EXCERPTED FROM

STEPHEN WOLFRAM
A NEW KIND OF SCIENCE

SECTION 8.7

Biological Pigmentation Patterns

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 biological organisms are instead generated by processes whose basic rules are extremely simple—and are often chosen essentially at random.

The pictures below shows some typical examples of patterns found on mollusc shells. Many of these patterns are quite simple. But some are highly complex. Yet looking at these patterns one notices a remarkable similarity to patterns that we have seen many times before in this book—generated by simple one-dimensional cellular automata.

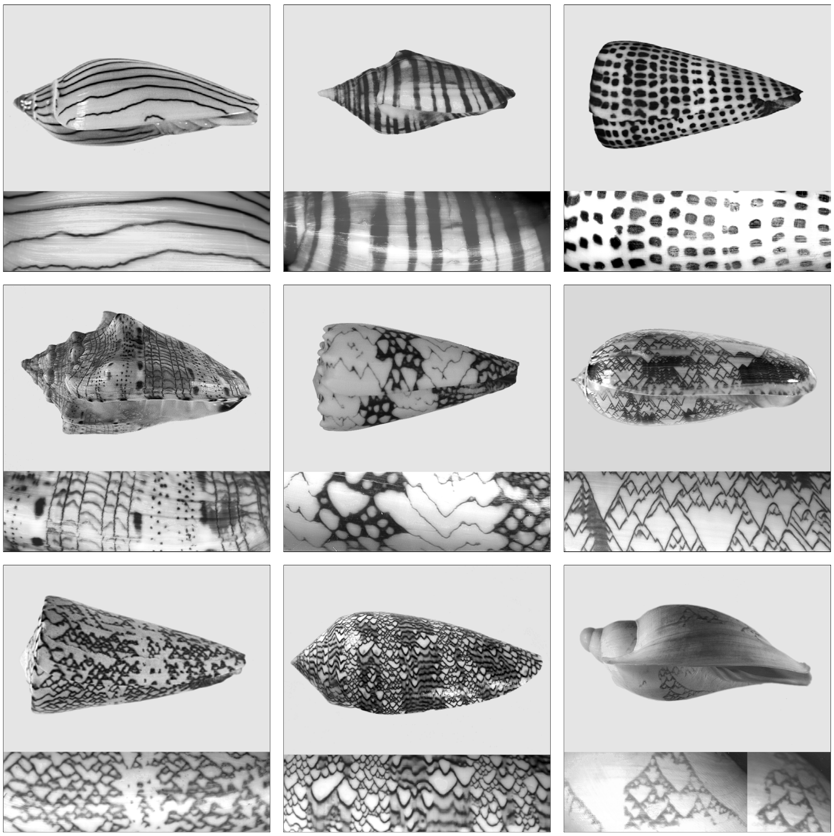

Typical examples of pigmentation patterns on mollusc shells. In each close-up the pattern grows from top to bottom, just like in a one-dimensional cellular automaton. Patterns with triangles are often said to have a "tent" or "divaricate" form. The shell on the bottom right is a slightly rare specimen where something close to an explicit nested pattern can be seen. Most of the shells are between one and four inches long; the one on the bottom right is nine inches long. The patterns are all various shades of brown on roughly white backgrounds. The shells are the following types: first row: Elliot's volute, vexillate volute, lettered cone; second row: music volute, banded marble cone, tent olive; third row: bough cone, textile cone, false melon volute (*Livonia mammilla*).

This similarity is, I believe, no coincidence. A mollusc shell, like a one-dimensional cellular automaton, in effect grows one line at a time, with new shell material being produced by a lip of soft tissue at the edge of the animal inside the shell. Quite how the pigment on the shell is laid down is not completely clear. There are undoubtedly elements in the soft tissue that at any point either will or will not secrete pigment. And presumably these elements have certain interactions with each other. And given this, the simplest hypothesis in a sense is that the new state of the element is determined from the previous state of its neighbors—just as in a one-dimensional cellular automaton.

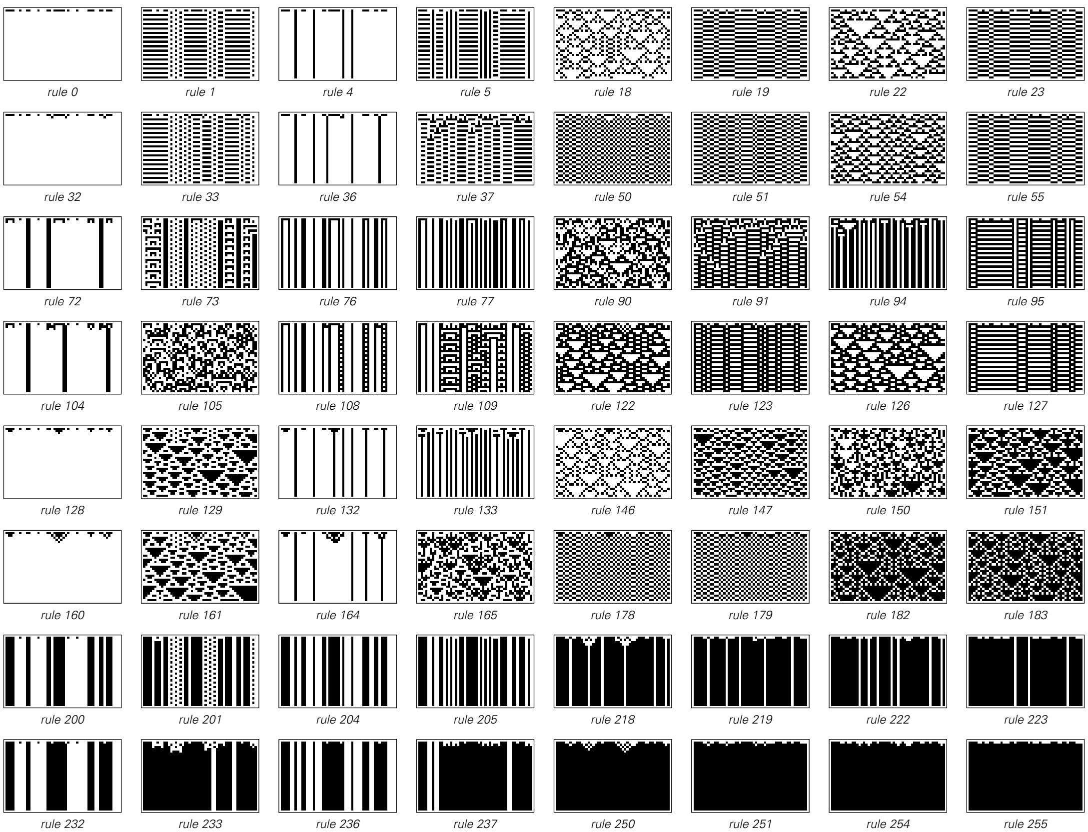

Examples of patterns produced by the evolution of each of the simplest possible symmetrical one-dimensional cellular automaton rules, starting from a random initial condition. The types of patterns obtained show striking similarities to those seen on mollusc shells from the previous page.

But which specific cellular automaton rule will any given mollusc use? The pictures at the bottom of the facing page show all the possible symmetrical rules that involve two colors and nearest neighbors. And comparing the patterns in these pictures with patterns on actual mollusc shells, one notices the remarkable fact that the range of patterns that occur in the two cases is extremely similar.

Traditional ideas might have suggested that each kind of mollusc would carefully optimize the pattern on its shell so as to avoid predators or to attract mates or prey. But what I think is much more likely is that these patterns are instead generated by rules that are in effect chosen at random from among a collection of the simplest possibilities. And what this means is that insofar as complexity occurs in such patterns it is in a sense a coincidence. It is not that some elaborate mechanism has specially developed to produce it. Rather, it just arises as an inevitable consequence of the basic phenomenon discovered in this book that simple rules will often yield complex behavior.

And indeed it turns out that in many species of molluscs the patterns on their shells—both simple and complex—are completely hidden by an opaque skin throughout the life of the animal, and so presumably cannot possibly have been determined by any careful process of optimization or natural selection.

So what about pigmentation patterns on other kinds of animals? Mollusc shells are almost unique in having patterns that are built up one line at a time; much more common is for patterns to develop all at once all over a surface.

Most often what seems to happen is that at some point in the growth of an embryo, precursors of pigment-producing cells appear on its surface, and groups of these cells associated with pigments of different colors then become arranged in a definite pattern. Typically each individual group of cells is initially some fraction of a tenth of a millimeter across. But since different parts of an animal usually grow at different rates, the final pattern that one sees on an adult animal ends up being scaled differently in different places—so that, for example, the pattern is smaller in scale on the head of an animal, since the head grows more slowly.

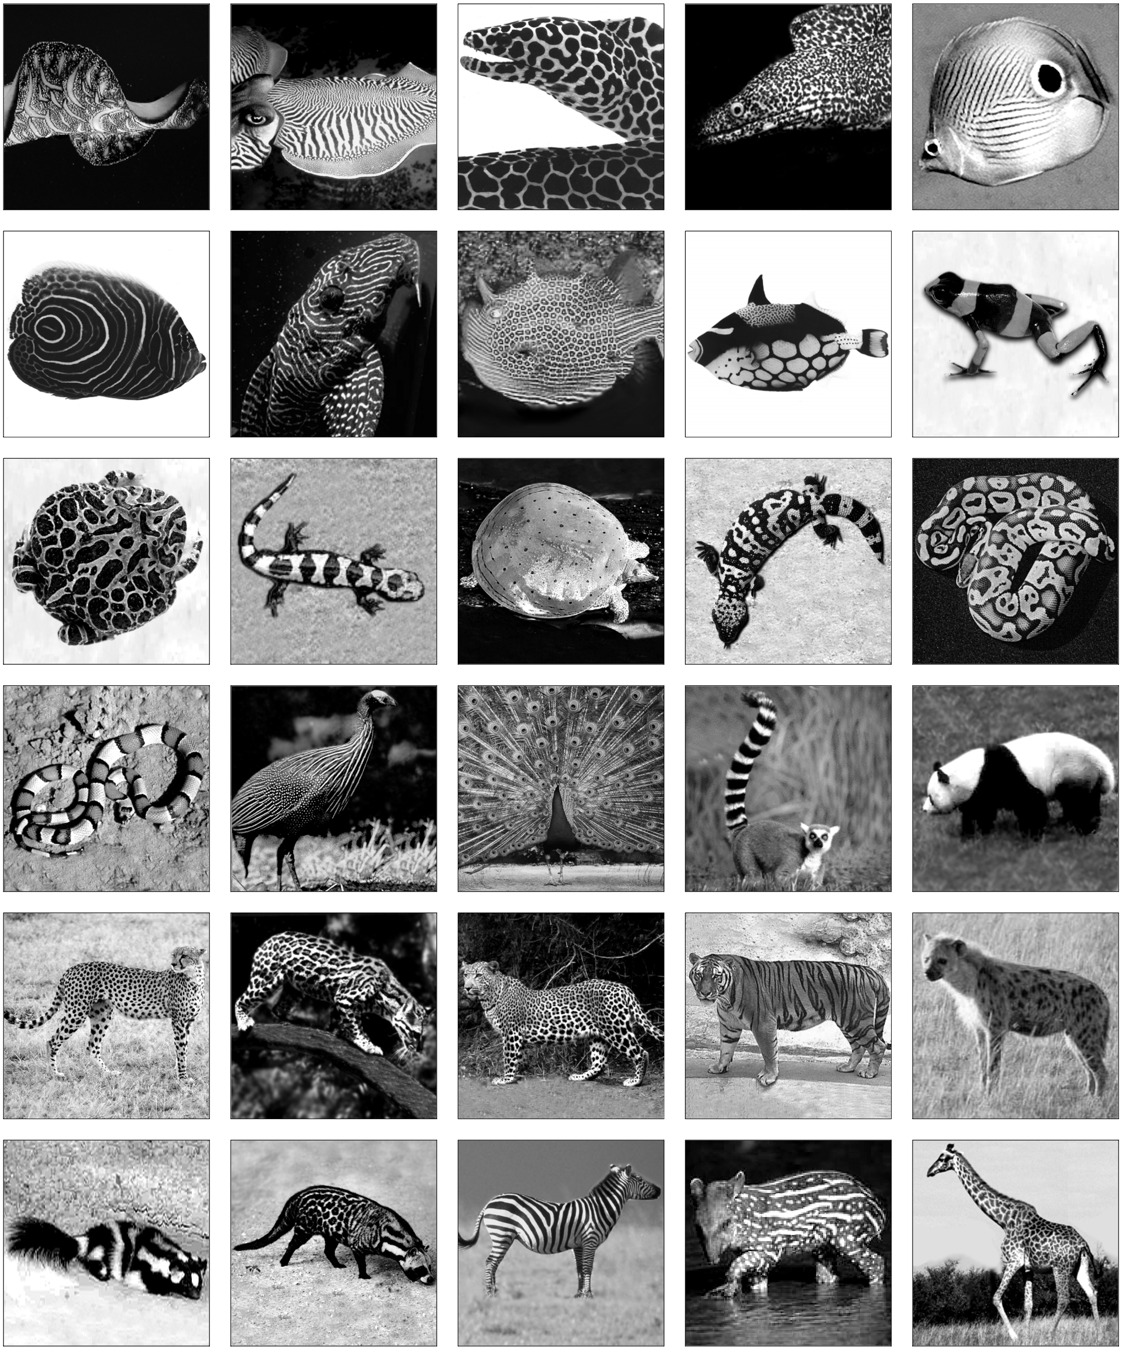

Typical examples of pigmentation patterns on animals. Note that many very different animals end up having remarkably similar patterns.

The pictures on the facing page show typical examples of pigmentation patterns in animals, and demonstrate that even across a vast range of different types of animals just a few kinds of patterns occur over and over again. So how are these patterns produced? Even though some of them seem quite complex, it turns out that once again there is a rather simple kind of rule that can account for them.

The idea is that when a pattern forms, the color of each element will tend to be the same as the average color of nearby elements, and opposite to the average color of elements further away. Such an effect could have its origin in the production and diffusion of activator and inhibitor chemicals, or, for example, in actual motion of different types of cells. But regardless of its origin, the effect itself can readily be captured just by setting up a two-dimensional cellular automaton with appropriate rules.

The pictures below show what happens with two slightly different choices for the relative importance of elements that are further away. In both cases, starting from a random distribution of black and white elements there quickly emerge definite patterns—in the first case a collection of spots, and in the second case a maze-like or labyrinthine structure.

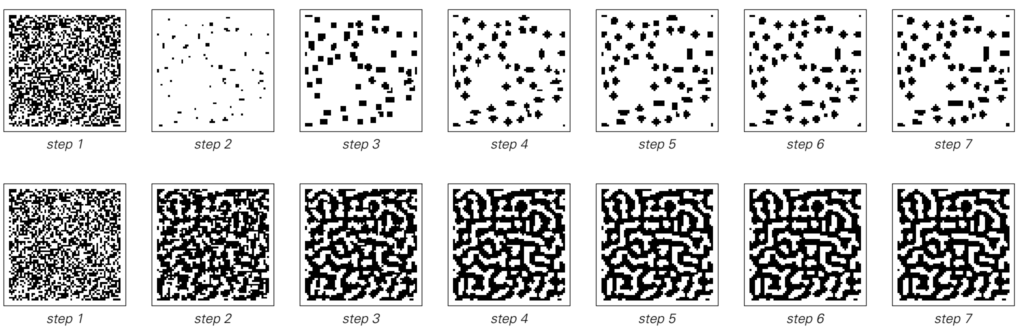

Evolution of simple two-dimensional cellular automata in which the color of each cell at each step is determined by looking at a weighted sum of the average colors of cells up to distance 3 away. In both rules shown the cell itself and its nearest neighbors enter with weight 1. Cells at distances 2 and 3 enter with negative weights— -0.4 per cell for the first rule, and -0.2 for the second. A cell becomes black if the weighted sum is positive, and white otherwise. Starting from random initial conditions, both rules quickly evolve to stationary states that look very much like pigmentation patterns seen in animals.

The next page shows the final patterns obtained with a whole array of different choices of weightings for elements at different distances. A certain range of patterns emerges—almost all of which turn out to be quite similar to patterns that one sees on actual animals.

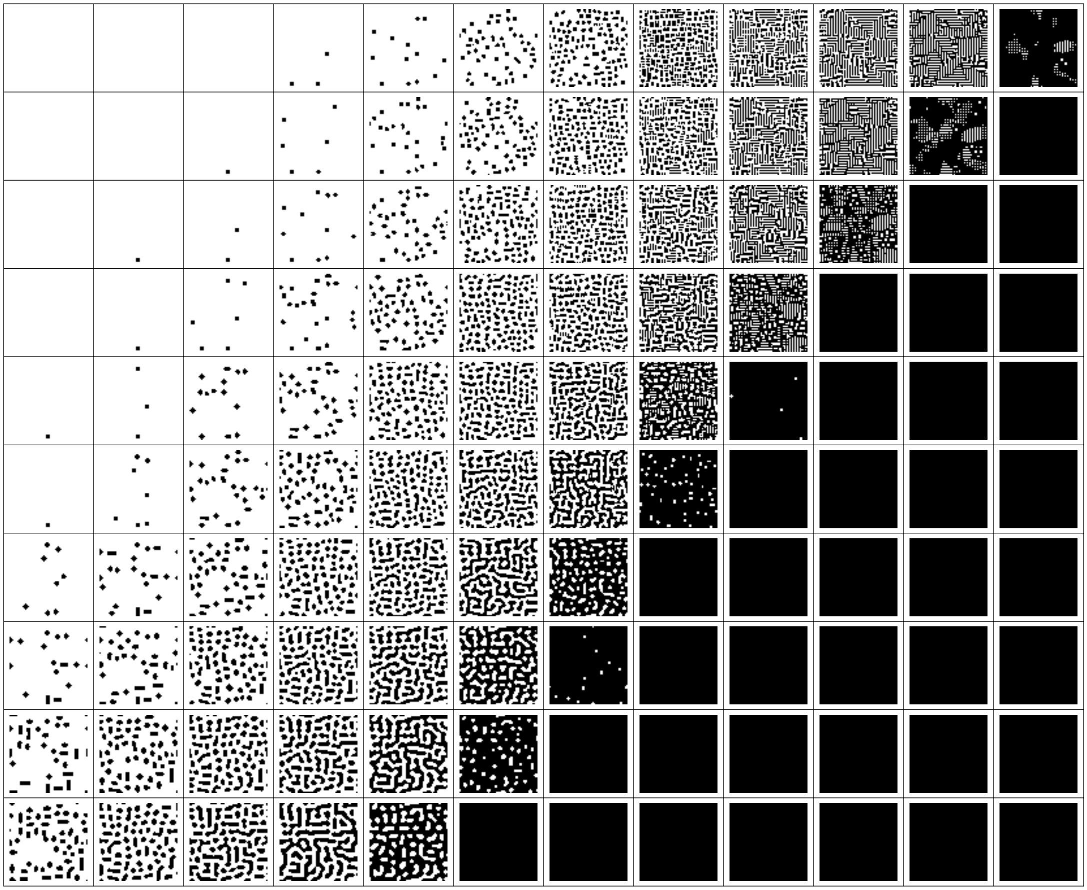

Patterns generated by rules of the type shown on the previous page, with a range of choices for the weights of cells at distances 2 and 3. Weights vary from -0.9 to 0 down the page for distance 2, and from -0.7 to 0.4 across the page for distance 3. In all cases the evolution starts from the same random initial condition, and is continued until it stabilizes. Note that pigmentation patterns for actual animals may contain either larger or smaller numbers of elements than the patterns shown here.

But all of these patterns in a sense have the same basic form in every direction. Yet there are many animals whose pigmentation patterns exhibit stripes with a definite orientation. Sometimes these stripes are highly regular, and can potentially arise from any of the possible mechanisms that yield repetitive behavior. But in cases where the stripes are less regular they typically look very much like the patterns generated in the pictures at the top of the facing page using a version of the simple mechanism described above.

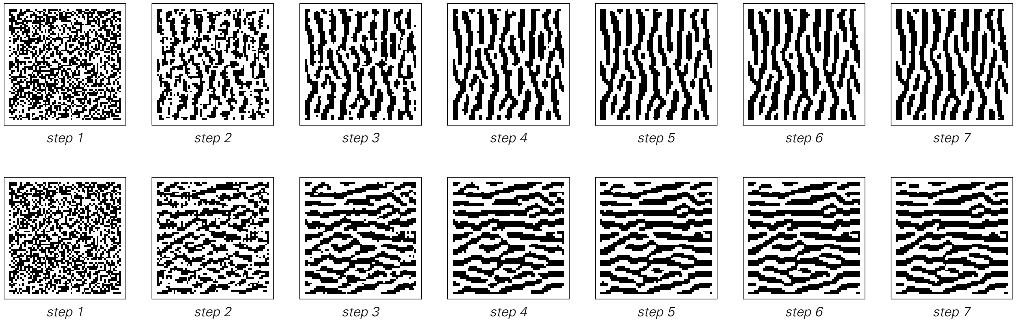

Examples of rules in which cells in the horizontal and vertical directions are weighted differently. In the first case, cells at distances 2 and 3 only have an effect in the vertical direction; in the second case, they only have an effect in the horizontal direction. The result is the formation of either vertical or horizontal stripes.

Financial Systems

During the development of the ideas in this book I have been asked many times whether they might apply to financial systems. There is no doubt that they do, and as one example I will briefly discuss here what is probably the most obvious feature of essentially all financial markets: the apparent randomness with which prices tend to fluctuate.

Whether one looks at stocks, bonds, commodities, currencies, derivatives or essentially any other kind of financial instrument, the sequences of prices that one sees at successive times show some overall trends, but also exhibit varying amounts of apparent randomness.

So what is the origin of this randomness?

In the most naive economic theory, price is a reflection of value, and the value of an asset is equal to the total of all future earnings—such as dividends—which will be obtained from it, discounted for the interest that will be lost from having to wait to get these earnings.

With this view, however, it seems hard to understand why there should be any significant fluctuations in prices at all. What is usually said is that prices are in fact determined not by true value, but rather by the best estimates of that value that can be obtained at any given time. And it is then assumed that these estimates are ultimately affected by all sorts of events that go on in the world, making random movements