In This Issue

Message from the Editor

Computational modeling of adaptive complex systems

Much of the research of the Santa Fe Institute centers on the study of complex adaptive systems. There is an immense diversity represented by the systems that we describe by the term "complex adaptive." Many examples of this diversity are present in the articles in this issue, which include reports on continuing or beginning programs in economics, global sustainability, cultural evolution, the interplay between climate and culture, learning and memory in the immune system, and evolution and learning in organizations. But, in spite of this wide-ranging diversity, we see common themes in these systems: we see the promise of discovering common principles, and there is commonality in our approach to their understanding.

The study of most of these systems is not new; many have been studied, in some sense, since the dawn of intellectual endeavor. But we are beginning to see some common features. They share a common architecture. All feature more or less independent agents, interacting with each other in an environment that is, at least in part, determined by the behavior of the other agents. In economics and in organizations those agents may be individuals, firms, or even nations. In prehistoric culture the agents may be individuals, families, clans, villages, or tribes. In the immune system the agents are different types of cells and molecules. We seek to understand the structure and behavior of these agent-based systems.

In addition to the search for features and principles that are common to several systems, what is new in our study of these systems is the use of new computational approaches to modeling. At SFI these approaches are called adaptive computation, a program that provides the mathematical and computational underpinnings for much of the modeling in other SFI activities. This program includes several thrusts: fundamental study of the principles of adaptive computation; the search for new techniques, in what may be a large universe of adaptive computational methods; and the development of general, agent-based, adaptive computational modeling platforms. These computational platforms and the application of adaptive computation to a variety of systems form a major theme of this issue and, indeed, of the current work at SFI.

Mike

L. M. Simmons, Jr.
Vice President, Academic Affairs

A New Home for the Institute

By the end of the year, and perhaps sooner, Santa Fe Institute expects to move to its new, and permanent, home. We have purchased the 32-acre site known locally as the Hurley Estate, after General Patrick Hurley, who built the 12,000-square-foot house on a hill in northeast Santa Fe for his family in the late 1950s. The house was designed as a place for entertaining large numbers of people, and the public spaces are large. We are particularly looking forward to what will be our new conference room, which will finally accommodate our workshops comfortably, as well as to a real library.

Researchers and staff will find the new space both inspiring and confining. The inspiration comes from the juniper- and pinon-covered hillsides that will surround us; the confinement comes from trying to fit into a house with fewer, but larger, rooms than we currently occupy in our office suites. We expect to build additional space as soon as finances permit. In the meantime we have our attention focused on installing a computer and telephone network, upgrading the utilities, and minimizing the disruption that any move entails. We ask our friends' tolerance for what will surely be some discontinuities during that period.

SFI was gratified to discover how much community support emerged on our behalf when we presented our case for a zoning variance before both the Santa Fe Zoning Board of Adjustment and, eventually, the Santa Fe City Council. Hundreds of people, many of whom we had not known before, stepped forward to say SFI would not only be a good neighbor, but that it is an important asset to the community. We intend to live up to their expectations.

The Hurley Estate, a 32 acre site in northeast Santa Fe, will be SFI's new home.

Undergraduates Join SFI Research Staff

For five undergraduates, summer vacation was given over to research at SFI. The students, who came to us from across the United States, worked with faculty mentors on individual projects focusing on some aspect of the computational properties of complex systems. These residencies are supported by the National Science Foundation through its Research Experiences for Undergraduates Program.

Jean Czerlinski is a senior majoring in physics at New College of the University of Florida. She worked with Mats Nordahl and Cris Moore on topics in dynamical systems and the dynamics of evolutionary systems. In particular, she investigated cellular automata where the rules can change the topology of the underlying lattice, and worked with Mats and Cris on evolutionary models that use Lindenmayer systems to take development into account.

Matthew Headrick, a senior in physics at Yale, worked with Melanie Mitchell, Peter Hraber, and Jim Crutchfield on their "Evolving Cellular Automata" project. This project involves applying genetic algorithms (GAs) to evolve cellular automata (CA) with desired dynamical and computational behaviors. Given the widespread applicability of CA as models of natural complex systems and as parallel computing devices, finding a way to automate the process of designing CA has great significance for a number of fields. They are also finding that studying the GA's behavior on this task is yielding insights into how the GA works, what impediments it can run

into, and what mechanisms underlie its ability to evolve a complex system by working directly only on simple, locally interacting components.

Nelson Minar, a mathematics senior at Reed College, worked on analyzing the effect of imposing spatial locality on emergent autocatalytic networks of lambda-calculus expressions. The basic model of Walter Fontana and Leo Buss, called "ALchemy," was implemented in a well-stirred soup model: all expressions interact with equal probability. Space has the effect of limiting the possible reactions: reactants are required to be near each other. They are exploring the effectiveness of spatial differentiation for encouraging speciation and reproduction.

With External Faculty member Alfred Hubler, Darren Pierre examined competitive control and modeling of a logistic map, a sim-

ple nonlinear system. Darren is a senior at Case Western Reserve University majoring in physics. The intent of this project was to design a control and modeling system simple enough to be studied analytically as well as numerically. This work is based on work

---

There's an explosion of information, but much of it is very specialized. At SFI there's an alternative. Instead of simply finding more "answers," one can also see how these answers fit together as a whole, across disciplines and methodologies.
JEAN CZERLINSKI

---

by Hubler and David Pines in "Prediction and Adaptation in an Evolving Chaotic Environment" where more elaborate modeling and control algorithms were used, giving extensive numerical results. However, the complexity of

the algorithms made analytic predictions difficult.

David S. Pieczkiewicz, also from Case Western Reserve, is a January 1993 graduate in anthropology. David worked with Bette Korber attempting to design a method for comparing and contrasting the technique of DNA signature analysis with more traditional phylogenetic tools in studies of the epidemiological linkage of HIV sequences. The technique gained popular attention in the recent case of a Florida dentist thought to have infected some of his patients with HIV. Signature analysis identified a characteristic pattern in the dentist's HIV samples that was rare in the background population, but was found in several of his HIV-positive patients. On the basis of these findings, the Centers for Disease Control in Atlanta concluded that the dentist was the source of the patients' infection. The CDC's conclusion, and the validity of the signature analysis technique itself, has been debated since. David's primary task was to design and implement simulations of HIV DNA mutation. The results of these simulations will be used as source data in the key question of this project: when can signature analysis techniques be advantageous and contribute new information, when used in conjunction with more traditional approaches? If new insights to this question are gained by the project, such information may be helpful for understanding the general applicability of signatures to epidemiological linkage.

In late summer Michael Lowenstein, a 1992 graduate in biology from the University of

Undergraduate interns Matt Headrick, David Pieczkiewicz, Jean Czerlinski, Nelson Minar and Darren Pierre. Not pictured is Michael Lowenstein.

Massachusetts at Boston, joined the REU program. Michael is working with Christian Burks, SFI/LANL, on solution space characterization. The process of sequence assembly can be thought of as taking the original DNA, cutting it into small pieces which can be sequenced, and figuring how the pieces fit back together. There is only one correct answer but a combinatorially large number of potential solutions. The work of Burks and Lowenstein involves positioning each potential solution in the solution space along with its score. This work may provide insight into which algorithms are best for large-scale sequence assembly given specific kinds of DNA data.

New Members of the SFI Family

"What intrigued me about the Santa Fe Institute," writes new SFI Trustee **Gordon K. Davidson**, "was the prospect of interdisciplinary exchanges that could yield insights into a variety of complex problems. What impressed me most is the prospect for applying the new learning that results from these exchanges to real-world problems." Gordon is Chairman of the Corporate Practice Group of Fenwick & West, a general-practice law firm based in Palo Alto with an office in Washington, D.C. Davidson specializes in providing legal advice to high-technology companies. His practice includes general corporate law as well as technology licensing and protection for intellectual property. He has represented many high-technology Silicon Valley companies including Apple Computer,

Gordon Davidson   **David B. Weinberger**   **John Koza**

Cadence Design Systems, Electronic Arts, Symantec, and 3DO. Prior to entering law school Mr. Davidson worked as an electronics engineer on an advanced computer research project at Stanford Research Institute and, subsequently, as a computer software engineer at Measurex Corporation.

A recent addition to the SFI Board of Trustees, **David B. Weinberger** is a General Partner of The O'Connor Partnerships and a Managing Director of the Swiss Bank Corporation. The O'Connor Partnerships are well-known internationally for their successful mathematical approach to proprietary trading. Weinberger currently plays an advisory role in the areas of mathematical research and trading technology, new product development, legal and regulatory matters, and securities industry relations. He received a Ph.D. in mathematics from Princeton and prior to his years in the securities industry he taught applied mathematics and computer science as a visiting lecturer at Yale University, and was a research mathematician at Bell Laboratories. His primary research activity was in the field of combinatorial optimization, and

he is the author of a number of research articles in that field. "SFI is a unique interdisciplinary research environment where a new generation of computer-based tools is being used to identify and understand 'emergent' behavior common to complex systems across many different disciplines," notes Weinberger. "As one who is extremely interested in fundamental unifying scientific principles, I find this work very exciting."

New SFI Science Board member **John R. Koza** is a Consulting Professor in the Computer Science Department at Stanford University and President, Third Millennium Venture Capital Limited in California. He is the author of *Genetic Programming: On the Programming of Computers by Means of Natural Selection* and the forthcoming *Genetic Programming II: Automatic Discovery of Reusable Subprograms*, (MIT Press, 1994). "The Santa Fe Institute has successfully combined the interdisciplinary approach needed to grapple with many of the really important problems of science while maintaining a very high quality of analysis and work," he says. ○

Climate, Culture, and Complexity

Climate and culture are two of the dominant complex systems of the earth. The climate system has evolved in tandem with the planet, continuously interacting with the changing composition and configuration of the earth's surface and the organic life that it supports. Human culture, on the other hand, is a relatively recent phenomenon, the product of the evolution of learned behavior as a primary adaptive mechanism. The climate system has always exerted a powerful influence on cultural systems. But it is also true that culture has now evolved to the point where human behavior is affecting the climate.

SFI is planning an exploratory workshop on the interface between climatic and culture systems. "The Medieval Warm Period of the 9th to 15th Centuries: Large-Scale Interaction of Climate, Ecosystems, and Human Behavior in North America" will be chaired by Henry Diaz (NOAA), Malcolm Hughes, and Jeffrey Dean (both from U. Arizona). To probe the complex interactions among climate, the geosphere, biosphere, and culture, the meeting will focus in-depth on the climate of the Medieval Warm Period and its effects on the prehistoric environments and populations of North America. This particular case study has been chosen for several reasons: the Medieval Warm Period (MWP) was a time when global climate was markedly different from the present, with temperatures in the Northern Hemisphere averaging perhaps 1°C warmer than the 20th century. This period is of particular current interest because it comprises a range of climatic variability that may serve as a useful analog to future climate conditions forced by increases in atmospheric greenhouse gas concentrations. Not only are considerable paleoenvironmental, historical, and archaeological data available for this period, but the relationships between environment and culture are especially distinct in prehistoric Southwest societies.

One aim is to develop a spatial model of climate in North America from 850 A.D. to 1450 A.D. and to relate the results to contemporary climate. A characterization of MWP environmental variability should provide a basis for investigating climate-biosphere-culture interactions and for assessing the potential effects of global warming in North America. A second objective is to investigate the degree to which human cultural changes as revealed by historical and archaeological evidence can be linked to changes in the physical environment. In addition to looking at systematic interactions during the MWP, the meeting will also consider the impact of the environmental change immediately following this era which inaugurated the Little Ice Age. Focus on behavioral adaptations to both low-impact, infrequent environmental changes and rapid "regime transitions" should shed light on the interactions among complex adaptive systems deriving from a wide range of time scales. It also promises an empirical basis for evaluating the consequences of changes in our current climate in light of projected greenhouse-gas-induced change. It is expected that a proceedings volume will result from the workshop.

○

Evolution of Human Organizations

SFI's research in prehistoric Southwest culture is leading to a broader approach that extends from the origins of human cultural behavior to state-level societies, including an exploration of the implications for advanced societies of the findings about simple societies. This broader-based approach will involve moving beyond the data from the Southwest, which is largely derived from a superb suite of artifacts, to combinations of prehistoric, historical, and physical records from much of North America (see Climate, Culture, and Complexity, in this issue) and from places like the Roman Empire, or more modern census records from China and India.

This approach to the evolution of culture may focus on the study of the emergence of collective behavior from independent agents whose actions are based on evolving individual schemata, leading to the evolution of social structures. Other important features might include the unintended consequences of the actions of independent agents, incorporation of biological traits in the models, cultural transmission of traits, the role of migration and the significance of linguistic diversity, and so on. (It is important to recognize that the agents in a cultural model might not be individuals but might be larger societal units such as clans, villages, or economic sectors.) This approach will require modeling platforms, like Swarm, that are sufficiently general to allow different researchers to view these very different problems.

This work is probably best done first in a general context, understanding the problems at a general-process level, and then pro-ceeding to a more detailed level. A working group on the evolution of culture, formed during a 1993 summer meeting of the SFI External Faculty, is currently formulating a plan on how to proceed with this approach.

Complementing this comprehensive approach is a parallel initiative that is focusing specifically on evolution and learning in modern organizations. In March 1992 the Institute held a workshop "Adaptive Processes and Organization," co-chaired by Michael Cohen (Institute of Public Policy Studies, U. Michigan) and David Lane (U. Minnesota). Two broad domains of discussion emerged, each with rich potential for further collaboration. The first centered on models that shared a common root metaphor of "adaptation as landscape exploration." Several participants interested in these landscape models found it promising to explore their common fascination with technological innovation in economies. Levinthal and Kauffman may work along these lines. Introduction of Hebbian rules into economic and political landscape models, suggested also by Gérard Weisbuch, turned out to be quite interesting to both Axelrod and Levinthal. The second, sometimes referred to as the "structuralist" concern, centered on how multiple agents interact to give rise to a superordinate entity with its own coherence, that may in turn constrain the subsequent actions of the lower-level agents. Weick, Padgett, Buss, Holland, and Axelrod are especially interested in pursuing questions about the emergence and persistence of organizational entities. Their empirical work involves systems as different as aircraft car-rier crews, state organizations in Europe such as the Medici in Florence, and multicellular organisms.

It has been noted that in the real-world, high-technology firms do not operate according to the classical theory of the firm (echoing much of Brian Arthur's work on why knowledge-based companies respond to markets differently than resource-based companies). SFI plans to further consider these modern innovative enterprises, which are examples of rapidly evolving complex adaptive systems and may be particularly amenable to study. Specific topics may include research into the relation of selection pressures to the evolution of an organization, particularly because selection pressures on individuals within firms can vary significantly from selection pressures on the firm itself; here there are obvious parallels with biological communities and ecologies. This work could shed light on our understanding of why modern organizations seem to break so catastrophically.

Possible case studies may be evolution and learning in government bureaucracies; change in military organizations, which tends to be introspective because it involves formal analyses of failures and successes; and perhaps the study of university organization. A segmented approach might be appropriate, one that would be specific to corporations and other organizations by class but that would then include cross-organizational aspects. Kenneth Arrow and Murray Gell-Mann have agreed to co-chair a founding workshop on adaptation and learning in modern organizations in the fall of 1993. ○
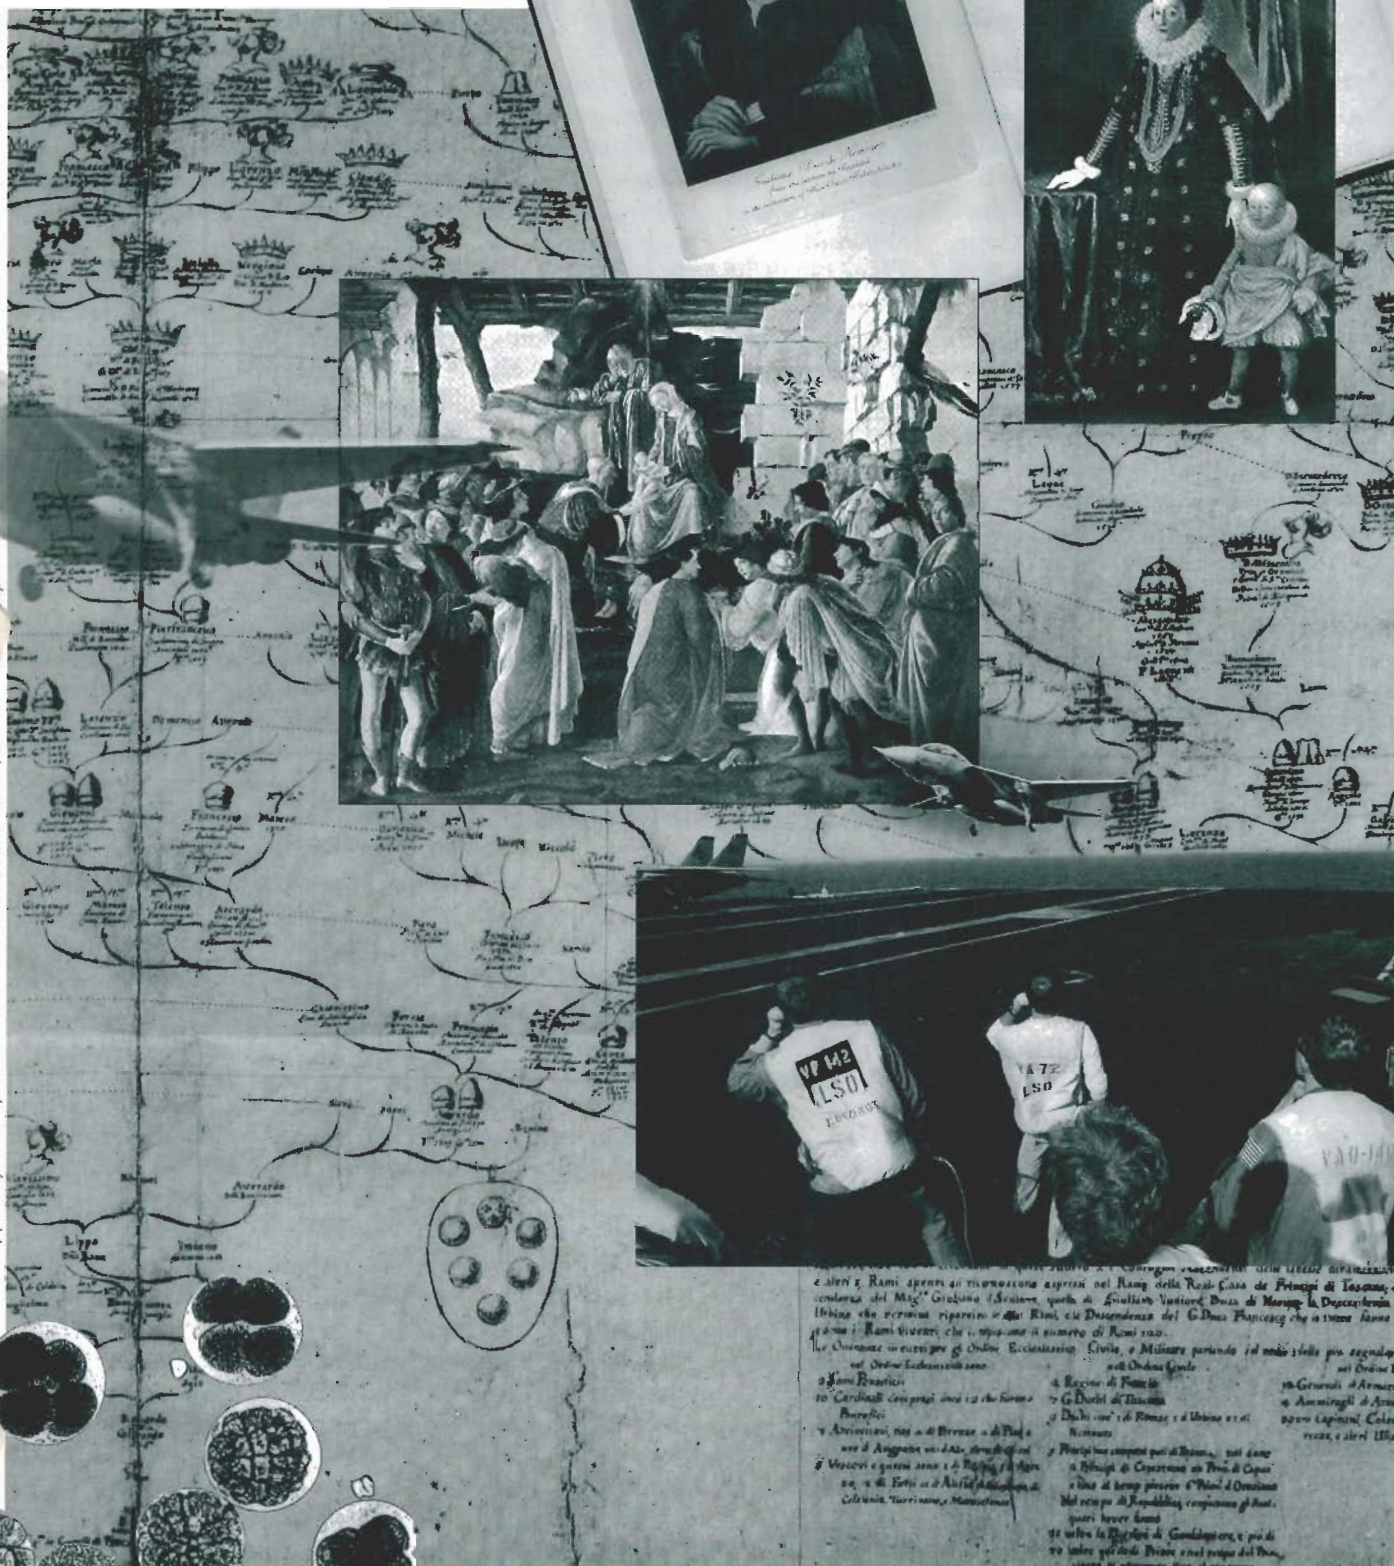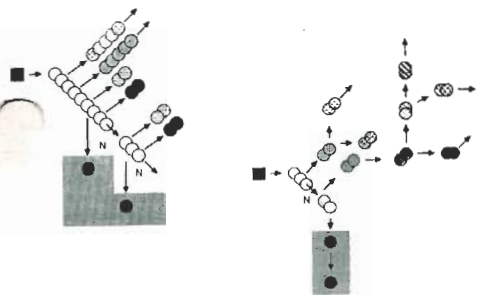
Understanding Scaling Phenomena

Scaling is a natural theme for the Institute since some manifestation of it occurs in many of the complex phenomena currently being investigated in its various programs. Indeed, in many of these areas, phenomenological scaling laws provide the only quantitative means for organizing the "data," thereby allowing some sense of "order" or "coherence" to be brought to the "system." The response of any such complex system to a change in scale typically reflects some deep underlying features that are often independent of its detailed dynamics or the model used to describe it.

Many researchers presently involved with the SFI have, at various times, explored aspects of scaling within the context of their own field. Such scaling phenomena occur over a wide and ubiquitous area ranging from biology and economics to classical nonlinear phenomena (under the guise of fractals, chaos, self-similarity, and the like) and quantum field theory. An initial SFI workshop in this field took place in 1991 when Philip Anderson and Sidney Nagel co-organized an interdisciplinary meeting focusing on self-organized criticality. Now Geoff West and Murray Gell-Mann plan to establish a program in scaling.

Classic examples of scaling include Pareto's law in economics, Zipf's law in linguistics, Barenblatt's problem in diffusion processes, Wilson-Kadanoff scaling in phase transition, asymptotic freedom and the unification of all forces in elementary particle physics, and the evolution of the universe since the Big Bang. The scaling phenomena associated with these sorts of problems are closely related to problems of pattern formation in systems such as snowflakes and earthquakes, river flow and the formation of landscapes, and the crucial concept of self-organized criticality. It extends, for example, to fundamental ecological questions as to how the leaf-level processes of photosynthesis and respiration might scale up to provide a better understanding of global atmospheric trends.

Generally speaking, scaling up from the small to the large is usually accompanied by a change from simplicity to complexity; typically this occurs while maintaining cer-

> *Generally speaking, scaling up from the small to the large is usually accompanied by a change from simplicity to complexity . . . .*

tain fundamental or rudimentary elements of the "system" invariant or conserved. We see this in architecture and engineering, in the various physical laws of nature and, most dramatically, in biological evolution itself. Perhaps all of these phenomena and trends have some sort of unified description and mathematical origin.

A possible paradigmatic structure for discussing some of these questions within a general framework is one that has arisen in the context of understanding scaling phenomena in quantum field theory and statistical physics. The analytic technique developed there is called the renormalization group; it is a rather powerful and general one and, like old-fashioned classi-cal dimensional analysis, it has potentially important applications in all areas of science beyond the original confines where it was discovered. It has a precise formulation and, consequently, precise quantitative statements concerning scaling in both elementary particle physics and phase transitions phenomena in statistical mechanics. For example, anomalous dimensions and the consequent power law behavior that appear in this context are quite analogous to fractal or Hausdorf dimensions that occur, for example, in classical complex systems. The response of such systems to a probe (which measures and quantifies its size, shape, and tempero-spatial structure) has much in common with self-similarity. This suggests that there may well be an underlying general unified approach along these, or analogous, lines to a variety of scaling phenomena occurring in many diverse areas. Some obvious questions that this line of thought raises are: Are there generalizations or analogues to the concept of a strength of "interaction" (be it in particle physics, molecular biology, economics, or the structure of language) that changes, or evolves, depending on the scale used to measure it? Are there "effective degrees of freedom" operating in such systems? Are there analogues to logarithmic scaling laws rather than the conventional power law behavior?

A program devoted to understanding scaling may provide a unifying intellectual structure for crossing some of the boundaries between programs and creating a common language. ○

Project 2050: Transition to Sustainability

Project 2050 is a collaborative effort involving SFI, the World Resources Institute, and the Brookings Institution to address the very hard question of how we can achieve a sustainable existence on this planet by early in the next century. Project 2050 will develop new techniques for making an integrative analysis of this issue using a combination of alternative "visions," base studies, scenario development, and computer modeling.

Rob Coppock is director of this four-year project which is supported by funding from the John D. and Catherine T. MacArthur Foundation.

The term "sustainable development" was brought into common use by the 1987 Brundtland Commission report. The Commission's definition of the term—meeting the needs of the present generation without compromising the ability of future generations to meet theirs—has been a powerful concept for focusing global attention on the need to integrate environmental stewardship and economic development.

Making the concept of sustainability precise, however, has proven difficult. At least some resources must be consumed in meeting today's needs, and the needs of future generations can only be estimated. What, then, constitutes a sustainable society? What time frame should be considered? What indicators tell us whether or not we are moving toward sustainability? What are the strategies and actions that would get us there?

A number of related efforts are interwoven within Project 2050 to address these questions. One major activity includes the development of an information base. This includes undertaking a series of studies to develop quantative characteristics of sustainability; examining a series of cross-cutting issues like information systems, governance, and economic development; and developing innovative computer-modeling techniques. In this last area, of course, SFI has special expertise. In June of this year a small workshop was held at the Institute to continue brainstorming, begun last fall, about how to contruct CLAW—"Crude Look at the Whole" models. Another meeting is planned next spring.

Since many of the linkages among topics in the 2050 Project involve a large number of elements and masses of data and many are nonlinear, the Institute's work in simulations involving artificial worlds is especially relevant. One such artificial world is ECHO, a general-purpose modeling platform designed by John Holland and being developed at the Institute. ECHO is designed to capture many of the features of models used by economists, evolutionary biologists, and ecologists as abstractions of real-world systems. As part of Project 2050, graduate student Terry Jones and External Faculty member Stephanie Forrest are currently developing ECHO as a general software tool. Among the applications being considered for this platform, once developed, are modeling the population levels of introduced species in Hawaii, and modeling observed changes in desert ecology under experimental conditions. Jones is right now looking at speciation in this model. Some of the questions he is asking are: When does speciation happen? How long does it take? With what frequency? Is geographical isolation required for speciation? How does the ECHO model relate to the real world in terms of speciation? Once speciation occurs, do the reproductive barriers between species persist—do the new species tend to diverge genetically, or do they hybridize? How does the mate recognition algorithm affect speciation? Do we see something akin to punctuated equilibrium? Is the model close enough to the real world that we should expect to see that sort of thing in the first place? Does phyletic evolution ever get us anywhere? Is there anything analogous to species sorting? Finally, is there a meaningful analogy in ECHO between economic and reproductive adaptations to that in the real world?

Another general-purpose modeling platform under development at SFI that will be used in Project 2050 is Swarm, a system designed for capturing the interactions among a large number of independent agents. The environment is in part determined by the agents, and therefore is modified as a result of their interactions. Swarm is being designed for modeling applications as diverse as ecology, economics, and the evolution of human cultural behavior. (For further discussion of ECHO and Swarm, see "Computational Platforms" in this issue.)

The human response to global changes in climate, population, economic situations, and other aspects of the environment is condi-

tioned by cultural tendencies of belief and action. How cultures evolve thus plays a powerful in role in how—or if—we attain sustainable development. This is the basis of a project led by External Faculty member Gérard Weisbuch which explores the dynamics of cultural evolution induced by and, in turn, affecting changes in environmental conditions. In one model being studied by Weisbuch, Guillemette Duchateau and Valerie Gremillion, neural network agents (individuals, firms, and governments) choose actions based on their predictions about the future state of the world. Operating on past experience, information from other agents, and their own interests, agents learn about their environment and its likely patterns while simultaneously affecting it. The role of cultural differences, government and economic influences, and the information that results in response to environmental change is being explored using these models. One major question involves the effects of information broadcasting—advertising, propaganda, education, opinion—on the dynamics of cultural evolution. Another explores how the interactions of resource depletion, technological efficiency, investment, and population can help define sustainability as well as attain it.

In another 2050 modeling project, Robert Axelrod at the University of Michigan is developing a simulation model to account for the emergence of new political actors. One of the main problems to attaining sustainability is the so-called tragedy of the commons. The tragedy of the commons arises when many independent actors (people, villages, states, or whatever) each "over graze" because there is no mechanism to enforce the collective interests of all against the private interests of each. This leads to resource depletion, elimination of biodiversity, overpopulation, war, and other major social problems. A major route to the prevention of this situation is the emergence of a political actor based upon the organization of previously independent actors. Today we have actors at the national level, but we do not have very effective political actors at the transnational level to regulate resource use at the global level.

Political scientists have a variety of theories to analyze the emergence of new political actors, but as yet, they do not have any formal models that account for this emergence endogeneously. Axelrod is developing a simulation model that takes as given the existence of lower level actors, and generates higher level actors from the interactions among them. His criteria for a new political actor are subordination of members of the group; collective action; and recognition by others of the group as an actor. A minimal goal is to provide an existence proof, that is, to show that a set of rules about individual choices along with set of rules about the physical world can lead in a natural way to a set of actors forming themselves into a group that meets the criteria for the emergence of a political actor at a higher level. A maximal goal is to understand the condition under which new political actors not only emerge, but are able to solve joint problems such as common defense and sustainable growth. ○

Cortical Plasticity in Human Development: Is There Reason for an Optimistic View?

Work on understanding the complexities of the brain is proceeding on several fronts at SFI. The study group on theoretical neurobiology, now two years old, met this summer to explore how the central nervous system becomes organized to be both flexible and robust. Added to this is a new research initiative in which researchers are looking at brain development and memory organization within a complex systems perspective. One aim is to establish whether recent discoveries of cortical malleability warrant a new look at learning and rehabilitation of human infants and adults.

In May George Cowan and Bela Julesz co-chaired a workshop that brought together participants from a variety of fields and research centers to seek new evidence of brain plasticity. Discussions covered the development of the cerebral cortex; cortical topographic organization and reorganization; activity-dependent processes; memory consolidation; behavioral plasticity; development of higher functions; and rehabilitation.

Participants considered new evidence of brain development and sensory processing, among which are evidence of dynamic reshaping of the somatosensory and visual cortex; textual learning and Hebbian chain-formation in early vision; the role of sleep in memory consolidation; and constraints on plasticity of higher functions. The emerging new principles should help us better understand the mechanism of learning and memo-

ry and the dynamics of nervous system disorders. The general principles of strengthened intracortical interactions via a cascade of local connections and the filling in of weak gaps by strong, surrounding activity unifies the various topics. Equally important, these principles suggest new techniques in rehabilitation, such as enhancing certain circuits and diminishing others.

It is presently planned that a proceedings volume from the workshop will be available in 1994. The book promises to be unique because it joins basic research with the work of rehabilitation, presenting experimental work from researchers from sensory, visual, and auditory fields using all accessible techniques in anatomy, physiology, psychophysics, behavioral methods, linguistics, and modeling. ○

Simplicity and Complexity in the Arts

To what extent does the process of composition in the arts function as an evolving complex adaptive system, one in which random variation and selection play significant roles? This was one of the questions recently discussed by a small working group as it considered plans for a full-scale SFI workshop on simplicity and complexity in the arts.

The core group participating in the one-day meeting in April con-

sisted of poets Alice Fulton and Marcia Southwick, the visual artist Hank DeLeo, and SFI researchers Charles Bennett, Murray Gell-Mann, John Holland, and Seth Lloyd. Southwick and Gell-Mann were the coordinators.

The session started with presentations by the SFI scientists on the meaning of simplicity and complexity and on the properties of complex adaptive systems, with some remarks about possible connections with the arts. The artists responded with discussions of the process of creation and of the evolution of individual work and the work of schools and movements. A number of questions arose. How much of the process of creation is conscious? Should the element of intention or direction be treated merely in terms of selection pressures or as affecting the search procedure itself? How are random elements introduced? In what ways do "frozen accidents" generate emergent regularities in form and content?

Theory compresses into brief laws vast, even indefinitely large bodies of data. Can artistic expression be compared to theory? Can works of art be regarded as compressing aspects of experience into brief statements (or images, musical passages, metaphors, etc.)? If so, is it possible to analyze that process? What is the relationship between abstraction and simplicity or complexity?

The core group plans to convene a workshop in the fall of 1993, with a more formal agenda and a broader range of participants, including artists and scientists from fields not so far represented, such as architecture, prose writing, music, and biology. ○

Computer screen view of
SWARM simulation system
by permission of Chris Langton.

AT SFI WORKERS ARE DEVELOPING A SET OF COMPUTATIONAL PLATFORMS
THAT WILL HELP SPEED THE PRODUCTION OF USEFUL SIMULATIONS
AND MODELS FOR A VARIETY OF COMPLEX SYSTEMS.

COMPUTATIONAL PLATFORMS:
SETTING THE STAGE FOR SIMULATION

SWARM and ECHO: Building the Basic Tool Kit

Scientists involved in creating computer simulations of complex adaptive systems know it's a painstaking, complicated process that not only requires setting up the parameters of the research but working out the required computer codes. At SFI workers are developing a suite of general and specific computational platforms that will allow researchers and applications-oriented developers to accelerate the production of usable simulations and models of a variety of complex systems. "We are basically providing people with a user friendly way to access an object-oriented programming language," says Chris Langton, the developer of the Swarm platform. "We're trying to provide a library of tools. That library will be defined in such a way that it will be easy for people to add their own tools."

Platform developers predict these tools will allow scientists to spend more time on individual scientific research and make research results more comparable.

"Reproducibility is not a given at the present," says Walter Fontana. "If everyone built his own computer model, this would be almost impossible," he added. "If I am interested in what you are doing and you hand me a complicated computer code, first of all, the odds that I will understand it are very slim. The odds that I am going to want to understand it are also very slim because I want to do science and not spend ten months trying to understand your code. Ten different people working on ten different areas could write ten different pieces of code. It would be a waste of time."

Will the use of computer platforms inhibit scientific creativity? Richard Palmer, for example, thinks it's better not to be constrained by a particular software platform. He concedes the establishment of such platforms will be useful for those who have "an idea and want to play with it," but who aren't terribly good at computer design. But he notes that "once

you have a platform you try to fit something into it, you tend not to think in terms that might not fall within that framework. People are more creative if they have to figure out a platform for themselves that is suited to their particular project rather than trying to fit into a general framework. I think its healthy for people to construct their own systems. One thing it forces you to do when you are turning a mathematical model on paper into a computer model is to fill in all the details." Palmer, who has developed a specific stock market model (see page 18), does think that input, output, and graphical display packages from general platforms could help scientists in simulation programs. "It would be great to have some packages around that could be used for different approaches," he says.

The Swarm system, currently under development at the Institute, is one of the platforms available for general-purpose applications. Swarm allows researchers to con-

duct simulations consisting of large numbers of interacting objects, or agents. The common architecture—a "swarm"—refers to large number of simple agents interacting, whether they are a swarm of bees, an ant colony, a flock of birds, or cars in city traffic. Agents—with their own internal data and rules—act by passing messages back and forth to each other. The system also provides a field object to associate the agents with coordinates in space. The agents can modify the environment and in turn their behavior is dictated by the state of the environment, providing a feedback loop.

"We are attempting to capture the architecture in a general-purpose way," says Langton. "Then people modeling insect behavior, the economy, the behavior of molecules getting caught up in complex dynamics, or the evolution of populations can go to the same simulator and not worry about a lot of very subtle computer science and engineering issues. Everything is accessible and reprogrammable," he said. "Different levels of users can implement the system at different levels of sophistication. The naive user can write his code and use all the predefined tools. The more sophisticated user can get in there, change the way the images are displayed, and define a new space."

Nearly three dozen simulation projects, at the SFI and elsewhere, already have been identified as possible beneficiaries of the Swarm programming framework. To date, two working simulations—heatbugs and traffic—have been implemented into Swarm.

An artificial economy model may be a "nice test bed for Swarm," says researcher David Lane. "The model creates an artificial world. It's designed to describe what the agents' interactions are and let it go. Swarm would help you watch it."

The Artificial Economy Project takes an engineering approach to the problem of how the economy coordinates itself as a coherent whole, with large-scale structure. The model includes five types of agents within a closed system. One firm manufactures machinery, which is in turn purchased and used by a second firm to produce a consumer product. Both firms employ individuals who use their paychecks to purchase the consumer product. The first firm employs researchers to design new machinery. A bank accepts savings from the workers and the firms and provides loans. The gross domestic product can be computed at any time on the basis of all the interactions. "All the individual agents are making decisions and can learn from minute to minute," explains Lane. "It's not a global decision-making process. The individuals interact within the institutions in the market for machinery, consumer goods and labor. The question is does order and growth emerge? What kind of conditions will let the world grow? Will it get wealthier or will it collapse?" The simulation looks for global regularities, he notes. "Can we describe conditions under which an economy will produce particular manifestations of coordination or stability at the macrolevel that will last much longer than the micro agents making decisions?"

The first attempts to implement the model weren't successful because the prototype was "over designed" and agent behavior was too restricted. The researchers then switched to an object-oriented version. The Swarm platform should make experimenting with the artificial economy simulation easier. "You can define what agents are in the system," says Lane, "what properties they have, and their modes of interaction. And you want an experimental interface that allows you to change things easily even during run time."

Another possible application of Swarm being developed by George Gumerman is a cultural evolution model which will observe how cultures change and evolve. "The simulation will allow researchers to study how earlier cultures became more complex, evolved, or collapsed," says Gumerman. "We haven't been very good at answering these questions in the past."

"We can't experiment with human societies, especially extinct societies," he notes. Computer simulation, however, will help in determining whether there are typical processes in evolving societies. The model will look at early cultures of hunters and gatherers, the simplest societies, from Southwest Indians to Australian aborigines. These examples present clans, gender differences, age differences, and religious societies. "What we are interested in is how these early cultures develop, evolve, and change," he adds. Gumerman hasn't decided which platform will be best for the project, but says that Swarm, with its agents acting on their own rules and forming a greater behavior, might provide an appropriate venue.

Another model that may port into Swarm successfully is an algo-

rithmic chemistry to study biological organization being developed by Leo Buss and Walter Fontana. This model views molecules as agents which have a structure and an action determined by that structure. When these agents interact with another agent, new agents are formed. "The point of two chemicals interacting," says Fontana, "is to produce a new chemical, a carrier of new patterns of interactions." As a result of two agents interacting a new expression is released into the system. To keep the system constrained, a randomly chosen agent is deleted. Initially almost every interaction is innovative, resulting in a new expression. Then a variety of different combinations start to produce the same expression. "Construction permits diversity with more than one pathway leading to the same thing," notes Fontana.

The algorithmic chemistry model allows researchers to draw conclusions about biology through abstractions from chemistry. "Porting our model onto the Swarm platform will give us a series of advantages," he said. "Swarm enables one to expand an existing model without having to rewrite major portions of the code. All you have to add are the specific parts."

As part of Project 2050, Gérard Weisbuch is developing a model of a sustainable world that can be implemented on a generic platform. Such a sustainable world requires a set population and new technologies. The agents can be individuals, governments, or firms. Their internal structures are represented by neural nets. Each agent has a brain—a neural net—and is able to learn from past experiences. But

the interests of the agents vary and interest matrices and neural nets are further changed by the learning process. Agents' decisions are made based on their interests and the information received. It will be at least another year before the model is complete and then more than likely the model will fit into the Swarm system, Weisbuch says. "Swarm is open. It's easy to bring changes and include more agents."

ECHO, a model originated by John Holland, is another generic platform. It too is based on a large number of diverse agents with internal models (rules) that direct behavior and form an aggregate behavior. The persistence of any agent depends directly on the context provided by the rest. The agents adapt, evolve, and reproduce over a geography with different inputs of renewable resources at various sites. Each agent has simple capabilities defined by a set of genes.

"The idea of Swarm is a general platform for modeling complex systems," says Terry Jones who is working on developing a model of speciation within the framework of ECHO. "In a sense ECHO is trying to do the same thing in a less ambitious way. It is more rigid than Swarm. It assumes certain things about what the world looks like, what kind of agents interact, and what happens when they do interact."

ECHO has a world with a set of two-dimensional sites inhabited by agents, explains Jones. The agents interact either through combat, trading, sexual reproduction, or asexual reproduction. Jones gives the real-world example of ants, flies, and caterpillars. The flies lay eggs on the caterpillar's back; if the

eggs aren't eaten by nearby ants, the fly larvae eventually will eat the caterpillar. Both the flies and ants and the flies and caterpillars have an antagonistic relationship. But there is a mutualism—a trading relationship—because the ants get food from the caterpillars and the caterpillars get protection from the ants.

There are typically a small number of resources, three or four, in ECHO's world. In order to reproduce, an agent must acquire enough resources. The agents also require resources to pay a tax. "You can't stay alive without expending energy at some rate," says Jones. "And, if an agent gets charged a tax he can't pay, he is considered bankrupt and essentially is taken out of the world. The resources are returned to the environment." Reproduction of the agents is based on genetic algorithms. Eleven genes determine various characteristics of the agents such as their ability to fight, who they will mate with, and their defense and trading strategies. All are subject to genetic mutations and crossovers.

Jones predicts ECHO will serve as a platform for biology, ecology, and perhaps economics. He personally is working on a speciation model. "The most widely held belief about speciation is that you need some type of geographic isolation dividing populations in the species," he notes. "Another theory is that speciation can occur within a population without separation." The modeling project will look at when speciation occurs, the time it takes, the frequency, and if geographic isolation is required, he said. And, if speciation does occur, do the reproductive barriers be-

tween species persist? "It will be interesting to find out how speciation happens in this model," Jones says. "It should be of interest to biologists and perhaps could be applicable to the real world. The hope is that an economist, ecologist, or biologist can come along and start to do realistic modeling in ECHO without being a computer programmer."

Chris Langton concurs. "What computers do for science is they often turn good scientists into bad programmers," he argues. "There are an awful lot of scientists who spend a good deal of their careers wallowing around in really bad code and not really getting what they want out of the computer. To do a really good job at whipping out some really good and efficient simulation, you have got to master a lot of computer engineering as well as mastering whatever specialty you are in. Most scientists just haven't had the proper training in how to write a good efficient bit of code."

The Swarm and ECHO platforms should go a long way toward helping researchers do good science without the hindrance of bad code.
○

Linda Little is a freelance science writer based in the Dallas area.

SFI Researchers Working on Simulations Mentioned in This Article

Brian Arthur, Economics and Food Research Institute, Stanford University

Robert Axtell, Brookings Institution

Leo Buss, Biology, Yale University

Joshua Epstein, Brookings Institution

Walter Fontana, Santa Fe Institute

George Gumerman, Santa Fe Institute

David Hiebeler, Santa Fe Institute, Computer Science, Harvard University

John Holland, Computer Science, Engineering and Psychology, University of Michigan

Terry Jones, Computer Science, University of New Mexico and Santa Fe Institute

Timothy Kohler, Anthropology, Washington State University

David Lane, Statistics, University of Minnesota

Christopher Langton, Complex Systems Group, LANL, and Santa Fe Institute

Blake LeBaron, Economics, University of Wisconsin

Richard Palmer, Physics, Duke University

To Learn More . . .

Arthur, B. 1992. On Learning and Adaptation in the Economy. SFI 92-07-038.

Fontana, Walter and Leo W. Buss, "The Arrival of the Fittest: Toward a Theory of Biological Organization," *Bulletin of Mathematical Biology*, in press.

Fontana, Walter and Leo W. Buss, "What would be conserved 'if the tape were run twice?'," *Proceedings of the National Academy of Sciences* submitted (1993).

Holland, John. 1993. Echoing Emergence. SFI 93-04-023.

Hiebeler, David. Implications of Creation. SFI 93-05-025.

Lane, D. 1993. Artificial Worlds and Economics, Parts 1 and 2. Journal of Evolutionary Economics, Vol. 3, no. 2 & 3.

Millonas, Mark. Swarms, Phase Transitions and Collective Intelligence. SFI 93-06-039.

Theraulaz, Guy, and J. L. Deneubourg. Swarm Intelligence in Social Insects and the Emergence of Cultural Swarm Patterns. In The Ethological Roots of Culture, NATO ASI Series.
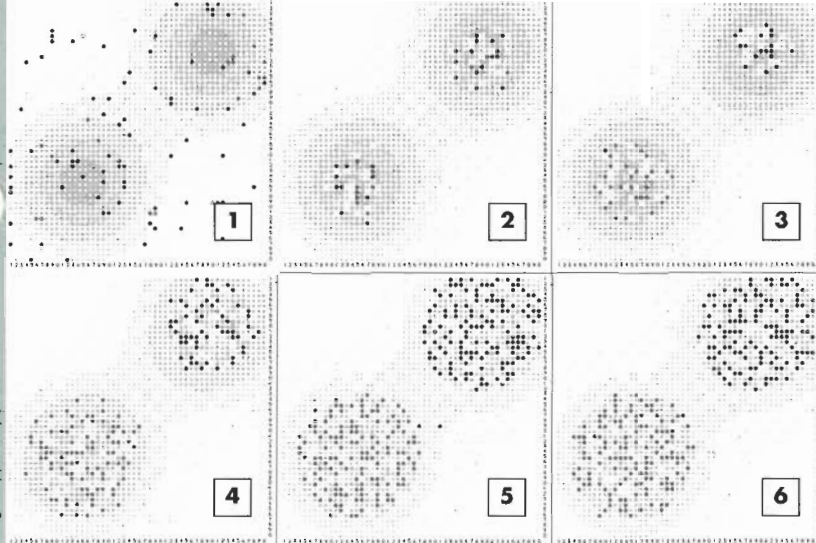
A Proto-History: A Set of Snapshots

1 Agents randomly arranged on the "sugarscape."

2 Goal-seeking behavior of agents leads them to the sugar mountains.

3 On the mountains the agents interact only with neighboring agents and the mountains start to become culturally homogeneous—dark greys in the NE, medium greys in the SW.

4 Cultural homogenization nears completion as the last dark grey agent in the SW is converted to medium grey. The civilizations are completely segregated and culturally distinct.

5 As the societies prosper they undergo population growth, taxing the limited resources available on the mountains. With resources scarce some agents begin foraging in the inter-mountain area.

6 The medium grey forager encounters the dark grey tribe and engages in either economic trade or combat.

SOCIETY ON THE SUGARSCAPE

The unequal distribution of wealth, epidemics, civil violence, famine, war—all are pressing social issues. Can such important collective behaviors be made to emerge from the interaction of individual agents obeying simple local rules? And, if so, how do changes in the local rules alter society as a whole?

Complexity researchers Joshua Epstein and Robert Axtell (Brookings Institution) have developed Artificial Social Life (ASL), a computer simulation that may give insight into these and other social questions. The development is part of Project 2050, a joint venture of the Santa Fe Institute, the Brookings Institution, and the World Resources Institute. "We are confronted with 'already emerged complex systems' such as economies, ecosystems, epidemics, social revolutions, arms races, and wars," comments Epstein. "Can we decode these collective structures and uncover simple local rules that would generate them?"

ASL simulates the interaction of adaptive agents that search for food and can mate, form groups, fight, trade, and transmit cultural attributes (e.g., tastes). Events unfold on an artificial resource topography—a sugarscape. Sugar can be eaten, stored indefinitely, and traded. There are sugar mountains separated by a sugar lowland, all surrounded by sugar badlands. Agents come into the world randomly distributed on the sug-

arscape with randomly distributed traits. Some of these (e.g., vision, metabolism, and sex) are fixed for each agent's life; they're "genetic." Others (e.g., tastes, or group membership) can change through social interaction and are interpreted as "cultural." In the simplest version, there is only the behavioral rule: From all sites within your vision, find the nearest unoccupied site with maximum sugar; go there and eat the sugar. The agent's stomach is then incremented by the value of that sugar, and decremented by her particular metabolic rate. If the result is negative, the agent is dead and is removed from play. Otherwise, the cycle repeats through all the agents.

"How rich is the collective behavior that emerges under this simple local rule?" ask Epstein and Axtell. The basic ecological principle of carrying capacity—that a given resource base cannot support an indefinite number of agents—is immediately evident. If seasons are introduced, migrators and hibernators emerge. But, what about the explicitly social issue of wealth? Epstein explains, "We wondered whether a stationary distribution of wealth (accumulated sugar) would emerge and, specifically, whether the agents would self-organize into a Pareto distribution," one in which a few have a large portion of the wealth and the many have very little. "We were shocked to see an uncanny agreement with Pareto,"

Axtell recalls.

The asserted tendency of market systems to achieve equilibrium—a state in which Supply equals Demand—is being challenged by ASL. Epstein explains that "if individual preferences evolve endogenously, or if information is imperfect, or if space is explicitly represented, the system can maintain itself far from equilibrium indefinitely."

He continued, "We have also grown entire little Proto-histories of society, in which cultural groups—Reds and Blues—emerge from a primordial 'soup,' and migrate to separate sugar peaks, where populations grow, forcing a diffusion back down into the lowland between the sugar mountains, where combat, and cultural assimilation (modeled as tag flipping) perpetually unfold." One such Proto-History is illustrated above

Axtell and Epstein hope to soon provide a version of ASL with easy to use interfaces with which to study epidemics, economics, ecosystems, international relations, and civil violence. The platform will allow scientists to do their own explorations and tailor the tools to their own lines of inquiry, they said. "ASL permits the study of society from the bottom up," Epstein emphasized. "With artificial social life, we can explore social behavior that is dynamic, evolutionary, and locally simple." ○

Image courtesy of Blake LeBaron.
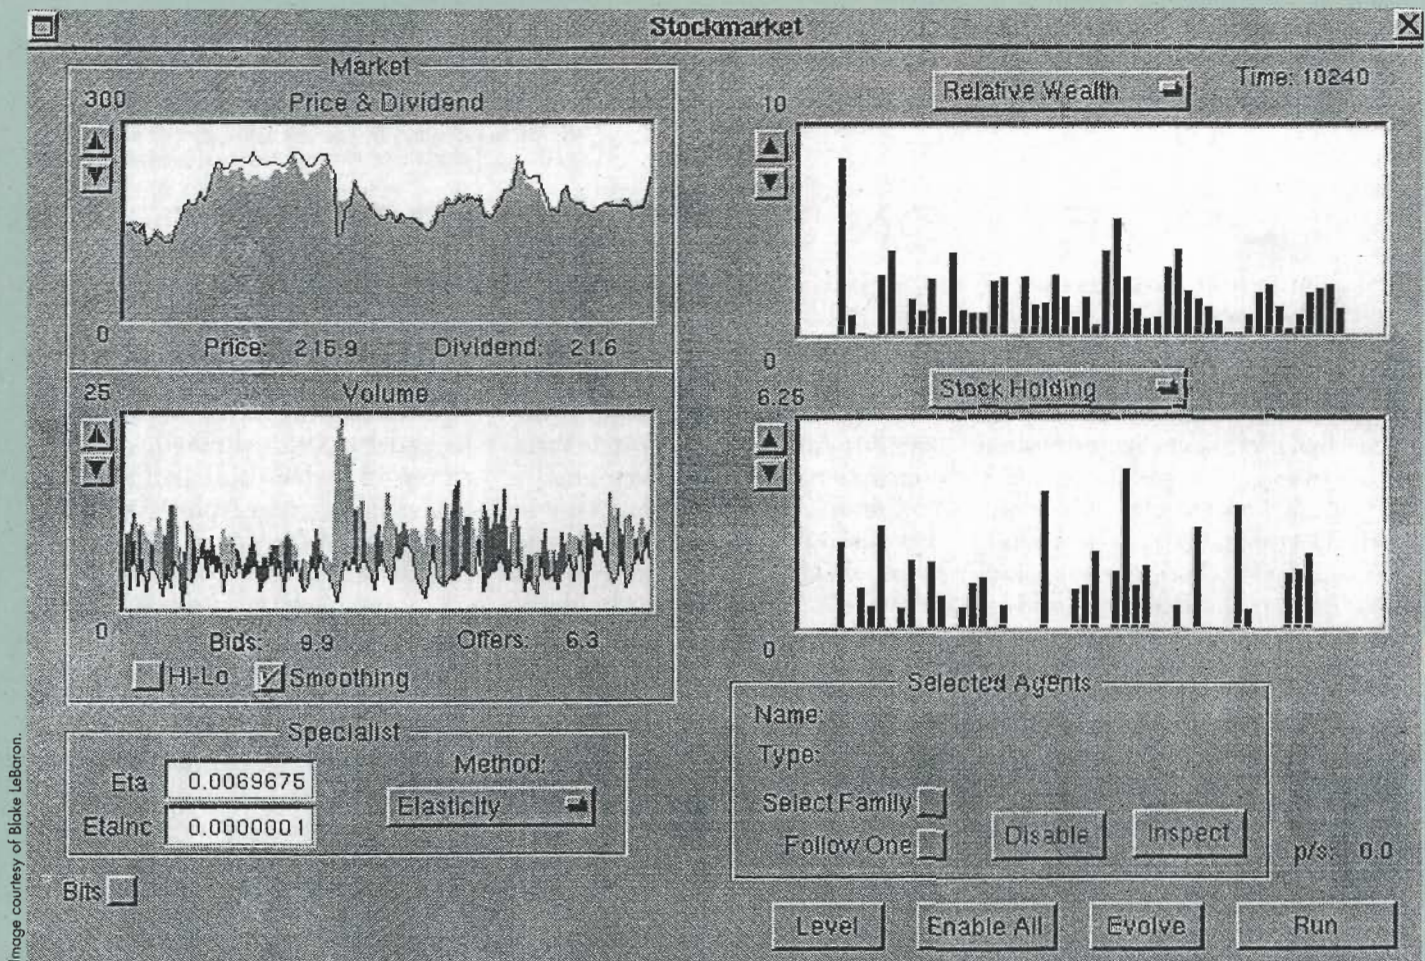
Screen Shot from Simulated Stock Market

Upper left: Actual price of stock.	**Upper right: Current wealth level of traders.**
Bottom left: Number of shares being traded.	**Bottom right: Number of shares held by traders.**

WALL STREET IN SANTA FE

One trader in the stock market is doing exceptionally well with a large volume. As the stock dividends go up, suddenly other traders attempt to buy stock. The brokers attempt to comply, but the orders far exceed the stock available. Some walk away empty handed.

The activity in the market has been busy. And, at the end of the day the brokers will check profits and losses. Some will return the next day, surviving to produce new traders. Others will lose their jobs and disappear.

There are periodic crashes and bubbles, but mostly the market is fairly stable, just like Wall Street. However, the traders in this market won't be taking the subway

home. They exist only in the memory of a NeXT computer—a part of a computer simulation of the stock market.

"This model gives us a view into what learning looks like in a financial setting," says Blake LeBaron. "It replaces the idea of a static, completely efficient market that absorbs all the information with one in which efficiency is more of a struggle. We are looking at that struggle."

The computer simulation of a virtual stock market was developed by LeBaron, Brian Arthur, Richard Palmer, John Holland, and Paul Tayler in hopes of learning why markets behave the way they do and how traders learn to make profits. While past models have been based on deductive reasoning, the economic behavior in this market is based on inductive reasoning, says Palmer. "This model has an inductive pattern recognition approach with competing ideas and hypotheses. If something works, then it gets more weight."

The stock market has a single stock. There are 100 or so agents, who can buy or sell the stock or place money in the bank. The bank pays a fixed interest rate.

"If the stock is going to pay more in dividends and go up, the agents are better off investing in the stock," said Palmer. "But if the stock is going to crash, they are better off investing in the bank. They have to make that judgment."

Similarly to the real stock market, if there is a high demand, the price goes up. If not, then the price comes down.

The agents start off as very dumb—zero intelligence agents—with actions controlled by sets of random rules. Over time they watch the effects of each of their rules, their successes and failures, and start giving the rules that brought success more weight in their decision making. Their learning is reflected in the dynamics of stock prices. Since the traders use the rules to make their forecasts, when a rule makes money for a trader, it's given preference. There's never a perfect rule of behavior and, as they change, the traders learn better how to exploit the way the market behaves, explains Palmer. "They keep changing and evolving." The simulation allows researchers to examine the individual traders and what makes them successful.

Through genetic algorithms they replace rules with more inventive ones or they combine and cross over rules to make a new hypothesis. "They are combining old ideas to make new ideas; keeping the most successful rules; and getting rid of the least successful," notes Palmer. "So we have mutation, recombination, and selection as the way that most of our agents form their hypotheses. It's not a static situation."

"There are moments of instability," says LeBaron. "But it's fairly orderly."

The researchers can change the number of traders, their rules, the money they have to work with, the price of stock, and any number of variables to determine trader and market reaction. The researchers can remove an agent from the market for a period of time, but when the trader is reintroduced he is playing by old rules. "When that agent is put back into the market, it doesn't do well at all," says Palmer. "The simulation has gone on and the other agents have learned more. The reintroduced agent is still adapted to another time and place. What was good then doesn't work now. The whole behavior has moved on, the market has changed, and the agents are more sophisticated."

And, just as in the real market, the less knowledgeable agent will be exploited by the others. "We can put in a mixture of agents; typically one type will do better than another type," says Palmer. "In a mixed market the smart agents can learn to exploit the mistakes of the dumb ones."

Traders that are found to be successful and have the highest profits survive and reproduce with other successful traders. The children will benefit from both parents' knowledge and successes. But because the chromosomes are recombined, the children's trading techniques will be new and different from the parents.

The model includes a variety of graphic tools for studying its behavior. All the characteristics of this artificial economic world can be displayed graphically, including the traders' holdings of shares, their wealth levels, their trading rules, the dividend history, the price history, and the balance of bids and offers.

This simulation of a stock market should give insight into the workings of Wall Street and trading behavior in general, says LeBaron. "This is one of the clearer areas where behavior is easy to study." ○

Linda Little

Mesopotamian scholar Norman Yoffee

considers what archaeologists and complexity theorists

may be able to teach each other and how the result could lead

to new thinking about the evolution of culture.

What *Does* Archaeology Have to Do with Complexity Theory?

(This letter is a synopsis of a chapter to appear in Evolutionary Approaches to Southwestern Prehistory *edited by Murrary Gell-Mann and George Gumerman.)*

Dear Murray:

Since I came to your 1992 conference "Adaptation and Organization in the Prehistoric Southwest" as a minor member of a working group, I was quite surprised when George asked me to write a "commentary" chapter for the volume. Of course, I recognized the syndrome: find an outsider to keep the Southwest experts honest and make them speak to a larger audience than their own brethren and sistren. Outsiders are also thought to be able to identify areas of disciplinary biases that are less visible to the assembly of initiates. I'm writing this memo to you, Murray, because, having been assigned my present role, I've got more than the usual discussant's problems and I think you're just the person to help me. After all, you're an outsider yourself to Southwest archaeology (although you know a lot, according to the insiders) and you're a REAL complexity humanoid in a conference that purports to consider what is "complex" about the Southwest and what such a designation might mean for future research.

My problems concern what I think might be some misunderstandings between the SFI complexity folk (mainly Stu Kauffman, Chris Langton, and perhaps your-
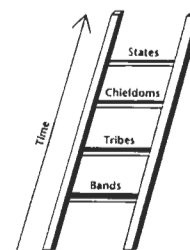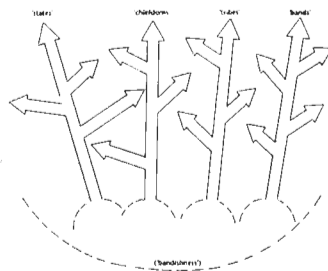
Figure 1. (left) Possible Evolutionary Trajectories

Figure 2. (right) Evolutionary Step-Ladder

self) and the collected Southwest archaeologists; on a more pessimistic note, I'm worried about one of the assumptions of the conference: if Southwest archaeologists can only learn about "complexity" from the SFI experts, they will become better archaeologists. I write this memo, then, presuming that you'll tell me where I'm off base on how the SFI investigates "complexity" and also whether my own notions of complexity in archaeology are of any interest to the gurus at the SFI.

Complexity and social evolutionary theory

The term "complex society" has become extremely fashionable, starting in the 1960s, to describe those socially and economically driven societies in which there is a politically centralized governmental system, namely states, or to those societies that were on their way to becoming states (usually assumed to be "chiefdoms"). One thus spoke of the "evolution of complex societies" or the "evolution of social complexity," while the study of hunter-gatherers and early agricultural societies ("middle-range societies") was marked off into a different domain.

It has been observed, however, that hunter-gatherers never lived in a Rousseauian, egalitarian state-of-nature; indeed, archaeologists now write about some Paleolithic and Neolithic societies as "complex." Furthermore, as archaeological skills in research design and recovery have improved and archaeological data have accumulated massively, especially with new quantitative and high-tech analyses, it is far from simple to interpret any archaeological remains.

"Big" pit-house villages in the Southwest, for example, present complex problems of digging, identifying formation processes, classifying data, and understanding ancient lifeways—of a score or so of people. It seems thus that complexity of modern recovery and analysis have become a shorthand for complexity of ancient social organization.

In the Southwest, furthermore, in the last decade, the classic investigations of modern puebloan societies that were used to model ancient pueblos have been heavily criticized. The understanding in these earlier studies was that the nature of the kinship system more-or-less accounted for the (limited) kinds of economic and political stratification in pueblos. Southwest societies (ancient and modern alike) were not "complex," i.e., states or on the way to becoming states. That model was initially attacked by Cordell and Plog and now it is generally accepted that ancient Southwest societies (especially Chaco, Hohokam, and Casas Grandes) were more "complex" than modern pueblos. However condescending (among other things) to modern puebloan peoples this revisionist view is, the Southwest has clearly joined in the archaeological trend toward "complexity inflation."

I think you can see from this brief discussion, Murray, that Southwest archaeologists brought a lot of baggage with them to the conference on "Complexity in the Prehistoric Southwest" at SFI. If you SFI types were only dimly aware of why some Southwest archaeologists had so much invested in the concept of "complexity," it was equally apparent that the

Southwest archaeologists had only dim notions of what you mean by "complexity."

What you may not realize, however, and what several of the Southwest archaeologists at the conference seemed to have missed, is that those of us who investigate ancient states have been talking your language (but with a different lexicon) in recent years. We're well past the idea that there were "control mechanisms" in ancient states, or that we can draw flow diagrams of neatly boxed social institutions, or that ancient states are homeostatic systems. Rather, ancient states are messes, filled with institutional struggle among classes and ethnic groups and there is conflict within such groups, too. Leaders half-understand what's happening in their societies. Bureaucracies consist of a mixture of patrimonial retainers and those recruited by ability. Finally, in ancient states, long-term stability is sacrificed regularly for short-term gain; against an old image of timeless rule by beneficent kings who are toppled by victorious barbarians or more militaristic neighbors, we now investigate internecine strife, the creation of wealth and misery, factional solidarities, and environmental degradation (in no particular order and sometimes all at once).

However, not only are ancient states and civilizations "complex systems," so are all human societies playgrounds for social negotiation, for the empowerment of the few, and in which the parts remain far from some equilibrium with each other and their environment. In terms of SFI, thus, it is perfectly clear that all human social systems, from Paleolithic hunter-gatherers to states are complex systems.

Toward a theory of relativity of complex social systems

Now that I hope you and I are agreed that Southwest prehistoric societies are (1) complex systems and (2) by that designation we know absolutely nothing more than we did before having declared them "complex," we need to get to the real work of social complexity theory. Are some societies more complex than others? In what ways can we measure complexity? Why are there different kinds of socially complex societies? And, is there a theory of relativity of social complexity that can fall within the kinds of phenomena SFI is interested in?

At the SFI conference (you'll recall), I distributed a draft of a paper entitled "Too Many Chiefs?" from which I wish to extract some points that relate to the above questions. The main purpose of Figure 1 is to offer an alternative to the well-known stage-level, step-ladder model (Figure 2) of social evolution. Figure 1 is meant to imply that not all human societies fall along one continuum, in which lower-rung examples (like modern pueblo societies) have failed to progress to the top of the ladder (as Figure 2 implies). Although I cannot pause here to discuss other implications of the diagram, I do wish to note that it does not claim that there are only four ideal (and epigenetic) types of societies, that social and economic political relations are fixed with a type, and that change within a type must occur in all relations at the same time and in the same direction.

What I do hope the diagram shows is that modern (ethno-graphic and ethnohistoric) "chiefdoms" may not be good models for those societies that preceded states in areas of the world in which states evolved. (Of course, chiefdoms could become states through contact with states, as the arrows in the diagram indicate; contact with big states might also reduce small states to chiefdoms). From our conversations in Santa Fe, I know you can see exactly why Figure 1 should be of interest to SFI theorists of complexity: it implies that there are different trajectories (or pathways) in social evolution and that, by identifying differing initial conditions (in the time of "bandishness"), we might predict (or, more modestly, begin to understand) what distinguishing qualities there are among the various trajectories.

Hallpike's distinction between "complex" and "complicated" provides a good beginning to describe the differences between trajectories to states and other evolutionary trajectories. "Complex societies" (*sensu* Hallpike) are those with various parts that are both interconnected and with significant degrees of economic inequality and subordination; social groups maintain separate identities and political rules and full-time specialists are developed to manage their interaction. "Complicated societies" are those in which social groups looks much like one another and, consequently, the societies are "mixed up" and "involved." Naturally, in the SFI and anthropological views of "complexity," we must amend Hallpike: all human social interactions are "complex" and different evolutionary trajectories simply exhibit different kinds of complexity.

The future of complexity at the SFI

While it's obvious, Murray, that the term "complexity," as a category that has been used in studies of the prehistoric Southwest, must be unpacked, I think it's also clear that the scientists at the SFI can benefit from an intensified program in social evolutionary studies. Archaeologists—trained to identify connectedness in social institutions, to measure patterns in social organizations and how these change over time, to delineate the feedback between "natural" and social environments, and to oppose reductionist formulas that ignore sources of variation—can contribute much to the general studies of complexity at SFI.

Reading the conference papers on Southwest prehistory from my high perch as a Mesopotamianist, it seems apparent that a social theory of relativity is critical and that we are poised at the edge of a rapid transition in our ability to construct it. The dynamic system of self-organized archaeologists exhibits the emergent property that arises from locally interactive schemata in which we can disarticulate kinds of complexities, speak clearly about the need to study initial conditions, and so explain the data in its different outcomes—outcomes that are not themselves ends, but part of historically contingent and continuous processes. Of course, some ideas are fitter than others. Over to you, Murray. ○

Norman Yoffee is Professor and Chair, Dept. of Near Eastern Studies, and Curator of Archaeology, Museum of Anthropology, University of Michigan.

The impact of the CSSS now extends to students
working in a broad range of fields beyond physics and math.
Faculty member David Campbell reflects on the changes in the school from a six-year perspective.

Nonlinear Science Redux:
The 1993 Complex Systems Summer School

When the idea of a "Complex Systems Summer School" was first discussed at the Santa Fe Institute back in 1987, I was among its strongest advocates, so I was both pleased and honored to be asked by Dan Stein to present a set of lectures on an "Introduction to Nonlinear Science" at the first Summer School six years ago. I was also honored—but less pleased—when it appeared that I would be asked to organize the second summer school. Fortunately for me (and for the summer school!), Erica Jen and then in subsequent years Dan Stein graciously and generously stepped forward to take responsibility. Hence when Dan and his co-director Lynn Nadel approached me last fall with the mission to "redo" my lectures for the 1993 Summer School, a sense of indebtedness (guilt?) reinforced my natural enthusiasm for the activity, and I chose to accept.

Admittedly, another factor in my decision was the belief that to "redo" the lectures would simply be nonlinear science "redux": namely, I would dust off the old overhead transparencies and slides, add a couple of new references and a spiffy computer demo or two, and

step once more into the breach. In the (probably apocryphal) phrase attributed to the brilliant but acerbic Swiss physicist, Wolfgang Pauli: "Schoen, aber falsch

> **Whereas the class of '88 was heavily slanted toward physicists and mathematicians, the class of '93 came from a much broader range of fields, including business, economics, biology, and medicine.**

(Beautiful, but false)." My weeklong tenure as a lecturer included several 2 A.M. "transparency finishes," as I endeavored to meet the changed requirements of a group of students with interests and experiences quite different from those in the first school. In the end, though, it was an exhilarating, al-

beit exhausting, experience.

I cannot resist a plug for the scientific perspective presented in the lectures. My thesis was that by focusing in an interdisciplinary manner on the "paradigms" of a new "nonlinear science"—where by paradigms I mean common experimental features of nonlinear phenomena, plus the associated theoretical concepts and analytic and numerical methods—one can obtain insight into problems that have resisted straightforward discipline-oriented approaches. The three paradigms I identified and discussed were (1) deterministic chaos and fractals; (2) solitons and coherent structures; and (3) pattern formation, competition, and selection. In a sense, this view of nonlinear science can be characterized as "complex systems minus adaptation." Although quite radical in the staid environment of the typical university or funding agency, nonlinear science definitely represents the conservative wing at the SFI. Of course, as any conservative would assert, it therefore must form the conceptual bedrock on which the SFI's intellectual agenda is based!

In 1988 I lectured only on chaos and solitons. The subject of

patterns was placed in the very capable hands of Alan Newell (theory) and Guenter Ahlers (experiment). Their resulting sets of lectures, preserved in the 1988 proceedings, are classics. This time I had to cover that additional material myself, and it wasn't easy, especially within the constraints of five lectures. Given the large number of students with medical or biological backgrounds, I de-emphasized the (important) subject of patterns in fluid systems and concentrated instead on some new results—related to the celebrated "Turing patterns"—on patterns in a reaction-diffusion equation model of glycolysis. Judging from the penetrating questions and considerable follow-up from the students, the choice was successful.

Reflecting on the character of the two Summer Schools, I think that the primary change between 1988 and 1993 was in the nature of training, interests, and backgrounds of the students. Whereas the class of '88 was heavily slanted toward physicists and mathematicians, the class of '93 came from a much broader range of fields, including business, economics, biology, and medicine. Although this breadth of interest engendered some problems related to jargon—we all cling to our shibboleths—it also fitted nicely with my expressly interdisciplinary approach to nonlinear phenomena.

As one would expect, the exceptional quality of the students continues to be a major factor in the success of the school. But equally important is the sense of intellectual community. Although the admissions process was the most competitive ever, from the moment the students arrived the atmosphere was open, nurturing, and non-competitive: the sense of sharing and excitement was tangible. The enthusiasm observed in previous schools quickly emerged here, as did the humor and sense of fun. Given the relative inexperience of some of the students with advanced mathematics, there was a clear problem in making some of the more technical concepts clear. We resolved this difficulty in part by setting up a system of "math buddies," pairing individuals with sophisticated mathematical skills

---

Given the relative inexperience of some of the students with advanced mathematics, there was a clear problem in making some of the more technical concepts clear. We resolved this difficulty in part by setting up a system of "math buddies" . . .

---

and advanced knowledge with those who were less experienced. Although there appeared to be some obvious instances in which criteria other than mathematical knowledge played a role in the selection of a "buddy," the program was nonetheless successful and resulted in greatly "leveraged" instruction. It was a real pleasure for me hearing the concepts I introduced in the lectures being re-explained—often more clearly—by the students in their lunch time and afternoon discussions.

Central to a proper understanding of nonlinear science is an appreciation of "experimental mathematics," the use of computer simulations to generate qualitative insights into analytically intractable nonlinear problems. With the help of David Mathews and Scott Yelich, a computer lab was set up at St. John's. I ran several laboratory sessions presenting demonstrations and problems focusing primarily on chaotic dynamics. (An aside to all you hardware vendors out there—the number of workstations available was far fewer than we need for effective instruction. Donations of equipment for future schools would be much appreciated.) The limited number of workstations forced the students to work in large groups, but again they made a virtue of necessity by establishing rules that the team member with greatest computer anxiety had to do most of the keyboarding and that true "hackers" were banned from touching the keyboard during the formal labs. This last restriction led to some amusing scenes as frustrated experts struggled to coax, rather than intimidate, their less experienced peers into entering the correct sequence of commands for a particular problem.

From my perspective, the most rewarding aspect of the school was watching the students' efforts to incorporate the new concepts of nonlinear science and complexity

into their own disciplinary frameworks, seeking examples of the paradigms and reinterpreting and extending their prior understanding. This is an essential part of the "intellectual technology transfer" that must take place if nonlinear science and complexity are to become recognized subjects in their own right. Among the students at the school, there were many outstanding examples of this process of "integrating" complexity; let me mention just a few.

Dr. Maureane Hoffman, Chief of the Hematology Laboratory at Duke University Medical Center, recognized from the lecture discussion of chemical oscillations, activator-inhibitor mechanisms, and reaction-diffusion equations that it might be possible to construct a detailed phenomenological model of her novel laboratory results on *in vivo* blood clotting. After several discussions, it appears that this is a modeling project on which substantial progress can be made in the short term and which, if successful, could have a tremendous long-term impact on the field, not least by reducing the number of live animal experiments needed to verify certain hypotheses.

Another medical problem, discussed by summer school participant Dr. Carrie Merkle of the Department of Physiology in Arizona, involved blood flow in (minute) renal capillaries. Given the non-Newtonian nature of the flow and the complications of the deformable capillary, it seems likely that at the moment one can only look with wonder at beautiful complexity, in the colloquial sense, of the sudden switching of the direction of flow in the "h"-crosslinks between capillaries. Nonetheless,

the intriguing connection with (possibly) chaotic temporal dynamics suggests a more careful look.

From the general discussion of testing limiting cases and asymptotic behavior, Andreas Wagner of the Biology Department in Yale recognized that a particular hypothesis he wished to test concerning a complicated biological network could, in fact, be tested by reducing the model to its simplest case.

Finally, Afshin Goodarzi from NYNEX sought to incorporate the latest developments in optimization and search algorithms into the enormously complex and technologically vital problems connected with static and dynamic resource allocation in the telephone network.

All in all, "nonlinear science redux" was for me an enjoyable, instructive, and valuable experience. My sole complaint was that after the week of my lectures, other responsibilities kept me from attending the lectures and from continuing to share the excitement and thrill of discovery that have become the hallmarks of the Summer School on Complex Systems. ○

David K. Campbell is Head, Department of Physics, University of Illinois, Urbana-Champaign.

Summer School Fact Sheet

Lectures

David Campbell, Physics, University of Illinois, Urbana-Champaign
Introduction to Nonlinear Dynamics

Sue Coppersmith, AT&T Bell Laboratories
Complex Structures and Dynamics in Condensed Matter Systems

Catherine Macken, Theoretical Biology and Biophysics, Los Alamos National Laboratory
Adaptive Evolution on a Rugged Landscape

Jay McClelland, Psychology, Carnegie-Mellon University
Connectionist Models of Cognition and Learning

Cris Moore, Santa Fe Institute
The Computation Paradigm of Complexity: An Introduction to Complexity Theory and Applications

Olaf Sporns, Neurosciences Institute, Rockefeller University
Neural Models of the Visual Cortex and Perception

Randall Tagg, Physics, University of Colorado, Denver
Instability and the Origin of Complexity in Fluid Flow

Matthew Wilson, Arizona Research Laboratories (Neural Systems, Memory and Aging Division)
From Network Models to Network Recordings

Summer School Fact Sheet

Support for the 1993 School

General support from the Center for Nonlinear Studies at Los Alamos National Laboratory, Department of Energy, National Institute of Mental Health, National Science Foundation, Office of Naval Research, Santa Fe Institute, and the University of California system.

Student support from Aexpert, Mitsubishi Electrical Corporation, NYNEX, Stanford University, and the Wharton School of Business

The Participants

44 graduate students, 8 post-doctoral fellows, 8 senior researchers, and 2 undergraduates representing the fields of Astronomy, Biology, Biomathematics, Biomedical Modeling, Chemistry, City and Regional Planning, Computer Science, Decision Sciences, Economics, Engineering, English, Finance, Genetics, Mathematics, Neuroscience, and Physics

To Apply to Next Year's School

Look for a call for applications in national science journals in November, 1993. Deadline for application is February, 1994. Applicants are expected to provide a current resumé along with a statement of current research interests and recommendations from two scientists.

"Titania's Tiara" ◆ Recursive art by Rollo Silver ◆ 505-586-1607
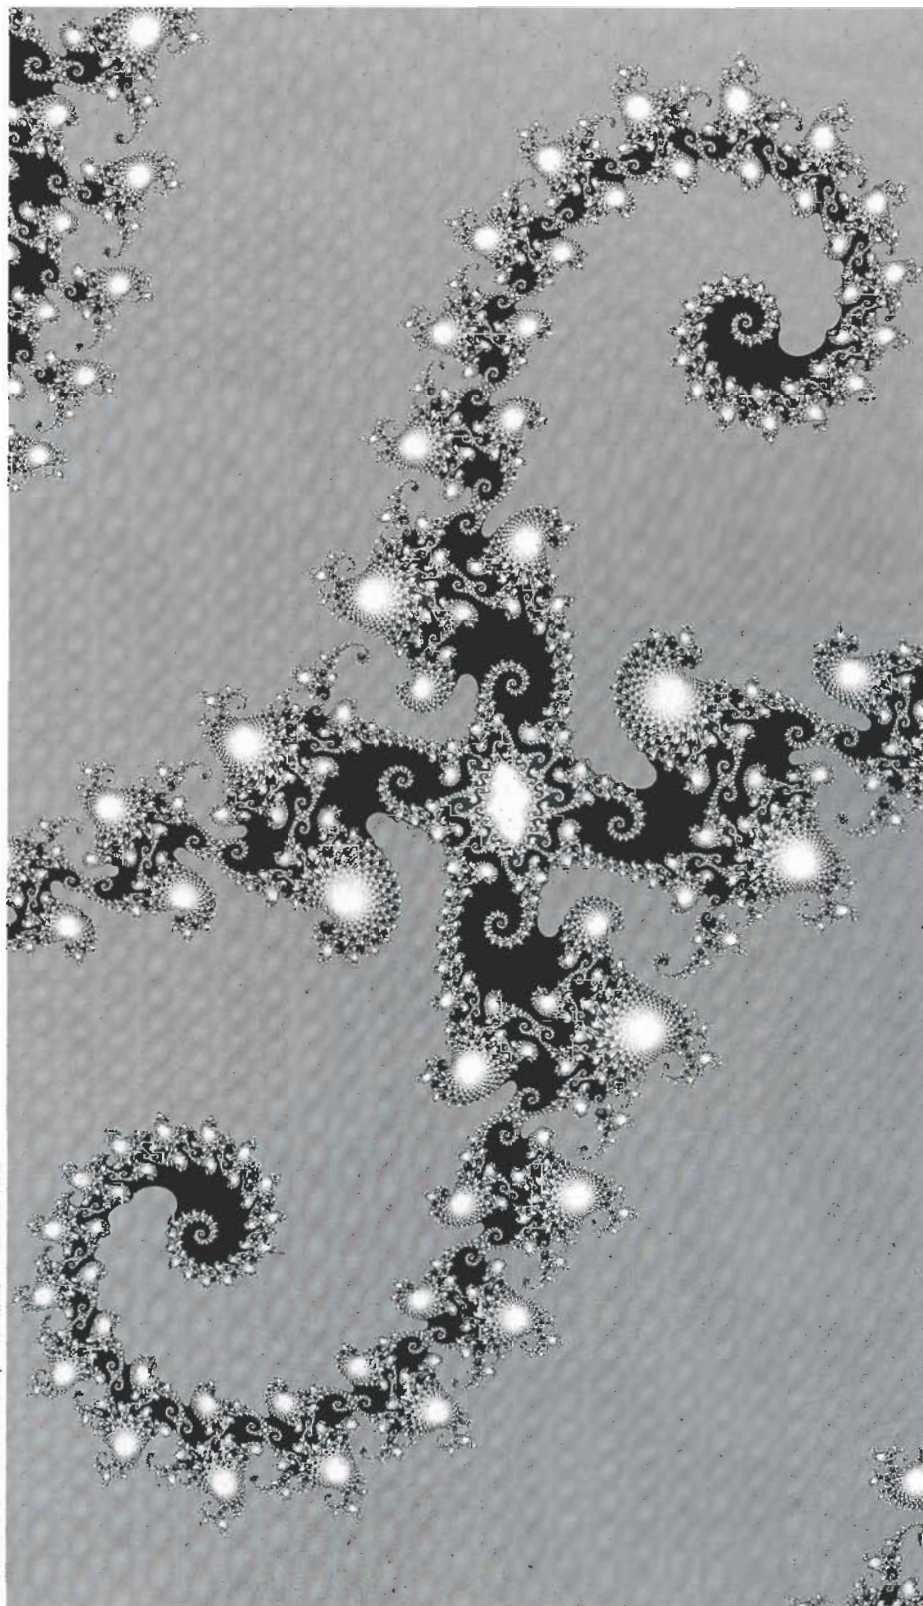
Marc Feldman

As Marcus Feldman articulates some theoretical aspects of biology—from population genetics, to demography, to modeling cultural evolution—it's genuinely hard to believe he didn't grow up with a biology book in tow, or even a budding glimmer of interest in the subject. He hadn't, in fact, taken a single biology course until well-immersed in a Ph.D. program. But biology was apparently very much in Feldman's destiny, despite his rooted affinity for mathematics in its purer form.

Born in Perth, Australia, Feldman earned his undergraduate degree in mathematics and statistics, and immediately joined IBM's newly opened Australian division as a systems analyst. It wasn't long, however, before his appetite for more education grew too hearty to dismiss; he was coaxed into quitting his post at IBM and taking a "75-percent drop in pay" to pursue a full-time master's fellowship in mathematics at The University of Melbourne.

It was there that biology began its insidious tumble into Feldman's life. As a freshly picked master's student, his thesis was selected for him. The topic? An application of probability theory to genetics. When Feldman completed his master's degree, he was urged by the professor who assigned the thesis to submit it for publication. Feldman did, to an editor named Samuel Karlin who happened to reside at Stanford University. Karlin had just begun studying the subject of applications of probability theory in genetics in collaboration with Sir Walter Bodmer—who today heads the Imperial Cancer Research Foundation, but at the time was a fledgling assistant professor in the genetics department. They were intrigued by Feldman's thesis.

"I was lucky to have both of them interested in what I was doing," Feldman says with modesty. At Karlin's invitation, he traveled to Stanford to begin a Ph.D. in mathematics.

Nearly two years into his doctoral program, Feldman attended a biology lecture "by accident." He was so taken with the subject that he altered his direction, opting to pursue a Ph.D. in mathematical biology. "I was exceptionally fortunate that Sam Karlin had such a wonderfully interdisciplinary perspective on science," Feldman recalls. With his doctorate, Feldman returned to Australia to become a lecturer in mathematics.

Two years later in 1971, "The Stanford biology department made me an offer I couldn't refuse," says Feldman. He's been in the biology department ever since. As destiny would have it, he has become increasingly drawn to the theoretical aspects of biology, especially as they pertain to population genetics and models of cultural evolution.

"I really like the idea of taking a biological problem," he notes, "and reconstructing it into a way that is amenable to modeling analysis."

Population Genetics and Cultural Evolution

As an emerging discipline, population genetics focused on the examination of experimental populations with genetic variants that change over time, and the application of mathematical or statistical models to predict those changes. Calculating the frequency of change in genetic variability is still a key ingredient of population genetics, but the discipline has grown to encompass a much broader range of topics. Among them: the causes for genetic variations (including within DNA itself) between individuals in a population and between populations, and within and between racial groups and subgroups; distinguishing between environmental and genetic variations; taxonomic ordering based on genetic differences; mathematical properties of such universal genetic operations as recombination and mutation; etc. Each topic, in turn, can be repeated for different species.

Regardless of the area of emphasis, the study of population genetics is based on the general premise that variation exists. And variation is transmitted from individual to individual either through genes—typically parent to offspring—or through what could be called epidemiological pathways—namely, non-genetic sources, such as learning from another member of the population.

This premise led Feldman and Stanford colleague Luigi Luca Cavalli-Sforza (a geneticist who arrived at the university at roughly the same time as Feldman) to initi-

ate some inquiries of their own. "We hit it off right away," says Feldman, and it wasn't long before the partners channeled their curiosity in one particular area: how to quantify changes in traits that are learned, using analogies to how one quantifies changes in traits that are acquired genetically.

Feldman and Cavalli-Sforza invented the "quantitative theory of cultural evolution," an approach to modeling the evolution of learned behaviors in cultures using mathematical applications. The pair have applied this theory to questions of the same sort that bewitch scientists in the field of evolutionary genetics. Questions like: how much variation can one expect to find within cultural subgroups or between cultural subgroups? What rules of transmission of traits lead to fast change or slow change? How does the rate of change affect observed variation?

In applying the theory practically, the partners narrow its scope using an approach similar to that employed by Gregor Mendel as he uncovered the basic principles of hereditary. Throughout Mendel's work in the field of genetics, he consistently focused on some of the tiniest particles that are transmitted biologically—genes. "In Mendel's time, people didn't know how things were transmitted from individual to individual," says Feldman. "They treated blue eyes as if they were transmitted from parent to offspring, but they didn't have any rules. Mendel discovered the basic rule of [genetic] transmission. This is how we try to treat the subject of cultural evolution."

To examine the myriad ways non-genetic (i.e., learned) traits are transmitted within a culture,

Feldman and Cavalli-Sforza begin by identifying and sorting them into behavioral components, or atoms. In culture, an atom can refer to a behavioral pattern of a particular group or to a specific trait of an individual's behavior—for instance, how a person pronounces words, depending on his birthplace or whether he watches television excessively. The pair studies these atoms to devise "sensible" ways of describing their rules of transmission among individuals.

In a cultural setting, non-genetic atoms can be transmitted in numerous ways, such as between work associates, via non-parental influencers such as teachers or other appointed mentors, or through such vehicles as books and electronic media. The cultural evolution model can be applied also to quantify changes in learned behaviors, and how variation of the changes and other conditions such as genetic predisposition can affect the evolution of a culture.

Applying the Theory of Cultural Evolution

Feldman and Cavalli-Sforza have applied their cultural evolution theory to an assortment of anthropological groups and situations. An area of emphasis has been in using the model to determine how cultural and genetic traits are related, and how the interrelationship affects evolution of both traits, as well as of the culture itself. Feldman offers two examples: dairying and sign language.

Many indigenous peoples use the milk of cows, sheep, or goats—a behavioral trait called dairying—while other people with access to these animals do not. Feldman and Cavalli-Sforza have applied

their theory to determine how the cultural atom of dairying relates to the genetic atom, or gene, responsible for a human's ability to digest milk. This has been an opportunity for the pair to study two separate, but interdependent, evolutionary processes: the learned cultural pattern of dairying, and the biological activity of individual genes whose evolution correlates with the cultural pattern.

The cultural evolution model has also been used to describe how the communication system of sign language within cultures or subcultures has and will continue to evolve. "A high proportion of deafness is hereditary," says Feldman. "So you have evolution of a cultural trait—changes in signing—that is restricted to a genetic minority." Feldman, Cavalli-Sforza, and their colleagues have written a book and several articles on this theme, using their quantitative theory to explain how sign language might evolve in a given culture.

Recombination, Algorithms, and SFI

Feldman dedicates roughly 50 percent of his working time to projects related to the cultural evolution model. The remainder of his time is spent predominantly on "classical" population genetics, which involves modeling the processes that occur within genes, such as mutation and recombination. Feldman's special interest is studying many genes simultaneously, particularly how genes along a chromosomal line break off and recombine to form new chromosomes.

Before Feldman became affiliated with the Santa Fe Institute, he had used computing as an occa-

sional tool for studying evolution; however, it wasn't until he started collaborating with associates at the Santa Fe Institute that his interest in computational models of evolution was ignited. "I started thinking of the subject in terms of algorithms and other things that computer scientists are concerned with," he says.

Genetic algorithms are computational models of evolution based on the principles of genetic variation. Ideally, the algorithms are used to identify solutions to problems by "evolving" a population of individuals via computational instructions. In the computer program, individuals are represented as bit strings (collections of ones and zeros) corresponding to chromosomes, and each bit on the string corresponds to an individual gene.

"If you have two such lines of computer instruction, they look like a pair of chromosomes. As with the genes on a chromosome, the bits on a piece of a computer program can be recombined . . . That is how computer algorithms are applied to biology. Likewise, it's how computer scientists apply the biological metaphor to computing. If combining the bits of a computer program [recombination] produces a better instruction, then it can be favored and the other one thrown away. As in the biological metaphor, there is natural selection."

Feldman's increasing area of interest is in utilizing computer analogies to study the mathematics of evolutionary theory in biology. He and his colleagues have already applied the mathematical theory to several practical areas of biology. For example, they've used it to study in humans the scores of

genes related to the major histocompatibility system, as part of the immune system. This block of genes is contained on one pair of chromosomes that has "hot spots" where recombination is frequent, and Feldman has used genetic algorithms to attempt to explain how these chromosomes maintain their variability.

Another practical application: the mathematical theory has been used to study evolution of the Rhesus (Rh) blood group system, specifically the existence and perpetuation of Rh-negative/positive genetic combinations that jeopardize the lives of many babies. If an Rh-negative female mates with an Rh-positive male, the probability is high that the pair's Rh system genes will combine to produce an Rh-positive baby; the condition of a "positive" fetus in a "negative" environment induces an immunological response in the mother; i.e., antibodies are produced that seriously jeopardize the baby. Historically this condition has resulted in countless "blue baby" births requiring immediate, complete blood transfusions, and a good many of these infants died. Today drugs that suppress production of the antibody can be administered to Rh-negative pregnant women, but the potential harm is still very real for those who do not receive the protective injection.

"What's interesting about the Rh blood group system is that it doesn't appear to be very adaptive in an evolutionary sense," notes Feldman. "Why shouldn't everybody evolve to be Rh positive since the presence of both factors [negative and positive] has caused a lot of children to die—and more will die?"

Feldman and his colleagues

have studied the maintenance of dangerous Rh combinations in populations from different perspectives. They've theorized, for instance, that these combinations persist to allow for reproductive compensation among humans—compensating for the loss of one child by reproducing another child. Other species, such as birds, "compensate" in this way; and many female mammals immediately return to estrus upon the death of an offspring.

The frequency of the Rh-negative factor in other populations has also been examined. As Feldman explains, there is only one population in the world—the very old, genetically distinct Basques—that has a higher Rh-negative frequency than Rh positive. This proportion leads evolutionists to speculate whether at one time the frequency of Rh negative was higher than the frequency of Rh positive and, if so, what happened to alter the frequency. Using the Basque population as a model, there exist several conditions, and combinations of conditions, that may explain the shift. It could be, says Feldman, migration of other populations into the Basque population, or it could be the impact of invading peoples from Asia or Eastern Europe who left genetic remnants that became the Basques.

Or this evolutionary pattern can be examined from both genetic and cultural perspectives. It may be that studying cultural factors such as the highly preserved, distinct Basque language can reveal some of the answers. By studying both their genes and linguistic traits (i.e., by applying the cultural evolution model), evolutionists may be able to obtain much more information than if they focused

their efforts solely on one of these areas.

Feldman is quick to note that applying mathematical theories to evolution shouldn't be viewed as a broad stroke or an end-all. "Every one of the questions we've raised can be presented in a quantitative way, but it's important that one is clear and doesn't make too many big claims about what he's doing. We're not solving the world's problems. We're contributing to qualitative knowledge about a subject that is too often left in vague or subjective terms."

Roots at the Santa Fe Institute

Nearly eight years ago, Feldman was asked by Murray Gell-Mann, one of SFI's founders, to visit the institute. "That was a really big high for me because the people were very interested in what I was doing," says Feldman. It was at SFI where he met computer scientist John Holland; soon after, the pair began pondering how their distinct approach to the mathematical study of evolution could be "recombined." (Holland, in fact, invented the genetic algorithms described earlier.)

"The founders (of SFI) were insightful in picking people who were naturally interdisciplinary in their inclinations," Feldman says. "They're not intolerant of other disciplines and the way they do things. The result has been very good cross-fertilization of disciplines."

Feldman notes, on the other hand, that this major strength of the institute—the diversity, intellect, and tolerance among the people it attracts—has fostered a weakness as well. "The Santa Fe Institute has attracted people who are very well settled in their universities," he says. "They want to participate in the intellectual stimulation and then take it back to their own institutions." Feldman's observation raises the question of whether SFI should maintain its present structure with very few resident scientists or build itself into an institute with a staff of resident, permanent scientists. As he points out, however, there aren't many choices when money is in limited supply.

"The chairman of the Board of Trustees at SFI, Jim Pelkey, has done a phenomenal job in not only bringing new money to the institute, but in bringing the institute to the attention of high-tech industries," according to Feldman. "People from those industries are starting to be interested in the institute. Pelkey has brought in people who have the advantage of not wanting to convert SFI into a subsidiary of their company; they're appreciating that it's important for the institute to maintain its independence and uniqueness."

Feldman visits SFI about eight times a year. In addition to serving on several boards and committees, he has taught in the summer school, given public lectures, and helped run workshops. Feldman also has a deep personal and professional interest in Project 2050, an enormous task devoted to developing rough models for what can be done to make a sustainable world by 2050 in terms of resources and development.

"It's a big endeavor that may not succeed, but the point about it is that at the Santa Fe Institute, you learn from problems that are soluble and insoluble," says Feldman. "It's impossible not to learn there. The kind of people that have been attracted to the institute are very open-minded about new approaches to old problems and old approaches to new problems."

Free Time?

When Feldman isn't in a classroom, behind a desk, sitting on a committee, conducting research, or visiting the Santa Fe Institute, he enjoys the company of family, athletics, and books of many genres. And Feldman's fascination with the lives and patterns of other cultures extends into his home as well. A collector of tribal art, Feldman has filled his house with aboriginal, New Guinea, and African artifacts. Some Santa Fe art galleries have excellent collections of tribal art from these areas, he notes. He's enchanted with this art form because, as he says, "it's so related to the daily lives of these people, and to their views of relationships with other beings— whether the beings have a concrete form or are imaginary, whether the beings are above or below them."

Athletics, too, is in Feldman's blood. Throughout his boyhood and college days, he devoted huge portions of time to cricket and soccer. In fact, he excelled at cricket and competed professionally, representing his state of Western Australia. In the San Francisco Bay Area, Feldman coached children's soccer for many years, and he passed his fondness for the sport on to his three sons. "My boys are very successful soccer players, and they coach soccer now, too. It's very culturally transmitted," Feldman says with a wide grin. ○

Jennifer McAllister is a freelance writer based in the San Francisco Bay Area.

Unfolding the Problem of Protein Folding
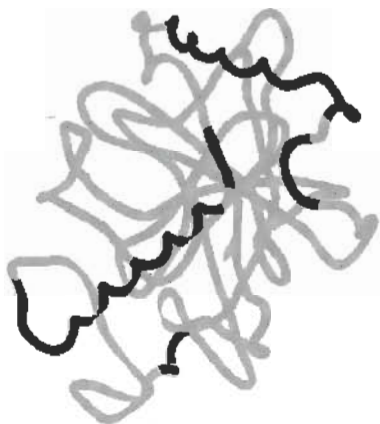
The new secondary structure class of protein (here illustrated in black) is yeast inorganic pyrophosphatase. This classification of protein secondary structure is a preliminary result of a new algorithm utilizing co-evolving, adaptive networks to define highly predictable classes of secondary structure.
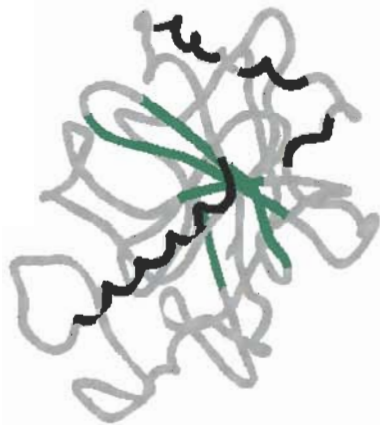
It is visually clear that the new "Xelix" class of secondary structure bears some relation to the conventional "alpha-helix" (illustrated here) class of secondary structure, but with the key distinction that it is significantly more predictable from amino acid sequence than are conventional alpha-helices.

Images appear by permission of Robert Farber.

SFI's current work on genetic data analysis explores protein structure and folding. Most genes serve as blueprints for the formation of proteins: if a gene incorrectly directs how a protein is made, it may threaten normal development and health. Genes determine protein structure by coding for the sequence of amino acids in the protein; this sequence in turn causes the protein to fold into its particular three-dimensional shape. Determining the structure of a protein is far more difficult than figuring out its sequence, so understanding the relationship between sequence and structure is especially important.

Some local structural "motifs," called secondary structure elements, reoccur in many different protein structures. Among them are alpha helices and beta sheets. Accurately predicting the location of secondary structure from amino acid sequence is thought to be a big step toward predicting full three-dimensional protein structure. To this end, Alan Lapedes, Evan Steeg, and Robert Farber are using coevolving adaptive networks to create new arrangements of protein secondary structure that are highly predictable from amino acid sequence. They have developed an algorithm that uses adaptive networks to simultaneously examine both sequence and structure data (such as that available from the Protein Data Base) to determine new secondary structure classes

that can be accurately predicted from sequences. A manuscript has been submitted to the Fifth Annual Neural Information Processing Systems Conference (December, 1993); the biological implications are outlined in a paper in preparation for the *Journal of Molecular Biology*.

Another SFI project attacks the problem of "inverse folding" in proteins. The classic approach to the protein-folding puzzle is to begin with a given sequence of amino acids and then search for the structure of the folded protein. The inverse folding approach begins with a given shape and then asks what amino acid sequence will fold into that shape. Alan Lapedes, Melanie Mitchell, Joseph Bryngelson, Jeff Inman, Jeffry Slotnick, and Adam Godzik are collaborating to develop a variety of algorithms—including "greedy" and genetic rules—to unravel this problem.

An energy function for a protein results from both amino acid sequence and protein shape; if the shape is fixed, then the energy becomes purely a function of the sequence. Isolating the sequence space makes it possible to search for the arrangement that has the lowest energy when threaded through the desired shape. In most situations, physical systems prefer states that correspond to a minimum of the energy. But actually, the lowest energy sequence is not quite the answer. If that same sequence has a similar ener-

gy on a different shape, then there is no particular reason for it to fold to the shape of interest. What is really needed is a large energy gap between the sequence when it is folded on the shape of interest versus other protein shapes.

Since scientists suspect that the present library of known protein shapes comes very close to exhausting the total number in nature, project members search sequence space looking for the biggest energy gap between the arrangement on the shape of interest and that on the known shape with the next lowest stable energy. Right now they are concentrating on evolving "ultimate" hemoglobin sequences. These sequences have a greater energy gap than that between native hemoglobin and any other known protein shape.

In a related project Joe Bryngelson is turning another standard approach inside out. Methods for predicting protein substructure typically seek structures that minimize some approximate energy function, given that this energy is a function of the three-dimensional structure of the protein. Most efforts in tertiary structure prediction have gone into creating approximate energy functions and especially into discovering algorithms for minimizing these functions. Bryngelson addresses a complementary issue: When is an energy function sufficiently accurate for protein structure prediction? Bryngelson has answered this question for a simple model of proteins and the manuscript has been submitted to the *Journal of Chemical Physics*. Meanwhile work at SFI is focusing on extending this idea to more realistic, complex protein models. ○

Believing Your Ears

Late last year ICAD '92—the First International Conference on Auditory Display—took place at the Santa Fe Institute. Plans are currently underway for another ICAD, to be convened in November, 1994. In the meantime the audification research network is working on the auditory representation of global dynamics in ECHO.

The ability to hear many sounds at once, and to hear one sound change in many ways simultaneously, may help scientists present complex information from high-dimensional systems. To date, computer interfaces have focused almost exclusively on vision as the information conduit. Researchers within the Audification project are extending this conduit with new uses of sound.

The major topics of ICAD '92 were sonification, audification, and auditory displays. Simply put, sonification is "data-controlled sound." It is the auditory equivalent of data visualization, the technique of "looking at data" to help analysts comprehend everything from weather information to financial or medical data. Audification is rendering audible such data as seismograms, radio telescope information, mechanical simulations, or equation-generated waveforms. Auditory displays are the sound aspect of general computer user interfaces, such as the icons and menu bars found in most computer software.

Participants described work on using sound to discern seismic events, such as earthquakes; diagnosing supercomputer software; looking for trends in financial data;

designing computer interfaces and chemistry lab equipment for vision-impaired persons; monitoring patients under anesthesia; understanding chaotic systems; reviewing census data; and using sound in virtual reality as well as other application areas. There was discussion of the pattern recognition capabilities of the human auditory system, long ignored (with the notable exception of sonar) in information presentation.

The proceedings of ICAD '92 are due to be published by Addison-Wesley in late 1993, as part of SFI's Studies in the Sciences of Complexity Series and will be the first book ever dedicated to the field of auditory display. According to Kramer, editor of the volume, "What little research has been published had to be culled from journals in fields as far ranging as chemistry, statistics, computer-human interface, music, and optical engineering. This proceedings brings this work together for the first time, and we hope it will help focus and energize research in auditory display."

ICAD '94 will be a single-track conference. Attendance is open to all, with no membership or affiliation requirements. For an automated response with general information, contact icad94@santafe.edu.

ECHO is a software system conceived by John Holland which provides for the study of populations of evolving, reproducing agents distributed over a geography with different inputs of renewable resources at various sites. (See the description in "Computation Platforms" in this issue.) The large number of possi-

ble interactions and resulting agents possible in ECHO make it very difficult to monitor and understand the system's dynamics. Of course, some of the system variables can be represented using graphical display techniques; however, both the global nature of the variables and the relative frequency of agent interactions in ECHO present formidable problems for graphic display. Auditory display offers a possible solution. Workers are focusing their initial efforts on developing a display that informs the system-user of the status of nine distinct variables, specifically the volume of trading, mating, and combat interactions of three agent species. ○
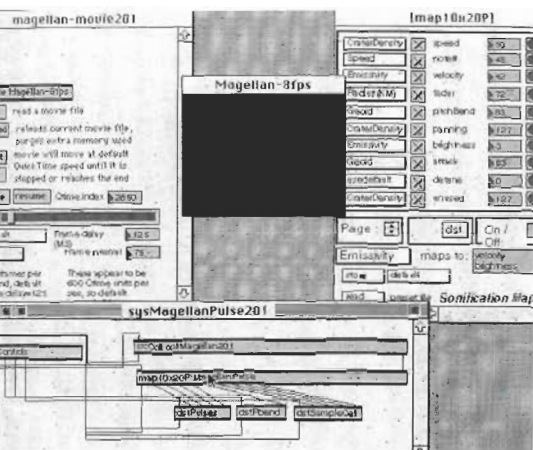
Screen shot of Gregory Krammer's sonification tool kit.

Adaptive Computation

SFI research in adaptive computation concentrates both on building computational models of adaptive systems and on using novel computational methods inspired by natural adaptive systems for solving practical problems. Work is proceeding on a number of fronts. John Holland, Stephanie Forrest, and Terry Jones—who is in residence at SFI full-time—are continuing work on the ECHO system, a simulation system for modeling ecological phenomena. Jones has ported ECHO from the Macintosh system to a Unix X-windows system, and has developed a detailed user interface. He, Forrest, and Holland are now approaching the problem of modeling speciation in ECHO. They have plans to collaborate with Jim Brown and other ecologists at the University of New Mexico on modeling data from real ecologies, such as data on bird extinctions. These projects will contribute to assessments of the biological plausibility of the system. (For further discussion of ECHO see "Computational Platforms: Setting the Stage for Simulation" and the reports on Audification and Project 2050 in this issue.)

Chris Langton's Swarm project is well underway as described elsewhere in this issue. Swarm is meant to be a general-purpose simulation tool for complex systems. Langton and David Hiebeler have implemented an initial version of Swarm, and it is now at the point where it is ready for applications. Nelson Minar, an undergraduate intern from Reed College, is working with Langton, Hiebeler, and Walter Fontana to implement Fontana's "ALchemy" system in Swarm. Howard Gutowitz visited SFI later during the summer to work with David Lane on implementing economic models, including a model of "information contagion," in Swarm. Several anthropologists working with SFI (including Jonathan Haas, George Gummerman, and Tim Kohler) have plans to use Swarm in their research as well.

Stephanie Forrest and Melanie Mitchell's work on foundations of genetic algorithms (GAs) continues. Melanie is staying on as Resident Director of the Adaptive Computation program through August, 1994 so the two will continue to have the opportunity to collaborate locally. They are extending their work on "royal road" functions—problems specially designed to study GAs in depth. The goal of this work is to understand in detail the workings of GAs and to characterize the class of problems on which they will work well. Forrest and Mitchell are also working with University of New Mexico graduate student Tim Preston on using GAs to study simple models of the evolution of recombination. The goal is to learn more about the class of environments in which a capacity for recombination will be selected by an evolutionary process, and how the evolutionary viability of recombination is related to its ability to improve fitness in a population.

In a related project Melanie Mitchell, Jim Crutchfield, and Peter Hraber are applying GAs to evolve desired behavior in cellular automata (CA). Their current

work concentrates on understanding the mechanisms by which the GA can produce complex computational behavior in CA and the impediments faced by the GA in attempts to do so. Matthew Headrick, an undergraduate intern from Princeton, joined their project for the summer.

A number of visiting researchers have joined the program this spring and summer:

Pentti Kanerva, from NASA's Research Institute for Advanced Computer Science, visited the Institute in February and worked on a proposal for extending his Sparse Distributed Memory model to larger-scale tasks such as real-world robotic navigation.

David Goldberg, from the General Engineering Department at the University of Illinois, visited SFI in March to work with Stephanie Forrest, Melanie Mitchell, and Terry Jones on foundations of genetic algorithms.

Una-May O'Reilly, a graduate student from Carlton University in Ottawa, spent two months at SFI working on aspects of genetic programming—a genetic algorithm technique developed by John Koza to evolve Lisp programs. She plans to return in the fall continue this research and begin writing her dissertation on this topic.

Richard Palmer participated in the economics and AC programs, working both on a stock market modeling project and the theory of genetic algorithms. He started a working group "Evolution on Landscapes," which studied evolutionary processes on fitness landscapes from the point of view of a number of disciplines.

Gregory Rawlins, from the Computer Science Department of Indiana University, visited SFI in August. He worked with a number of people in the AC program on aspects of genetic algorithms and evolutionary robotics.

This spring the program sponsored two working groups related to learning and adaptation in artificial systems. Both were supported by funds from the National Science Foundation. "Reinforcement Learning in Robotics: The Challenge of Scaling Up," was organized by Nils Nilsson of Stanford University and Melanie Mitchell. Reinforcement learning is an approach to machine learning in which learning agents act in an environment and are intermittently given reinforcements for certain actions. This approach has in recent years emerged as a central area of machine learning, but has not yet achieved success on large-scale learning problems, such as learning in real robots in complex environments. This small working group brought together some of the most prominent researchers in this field to discuss what is necessary to "scale up" reinforcement learning techniques to larger-scale problems. Progress was made on isolating some central issues, especially those related to representation and credit assignment, and on means for addressing these issues. There was also much discussion on what are appropriate problems for reinforcement learning.

"Learning and Adaptation in Robots and Situated Agents" brought together people working on "bottom up" approaches to learning and adaptation, especially in the context of robots and other artificial agents. These approaches include neural networks, classifier systems, subsumption architectures, behavior networks, "schema mechanisms," certain reinforcement learning techniques, and others. It brought together a number of central people from each of these areas. There was much discussion and new insight into how these various approaches compare with each other, what the strengths and weaknesses of the various approaches are, and how to make progress in developing the different approaches.

In July, the AC program sponsored "Plastic Individuals in Evolving Populations: Models and Algorithms," a meeting organized by Rik Belew (UC San Diego) and Melanie Mitchell. The purpose of this workshop was to bring together a number of leading researchers in various fields, all studying aspects of the interaction between learning (or more generally plasticity) and evolution. Some are computational researchers who have modeled such phenomena; some are biologists or psychologists who have studied them theoretically or experimentally. ○

Economic Realities

During his tenure this spring as resident research director of the economics program, Blake LeBaron has focused in particular on steering the program in an empirical direction, one that stresses the application of SFI techniques to explain real-world economic facts. In tandem with theoretical work, under Blake's direction the program looked at things like technical trading patterns in financial markets and the dynamics of trading volume from the SFI perspective. The economics work has been integrated with other empirical programs at the Institute, especially research in low-dimensional chaos. This encourages information transfer about methods that can be equally useful for a stock-price time series as well as, for example, observations coming from brain wave data.

In one project LeBaron, along with Brian Arthur, John Holland, Richard Palmer, and Paul Tayler, have developed an artificial stock market model. This simulation provides an environment for studying the behavior of a collection of artificially intelligent agents trying to forecast the future of a traded asset that pays a varying dividend. The agents develop trading rules (based on a variety of market data) by which they determine when to buy, sell, or hold the asset. The aim is to understand what phenomena result from the interactions of different learning algorithms working in a simple stock market trading environment. This model is being used to generate a simulated time series that will be compared to actual time series to see what kinds of real-world phenomena are replicated in this computer-generated market.

(For further discussion see "Wall Street in Santa Fe" in this issue.)

Some traders of financial assets claim that certain simple rules can be used to help forecast future prices, although this claim has long been disputed. Another current SFI project is shedding light on this debate. Buz Brock, Josef Lakonishok, and Blake LeBaron have used new statistical tests to show that certain simple technical trading rules do find statistically significant patterns in both stock prices and foreign exchange rates. The SFI trio used a simple moving average and trading range types of rules as specification tests for several different stochastic processes of stock returns. But while the trading rules behaved as technical traders would have predicted, none of the tested processes were able to replicate the unusually large returns during buy periods and small returns during sell periods that were found in the original series. Nor did the study adjust for transactions costs and risks. This remains an important empirical component that still needs to be explained.

Human experiments on auctions find that subjects make systematic bidding errors that cannot be explained by standard bidding models. James Andreoni and John Miller are considering these errors using a model of adaptive learning based on a genetic algorithm. They are finding that artificial adaptive agents exhibit many of the same bidding patterns as those observed in auctions with humans. Andreoni and Miller think that adaptive learning may provide an explanation for the divergence between theoretical and experimental results in auction markets.

Their findings suggest that adding adaptive benchmarks to experimental and theoretical results has broad potential for economic analysis in general.

Another puzzle for economic theory is explaining the observed instability of economic aggregates. A number of reasons exist for variation in the pace of production and consumption at the local level, but it is hard to see why there should be large variations in those factors that are synchronized across the entire economy. In an effort to understand this real-world phenomenon Per Bak, K. Chen, José Scheinkman, and Michael Woodford have constructed a simple model of a multi-sector, multi-stage production process. It illustrates that, in fact, small shocks do not "cancel out" in the aggregate. Conventional reasoning fails as a result of significantly nonlinear, strongly localized interactions between different parts of the economy. The type of macroscopic instability that can result has been studied in a variety of other contexts under the name of "self-organized criticality."

Having made considerable progress toward solidifying the nexus from SFI tools to observed economics phenomena, LeBaron returns to U. Wisconsin this fall. In September SFI will host a full-scale gathering—on the scale of the ground-breaking 1987 meeting which brought together economists and physical and biological scientists to forge a new conceptual framework for economics—to provide renewal for the program. The meeting will be led by Kenneth Arrow and Philip Anderson, co-chairs of the 1987 workshop. ○
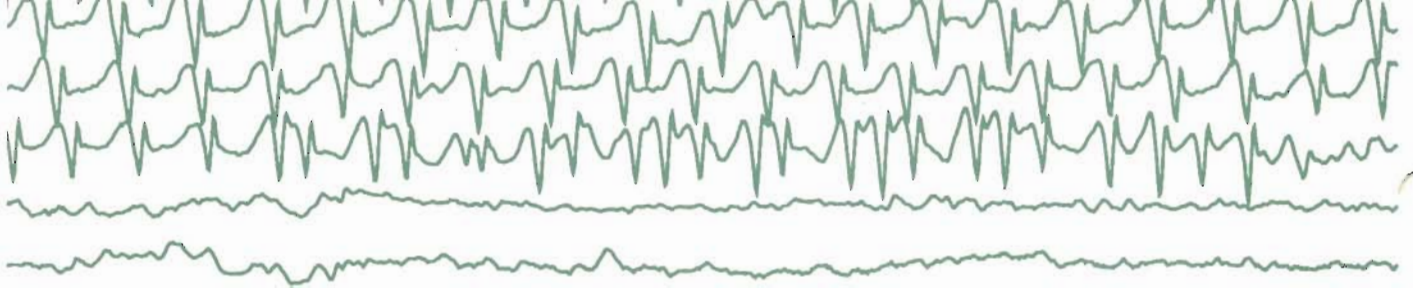
Patterns in Time Series Data

Time series analysis problems are central to a wide range of disciplines including physics, biology, and economics. SFI work in this field continues on several fronts.

Supported by a National Institutes of Health grant James Theiler and collaborators Doyne Farmer, Dante Chialvo, and Andre Longtin with graduate fellow Brandt Hinrichs are currently investigating allegations that the electroencephalogram (EEG) is chaotic and, in general, are exploring the use of nonlinear time series methods in the characterization of the dynamical behavior of the nervous system. The aim of this work is to discover if there is any underlying deterministic structure that may lie hidden in apparently random neural phenomena. Postdoctoral Fellow Milan Palus is also working on this issue, supported by a separate NIH award. Although the group has seen that unambiguous evidence for nonlinearity has been found in some EEG time series, in no case have they found good evidence for chaos in either normal EEG or in epileptic EEG, despite widespread claims to that effect.

While the evidence for nonlinearity is relatively straightforward, the issue of chaos is still unresolved. By looking at a spectrum of statistics, including one that compares forward to backward prediction, they are addressing this issue in the context of their EEG time series. The group has recently acquired some EEG data taken not from the scalp, but directly from the brain's surface, and they are now analyzing these sets for evidence of nonlinearity. Because these are also multivariate time series, they will be able to use recently developed tools for generating multivariate surrogate data as part of the tests.

The collaborators are also compiling a software library for nonlinear time series analysis. The package—called **Ts/tools**—will soon be released to the public domain (parts of it are already available). Currently the package includes routines for dimension estimation, nonlinear forecasting, surrogate data generation, recurrence plots, and a host of general-purpose tools such as Fourier transforms and convenient random number generators.

SFI's other current major arena for time series analysis is within the economics program; the analysis of large-scale economic databases provides a convincing means of comparing the Institute's modeling efforts with actual market performance. Buz Brock, Blake LeBaron, Jose Scheinkman, Dee Dechert, and David Hsieh, for example, have developed a new type of nonlinear diagnostic test that can be applied to different financial time series. They have also applied the test (called the BDS test) to simulated time series. Applying this test to financial markets, they have found strong evidence for nonlinearities, but little evidence for low-dimensional chaotic dynamics. LeBaron has also looked for qualitative properties of nonlinearities along with multivariate properties including trading volume. His results show there are connections between volume, volatility, and the overall predictability of price movements, and he has applied this information within a forecasting framework to improve sample predictions for some time series. This result is closely related to similar properties connecting volume movements and predictability. During certain periods volume movements indicate more or less predictability in a financial series. These results suggest that an important nonlinear feature of financial time series is that their near-term forecastability changes across different states of the dynamical system. This feature is common to many nonlinear dynamical systems.

Time series work in economics progresses in close contact with Theiler's low-dimensional chaos group. In fact, in May a nominal "time series" working group was formed, drawing interested workers from each program as well as additional visitors. The group met regularly throughout the month to address general issues in nonlinear time series as they relate to economics and finance, physics, biology, and medicine. It ended the month by zeroing in on some important unresolved questions in the field. For example: How useful is the concept of low-dimensional chaos for analysis of nonlinear time series? Can we answer questions about whether a system is inherently stable and driven by outside shocks, or driven by small shocks that get amplified by the structure of the system? What happens when, given a finite amount of data, an incorrectly specified model does better at prediction than the true model? The group agrees there needs to be more work on connecting data to theoretical models and estimation of structural models. Further research also needs to be done in understanding and comparing bootstrapping and surrogate data methods. Each of these issues suggests multiple paths

for future work at SFI.

Look for the publication soon of *Time Series Prediction: Forecasting the Future and Understanding the Past*, edited by Neil Gershenfeld and Andreas Weigend, which details the results of a 1992 SFI workshop on time series forecasting. The results of an earlier workshop are contained in the proceedings volume *Nonlinear Modeling and Forecasting*, edited by Martin Casdagli and Stephen Eubank. ○

Complexity, Memory, and Learning in the Immune System

For a number of years Alan Perelson, Gérard Weisbuch, Lee Segel, Rob De Boer, and Avidan Neumann have been developing network models of the immune system that dealt with the class of antibody producing white blood cells, known as B cells. These models predicted the population dynamics of B cells and in some cases the antibodies they produced. Sullivan Scholar Avidan Neumann and SFI visitor Bernhard Sulzer have worked on improving this class of models by incorporating additional details about the various stages in the life history of a B cell. In particular, they have developed models that explicitly take into account the differentiation of a B cell from a proliferating state into a more differentiated stage known as a plasma cell. Plasma cells are very efficient secretors of antibody.

These network models, even with the improvements included by Neumann and Sulzer, are still incomplete. For reasons of simplicity, one important element of the immune system, a class of regulatory white blood cells, known as T cells, was never explicitly considered in the models; rather they were simply assumed to be present at sufficient concentration to mediate all immune phenomena. T cells are the cells infected by HIV in AIDS and, as the disrup-

tion of the immune system in AIDS shows, they are essential to the proper functioning of the immune system. For the past year Neumann and Alan Perelson have been developing a new class of network models in which T cells are explicitly considered and their functions modeled.

One important application of the new class of models is to the problem of self-nonself discrimination. The immune system needs to be able to distinguish self from nonself in order to avoid autoimmune disease. One important process in self-nonself discrimination is the elimination in the thymus of T cells reactive against self. T cells, like all blood cells, are made in the bone marrow. When produced these cells are undifferentiated and incapable of performing the functions of a T cell. The cells migrate to the thymus and under the influence of thymic hormones differentiate into fully functional T cells. However, while in the thymus, T cells are tested for anti-self reactivity. If a T cell interacts too strongly with self, it dies in the thymus. This process of elimination of anti-self T cells in the thymus is not 100% efficient and some anti-self T cells survive. Neumann and Perelson are using their model to show under what conditions immune networks can control the potentially dangerous anti-self re-
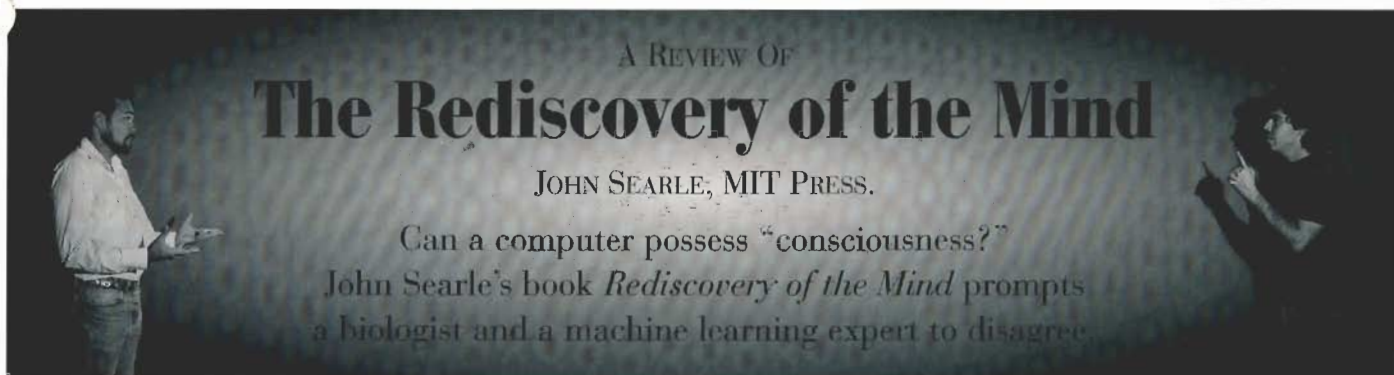
sponse caused by the escape of anti-self T cells from the thymus.

In related work, Perelson and Rob De Boer have been studying the effects of the elimination of anti-self lymphocytes on the diversity of the immune system. It has been commonly thought that the immune system needs to be very diverse in order to recognize all possible pathogens. What De Boer and Perelson have shown with a mathematical model is that the major factor in determining the diversity of the immune system is not the need to recognize foreign pathogens but rather the need to avoid reacting with self. For example, one could create a sticky lymphocyte that recognized all other cells. Such a lymphocyte would recognize all foreign cells, but unfortunately it would also recognize all self cells. Since cells that recognize self are killed, one needs to generate a sufficiently diverse population of cells in the bone marrow so that after removal of the anti-self cells the cells that remain can still recognize foreign pathogens. The theory that De Boer and Perelson generated showed that the number of different types of cells that need to be generated in the bone marrow is strongly dependent upon the number of self-antigens and only influenced weakly by the number of foreign pathogens. ○

Residential Researchers and Visitors, January to June 1993

Donald Adolphson, Brigham Young University

Philip Anderson, Princeton University, SFI External Faculty

James Andreoni, University of Wisconsin

Michael Angerman, Los Alamos National Laboratory

W. Brian Arthur, Stanford University, SFI External Faculty

Michel Baranger, Massachusetts Institute of Technology/Los Alamos National Laboratory

Aviv Bergman, Stanford University/Interval Research

Bill Bruno, Los Alamos National Laboratory

Leo Buss, Yale University

Joe Bryngelson, Los Alamos National Laboratory

David Campbell, University of Illinois, Urbana-Champaign, External Faculty

Dante Chialvo, State University of New York, Syracuse

Jim Crutchfield, University of California, Berkeley, SFI External Faculty

Jean Czerlinski, New College of the University of Southern Florida, Undergraduate Intern

Rob De Boer, University of Utrecht, SFI External Faculty

Michael Diehl, State University of New York, Buffalo, Graduate Fellow

Guillemette Duchateau, Ecole Normale Superieure

David Easley, Cornell University

Rob Farber, Los Alamos National Laboratory, SFI External Faculty

Doyne Farmer, Prediction Company, SFI External Faculty

Walter Fontana, SFI Postdoctoral Fellow

Stephanie Forrest, University of New Mexico, SFI External Faculty

Dan Friedman, University of California Santa Cruz

Murray Gell-Mann, California Institute of Technology, SFI External Faculty

Leon Glass, McGill University

David Goldberg, University of Illinois

Valerie Gremillion, Los Alamos National Laboratory

George Gumerman, SFI Science Board

Brosl Hasslacher, Los Alamos National Laboratory

Matthew Headrick, Princeton University, Undergraduate Intern

Stefan Helmreich, Stanford University, Visiting Graduate Fellow

David Hiebeler, Thinking Machines Corporation

Brant Hinrichs, Beckman Institute, University of Illinois, Graduate Fellow

Alfred Hubler, University of Illinois, SFI External Faculty

Jeff Inman, SFI Member

Julian Jack, Oxford University

Atlee Jackson, University of Illinois

Erica Jen, Los Alamos National Laboratory, SFI External Faculty

Terry Jones, University of New Mexico, Graduate Fellow

Pentti Kanerva, NASA Ames Research Center

Stuart Kauffman, University of Pennsylvania, SFI External Faculty

Tim Kohler, Washington State University

Nancy Kopell, Boston University, SFI Science Board

Bette Korber, Los Alamos National Laboratory/SFI Postdoctoral Fellow

Greg Kramer, Clarity, SFI Member

David Lane, University of Minnesota, SFI External Faculty

Chris Langton, Los Alamos National Laboratory, SFI External Faculty

Alan Lapedes, Los Alamos National Laboratory, SFI External Faculty

Blake LeBaron, University of Wisconsin, Economics Program Resident Director

Martin Lettau, Princeton University, Graduate Fellow

Angela Linse, University of Washington

Seth Lloyd, Los Alamos National Laboratory, SFI External Faculty

Bill Macready, University of Toronto, SFI Postdoctoral Fellow

Franco Malerba, Bocconi University

Robert Marks, Stanford University

Dan McShea University of Michigan

John Miller, Carnegie-Mellon University, SFI External Faculty

Mark Millonas, Los Alamos National Laboratory

Nelson Minar, Reed College, Undergraduate Intern

Melanie Mitchell, University of Michigan, Adaptive Computation Resident Director

John Moody, Oregon Graduate Institute

Cris Moore, SFI Postdoctoral Fellow

Benoit and Penelope Morel, Pittsburgh Cancer Institute

Lynn Nadel, University of Arizona, Codirector, Summer School

Avidan Neumann, SFI Postdoctoral Fellow

Mats Nordahl, SFI Postdoctoral Fellow

Mihaela Oprea, Washington University

Una-May O'Reilly, Carleton University, Graduate Fellow

Luigi Orsenigo, Bocconi University

Richard Palmer, Duke University, SFI External Faculty

Milan Palus, Prague Psychiatric Center, SFI Postdoctoral Fellow

Luca Peliti, University of Naples

Jerry Percus, Courant Institute

Alan Perelson, Los Alamos National Laboratory, SFI External Faculty

David Pieczkiewicz, Case Western Reserve, Undergraduate Intern

Darren Pierre, Case Western Reserve, Undergraduate Intern

David Pines, University of Illinois, Urbana-Champagne, SFI External Faculty

Alexis Manaster Ramer, Wayne State University

Steen Rasmussen, Los Alamos National Laboratory

Greg Rawlings, University of Indiana

Tom Ray, University of Delaware

David Rogers, Molecular Simulations

Matt Root, Washington State University

Randall Rose, St. John's College, Undergraduate Intern

Robert Savit, University of Michigan

Bruce Sawhill, St. John's College

Peter Schuster, Institute for Molecular Biology, SFI External Faculty

Lee Segel, Weizmann Institute

Roger Shepard, Stanford University

Martin Shubik, Yale University

Josh Smith, Massachusetts Institute of Technology/Media Lab

Peter Stadler, University of Vienna

Daniel Stein, University of Arizona, External Faculty

Charles Stevens, Salk Institute

Paul Stolorz, SFI Postdoctoral Fellow

William Sudderth, University of Minnesota

Bernard Sulzer, University of Munich

James Theiler, Los Alamos National Laboratory/SFI Postdoctoral Fellow

Bill Tozier, University of Pennsylvania

Joe Traub, Columbia University

Miguel Virasoro, University of Rome

Frederik Wiegel, Twente University of Technology

Gérard Weisbuch, Ecole Normale Supérieure, SFI External Faculty

David Wolpert, SFI Postdoctoral Fellow

Larry Wood, GTE

Wojciech Zurek, Los Alamos National Laboratory, SFI External Faculty

A REVIEW OF

The Rediscovery of the Mind

JOHN SEARLE, MIT PRESS.

Can a computer possess "consciousness?"
John Searle's book *Rediscovery of the Mind* prompts
a biologist and a machine learning expert to disagree.

Laury Alexander, JILA Photo/Design

Thomas B. Kepler

Aristotle botched physics so badly that contemporary physicists simply ignore philosophers as a matter of course. Descartes' handling of the mind-body question was equally egregious, but scientists who study the mind seem actually to encourage, and often participate in, philosophizing about their subject. Over the last decade, working scientists and the educated public have been pelted under a hailstorm of books devoted to the philosophical elucidation of consciousness. Most of these disappoint rather than illuminate; in spite of their clever ideas and technical sophistication they just don't seem to grapple adequately with the definitive characteristics of consciousness. Instead, they invent elaborate schemes to avoid it. A notable exception is John Searle's *The Rediscovery of the Mind*, which delivers a long overdue dose of common sense within a deliciously wry and enjoyable text.

OK then, why philosophy? Because there is a deep sense in which the study of the mind cannot be like physics. Physics has succeeded by shifting attention away from subjective sensation: heat is thus treated as molecular agitation, color as photon wavelength. But subjectivity is the essence of the mental. If you shift attention away from subjectivity here, you miss the whole point. Descartes was so taken with the dichotomy that he decided there had to be two different substances making up the world, physical stuff and mental stuff. This opinion is now so thoroughly discredited as mystical and anti-scientific that the epithet "Cartesian Dualist" provokes rows among ordinarily gentle academicians. Our desperate rush from dualism has been so frenzied, says Searle, that even while distancing ourselves from its particular conclusions we have embraced the terms in which it was framed, and the resulting theories have been just as incoherent. Artificial Intelligence, as a paradigm for the study of real intelligence, is one of the most prominent symptoms.

David H. Wolpert

For a long time John Searle has had a reputation as a purveyor of frustratingly wooly arguments concerning the nature of the mind. Now he has written a new book, *The Rediscovery of the Mind*, which argues anew that the modern scientific approach to the mind and consciousness is ill-conceived on a foundational level and that, for foundational reasons, computers can't be conscious. Unfortunately, these elaborations should do little to change Searle's hard-won reputation.

First, it must be said that Searle demolishes his straw men with aplomb. Indeed, if there really are people out there prattling on in the doltish way he describes (e.g., "there is no subjective phenomenon of 'consciousness' which needs explaining in a theory of the mind"), those people should be hung, drawn, and sixteenthed, without delay, and in as public and humiliating a setting as possible. More power to Searle for working towards this end.

Searle also writes quite readably. Like the vast majority of philosophers, he has a tendency to substitute contorted and filigreed language for substance. But he keeps this tendency reasonably well in check.

Unfortunately though . . .

The problem is that Searle never takes the time to scrutinize his own views with the care he lavishes on the views of others. Yet it is precisely when one is far removed from experiment—as in the issues that Searle addresses—that it is absolutely crucial that one examine one's own arguments extremely critically. If you don't, you risk descending into the silliest (and most prevalent) kind of philosophy: the projecting of one's personal foibles and prejudices onto the universe. Unfortunately, just like his colleague-in-arms Roger Penrose, Searle is unabashedly this kind of philosopher, admonishing the universe to pay strict attention to his personal notions and fancies, all so heart-felt and solemn.

Searle lists the axioms upon which contemporary discourse is based, among them: "every fact in the universe is in principle knowable and understandable by human investigators" and "the only things that exist are ultimately physical, as the physical is traditionally conceived, that is, as opposed to the mental." From this perspective, the pertinent question may be phrased, "How is it possible for unintelligent [or nonconscious] bits of matter to produce intelligence [or consciousness]?" or "[how can we] give an account of the mental that makes no reference to anything intrinsically or irreducibly mental?" The consensus answer is organization. (What else is there?) Then the fatal extrapolation: if it is organization that counts, then its substrate need not be considered; we cut through messy detail and concentrate on the abstract (hence, important) structure. Finally, then, we can duplicate the relevant organization in whatever machinery we elect, i.e., a computer. Then "the computer produces intelligence because it is implementing the right computer program."

This approach does have immediate appeal (even if we have been bullied into accepting it) but, when the ideas are further elaborated, one discovers a more furtive dualism, no less objectionable for its stealth. Again, consciousness has simply been stashed between the mattresses. Of course, you might declare that "intelligent behavior is the essence of intelligence," thereby validating your disregard of the subjective. This yields the Turing Test and, unfortunately, repeats the mistakes of behaviorism. Searle expounds at length on the technical difficulties of this functionalism, but I think the gut-level criticism is decisive: I know that subjectivity is a real and important part of my mental life; if you ignore it, you have not understood me. The response to this is the fiendish, incredible, yet entirely serious assertion that consciousness does not really exist at all. It's easy. First, grant that all understanding is based on some underlying theory. Then recognize that the theory implicit in the introspection of our own mental states is a particularly naive one, and call it "Folk Psychology," within which "consciousness" plays the role of "caloric fluid." It is simply a mistaken construct which will vanish once the real statistical mechanics of the mind is found.

Against these and other contortions, Searle pleads that:

> one can accept the obvious facts of physics—for example, that the world is made up entirely of physical particles in fields of force without at the same time denying the obvious facts about our own experiences—for example, that we are all conscious and that our conscious states have quite specific irreducible phenomenological properties.

Does this seem inconsistent to you? Aha! You are still embracing the Cartesian categories. Searle argues, as have many before him, that "reduction" needs careful attention and, furthermore, that one can consistently believe that "mental phenomena are caused by neurophysiological processes in the brain" while denying that mental states are identical to specific, recognizable neurophysiological states. That is, physical events cause mental ones, but we are barred from ever (Searle hedges, unnecessarily I believe; without a "major intellectual revolution") knowing just how.

You may remember Searle as the drudge who would burst the giddy bubble of AI in the early '80s with his "Chinese Room" argument, whose bottom line, briefly, was that syntax does not suffice for semantics: meaning does not occur simply by virtue of "running the right program." Now he is back and making the more radical statement that:

> notions such as computation, algorithm, and program do not name intrinsic physical features of systems. Computational states are not discovered within the physics, they are assigned to the physics.

Physics is not sufficient to syntax. The Chinese Room has acquired a profound new motive.

There is much to argue about in this book (I think he backs away from some of the more interesting conclusions to which his arguments might otherwise lead), but it finally asks the right questions and confronts us with the task of redefining the terms of our investigations. The debate will no doubt continue for as long as there are minds to conduct it: we will never be rid of the philosophers on this one.

Because of this, it's hard to write a coherent review of the shortcomings of Searle's arguments. In essence, Searle implores us to accept his wooly, vaguely defined notions, and then to acquiesce in his telling God to get out of the way, that those wooly notions are the basis for arguments that fix the way things must be. But if we never buy into Searle's wooly notions in the first place . . .

A perfect example of all this is Searle's notorious Chinese room argument:

> "*. . . imagine that someone who understands no Chinese follows an algorithm to converse in Chinese . . . no part of the system literally understands Chinese . . . so a programmed computer does not understand Chinese either*" (emphasis mine).

The obvious question is, in what sense does Searle mean that the system does not "understand" Chinese? He never says; he only excludes any property that is based on behavior. But without using behavior, what's left, this side of solipsism? How can Searle infer that anyone other than himself "really understands" Chinese? (Curiously, Searle seems to be implicitly aware of this difficulty with his view. He claims that behavior must be combined with "appropriate causation" for one to infer that someone "understands" something. What in the world this means Searle never really explains. And, rather tiresomely, it all just begs the question of whether a Chinese room exhibiting the proper behavior combined with "appropriate causation" would "understand" Chinese.)

This particular breed of terminology sloppiness is exasperating almost to the point of anger. It suggests all manner of silly questions similar to that of the Chinese room: Does a calculator "really calculate"? Does a four-legged mobile robot "really walk"? We always have the freedom to define (or not to define) our terms any way we want. But—obviously—if no theorems or practical insights arise from one particular (non)definition, there's no reason to use it. If the distinction between an algorithm "simulating the understanding" of Chinese and someone "really understanding" Chinese makes no possible difference, why make the distinction in the first place? Indeed, what formal (!) meaning could such a distinction possibly have?

In fact, one can go further. One interpretation of the Chinese room argument is that what it really shows is how Searle's intuitive notion of the word "understand" breaks down when stressed. Yet Searle refuses to draw this conclusion. Instead he insists that his intuitive notion is valid, and that, in essence, it's the universe which must conform to him, by accurately reflecting in its grand design Searle's personal preconceptions. This is a perfect example of Searle's projecting his personal notions and fancies onto the universe. It is, in the profoundest sense, anti-scientific.

Searle muddies all this up with many other issues. For example, he belabors the obvious point that one's subjective impressions are not "directly accessible" to an outside observer. Of course, a computer's subjective impressions are not "directly accessible" to such an observer either. (The situations are identical, from a philosophical viewpoint: in both cases, we can observe the behavior of the black box, and we can look inside the black box, but we can not "directly access" the subjective impressions of that black box.) He blithely bounces across many other very subtle issues as well, for example, by arguing about what it would "be like" to experience someone else's mental processes. (Searle ignores the obvious objection, that this question verges on being tautologically meaningless, since it implicitly assumes some kind of duality between the "you" of certain mental processes and the "you" experiencing those processes.) He also tries to overload the distinction between syntax and semantics, unleashing thunderbolts like "syntax is not intrinsic to physics." (Of course, many physicists [not to mention Zen Buddhists] would dispute this statement in the strongest possible terms. Searle should skim an article or two on string theory to see just how "syntactic" physics can be.)

In the end, one cannot help but imagine Searle sitting down with some latter-day Robbie the robot, insisting to poor Robbie that he cannot possibly be conscious, cannot possibly "understand" what Searle is saying. The affronted Robbie objects, naturally enough, but Searle persists, weaving finely tortured arguments in the vapor. Eventually Robbie, getting increasingly frustrated, warns Searle that he should stop being so insulting. Searle blithely ignores Robbie, chirping "Not conscious, not conscious, not conscious." At which point Robbie the robot, exasperated and annoyed, reaches across the table and bops him one.

Tom's Reply to David

Straw Men: I know that David is familiar with Pat Churchland's *Neurophilosophy* and Daniel Dennett's *Consciousness Explained*, for example. Neither "prattles doltishly": both are very sharp and cleverly argued as well as tremendously popular. They are also, however, strong in precisely the ways that Searle indicates. Wooliness: Where is the wooliness, in Searle's arguments or in reality? To assume that reality is all sharp edges and light, with respect to mental states in particular, is not necessarily useful and perhaps is even incoherent. "Understanding" is bound up in the peculiarly inaccessible realm of subjectivity: questions about novel forms of understanding are not analogous to questions about novel forms of "walking." To call Searle's attitude "profoundly anti-scientific" is to raise hyper-physicalism to a status unsupported by and unnecessary for the actual practice of science. Does David object to Searle's "personal preconceptions" or to his refusal to abandon his personal, i.e., subjective, conceptions? Robby: Searle has very carefully limited his assertions about conscious machines. He has stated his belief that we and our brains are machines (I myself am not so sure), and that therefore, quite clearly, machines can be conscious. He does not rule out the possibility that artifacts could someday be conscious, but insists that consciousness cannot arise by simply executing the right program, and that the computational paradigm is misguided. But what an inspired twist on the Turing test: the machine, programmed to detect hostility, slaps the skeptical researcher into submission. I'm convinced.

David's Reply to Tom

I'm afraid I must disagree with much of Tom's review. For example, Tom agrees with Searle that "subjectivity is the essence of the mental. If you shift attention away from subjectivity here, you miss the whole point." (The implication is that the objective approach of traditional science "misses the whole point," as far as the mind is concerned.) Well, this sounds good, but things are not so simple. I can give you psycho-active drugs which will, reliably, produce any one of a number of different pronounced effects on your "subjectivity." In what sense is such technology and the theoretical/empirical edifice behind it not part of traditional science?

I think much of the weakness of Searle's argument can be illustrated by replacing in Tom's review the loaded term "conscious" with a term which, two centuries ago, was equally loaded: "alive." For example, Tom seconds Searle in asking "How is it possible for unintelligent [or nonconscious] bits of matter to produce intelligence [or consciousness]?" The analogous question was "How is it possible for unliving matter to come together and be alive?" Searle's "how can we give an account of the mental that makes no reference to anything intrinsically or irreducibly mental?" becomes "How can we give an account of the living that makes no reference to anything intrinsically or irreducibly alive?" In sum, Searle's word games have been played many times before. There are differences this time, of course: "conscious" is not identical to "alive." But in both cases, there is a bad tendency to project one's subjective notions onto the universe. The lesson of history is that such projections die a squishy death.

Tom's Final Word

First, is "conscious" really analogous to "alive" in the sense David means? Again, I hold that the fundamental subjectivity of consciousness bars its reduction to purely objective terms. The same is not a priori true of life. Second, do we really know how "non-living bits of matter interact to produce life?" Not really, not yet. Appeal to the reducibility of living systems to physics is still premature. ○

David's Final Word

(1) "Understanding" might not be analogous to walking, but it's closely analogous to calculating. Why does Searle ignore the subjectivity of "calculating" in favor of that of "understanding"? I submit it's because technology has advanced to the point of pocket calculators but not yet to the point of pocket understanders. And anyone who's used a pocket calculator would laugh Searle out of the room if he based his argument on calculation. (2) As far as Searle's beliefs that "our brains are machines," I really can't find a coherent viewpoint here; if our brains are machines, then what is it about a machine, if not the program, that can confer consciousness? Searle never answers. It really does look to me like he's splitting nonexistent hairs, solely to (try to) salvage his preconceptions. ○

Tom Kepler, a former SFI Postdoctoral Fellow, teaches biomathematics at North Carolina State U.
David Wolpert is a Postdoctoral Fellow at the Santa Fe Institute.

1993 SFI Working Papers

*An * before the paper number means the paper is out of print as a working paper; when possible, we have given an alternate source for the paper.*

93-01-001
Smooth Maps of the Interval and the Real Line Capable of Universal Computation
Cristopher Moore

93-01-002
The Evolution of Secondary Organization in Immune System Gene Libraries
Ron Hightower, Stephanie Forrest, and Alan S. Perelson

93-01-003
Political Parties and Electoral Landscapes
Ken Kollman, John H. Miller, and Scott E. Page

*93-01-004
Aggregate Fluctuations from Independent Sectoral Shocks: Self-Organized Criticality in a Model of Production and Inventory Dynamics.
Per Bak, Kan Chen, José Scheinkman, and Michael Woodford
[in print: *Ricerche Economiche* 47 (1993): 3–30]

93-02-005
The Quasi-Periodic Oscillations and Low-Frequency Noise of Scorpius X-1 as Transient Chaos: A Dripping Handrail?
Jeffrey D. Scargle, David L. Donoho, James P. Crutchfield, Thomas Steiman-Cameron, James Imamura, and Karl Young

*93-02-006
Pathways to Randomness in the Economy: Emergent Nonlinearity and Chaos in Economics and Finance
W. A. Brock

93-02-007
On the Use of Evidence in Neural Networks
David H. Wolpert

93-02-008
The Santa Fe Art Market
Martin Shubik

93-02-009
Combining Generalizers Using Partitions of the Learning Set
David H. Wolpert

93-03-010
Turbulent Pattern Bases for Cellular Automata
James P. Crutchfield and James E. Hanson

93-03-011
The Evolution of Interface: Reduction of Recombination Among Three Loci
David B. Goldstein, Aviv Bergman, and *Marcus W. Feldman*

93-03-012
Some Advantages and Disadvantages of Recombination
Sarah P. Otto, Marcus W. Feldman, and Freddy B. Christiansen

93-03-013
Generic Excitable Dynamics on a Two-Dimensional Map
Dante R. Chialvo

93-03-014
Revisiting the Edge of Chaos: Evolving Cellular Automata to Perform Computations
Melanie Mitchell, Peter T. Hraber, and James P. Crutchfield

93-03-015
Memory in Idiotypic Networks due to Competition Between Proliferation and Differentiation
Bernhard Sulzer, Avidan U. Neumann, J. Leo van Hemmen, and Ulrich Behn

93-03-016
On Overfitting Avoidance as Bias
David H. Wolpert

93-03-017
Chaotic Time Series Analysis: Identification and Quantification Using Information-Theoretic Functionals
Milan Palus

93-04-018
Real-Valued, Continuous-Time Computers: A Model of Analog Computations, Part I
Christopher Moore

93-04-019
Predicting the Size of the T-Cell Receptor and Antibody Combining Region from Consideration of Efficient Self-Nonself Discrimination
Jerome K. Percus, Ora E. Percus, and Alan S. Perelson

93-04-020
How Diverse Should the Immune System Be?
Rob J. De Boer and Alan S. Perelson
[in press: *Proc. Roy. Soc. London* B (1993)]

93-04-021
News from the Northern American Southwest: Prehistory on the Edge of Chaos
Timothy A. Kohler

93-04-022
A Model for the Immune System Response to HIV: AZT Treatment Studies
Denise E. Kirschner and Alan S. Perelson

93-04-023
Echoing Emergence
John H. Holland

93-04-024
Reverse Hillclimbing, Genetic Algorithms and the Busy Beaver Problem
Terry Jones and Gregory J. E. Rawlins

*93-05-025
Implications of Creation
David E. Hiebeler

93-05-026
Don't Bleach Chaotic Data
James Theiler and Stephen Eubank

93-05-027
Evolution of Recombination Among Multiple Selected Loci: A Generalized Reduction Principle
Lev A. Zhivotovsky, Marcus W. Feldman, and Freddy B. Christiansen

93-05-028
Fluctuation Spectroscopy
K. Young and J. P. Crutchfield

93-05-029
The Escape Problem for Irreversible Systems
R. S. Maier and D. L. Stein

*93-05-030
The Cognitive Revolution?
John L. Casti

93-06-031
Random Walks and Orthogonal Functions Associated with Highly Symmetric Graphs
Peter F. Stadler

*93-05-032
The Theoretical Grapples with Complex Systems
Peter A. Segel

93-06-033
The Calculus of Self Interest in the Development of Cooperation: Sociopolitical Development and Risk among the Northern Anasazi
Timothy A. Kohler and Carla R. Van West

93-06-034
Mogollon Manos, Metates and Agricultural Dependence: Pithouse Villages, A.D. 200–1000
Michael W. Diehl

93-06-035
Observing Complexity and The Complexity of Observation
James P. Crutchfield

93-06-036
The Effect of Focusing and Caustics on Exit Phenomena in Systems Lacking Detailed Balance
Robert S. Maier and D. L. Stein

93-06-037
When Will a Genetic Algorithm Outperform Hill-Climbing?
Melanie Mitchell and John Holland

93-06-038
A Biologist's Guide to Internet Resources
Una Smith

93-06-039
Swarms, Phase Transitions, and Collective Intelligence
Mark M. Millonas

93-06-040
Dynamics, Computation, and the "Edge of Chaos": A Re-Examination

Melanie Mitchell, James P. Crutchfield, and Peter Hraber
93-06-041
Prediction and Adaptation in an Evolving Chaotic Environment
Alfred Hubler and David Pines
93-06-042
Linear Operators on Correlated Landscapes
Peter F. Stadler
93-06-043
Statistics of RNA Melting Kinetics
Manfred Tacker, Walter Fontana, Peter F. Stadler and Peter Schuster
93-07-044
Fast Folding and Comparison of RNA Secondary Structures
Ivo L. Hofacker, Walter Fontana, Peter F. Stadler, L. Sebastian Bonhoeffer, Manfred Tacker, and Peter Schuster
93-07-045
From Sequences to Shapes and Back: A Case Study in RNA Secondary Structures
Peter Schuster, Walter Fontana, Peter F. Stadler, and Ivo L. Hofacker
93-07-046
Estimating Functions of Probability Distributions from a Finite Set of Samples. Part I: Bayes Estimators and the Shannon Entropy
David H. Wolpert and David R. Wolf
93-07-047
Estimating Functions of Probability Distributions from a Finite Set of Samples. Part II: Bayes Estimators for Mutual Information, Chi-Squared, Covariance, and Other Statistics
David R. Wolf and David H. Wolpert
93-07-048
Immune Networks Modeled by Replicator Equations
P. Stadler, P. Schuster, and A. Perelson
93-07-049
Perspectives and Growth Areas of Basic Research in Germany
Siegfried Grossmann
93-08-050
Modeling Defective Interfering Virus Therapy for AIDS: Conditions for DIV Survival
G. Nelson and A. Perelson
93-08-051
Th1/Th2 Cross Regulation
M. Fishman and A. Perelson
93-08-052
Approaches to Artificial Intelligence
N. Nilsson and D. Rumelhart

New Research Awards January to June 1993

Bay Foundation, general support of the Institute's research program.

J. C. Downing Foundation, general support of the Institute's research program.

Ann and Gordon Getty Foundation, general support of research.

John D. and Catherine T. MacArthur Foundation, three-year grant to SFI which continues support for core research at current levels.

National Science Foundation, a three-year grant to establish a Research Experiences for Undergraduates Site at SFI.

National Science Foundation, renewal of a three-year award to support core research in complex systems at SFI.

National Science Foundation, support for Stuart Kauffman's research on "Adaptation to the Edge of Chaos in Cells and Ecosystems."

Carol O'Donnell Foundation, for the "Swarm" Simulation System project under the direction of Chris Langton; these funds support the work of David Hiebeler, a programmer from Thinking Machines Corporation.

New 1993 members of the **SFI Business Network for Complex Systems Research** are Maxis, the SimCity and SimLife software producer; the Electric Power Research Institute (EPRI) in Palo Alto, California; and Legg-Mason Mutual Fund Investors, headquartered in Baltimore. The newest member of the Business Network is ARPA (the DOD Advanced Research Projects Agency). ARPA's presence will put us in touch with a broad community of people working in areas of advanced computation, and it constitutes a strong endorsement of our program to other companies. BusNet membership now includes ARPA, Batterymarch Financial, Citibank, DEC, EPRI, Interval Research, John Deere, Legg-Mason Mutual Fund Investors, and Maxis.

Alfred P. Sloan Foundation, funding for Melanie Mitchell and Stephanie Forrest to study the foundations and applications of genetic algorithms.

Alfred P. Sloan Foundation, for support of theoretical neurobiology. This funding, along with funds carried over from last year's Pew Charitable Trusts grant for neurobiology, supported two working groups in July and August, led by Nancy Kopell (Boston U.) and by Michael Stryker (UC San Francisco).

Alex C. Walker Educational and Charitable Foundation, support for economics research.

New From the Santa Fe Institute Series

Recent Volumes

Thinking About Biology
Edited by W. Stein and F. Varela
Lecture Notes Volume III, 1993
Order numbers 62653-2 (hardcover) and 62454-0 (papercover)

Time Series Prediction: Forecasting the Future and Understanding the Past
Edited by Andreas S. Weigend and Neil A. Gershenfeld
Proceedings Volume XV [mid-October] Order numbers 62601 (hardcover) and 62602 (papercover)

The Double Auction Market: Institutions, Theories, and Evidence
Edited by Daniel Friedman and John Rust
Proceedings Volume XIV, 1993
Order number 62263-7 (hardcover) and 62459-1 (papercover)

The Principles of Organization in Organisms
Edited by Jay E. Mittenthal and Arthur B. Baskin
Proceedings Volume XIII, 1992
Order numbers 52765-0 (hardcover) and 58789-0 (papercover)

Nonlinear Modeling and Forecasting
Edited by Martin Casdagli and Stephen Eubank
Proceedings Volume XII, 1992
Order numbers 52764-2 (hardcover) and 58788-2 (papercover)

Evolution of Human Languages
Edited by John A. Hawkins and Murray Gell-Mann
Proceedings Volume XI, 1992
Order numbers 52572-0 (hardcover) and 52573-9 (papercover)

Artificial Life II
Edited by Christopher G. Langton, Charles Taylor, J. Doyne Farmer, and Steen Rasmussen
Proceedings Volume X, 1992
Order numbers 52570-4 (hardcover) and 52571-2 (papercover)

1991 Lectures in Complex Systems
Edited by Lynn Nadel and Daniel Stein
Lecture Volume IV, 1992 Order number 57834-4 (hardcover only)

Upcoming Volumes

1992 Lectures in Complex Systems
Edited by Lynn Nadel and Daniel L. Stein
Lecture Volume V

Auditory Display: The Proceedings of ICAD '92, the International Conference on Auditory Display
Edited by Gregory Kramer
Proceedings Volume XVI (late 1993)
This volume will include a CD.

Artificial Life III
Edited by Christopher G. Langton
Proceedings Volume XVII (late 1993)

Evolutionary Approaches to Southwestern Prehistory
Edited by George Gumerman and Murray Gell-Mann
Proceedings Volume XVIII

Integrative Themes
Edited by George A. Cowan, David Pines, and David Meltzer
Proceedings Volume XIX

STAFF

Santa Fe Institute
1660 Old Pecos Trail, Suite A
Santa Fe, NM 87501
U.S.A.

NON-PROFIT
ORGANIZATION
U.S. POSTAGE
PAID
PERMIT NO. 415
Santa Fe, NM