

EMERGENCE

John H. Holland

ABSTRACT

Emergence is a pervasive phenomenon - found in contexts as different as games, seeds, and scientific models - but it has been little studied scientifically. An initial examination makes it clear that scientific models, particularly computer-based models, are essential to a deeper understanding of emergence. As a preamble to a closer look at the process of modeling emergent phenomena, the paper discusses different uses of scientific models. Then it looks at the critical role of building blocks (mechanisms, agents) in constructing such models. Many of the mysteries attending emergence can be reduced to the study of nonlinear interactions between building blocks. Particular emphasis is placed on the creative aspects of such reduction. The paper concludes by outlining the basic elements of models that increase our understanding of emergence.

1. *Emergence - prologue.*

Emergence appears in contexts that range from board games and seeds to the scientific theories embodied in Newton's laws of gravity and Maxwell's equations. Board games provide the easiest example: In the case of chess, agreement on fewer than two dozen rules, provides a game in which new possibilities are regularly discovered after centuries of intensive study. Seeds are much more complicated than board games, but they are the very embodiment of emergence: Somehow these small capsules enclose specifications for structures as complicated and distinctive as giant redwoods, orchids, and lilacs. Nowadays, we know that genes in the seed specify a step-by-step unfolding of biochemical interactions, but only fragments of this complex process are clearly understood. Indeed, we will not truly understand genes and chromosomes until we understand the gene-specified interactions that take a seed, or a fertilized egg, to a

mature organism. In short, we will not understand life and living organisms until we understand emergence.

Newton's laws of gravity, or Maxwell's equations describing electromagnetic phenomena, provide still different examples of emergence. The "laws" so described have much in common with the rules of a game in which "moves" are made with the help of mathematical tools. These moves take us to new equations and mathematical statements that are consequences of the defining equations. As in the case of games, we uncover possibilities quite unsuspected by the authors. Newton could not suspect that his equations would reveal the gravity-assisted boost that takes space probes to the outer planets, and Maxwell, for all his insight, could not anticipate that his equations would make possible the exquisite control of electrons that is the sine qua non of electronic devices. Much of our understanding of the physical world emerges from a small corpus of fundamental equations built on the foundations laid by Newton and Maxwell.

Emergent phenomena are still more common than these scenarios would suggest. Emergence is a common feature of complex adaptive systems (cas) - ant colonies, networks of neurons, the immune system, the Internet, and the global economy, to name a few - where the behavior of the whole is much more complex than the behavior of its parts. Many deep questions about the human condition depend upon understanding the emergent properties of complex adaptive systems: How do living systems emerge from the laws of physics and chemistry? Can we explain consciousness as an emergent property of the central nervous system? Are there economic systems that both encourage innovation and assure a reasonable distribution of goods? We will not know the limitations of scientific answers to questions like these until we understand the whys and wherefores of emergent phenomena.

2. Barriers to the study of emergence.

Emergence, despite its ubiquity and importance, is an enigmatic, recondite topic, more wondered at than analyzed. The hallmark of emergence - "much coming from little" - gives it a paradoxical, almost fraudulent, character smacking of "get rich quick" schemes. There are also philosophers, and some scientists, who take emergence seriously but think that

it *cannot* be explained in scientific terms. Scholars of this persuasion hold that emergent phenomena are holistic phenomena irreducible to the interactions of well defined mechanisms. Specifically, this view holds that a machine cannot generate extensions and improvements unless they are explicitly designed into the machine at the time of its construction.

This stance is similar to a stance widely held until the middle of the 20th century: Machines cannot reproduce themselves. The reasoning was based on the idea that a machine, to reproduce itself, would need a description of itself. But then, that description would have to include a description of the description, and so on, ad infinitum. "Clearly" this is an impossibility, not so very different from "getting more out than you put in". Because living organisms obviously reproduce themselves, this "impossibility" was taken as a major distinction between machines and living organisms. This stance on self-reproduction collapsed in the 1950's when John von Neumann, working with an idea provided by Stan Ulam, provided a description of a self-reproducing machine (von Neumann, 1966).

In the case of emergence, the compact definitions of games and physical laws, with their ever-expanding consequences, seem to belie the view that emergence cannot be described scientifically. Indeed, I think the barriers to developing a mechanical explanation of self-generated enhancement, and emergence, are not ones of principle. The difficulty, it seems to me, stems more from the daunting diversity of emergent phenomena. Like consciousness, life, or energy, emergence is ever-present, but protean in form. In part, too, the difficulty stems from the similarities between emergent phenomena and serendipitous novelty. The play of light on waves produces an ever-changing scintillation, but there is little of the organization we would expect of emergence in a rule-governed system. The false trails of serendipitous novelty, alongside the widely different examples of emergence, make it hard to isolate the elements of emergence.

There is another aspect of emergence that can divert investigation onto a false trail: It is tempting to take the inability to anticipate - surprise - as a critical aspect of emergence. It is true that surprise, occasioned by the antics of a rule-based system, is often a useful psychological guide, directing attention to emergent phenomena. However, I do not look upon surprise as an essential element in staking out the territory. In short, I do not think emergence is an "eye of the beholder" phenomenon

that goes away once it is understood.

Our current understanding of emergence, so far as it goes, comes to us mostly through a catalog of instances, augmented in some cases by rules-of-thumb such as “place the seed in damp soil” or “get your major pieces in action”. In many cases our understanding of emergence is often little better than the child’s invocation of Jack Frost to explain the wondrous colors of autumn. Such an explanation stirs the imagination, but it is ultimately unsatisfying. The scientist’s instinct is to start looking for a deeper explanation, an explanation that may go as far as the molecular biologist’s contemplation of the tangled bio-molecular interactions that produce autumn changes. The deeper explanation, once understood, inevitably gives imagination an exhilarating boost. But just what should we look for in trying to understand emergence?

3. *A scientific approach to emergence.*

It is unlikely that a topic as complicated as emergence will submit meekly to a concise definition, and I have no such definition to offer. I can, however, provide some markers that stake out the territory, along with some requirements for studying the terrain.

In what follows, I’ll restrict the discussion to systems that can be defined with rules or laws. Games, systems made up of well-understood components (e.g. molecules composed of atoms), and systems defined by scientific theories (e.g. Newton’s theory of gravity) are prime examples. Emergent phenomena also occur in domains for which we presently have few agreed upon rules: ethical systems, the evolution of nations, and the spread of ideas come to mind. Most of the ideas developed here have relevance for such systems, but precise application will depend upon better conjectures about the laws (if any) that govern the development of such systems. There may also be other valid scientific uses for the term “emergence”, but the rule-governed domain is rich enough to keep us fully occupied.

The first step in staking out the territory is simply noting, again, that small numbers of rules or laws can generate systems of surprising complexity. Moreover, this complexity is not just the complexity of random patterns. There are recognizable features, as in a pointillist painting. In addition, the systems are animated - *dynamic*. Though the laws are in-

variant, the things they govern are changing. The changing patterns of the pieces in a board game, or the trajectories of baseballs, planets, and galaxies under Newton's laws, show the way. The rules or laws generate the complexity, and the ever-changing flux of patterns that follows leads to perpetual novelty and emergence.

Recognizable, repeating features or patterns are pivotal in understanding the dynamics of these systems. I'll call a recognizable, repeating phenomenon *regular*, and I'll not call a phenomenon *emergent* unless it is *regular*. That a phenomenon is regular does *not* mean that it is easy to recognize or explain. The task can be difficult even when the laws underpinning the dynamics are known. In chess it took centuries of study to recognize certain patterns of play, such as the control of pawn formations, yet these patterns greatly enhance the possibility of winning the game. Similarly, it took centuries of study to extract some of the regular dynamic patterns inherent in Newton's laws, such as the gravitational boosts used in planetary exploration, and still we learn.

Given the lack of an over-arching definition, along with the complexity and subtleness of emergent patterns, how do we approach the problem scientifically? At present, with some notable exceptions, we are still collecting examples of emergent phenomena, much like a butterfly collector. Collecting is valuable, but to develop a general understanding we must discard the idiosyncratic features of particular examples. If we can extract core features, then we can go on to meld those features into a general setting that guides our exploration. These considerations lead us directly to model-building.

It may not be obvious at first, but the study of emergence and model-building go hand in hand. The essence of model-building is shearing away detail to get at essential elements. A model, by concentrating on selected aspects of the world, makes possible the prediction and planning that reveal new possibilities. That is exactly the problem we face in trying to develop a scientific understanding of emergence.

Models and model-building are more than a scientific craft. In fact the word "model" has been used with broader connotations from the outset:

"When we meane to build, We first survey the Plot, then draw the Modell." [Shakespeare]

In this broader usage models include such things as maps, architectural diagrams, scale models, games, flight simulators, mathematical models, cartoons, and mental strategies and even metaphors.

Among living forms on earth, the construction of objects and scripts that serve as models is a uniquely human activity. The models may be small - the early Egyptians produced exquisite miniatures of animals and boats - or they may be large - that huge immobile arrangement of monoliths, Stonehenge, models the passage of seasons. The process of model building has an element of mystery, often displaying emergence in a literal way. It is more than coincidence that early modeling efforts, such as the Stonehenge and the Egyptian boats of passage, were under the control of a priesthood. From earliest times, human endeavor has been directed toward discovering ways to channel a chaotic world through rules and models. This starts with rule-bound sacrifices to the gods - we model the world in terms of personalities and ways of propitiating those personalities. Later, we discover mechanisms and ways of using them to control parts of the world (e.g., gates, pumps, and wheels), and we begin to model the world with mechanisms instead of personalities. Eventually, we come to such things as complex computer-controlled devices and scientific models that employ abstract mechanisms such as quarks and gluons.

At another level, models are such an automatic feature of day-to-day existence that we rarely stop to think how ubiquitous, various, and important they are. Driving home from work is model-directed - we have a kind of internal map of the major landmarks and turning points along the way. We are typically unaware of this map, until we have to search for an alternate route because of construction or traffic. Similarly, when we encounter an unfamiliar scene, we automatically parse it into something recognizable by constructing a model on the fly. We use familiar building blocks - trees, buildings, automobiles, other humans, specific animals, and so on - to build a model that lets us anticipate the dynamics of the scene. Such everyday models give us the advantage of executing virtual (usually unconscious) experiments, greatly reducing the need for overt, time-consuming, possibly dangerous, actions.

For most of us model-building starts at an early age. As children we use building blocks to generate concrete realizations of our imagination - castles and space stations. This facility for recombining standard objects to make new things carries over into later occupations. A watchmaker

uses familiar mechanisms - gear wheels, springs, pinions, and so on - to generate marvels of timekeeping, and a scientist does the same thing at a more abstract level, generating complex objects, e.g. molecules, from simpler objects, e.g. atoms. By selecting building blocks and the ways of recombining them, we set up the rules that make rule-governed systems comprehensible. Model planes grow into the models used in wind tunnels to determine flight characteristics. Still later, we used computer-based models of planes to test their performance envelope under dynamic conditions, both normal and abnormal (such as how a 747 performs with two engines out). A well-conceived model exhibits the complexity, and emergent phenomena, of the system being modeled, without the obscuring effects of incidentals.

In a sense, all of science is based on model construction. But, in this role, models need bear no obvious resemblance to the thing be modeled. Newton's equations, as symbols confined to a sheet of paper, bear no resemblance to the orbits of planets around the sun. And Maxwell's equations bear no relation to the patterns of iron filings that inspired them. Yet they model this physical reality in ways that no scale model could ever achieve - think of all the manifestations of what we now call "gravity" and "electromagnetism". The unanticipated predictions and marvels tied up in these equations provide some of our best examples of emergence. A great deal more comes out than the authors anticipated, even allowing for their superb intuition. To understand emergence, we must understand the way in which models in science, and elsewhere, allow us to transcend the knowledge that went into their construction.

4. Board games, number, and maps - precursors of scientific models.

Despite the pervasive use of models in the sciences and elsewhere, the art of model-building is not a familiar topic, even to many practicing scientists. Fortunately, scientific models rest upon cornerstones that have long been a familiar part of human culture: board games, numbers, and maps.

Board games.

Board games are a singular human construct, already a common feature of the early Egyptian Dynasties (3000 B.C. and earlier). Board games are typified by pieces arrayed on a partitioned board, with rules that set the

legal ways for placing or moving pieces on that board. It takes only a few rules to define a game as complex as chess or Go. The rules constrain possibilities: not all board configurations are legal, and new configurations follow from legal changes in configurations already achieved. Though the rules do forbid many configurations, the number of legal configurations remains large, and the ways of getting from one configuration to another are intricate.

Board games provide a particularly simple example of the emergence of great complexity from simple rules or laws. Even in traditional 3-by-3 tic-tac-toe the number of distinct legal configurations exceeds 50,000, and the ways of winning are not immediately obvious. The play of 4-by-4-by-4, three-dimensional tic-tac-toe offers surprises enough to challenge an adult. Chess and Go have enough emergent properties that they continue to intrigue humans and offer new discoveries after centuries of study. And it is not just the sheer number of possibilities. There are lines of play and regularities that continue to emerge after years of study, enough so that a master of this century would handily beat a master of the previous century.

As we will see, the rules of a board game hold much in common with the rules of logic. And, from there, it is not a long distance to the axiomatic and equation-based models of science. Much of our modern outlook is conditioned on the discoveries that emerge from this way of looking at the world, from atoms and genes to superconductivity and antibiotics. Mathematical models provide an unusually penetrating way of discovering unexpected aspects of our world. That a modeling technique as abstract as mathematics should be so efficacious is a mystery often remarked by scientists, but it is less a mystery when we put it in this context of games and rules.

Number

The other ancient cornerstone for model-building is the concept of number, the foundation of mathematics. Number may seem to be the very embodiment of concreteness. After all what could be more concrete than saying, "There are three busses in the parking lot", or "I have two children". However, it is another of those concepts that is at once familiar and mysterious. A careful look at number starts with abstraction - shearing away detail.

Numbers go about as far as we can go in shearing away detail. When

we talk of numbers, there's nothing left of shape, or color, or mass, or anything else that identifies an object, except the very fact of its existence. Another way to say the same thing is to say that, when we are talking about number, all collections that have the same number of objects, say three, are to be treated as equivalent. Three busses, three storks, and three mountains are equivalent "realizations" of the number three.

Shearing away detail is the very essence of model building. Whatever else we require, a model must be simpler than the thing modeled. In certain kinds of fiction, a model that is identical with the thing modeled provides an interesting device, but it never happens in reality. Even with virtual reality, which may come close to this literary identity one day, the underlying model obeys laws which have a compact description in the computer - a description that generates the details of the artificial world.

As we move beyond number, we can of course change the details sheared away. The color "red" treats as equivalent all collections of objects that have that color. Similarly, we throw away masses of detail when we invent concepts such as "trees", "grandmothers", and "airplanes". An individual tree, for instance, has a plethora of detail about leaf shape, placement of branches, and so on, and trees of different species can be quite different in most of their details - compare an oak to a pine. Still, there are certain things held in common by all scenes containing trees, and it is this common part that enables us to build up the "tree" classification. The same holds true for something as specific and unique as "my friend, Alice", where details of dress, hairstyle, etc., are set aside in order to recognize the person. By ignoring selected details we obtain "building blocks" - regular phenomena - that appear repeatedly in our experience of the world.

Maps

We can go a step further toward an understanding of models by considering maps. Maps eliminate detail in a straightforward way and, like games, they are among the earliest model-artifacts. Think first of a simple roadmap. If it is fairly complete, as in the case of most roadmaps, then cities, towns, and villages are represented by dots or squares of varying size, and the roads connecting these population centers are represented by lines of various colors representing road quality. There may be some lakes and rivers indicated, but in general the map concentrates on

population centers and roads. There are two kinds of relations preserved: (1) There is a one-to-one relation between the population centers and the "dots" on the map. Each city, town, and village is represented by a dot on the map.

(2) The dots are arranged on the map in the same configuration as population centers have in the actual geography of the state. That is, larger cities that are close together in the state are represented by large dots that are close together on the map, a town that is close to the state boundary is represented by a smaller dot close to the edge of the map, and so on. However, all distances have been scaled down, so that cities that are 20 miles apart in reality are separated by 2 inches on the map. The curves, straightaways, and intersections of the roads are represented on the same scale.

A moment's thought shows that a map retains *few* details. We learn little about what we will see by the roadside in driving down one of the roads, nor even much about minor zigs and zags in the road (those changes in direction too small to show up at the scale of the map), let alone any details about what the towns look like. What *is* retained is just the essential information about getting from one place to another *under normal circumstances*. Road construction or windstorms can make the route suggested by the map infeasible or impossible.

The map *does* provide a scaled correspondence between salient points in the world and points on the paper. Scale also asserts itself when we extend our view beyond maps to other kinds of model. We at once encounter a whole class of models called *scale models*: scale ships, scale railroads, scale planes, etc. We also expect scale in most statues and representational sculpture, though a monument like Mount Rushmore may be scaled to be *larger* than the original. However, if we look still further afield, we encounter models in which scaling plays little or no role.

Scaling is a special case of a deeper concept, *correspondence*. We automatically get correspondence when we produce a scaled model, but correspondence is possible without scaling. To construct a model using correspondence, we first select the details or features to be represented, then construct the model so that some part of the model *corresponds* to each selected detail. Think of a cake recipe. It models the steps we actually use to produce a cake. Each step in the recipe, e.g. "add a cup of sugar", corresponds to a complex activity involving a series of physical movements and measurements. Computer-based models have much in

common with recipes.

5. *Computer-based models.*

Computer-based models greatly enhance our possibilities for understanding emergence by providing accessible, controllable instances of the phenomenon. A computer-based model can be started, stopped, examined, and restarted under new conditions, in ways impossible for most real dynamic systems (e.g., an ecosystem or an economy). In examining computer-based models, we also come back to the element of surprise as a clue that suggests emergence: It is commonplace for a scientist-programmer to provide the computer with a program that is fully capable of surprising its designer. Though the program is fully reducible to the rules (instructions) that define it, so that nothing remains hidden, the behaviors generated are not easily anticipated from an inspection of those rules. Indeed that is the very purpose of the model: to explore the consequences of its assumptions.

A model defined by a computer program, as mentioned, is like a recipe, and the computer is like an automated stove: Once the recipe is inserted, the delicacy described emerges. The computer automatically reveals the behavior implicit in the model's defining program. In this the computer-based model differs from the more familiar mathematical models defined by equations. It may take years of sophisticated mathematical analysis to reveal the consequences bound up in the defining equations.

Computer-based models present the modeler with a rigorous challenge. The claims of verbally described models are often established by rhetoric. What appear to be equally good arguments often back competing claims for the same model - consider claims about global warming or species preservation. The same can sometimes be said for traditional mathematical models, where even the most rigorous mathematical proofs skip "obvious" steps. There is no skipping of steps in a computer program. The computer executes each and every instruction in the sequence given. A missing or incorrect instruction will send the program careening away from the modeler's intent. In this, a computer-based model is much like the *working* mechanical model the U.S. Patent Department required in an earlier age. No matter how clever and convincing the descriptions, if the working mechanical model didn't produce the results claimed, the

patent was not allowed. Similarly, a computer-based model is both *rigorously described* - it is presented as a program that can be examined in detail - and it is *executable*.

Computer-based models are at once abstract and concrete. They are abstractly defined in terms of numbers, relations between numbers, and changes in numbers over time - a feature they share with mathematical models. At the same time, the numbers are actually "written down" in the computer's registers, rather than being represented symbolically. Moreover, the numbers are overtly manipulated by the computer's instructions, much as a grain mill produces flour. We can produce quite concrete records of these manipulations. These records are closely related to the laboratory notebook records of a carefully run experiment. Computer-based models, then, partake of features of both theory and experiment. As we'll see, this combination of the abstract and concrete offers both advantages and disadvantages.

It is, at first, surprising that a wide range of concrete objects and processes can be represented in computers by numbers and the manipulation of numbers. Both computer-based models and mathematical models share this rather mysterious ability. How do we use numbers to simulate the flight of an airplane over Chicago in a summer thunderstorm? Such numerical representations have become so common that you can run flight simulations on your home computer and, with a bit more effort, we get full-fledged industrial flight simulators that wring sweat from experienced pilots when they "fly" in simulated emergency situations. How can this be?

The starting point is a basic concept in the study of dynamic systems, the concept of *state*. The natural question is, "What can we possibly mean by the state of a jet airplane flying over Chicago?" The answer to this question is closely connected with the information the pilot uses to fly the jet.

To get at this connection between information and state, let me start with a simpler system: the control panel of the family car. The car's control panel is not in principle much different from that in the jet, it is just much, much simpler. It tells us only the essentials that we need to know when driving: the speed of the car, the fuel level, the engine temperature, the battery charge, and the oil pressure, are typical. These readings model the state of the car, at a certain level of detail, when it is underway. We could add more readings, such as the air pressure in the

tires, or the amount of antifreeze in the radiator, to get a more detailed state. This more detailed state would provide the wherewithal for a more sophisticated model; however, decades of experience have shown that the gauges first mentioned are sufficient for operating the car in most situations.

A jet in flight is a much more complicated dynamic system than the family car, so the pilot's compartment is filled with a panoply of displays, gauges, dials, and warning lights that provide information about the conditions that affect the jet's flight. They tell about the plane's speed and position, the amount of fuel in its various fuel tanks, the operating condition of the engines, the position of the landing gear, and on through hundreds of other pieces of information. Indeed, there is enough information for the pilot to fly the plane "blind", using instrument readings alone.

For both the car and the jet, the displays and gauges produce readings that either are numbers or are easily reduced to numbers. A warning light can be either "on" or "off", which can be represented as a 1 or 0, and even the sophisticated positional display is presented by an array of dots (called "pixels") which can be represented as an array of 1's and 0's. In other words, it is easy to reduce the information on the control panels to numbers. These numbers can, as usual, be stored in registers in the computer. Together they define the state of the model, much as the arrangement of pieces defines the state of a board game.

We give the computer a representation of the state of the model by entering these numbers into its storage registers. Then, we enter instructions (a program) that cause these numbers to change over time as specified by a *transition function*. This is the counterpart of defining the rules of the game. The numbers in the registers change in a way that mimics the state changes in the object being modeled. The *universality* of the general-purpose computer assures that any transition function defined by a finite number of rules can be so-mimicked.

As in the case of games, we now confront the notion of choice. The driver or the pilot can choose among alternatives, e.g. making the car or jet go faster or slower. Phrased in terms of states this means that, once again, from any state we can construct a tree of legal alternatives. In games, these alternatives are the legal moves allowed by the rules. In the case of the car or the jet, the laws are those imposed by nature and the technology. Executing a sequence of controlling actions is the counterpart

of making a sequences of moves in the game. In both cases, we choose a path through the tree of possibilities.

When both the numbers and the program have been stored in the computer, we simply start the computer executing its instructions. Think again of a video game or flight simulator. The instructions, acting on the stored numbers defining the model's state, determine what happens instant by instant. What we see on the computer screen, is a back-translation of the numbers to gauge-readings, displays, etc., that capture the look and feel of the original machine. Controlling actions amount to input to the program at various stages of the calculation. The input is supplied by typing, or by the video game's joystick, or by realistic controls in a full fledged flight simulator. The result is a dynamic, computer-based model - a major vehicle for the scientific investigation of models and emergence.

6. *The uses of models.*

Even when we restrict ourselves to the sciences, models serve several purposes. There are three broad categories that include most kinds of models. (I warn the reader that different scientists would provide different divisions or characterizations - to my knowledge there is no widely accepted categorization). In each category, there is a characteristic claim for the model, an accompanying validation criterion, and one or more examples. Here's my list:

Predictive Models

Claim: Starting from a limited set of mechanisms and constraints, and an initial state that corresponds to current conditions, the model predicts conditions in the future or under a different regime, at some useful level of detail and reliability .

Validation: Data from experiments confirm the models predictions.

Examples: Weather models; traditional scientific models (e.g. the PVT relation for gases).

Existence Proof Models

Claim: The model provides a rigorous demonstration that some process or phenomenon is possible (e.g., a machine *can* reproduce itself) or impossible (e.g. material bodies *cannot* exceed the speed of light).

Validation: The model, when executed, works as claimed (much like

validating a patent).

Examples: Von Neumann's self-reproducing automaton; the classic gedanken experiments of physics.

Exploratory Models [called "de-mystifying models" by C.F. Stevens, and "models for ideas" by J. Roughgarden]

Claim: The model provides an "explanation" of complex phenomena in terms of a limited set of mechanisms and constraints; the model often suggests "places to look" for salient phenomena, regularities hidden in complex data, etc.

Validation: The model suggests new avenues to scientists familiar with the area.

Examples: Maxwell's demon; Schrödinger's quasi-crystal model of life; Simon's limited rationality model in economics; "lock and key" models in immunology.

Exploratory models can go through a series of stages. They usually start by helping to formulate relevant questions about complex phenomena. As has often been remarked, arriving at the "right" questions is 90% of the scientific effort. As in the construction of metaphors, and other new ways of looking at the world, taste and discipline are critical elements in formulating good exploratory models. With time, an exploratory model may take on aspects of an existence proof or predictive model.

The methods for specifying a scientific model are, almost always, either simultaneous equations or, more recently, computer programs (though gedanken experiments may be more informal, resting on shared axioms). Both methods of specification are equally rigorous, but they have different degrees of generality and different ways of abstracting from observation. With the advent of the programmed digital computer, we can critically examine existence proof and exploratory models (and some predictive models) that are several orders of magnitude more complex than was possible earlier in this century.

7. Building blocks, mechanisms and agents.

Models, and particularly computer-based models, nicely integrate the themes exemplified by games, numbers, and maps. To implement a model on a computer we first determine the model's major components - the model's building blocks. Then we implement these components as

sets of instructions in the computer called subroutines. Finally, the subroutines are combined in the computer in a way that determines their interactions, yielding the overall program that defines the model. The result is a computer-based realization of the transition function (rules) that define the model's behavior.

Building blocks play a ubiquitous role in our understanding of the world. Any human can, with the greatest of ease, parse an unfamiliar scene into familiar objects - trees, buildings, automobiles, other humans, specific animals, and so on. This quick decomposition of complex visual scenes into familiar "building blocks" is something that we *cannot* yet mimic with computers. The task is too complex to be carried out by brute force, despite the computer's tremendous advantage in speed, and we have no plausible computer-based models of human parsing procedures. This lack of an adequate model is closely related to our lack of understanding of the activities of neurons in the central nervous system.

Whatever the parsing process, it *is* clear that we can use small numbers of building blocks to construct, or reconstruct, complex scenes and configuration. If we consider vision, we can see the importance of the generative character of building blocks. The actual projection of external scenes on the millions of sensory cells in our eye is never twice the same; nevertheless, every scene has some aspects that have appeared before. Over the years we get better and better at discerning and classifying these common elements - the building blocks. Moreover, because we see the building blocks over and over again, we gain facility in determining their essence, learning just what details are relevant. The same considerations apply, at a higher level, when we consider the tremendous range of expression provided by stringing together copies of the few thousand building blocks we call words. It is our ability to discern and use building blocks that makes the perpetual novelty of our world understandable, and even predictable.

The process of discovering building blocks goes on throughout one's life, and in science it goes on from generation to generation. Though the number of building blocks in our repertoire may be small relative to the number of configurations in which they appear, we can always acquire more. Part of this is simply refining extant classifications, moving from the general to the more specific. A young child may confuse a cow and a horse, calling both "horsy", while an experienced farmer will distinguish different breeds of cow and will know that Betsy, as an individual

in his herd, gets restive when she is milked. Even an experienced camper will learn new building blocks if he or she takes up animal tracking (the newly turned leaf or the displaced pebble) or cross-country trekking in the arctic (the kinds of snow). Occasionally there is a major addition to the repertoire of building blocks. In most human activities, the discovery of a major new building block causes a "revolution", opening new realms of possibility. Think of, say, "perspective" in the arts, or "gravity" in the sciences.

As time goes on, humans get better and better at knowing what details to discard. We learn what is irrelevant to "handling" or understanding situations, and we refine our building blocks accordingly. We also learn to use rules - sometimes called "laws" when they're used this way - to project the way in which the blocks will shift and recombine as the future unfolds. That is, we build models that help us anticipate the future. We even rerun the projections with variations and modifications to see what the possibilities are, with particular emphasis on not "falling off cliffs". This use of models is particularly obvious in playing sophisticated board games, but it comes into play in everything from the mundane task of finding an alternate route when roadwork blocks the usual route home, to the generation of sophisticated hypotheses in science.

In the sciences, building blocks have had a central role from the outset. The Greeks developed the idea that all machines can be constructed by combining (copies of) six elementary mechanisms (the lever, the screw, the inclined plane, the wedge, the wheel, and the pulley). The idea of explaining the different properties of matter in terms of elementary building blocks called atoms also originated with the Greeks. This idea was progressively refined until we get such things as the periodic table of the elements and the modern conception of atoms in terms of nuclei and orbital electrons.

In 1969, Herbert Simon used the combination of elementary mechanisms to illustrate a key point about the construction (and evolution) of complex systems. He tells a tale of two watchmakers: One watchmaker constructs each watch piece-by-piece using the elementary mechanisms known to the Greeks - levers, wheels, and so on. The other watchmaker works in terms of sub-assemblies constructed from the elementary mechanisms - a mainspring sub-assembly, the gear train for the watch hands, and so on. The sub-assemblies are then combined into more complex assemblies, until finally the watch is formed. If the structures are unstable

until fully assembled (the whole watch in the first case, the sub-assemblies in the second case), then any interruption or untoward event will mean starting over on that particular structure. When interruptions are frequent, the second watchmaker - the one using a hierarchy of building blocks - has a clear advantage. Simon's tale offers substantial insight into the prevalence of hierarchical, building-block structures in the natural world, a world in which untoward events are commonplace.

Indeed, to understand and manipulate complex systems, be they biological cells or computers, we almost always develop hierarchical descriptions with successive levels of building blocks. In a general setting, this means looking at complexity and emergence in terms of *mechanisms* and procedures for *combining* them. To make this work, we have to extend the idea of mechanism beyond the overtly mechanical. We come closer to the physicist's notion of elementary particles as mechanisms for mediating interactions, as when a photon causes an electron to jump from its orbit around an atom. Mechanisms, so described, provide a precise way of describing the elements, rules, and interactions that define complex systems. The resulting descriptions of the diverse rule-governed systems that exhibit emergence gain considerably in uniformity. We can then compare quite different systems. Therein lies our hope of finding similarities and common rules or laws. With diligence, and good fortune, we should be able to extract some of the "laws of emergence".

When we look at complex adaptive systems in this way, we find that many of them are naturally described in terms of agent-based models, where mobile "mechanisms" (agents) interact and adapt to each other. The classic description of agent-based emergence is Douglas Hofstadter's 1979 metaphor of the ant colony: An individual ant (agent) has a limited, reflex-driven repertoire - a large fraction of that repertoire can be modeled with twenty or so rules. The colony of ants, on the other hand, exhibits remarkable flexibility in probing and exploiting its surroundings. It reacts adaptively to disasters (invasions by other ant colonies, downpours, and so on), it searches out and exploits changing food sources, and it persists over many, many worker ant generations. Somehow the simple laws of the agents generate an emergent behavior far beyond the capacities of individual agents. It is noteworthy that this emergent behavior occurs without directives from a central executive.

Emergence in agent-based models usually involves patterns of interaction - regular patterns - that persist despite a continual turnover in the

agents generating the pattern. A simple example of such a pattern is the standing wave in front of a rock in a white-water river. The water molecules making up the wave change, instant by instant, but the wave persists as long as the rock is there and the water flows. Ant colonies, cities, and the human body (which turns over *all* of its constituent atoms in less than two years) provide more complex examples. These persistent patterns can themselves become building blocks for still more complicated persistent patterns. Within such a regime, hierarchical organization is a natural outcome. Emergent macro-patterns that depend upon shifting micro-patterns make emergence fascinating, and difficult to study.

8. *Reduction.*

At this point we encounter a topic of some controversy, though more a controversy among philosophers and post-modern writers than among practicing scientists. Much of our discussion has centered on the construction of new levels of description through the combination and interaction of mechanisms (building blocks). If we turn this discussion on its head, explaining behavior at one level in terms of the interactions of mechanisms at a deeper level (e.g., the description of molecular dynamics in terms of the interaction of atoms), we encounter the concept of *reduction*.

Reduction - the technique of describing complicated systems in terms of *interactions* of simpler systems - is the usual, almost universal, scientific approach to a new area. Indeed, reduction motivates most of the work in basic science. Over the centuries it has produced an interlocking hierarchy of structures that leads from strings and quarks through nucleons, atoms, molecules, molecular biology, and onward. In one sense this hierarchy implies that all phenomena in the universe are ultimately reducible to the laws of physics. However, most scientists would state this a bit more cautiously, saying that all phenomena are *constrained* by the laws of physics. Just what is implied by such a view?

First of all, even if one holds strictly to this view, it does not follow that all explanations should be couched *directly* in terms of the laws of physics. It would be both tedious, and unenlightening, to explain every chemical reaction by using the apparatus and time-scales of quantum mechanics. It is enough to relate various kinds of chemical bonds to

quantum mechanical features, using bonds for the rest of the explanation. Even in a model universe, like chess, where the defining laws are completely known and simple, much that is observed is determined by large-scale phenomena, like cooperative pawn formations. Unless we can formulate macro-laws that deal directly with these large-scale phenomena, it is difficult to catalog possibilities.

As we move up this hierarchy, we see that new levels of description are imposed on the basic description. But these new levels must *not* contradict the constraints imposed by the earlier levels. We add new laws that satisfy the constraints imposed by laws already in place. Equally important, the new laws are consequences of those laws. Moreover, these new laws provide a new level of description of complex phenomena that are consequences of the original laws. We will gain a deeper understanding of emergence, if we can deepen our understanding of this idea of levels of definition.

We can develop a more precise notion of the relation between level and consequence by looking at the axioms of Euclidian geometry. It was long thought (hoped) that Euclid's fifth "axiom of parallel lines" could be proved from the other four axioms. However, in the 19th century, it was shown that one could add a fifth axiom that contradicted Euclid's fifth, while still retaining a consistent axiom system. This discovery led to a whole new range of non-Euclidian geometries, ultimately leading to such things as Einstein's theory of relativity. The point, for present purposes, is that the first four axioms completely constrain what can be achieved by adding additional axioms, but they do not foreclose different options.

Anything that can be accomplished by adding axioms to Euclid's first four axioms, say adding Euclid's fifth or one of the axioms that contradict it, can be accomplished within the system of four axioms alone: In the four-axiom system, we can always prove a set of theorems of the form IF (new axiom) THEN (derivation of theorem based on extant axioms). That is, we treat the new axiom as a conditional assumption, and carry out derivations based on that assumption. The resulting theorems exactly parallel the theorems that can be derived in a five axiom system that incorporates the new axiom.

Note that we could equally well add other axioms that have nothing to do with parallel lines. There is an endless range of assumptions (axioms) that *could* be used. Most of these assumptions would yield theorems

that are uninteresting or trivial vis-a-vis questions about geometry. Still what can be proved or studied with the addition of these new axioms, can equally be studied in the original system without them. The reason for highlighting some assumptions as axioms comes from an entirely different direction. The highlighted axioms define the direction of the study. It required a deep understanding of geometry to formulate an axiom that *both* contradicted Euclid's fifth axiom *and* contributed a set of theorems that enlarged our conception of geometry.

We can look upon macro-laws at higher levels in the hierarchy of scientific laws as axioms added to the original axioms (the basic laws of physics). Typically, the added macro-laws will have premises that pick out a range of situations that occur frequently or involve possibilities that lever the system onto new paths. The overall system is still constrained by the original laws and we could, in principle, derive everything in terms of these original laws, as in the example of Euclid's geometry. But, there are many possible conditions (macro-laws), and the trick is to pick those that offer possibilities not apparent from direct inspection of original laws. Said another way, we must "tune" the constraints supplied by the new laws so that the study concentrates on interesting domains not easily apprehended or explored in the original setting. That's really the reason for highlighting carefully selected assumptions as macro-laws - as with Euclid's axioms, they define the direction of the study.

When we observe regularities (e.g., the usual valence laws of chemical reactions) we carry out operations at that level, replacing what may be difficult or even infeasible calculations from first principles (the laws of quantum mechanics). These regularities still satisfy the constraints of the underlying micro-laws, but they involve additional conditions, usually called "normal" or "natural" conditions. Under these assumptions, the regularities persist and a simpler, "derived" dynamics can be used. When these conditions are absent, we abandon the macro-level, and return to the micro-level for the more detailed considerations then required. Kirchhoff's laws for the conduction of electricity work well under normal conditions, but under low temperature regimes we get the "abnormal" superconductive regime which requires a return to basic quantum mechanical considerations.

In the phrase describing reduction at the outset of this section, I italicized *interaction*. I did this because there is a common misconception about reduction: To understand the whole, you analyze a process into

atomic parts, and then study these parts in isolation. Such analysis works when the whole can be treated as the *sum* of its parts, but it does *not* work when the parts interact in less simple ways. Sums work when we analyze a complex sound wave, say an instant from a symphony, into its component frequencies. We can then reconstruct the whole by adding these components together; some kinds of digital recording depend upon this instant-by-instant ability to recombine component frequencies into a sustained performance. However, when the parts interact in less simple ways, as when ants in a colony encounter each other, knowing the behaviors of the isolated parts (ants) leaves us a long way from understanding the whole (the colony). The simple notion of reduction - studying the parts in isolation - does not work then. We have to study the *interactions* as well as the parts.

Emergence, in the sense used here, only occurs when the activities of the parts do not simply sum up to give the behavior of the whole. That is, emergent phenomena only occur when the whole is indeed more than a sum of its parts. Chess provides a good example: We *cannot* get a good picture of a chess game in progress by simply adding up the values of the pieces on the board. The pieces interact to support one another and to control various parts of the board. This interlocking power structure, when well conceived, can easily overwhelm an opponent with higher valued pieces that are poorly arrayed. A good analysis of the game's setting must provide a direct way of describing these interactions. The same holds, a fortiori, for more sophisticated versions of emergence. A reduction that does not provide for the study of interactions will not be of much help in the study of emergence.

9. *The creative obverse of reduction*

The insights that lead to interesting choices for macro-laws often depend upon a careful use of metaphor and cross-disciplinary comparisons, particularly in the study of emergence. The constraints so imposed play a role similar to the constraints imposed by meter and rhyme when composing poetry. Such constraints are as likely to enhance imagination as to inhibit it.

This creative side of reduction involves what, at first, seems a conundrum. The building blocks of a watch have been familiar since the time

of the Greeks, but the watch is an innovation that has been with us for less than two centuries. Why was the watch so slow to emerge when the building blocks were so familiar? Here we come upon a central point about innovation and the study of emergence: Building a model or developing a theoretical construct in science is *not* a matter of deduction. It's important to distinguish the finished product in science from the process that produces that product.

The finished product in science, usually a published scientific paper or book, is presented with careful, step-by-step reasoning. Each step follows directly and clearly from the previous step, at least for the cognoscenti. The whole presentation strives for inevitability, wherein the conclusions are an irrefutable consequence of the starting point. In practice, this inevitability is an ideal only approximated, but the best scientific publications are quite convincing in this respect. This widely accepted scientific standard gives rise to a view, held by some scholars and a few scientists, that science is actually conducted in this step-by-step, almost mechanical way. Imagination and creation are marginalized. However, few scientists, if any, actually carry out their research in this fashion.

Scientists rarely discuss this metaphor-driven aspect of their work, but James Clerk Maxwell provides a wonderful exception. In his collected papers (Maxwell, 1890) you can read how he used a mental model of floating gear wheels to enhance his intuition about electromagnetic fields.

We must therefore discover some method of investigation which allows the mind at every step to lay hold of a clear physical conception, without being committed to any theory founded on the physical science from which that conception is borrowed, so that it is neither drawn aside from the subject in pursuit of analytical subtleties, nor carried beyond the truth by a favourite hypothesis.

He goes on to give a more specific example.

[Refer] everything to the purely geometrical idea of the motion of an imaginary fluid [which is] merely a collection of imaginary properties which may be employed for establishing certain theorems in pure mathematics in a way more intelligible to many minds and more applicable to physical problems than that in which algebraic symbols alone are used.

Maxwell's other writings make it clear that his "clear physical conception" is exemplified by the mechanism-oriented fluid mechanical model that he used to arrive at his famous equations for electromagnetic fields. Thus, Maxwell moves from a specific mechanical model to the greatest feat of abstraction since Newton formulated his equations for gravitation.

The construction of a mental model of this kind closely resembles the construction of a metaphor:

- (i) There is a source system with an established aura of facts, interpretation and practice.
- (ii) There is a target system with a collection of observed phenomena that are difficult to interpret or explain.
- (iii) There is a translation from source to target that suggests a means of transferring inferences for the source into inferences for the target.

Both models and metaphors enable us to see new connections. For most who are heavily engaged in creative activities, be it in literature or the sciences, metaphor and model lie at the center of their activities. In the sciences, both the source and the target are best characterized as systems rather than isolated objects. Typically, these are systems of interacting (copies of) mechanisms. The mechanisms may be literal, as in Maxwell's use of gears, or they may be figurative, as in the use of quarks and gluons to explain the construction of nucleons. The scouting expedition that determines the mechanisms appropriate to source and target requires considerable insight and intuition. The result distinguishes pedestrian science from innovative science.

In the sciences, decisions about which properties of the source system are central for understanding the target, and which are incidental, are resolved by careful testing against the world. As a result of testing and deduction, a well-established model in the sciences accumulates a complicated aura of technique, interpretation, and consequences, much of it unwritten. One physicist will say to another "this is a conservation of mass problem" and immediately both will have in mind a whole array of knowledge associated with problems modeled in this way. This use of sources already well-tested to gain insight into new problems has much to do with the cumulative nature of the scientific enterprise.

There is a close relation between this construction of metaphor and our earlier discussion of building blocks. In the sciences, new building blocks are usually constructed by combining building blocks from a level of greater detail: proteins from amino acids, amino acids from atoms,

atoms from nuclear particles and electrons, and so on. A new building block opens up whole realms of possibility because it can be combined in so many ways with extant building blocks. And even a fixed set of building blockss can be used over and over again without seriously impairing the chances for original discoveries: Think of the words in a dictionary or folk themes in music. To use these combinatorial possibilities one must select and build upon salient, regular patterns.

It is a matter of speculation, but worth examining, that the mechanisms of selection in the creative process are akin to those of evolutionary selection, simply running on a much faster time-scale. Speed-ups, simple though they are in concept, can sometimes radically revise our understanding. A lapse-frame movie of a wild grapevine moving up a tree looks remarkably purposeful, and lapse-frame animation of geological evolution shows the fluid, responsive, coherent movement of clouds in the sky. A lapse-frame animation of the evolution of some family of organisms shows the tentative probes, withdrawals, redirections, and cumulative construction we associate with creative activity. In both evolutionary and creative exploration we encounter patterns and lines of development (strategies) that emerge under selection. And, in both cases, emergent building blocks propagate their effects in cumulative ways, through recombination and interaction. There's not room here for an extended discussion, but the interested reader can learn more by perusing *Hidden Order* (Holland, 1995). There is much to be learned, I think, by modeling cognition via a translation of the mechanisms of natural selection, mimicking Maxwell's translation from gears to fields.

There are those who argue that "evolution is too slow to produce the complex mechanisms we observe in living organisms", or that "there is not enough time and experience to produce the complex grammar employed by young humans (so it must be 'wired in')". I would say these arguments fail to appreciate the speed-ups offered by building blocks. Grammars offer a good example because the "not enough time" argument has been forcefully used by well-known scientists. Yet, we know from psychology and physiology that there is a hierarchy of building blocks: phrases formed of word combinations ← words formed of phonemes ← phonemes formed of common sound elements ← expressed sounds formed by combining short muscular routines ← muscular routines formed by the repetitive firing of assemblies of neurons. The neural system produces *hundreds of thousands* of tests *each day* of the lowest level building

blocks - the neural assemblies. Even very simple adjustment procedures can produce exquisite, sophisticated adjustments of the interactions at such sampling rates. These adjusted interactions provide building blocks at the next level, which can then be combined into progressively more sophisticated routines - routines that ultimately play a role comparable to that of grammars. At this point in time, no one has produced a careful argument that shows that "there is not enough time" for such a hierarchy to develop under these sampling rates. Indeed, if we draw an analogy between generations in natural selection and successive tests of neural assemblies, we have reason to believe just the opposite - the time is more than adequate.

10. *Recapitulation*

Earlier, we examined two early human inventions - numbers and board games. These ancient pursuits were contrived long before humans began recording their intellectual achievements and they are simply described, though their discovery was far from simple. In both cases, the short, intuitive definitions generated objects that have been fruitfully studied to the present day. Both easily illustrate the "much from little" hallmark of emergence.

In the broader arena of metaphor and innovation, inventions like numbers and board games epitomize our human ability to reorganize perception through the use of abstraction and induction. Numbers, in particular, point up the uses of abstraction. To come to the concept of number almost all details must be dropped from multitudes of observations to arrive at regularities like "two-ness", "three-ness", and so on. We do this with the greatest of ease, once taught the trick, but it is no mean feat to discover the trick. It is even more of a feat to recognize the organizing powers of numbers. Over the centuries, numbers have moved from the counting of herds, to a basis for trade, to the Pythagorean and Archimedian theories of the world that replaced myths, to current practice that puts number at the center of the human scientific endeavor (for example, see Newman, 1956). This progression was far from obvious at the outset.

Board games are not usually accorded the same primacy as numbers, but I think they are an equally important cornerstone in the scientific

endeavor. In particular, I think board games, as well as numbers, mark a watershed in human perception of the world. A board game, qua game, only exists because the players act within the agreed constraints set by the game's rules. Though the rules must be fully and compactly specified for the game to be "playable", they can be contrived freely relative to the real world, subject only to incidental physical constraints involved in movement of the playing pieces. This freedom from direct physical constraints encourages modifications in the rules, accompanied by empirical judgments as to which rules yield a better game. Each new try amounts to a new miniature universe governed by fully defined laws. It is not a long step from such an outlook to the idea that the world itself might be rule-governed.

Above all, board games, unlike numbers in their raw form, capture the dynamic of unfolding actions and their consequences. There is an initial position and the successive actions of the players gives rise to a succession of positions, all within the constraints provided by the rules. Different actions causes different successions. "Cause and effect", as well as the possibility of controlling the outcome, become obvious in this context.

With board games there is a progression over time, similar to the progression for numbers. As we move forward in time from the board games of the early Egyptian dynasties, the rules of a game expand to become the "rules of logic". Thales' advocacy of "logical speculation" - the counterpart of our search for rules to explain systems exhibiting emergence - moved rule-making to a broader interpretation. This "logical speculation" required adherence to agreed upon rules of reasoning, followed by a comparison of the results with the real world. Thales specifically supported "logical speculation" as an alternative to traditional myths as a way of understanding nature. From Thales onward we have increasingly sophisticated attempts to model the world within a logical framework encapsulating cause and effect. Euclid's geometry evolves into such triumphs as Kepler's model of the solar system and Newton's laws of the universe. The sine qua non of these models is a small, easily comprehended set of laws that yields a wide range of testable consequences.

In the 19th century, Lyell and his compatriots developed models based on the rate of weathering of mountains and sediment deposition. A whole new conception of the world and its age came into being. Suddenly there was room enough and time for things to evolve, allowing an expla-

nation for the hitherto mysterious skeletons of “monsters” that had been encountered in quarrying. The laws were simple and the conception was testable. Even more important, the laws fit the constraints imposed by Newton’s laws. The effect was a cumulative extension of science. Each test of any part of the framework added credibility to the developing whole. When Darwin comes along with his astute observations and connections, embedding all within the constraints of Lyell’s geology, he gains a credibility that his grandfather, Erasmus Darwin, never gained, though Erasmus’ imaginative insight was of an equal order.

The conjunction of the logical dynamic offered by games with the universal measurability offered by number culminates in a form of modeling that typifies modern science. We see intimations of this relation in the Pythagorean theory tying numbers to the musical scale. Number, because of its extreme abstraction, can be attached to almost anything, and the laws of arithmetic nicely reflect various cumulative effects such as the merging of herds, the increase in height as standard blocks are added to a pyramid, the distance traveled at a steady pace, the relation between orbital distance and orbital velocity, the accumulation of sediment under regular weathering, and so on.

These progressions are the very essence of emergence. The terrain is convoluted, but there are landmarks:

mechanisms (building blocks, generators, agents) leading to *perpetual novelty* (very large numbers of generated configurations);

dynamics and regularities (persistent, recurring structures or patterns in the generated configurations);

hierarchical organization (configurations of generators become generators at a higher level of organization).

And, underpinning the whole venture, we have *models* and *model-building*.

Emergence is a matter of macro-laws, the obverse of reduction. Emergence is compounded when the macro-laws serve as building blocks for another layer of macro-laws. It is possible to formalize these relations (see *constrained generating procedures* in Holland, 1998), and in time we may be able to construct a genuine theory on the basis of some such formalism.

Interactions play a central role in the study of emergence. A detailed knowledge of the repertoire of an individual ant does not prepare us for the remarkable flexibility of the ant colony. The capabilities of a com-

puter program are hardly revealed by a detailed analysis of the small set of instructions used to compose the program. We will soon know the complete set of genes (or, at least, some of the alleles of each gene) coded in human DNA, but we will be far from understanding the program those genes specify - the program that takes a fertilized egg to the complicated 100 billion cell mature organism. The interactions of the cells in this vast ecosystem, the stuff of biology and medicine, are difficult to understand. But there is more. In that array of more than 100 billion cells there is a network consisting of several tens of billions of specialized cells called neurons. Understanding the behaviors mediated by these cells, the stuff of psychology, is much more than a matter of understanding the properties of isolated neurons. In all of these cases we have to develop an understanding of the constraints imposed by one part of the system on other parts. Typically, these constraints evolve as the system develops, with each part adapting to other parts.

That systems exhibiting emergence require studies that go beyond the simple reduction of studying isolated parts does not mean that they are beyond our grasp. After all, chemistry is a very successful science, even though we cannot understand that science via a direct investigation of the laws of physics. Patience is required. Games like chess and Go, with defining rules so simple they are quickly comprehended by a young child, have been studied for centuries, and still we learn. Why should we expect it to be different for the more intricate rules that define complex adaptive systems and other systems that exhibit emergence?

University of Michigan

Acknowledgements: This article draws heavily on my recent book, *Emergence: From Chaos to Order*. In places it incorporates sections of that book, sometimes slightly rewritten, othertimes substantially rewritten. I thank the book's publisher, Addison-Wesley Longman, for permission to use those sections.

REFERENCES

- Hofstadter D.R. (1979), *Gödel, Escher, Bach: an Eternal Golden Braid*. New York: Basic Books.
- Holland J.H. (1995), *Hidden Order: How Adaptation Builds Comple-*

- xity. Reading: Addison-Wesley.
- Holland J.H. (1998), *Emergence: From Chaos to Order*. Reading: Addison-Wesley.
- Maxwell J.C. (1890), *The Scientific Papers of James Clerk Maxwell*. London: Cambridge University Press.
- Newman J.R. (1956), *The World of Mathematics*. New York: Simon & Schuster.
- Simon H.A. (1969), *The Sciences of the Artificial*. Cambridge, Massachusetts: MIT Press.
- von Neumann J. (1966), *Theory of Self-Reproducing Automata*. Urbana: Univ. of Illinois Press.