EXCERPTED FROM

STEPHEN WOLFRAM

A NEW KIND OF SCIENCE

General Notes

General Notes

■ **Website.** A large amount of additional material related to this book and these notes will progressively be made available through the website www.wolframscience.com. (See also the copyright page at the beginning of the book.)

■ **The role of these notes.** The material in these notes is intended to be complementary to the main text, and is not always self-contained on its own. It is thus important to read these notes in parallel with the sections of the main text to which they refer, since some necessary points may be made only in the main text. Captions to pictures in the main text also often contain details that are not repeated in these notes.

■ **Writing style.** This book was not easy to write, not least because it contains many complex intellectual arguments presented in plain language. And in order to make these arguments as easy to understand as possible, I have had to adopt some rhetorical devices. Perhaps most annoying to those with a copyediting orientation will be my predilection for starting sentences with conjunctions. The main reason I have done this is to break up what would otherwise be extremely long sentences. For the points that I make are often sufficiently complex to require quite long explanations. And to make what I have written more readable than, say, a typical classic work of philosophy, I have broken these explanations into several sentences, necessarily with conjunctions at the beginning of each. Also annoying to some will be my widespread use of short paragraphs. In the main text I normally follow the principle that any paragraph should communicate just one basic idea. And my hope is then that after reading each paragraph readers will pause a moment to absorb each idea before going on to the next one. (This book introduces the third major distinct style of writing that I have used in publications. The first I developed for scientific papers; the second for documents like *The Mathematica Book*.)

■ **Billions.** Following standard American usage, billion in this book means 10^9, trillion 10^{12}, and so on.

■ **Clarity and modesty.** There is a common style of understated scientific writing to which I was once a devoted subscriber. But at some point I discovered that more significant results are usually incomprehensible if presented in this style. For unless one has a realistic understanding of how important something is, it is very difficult to place or absorb it. And so in writing this book I have chosen to explain straightforwardly the importance I believe my various results have. Perhaps I might avoid some criticism by a greater display of modesty, but the cost would be a drastic reduction in clarity.

■ **Explaining ideas.** In presenting major new ideas in a book such as this, there is a trade-off between trying to explain these ideas directly on their own, and using previous ideas to provide a context. For some readers there is a clear short-term benefit in referring to previous ideas, and in discussing to what extent they are right and wrong. But for other readers this approach is likely just to introduce confusion. And over the course of time the ideas that typical readers know will tend to shift. So to make this book as broadly accessible as possible what I mostly do is in the main text to discuss ideas as directly as I can—but then in these notes to outline their historical context. Occasionally in the main text I do mention existing ideas—though I try hard to avoid fads that I expect will not be widely remembered within a few years. Throughout the book my main goal is to explain new ideas, not to criticize ones from the past. Sometimes clarity demands that I say explicitly that something from the past is wrong, but generally I try to avoid this, preferring instead just to state whatever I now believe is true. No doubt this book will draw the ire of some of those with whose ideas its results do not agree, but much as I might like to do so, I cannot realistically avoid this just by the way I present what I have discovered.

■ **Technology references.** In an effort to make the main text of this book as timeless as possible, I have generally avoided referring to everyday systems whose character or name I expect will change as technology advances. Inevitably, however, I do discuss computers, even though I fully expect that some of the terms and concepts I use in connection with them will end up seeming dated in a matter of a few decades.

▪ **Whimsy.** Cellular automata and most of the other systems in this book readily admit various kinds of whimsical descriptions. The rule 30 cellular automaton, for example, can be described as follows. Imagine a stadium full of people, with each person having two cards: one black and one white. Make the person in the middle of the top row of seats hold up a black card, and make everyone else in that row hold up a white card. Now each successive person in each successive row determines the color of the card they hold up by looking at the person directly above them, and above them immediately to their left and right, and then applying the simple rule on page 27. A photograph of the stadium will then show the pattern produced by rule 30. Descriptions like this may make abstract systems seem more connected to at least artificial everyday situations, but if the goal is to focus on fundamental ideas, as in this book, then such whimsy is, in my experience, normally just a major distraction.

▪ **Timeline of writing.** I worked on the writing of this book with few breaks for a little over ten years, beginning in June 1991, and ending in January 2002. The chapters were written roughly as follows: Chapter 1: 1991, 1999, 2001; Chapter 2: 1991–2; Chapter 3: 1992; Chapter 4: 1992–3; Chapter 5: 1993; Chapter 6: 1992–3; Chapter 7: 1994–6; Chapter 8: 1994–5, 1997; Chapter 9: 1995–8, 2001; Chapter 10: 1998–9; Chapter 11: 1995; Chapter 12: 1999–2001. Some sections of chapters (usually later ones) were added well after the rest. These notes were also sometimes written well after the main text of a given chapter.

▪ **Identifying new material.** The vast majority of results in this book have never appeared in published form before. A few were however included—implicitly or explicitly—in publications of mine from the early 1980s (see page 881). Whenever I am aware of antecedents to major material in the main text I have indicated this in the notes. Within the notes themselves, results that are given without historical discussion and without statements such as "it is known that" are generally new to this book. Researchers seeking further information should consult the website for the book.

▪ **Citations and references.** In developing the ideas described in this book I have looked at many thousands of books, papers and websites—and have interacted with hundreds of people (see page xiii). But rather than trying to give a huge list of specific references, I have instead included in these notes historical information tracing key contributions. From the names of concepts and people that I mention, it is straightforward to do web or database searches that give a vastly more complete picture of available references than could possibly fit in a book of manageable size—or than could be created correctly without immense scholarship. Note that while most current works of science tend to refer mainly just to very recent material, this book often refers to material that is centuries or even millennia old—in some ways more in the tradition of fields like philosophy.

▪ **Historical notes.** I have included extensive historical notes in this book in part out of respect for what has gone before, in part to provide context for ideas (and to see how current beliefs came to be as they are) and in part because the steps one goes through in understanding things may track steps that were gone through historically. Often in the book my conclusions in a particular field differ in a fundamental way from what has been traditional, and it has been important to me in confirming my understanding to study history and see how the conclusions I have reached were missed before. My discussion of science in this book is generally quite precise, being based among other things on computer experiments that can readily be reproduced. But my discussion of history is inevitably less precise. And while I have gone to considerable effort to ensure that its main elements are correct, ultimate objective confirmation is usually impossible. I have always tried to read original writings—for I have often found that later characterizations drop elements crucial for my purposes, or recast history to simplify pedagogy. But even for pieces of history where the people involved are still alive there are often no primary written records, leaving me to rely on secondary sources and recollections extracted in personal interviews—which are inevitably colored by later ideas and understanding. And while with sufficient effort it is usually possible to give fairly simple explanations for fundamental ideas in science, the same may not be true of their history. Looking at the historical notes in this book one striking feature is how often individuals of significant fame are mentioned—but not for the reason they are usually famous. And perhaps the explanation for this is in part that most of those who one can now see made contributions to the kinds of foundational issues I address were capable enough to have been successful at something—but without the whole context of this book they tended to view the types of results I discuss largely as curiosities, and so never tried to do much with them. Note that in mentioning people in connection with ideas and results, I have tried to concentrate on those who seemed to make the most essential contributions for my purposes—even when this does not entirely agree with traditions or criteria in particular academic fields.

■ **Dates.** Rather than following the usual academic practice of giving years when the discoveries were first published in books or journals, I have when possible given years when discoveries were first made. Note that I use a form like 1880s to refer to a decade, and 1800s to refer to a whole century.

■ **Autobiographical elements.** Every discovery in this book has some kind of specific personal story associated with it. Sometimes the story is quite straightforward; sometimes it is convoluted and colorful. But much as I enjoy recounting such stories, I have chosen not to make them part of this book.

■ **Cover image.** The image on the cover of this book is derived from the first 440 or so steps (with perhaps 10 at each end cut off by trimming) of the pattern generated by evolution according to the rule 110 cellular automaton discussed on page 32, with an initial condition consisting of repeats of □□□■■■■■ followed by repeats of ■■■■■□□□. The picture on the right shows 3000 steps in this evolution. The central region grows by 1 cell every 2 steps on the left and 22 cells every 340 steps on the right. Many persistent structures emanate from the right-hand edge of the region. After just 29 steps, this edge takes on a form that repeats every 1700 steps. During each such cycle, a total of 65 persistent structures are produced, of 11 of the 15 kinds from page 292, and their interactions make the full repetition period 6800 steps.

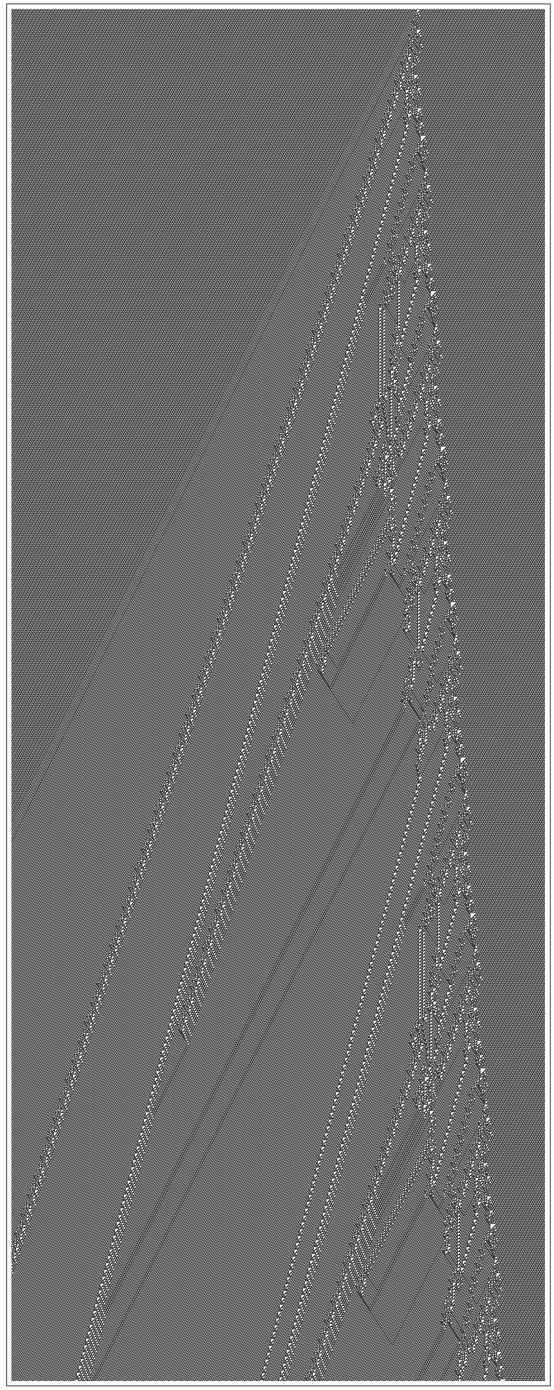

■ **Endpapers.** The goldenrod pages inside the front cover show the center 900 or so cells of the first 500 or so steps in the evolution of the rule 30 cellular automaton of page 29 from a single black cell. The pages inside the back cover show the next 500 or so steps.

■ **Using color.** Aside from practicalities of printing, what made me decide not to use color in this book were issues of visual perception. For much as it is easier to read text in black and white, so also it is easier to assimilate detailed pictures if they are just in black and white. And in fact many types of images in this book show quite misleading features in color. In human visual perception the color of something tends to seem different depending on what is around it—so that for example a red element tends to look purple or pink if the elements around it are respectively blue or white. And particularly if there are few colors arranged in ways that are not visually familiar it is typical for this effect to make all sorts of spurious patterns appear.

■ **Pictures in the book.** All the diagrammatic pictures in this book were created using *Mathematica*. (The photographs were also laid out and image-processed using *Mathematica*.) The ability of *Mathematica* to manipulate graphics in a symbolic way was crucial—and was what ultimately made it possible for the book to have so many elaborate pictures.

To those familiar with book layout it may seem surprising that I was able to include so many pictures of so many different shapes and sizes without having to resort to a device like figure numbers. And indeed it required solving innumerable small geometrical puzzles to do so. But what ultimately made it possible was that the *Mathematica* programs for generating the pictures were almost always general enough that it was straightforward for me to get, say, a picture with a different number of cells or steps.

▪ **Hyphenation.** An unusual feature of the text in this book is that it almost never uses hyphenation; from seeing so much word-wrapped text on computers I at least have come to view hyphenation as an ugly and misleading device.

▪ **Book production system.** Beyond its actual content, the production of this book was a highly complex process that relied on the methodology for software releases developed at my company over the past fifteen years. Had I been starting the book now I would likely have authored all of it directly in *Mathematica* and *Publicon*. But a decade ago I made the decision to compose all the original source for the book in FrameMaker. This source was then processed by an elaborate automated procedure much like a standard software build. The first step involved converting a MIF version of the complete source into a *Mathematica* symbolic expression. Then within *Mathematica* various transformations and tests were done on this expression—with for example every program in these notes being formatted and broken into lines using rules similar to *Mathematica* StandardForm. The resulting symbolic expressions were then converted back to MIF, formatted in FrameMaker, and automatically output as PDF. (Note that special characters in programs are rendered using the new Mathematica-Sans font specifically created for the book.) (See also the colophon at the very end of the book.)

▪ **Printing.** Many of the pictures in this book have a rather different character from things that are normally printed. For unlike traditional diagrams consisting of separate visible elements—or photographs involving smooth gradations of color—they often for example contain hundreds of cells per inch, each in effect independently black or white. And to capture this properly required careful sheet-fed printing on paper smooth enough to avoid significant spreading of ink. (See also the colophon at the very end of the book.)

▪ **Index.** In the index to this book I have tried to cater both to those who have already read the book in detail, and to those who have not. My approach has generally been to include any term that might realistically come to mind when thinking of a given topic—or remembering what the book says about that topic. And this means that even if the book mentions a term only in passing, I have tended to include it if for one reason or another I think it is likely to be memorable to people with certain experience or interests. Note that looking up *Mathematica* functions used in connection with some issue is often a good way to identify related issues. In the actual building of the index in this book, sorting, processing and checking were done using a variety of automated *Mathematica* procedures, operating on a symbolic representation of the full text and index of the book. Often it is possible by reading an index to identify the important issues in a book. And to some extent that is possible here, though often the presence of more subentries just reflects material being more spread out, not more important.

▪ **People in the index.** Conventions for personal names vary considerably with culture and historical period. I have tried in the index to give all names in the form they might be used on standardized documents in the modern U.S. I have done standard transliterations from non-Latin character sets. I give in full those forenames that I believe are or were most commonly used by a particular individual; for other forenames (including for example Russian patronymics) I give only initials. I normally give formal versions of forenames—though for individuals I have personally known I give in the text the form of forenames I would normally use in addressing them. I have dropped all honorifics or titles, except when they significantly alter a name. When there are several versions of a name, I normally use the one that was current closest to the time of work I mention. For each person in the index I list the country or countries where that person predominantly worked. Note that this may not reflect where the person was born, educated, did military service, or died. Rather, it tries to indicate where the person did the majority of their work, particularly as it relates to this book. I generally refer to countries or regions by the names of their closest present-day approximations, as these might appear in postal addresses. When borders have changed, I tend to favor the country whose language is what the person normally speaks or spoke. I usually list countries in the order that a person has worked in them, ignoring repeats. Note that while many of the people listed are well known, extensive research (often through personal contacts, as well as institutional and government records) was required to track down quite a few of them. (Ending dates are obviously not included for people who died after the writing of this book was finished in January 2002.)

▪ **Notation.** In the main text, I have almost entirely avoided any kind of formal symbolic notation—usually relying instead on diagrammatic pictures. In these notes, however, it will often be convenient to use such notation to give precise and compact representations of objects and operations. In the past, essentially the only large-scale notation available for theoretical science has been traditional mathematical notation. But on its own this would do me little good—for I need to represent not only traditional mathematics, but also more general rules and programs, as well as procedures and algorithms. But one of the reasons I created the *Mathematica* language was precisely to provide a much more general notation. So in these notes I use this language throughout as my notation. And this has many important advantages—and indeed it is hard to imagine that I would ever have been able to write these notes without it. One point is that it is completely uniform and standardized: there can never be any hidden assumptions or ambiguity about what a particular piece of notation means, since ultimately it is defined by the actual *Mathematica* software system and its documentation (see below). In cases where there is traditional mathematical notation for something, the corresponding *Mathematica* notation is normally almost identical—though occasionally a few details are changed to avoid ambiguity. The concept that everything is a symbolic expression allows *Mathematica* notation, however, to represent essentially any kind of abstract object. And when it comes to procedures and algorithms, the primitives in the *Mathematica* language are chosen to make typical steps easy to represent—with the result that a single line of *Mathematica* can often capture what would otherwise require many paragraphs of English text (and large amounts of pseudocode, or lower-level computer language code). Another very important practical feature of *Mathematica* notation is that by now a large number of people are familiar with it—certainly more than are for example familiar with sophisticated traditional notation in, say, mathematical logic. And the final and very critical advantage of *Mathematica* notation is that one can not only read it, but also actually execute it on a computer, and interact with it. And this makes it both vastly easier to apply and build on, and also easier to analyze and understand.

▪ ***Mathematica.*** I created *Mathematica* to be an integrated language and environment for computing in general, and technical computing in particular. Following its release in 1988, *Mathematica* has become very widely used in science, technology, education and elsewhere. (It is now also increasingly used as a component inside other software systems.)

Mathematica is available from Wolfram Research for all standard computer systems; much more information about it can be found on the web, especially from www.wolfram.com. There are many books about *Mathematica*—the original one being my *The Mathematica Book*.

The core of *Mathematica* is its language—which is based on the concept of symbolic programming. This language supports most traditional programming paradigms, but considerably generalizes them with the ideas of symbolic programming that I developed for it. In recent years there has started to be increasing use of the language component of *Mathematica* for all sorts of applications outside the area of technical computing where *Mathematica* as a whole has traditionally been most widely used.

The programs in these notes were created for *Mathematica* 4.1 (released 2000). They should run without any change in all subsequent versions of *Mathematica*, and the majority will also run in prior versions, all the way back to *Mathematica* 1 (released 1988) or *Mathematica* 2 (released 1990). Most of the programs require only the language component of *Mathematica*—and not its mathematical knowledge base—and so should run in all software systems powered by *Mathematica*, in which language capabilities are enabled.

Here are examples of how some of the basic *Mathematica* constructs used in the notes in this book work:

▪ Iteration

```
Nest[f, x, 3]⟶f[f[f[x]]]
NestList[f, x, 3]⟶{x, f[x], f[f[x]], f[f[f[x]]]}
Fold[f, x, {1, 2}]⟶f[f[x, 1], 2]
FoldList[f, x, {1, 2}]⟶{x, f[x, 1], f[f[x, 1], 2]}
```

▪ Functional operations

```
Function[x, x + k][a]⟶a + k
(# + k &)[a]⟶a + k
(r[#1] + s[#2] &)[a, b]⟶r[a] + s[b]
Map[f, {a, b, c}]⟶{f[a], f[b], f[c]}
Apply[f, {a, b, c}]⟶f[a, b, c]
Select[{1, 2, 3, 4, 5}, EvenQ]⟶{2, 4}
MapIndexed[f, {a, b, c}]⟶{f[a, {1}], f[b, {2}], f[c, {3}]}
```

▪ List manipulation

```
{a, b, c, d}⟦3⟧⟶c
{a, b, c, d}⟦{2, 4, 3, 2}⟧⟶{b, d, c, b}
Take[{a, b, c, d, e}, 2]⟶{a, b}
Drop[{a, b, c, d, e}, -2]⟶{a, b, c}
Rest[{a, b, c, d}]⟶{b, c, d}
ReplacePart[{a, b, c, d}, x, 3]⟶{a, b, x, d}
Length[{a, b, c}]⟶3
Range[5]⟶{1, 2, 3, 4, 5}
Table[f[i], {i, 4}]⟶{f[1], f[2], f[3], f[4]}
```

```
Table[f[i, j], {i, 2}, {j, 3}] ⟶
  {{f[1, 1], f[1, 2], f[1, 3]}, {f[2, 1], f[2, 2], f[2, 3]}}
Array[f, {2, 2}] ⟶ {{f[1, 1], f[1, 2]}, {f[2, 1], f[2, 2]}}
Flatten[{{a, b}, {c}, {d, e}}] ⟶ {a, b, c, d, e}
Flatten[{{a, {b, c}}, {{d}, e}}, 1] ⟶ {a, {b, c}, {d}, e}
Partition[{a, b, c, d}, 2, 1] ⟶ {{a, b}, {b, c}, {c, d}}
Split[{a, a, a, b, b, a, a}] ⟶ {{a, a, a}, {b, b}, {a, a}}
ListConvolve[{a, b}, {1, 2, 3, 4, 5}] ⟶
  {2 a + b, 3 a + 2 b, 4 a + 3 b, 5 a + 4 b}
Position[{a, b, c, a, a}, a] ⟶ {{1}, {4}, {5}}
RotateLeft[{a, b, c, d, e}, 2] ⟶ {c, d, e, a, b}
Join[{a, b, c}, {d, b}] ⟶ {a, b, c, d, b}
Union[{a, a, c, b, b}] ⟶ {a, b, c}
```

- Transformation rules

```
{a, b, c, d} /. b → p ⟶ {a, p, c, d}
{f[a], f[b], f[c]} /. f[a] → p ⟶ {p, f[b], f[c]}
{f[a], f[b], f[c]} /. f[x_] → p[x] ⟶ {p[a], p[b], p[c]}
{f[1], f[b], f[2]} /. f[x_Integer] → p[x] ⟶ {p[1], f[b], p[2]}
{f[1, 2], f[3], f[4, 5]} /. f[x_, y_] → x + y ⟶ {3, f[3], 9}
{f[1], g[2], f[2], g[3]} /. f[1] | g[_] → p ⟶ {p, p, f[2], p}
```

- Numerical functions

```
Quotient[207, 10] ⟶ 20
Mod[207, 10] ⟶ 7
Floor[1.45] ⟶ 1
Ceiling[1.45] ⟶ 2
IntegerDigits[13, 2] ⟶ {1, 1, 0, 1}
  IntegerDigits[13, 2, 6] ⟶ {0, 0, 1, 1, 0, 1}
DigitCount[13, 2, 1] ⟶ 3
FromDigits[{1, 1, 0, 1}, 2] ⟶ 13
```

The *Mathematica* programs in these notes are formatted in *Mathematica StandardForm*. The following table specifies how to enter these programs in *Mathematica InputForm*, using only ordinary keyboard characters:

<table>
<tr><td>π</td><td>Pi</td><td>∞</td><td>Infinity</td><td>ⅇ</td><td>E</td><td>ⅈ</td><td>I</td></tr>
<tr><td>x°</td><td>x Degree</td><td>x^y</td><td>x ^ y</td><td>√x</td><td>Sqrt[x]</td><td>x → y</td><td>x -> y</td></tr>
<tr><td>x ≠ y</td><td>x != y</td><td>x ≤ y</td><td>x <= y</td><td>∂_x y</td><td>D[y, x]</td><td>¬ x</td><td>! x</td></tr>
<tr><td>x ∧ y</td><td>x && y</td><td>x ∨ y</td><td>x || y</td><td>x ⊻ y</td><td>Xor[x, y]</td><td>x ⊼ y</td><td>Nand[x, y]</td></tr>
</table>

■ **About the programs.** Like other aspects of the exposition in this book, I have gone to considerable effort to make the programs in these notes as clear and concise as possible. And I believe the final programs will be useful both to execute, and to read and study—if necessary without a computer. Most of the programs involve only built-in *Mathematica* functions, and so can be run in *Mathematica* without setting up any further definitions. (Many programs nevertheless contain variables that need to be assigned their values before the programs are run—as can be done for example with *Block[{k = 2}, program]*. When subsidiary functions are used, these functions also typically need to be defined before the programs are run—even though in these notes I often show the necessary definitions after the programs. Note that most of the programs do not explicitly do input checking or error generation. Only occasionally do the programs significantly sacrifice efficiency for elegance.) A good first step in understanding any program is to run it on a few inputs. The symbolic character of the *Mathematica* language also allows programs to be taken apart, so that their pieces can be run and analyzed separately. Careful study of the various programs in these notes should provide good background not only for implementing what I discuss in the book, but also for doing high-level programming of any kind. Many of the programs use several of the programming paradigms available in *Mathematica*—making it essentially impossible to capture their essence in any lower-level language. Note that a given program can essentially always be written in *Mathematica* in many different ways—though often other ways end up being vastly longer than the ones presented here. Material about the programs should be available at the book website—including for example some of the automated tests run to check the programs, as well as annotations about how the programs work.

■ **Computer experiments.** Essentially all the computer experiments for this book were done using *Mathematica* running on a standard workstation-class computer, and later PC (initially on a 33 MHz NeXTstation, then on a 100 MHz HP 700 running NeXTSTEP, then on a 200 MHz P6 PC running Windows 95, and finally on 450 MHz, 700 MHz and faster PCs running Windows 95, and later Windows NT—with a Linux fileserver). For some larger searches earlier in the project, I wrote special-purpose C programs connected to *Mathematica* via *MathLink*. (Increasing computer speed and greater efficiency in successive versions of *Mathematica* have gradually almost eliminated my use of C.) In some cases I have run programs for many days or weeks, sometimes distributed via *MathLink* across a few hundred computers in my company's network. So far in my life the primary computer hardware systems I have used have been: Elliott 903 (1973–6); IBM 370 (1976–8); CDC 7600 (1978–9); VAX 11/780 (1980–2); Sun-1, 2, Ridge 32 (1982–4); CM-1 (1985); Sun-3 (1985–8); SPARC (1988–91); NeXT (1991–4); HP 700 (1995–6); PC (1996–). The primary languages have been: assembler (1973–6); FORTRAN (1976–9); C (1979–~1994); SMP (1980–6); *Mathematica* (1987–). (See also page 899.)

■ **Educational issues.** The new kind of science in this book represents a unique educational opportunity. For it touches an immense range of important and compelling everyday phenomena and issues in science, yet to understand its key ideas requires no prior scientific or technical education. So this means that it is potentially realistic to use as the basis for an overall introduction to the ideas of science. And indeed having understood its basic elements, it becomes vastly easier to understand many aspects of traditional science, and to see how they fit into the whole framework of knowledge.

No doubt there will at first be a tendency to follow the progression of scientific history and to present the ideas of this book only at an advanced stage in the educational process, after teaching many aspects of traditional science. But it is fairly clear that it is vastly easier to explain much of what is in this book than to explain many ideas in traditional science. For among other things the new kind of science in this book does not rely on elaborate abstract concepts from traditional mathematics; instead it is based mostly just on pictures, and on ideas that have become increasingly familiar from practical use of computers. And in fact, in my experience, with good presentation, surprisingly young children are able to grasp many key ideas in this book—even if their knowledge of mathematics does not go beyond the simplest operations on numbers.

Over the past fifty or so years traditional mathematics has become a core part of education. And while its more elementary aspects are certainly crucial for everyday modern life, beyond basic algebra its central place in education must presumably be justified more on the basis of promoting overall patterns of thinking than in supplying specific factual knowledge of everyday relevance. But in fact I believe that the basic aspects of the new kind of science in this book in many ways provide more suitable material for general education than traditional mathematics. They involve some of the same kinds of precise thinking, but do not rely on abstract concepts that are potentially very difficult to communicate. And insofar as they involve the development of technical expertise, it is in the direction of computing—which is vastly more relevant to modern life than advanced mathematics.

The new kind of science in this book connects in all sorts of ways with mathematics and the existing sciences—and it can be used at an educational level to place some of the fundamental ideas in these areas in a clearer context. In computer science it can also be used as a rich source of basic examples—much as physics is used as a source of basic examples in traditional mathematics education.

A remarkable feature of the new kind of science in this book is that it makes genuine research accessible to people with almost no specific technical knowledge. For it is almost certain that experiments on, say, some specific cellular automaton whose rule has been picked at random from a large set will never have been done before. To conclude anything interesting from such experiments nevertheless requires certain scientific methodology and judgement—but from an educational point of view this represents a uniquely accessible environment in which to develop such skills.

In many fields, advanced education seems useful only if one intends to pursue those specific fields. But a few fields such as physics are notable for being sources of individuals with broadly applicable skills. I believe that the new kind of science in this book will in time serve a similar role.

■ **Reading this book.** This is a long book densely packed with ideas and results, and to read all of it carefully is a major undertaking. The first section of Chapter 1 provides a basic—though compressed—overview of some of key ideas. Chapter 2 describes some of the basic results that led me to develop the new kind of science in the book. Every subsequent chapter in one way or another builds on earlier ones. Some people will probably find the sweeping conclusions of the final chapter of the book the most interesting; others will probably be more interested in specific results and applications in earlier chapters.

These notes are never necessary for the basic flow of any of the arguments I make in the book—though they often provide context and important supporting information, as well as considerable amounts of new primary material. Specialists in particular fields should be sure to read the notes that relate to their fields before they draw any final conclusions about what I have to say.

I have written this book with considerable care, and I believe that to those seriously interested in its contents, it will repay careful and repeated reading. Note that in the main text I have tried to emphasize important points by various kinds of stylistic devices. But in packing as much as possible into these notes I have often been unable to do this. And in general these notes have a high enough information density that it will be rare that everything they say can readily be assimilated in just one reading, even if it is quite careful.

■ **Learning the new kind of science.** There will, I hope, be many who want to learn about what is in this book, whether out of general intellectual interest, to apply it in some way or another, or to participate in its further development. But regardless of the purpose, the best first step will certainly be to read as much of this book as possible with care. In time there will doubtless also be all sorts of additional material and educational options available. But ultimately the key to a real understanding is to experience ideas for oneself. And for the new kind of science in this book this is in a sense unprecedentedly easy, for all it requires is a standard computer on which to do computer experiments.

At first the best thing is probably just to repeat some of the experiments I describe in this book—using the software and resources described at the website, or perhaps just by typing in some of the programs in these notes. And even if one can already see the result of an experiment in a picture in this book, it has been my consistent observation that one internalizes results of experiments much better if one gets them by running a program oneself than if one just sees them printed in a book. To get a deeper understanding, however, one invariably needs to try formulating experiments for oneself. One might wonder, for example, what would happen if some particular system were run for more steps than I show in this book. Would the system go on doing what one sees in the book, or might it start doing something quite different? With the appropriate setup, one can immediately run a program to find out. Often one will have some kind of guess about what the answer should be. At first—if my own experience is a guide—this guess will quite often be wrong. But gradually, after seeing what happens in enough cases, one will begin to develop a correct and robust new intuition. Realistically this seems to take several months even for the most talented and open-minded people. But as the new intuition matures, ideas in this book like the Principle of Computational Equivalence, that may at first seem hard to believe, will slowly come to seem almost obvious.

For someone to assimilate all of the new kind of science I describe in this book will take a very significant time. Indeed, in a traditional educational setting I expect that it will require an investment of years comparable to learning an area like physics. How long it will take a given individual to get to the point of being able to do something specific with the new kind of science in this book will depend greatly on their background and particular goals. But in almost any case a crucial practical step—if it has not already been taken—will be to learn well *Mathematica* and the language it embodies. For although most simple programs can be implemented in almost any computational environment, not using the capabilities of *Mathematica* will be an immediate handicap—which, for example, would certainly have prevented me from discovering the vast majority of what is now in this book.

■ **Developing the new kind of science.** Up to this point in its history the science in this book has essentially been just my personal project. But now that the book is out, all sorts of other people can begin to participate—adding their own personal achievements to the development of the intellectual structure that I have built in this book.

The first obvious but crucial thing to do is to explain and interpret what is already in the book. For although this is a long book that I have tried to write as clearly as possible, there is immensely more that can and should be said—in many different ways—about almost all the ideas and results it contains. Sometimes a more technical presentation may be useful; sometimes a less technical one. Sometimes it will be helpful to make more connections to some existing area of thought or scholarship. And sometimes particular ideas and results in this book will just benefit from the emphasis of having a whole paper or book or website devoted to them.

One of my goals in this book has been to answer the most obvious questions about each of the subjects I address. And at this I believe I have been moderately successful. But the science I have developed in this book opens up an area so vast that the twenty years I have spent investigating it have allowed me to explore only tiny parts. And indeed from almost every page of this book there are all sorts of new questions that emerge. In fact, even about systems that I have studied as extensively as cellular automata I am always amazed at just how easy it is to identify worthwhile questions that have not yet been addressed. And in general the ideas and methods of this book seem to yield an unending stream of important questions of a remarkable range of different kinds.

On the website associated with this book I plan to maintain a list of questions that I believe are of particular interest. The questions will be of many kinds and at many levels. Some it will be possible to address just by fairly straightforward but organized computer experimentation, while others will benefit from varying levels of technical skill and knowledge from existing areas of science, mathematics or elsewhere.

Like any serious intellectual pursuit, doing well the new kind of science in this book is not easy. In writing the book I have put great effort into explaining things in straightforward ways. But the fact that in some particular case I may have succeeded does not mean that the underlying science was easy. And in fact my uniform experience has been that to make progress in the kind of science I describe in this book requires at a raw intellectual level at least as much as any traditional area of science. The kind of extensive detailed technical knowledge that characterizes most traditional areas of science is usually not needed—though it can be helpful. But if anything, greater clarity and organization of thought is needed than in areas where there is existing technical formalism to fall back on. At a practical level the most important basic skill is probably *Mathematica* programming. For it is crucial to be able to try out new ideas and experiments quickly—and in my experience it is also important to have the discipline of formulating things in the precise language of *Mathematica*.

One feature of this book is that it covers a broad area and comes to very broad conclusions. But to get to the point of being able to do this has taken me twenty years of gradually building up from specific detailed results and ideas. And I have no doubt that in the future essentially all significant contributions will also be made by building on foundations of specific detailed facts. And indeed, what I expect to be the mainstay of the science that develops from this book is the gradual accumulation of more and more knowledge of a variety of detailed concrete kinds.

I have tried in this book to lead by example in defining the way I believe things should be done. Probably the single most important principle that I have followed is just to try to keep everything as simple as possible. Study the simplest systems. Ask the most obvious questions. Search for the most straightforward explanations. For among other things, this is ultimately how the most useful and powerful results are obtained. Not that it is easy to do this. For while in the end it may be possible to get to something simple and elegant, it often takes huge intellectual effort to see just how this can be done. And without great tenacity there is a tremendous tendency to stop before one has gone far enough.

In most existing fields of science there are so many technicalities to learn and keep current on that it is rare for anyone but a professional scientist to be able to make any significant contribution. But in the new kind of science that I describe in this book I believe that at least at first there will be opportunities for a much broader range of people to make contributions. In existing fields of science their largely closed communities tend to maintain standards of quality mostly through direct institutional and personal contact. Yet particularly when there are technical aspects to a field it is also comparatively easy for practitioners to assess a piece of work just from the overall way it handles and presents its technicalities. And in fact there are obvious analogs of this in the new kind of science that I describe in this book. First, there is the issue of whether tools like computers are used in effective ways. But in many ways more central is whether there is a certain basic level of clarity and simplicity to a piece of work. Often it is difficult to achieve this. But the point is that the skills necessary to do so correspond rather directly to the ones necessary to carry out the actual science itself well.

▪ **Applications.** At the core of this book is a body of ideas and results that define a new kind of basic science. And I have no doubt that in time this will yield a remarkably broad range of applications. And sometimes—particularly in technology—these applications may be quite straightforward and direct. But if the objective is to develop a model for some specific system in nature or elsewhere it is almost inevitable that this will not be easy. For while I believe that the basic science that I develop in this book provides a remarkably powerful new framework, coming up with an actual model requires all sorts of detailed work and analysis. Certainly it would be wonderful if one could just take the ideas and results in this book and somehow immediately use them to create models for all sorts of systems. And indeed—particularly from the examples I give in Chapter 8—there will probably be at least a few cases where this can be done. But most of the time nothing like it will be possible. And instead—just as in any other framework—there will be no choice but first to learn all sorts of details of a system, and then to use judgement and creativity to see which of them are really essential to a model and which are not. (See also page 364.)