves and shells that we discussed earlier in this section—differences between the shapes and forms of animals may often be due in large part merely to different patterns in the rates of growth for their different parts.

Needless to say, just like with leaves and shells, such differences can have effects that are quite dramatic both visually and mechanically turning, say, an animal that walks on four legs into one that walks on two. And, again just like with leaves and shells, it seems likely that among the animals we see are ones that correspond to a fair fraction of the possible choices for relative rates of growth.

We began this section by asking what underlying rules of growth would be needed to produce the kind of diversity and complexity that we see in the forms of plants and animals. And in each case that we have examined what we have found is that remarkably simple rules seem to suffice. Indeed, in most cases the basic rules actually seem to be somewhat simpler than those that operate in many non-biological systems. But what allows the striking diversity that we see in biological systems is that different organisms and different species of organisms are always based on at least slightly different rules.

In the previous section I argued that for the most part such rules will not be carefully chosen by natural selection, but instead will just be picked almost at random from among the possibilities. From experience with traditional mathematical models, however, one might then assume that this would inevitably imply that all plants and animals would have forms that look quite similar.

But what we have discovered in this book is that when one uses rules that correspond to simple programs, rather than, say, traditional mathematical equations, it is very common to find that different rules lead to quite different—and often highly complex—patterns of behavior. And it is this basic phenomenon that I suspect is responsible for most of the diversity and complexity that we see in the forms of plants and animals.

Biological Pigmentation Patterns

At a visual level, pigmentation patterns represent some of the most obvious examples of complexity in biological organisms. And in the past it has usually been assumed that to get the kind of complexity that one sees in such patterns there must be some highly complex underlying mechanism, presumably related to optimization through natural selection.

Following the discoveries in this book, however, what I strongly suspect is that in fact the vast majority of pigmentation patterns in