The MIT Press

Complex Adaptive Systems

Author(s): John H. Holland

Source: *Daedalus*, Vol. 121, No. 1, A New Era in Computation (Winter, 1992), pp. 17-30

Published by: The MIT Press on behalf of American Academy of Arts & Sciences

Stable URL: http://www.jstor.org/stable/20025416

Accessed: 04-05-2017 09:42 UTC

John H. Holland

Complex Adaptive Systems

O NE OF THE MOST IMPORTANT ROLES a computer can play is
to act as a simulator of physical processes. When a com-
puter mimics the behavior of a system, such as the flow of
air over an airplane wing, it provides us with a unique way of
studying the factors that control that behavior. The key, of course, is
for the computer to offer an accurate rendition of the system being
studied. In the past fifty years, computers have scored some major
successes in this regard. Designers of airplanes, bridges, and Ameri-
ca's Cup yachts all use computers routinely to analyze their designs
before they commit them to metal. For such artifacts, we know how
to mimic the behavior quite exactly, using equations discovered over
a century ago.

However, there are systems of crucial interest to humankind that
have so far defied accurate simulation by computer. Economies,
ecologies, immune systems, developing embryos, and the brain all
exhibit complexities that block broadly based attempts at compre-
hension. For example, the equation-based methods that work well
for airplanes have a much more limited scope for economies. A
finance minister cannot expect the same accuracy in asking the
computer to play out the impact of a policy change as an engineer can
expect in asking the computer to play out the implications of tilting
an airplane wing.

Despite the disparities and the difficulties, we are entering a new
era in our ability to understand and foster such systems. The grounds
for optimism come from two recent advances. First, scientists have
begun to extract a common kernel from these systems: each of the

*John H. Holland is Professor of Psychology and Professor of Computer Science and
Engineering at the University of Michigan and Maxwell Professor at the Santa Fe Institute.*

17

systems involves a similar "evolving structure." That is, these systems change and reorganize their component parts to adapt themselves to the problems posed by their surroundings. This is the main reason the systems are difficult to understand and control—they constitute a "moving target." We are learning, however, that the mechanisms that mediate these systems are much more alike than surface observations would suggest. These mechanisms and the deeper similarities are important enough that the systems are now grouped under a common name, *complex adaptive systems*.

The second relevant advance is the new era in computation that is the theme of this issue of *Dædalus*. This advance will allow experts who are not computer savvy to "flight-test" models of particular complex adaptive systems. For example, a policy maker can directly examine a model for its "reality," without knowing the underlying code. That same policy maker can then formulate and try out different policies on the model, again without becoming involved in the underlying coding, thereby developing an informed intuition about future effects of the policies.

It is the thesis of this article that these new computation-based models, when constructed around the common structural kernel of complex adaptive systems, offer a much-needed opportunity: They enable the formulation of new and useful policies vis-à-vis major problems ranging from trade balances and sustainability to AIDS.

COMPLEX ADAPTIVE SYSTEMS

To arrive at a deeper understanding of complex adaptive systems—to understand what makes them complex and what makes them adaptive—it is useful to look at a particular system. Consider the immune system. It consists of large numbers of highly mobile units, called *antibodies,* that continually repel or destroy an ever-changing cast of invaders (bacteria and biochemicals), called *antigens*. Because the invaders come in an almost infinite variety of forms, the immune system cannot simply develop a list of all possible invaders. Even if it could take the time to do so, there is simply not room enough to store all that information. Instead, the immune system must change or adapt ("fit to") its antibodies as new invaders appear. It is this ability to adapt that has made these systems so hard to simulate.

The immune system faces the additional complication that it must distinguish self from other; the system must distinguish the legitimate parts of its owner from the ever-changing cast of invaders. This is a herculean task because the owner's cells and their biochemical constituents number in the tens of thousands of kinds. Mistakes in identification do occur in some individuals, giving rise to the usually fatal autoimmune diseases, but they are rare. The immune system is so good at self-identification that, at present, it provides our best scientific means of defining individuality. An immune system will not even confuse its own cells with those in a skin graft from a sibling, for example.

How does the immune system manage the ongoing process of adaptation that enables it to achieve such remarkable levels of identification? We do not really know, though there are interesting conjectures with varying degrees of evidence. Models of this complex adaptive system are hard to formulate. It is particularly difficult to provide experts in the area with models that allow "thought experiments," models that enable the expert to develop intuition about different mechanisms and organizations.

We face similar problems when dealing with the other complex adaptive systems.[1] All of them involve great numbers of parts undergoing a kaleidoscopic array of simultaneous interactions. They all seem to share three characteristics: *evolution, aggregate behavior,* and *anticipation.*

As time goes on, the parts evolve in Darwinian fashion, attempting to improve the ability of their kind to survive in their interactions with the surrounding parts. This ability of the parts to adapt or learn is the pivotal characteristic of complex adaptive systems. Some adaptive systems are quite simple: a thermostat adapts by turning the furnace on or off in an attempt to keep its surroundings at a constant temperature. However, the adaptive processes of interest here are complex because they involve many parts and widely varying individual criteria (analogous to the constant temperature sought by the thermostat) for what a "good outcome" would be.

Complex adaptive systems also exhibit an aggregate behavior that is not simply derived from the actions of the parts. For the immune system this aggregate behavior is its ability to distinguish self from other. For an economy, it can range from the GNP to the overall network of supply and demand; for an ecology, it is usually taken to

be the overall food web or the patterns of flow of energy and materials; for an embryo, it is the overall structure of the developing individual; for the brain, it is the overt behavior it evokes and controls. Generally, it is this aggregate behavior that we would like to understand and modify. To do so, we must understand how the aggregate behavior *emerges* from the interactions of the parts.

As if this were not complex enough, there is a further feature that makes these systems still more complex—they anticipate. In seeking to adapt to changing circumstance, the parts can be thought of as developing rules that anticipate the consequences of certain responses. At the simplest level, this is not much different from Pavlovian conditioning: "If the bell rings, then food will appear." However, even for simple conditioning, the effects are quite complex when large numbers of parts are being conditioned in different ways. This is particularly the case when the various conditionings depend upon the interactions between parts. Moreover, the resulting anticipation can cause major changes in aggregate behavior, even when they do not come true. The anticipation of an oil shortage, even if it never comes to pass, can cause a sharp rise in oil prices, and a sharp increase in attempts to find alternative energy sources. This emergent ability to anticipate is one of the features we least understand about complex adaptive systems, yet it is one of the most important.

There is one final, more technical point, that needs emphasis. Because the individual parts of a complex adaptive system are continually revising their ("conditioned") rules for interaction, each part is embedded in perpetually novel surroundings (the changing behavior of the other parts). As a result, the aggregate behavior of the system is usually far from optimal, if indeed optimality can even be defined for the system as a whole. For this reason, standard theories in physics, economics, and elsewhere, are of little help because they concentrate on optimal end-points, whereas complex adaptive systems "never get there." They continue to evolve, and they steadily exhibit new forms of emergent behavior. History and context play a critical role, further complicating the task for theory and experiment. Though some parts of the system may settle down temporarily at a local optimum, they are usually "dead" or uninteresting if they remain at that equilibrium for an extended period. It is the process of becoming, rather than the never-reached end points, that we must study if we are to gain insight.

MASSIVELY PARALLEL COMPUTERS

The introduction of the digital programmed computer profoundly changed our view of what could be accomplished with computation. Massively parallel computers—computers made up of hundreds of thousands of interconnected microcomputers—will produce changes that are equally profound. It is not just a matter of speed, though that is important. Because a massively parallel computer can handle large numbers of actions simultaneously, it offers new ways of displaying and interacting with data. It provides ways of studying complex adaptive systems as far beyond the reach of a current workstation as that workstation's capacities are beyond the reach of an adding machine or a slide rule. Indeed, massively parallel computers should produce a revolution in the investigation of complex adaptive systems comparable to revolution produced by the introduction of the microscope in biology.[2]

The longer-range effects of massive parallelism are not easy to predict at this early stage, but a little hindsight offers some clues. At the beginning of the computer era, in the 1940s and early 1950s, most computer scientists foresaw increasing speed and storage, along with an ever-increasing ability to tackle scientific and business problems. But the magnitude of those increases as they unfolded, coupled with precipitous decreases in price, amazed us. They made possible widespread word processing, electronic mail, the personal work station, and related sets of activities, such as personal video games and simulations. This has produced new major sectors of the economy and has altered both the work and play of large numbers of people. This process of headlong increases in speed and storage, accompanied by decreasing prices, is already underway for massively parallel machines. The new "microscope" will soon be as pervasive as the personal workstation is today.

MODELS OF COMPLEX ADAPTIVE SYSTEMS

A complex adaptive system has no single governing equation, or rule, that controls the system. Instead, it has many distributed, interacting parts, with little or nothing in the way of a central control. Each of the parts is governed by its own rules. Each of these rules may participate in influencing an outcome, and each may influence the

actions of other parts. The resulting rule-based structure becomes grist for the evolutionary procedures that enable the system to adapt to its surroundings.[3] We can develop a better understanding of these evolutionary procedures if we first take a closer look at this idea of a distributed, rule-based structure.

Most rules can be parsed into simple *condition/action* rules: If [condition true], then execute [action]. The simplest rules in this form look much like specifications for psychological reflexes: If [the surface feels hot], then execute [a backward jerk of the hand]; if [there is a rapidly moving object in peripheral vision], then execute [a movement of the eyes until the object is in the center of the visual field]. More complicated rules act on messages sent by other rules, in turn sending out their own messages: If [there is a message X], then execute [transmission of message Y]. Quite complicated activities can be carried out by combinations of such rules; in fact, any computation that can be specified in a computer language can be carried out by an appropriate combination of condition/action rules.

This distributed, many-ruled organization places strong requirements on computer simulation of complex adaptive systems. The most direct approach is to provide a simulation in which many rules are active simultaneously—a "natural" for massively parallel computation.

When many rules can be active simultaneously, a distributed, rule-based system can handle perpetual novelty. On encountering a novel situation, such as "red car by the side of the road with a flat tire," the system activates several relevant rules, such as those for "red," "car," "flat tire," and so on. It builds a "picture" of the situation from parts rather than treating it as a monolithic whole never before encountered. The advantage is similar to that obtained when one describes a face in terms of component parts, rather than treating it as an indecomposable whole. Select, say, 8 components for the face—hair, forehead, eyebrows, eyes, cheekbones, nose, mouth, and chin. Allow 10 variants for each component part—different hair colors and textures, different forehead shapes, and so on. Then $10^8 = 100,000,000$ faces can be described by combining these components in different ways. This at the cost of storing only $8 \times 10 = 80$ "building block" components. Moreover, when a building block is useful in one combination, it is at least plausible that it will prove useful in other, similar combinations. Building blocks thus give the system a capacity for transferring previous experience to new situations.

Massive parallelism is clearly an advantage in simulating a complex adaptive system conceived of in terms of simultaneously acting rules. An individual processor can be allocated to each rule, while the connections between the processors provide for rule interactions. The resulting model is both natural and rapidly processed.

To provide for adaptation, the system must have ways of changing its rules. Such procedures give the system its characteristic "evolving structure." There are two kinds of computational procedures that are relevant: *credit assignment* procedures and *rule discovery* procedures.

Credit assignment is necessary because one wants the system, and its rules, to evolve *toward* something. Credit assignment first requires a sense of what "good" performance is, then it requires a way to pick out and "reward" those parts of the system that seem to be causing good performance. A system that rewards good performance may never become optimal, but it can get better and better.

Credit assignment is a traditional problem in artificial intelligence research. In a rule-based system, the object is to assign credit to individual rules in proportion to their contribution to the system's overall (aggregate) performance. We can think of this credit as a *strength* assigned to the rule: The more a rule contributes to good performance, the stronger it becomes, and vice versa.[4] By "stronger" we mean that the rule, based on its past successes, is given a stronger voice in future decisions. As successive situations are encountered, the relevant rules compete to control behavior, the stronger rules being the likely winners. That is, if a rule has produced a good outcome in some situation in the past, then it is more likely to be used in similar situations in the future.

Credit assignment can enable a system to select the best from the rules it has, but it cannot supply the system with new rules. If it is to evolve to deal with new situations, the system will have to create new rules. For this the system requires some kind of rule discovery procedure. Rule discovery is a subtle process, because it is important that the discovery process generate *plausible* rules, rules that are not obviously wrong on the basis of past experience. The philosopher C. S. Pierce is quite informative on this matter.[5] To apply Pierce's reasoning to this model, it is convenient to think of rules as made up of smaller pieces, or building blocks. My own version of Pierce's commentary, then, is that the discovery and recombination of

building blocks is an important step toward assuring the plausibility of newly invented rules.[6]

To approach rule discovery in terms of building blocks, it is useful to think of "breeding" strong rules. That is, strong rules are selected as "parents," and new offspring rules are produced by *crossing* the parents. The assumption is that strong rules have valuable building blocks inside them that should be incorporated into new rules. This process mimics the process whereby a breeder crosses horses or a farmer produces new varieties of hybrid corn. Here the object is to produce offspring rules that amount to plausible new hypotheses. Rule discovery procedures of this kind are called *genetic algorithms.*[7] A genetic algorithm "learns" automatically by biasing future generations of rules toward combinations of above-average building blocks (as, in genetics, coadapted sets of genes appear ever more frequently in successive generations). It can be proved that genetic algorithms find and recombine useful building blocks. They have counterparts in each of the known complex adaptive systems. Of course, many of the new rules generated by this process are nonsense, but nonsense rules do not promote "good" behavior and are systematically weeded out.

This rule discovery procedure, once again, lends itself to massively parallel computation. Crossing strong parents is a simple operation that imposes low processing requirements on the computer. Because the whole set of rules can be treated as a population, with mating going on simultaneously throughout the population, parallelism is easily exploited.

INTERNAL MODELS: THE FUNDAMENTAL ATTRIBUTE OF COMPLEX ADAPTIVE SYSTEMS

There is still one property of complex adaptive systems that we have to examine more closely. Complex adaptive systems form and use internal models to *anticipate* the future, basing current actions on expected outcomes.[8] It is this attribute that distinguishes complex adaptive systems from other kinds of complex systems; it is also this attribute that makes the emergent behavior of complex adaptive systems intricate and difficult to understand.

It is interesting to note that we rarely think of anticipation, or prediction, as a characteristic of organisms in general, though we

readily ascribe it to humans. Still, a bacterium moves in the direction of a chemical gradient, implicitly predicting that food lies in that direction. The butterfly that mimics the foul-tasting Monarch butterfly survives because it implicitly forecasts that a certain wing pattern discourages predators. A wolf bases its actions on anticipations generated by a mental map that incorporates landmarks and scents. The science of computer simulations itself represents man's attempt to make predictions ranging from the flight characteristics of yet untried aircraft to future GNP, but we have only recently been able to endow programs themselves with model-building capabilities. It is important that we understand the way in which complex adaptive systems build and use internal models, because so much of their behavior stems from anticipations based on these internal models.

An internal model allows a system to look ahead to the future consequences of current actions, without actually committing itself to those actions. In particular, the system can avoid acts that would set it irretrievably down some road to future disaster ("stepping off a cliff"). Less dramatically, but equally important, the model enables the agent to make current "stage-setting" moves that set up later moves that are obviously advantageous. The very essence of attaining a competitive advantage, whether it be in chess or economics, is the discovery and execution of stage-setting moves.

An internal model may, of course, be incorrect in some or many ways. But then hindsight can be used to improve the model; the model is modified whenever its predictions fail to match subsequent outcome (credit assignment again). The system can thus make improvements without overt rewards or detailed information about errors. This is a tremendous advantage in most real-world situations, where rewards or corrective information occur only at the end of long sequences of action. Whether one is playing a game of chess or making a long-term investment, the rewards for current action are usually much delayed. Internal models enable improvement in the interim.

AN INTERIM SUMMARY

Here's a condensed view of the description of complex adaptive systems presented so far. The systems' basic components are treated as sets of rules. The systems rely on three key mechanisms: *parallelism, competition,* and *recombination.* Parallelism permits the system

to use individual rules as building blocks, activating sets of rules to describe and act upon the changing situations. Competition allows the system to marshal its rules as the situation demands, providing flexibility and transfer of experience. This is vital in realistic environments, where the agent receives a torrent of information, most of it irrelevant to current decisions. The procedures for adaptation—credit assignment and rule discovery—extract useful, repeatable events from this torrent, incorporating them as new building blocks. Recombination plays a key role in the discovery process, generating plausible new rules from parts of tested rules. It implements the heuristic that building blocks useful in the past will prove useful in new, similar contexts.

Overall, these mechanisms allow a complex adaptive system to adapt, while using extant capabilities to respond, instant by instant, to its environment. In so doing the system balances exploration (acquisition of new information and capabilities) with exploitation (the efficient use of information and capabilities already available). The system that results is well founded in computational terms, and it does indeed get better at attaining goals in a perpetually novel environment.

ACCESS TO SIMULATIONS OF COMPLEX ADAPTIVE SYSTEMS

Simulations of complex adaptive systems, executed on computers, produce floods of data. The result is reminiscent of the early days of "batch processing" on computers: When the output appears as interminable pages of printout and numerical tables, it is difficult to uncover significant or surprising interactions, much less react to them. The user can be reduced to observing, rather than experimenting and controlling. This need not be.

If we are to make parallel simulations of complex adaptive systems accessible, two criteria must be satisfied. First, the parallel simulation must directly mimic the ongoing parallel interactions of the complex adaptive system.[9] Second, there must be a visual, game-like user interface that provides natural controls for experts not used to exploring systems via computers. For example, a policy maker should be able to try out an economic model in much the way that a pilot tries out a flight simulator. Actions and decisions should be made in the usual way, without requiring any cognizance of the

underlying computations. It should also be easy to see if the model behaves in realistic ways in well-known situations. This has the additional value of allowing experts to feed back "reality checks" to the simulation designers. Research initiatives at the Santa Fe Institute, in cooperation with a commercial firm, SimLabs, lead us to believe that powerful interfaces of this kind are possible for complex adaptive systems.

CURRENT SIMULATIONS OF COMPLEX ADAPTIVE SYSTEMS

We are only in the earliest stages of developing simulations of the kind just discussed, but there are some suggestive results. The work of Marimon, McGrattan, and Sargent on the evolution of money provides an early example. It was initiated in 1989 as part of the economics program at the Santa Fe Institute.[10] This study uses a combination of theory and simulation to study the effect of adaptive, rule-based agents in a classical trading model from economics, Wicksell's triangle. It shows that even when the artificial agents start with randomly generated rules, they soon decide upon a medium of exchange and reach close-to-optimal trading patterns. Among other studies, there is a new approach to understanding the immune system using a massively parallel computer,[11] and an actual policy study using data from the office of management and budget in Milan, Italy.[12] The latter is directly concerned with increasing the efficiency of decision making in the 730 offices scattered throughout the Lombardy region. The study's major objective, which it attained, was to discover which factors, from a very large number, were relevant to the various decisions made by the local offices. By using this information, the director structured decision procedures that would lead to increased efficiency in the local offices.

These early results are really only accessible to the computer savvy, but they point the way. In all three of the models cited, the study of the mechanisms providing evolutionary changes in the system's structure will encourage more realistic, more accessible models. We can then expect current exploratory research to expand into substantial advances available to a wide range of users.

MATHEMATICS AND THEORY

Complex adaptive systems are so intricate that there is little hope of a coherent theory without the controlled experiments that a mas-

sively parallel computer makes possible. At the same time, in an area this complex, experiments unguided by an appropriate theoretical framework usually amount to little more than "watching the pot boil." Sustained progress outside the guidelines of a theory is as unlikely as attempting modern experimental physics outside the framework of theoretical physics. After all, no system currently under investigation in physics is as complex as a full-fledged complex adaptive system. We need experiments to inform theory, but without theory all is lost.

Fortunately, there are several points at which we can bring mathematics to bear on the approach outlined above. We can show that, under certain conditions, appropriate credit assignment procedures do indeed strengthen the relevant stage-setting rules. We can also show that recombination, mediated by a genetic algorithm, does progressively bias the population of rules toward the use of above-average building blocks.[13] There are also formal frameworks that apply to the process of generating internal models, with accompanying proofs that establish some of their elementary properties.[14] On a broader perspective, there are relevant pieces of mathematics from mathematical economics and mathematical ecology that can be generalized to apply to all complex adaptive systems.[15]

The challenge is to weld these disparate pieces into a theory, a theory that explains the pervasiveness of the evolutionary processes forming the common kernel of all complex adaptive systems. The theory should elucidate the mechanisms that assure the emergence of internal models. Coordinated computer simulations should provide critical tests of the unfolding theory. The simulations should also suggest well-informed conjectures that offer new directions for theory. The broadest hope is that the theoretician, by testing deductions and inductions against the simulations, can reincarnate the cycle of theory and experiment so fruitful in physics.

To my knowledge there is only one organization, the Santa Fe Institute, that has taken the general mathematical study of complex adaptive systems as its central mission.[16] The institute has drawn to its campus a unique range of experts in physics, economics, and related mathematical disciplines. It has formed a working alliance with the University of Michigan to take advantage of that university's particular strengths in psychology, sociology, and business administration. Even though the institute is only five years old, these

interactions have already produced substantial changes in the study of complex adaptive systems.

SUMMARY

Complex adaptive systems represent the kernel of some of our most difficult problems, ranging from trade balances to control of the AIDS epidemic. They can be simulated on massively parallel computers by defining a network of interacting rule-based components. By providing natural "flight-simulator-like" interfaces for such simulations, we can open these systems to exploration by policy makers and other experts who do not have the time to become computer savvy. This has the double value of giving the designers "reality checks," while allowing policy makers to explore the differences effected by different policies. By looking for pervasive phenomena in such experiments, we can implement the classic hypothesize-test-revise cycle for the study of complex adaptive systems. The experimental part of this cycle is particularly important, because such systems typically operate far from equilibrium, continually undergoing revisions and improvements. They do *not* yield to classic, equilibrium-based mathematical approaches that rely on linearity, attractors, fixed points, and the like. A new kind of mathematical framework is required, one that emphasizes continuing adaptation through recombination of building blocks.

Without such a framework, the computer-based experiments will be little more than uncoordinated forays into an endlessly complex domain. With such a framework, we can greatly expand our understanding of these important, difficult questions.

ACKNOWLEDGEMENTS

ENDNOTES

[1] See D. L. Stein, ed., *Lectures in the Sciences of Complexity* (Reading, Mass.: Addison-Wesley, 1989); and P. W. Anderson, K. A. Arrow, and D. Pines, eds., *The Economy as an Evolving Complex System* (Reading, Mass.: Addison-Wesley, 1988).

[2] J. H. Holland, "A Universal Computer Capable of Executing an Arbitrary Number of Sub-Programs Simultaneously," *Proc. 1959, Eastern Joint Computer Conference* (New York: Institute of Electrical and Electronic Engineering, 1959), 108–13; and W. D. Hillis, *The Connection Machine* (Cambridge: MIT Press, 1985).

[3] See J. H. Holland, K. J. Holyoak, R. E. Nisbett, and P. R. Thagard, *Induction: Processes of Inference, Learning, and Discovery* (Cambridge: MIT Press, 1989).

[4] Ibid., 70–75.

[5] C. S. Pierce, *Collected Papers, Science and Philosophy,* vol. 7, ed. A. W. Burks (Cambridge: Harvard University Press, 1958).

[6] J. H. Holland, *Adaptation in Natural and Artificial Systems* (Ann Arbor: University of Michigan Press, 1975).

[7] Ibid., 89–140.

[8] J. H. Holland, "Emergent Models," in A. Scott, ed., *Frontiers in Science* (Cambridge: Blackwell, 1990).

[9] S. Forrest, "Emergent Computation," in S. Forret, ed., *Emergent Computation* (Amsterdam: Elsevier [North-Holland], 1990).

[10] R. Marimon, E. McGrattan, and T. J. Sargent, *Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents,* Santa Fe Working Paper 89–004 (Santa Fe: Santa Fe Institute, 1989).

[11] S. Forrest and A. Perelson "Genetic Algorithms and the Immune System," in H. Schwefel and R. Maenner, eds., *Parallel Problem Solving from Nature* (Berlin: Springer Verlag, 1991), 320–25.

[12] N. H. Packard, "Genetic Learning Algorithm for the Analysis of Complex Data" *Complex Systems* 4 (1990): 543–72.

[13] Holland, *Adaptation in Natural and Artificial Systems,* chap. 7.

[14] Holland, Holyoak, Nisbett, and Thagard, *Induction,* appendices 2A and 2B.

[15] Anderson, Arrow, and Pines, eds., *The Economy as an Evolving Complex System,* 75–97, 205–41.

[16] D. Pines, ed., *Emerging Syntheses in Science* (Reading, Mass.: Addison-Wesley, 1987).