EXCERPTED FROM

STEPHEN WOLFRAM
A NEW KIND OF SCIENCE

SECTION 2.3

Why These Discoveries Were Not Made Before

Why These Discoveries Were Not Made Before

The main result of this chapter—that programs based on simple rules can produce behavior of great complexity—seems so fundamental that one might assume it must have been discovered long ago. But it was not, and it is useful to understand some of the reasons why it was not.

In the history of science it is fairly common that new technologies are ultimately what make new areas of basic science develop. And thus, for example, telescope technology was what led to modern astronomy, and microscope technology to modern biology. And now, in much the same way, it is computer technology that has led to the new kind of science that I describe in this book.

Indeed, this chapter and several of those that follow can in a sense be viewed as an account of some of the very simplest experiments that can be done using computers. But why is it that such simple experiments were never done before?

One reason is just that they were not in the mainstream of any existing field of science or mathematics. But a more important reason is that standard intuition in traditional science gave no reason to think that their results would be interesting.

And indeed, if it had been known that they were worthwhile, many of the experiments could actually have been done even long before computers existed. For while it may be somewhat tedious, it is certainly possible to work out the behavior of something like a cellular automaton by hand. And in fact, to do so requires absolutely no sophisticated ideas from mathematics or elsewhere: all it takes is an understanding of how to apply simple rules repeatedly.

And looking at the historical examples of ornamental art on the facing page, there seems little reason to think that the behavior of many cellular automata could not have been worked out many centuries or even millennia ago. And perhaps one day some Babylonian artifact created using the rule 30 cellular automaton from page 27 will be unearthed. But I very much doubt it. For I tend to think that if pictures like the one on page 27 had ever in fact been seen in ancient times then science would have been led down a very different path from the one it actually took.

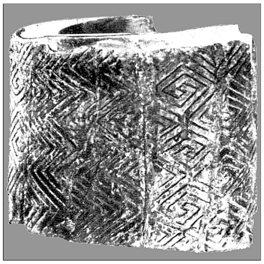

22,000 BC (Paleolithic)

3500 BC (Sumerian)

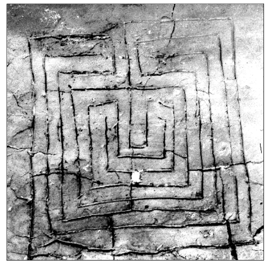

1200 BC (Greek)

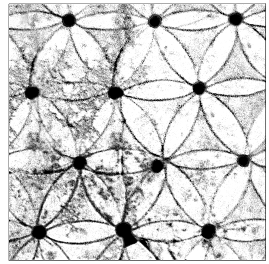

9th century BC (Phoenician)

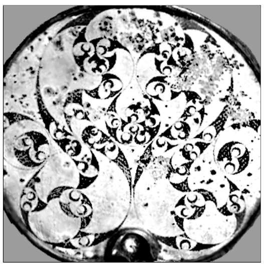

1st century BC (Celtic)

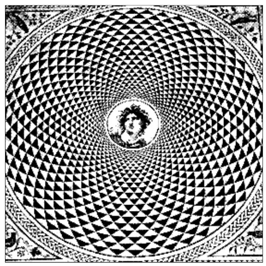

2nd century (Roman)

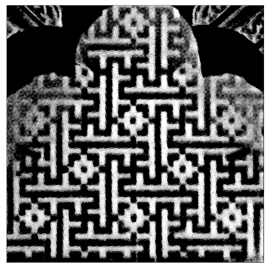

8th century (Islamic)

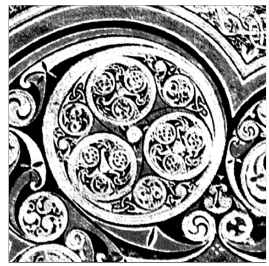

8th century (Celtic)

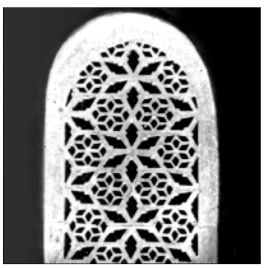

12th century (Italian)

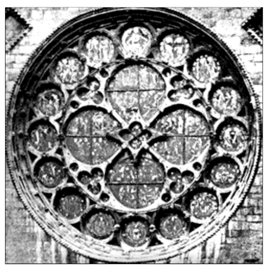

13th century (English)

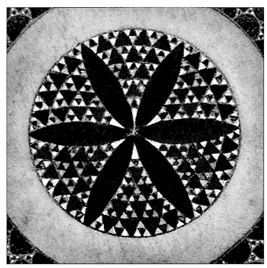

13th century (Italian)

13th century (Italian)

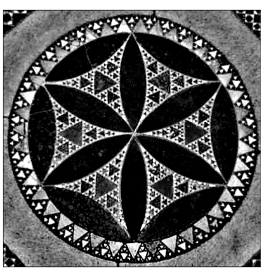

13th century (Italian)

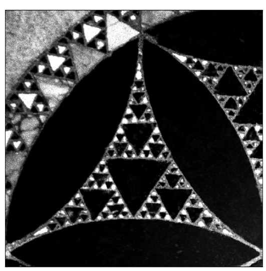

13th century (Italian)

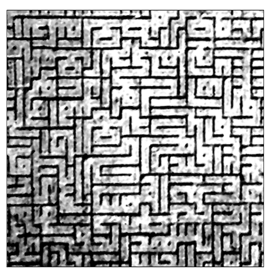

14th century (Islamic)

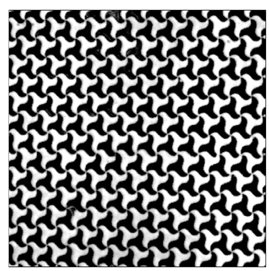

14th century (Islamic)

Historical examples of ornamental art. Repetitive patterns are common and some nested patterns are seen, but the more complicated kinds of patterns discussed in this chapter do not ever appear to have been used. Note that the second-to-last picture is not an abstract design, but is instead text written in a highly stylized form of Arabic script.

Even early in antiquity attempts were presumably made to see whether simple abstract rules could reproduce the behavior of natural systems. But so far as one can tell the only types of rules that were tried were ones associated with standard geometry and arithmetic. And using these kinds of rules, only rather simple behavior could be obtained—adequate to explain some of the regularities observed in astronomy, but unable to capture much of what is seen elsewhere in nature.

And perhaps because of this, it typically came to be assumed that a great many aspects of the natural world are simply beyond human understanding. But finally the successes based on calculus in the late 1600s began to overthrow this belief. For with calculus there was finally real success in taking abstract rules created by human thought and using them to reproduce all sorts of phenomena in the natural world.

But the particular rules that were found to work were fairly sophisticated ones based on particular kinds of mathematical equations. And from seeing the sophistication of these rules there began to develop an implicit belief that in almost no important cases would simpler rules be useful in reproducing the behavior of natural systems.

During the 1700s and 1800s there was ever-increasing success in using rules based on mathematical equations to analyze physical phenomena. And after the spectacular results achieved in physics in the early 1900s with mathematical equations there emerged an almost universal belief that absolutely every aspect of the natural world would in the end be explained by using such equations.

Needless to say, there were many phenomena that did not readily yield to this approach, but it was generally assumed that if only the necessary calculations could be done, then an explanation in terms of mathematical equations would eventually be found.

Beginning in the 1940s, the development of electronic computers greatly broadened the range of calculations that could be done. But disappointingly enough, most of the actual calculations that were tried yielded no fundamentally new insights. And as a result many people came to believe—and in some cases still believe today—that computers could never make a real contribution to issues of basic science.

But the crucial point that was missed is that computers are not just limited to working out consequences of mathematical equations. And indeed, what we have seen in this chapter is that there are fundamental discoveries that can be made if one just studies directly the behavior of even some of the very simplest computer programs.

In retrospect it is perhaps ironic that the idea of using simple programs as models for natural systems did not surface in the early days of computing. For systems like cellular automata would have been immensely easier to handle on early computers than mathematical equations were. But the issue was that computer time was an expensive commodity, and so it was not thought worth taking the risk of trying anything but well-established mathematical models.

By the end of the 1970s, however, the situation had changed, and large amounts of computer time were becoming readily available. And this is what allowed me in 1981 to begin my experiments on cellular automata.

There is, as I mentioned above, nothing in principle that requires one to use a computer to study cellular automata. But as a practical matter, it is difficult to imagine that anyone in modern times would have the patience to generate many pictures of cellular automata by hand. For it takes roughly an hour to make the picture on page 27 by hand, and it would take a few weeks to make the picture on page 29.

Yet even with early mainframe computers, the data for these pictures could have been generated in a matter of a few seconds and a few minutes respectively. But the point is that one would be very unlikely to discover the kinds of fundamental phenomena discussed in this chapter just by looking at one or two pictures. And indeed for me to do it certainly took carrying out quite large-scale computer experiments on a considerable number of different cellular automata.

If one already has a clear idea about the basic features of a particular phenomenon, then one can often get more details by doing fairly specific experiments. But in my experience the only way to find phenomena that one does not already know exist is to do very systematic and general experiments, and then to look at the results with as few preconceptions as possible. And while it takes only rather basic computer technology to make single pictures of cellular automata, it requires considerably more to do large-scale systematic experiments.

Indeed, many of my discoveries about cellular automata came as direct consequences of using progressively better computer technology.

As one example, I discovered the classification scheme for cellular automata with random initial conditions described at the beginning of Chapter 6 when I first looked at large numbers of different cellular automata together on high-resolution graphics displays. Similarly, I discovered the randomness of rule 30 (page 27) when I was in the process of setting up large simulations for an early parallel-processing computer. And in more recent years, I have discovered a vast range of new phenomena as a result of easily being able to set up large numbers of computer experiments in *Mathematica*.

Undoubtedly, therefore, one of the main reasons that the discoveries I describe in this chapter were not made before the 1980s is just that computer technology did not yet exist powerful enough to do the kinds of exploratory experiments that were needed.

But beyond the practicalities of carrying out such experiments, it was also necessary to have the idea that the experiments might be worth doing in the first place. And here again computer technology played a crucial role. For it was from practical experience in using computers that I developed much of the necessary intuition.

As a simple example, one might have imagined that systems like cellular automata, being made up of discrete cells, would never be able to reproduce realistic natural shapes. But knowing about computer displays it is clear that this is not the case. For a computer display, like a cellular automaton, consists of a regular array of discrete cells or pixels. Yet practical experience shows that such displays can produce quite realistic images, even with fairly small numbers of pixels.

And as a more significant example, one might have imagined that the simple structure of cellular automaton programs would make it straightforward to foresee their behavior. But from experience in practical computing one knows that it is usually very difficult to foresee what even a simple program will do. Indeed, that is exactly why bugs in programs are so common. For if one could just look at a program and immediately know what it would do, then it would be an easy matter to check that the program did not contain any bugs.

Notions like the difficulty of finding bugs have no obvious connection to traditional ideas in science. And perhaps as a result of this, even after computers had been in use for several decades, essentially none of this type of intuition from practical computing had found its way into basic science. But in 1981 it so happened that I had for some years been deeply involved in both practical computing and basic science, and I was therefore in an almost unique position to apply ideas derived from practical computing to basic science.

Yet despite this, my discoveries about cellular automata still involved a substantial element of luck. For as I mentioned on page 19, my very first experiments on cellular automata showed only very simple behavior, and it was only because doing further experiments was technically very easy for me that I persisted.

And even after I had seen the first signs of complexity in cellular automata, it was several more years before I discovered the full range of examples given in this chapter, and realized just how easily complexity could be generated in systems like cellular automata.

Part of the reason that this took so long is that it involved experiments with progressively more sophisticated computer technology. But the more important reason is that it required the development of new intuition. And at almost every stage, intuition from traditional science took me in the wrong direction. But I found that intuition from practical computing did better. And even though it was sometimes misleading, it was in the end fairly important in putting me on the right track.

Thus there are two quite different reasons why it would have been difficult for the results in this chapter to be discovered much before computer technology reached the point it did in the 1980s. First, the necessary computer experiments could not be done with sufficient ease that they were likely to be tried. And second, the kinds of intuition about computation that were needed could not readily have been developed without extensive exposure to practical computing.

But now that the results of this chapter are known, one can go back and see quite a number of times in the past when they came at least somewhat close to being discovered.

It turns out that two-dimensional versions of cellular automata were already considered in the early 1950s as possible idealized models for biological systems. But until my work in the 1980s the actual investigations of cellular automata that were done consisted mainly in constructions of rather complicated sets of rules that could be shown to lead to specific kinds of fairly simple behavior.

The question of whether complex behavior could occur in cellular automata was occasionally raised, but on the basis of intuition from engineering it was generally assumed that to get any substantial complexity, one would have to have very complicated underlying rules. And as a result, the idea of studying cellular automata with simple rules never surfaced, with the result that nothing like the experiments described in this chapter were ever done.

In other areas, however, systems that are effectively based on simple rules were quite often studied, and in fact complex behavior was sometimes seen. But without a framework to understand its significance, such behavior tended either to be ignored entirely or to be treated as some kind of curiosity of no particular fundamental significance.

Indeed, even very early in the history of traditional mathematics there were already signs of the basic phenomenon of complexity. One example known for well over two thousand years concerns the distribution of prime numbers (see page 132). The rules for generating primes are simple, yet their distribution seems in many respects random. But almost without exception mathematical work on primes has concentrated not on this randomness, but rather on proving the presence of various regularities in the distribution.

Another early sign of the phenomenon of complexity could have been seen in the digit sequence of a number like $\pi \simeq 3.141592653\ldots$ (see page 136). By the 1700s more than a hundred digits of π had been computed, and they appeared quite random. But this fact was treated essentially as a curiosity, and the idea never appears to have arisen that there might be a general phenomenon whereby simple rules like those for computing π could produce complex results.

In the early 1900s various explicit examples were constructed in several areas of mathematics in which simple rules were repeatedly applied to numbers, sequences or geometrical patterns. And sometimes nested or fractal behavior was seen. And in a few cases substantially more complex behavior was also seen. But the very complexity of this behavior was usually taken to show that it could not be relevant for real mathematical work—and could only be of recreational interest.

When electronic computers began to be used in the 1940s, there were many more opportunities for the phenomenon of complexity to be seen. And indeed, looking back, significant complexity probably did occur in many scientific calculations. But these calculations were almost always based on traditional mathematical models, and since previous analyses of these models had not revealed complexity, it tended to be assumed that any complexity in the computer calculations was just a spurious consequence of the approximations used in them.

One class of systems where some types of complexity were noticed in the 1950s are so-called iterated maps. But as I will discuss on page 149, the traditional mathematics that was used to analyze such systems ended up concentrating only on certain specific features, and completely missed the main phenomenon discovered in this chapter.

It is often useful in practical computing to produce sequences of numbers that seem random. And starting in the 1940s, several simple procedures for generating such sequences were invented. But perhaps because these procedures always seemed quite ad hoc, no general conclusions about randomness and complexity were drawn from them.

Along similar lines, systems not unlike the cellular automata discussed in this chapter were studied in the late 1950s for generating random sequences to be used in cryptography. Almost all the results that were obtained are still military secrets, but I do not believe that any phenomena like the ones described in this chapter were discovered.

And in general, within the context of mainstream science, the standard intuition that had been developed made it very difficult for anyone to imagine that it would be worth studying the behavior of the very simple kinds of computer programs discussed in this chapter. But outside of mainstream science, some work along such lines was done. And for example in the 1960s early computer enthusiasts tried running various simple programs, and found that in certain cases these programs could succeed in producing nested patterns.

Then in the early 1970s, considerable recreational computing interest developed in a specific two-dimensional cellular automaton known as the Game of Life, whose behavior is in some respects similar to the rule 110 cellular automaton discussed in this chapter. Great effort was spent trying to find structures that would be sufficiently simple and predictable that they could be used as idealized components for engineering. And although complex behavior was seen it was generally treated as a nuisance, to be avoided whenever possible.

In a sense it is surprising that so much could be done on the Game of Life without the much simpler one-dimensional cellular automata in this chapter ever being investigated. And no doubt the lack of a connection to basic science was at least in part responsible.

But whatever the reasons, the fact remains that, despite many hints over the course of several centuries, the basic phenomenon that I have described in this chapter was never discovered before.

It is not uncommon in the history of science that once a general new phenomenon has been identified, one can see that there was already evidence of it much earlier. But the point is that without the framework that comes from knowing the general phenomenon, it is almost inevitable that such evidence will have been ignored.

It is also one of the ironies of progress in science that results which at one time were so unexpected that they were missed despite many hints eventually come to seem almost obvious. And having lived with the results of this chapter for nearly two decades, it is now difficult for me to imagine that things could possibly work in any other way. But the history that I have outlined in this section—like the history of many other scientific discoveries—provides a sobering reminder of just how easy it is to miss what will later seem obvious.