THE WORLD ACCORDING TO WOLFRAM

Brian Hayes

A reprint from

American Scientist

the magazine of Sigma Xi, the Scientific Research Society

Volume 90, Number 4
July–August, 2002
pages 308–312

THE WORLD ACCORDING TO WOLFRAM

Brian Hayes

Arthur C. Clarke wrote: "Any sufficiently advanced technology is indistinguishable from magic." In the same vein, perhaps any sufficiently advanced science risks being mistaken for the raving of a crackpot.

This uncomfortable thought is prompted by a book that has landed on my desk with a five-pound thud. The title promises *A New Kind of Science*, and inside are claims no less extravagant. "I have discovered vastly more than I ever thought possible," the author's preface boasts, "and in fact what I have done now touches almost every existing area of science, and quite a bit besides." In 1,200 pages, this one volume—the work of one person—undertakes to explain the structure of space and time at the deepest level, clears up the mysteries of entropy and randomness, and elbows aside mathematics and statistics as central tools of scientific analysis. Along the way, the book also corrects the errors of Darwinism, shows where chaos theory went wrong, explains the forms of seashells, tree leaves and snowflakes, and puts human free will on a proper philosophical footing.

The author acknowledges that these grand ambitions may be met with a measure of skepticism. "If I myself were just to pick up this book today without having spent the past twenty years thinking about its contents, I have little doubt that I too would not believe many of the things it says." But he urges patience and perseverance. To assimilate the new ideas "will require an investment of years comparable to learning an area like physics." Those mired in the old kind of science may resist at first, but, "In time I expect that the ideas in this book will come to pervade not only science and technology but also many areas of general thinking. And with this its methods will eventually become a standard part of education—much as mathematics is today."

Such grandiose visions are often a telltale sign of the crank, and there are other reasons to be wary here. The author has been working in seclusion and secrecy for 10 years, and during that time has submitted none of his results to peer review. The publisher of the volume is the author's own company, so that even now there was been no opportunity for independent editorial judgment. The circumstances suggest a person cut off from the social context of science.

And yet the author of *A New Kind of Science* is not an outsider, and he is not a crank or a crackpot. He is Stephen Wolfram, physicist, computer scientist and entrepreneur. Wolfram published his first paper in particle physics at age 15 and earned a Ph.D. at 20. He was a precocious professor at Caltech, then moved to the Institute for Advanced Study and later the University of Illinois before leaving the academic world to create the software called Mathematica. His ideas deserve a reading even if they are presented in a manner that raises both eyebrows and hackles.

Before proceeding further I should disclose my own occasional interactions with Stephen Wolfram. As a magazine editor I once invited him to write an article; some years later I wrote a review of Mathematica (and accepted a complimentary copy of the software, as well as a T-shirt); on another occasion I wrote for *The Mathematica Journal*, which was sponsored by Wolfram Research, Inc. More recently, as a condition of being allowed to see *A New Kind of Science* in advance of publication, I signed a nondisclosure agreement, which has now expired. And a few weeks ago Wolfram and I met to talk about the book.

Playing by the Rules

Science is usually viewed as an inductive art, which starts with observations or experiments and proceeds toward laws of nature. Wolfram stands this process on its head, suggesting that we start by listing all possible laws of nature and then see which of them might correspond to observed events.

Is such a scheme feasible? If we set out to compile a catalogue of all conceivable natural laws, where would we begin, and how would we know when to stop? Wolfram answers by introducing abstract systems small enough and simple enough that we can enumerate all the rules or programs that might possibly govern their behavior. In these systems you don't have to go very far down the list before the rules start making toy worlds with interesting properties.

Wolfram first formulated these ideas in studies of cellular automata—geometric arrays of minimalist computing elements, called cells. Each cell

Brian Hayes is Senior Writer for American Scientist. *Address: 211 Dacian Avenue, Durham, NC 27701; bhayes@amsci.org*

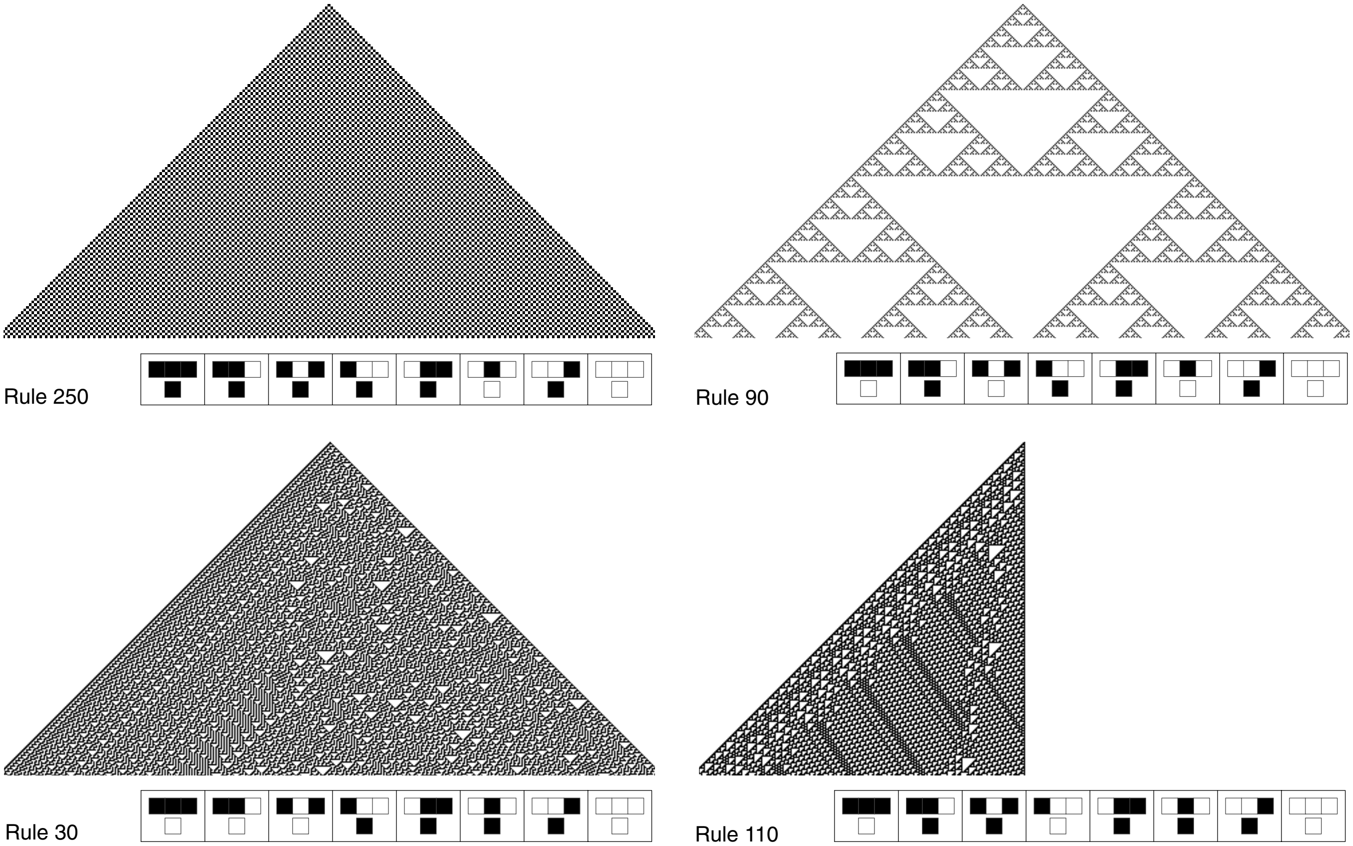

Figure 1. Four kinds of behavior seen in one-dimensional cellular automata were classified by Stephen Wolfram, who also devised the scheme for numbering the rules. Each panel is a spacetime diagram, with space extending horizontally and time proceeding downward, so that successive configurations of the system are drawn one below the other. In each case the initial configuration is a line of white cells with either one or two black cells in the center. All of the systems shown are among the simplest nontrivial cellular automata, with just two states per cell (*black* and *white*) and a neighborhood of three cells. Rule 250 yields simple alternation of black and white cells. Rule 90 produces a nested, recursive pattern. Parts of the pattern generated by Rule 30 show evidence of randomness, and Rule 110 gives rise to objects that move across the space and interact with one another. (The illustration of Rule 110 is shown at half the scale of the others.) Both the illustrations that accompany this article are taken from *A New Kind of Science* and are reproduced courtesy of Stephen Wolfram, LLC.

in an automaton has a finite number of possible states, and it spends its time continually recomputing this state. To do so it looks at its own present state and at the states of a few neighboring cells, then applies a fixed rule to determine its next state. As soon as all the cells have updated their states in this way, the process starts over. All the cells use the same rule and operate in synchrony.

In the early 1980s Wolfram began a systematic survey of one-dimensional cellular automata, where the cells are arranged in a line or a ring. The operation of a one-dimensional automaton is particularly easy to visualize: If you draw the line of cells horizontally and show successive states of the system going down the page, the result is a two-dimensional spacetime diagram that reveals the entire evolution of the system.

Consider a one-dimensional automaton where the cells have just two possible states (*black* and *white*) and where each cell interacts only with its nearest neighbors. There are eight possible configurations of a cell and its two neighbors, and for each of these configurations the next state can be either *black* or *white*. Hence there are 2^8 possible rules for the evolution of the system. These 256 rules are all the available candidates for laws of nature in this tiny universe.

Can anything interesting ever happen in a world of such meager resources? Wolfram identifies four broad categories of behavior, which are illustrated in Figure 1. First are simple repetitive patterns, such as all white cells or all black cells, or alternation in stripes and checkerboards. The second category includes self-similar, fractal patterns, whose characteristic feature Wolfram describes as "nesting." In the third category are rules that generate apparent randomness; although the output is not a totally patternless spatter of black and white cells, the sequence of states of a single cell can have the statistical properties of a random series. Finally, a few rules yield localized structures that maintain their coherence while moving through the cellular space. These persistent structures are reminiscent of particles being created and annihilated in a quantum field theory.

It was the random patterns and the persistent structures that caught Wolfram's attention. Where does all that complexity come from? The naive assumption might have been that simple rules would yield only simple behavior, and the complexity of the pattern would grow steadily along with the complexity of the underlying rules. But Wolfram observed a different relation: There is a threshold of complexity. If the rules are kept simple enough to stay below the threshold, the outcome is always rather dull. (For example, a cellular automaton in which each cell's next state depends only on that same cell's present state is inevitably repetitious.) But once the rules get just a little more intricate, all four categories of behavior appear. What's more, once the system has crossed the threshold, further embellishments of the rules have little effect. New details may appear, but all the patterns can still be assigned to the same four basic categories.

To provide evidence in support of this last assertion, Wolfram surveys an exuberant variety of systems—some familiar, some novel, all interesting in their own right. He looks at cellular automata with more than two states per cell, including systems where the range of possible states is a continuum. He considers cellular automata in two and three dimensions. He studies automata where the cells are updated one at a time rather than simultaneously. Leaving behind the theme of cellular automata, he studies Turing machines and grammar-driven systems where substitution rules allow a string of symbols to grow and change. He looks at differential equations and at iterated numerical functions. He even examines the sequence of prime numbers. In each case his conclusion is the same: Simple sets of rules can generate complex outputs, but piling further complications onto the rules leads to little additional complexity in the outcome.

A Cellular Universe

It's all very well to dabble in disembodied bits and streams of digits, but what has all this to do with science in our world of atoms and energy? Wolfram's answer is that the same kinds of simple rules or programs are found at work everywhere in the universe.

One case where the mapping between the two realms is quite direct comes from biology, where certain mollusk shells are decorated with patterns similar to those produced by some cellular automata. And the resemblance is surely not a coincidence: The shell patterns are deposited by a row of pigment-secreting cells (that is, *biological* cells) that appear to act much like a one-dimensional cellular automaton, with neighboring cells communicating through the exchange of chemical signals. These resemblances have been noted before, but Wolfram argues for an unusually strong version of the idea, claiming that all possible cellular automaton rules in a certain class are observed on mollusk shells.

Elsewhere in biology, Wolfram applies similar methods of analysis to other pigment patterns, to the arrangement of stems and branches in plants, and to the shapes of leaves. In physics he treats the growth of snowflakes and other crystals, the fracturing of solids and the onset of turbulence in fluids. There is even a brief discussion of economics, suggesting that the kind of randomness observed in some cellular automata could account for price fluctuations in stock markets.

In another chapter of *A New Kind of Science* Wolfram presents his version of the thesis that the universe as a whole is something like a cellular automaton. The model looks below the level of everyday experience and even beyond the events studied in high-energy physics, where the world seems to be made up of electrons and quarks and other "elementary" particles, moving through a continuum of space. In Wolfram's view of the universe there is no continuum, and particles are a mere epiphenomenon; indeed, motion and geometry are also little more than illusions. In this cosmology, space and time are both assumed to be discrete, just as they are in a cellular automaton,

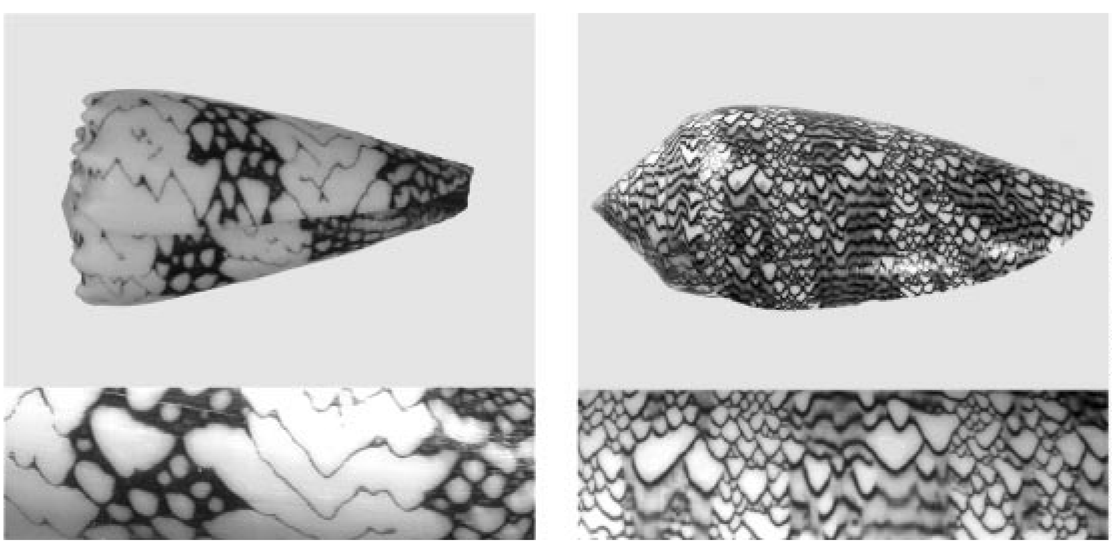

Figure 2. Pigment patterns on snail shells might be products of a biological system working something like a one-dimensional cellular automaton. Wolfram argues that various mollusks illustrate all possible patterns generated by a specific class of automata. The shells shown are identified as the banded marble cone *(left)* and the textile cone *(right)*.

and the only things that move are signals passing from cell to cell.

The most obvious way of implementing this idea would be to partition space into tiny cubical volumes, creating a cellular automaton on a three-dimensional grid. Wolfram looks with disfavor on this simplest solution because it imposes a particular geometry on space and also requires some kind of master clock to synchronize the updating of all the cells throughout the grid. His alternative is a model where the cells are nodes of a free-form network that has a well-defined topology but no specific geometry. In other words, the connections between nodes are all determined beforehand, but the spatial coordinates of the nodes are left unspecified. Concepts such as shape and position have no meaning at this level: The geometry of space emerges from the model rather than being built into it. Specifically, our world is perceived as being three-dimensional because each node of the network has three incoming links from other nodes and three outgoing links, creating the same connectivity as a three-dimensional lattice. Another feature of the model is that updating the cells requires no synchronizing master clock; updating events propagate along the links of the network itself. Making each link a one-way conduit defines a direction for time and causality; the whole structure is called a causal net.

This theory of everything is one of the book's wilder flights of fancy; there is no immediate prospect of testing it by experiment. But the same is true of all other attempts to explain the structure of the universe at this level of detail. In any case, Wolfram is not coy in his manner of proposing the model. When I asked him how seriously he intends it to be taken, he said he would be quite surprised if something very much like it doesn't turn out to be right.

What to Make of It

Wolfram warns that developing an intuition for his new kind of science will take months, "even for the most talented and open-minded people." For me I suppose it may take years. But even if I am premature in passing judgment, I want to give a preliminary assessment of a few major themes.

The core notion—that simple rules or programs can yield complex behavior—is surely both true and important. Whether it constitutes "a new kind of science" remains to be seen.

A closely related point is Wolfram's insistence that programs or algorithms are the best way of expressing ideas in the sciences. In particular, explicit rules of evolution are preferable to equations, which merely state the constraints that a system must satisfy without necessarily showing how to satisfy them. Perhaps he's right; my own experience is that I understand best what I can program. But Wolfram would expand this observation into a broad indictment of "the mathematical framework traditionally used in the exact sciences," which is reckless overkill. Wolfram obviously needs that framework (and uses it expertly) in his own work.

The concepts of randomness and complexity are central to the argument of this book, and yet Wolfram is curiously lax about defining them. He doesn't address the issue directly until 550 pages into his narrative, after many references to "intrinsic generation of randomness" in cellular automata and other simple systems. If you are accustomed to thinking of randomness as an inherent property of a pattern—something that can be traced back to the way it was created—"intrinsic generation" makes no sense in this context, because the mechanism that created the pattern is totally deterministic. It turns out that Wolfram defines randomness and complexity in terms of how patterns are perceived rather than how they are created. Roughly speaking, if it looks random, it *is* random. Fair enough, but it remains unclear just what is being generated intrinsically.

With the exception of the cosmic causal net, the examples that illustrate applications of Wolfram's ideas are strangely bland. Snowflakes, fluid turbulence, branching in plants, pigment patterns in animals—these are all rather shopworn specimens, which have long been explained by models of the same general type. (Alan Turing gave a computational account of leopard spots and zebra stripes 50 years ago.) If the new kind of science is to have much generality, it will need to show its worth in other areas. In developmental biology, for example, can we write a simple program that explains the complex structure of *Caenorhabditis elegans*? Anatomists have traced the paths of all 959 somatic cells in this worm, but expressing the underlying algorithm in terms of a few simple rules looks like a challenge.

Wolfram's comments on evolutionary biology are perhaps the lamest passages of the entire book. Noting that some traits of some organisms seem to explore the entire space of available variations, he concludes that Darwinian selection can't be acting on those traits. "It is my suspicion," he writes, "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." He writes as if he were unaware that a debate between neutralists and selectionists had ever entered biology.

All of the programmable systems explored in *A New Kind of Science* have a distinctive trait in common: They have a densely occupied space of programs. In a cellular automaton, for example, any rule relating a neighborhood configuration to a next state is a valid program. Other programmable systems—such as desktop computers and the DNA-reading apparatus of the living cell—are much choosier about what they will recognize as a valid program. If you try feeding them random strings of bits or random sequences of nucleotides, you're in for a frustrating experience. It's not obvious to me how the paradigm of complex behavior from simple rules can be extended to such systems.

Clarity and Modesty

So much for the content of *A New Kind of Science.* The book also has a personal and historical context that requires comment.

In a note titled "clarity and modesty," Wolfram explains that he had to sacrifice the latter to attain the former. But immodesty is only half the issue here. The problem is not just the rosy spotlight that Wolfram shines upon himself at center stage; it's also the utter darkness that enshrouds all the other actors in this drama. The main text of *A New Kind of Science* (850 pages) names no names at all; the only work attributed to a specific individual is Wolfram's. The notes at the end of the book (another 350 pages in smaller type) do mention names of people, but briefly, grudgingly and often dismissively. (It's remarkable how many discoverers failed to appreciate the significance of their own work.) And many of the historical notes manage to present an anonymized history, written in the passive voice: "By the end of the 1950s it had been noted that...." "Over the course of the 1960s constructions were found...."

The book has no bibliography; the only references listed are Wolfram's own publications.

Here is a precis of Wolfram's history of cellular automata. He discovered them in 1981 (although he had made important precursor experiments earlier, as a teenager). Some time later he learned that John von Neumann had had the idea 30 years earlier. But von Neumann missed making the crucial discovery that simple rules could produce complex behavior, and so did others who toyed with the systems in the intervening years. By the late 1970s, "research on systems equivalent to cellular automata had largely petered out." But then the publication of Wolfram's papers redefined and reinvigorated the field, drawing in many followers—although most of them went off on the wrong tangent or wasted their time on trivial details.

It didn't happen that way. To begin with, interest in cellular automata did *not* peter out in the 1970s; it was thriving. The field continued to grow in the 1980s, when Wolfram's participation doubtless helped, but more important was work on the physics of computation and reversible cellular automata. Wolfram took no part in that. The main actors were Edward Fredkin, Charles Bennett, Tommaso Toffoli and Norman Margolus—who were also, as it happens, the ones who explained to Wolfram in 1981 the nature and historical context of his own work.

Wolfram's telling of the mollusk-shell story is another notable example of delusional history. The shells have been known since antiquity, he says, "but almost no efforts to understand the origins of such patterns seem ever to have been made." He then mentions one such effort, by C. H. Waddington and Russell Cowe, but says it was too narrow. He ignores entirely a cellular automaton devised in 1973 by G. T. Herman and W. H. Liu to model one of the shell patterns. According to Wolfram, only his 1982 discovery brought the matter to wide attention and led others to take it up. Among those others was Hans Meinhardt, but Wolfram disparages his models as being too elaborate. In fact, Meinhardt came to the problem independently of Wolfram; what brought the shell to his attention was seeing it on his dinner plate in an Italian restaurant. And his models are elaborate because they account for the specific features observed on actual shells, rather than simply declaring that all shells look like one cellular automaton or another.

At the end of a long lunch with Wolfram, our conversation turned to these matters of history and attribution, then drifted on to a related topic. I asked why he had not given his new kind of science a name. In the book, he refers constantly to "the new kind of science I describe in this book," which gets cumbersome after the first 50 repetitions. He said he'd struggled to come up with a suitable name, but nothing quite fit the bill. At Wolfram Research, people spoke of "NKS," he said, or sometimes "Wolfram science."

"Are there any sciences named for people?" I wondered aloud.

"Well, there's Newtonian physics," he replied.

It was not the first time the names Wolfram and Newton have been mentioned in the same breath, and I suppose it might be taken as further evidence of an ego bursting all bounds. But I see it in another light: He is simply too modest to name the field Wolframian science. That part is left to us.

Bibliography

Fredkin, Edward. 1992. A new cosmogony. http://www.digitalphilosophy.org/new_cosmogony.htm

Herman, G. T., and W. H. Liu. 1973. The daughter of Celia, the French flag and the firing squad: Progress report on a cellular linear iterative-array simulator. *Simulation* 21:33–41.

Meinhardt, Hans. 1995. *The Algorithmic Beauty of Sea Shells.* With contributions and images by Przemyslaw Prusinkiewicz and Deborah R. Fowler. New York: Springer-Verlag.

Toffoli, Tommaso, and Norman Margolus. 1990. Invertible cellular automata: A review. *Physica D* 45:229-253

Turing, Alan M. 1952. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society,* Series B 237:37–72.

von Neumann, J. 1966. *Theory of Self-Reproducing Automata.* Edited and completed by A. Burks. Champaign, Ill.: University of Illinois Press.

Waddington, C. H., and Russell J. Cowe. 1969. Computer simulation of a molluscan pigmentation pattern. *Journal of Theoretical Biology* 25:219–225.

Wolfram, Stephen. 1986. *Theory and Applications of Cellular Automata.* Singapore: World Scientific.

Wolfram, Stephen. 2002. *A New Kind of Science.* Champaign, Ill.: Wolfram Media, Inc.