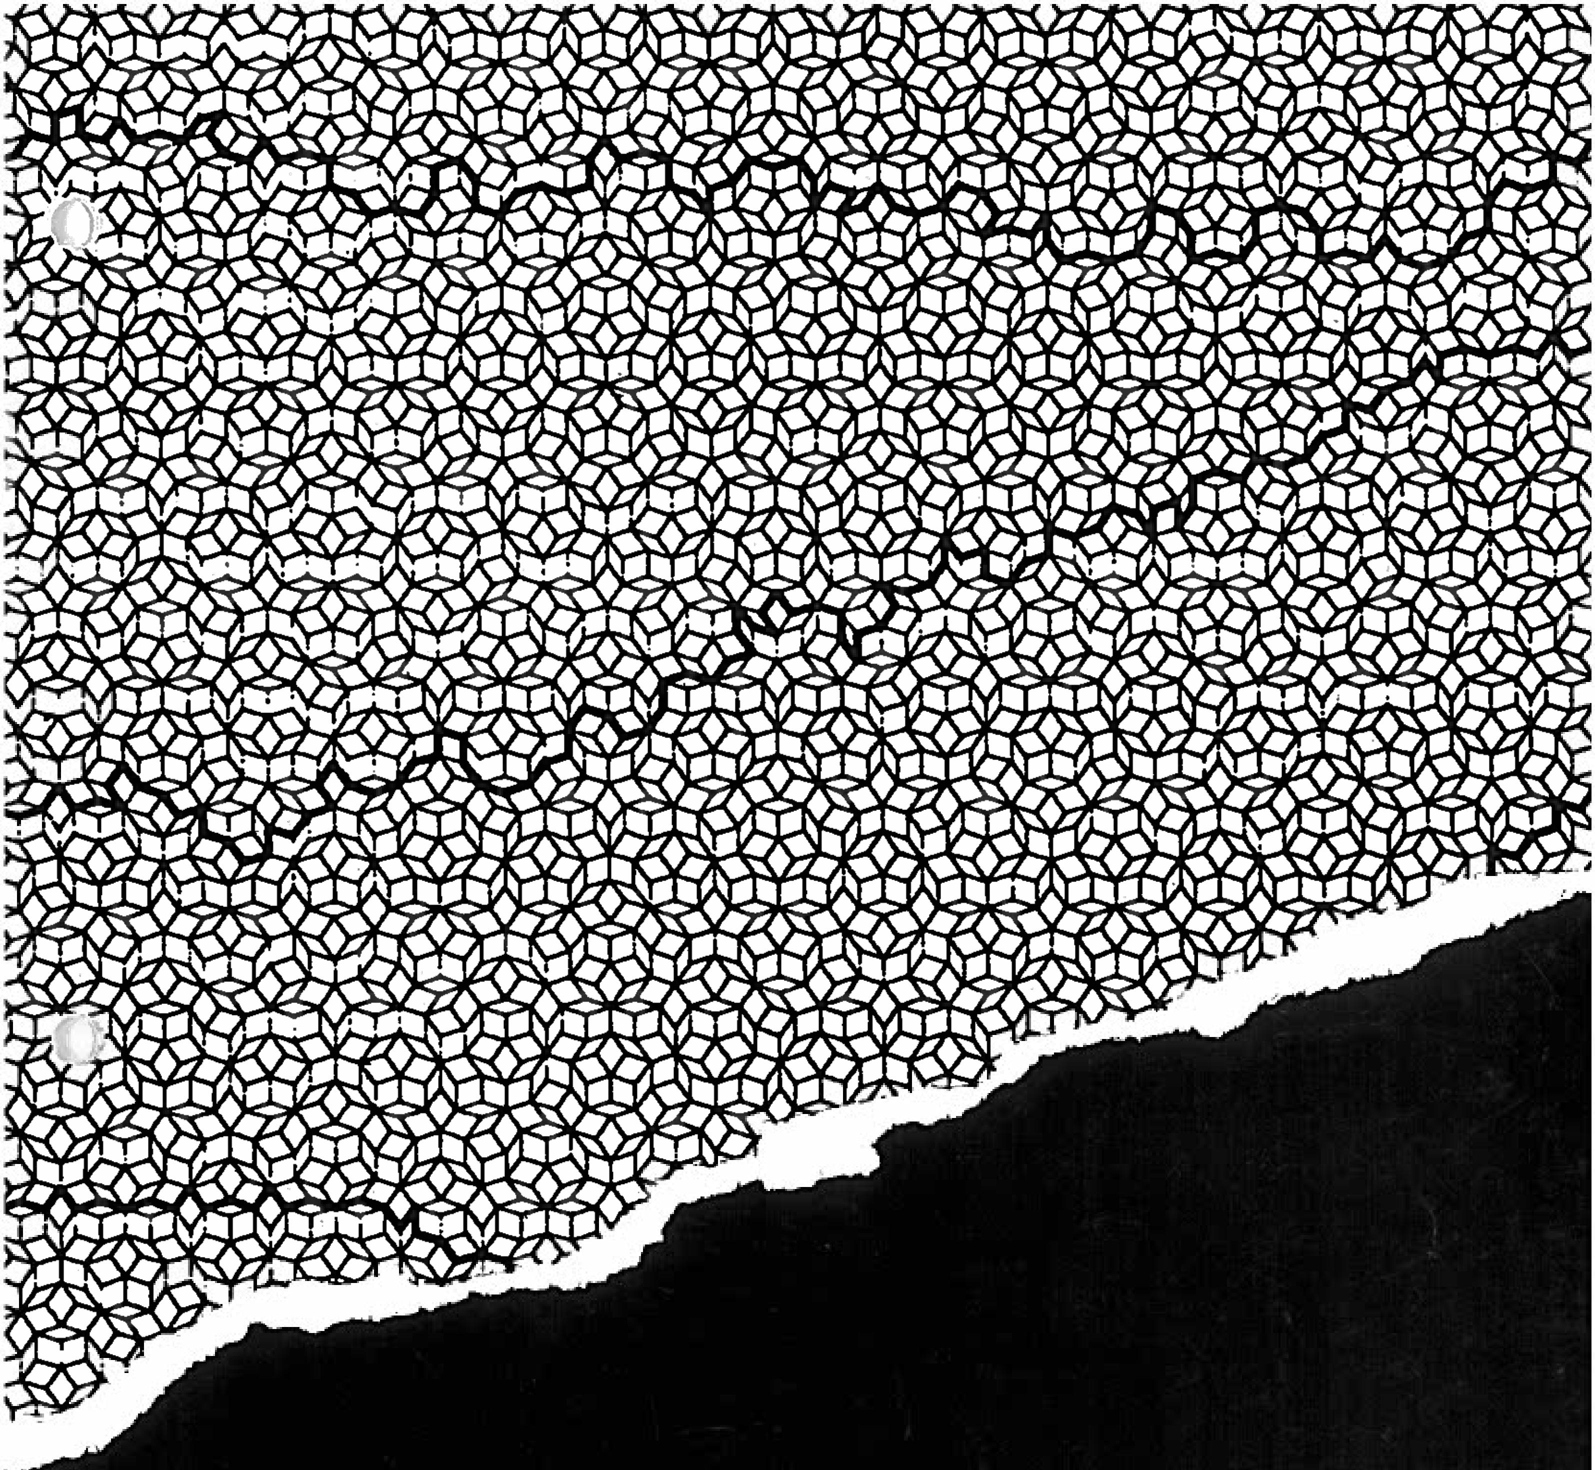

SANTA FE INSTITUTE

WINTER-SPRING, 1990
VOLUME 5, NUMBER 1

In This Issue

On the Cover

On the front cover, the patterned background shows stages of growth of an aperiodic Penrose tiling. (The black and white "hills" in the foreground were added by the SFI staff for graphic purposes.) As Charles Bennett, IBM, explains:

"Aperiodic structures and dynamical rules for their growth were among the topics considered in the October 1989 Santa Fe Institute workshop "Coherence and Complexity in Homogeneous Stochastic Media." Perhaps the most famous aperiodic structure is the tiling devised by Roger Penrose of Oxford University, defined by imposing certain local rules on the placement of two kinds of tiles (small rhombuses in the illustration). Here the Penrose tiling is grown starting from a one-dimensional boundary condition, a band of tiles with a simple self-similar structure. Successive heavy lines show the growth front at later stages, as tiles are added to sites randomly chosen from among the "live" sites along the existing growth front, i.e., sites at which the local rule uniquely specifies how the next tile can be added.

"From time to time parts of the growth front become dormant (flat regions). In these regions growth is temporarily inhibited because the local rule fails to specify the next tile uniquely. Growth must then proceed laterally from the remaining live sites until the intervening dormant regions are reactivated. When the Penrose tiling is grown, as here, from an extended one-dimensional boundary, growth can continue indefinitely, since some live sites will always remain to reactivate the dormant regions. By contrast, if a Penrose tiling is grown from a finite seed, the entire growth front eventually becomes dormant, and can only be reactivated by invoking non-local rules."

This figure was generated by David DiVincenzo of IBM Yorktown and reprinted with his permission.

Bulletin of the Santa Fe Institute

Volume 5, No. 1
Winter-Spring, 1990

Editor:
L. M. Simmons, Jr.
Associate Editor:
Ginger Richardson
Publication:
Ronda K. Butler-Villa

The Bulletin of the Santa Fe Institute is published biannually by SFI to keep our friends and supporters informed about the scientific and administrative programs. The Bulletin is free of charge and may be obtained by writing to the Editor, 1120 Canyon Road, Santa Fe, New Mexico 87501.

The Santa Fe Institute is a multidisciplinary scientific center formed to nurture deeper examination of complex systems and their simpler elements. A private, independent institution, SFI was founded in 1984.

Its primary concern is to focus the tools of traditional disciplines and emerging new computer resources on the problems and opportunities that are involved in the multidisciplinary study of complex systems—those fundamental processes that shape almost every aspect of human life and experience.

President's Message

We have completed a year of great progress toward our ambitious goals. Commitments made to the Institute by our very able academic staff and visitors means that the rate of increase in the Institute's research programs will be limited only by our financial resources. Planned growth in our academic program must now be matched with firm plans to increase the size of our physical facilities. Our schedule will commit us, under the terms of our renewed lease with the Cristo Rey parish authorities for our present quarters in the "convent," to move to a permanent or interim campus no later than the end of calendar 1991. This means that firm decisions must be taken before next summer concerning the location and size of our new campus. A reserve fund has been established to help us undertake preliminary planning for our new facility. A campaign for funds to establish this facility will be organized early in 1990.

An overall measure of the increased level of activity of the Institute is the fact that total expenditures in the current year, over $1.5 million, exceed the previous year's level by nearly 90%, the major part of the increase going to an expanded academic program. We plan to increase again in 1990 by at least one-third over the 1989 level. With growing emphasis on our development efforts we expect that, in fact, the Institute will be able to match last year's growth rate.

The Institute is increasingly able to base its appeal for support on a demonstrated record of performance. This record, which would be remarkable for a long-established institution, is even more impressive when we recall that 1989 marks the second year of our full-time academic program. Our achievements include:

Seth Lloyd, George Cowan, and Philip Anderson take a lunch break in the Institute courtyard. Photo by Cary Herz © 1989.

- Professor Brian Arthur completed his sabbatical leave and returned to Stanford at the end of the year, leaving behind at SFI a vigorous program in economics. During this year, the program included three meetings and thirty-five seminars, residency by forty scholars for periods of time ranging from several days to several weeks, a two-day fall meeting in New York City in conjunction with the New York Academy of Sciences to present some initial research results, and a program review at the third annual workshop in the fall, co-chaired by Kenneth Arrow and Philip W. Anderson. A major activity in 1990 will be the sponsorship of a national competition to investigate the ability of computer algorithms to duplicate or exceed the efficiency of human agents in a trading auction. It is organized under the title "Double Auction Tournament." The competition will be followed by a workshop later in the spring to assess the results.
- Other programs pursued in 1989 at SFI included workshops and network meetings on integrative aspects of adaptive, complex systems co-chaired by Murray Gell-Mann and David Pines; molecular evolution co-chaired by Stuart Kauffman and Alan Perelson; parallel computer systems co-chaired by Margaret Simmons, Rebecca Koskela, and Ingrid Bucher; hierarchical energy landscapes co-chaired by Hans Frauenfelder and Robert D. Young; complexity, entropy, and the physics of information chaired by Wojciech Zurek; the evolution of human languages chaired by Jack Hawkins and Murray Gell-Mann; prehistoric southwestern society co-chaired at the School of American Research by George Gumerman and Murray Gell-Mann; the relation between human cognition and emotion co-chaired by David Rumelhart and Jerome Singer; organization and complexity in stochastic media chaired by Charles Bennett; a contractor/grantee meeting on the human genome chaired by Sylvia Spengler; developmental biology chaired by Richard Burian, Stuart Kauffman, and William Wimsatt; and environmental issues, economics, and public policy co-

(continued)


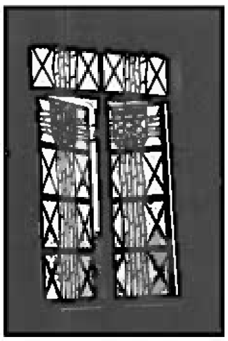

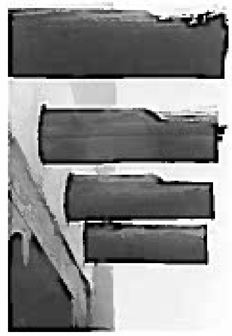


SFI photos by Cary Herz © 1989.

President's Message (continued)

chaired by Brian Arthur and Murray Gell-Mann.

- As a member of an academic consortium, the Institute managed the second annual Complex Systems Summer School, directed by Dr. Erica Jen, in Santa Fe. It was attended by more than fifty students and was greeted with as much enthusiasm by the attendees and lecturers as was the first Santa Fe summer school in 1988.
- The publications program added three new proceedings volumes to its series on the sciences of complexity and an additional volume of edited lectures summarizing the 1988 Summer School lectures. Four more volumes are currently in progress. Sales of these volumes is encouraging and they will be promoted with additional publicity by Addison-Wesley in 1990.
- Gifts and grants to the Institute from corporations, foundations, government agencies, and individuals increased in 1989 by 50% to more than $1.5 million from the 1988 level of slightly more than $1 million.
- Strengthening of the Board of Trustees is continuing in order to increase the Institute's ability to manage its responsibilities and to attract the added sponsorship that will be required to fund its growth.
- Planning for a major economics program in 1990 has been essentially completed, based on core support from Citicorp and the Institute, at last year's levels. We anticipate additional funding from other sources.

At year's end it is a pleasure to acknowledge the efforts of a dedicated staff and a host of tireless Board members, visiting scientists, and generous supporters of the Institute. Our hopes and wishes for the New Year and the new decade are very great but are securely based on the accomplishments of the past. Best wishes for 1990.

—GAC

Postdoctoral Fellowship in Complex Systems Studies at the Santa Fe Institute

The Santa Fe Institute may have one or more openings for postdoctoral fellows beginning in September, 1990.

The Institute's research program is devoted to the study of complex systems, especially complex adaptive systems. Systems and techniques under study include the economy; the immune system; the origin of life; artificial life; models of evolution; neural networks; genetic algorithms and classifier systems; complexity, entropy, and the physics of information; nonlinear modeling and prediction; cellular automata; and others.

Candidates should have or expect to receive soon a Ph.D. and should have backgrounds in theoretical physics or chemistry, computer science, mathematics, economics, game theory, theoretical biology, dynamical systems theory, or related fields. An interest in interdisciplinary research is essential.

Applicants should submit a curriculum vitae, list of publications, and statement of research interests, and arrange for three letters of recommendation.

Send applications to:

Postdoctoral Committee
Santa Fe Institute
1120 Canyon Road
Santa Fe, NM 87501

Applications or inquiries may also be sent by electronic mail to:

postdoc@sfi.santafe.edu

The Santa Fe Institute is an equal opportunity employer.

Letters to The Editor

Editor:

Questions of perception, cognition, language, scientific knowledge, evolution, and complexity are important issues for SFI. A new research project, the Biological Knowledge Laboratory (BKL), is pursuing these as basic research topics but at the same time is looking toward a practical goal which unifies them. The practical goal is to build a computerized "Scientist's Assistant" (SA) that can read the biological literature and then discuss it with a biologist in an informative and intelligent manner. The application field is bacterial chemotaxis, the simplest sensory-motor system known.

A machine that reads the literature must deal with the syntax and semantics of natural language and diagrams. The difficult problem in natural language analysis is semantics, understanding the meaning of a text. To understand a scientific article the reader must already have some knowledge of the scientific area. The reader must go through a reasoning process to discover the author's intent. This means that we have to "bootstrap" the SA by building into the system, by hand, the basic scientific knowledge and cognitive skills that the machine needs to process and store this information in a "knowledge base." Many of the same considerations apply when the computer has to analyze the questions that a scientist asks and then compose a response. The primary response of the SA will be to display the most appropriate paragraphs and diagrams from the documents in the knowledge base as well as its reasons for choosing them.

The exciting thing about this task is that we have to think long and hard about the nature of scientific knowledge itself, because we have to explicitly represent it in a machine. This representational imperative will, we hope, push us to a far deeper understanding of scientific knowledge and the process of scientific discovery and validation.

Diagrams in papers are an efficient way, and often the only way, to encode and present certain types of information. Unfortunately, little has been done by workers in computer vision to build systems that can understand and store diagrams, much less find and return the ones most relevant to a particular query. Building a diagram-understanding system is a major project in the BKL. In our system the visual "focus of attention" of the system is moved about on the diagram following lines and contours, noticing nearby objects, and jumping to "interesting" areas as suggested by peripheral vision. The overall organizing principle for diagrams seems to be a gestalt principle that we call "generalized equivalence." For example, the tick marks on a graph can be recognized because they are all the same length (length equivalence). Or for that matter, each line of text in this article that you are reading is really an equivalence class of characters which lie on (define) a horizontal line segment.

In artificial intelligence it is normally assumed that knowledge and reasoning should be dealt with by discrete mathematical logic. But looking at the biological and psychophysical aspects of human and animal behavior suggests an alternate approach, analog simulation, which the BKL is pursuing. With our ZOE system we are modeling the life and evolution of cells containing genes and containing molecules that react and diffuse inside and outside them. This simulation of living systems will be the basis for our system's "thinking" about biology, and for that matter, our model of how biologists think about biology—directly, as three-dimensional physical objects with appearance and properties—working schematic versions of the real thing, programs running on the brain's computational machinery.

Our project has some similarities to the work on the Matrix of Biological Knowledge, but the Matrix program's emphasis is more on tying together disparate databases, whereas ours is focused on discovering and encoding the conceptual contents of collections of scientific papers. The ideas behind the BKL were described at the first Biomatrix meeting in Santa Fe in 1987.

Those wanting further specific information can contact me here at the Biological Knowledge Laboratory, Northeastern University, Boston, MA 02115.

Sincerely,
Robert P. Futrelle

Editor's Note: In the summer of 1987, Prof. Robert P. Futrelle participated in SFI's summer workshop on the "Matrix of Biological Knowledge." He is now head of the Biological Knowledge Laboratory at Northeastern University.

Dear Editor:

I am preparing a workshop on Artificial Intelligence, Expert Systems, and Modern Computer Methods in Systematic Biology (ARTISYST Workshop), and it occurred to me that this program may be of interest to the readers of your Bulletin.

The workshop, funded by the National Science Foundation, will take place September 9–14, 1990, at the University of California, Davis, with about 45 participants, an even mixture of biologists and computer scientists. The Workshop subject areas are: 1) expert workstations for systematics; 2) expert systems, expert workstations, and other tools for identification; 3) phylogenetic inference and mapping characters onto tree topologies; 4) literature data extraction and geographical data; and 5) machine vision and feature extraction. Hotel and per diem expenses will be paid for participants, and travel will be paid to a maximum of U.S. $500. Attendance by invitation only.

Interested persons can apply by sending me: 1) their name, address, and phone number; 2) whether they apply as a computer scientist or as a biologist; 3) a short resumé; 4) a description of their previous work and planned research, and how they relate to the workshop; and 5) whether they, as a biologist (or computer scientist), plan to establish permanent collaboration with computer scientists (or biologists). **Applications must be received by April 15, 1990.**

Thanks for your help.

Sincerely yours,
Dr. Renaud Fortuner
Calif. Dept. of Food & Agric., Rm. 340
P. O. Box 942871
Sacramento, CA 94271-0001
Phone: (916)445-4521
Fax: (916)322-5913
E-mail: rfortuner@ucdavis.edu

Book Reviews

Read On

The Emperor's New Mind
By Roger Penrose
Oxford Univ. Press, 1989, 466 pages

This book is an essay on the nature of consciousness. In particular, it addresses the question whether there is a crucial difference between computers, as they are currently conceived, and our brains. Many books and papers have been written on this subject, and someone who has read a few of these may be inclined to think that he or she knows what issues will be brought up in such an essay, and what the arguments will be on each side. But Penrose has some surprises for such a reader. In order to explain his thesis, he must first tell the reader about all of the following: Gödel's incompleteness theorem, the quantum measurement problem, the second law of thermodynamics, the big bang, black holes, and especially, quantum gravity. It may sound crazy that quantum gravity (not to mention black holes) could have anything whatever to do with consciousness—indeed, Penrose himself points out that it sounds crazy—but once the argument is finished, it is clear that these matters are central to his thesis.

Let me try to summarize his thesis briefly, even though the summary will not do justice to his carefully written essay.

Computers operate on an algorithmic basis. Our brains also do a lot of algorithmic information processing, but most of this is unconscious. Why, then, is anything that our brains do associated with consciousness? Penrose's answer: Perhaps consciousness has evolved for the purpose of making judgments which cannot be made algorithmically. (A computer would then not be conscious.) But (one will counter) consciousness is, after all, associated with brains, and brains are physical systems, and the laws of physics *provide* an algorithm for computing anything that goes on in the physical world. So how can consciousness have anything non-algorithmic about it? This is where quantum gravity comes in. The fact is that there are problems with the laws of physics as they are currently understood. In particular, there is the problem of the collapse of the quantum wave function: exactly *when* do the many potentialities contained within a quantum description collapse into a single actuality? Penrose argues that this question cannot be made to go away just by finding the right interpretation of quantum mechanics. We need a new physics. Where is this new physics going to come from? Quantum gravity. He argues that the correct quantum gravity theory, when it is finally found, may solve the quantum measurement problem, and what is more important, it is likely to have a non-algorithmic character. The frequent quantum decisions that are made in the brain would then be made in accordance with a set of laws whose consequences are not computable. Hence the possibility of consciousness.

Many readers will not agree with Penrose's argument to the end. But such readers will still find in this book a wealth of interesting ideas to ponder. Many of these ideas come from physics, many from mathematics and the theory of computation, and many from empirical results on the functioning of the human brain. For example, one learns about "blindsight," the phenomenon in which a person whose visual cortex has been damaged can correctly *guess* what object has been placed in his blind area, even though he is not conscious either of seeing the object or of knowing what the object is. The book is intended for the informed layperson, but it goes more deeply into the physics and mathematics than one might expect of an essay aimed at such an audience. The reader will have to work hard to follow everything, but the questions are worth the effort.

—William Wooters

William Wootters, on sabbatical leave from the Department of Physics, Williams College, is in residence as a Visiting Associate Professor at SFI for the 1989-1990 academic year. His primary interest is in the program on complexity, entropy, and the physics of information and he is also working on cellular automata.

From Clocks to Chaos: The Rhythms of Life
By Leon Glass and Michael Mackey
Princeton Univ. Press, 1988, 248 pages

From its title, one might suspect that *From Clocks to Chaos* is yet another popularization of chaos for the general reader. Such popularizations have tended to review the beautiful computer-generated images of chaos, and the elegant mathematics of chaos, and often contain wild speculations on the relevance of chaos to science in general. By contrast, this book is about the concrete application of dynamical systems theory to the understanding of a wide range of biological phenomena. Questions addressed include: What are the mechanisms generating heart beats and other oscillatory behavior? What are the effects of externally administered stimuli on such oscillations? How are oscillations organized spatially within biological tissue? Wild speculations are avoided in favor of detailed presentations of experimental results and predictions from theory.

The authors are two of the most active researchers in this promising new field. Their stated objective in writing this book is to interest both theoreticians and experimentalists in the above questions. As I am one of the former, this review will reflect my reactions as a theoretician to this book.

Surprisingly little of this book is devoted to applying chaos to biology. A more dominant role is played by periodic behavior, as follows. Several theoretical mechanisms for periodic behavior are proposed: integrate and fire models, supercritical Hopf bifurcations (soft excitation), and subcritical Hopf bifurcation (hard excitation). The authors point out how these different mechanisms may be distinguished by their distinctive responses to external stimuli. Oscillations can be annihilated, or phase locked. A topological theory of phase resetting is explained. All these phenomena are illustrated by experimental data from biological experiments, and implications for the treatment of "dynamical diseases" are discussed. It is impressive to see theory and experiment so closely linked. The novel theoretical challenge is largely to understand how hypothesized dynamical models respond to external forcing, and consequently to design intelligent experiments to test the model. Although remarkably clear evidence for chaotic behavior in stimulated chick heart cells was presented, this seemed to be of little practical interest. The authors find no hard evidence for chaotic behavior in essential biological functions, in contrast to a well-known speculation that "life is transient chaos."

This book is unique in containing such a wealth of descriptions of biological experiments together with graphically displayed data. There is even a graph displaying data from male orgasm, complete with arrow and literature citation. As luck would have it, the theoretical explanation proposed in this case is the hard excitation mechanism. To keep the book to its modest length, no one topic was explored in much depth. For example, very few details are given for the derivation of the famous Hodgkin-Huxley equations for the squid axon. In fact there are very few equations at all in this book. Fortunately the book contains an extensive bibliography to the original sources for further reading.

A challenging theme left for future research is the large-scale application of computers to fitting much more complex models to data, and to the analysis of such complex models. I share the author's opinion that close collaboration between experimentalists and theoreticians is essential to meet this challenge, and that long-term funding must be found.

—Martin Casdagli

Martin Casdagli received his Ph.D. in mathematics from Warwick University and has worked at La Jolla Institute, the University of Arizona, Queen Mary's College, and Los Alamos National Laboratory. He is a Post-doctoral Fellow at the Institute. His primary interest is in nonlinear modeling and forecasting.

An Interview with L. M. Simmons

Creating the Right Environment

L. M. Simmons, Jr.

Theoretical physicist Mike Simmons became vice president of the SFI in 1986 and since then has devoted increasing amounts of his time to helping build the new organization and its research program. In 1988 he took leave from the Theoretical Division at Los Alamos to serve as executive vice president of the Institute. As deputy to the president, he shares oversight of all Institute activities, but has particular responsibility for the research programs, computing, and publications. The author of numerous articles on mathematical physics, elementary particle theory, and quantum theory, Simmons has also spent much of his career building and nurturing innovative research organizations. At Los Alamos he was Assistant and Associate Theoretical Division Leader and Deputy Associate Director for Physics and Mathematics. He originated, and was instrumental in founding, the Center for Nonlinear Studies. He is an Honorary Trustee, a member of the Advisory Board, a past Trustee, and past President of the Aspen Center for Physics. He recently shared his views on the SFI and on the process of nurturing an innovative multidisciplinary research institution.

How did you get involved with the Institute?

I wasn't present at the beginning, nor when the Institute was incorporated in 1984. One day early in 1986 George Cowan and Nick Metropolis came to my office at Los Alamos and asked if they could discuss the Santa Fe Institute with me. Now, I had been devoting most of my energy at Los Alamos to an ambitious project to create a new national laboratory for basic research, to be located deep underground and devoted to searches for rare and exotic events in particle physics and astrophysics. Although that project was not to succeed, I hadn't paid much attention to other things for a while. Although I count all of the founders among my friends, I knew about the Institute only through hearsay. Frankly, what I did know did not impress me. The Institute had a reputation for exclusivity; it had held only a couple of workshops; it had almost no money, no campus, and no employees; and everyone involved was working at least full time on other things. I simply told George that I didn't know anything about the Institute, and I probably didn't look very interested.

George said it was just as well that I didn't know much about the Institute because he wanted to describe the concept. He and Nick proceeded, for the next hour or two, to explain the rather grand concept being worked out by the founders. I thought it was fantastic. A new multidisciplinary institution that would try to overcome the artificial barriers between the disciplines was exciting. National laboratories are pretty good at interdisciplinary research and I understood that important things could result. Facilitating the emergence of new sciences was clearly of great importance, as was training a new generation of young scientists in these fields. I liked the idea of delocalized research networks that would regularly bring in a lot of carefully selected visitors and encourage continuing cooperation on a common research program. I had had a lot of responsibility for the very vigorous visitors program in the Theoretical Division; I'd had a hand in designing the Center for Nonlinear Studies with its strong visitors program; and I knew a lot about what produced the great successes of the summer research program at Aspen. So I understood a bit about how these things could be organized and encouraged. I also knew that bringing scientists together from different institutions and different fields could produce really new results if the surroundings were right.

By the end of our first discussion I said, "George, this is important; how can I help?" We had several discussions over the next few days or weeks. George described his concept of the Institute and I offered a lot of suggestions. I also spent many hours in discussion with Pete Carruthers, David Campbell, David Pines, and a few others. Not long after this I agreed to become one of the Institute's vice presidents.

"The Santa Fe Institute is small enough and flexible enough to be very quick on its feet. It's very important that it retain that adaptability."

It doesn't sound as if there was much to be vice president of.

That's right, but we had three vice presidents anyway. We had only one rented room in Santa Fe. We had more officers and committees than we needed but we were also enthusiastic and full of ideas. Everyone was committed to the concept of the Santa Fe Institute and we just worked very hard to make it happen. George Cowan and Art Spiegel began to have some success in raising modest amounts of money. Most of the people associated with the Institute were elsewhere but they were able to work on plans for workshops and funding. Locally there were really only a few volunteers trying to organize things. We formed a lot of committees to think about the program, computer needs, space needs, finances, governance, publications, and so on. But these were mostly just rearrangements of the same small group of people. David Pines spent a leave at Los Alamos and shared the jobs of chairman and president with George Cowan. They, together with David Campbell, Pete Carruthers, Dick Slansky, and I, plus a few others, kept meeting and planning and writing. That's when David Pines, Pete, and I produced the first issue of the Bulletin...pretty much by hand, or rather, by xerox.

You mentioned that national laboratories are good at interdisciplinary research. Why, from the perspective of a national laboratory, did you think the Institute was needed?

They are good at interdisciplinary research. To a large extent that's why Los Alamos exists, to bring the tools of many different disciplines to bear on very hard problems. It's a very creative place and people move across disciplinary lines very freely. Under Pete Carruthers with, I hope, some help from me, the Theoretical Division became one of the best places in the world for theoretical physics. It has an academic atmosphere and many excellent researchers have an opportunity to pursue their own ideas. It is also multidisciplinary with flourishing programs in most of theoretical physics and chemistry, much of mathematics, and theoretical biology. But national laboratories quite properly have specific missions and it can be difficult to attack a problem that doesn't fit within the mission. A laboratory is also embedded within a large bureaucracy which has its own disciplinary or programmatic barriers and which doesn't change very rapidly. The Santa Fe Institute is small enough and flexible enough to be very quick on its feet. It's very important that it retain that adaptability.

Aren't universities open to new ideas?

Of course they are, and they are wonderful and important institutions. They're just not very good at changing their structure. The disciplinary structure of the modern research university exists for good reasons but real cooperation across that structure is very difficult. People from biology don't often work in the economics department even though, here at SFI, we find that the best people in the two disciplines have a lot to say to each other. The SFI can demonstrate to the universities that cross-disciplinary interactions such as these are extraordinarily productive. Maybe we can export some techniques.

You mentioned your experience with the origins of the Center for Nonlinear Studies and its influence on your reaction to the Santa Fe Institute.

Yes. There are a couple of points to be made. The first has to do with recognizing and exploiting opportunity. The second has to do with the visitors program, or the use of borrowed personnel to develop a research activity and simultaneously to strengthen the lending institutions. In the late '70's I noticed that there were scientists in several T-Division groups and in other parts of the laboratory using the same newly emerging tools of nonlinear science to address very different problems. I began to think about how to encourage cooperation among these people. It required a new structure to succeed. As I began to discuss this with Pete Carruthers, Basil Nichols, David Campbell, and others it became clear that the laboratory had much to gain by the creation of a new entity that would provide a medium for interchange within the existing structure of the laboratory and with academia. Under the outstanding leadership of David Campbell and Al Scott the Center became a great success. It operates solely with borrowed personnel, and with a large number of postdoctoral fellows and external visitors. The existing organizations in the laboratory benefit from this success and are strengthened, not weakened by the Center.

The founders of the Santa Fe Institute recognized another larger and more important target of opportunity in the newly emerging and closely related sciences of complexity. I strongly believe that the existence of the SFI will strengthen the educational and research institutions of the nation. We are encouraging collaborations, hastening the development of new ideas, attracting young scientists to these fields, and returning the scientists whom we borrow to spread these new ideas on their home campuses. Even if we take something from the system, money and an occasional professor, for example, the Institute will return more to the system than it takes. The analogy with the CNLS could scarcely be stronger. This is one of the things that I recognized immediately in the concept that George laid out for me and it's one of the things I found exciting about the idea of the SFI.

Are there similarities between the Institute and the Aspen Center?

Yes, there are a few, and some differences. The Center concentrates only on physics whereas the Institute is broadly interdisciplinary. Both, however, are highly selective and are characterized by intense and informal interchange of ideas out of which new collaborations grow. And both seek to provide an atmosphere free from the distractions of home institutions where researchers can concentrate on working out their ideas. As the Center has done for theoretical physics, I hope the Institute will provide a home for a broad network of scientists working on complex systems, a place where they will return from time to time to renew collaborations and to find fresh stimulation.

"One of our goals is to keep a rich mix of specialists here, but to choose specialists with broad interests outside their own fields."

You have devoted considerable time and effort to the care and feeding of scientific organizations. Why?

For several reasons, I suppose. It can be fun, and I think it is very important. No one who has tried to work with a badly organized or managed institution would argue that structure and management are unimportant; they can make all the difference. Of course, the most important thing is getting very good people. That's essential, but it's not quite enough. The structure must be carefully thought out and arranged to encourage the most productive scientific interactions, but it must also be unobtrusive, almost invisible to the scientific participants. The right environment is critical; that's what we're trying to create here at the Institute. There is a unique spirit associated with a successful scientific organization and it is very fragile. With inspiration and effort a scientific institution can be designed and nurtured to encourage creativity and communication and to make it possible for the participants to work more productively. There is a special satisfaction in participating in that effort. In the end, contributing to the creation or growth of an important research institution produces effects, through the work that is enabled, out of all proportion to the effort one expends.

Let's return to the Institute. How does the Institute break down the barriers you referred to and encourage collaboration?

What we are doing is unique and it's quite new. Honestly I don't think we fully understand the process yet, but some of the necessary elements we do know. We begin a new program with a workshop. The participants are carefully selected for excellence, for their accomplishments, and for an open-minded approach. Communication across the boundaries of the established fields is difficult because of the barriers of specialized vocabularies. The same word may be used with slightly different meanings in different fields. It takes time and mutual respect to tease this all out. We get people together for as long as possible, a week or even two, in an atmosphere of considerable informality. The talks are "highly interruptible" and stimulate lively multi-party exchanges. I imagine this is a bit frustrating for the speaker, at least at first. But if we've chosen a sufficiently fascinating and important problem, and the right mix of people, then something important emerges. This process itself is a complex system.

As projects and the SFI itself expand, do you find there is a certain amount of tension between the pull of specialization and the attempt at integration?

Yes. And this aspect of the program needs careful tending. But many of the scientists at the SFI are associated with more than one program. This helps to assure that the ideas and problems of one field are seen in the light of knowledge and techniques from several other fields. One of our goals is to keep a rich mix of specialists here, but to choose specialists with broad interests outside their own fields. A flow of visitors through the Institute helps to assure that the insights and tools of several different fields are always in play.

Another problem is that as programs grow and as our budget grows we need some management structure. Authority and responsibility for research programs and their budgets must be delegated. But as the overall research program is divided by these practical necessities it may become fragmented. All this can work counter to the integration we seek. This is destructive to the creative process of synthesis that we want to nurture.

How do you combat this?

Ultimately it may not be possible but I don't think it's an imminent danger. The synthesis of new sciences we seek may, a generation or two hence, result in artificial barriers between these new fields. In the interim we will have solved some very important problems. I intend to be vigilant in guarding against compartmentalization of our program. One of my tasks is to watch over what we call "Other Research"; that is, research that does not fit within one of our programs, such as economics or immunology. This is the category into which a core cadre of scientists with broad interests in complex systems fits. We also need bright young postdocs and graduate students to shake things up. It's absolutely essential to our success.

Another thing that helps a great deal is the continuing series of integrative workshops on adaptive complex systems that Murray Gell-Mann and David Pines lead. These draw participants from each of the more specialized programs and keep us all thinking about the broad themes, the common threads that run through the entire program of the Institute.

All we can really do is to be constantly watchful and to resist the many temptations to place a foot on the slippery slope that leads toward compartmentalization of the program.

What are you doing about postdoctoral fellows?

Last year we had John Miller in economics. This year we have Martin Casdagli, a mathematician and dynamical systems theorist, and Wentian Li, a specialist in cellular automata. I hope we can add more postdocs in the future. The postdoctoral program is one of my real enthusiasms. I keep trying to insert extra postdocs into the budget during our planning process. It's the young scientists who can move quickly into these new fields and who will help define them. Certainly we need senior people here; they provide leadership and they provide advocacy for the young people. But my own inclination is to keep the number of senior scientists down to the minimum necessary and devote more of our resources to young scientists.

Does that include graduate students?

Certainly. The goal of the founders was to start an institution for graduate education in these new sciences. Without a permanent faculty, or at least faculty continuously in residence for some years, it's not easy to see how to get started. We struggled with this problem and then the answer found us. Graduate students

> "If the newly emerging sciences of complexity are as important as we say they are for the next century, then we have an obligation to train young people to work on these problems."

from established universities are beginning to spend more time here. Many come to the complex systems summer school, but a smaller number come to engage in research projects. I think their numbers will grow rapidly as we find the resources to support more resident researchers. The only limitation now is money. As more research scientists are able to spend significant amounts of time here more graduate students will come to join research groups, and some will begin to do thesis research here.

How does the summer school fit into the program of the Santa Fe Institute?

As I said, education has always been a high priority. From the beginning Pete Carruthers, David Campbell, George Cowan, David Pines, and I have been fans of the summer school. It is very much a consortium effort, the principal players being Arizona, Los Alamos, and SFI, but with important support from a larger group of sponsoring institutions. Within the SFI I've tried to keep it as a high-priority item. In the hands of its first two directors, Dan Stein and Erica Jen, the school became an instant success. It's our joint beginning on the task of systematically educating a new generation in the foundations of the sciences of complexity. Its published lectures, we hope, will form a pedagogical basis on which others can build. In the future, I'd like to see the summer school spin off some more specialized schools, summer and winter, and encourage the students to use the summer school as a basis for starting research projects that can continue at SFI or elsewhere. This research component has already begun to develop within the 1989 and 1990 schools. It's out of our joint experience with the summer school that the cooperating institutions, SFI included, will develop a real curriculum in complex systems. Then we can begin seriously to train graduate students.

But the Institute does not intend to admit undergraduates?

We're not very interested in the formal academic credentials. It's the talent and the motivation that matter. Undergraduate students have begun to discover us. Driven by their demand to be exposed to what we are doing, I started a program of undergraduate internships. The program is experimental and of necessity is small, but so far it's working well. Students may engage in reading and study under a scientific mentor or take on a research problem. The only constraint is that the internship must further the student's education.

Education has always been seen as an important, even an essential, part of the Institute's program. If the newly emerging sciences of complexity are as important as we say they are for the next century, then we have an obligation to train young people to work on these problems.

What do you see as the Institute's greatest strengths?

The first is the incredibly interesting and important array of ideas that forms our program. The themes of adaptation and learning, information processing, adaptive searches in high-dimensional spaces, self organization, and nonlinearity in all its ramifications permeate our programs and are a part of the common fabric of the emergent sciences of complexity. Our second asset is an outstanding Science Board and External Faculty, the source of that rich set of ideas. And third, we've been very fortunate in attracting an exceptionally talented and committed support staff. The Institute could not have accomplished so much scientifically in such a short time without the outstanding achievements of our staff.

The main threat I see facing the Institute is that the scientific challenges we face will outstrip the resources we can muster to meet them.

What does the future hold for the Institute?

We are in the process of growing a very strong research program based on the collaboration of a rich mix of junior and senior scientists from a broad set of disciplines, ranging from the biological sciences to mathematics, economics, computation, and physics.

At the moment our program concentrates heavily on workshops. We use these both to define new areas of work and to keep the interchange going in established networks, such as that in economics. Workshops will likely always be a feature of SFI activities, but I expect the number to decrease somewhat as we develop a stronger residential research program.

Our short-term visitor program is healthy, but I'd like to see a larger number of people appointed to long-term visiting positions, a year or so. Brief visits keep that important flux of ideas going but longer visits allow the time for a visiting scientist to really develop new research ideas. That's important to the Institute and will also enrich the home campus.

Within the next couple of years the Institute is likely to make its first faculty appointments and thereby to establish points of stability on our research networks. This will also provide the continuity to begin graduate education in earnest. It's reasonably clear that we'll begin, as we already have, by hosting advanced graduate students for thesis research. The granting of degrees of our own is in the plan but I can't predict how that will work out.

Under the guidance of Ronda Butler-Villa, the book publications program is developing well. I don't foresee a natural limit there. As the sciences of complexity emerge, more material will continue to emerge and we can look forward to a continual stream of influential books. At some stage it will be appropriate to begin a journal of complex systems studies. I think that is definitely needed, but I don't know how soon it will begin.

We have outgrown our current quarters. That's likely to require a somewhat disruptive move to temporary space before we can acquire and develop a permanent campus. With a bit of luck we may be able to announce in the next issue of the Bulletin the location of that permanent campus and I very much hope that we'll be there in a few years.

If this sounds very optimistic, it's an attitude influenced by my enthusiasm for what we're doing at SFI.

Brian Arthur on the SFI Economics Research Program

Developing a "Santa Fe Approach"

After directing the Economics Research Program for the last year and a half, Brian Arthur is returning to Stanford. Arthur has been with the Institute continuously since June 1988, when the initial loose network of researchers that had grown out of the 1987 Arrow-Anderson-Pines workshop on the economy as a complex, evolving system transformed itself into a residential research program. In that time, involvement in the program has grown from an initial twenty or so economists, physicists, and other specialists to about forty researchers. Some have spent only a few weeks at the Institute; others six months or more.

The Santa Fe Approach

The program has generated a lot of attention. "I was surprised," Arthur says, "that the program turned out to generate so much interest in the economics community. I believe part of the reason that the Santa Fe ideas have struck some resonance in the profession is lucky timing. Economics in the last few years has been changing slowly—discovering the importance of nonlinearities or what economists would call non-convexities. The other part is that Santa Fe was able provide both the experts and the techniques to work some of these ideas through in economics."

Arthur believes he can see the beginnings of a "Santa Fe approach" to economics. "If there is a theme in our work around here" he says, "it is that we are looking at the economy not as a static entity that displays certain equilibrium quantities and prices, but as an evolving system where new patterns and structures are constantly emerging, often in part by chance. These patterns in the economy change, die away, and sometimes re-emerge much as in a physical fluid. What excites me is that this viewpoint can be made rigorous. We are beginning to develop the mathematical techniques to build this process-view of the economy into the heart of the subject."

Reaching Equilibrium

Much of the time Arthur was here participants in the program spent exploring themes of "learning" in the economy. "We needed to do that as our first priority," explains Arthur. "In economics we have very sophisticated notions of what constitutes a solution—an equilibrium. But how an equilibrium is arrived at is for the most part unexplained. We assume for example that the institutions and traders in the economy can perform extremely far-sighted calculations, that they know that others can do so, and therefore that a certain solution can be foreseen, coordinated upon, and realized. What happens when traders are not infinitely smart, or when they can't assume that others are infinitely smart and therefore predictable, is simply unknown."

The program has been looking at how equilibria are reached, how one solution is selected when several alternatives are possible, and what might happen when the players in the economy are not superrational, by assuming that agents in the economy who are poorly informed and capable of less than perfect calculations can acquire information and learn from their actions as time progresses.

Projects within the program concerned with learning and adaptation in the economy include the work of Kenneth Arrow and Frank Hahn on bounded rationality; the experimental work of John Rust, Richard Palmer and others on learning in double auctions; the work of John Holland and others on genetic algorithms; the work of Tom Sargent, Ramon Marimon, and Ellen McGrattan on machine learning applied to economic problems; the work of William Brock, Sargent, Marimon, and Rust on learning in games; and the work of Arthur and David Lane on product choice under incomplete information.

The analysis of learning problems is challenging. Most learning models are stochastic in that actions taken and observations gained are to some degree random, and they are nonlinear in that behavior that is frequently taken is thereby further reinforced. Analysis of the dynamics of learning therefore requires sophisticated methods from nonlinear stochastic process theory.

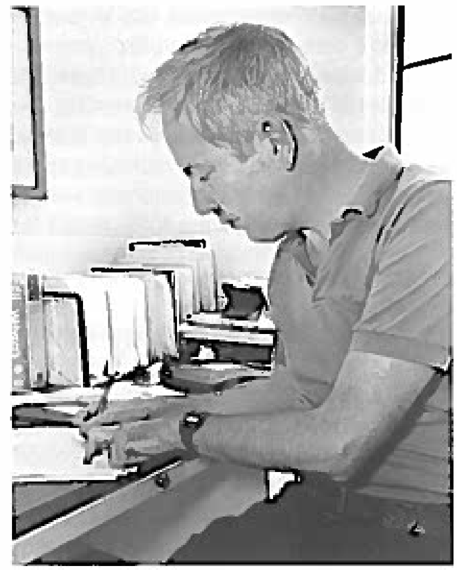

W. Brian Arthur. Photo by Cary Herz © 1989.

Holland on Learning

"Several of us were particularly interested in John Holland's ideas on learning," says Arthur. "By modeling agents' actions by sets of behavioral rules that change and adapt, we can apply machine-learning algorithms from automata theory. We are therefore not restricted to describing behavior by parametric systems that are recursively updated. Instead, we can model adaptive processes as the interaction and change of massively parallel, interlinked, competitive, and possibly inconsistent sets of rules." Holland's classifier system uses particular rules more heavily if they lead to appropriate outcomes. His genetic algorithm also "discovers" new, possibly useful rules, and adds these to the system. Useful rules then form into "default hierarchies" and causal chains

Economic Research Program Papers

There is a limited supply of the following papers available to students and researchers. If you would like a copy of one of these papers, please send a postcard with your name, mailing address, and the number of the paper that you would like to receive to Publications, Santa Fe Institute, 1120 Canyon Road, Santa Fe, NM 87501.

89-001	"A Double Auction Market for Computerized Traders" John Rust, Richard Palmer, and John H. Miller
89-002	"Communication, Computability and Common Interest Games" Luca Anderlini
89-003	"The Coevolution of Automata in the Repeated Prisoner's Dilemma" John H. Miller
89-004	"Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents" Ramon Marimon, Ellen McGrattan, and Thomas J. Sargent
89-005	"The Dynamical Behavior of Classifier Systems" John H. Miller and Stephanie Forrest
89-006	"Nonlinearities in Economic Dynamics" José A. Scheinkman
89-007	"'Silicon Valley' Locational Clusters: When Do Increasing Returns Imply Monopoly?" W. Brian Arthur

less useful rules and chains are gradually dropped. "In applying such machine learning ideas," says Arthur, "we are further developing them and streamlining them for use in economics where prediction and strategy are major components. We are also interested in neural networks—structures that again allow parallel action, adaptation, and the emergence of behavioral patterns, but have done very little with them as yet."

Holland's system suited many learning situations in economics. Economic agents must choose, whether in an auction-bidding process or in a normal-form game, among several competing decisions or actions leading to outcomes (or payoffs) that are in advance to some degree unknown. Learning in this context takes place as the agents update their probabilities of triggering each action on the basis of the payoffs or outcomes they experience. These actions need not be simple rules of thumb, they may be highly sophisticated "subroutines" of behavior, and they may be concatenated, with one action taken triggering choice among others. Such a scheme can then be general and sophisticated enough to model learning in a variety of neoclassical economic problems.

An algorithm for updating probabilities of choosing among competing actions is called a *learning automaton.* In the last year, Arthur has developed and studied a class of learning automata based on Holland's ideas in machine learning. He has been able to show theoretically that these automata can learn their way to choosing consistently the "optimal" action in the "multi-arm bandit" problem of choosing among actions with unknown, random rewards. Where several such learning automata are used to model players in an iterated game, his analysis shows that when the resulting set of strategies converge, they converge to the standard game theory solutions—Nash equilibria. Holland, Palmer and Arthur have used a collection of such automata to represent agents in a learning version of Robert Lucas's stock market. Starting from random behavior, the computer-algorithm "agents" learned when to buy and sell stock appropriately, and the model converged rapidly to price fluctuations around fundamental value.

"What fascinates me," says Arthur, "is these learning models do not merely reproduce standard perfect-rationality equilibria. In our stock market model, our agents quickly learned to buy at a price below fundamental value, and sell above it, as standard theory would predict. But I was surprised one morning to find that they had discovered a form of technical analysis. They were buying in at a *high* price, which then became a 'support price.' Every time the market hit this price, our computer agents had learned that the price would go up again. So they bought in, and enough of them bought in to make this a self-realizing prophesy. I suspect that such non-standard 'psychological' outcomes will appear in other models that include learning."

Plans for 1990

Work on the theme of learning and adaptation in the economy will continue in the program when Arthur leaves. Planned for 1990 is a second summer working group that will take up some of these issues. The Program, under a steering committee consisting of Kenneth Arrow, Philip Anderson, Brian Arthur and John Holland, will appoint a new director for the academic year 1990-1991. Other themes—non-convexity in the economy, the emergence of "webs" of technologies, self-organized criticality in industrial development—are also on the docket for closer exploration.

Looking at 1990

The Institute's academic program will continue to grow sharply throughout 1990. Half a dozen research networks are underway, resulting in a variety of small meetings and conferences already scheduled for the new year. In March SFI's Science Board—the Institute's interim senior faculty and academic advisory body—will meet to review current programs as well as to consider new initiatives. Based on past years' experience, SFI's academic schedule will become even more crowded after this annual Spring assessment. In the meantime, here is how the next year looks from the current perspective.

Mathematical Approaches to DNA

Things will be off to a quick start with the January workshop "Mathematical Approaches to DNA," co-chaired by Nicholas Cozzarelli, University of California, Berkeley, and Sylvia Spengler, Lawrence Berkeley Laboratory, and hosted by SFI. The chairpersons are part of the Mathematics and Molecular Biology Program funded by the National Science Foundation. More than one hundred mathematicians and biologists are expected to attend this four-day meeting, which will focus on the topology and geometry of DNA including energetics and alternative structures. Other sessions concern questions of algorithms and other concerns in mapping and matching DNA sequences, that have come out of sequencing large pieces of DNA, including the human genome.

Artificial Life

In early February the second Artificial Life workshop—co-chaired by Doyne Farmer, Christopher Langton, Steen Rasmussen, and Chuck Taylor, all from LANL, and co-sponsored by SFI and the Center for Nonlinear Studies—is expected to draw more than three hundred scientists to Santa Fe. Artificial life is the study of physical or computational systems that exhibit behavior characteristic of natural living systems, a quest to explain life in any of its possible manifestations, without restriction to the particular examples that have evolved here on earth. The first workshop, co-sponsored by SFI and CNLS and held at Los Alamos in September of 1987, identified the essential theoretical and practical problems to be solved in order to bring about artificial life. This second workshop will include a mixture of lectures, live demonstrations, posters, panel discussions, and a contest for artificial life forms; topics include self-organizing structures, collective phenomena, and emergent behavior; the origin of life; natural, artificial, and cultural evolution; extra-terrestrial life; computer viruses; and the social and philosophical implications of artificial life. As at the first meeting, one of the central events will be an artificial "4-H" show, with prizes for the "liveliest" artificial life forms.

Economic Meetings

The Institute's work in economics, adaptive computation, and time-series forecasting is described in detail elsewhere in this issue. Each of these programs will feature both residential research and network meetings throughout 1990.

In May Daniel Friedman, University of California, Santa Cruz, and John Rust, University of Wisconsin, will co-chair

Researchers in Residence, Spring 1990

Visitors/External Faculty Members

LUCA ANDERLINI, Economics, St. John's College, Cambridge
KENNETH ARROW, Economics, Stanford
TED BERGSTROM, Economics, University of Michigan
MICHELE BOLDRIN, Economics, University of California, Los Angeles
DANIEL FRIEDMAN, Economics, University of California, Santa Cruz
FRANK HAHN, Economics, Cambridge University
GARY HANSEN, Economics, University of California, Los Angeles
JOHN HOLLAND, Computer & Engineering Sciences, Psychology, University of Michigan
STUART KAUFFMAN, Biophysics & Biochemistry, University of Pennsylvania
CATHERINE MACKEN, Mathematics & Statistics, University of Auckland
JOHN MILLER, Social & Decision Sciences, Carnegie-Mellon University
NORMAN PACKARD, Center for Complex Systems Research, University of Illinois, Urbana
DAVID PINES, Physics, University of Illinois, Urbana
JOHN RUST, Economics, University of Wisconsin
JOSÉ SCHEINKMAN, Economics, University of Chicago
PAUL TAYLOR, Deloit, Haskins and Sells, London
GÉRARD WEISBUCH, Laboratoire de Physique, France
WILLIAM WOOTTERS, Physics, Williams College

Postdoctoral Fellows

MARTIN CASDAGLI
WENTIAN LI

Graduate Students

DAVID CAI, Northwestern University
MICHAEL ANGERMAN, New Mexico State University

Undergraduate Interns

JULIE PULLEN, Macalester College
JULIE REHYMEYER, St. John's College

"Price Dynamics and Trading Strategies in Double Auction Markets." The meeting, following shortly after the Institute's Double Auction Tournament, will bring together an interdisciplinary group of economic theorists, experimental economists, econometricians, learning theorists, and practitioners from real world markets such as the Chicago exchanges to better understand the "invisible hand" in the double auction institution.

The Economics program again will convene a summer study group lasting from early July through mid-August. The framework for this year's summer program will be somewhat different from last year's. In 1989 there was quite a bit of research collaboration going on, but the format was dominated by a large number of talks. In 1990 the group will have fewer talks and a smaller group of participants and will concentrate on research projects. Topics to be considered are learning (in particular using classifiers and genetic algorithms) in markets and games; nonlinear and non-convex phenomena in the economy; self-organized criticality; the evolution of economic structure; neural networks in economics; and nonlinear stochastic dynamics. In September the Economics network will hold its fourth annual program review.

Nonlinear Modeling and Forecasting

The Institute's initiative in adaptive computation is off to a strong start. Beginning in January SFI External Faculty member John Miller, Carnegie-Mellon University, will be at SFI for eight months as resident director for this effort. Miller will be working with John Holland, Stephanie Forrest, and additional visitors throughout the year. It is also possible that at least one program network meeting will be held at SFI in 1990.

In the Fall of 1990 SFI will sponsor a workshop "Nonlinear Modelling and Forecasting" as part of its computer forecasting and modelling techniques network. It will be chaired by Martin Casdagli,

Russell Mittermeier, President of Conservation International. Dr. Mittermeier participated in the recent SFI workshop "Environmental Issues, Economics and Public Policy" and gave, with Mark Plotkin, a public lecture on conserving tropical rain forests.

Stephen Eubank, and Doyne Farmer. This meeting will be one of the first to bring together experts in nonlinear modeling and forecasting from a wide variety of disciplines including statistics, information theory, neural nets, function approximation, and signals and radar establishment.

Information and Entropy

Two other SFI research networks have meetings currently scheduled for 1990. In mid-April the "Complexity, Entropy, and Physics of Information" (CEPI) network will hold a workshop at SFI, chaired by the network leader Wojciech Zurek, Los Alamos National Laboratory. Discussions will focus on the information-entropy connection; quantum measurements and the interpretation of the wave function of the universe; physics of computation and its implications for the nature of physical laws; algorithmic complexity as well as alternative measures of complexity; and the evolution of complexity in the course of the history of the universe.

Pattern Recognition

The Institute has formed a new research network "Pattern Recognition in Biological Sequences." It will consider questions such as "How is it possible to understand even relatively simple aspects of the human genome from mere strings of symbols?" or "How can one understand the structure and function of proteins given their amino acid sequence?" The network will promote interaction among an interdisciplinary group of people interested not so much in the mechanics of the sequencing effort, but rather with what will actually be done with the billions of base sequences as they are obtained. This effort requires close interaction among disciplines that have previously not had a record of close collaborative efforts. The Steering Committee for this effort is Alan Lapedes (Los Alamos National Laboratory and Santa Fe Institute), Christian Burks (Los Alamos National Laboratory), and Gary Stormo (Molecular, Cellular and Developmental Biology Department, University of Colorado). A small workshop is scheduled in the spring of 1990 with the aim of introducing researchers in separate disciplines (molecular biology, neural networks, computer science) to each other and setting the stage for future collaborations.

Glasses

In May, 1990, a "Glasses, Biomolecules, and Evolution" workshop, cochaired by Hans Frauenfelder and Robert Young, University of Illinois, will examine mounting theoretical and experimental evidence that many, maybe all, complex systems exist in a very large number of conformational substates, valleys in the conformational landscape corresponding to slightly different structures. Glasses, spin glasses, proteins, nucleic acids, evolutionary systems, and neural nets all exhibit conformational substates. The unifying concept of a rugged landscape suggests deep connections among complex systems. At this workshop scientists from various fields will explore the energy or fitness landscape of typical systems. Most of the week-long workshop will be spent on in-depth discussions of novel approaches and techniques in the various fields. The outcomes of the workshop ultimately depend on the participants. Possibilities include formation of a research network; future workshops, symposia, or schools; and a workshop proceedings as part of the series published by the Santa Fe Institute.

Organization of Organisms

In June, 1990, the workshop "Organization of Organisms," chaired by Jay Mittenthal, University of Illinois, is scheduled. Recent progress in diverse fields suggests that a small set of principles may underlie the organization of complex adaptive systems, such as organisms and artificially intelligent systems. To persist, a complex adaptive system must meet many constraints—relatively invariant conditions for its survival. In an organism, structures perform processes that meet constraints. The processes are coupled in a network. The overall aim of the workshop is to characterize principles that relate constraints to the architecture of organismal networks of processes. Approximately thirty scientists—fifteen senior and fifteen junior level researchers—will participate in the workshop, bringing expertise in diverse fields, including theoretical biology, developmental biology, physiology, functional morphology, evolution, and computer science. It is anticipated that a proceedings volume will result; another potential product is an article in the *Quarterly Review of Biology*.

Summer School

The third annual Complex Systems Summer School will again operate for four weeks beginning in June, on the St. John's College campus. Co-Directors for this year's school are Lynn Nadel and Daniel Stein, both of the University of Arizona. The themes for this year's school will include the nature of chaos, disorder or randomness in a variety of physical and mathematical systems; pattern formation in physical, chemical and biological systems; and cognition and computational approaches to brain and cognitive function. The format will be slightly different from past years. The program will consist of a less intensive schedule of lectures, leaving more time for study, research, and work in the computer laboratory. Student research projects will be spread out over the four-week course of the school. An edited lectures volume is once again expected to result from the School.

Prehistoric Southwestern Society

In the fall of 1990 SFI plans a full-scale workshop "The Organization and Evolution of Prehistoric Southwestern Society," co-chaired by George Gumerman, Southern Illinois University, and Murray Gell-Mann, California Institute of Technology. The planning workshop for this program was held in September 1989 in conjunction with the School of American Research in Santa Fe. That gathering reviewed, from virtually all relevant points of view, the evidence on abandonments and aggregations, increases and decreases in complexity in the prehistoric Southwest as a whole. Emphasis was on behavioral topics, rather than on sub-regions within the Southwest, each scholar working with several other researchers to address a specific behavioral topic since no single individual is knowledgeable about a subject for all the prehistoric Southwestern traditions. A book will result, to be published in the School of American Research series.

That book, in draft form, will be avail-

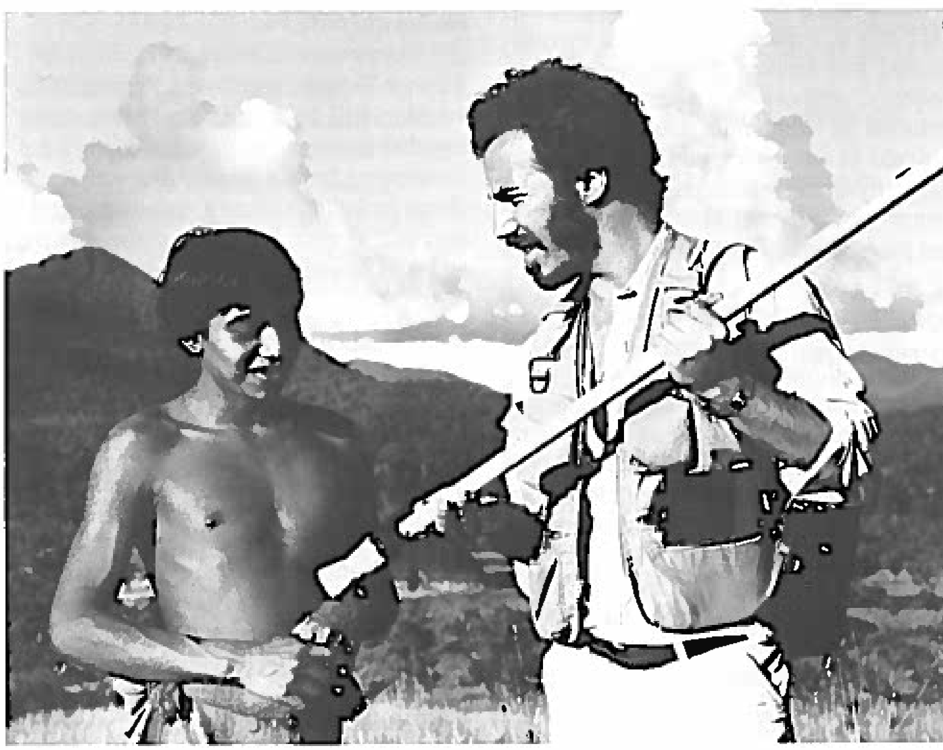

Mark Plotkin, Vice-President of Conversation International, a participant in the Fall, 1989, meeting "Environmental Issues, Economics, and Public Policy." The workshop considered how economics might be integrated into public policy studies on environmental questions using neotropical forest conservation as a case study.

SFI Undergraduate Intern Program

SFI has initiated an Undergraduate Internship program, the aim of which is to inform talented undergraduate students about the sciences of complexity and to begin to educate them in these subjects. As part of this effort the Institute will bring in a limited number of students to work on SFI programs or to participate in a reading and study program under the guidance of a visiting faculty member, postdoctoral fellow, or external faculty member. There are many avenues open: a course of supervised reading and study; investigating one or more carefully chosen research problems; developing a piece of computer code for a current research program; or become competent on the Unix system and in a computing language so that research participation is an option. The only constraint is that the internship must contribute to the education of the student.

With the inauguration of the undergraduate internship program, participation in the graduate-level Complex Systems Summer School, and graduate student residencies for thesis research, along with its sponsorship of postdoctoral residential research, SFI is taking an active role in complex systems education at all post-secondary levels.

We are pleased to welcome two undergraduate interns to the Institute during the Spring term, 1990.

Julie Pullen is a junior mathematics and physics major at Macalester College, St. Paul, MN. Last summer she received a Pew Memorial Trust undergraduate research grant for work at the University of Chicago, where she worked in the solar energy group. On her own she has been reading about chaos, synergetics, and complex systems. She joined the SFI undergraduate intern program for a month of reading and study under the Macalester Interim program. Her mentor is External Professor Stuart Kauffman.

Julie Rehymeyeris a sophomore at St. John's College in Santa Fe. From her courses she has developed a deep interest in science and is using the SFI undergraduate intern program to broaden her scientific education and to gain some first-hand experience with research. Ms. Rehmeyer joined the program in December and works part-time at the SFI. Under the tutelage of External Professor Doyne Farmer and others, including Richard Bagley, John Gibson, and Michael Angerman, she is studying dynamical systems theory, learning the Unix system, and learning the C and C++ languages to prepare for work on the nonlinear modeling and forecasting project.

—LMS

able to members of the autumn 1990 workshop on the prehistoric southwest as a complex evolving system. The general goal of this conference is to use the results of the earlier meeting as a basis for addressing general questions regarding the organization and evolution of relatively simple societies. The program will concentrate more on the theoretical aspects of the subject and will seek to draw ideas from the biological, social, economic, and computational sciences for use in model building. The aim is to provide some explanations for the nature of culture change and to set the agenda for future data collection, use of new analytical techniques, and the generation of models applicable to archaeological testing.

Toward a Sustainable Human Society

During 1989 the Institute held several informal planning meetings in preparation for what will probably be its most broadly cast topical program to date, "Multiple Paths Toward a Sustainable Human Society," a research program chaired by SFI President George Cowan. The Institute plans a major conference on this topic in mid-1990. Its purpose will be to consider the various means that are available and that are potentially most useful for defining policies and implementing them at every level of social, political, and economic organization. The idea is to take a crude look at the whole problem of how human beings might approach a sustainable human society that avoids catastrophic war and degradation of the biosphere, while satisfying human needs and aspirations. Are there institutions, forces, and tendencies in the world today that, if suitably encouraged, can lead us to sustainability in population, the environment, and economic, social and political relations among individuals and groups?

As presently planned, the first part of the conference may consist of one or more talks that outline a minimum consensual list, a "world view," of policies and actions necessary to achieve a sustainable and desirable human society. It would be followed by a panel discussion of the proposed world view with representation by participants who might depart from the presumed consensus. Modifications to the world view would be incorporated with consent of the initial presenters during the critique. The world view would address the following issues: individual, societal, and political beliefs and aspirations; economic interactions, including scientific, technological, and demographic factors; and environmental and health issues.

Integration of Programs

In addition to these meetings on specific aspects of complexity, the Institute also will continue its series of meetings on the "Nature of Complex Adaptive Systems." These workshops, led by SFI Science Board Co-Chairmen Murray Gell-Mann and David Pines, bring together key participants in Institute programs and workshops to focus on the common threads of complexity running through the various programs and the overarching themes of the Santa Fe Institute. It it likely that several of these meetings will take place this year, one in the form of an extended meeting of the Science Board of the Santa Fe Institute in March, 1990.

—GR

A New Initiative
Adaptive Computation

Large-scale, complex systems, such as those that have caught the collective attention of the Santa Fe Institute, provide researchers with a multiplicity of challenges. They must establish new, precise, cross- disciplinary vocabularies and use them to clearly define problems for study. The research programs developed to address these problems must be broad enough to admit a range of possible solutions. Given the highly complex nature of the systems under scrutiny and their tendency to exhibit nonlinear dynamics and constantly evolving characteristics, it is not surprising that the emerging potential of adaptive computation is already being harnessed to support integrated research efforts in several fields. Adaptive computation is also emerging as a fundamental field, providing material for a developing research program of its own.

Most computer programs, however complicated, are linear and sequential by nature and tend to follow a deterministic set of rules from beginning to end. Thus "adaptive computation" may appear oxymoronic to some. But the adaptive aspect allows for the innovative research approach of building computer programs capable of fundamentally altering their own rules and structure while exploring various paths to problem solution. In the study of complex systems, this is very useful.

Broad Application

"It is already clear that adaptive computation underlies very much of what we want to do at SFI. As the new program on adaptive computation unfolds, I believe we will discover that it is even more broadly applicable," comments L. M. Simmons, SFI Executive Vice President. Initially research efforts in the adaptive computation program will center on the genetic algorithms and classifier systems developed by SFI Science Board Member John Holland. "This area of adaptive computation has already had a deep influence on our thinking, in the economics program in particular," Simmons says.

He adds that there is other SFI work on computational systems capable of learning and adapting. Among the best developed are the neural networks of Alan Lapedes and Robert Farber, SFI external faculty members from the Complex SystemsGroup at Los Alamos. Another example is embodied in the time-series forecasting work of Los Alamos staffer and SFI External Professor Doyne Farmer and his collaborators; their system for nonlinear modeling and forecasting is a powerful technique for predicting the future behavior of a system. These systems can be regarded as very general function-fitting techniques and in this sense are related to each other. As Simmons points out, "In a real sense, any computational system capable of sufficient revision of its structure and output in response to inputs is engaging in adaptive computation."

Simmons suspects that there are other computational systems yet to be explored that will "exhibit some of the same characteristics of learning and adaptation in response to information provided by an external environment. Perhaps we will recognize the current systems as limited realizations of a much more general adaptive computational system." The hope of learning more about the fundamental principles underlying adaptive, complex systems is one reason Simmons feels it is important for the SFI to engage in systematic research in the area of adaptive computation.

Beginning Efforts

Certain aspects of the program are already in place. In Santa Fe in 1989, John Holland and economist Brian Arthur initiated work on a classifier system model of the stock market. During 1990, Holland plans to be back in residence at the SFI to work on classifier systems with External Faculty Member John Miller, from Carnegie-Mellon University. Providing on-site continuity to the program, Miller will spend eight months at the SFI beginning in January 1990. Stephanie Forrest, a postdoctoral fellow at Los Alamos National Laboratory will also be working with Miller. As the program progresses, additional visitors will be brought in to participate and, as is often the case at the SFI, participants in other programs will undoubtedly contribute as well.

Some of the SFI programs closely related to adaptative computation include research programs in economics, theoretical immunology, evolutionary biology, and artificial life. The integrative workshops on adaptive, complex systems that lie at the core of the institute's efforts also stand to gain from fundamental developments in the mathematics of adaptive computation. The value of adaptive computation in creating simulations of complex systems is obvious in the work of John Holland, a pioneer in the field. With Brian Arthur, Holland is using genetic algorithms and classifier systems to build a model of the stock market in which the individual agents are initially ignorant, are not perfectly rational, do not have perfect foresight, and must "learn" over time. By observing the dynamics at work in the simulations, the researchers hope to better understand some of the mechanisms that operate in the stock market, including such phenomena as speculative bubbles and crashes that do not come from standard economic theory.

Simulating a Closed Ecology

Holland's current research interests also include a model of a closed ecology, "a little world with a trivial physics and chemistry and organisms that adapt over time and, in the end, exhibit co-evolving species of increasing complexity. Already this miniature universe exhibits counterparts of the "biological arms race" in which a bush in a tropical forest may rid itself of predators by producing a poison, perhaps containing "quinine"; a predator develops an enzyme that digests quinine; the bush invents a more sophisticated poison, "quinine-b"; then a predator invents a more efficient enzyme; and so the race goes.

One of the advantages of such simulations is that the researcher can back up to an earlier point, slightly alter the conditions, and resume operations. It is then possible to observe why the evolution takes place and under what circumstances.

Large-Scale Systems

Although large systems in constant change may be intricate and difficult to analyze, let alone anticipate, they are not necessarily random; certain properties can be observed and described over time. A system such as the weather, for example, may never be perfectly described nor predicted. But, as Holland points out, the concept of a front in weather prediction enables us to understand, even predict, aspects of that system. While it may never be possible to say with certainty that it will rain in London a year from today, even limited predictive capability can be of use in decision making. Planting can be delayed; drought resistant crops used; and, on another level, people planning to walk to work can leave home with an umbrella.

In contributing innovative methods and concepts to the study of constantly evolving, complex systems in the world around us, adaptive computation may be a tool that leads to new levels of predictive capacities. But the tool itself is still evolving.

Self Reprogramming

In simulations of large-scale systems that display varying dynamics, persistent revision, and the co-existence of numerous conflicting possibilities, Holland stresses that the computer must be able to reprogram itself: "Because of this complexity, machine learning is vital. No designer can hope to anticipate the variety of conditions that will be encountered. A system designed without the possibilities of automatic revision is bound to be brittle, producing implausible actions in unanticipated situations and requiring extensive expert reprogramming on each occasion." In the field of artificial intelligence, machine learning goes a step beyond the expert systems approach. That approach imitates the diagnostic ability of an expert in one, usually rather narrow, domain. It works well within that domain but, Holland cautions, in cases that lie outside that narrow domain "the expert system can make very foolish responses while giving no indication that the diagnosis is not as secure as the standard ones." If, for example, an expert system designed to diagnose blood diseases is given the blood of a pregnant woman, it may attribute a disease to her. Such a system is "brittle."

To handle this brittleness, there must be a way for the system to reprogram itself by using its experience to correct its rules. The machine must not only evaluate the rules—or "hypotheses" as Holland prefers to call them—and eliminate the weaker ones, but it must generate new hypotheses that are plausible.

Highly Parallel Systems

A "classifier system" is needed. These systems are rule based and operate in an IF...THEN form; they are also highly parallel, meaning many rules are simultaneously active. Instead of trying to describe a complex situation by one rule, a cluster of rules can be used, and any given rule may also serve as part of a description of a wide range of situations. To illustrate, Holland uses an everyday example: "Say that you have a rule for handling flat tires. The same rule applies to your rented car, to someone else's Saab, or to a truck by the side of the road. The same rule has many uses, but in different contexts. This is an important feature of classifier systems."

> *"Because of this complexity, machine learning is vital. No designer can hope to anticipate the variety of conditions that will be encountered."*
> — John Holland

Until the process of evaluating hypotheses begins, the classifier system still looks very much like a standard expert system. But with the initiation of "credit assignment" at the mid-level of the model-building process, begins the first level of adaptation, as Holland explains:

"One way to decide which rules are strong and which are weak is to put them into competition. The stronger hypotheses win the competition. Algorithms for assigning credit must somehow give rankings to rules, so that rules that have worked in the past are strengthened and favored for the future.

"Some rules are easier to evaluate than others. In playing a game of checkers, for example, it is easy to credit a rule that says taking a triple jump is good. What is not so easy, and what distinguishes good checker players from poor checker players, is to realize that a certain activity four moves earlier can set up a triple jump. This is a 'credit assignment' problem. Stage setting is very important in games, in economics, in many fields, but assigning the credit back to stage-setting rules is difficult, particularly for the computer by itself."

One procedure for assigning credit, termed a "bucket-brigade algorithm," assists in solving such problems. The machine constantly sifts through waves of information, revising, sorting, and evaluating. In the process, the many rules coupled to "winning" rules also benefit, and as activity increases, benefits will eventually reach back to the stage-setting rules. As certain hypotheses demonstrate their strength, the bucket brigade seeks out the antecedents and assigns credit to them as well.

(continued)

Rule Discovery

The second level of adaptation, and the highest level of the model generation process, is "rule discovery." It occurs when the weak rules are eliminated from the system and must be replaced by new hypotheses.

"The critical aspect of rule discovery is that the new hypotheses must be plausible in terms of past experience," Holland says. "Plausibility, an interesting but rather vague concept, means that, in terms of past experience, there is a reasonable chance that the hypothesis will prove true—although there is certainly no guarantee of this." Like an innovation occuring at a timely moment, it has possibilities for finding a niche.

"We are still in the early stages of feasibility in this area," Holland cautions of the self-reprogramming systems, though there are working examples that indicate encouraging results. A model of a gas pipeline built by civil engineer David Goldberg (now of the University of Alabama at Tuscaloosa) started with randomly generated rules and proceeded with part of the computer simulating a real gas pipeline and part acting as controller. When finished, the program was able to control the system as well as an experienced operator.

Genetic Algorithms

In the difficult but crucial rule generation process, genetic algorithms are used. By treating strong rules as "parents" whose hybrid "offspring" are produced by recombining parts of the parents, the algorithms serve to replace discarded weak rules with stronger new ones. Since it builds on experience, the exploration process becomes less random as "weak" parents and their potential offspring are eliminated and successful rules gain strength.

Los Alamos Postdoctoral Fellow and SFI Member Stephanie Forrest works with genetic algorithms. This work is part of a project on theoretical immunology with SFI External Professor Alan Perelson, also of Los Alamos National Laboratory. Forrest defines a system using genetic algorithms as an "abstract computerized model of genetic evolution." It can be important, she says "in modeling biological systems, as a method for optimizing nonlinear functions using randomized search, and as a paradigm for machine learning in which the genetic algorithm is used to construct programs that improve their own performance."

Genetic algorithms are also part of her work in the study of classifier systems with John Miller, and in a Los Alamos research project with Los Alamos Consultant and SFI Member Gottfried Mayer-Kress on the evolution of co-operative behavior. "The problems are so hard," she says, "that people are working on little bitty pieces of them." But she sees cause for optimism in the possibilities.

The Early Stages

The potential for adaptive computation in the study of complex systems is invaluable, though the applications and theory are still in the early stages. Research efforts are assisted by the availability of massively parallel computers which make it feasible to track the hundreds of thousands of rules that may be simultaneously active in a system.

As a complement to the experimental components, Holland is looking ahead to the need for further development of theoretical aspects, a process that may eventually require a new mathematics. It will be "something very symbolic, that to the layman will be an arcane language," he says. "It will differ from the standard mathematics used in physics and many other sciences, which concentrates on end points, equilibria, or stable solutions. The systems that we're studying—biological systems, economic systems under evolution, the immune system—do not stabilize. (In living systems, for example, stability is synonymous with death). These systems are always unfolding, and the constantly changing part is what we want to study."

The theory will be necessary, he notes, because in the complex systems to be studied, there are so many areas to be developed that researchers could spend whole lifetimes looking in fruitless places. As the theory unfolds, it should direct attention to areas that are likely to be productive. Benefits are cyclical: as the experiment suggests changes to the theory, the theory in turn suggests where to try the experiments.

"You go around this loop in an indefinite fashion," Holland says, "and if you're lucky, you make advances."

Holland stresses that the programs are exploratory by nature. "They need people and funding. In 10 or 20 years we will have a deeper understanding of these systems. We cannot expect great breakthroughs—areas this complex don't generally yield to this—but there are likely to be substantial advances in understanding.

"There will be some insights and certainly more understanding, perhaps enough to guide policy in real situations. It is unlikely that we will be in complete control of such complex systems, but neither will we be in complete ignorance. Our hope is for improved understanding within well-recognized limits: a new level of education. Much as a flight simulator is used to train pilots—it lets them practice in risky situations—it may be possible for us to explore risky situations in the larger world."

As in the example of the weather, perfect predictability or control may be out of the question. If, however, study of large, complex systems can produce general insights and limited predictive abilities, then, like the farmer who decides to plant a different crop or the pedestrian who decides to leave home with an umbrella, we could use the information to plan our own alternatives in order to adapt.

"Adaptation is the key word," concludes Holland. "In fact, most of my work is an attempt to understand that term."

—Jeanie Puleston Fleming

Jeanie Fleming is a local freelance writer.

Doyne Farmer

A Profile of Doyne Farmer

Exploiting Chaos to Predict the Future

A 1973 graduate of Stanford University, Doyne Farmer has been interested in complexity since his graduate studies in physics at the University of California, Santa Cruz. After obtaining his doctorate from UCSC in 1981, he was an Oppenheimer Fellow at Los Alamos National Laboratory, where he is presently the Group Leader of the Complex Systems Group, Theoretical Division, and an affiliate of the Center for Nonlinear Studies. The author of numerous articles on chaos, Farmer is also the co-editor (with A. Lapedes, N. Packard, and B. Wendroff) of Evolution, Games, and Learning: Models for Adaptation in Machines and Nature. *He has been active in the Santa Fe Institute since it was founded and was a faculty member of the 1988 and 1989 Complex Systems Summer Schools. He is presently a member of the SFI Science Board.*

Working at the Institute and at Los Alamos National Laboratory Farmer, along with colleagues Martin Casdagli, Stephen Eubank, John Gibson, Stephen Pope, and David Wolpert, has been developing a new methodology for nonlinear modeling and prediction.

Could you summarize the purpose and goals of your work?

The main goal of our work is to develop efficient numerical techniques that construct nonlinear models directly from large databases. We are doing automatic model building, where we look at the data and build a model directly off it, assuming we know nothing about the system. In practice this involves programming a computer to search data sets for patterns that have predictive value. In this modern age of information explosion, there is no shortage of large databases to analyze. For example, weather satellites, scientific experiments, and the stock market generate large amounts of data every day.

Our approach is very different from the traditional scientific methodology of building a model from first principles, which goes back to Newton's laws of gravitation, modeling planetary motion. Interestingly, even here, the elliptical *patterns* of planetary motion were first discovered by Kepler from an exhaustive search through data on planetary motions. Without Kepler's empirical discovery of patterns, Newton might never have been able to formulate fundamental laws. Our approach is similar to Kepler's. When the data are much more complex and irregular, it can be much more difficult to build models from first principles, or develop intuition for what is going on. We use computers to try and develop such "intuition."

We pride ourselves on our intuition. People are always saying, "Human beings can do these wonderful things." That is probably true, but nonetheless, it is my conviction that human beings have a very poor intuition for understanding nonlinear behavior. It is easy to write down lists of numbers that, to the eye, look random, but which nonetheless are described by very simple deterministic relationships. That's what chaos is all about. Chaos shows that nonlinearity can generate apparent randomness, phenomena that look incomprehensible to us at first glance, but about which we can gain much intuition only if we look at the data in the right way.

One of the main goals of our research is to discover whether certain data sets are deterministic or random. If our modelling techniques give rise to good predictions, this is evidence that the data, although apparently random, are in fact deterministic. Thus our approach is pragmatic at heart.

In the case of weather satellites data, we would like to translate the information gathered into the things that we would like to know, such as surface temperatures of places on Earth where we do not have weather stations, pressures, wind velocities, and vegetation types on Earth.

In the case of stock market data, we would like to sift through the information and discover what influences what, and ultimately find out if our techniques can give better forecasts than other competing techniques.

So what the model does is codify the structure derived from the data.

Right. Our approach, when confronted with some data, is to try to understand what's in the data. The first question one has to address is how to represent the data, to come up with the proper language to describe it. The next thing to decide is: What do you want your model to do? So you write down some objective functions that will tell whether you're doing a good job or not. Then you have to extract the features in that representation that you really want to pay attention to, and you build a model that tells you: given the cause, this is going to be the effect, a model that actually makes the predictions. Now, that last part is a map. A map is just a graph. Just think of drawing an x axis and sticking a point on it. What's the corresponding point on the y axis? A map is just a little rule that tells you how to go from x to y. x in this case is the cause, and y is the effect. We want to build maps from cause to effect.

Maps will tell us how to predict what is going to happen. The weather is an example. x would be the weather today, and y would be the weather tomorrow. Now nonlinear just means that the graph is not straight. It would be naive to expect that all the graphs that you're ever going to draw of cause-and-effect relationships in the world are going to be straight lines. That actually has been the approach most people have taken up until fairly recently—the reason simply being, that was what they could do. That was what was tractable and could be treated with a pencil and paper. But increasingly, computers allow us to break away from that. We've come up with some new tricks, local approximation of these functions. We've come up with relationships that link ideas on function approximations to ideas in chaos. So now we can make predictions, once we know something about the system we're looking at, which can be encapsulated in a couple of numbers. If we know we have a certain number of data points, we can tell you how good we expect our model to be.

Since I came out of the field of dynamical systems and chaos, I have seen a lot of examples of things that looked random but were deterministic. So, my attitude is to make my nonlinear model more accurate so that it can map the nonlinearities in the system, but still to assume that as much of the world as possible is really deterministic. It is that difference in attitude that has made a lot of difference in the way we approach doing real world problems. What occurs to us as being an obvious way to approach the problem is very different from what occurs to a statistician as the obvious way to approach the problem.

What about software?

I've spent a lot of my life building software models and writing software. Several years ago I realized this was going to develop into a big project, so that it was worth doing the software right. We made a considerable detour and learned a new object-oriented programming language, because we wanted to build a package that could be easily worked on by a lot of different people, one that was an integrated tool that could serve as a connection point for a lot of work. We have a group of people who have been working on this. Stephen Eubank, John Gibson, Stephen Pope, and I have a working software package that allows us to build a variety of different types of nonlinear models and use them to make predictions.

...things that look random may actually have deterministic structure, and if you can capture that deterministic structure, you should be able to make short-term predictions. Chaos implies short-term predictablity, but long-term unpredictability. We are trying to exploit the short-term predictability.

What in your research up to this point convinced you that a new methodology was needed?

I came into this area through the back door. My background was in dynamical systems and chaos. When I was a graduate student, I was already working on the notion that chaos ought to allow you to make better predictions. Chaos is a double-edged sword, because on one side it says that a deterministic system with a tiny bit of uncertainty in it can amplify that uncertainty and generate something that looks random. On the other side, it says that things that look random may actually have deterministic structure, and if you can capture that deterministic structure, you should be able to make short-term predictions. So, chaos implies short-term predictability, but long-term unpredictability. We are trying to exploit the short-term predictability.

Now, as I said, I came at this from a background in physics and dynamical systems. But, as I plunged into it, I became increasingly aware that there was a big field out there called time-series analysis, where people have been examining this problem since the 1920's. Their perspective on the problem came out of statistics and the theory of probability, which is a different view of the way the world works. The deterministic view of the world in terms of dynamical systems is to model the state of a system as a point which follows a deterministic trajectory. The opposite view of the world is that the state of a system is described by a probability distribution which describes degrees of uncertainty. Traditional time-series analysts attempt to find rules for how these degrees of uncertainty transform into one another.

In other words, you're saying there's not really any kind of a viable overview?

Yes. This is a ubiquitous problem with science, and particularly in interdisciplinary activities like this. Independent communities arise, and they come up with their own jargon. There are the statistics people, engineering and control theory people, pure mathematicians, and economists, and then there is the machine-learning community using neural nets, classifier systems, and a host of other approaches. A lot of the models these people come up with have striking similarities. They're just describing the same things in different languages.

How does your work compare with that of John Holland in classifier systems?

The classifier system is an attempt to have the system automatically learn. You just throw information at it and it builds its own internal model and makes its own predictions. It adjusts as it goes along, and it is automatic. The classifier system works well in certain domains. However, the domains where the kind of things we're doing work well, and those where the classifier system works well are different. We're working in a domain where we have more precise information and feedback. We also work in a domain where we can exploit smoothness and continuity, and where we assume that the world is described by continuous real numbers.

It seems that all of this research is predicated on the idea that the world is far less random than we had ever thought.

That's right. All these things are attempts to exploit the underlying simplicity of the world in an automatic fashion. Neural nets are another approach to this learning problem. It's clear that neural nets are actually doing a nonlinear function approximation problem. And that's what we're doing. Our approach is simply more overt. There have been a lot of interactions and cross-fertilization in the last few years between the neural net people and us. I'm pleased to see that some of the synthesis is already happening. We've had a statistician, Irene Poli from the University of Bologna, who has been visiting us for the last four months or so, and she has come up with a nice synthesis of these three points of view. There are also a lot of analogies between classifier systems and neural nets. I think the analogies are in fact closer than anybody realized at first, because the terminology varied, and also, I think, because John was originally motivated much more by traditional expert systems than he was by neural nets.

What are some of the practical applications of your techniques?

Fluid flows are something that we've had our most dramatic success with. We've taken time series that came out of fluid flows—in an experimental context, not geophysical phenomena. In comparing our models to regular linear models, in some cases we've done forty times as well as the linear models in predicting points we haven't seen yet. In this area, improvements by a factor of forty are really big improvements.

What about chemical and industrial processes?

These are promising areas, because chemical reactions are known to have a lot of nonlinear simplicity of the type that we should be able to do well on. A lot of industrial processes, like steel manufacturing, are basically chemical reactions going on in a big vat. You pour a lot of stuff into pots, stir the pot, and something runs out the other side. Problems exist in trying to predict what's going to come out, given what you put in. One of the limiting factors in steel production is that they can't predict perfectly what kind of steel they are going to make on a given day.

Weather forecasting is another area of promise. There are occasionally times where it snows or rains throughout, for example, New Mexico, but frequently the weather in one part of the state is completely different from weather in another part. We may be able to take all the information that is out there, break it down, and make some more useful and codified predictions based on the information that is actually available. One of the big differences between our techniques and something like a global circulation model is that the latter has to have information to put on a grid; it uses physics to make predictions. Our system takes whatever information it's got, and builds a model that's directly linked to that set of information.

All these things are attempts to exploit the underlying simplicity of the world in an automatic fashion.

Cardiac arrhythmia?

There is a lot of evidence that various cardiac arrhythmias are described by low-dimensional chaos. We would like to reanalyze the data and try to find out whether or not it is really accurate.

There is also a lot of work suggesting that low-dimensional chaos occurs in neuro-physiological systems. We are really lucky that a new postdoctoral fellow, Andre Longtin, who has just joined us, has a strong background in neurophysiology. We are hoping to utilize what he has done and what he knows to analyze various data sets and figure out whether chaos is really playing a role in these things, and whether we can use it as a classifier tool.

(continued)

Economic models?

In many respects, I view economics as being the hardest of these applications, because it is a very complicated system. On the other hand, it is also one of the most fruitful if you can actually predict anything. We are lucky that Dow Jones has given us access to their database and archives. One of the things we are working on now is trying to see whether there are indicators in the media of the collective mood of humanity. In other words, are there mood indicators that move up and down as a function of time and can we detect and measure them from information in the media? One of the wonderful things about Dow Jones is that they have an enormous number of journals in electronic form, from the Wall Street Journal to some fairly obscure magazines. We are looking at things like word counts—for example, the number of positive as opposed to negative adjectives—to see whether there are statistically significant swings that indicate society's moods.

> *We are working on . . .trying to see whether there are indicators in the media of the collective mood of humanity. In other words, are there mood indicators that move up and down as a function of time and can we detect and measure them from information in the media?*

How does your current work tie-in with your long-standing interest in predicting the future?

Well, one of the reasons I like working in this area is because it is a synthesis of the ideas that originally led me into chaos to begin with. Norman Packard and I were led indirectly into chaos through the business of building and programming computers to beat the game of roulette. Roulette made me realize that prediction is a difficult and profound problem. It made me think about what it means for something to be random—so much so that when chaos came along, I was primed to embrace and pursue it. We initially used physical models based on first principles. These models worked very well as long as the parameters were adjusted properly. However each roulette wheel is different and adjusting the parameters could be quite difficult. By moving to empirical models that were not based on first principles we were able to make the parameter adjusting automatic. This parallels the empirical model building we are doing now.

—Anne Pedersen

Anne Pedersen is a Santa Fe-based writer and editor.

1990 Complex Systems Summer School

June 4–28
Santa Fe, New Mexico

The School is intended to provide graduate students and postdoctoral scientists with an introduction to the study of "complex" behavior in mathematical, physical, and living systems. The program includes four weeks of course lectures together with seminars, computer workshops, and experimental labs. Individual/group research projects will be encouraged. Students are expected to have graduate level training in one of the mathematical, physical, biological or information sciences. Students will be supported with housing, meals and travel funds as necessary, subject to funding availability. No tuition fees. Co-Directors are Lynn Nadel and Daniel Stein, both of the University of Arizona.

Course Lecturers

ANDREW BARTO, *Computer Science, University of Massachusetts:* learning algorithms
BRUCE BAYLY, *Mathematics, University of Arizona:* complexity in fluid flow
CHARLES DOERING, *Physics, Clarkson University:* stochastic processes in the physical and biological sciences
LEIF FINKEL, *Biomedical Engineering, University of Pennsylvania:* visual system function and development
ZIAUL HASAN, *Physiology, University of Arizona:* movement complexity
RALPH LINSKER, *IBM:* neural computation and pattern formation
RICHARD MICHOD, *Ecology and Evolutionary Biology, University of Arizona:* evolutionary biology
SIDNEY NAGEL, *Physics, University of Chicago:* experimental anaylsis of disordered systems
THOMAS ROSENBAUM, *Physics, University of Chicago:* experimental analysis of disordered systems
DAVID RUMELHART, *Psychology, Stanford University:* connectionist models of cognition
WING TAM, *Physics, University of Arizona:* pattern formation in chemical systems
JOSEPH TRAUB, *Computer Science, Columbia University:* computational complexity

Contact

Andi Sutherland, Santa Fe Institute, 1120 Canyon Road, Santa Fe, New Mexico 87501; email: summerschool@sfi.santafe.edu (505) 984-8800; fax: (505) 982-0565

Sponsors

Center for Nonlinear Studies, Los Alamos National Laboratory
Sandia National Laboratories
Santa Fe Institute
University of Arizona
University of California system
University of Illinois
University of Texas

Application Deadline: March 1, 1990

Publications

Santa Fe Institute Studies in the Sciences of Complexity

Books

Lectures in the Sciences of Complexity, Lectures Volume I, edited by D. Stein. Redwood City, CA: Addison-Wesley, 1989. Lectures from the 1988 Complex Systems Summer School.

Lattice Gas Methods for Partial Differential Equations, Proceedings Volume IV, edited by G. Doolen et al. Redwood City, CA: Addison-Wesley, 1989. Proceedings from the 1987 workshop of the same title.

Computers and DNA, Proceedings Volume VII, edited by G. Bell and T. Marr. Redwood City, CA: Addison-Wesley, 1989. Proceedings from the 1988 workshop "The Interface between Computational Science and Nucleic Acid Sequencing."

Second Printings (Reprints)

Artificial Life, Proceedings Volume VI, edited by C. Langton. Reading, MA: Addison-Wesley, 1988.

Lectures in the Sciences of Complexity, Lectures Volume I, edited by D. Stein. Redwood City, CA: Addison-Wesley, 1989.

Proposed Volumes

A number of volumes are in production or proposed for publication in 1990. These include the following tentatively titled volumes:

Complexity, Entropy, and the Physics of Information, Proceedings Volume VIII, edited by W. H. Zurek; proceedings from the 1989 workshop of the same title.

Molecular Evolution on Rugged Landscapes: Proteins, RNA, and the Immune System, Proceedings Volume IX, edited by A. S. Perelson and S. A. Kauffman; proceedings of the 1989 workshop "Applied Molecular Evolution and the Maturation of the Immune Response."

The Evolution of Human Languages, Proceedings Volume X, edited by J. Hawkins and M. Gell-Mann; the proceedings of the 1989 workshop of the same title.

Complex Adaptive Systems, Proceedings Volume XI, edited by M. Gell-Mann and D. Pines; the proceedings of the integrative workshop on Complex Adaptive Systems.

Artificial Life II, Proceedings Volume XII, edited by J. D. Farmer, S. Rasmussen, and C. Langton; the proceedings of the second co-sponsored workshop on artificial life.

Introduction to the Theory of Neural Computation, Lecture Notes Volume I, by John Hertz, Richard Palmer, and Anders Krogh; based on a course on Advanced Topics in Statistical Mechanics given at Duke University.

1989 Lectures in Complex Systems, Lectures Volume II, edited by Erica Jen; lectures from the 1989 Complex Systems Summer School

1990 Lectures in Complex Systems, Lectures Volume III, edited by D. Stein and L. Nadel; lectures from the 1990 Complex Systems Summer School.

Dynamics of Complex Systems, a monograph by Gérard Weisbuch translated from the French (*Dynamique des systèmes complexes*, InterEditions, 1989).

Other Books

Books published outside of the SFI series, but based on work done in part at SFI include:

Induction by John Holland (a new paperback version is being prepared by MIT Press)

The Origins of Order: Self-Organization and Selection in Evolution by Stuart Kauffman (Oxford University Press, 1990).

Research Papers Based on Work Done in Part at SFI

Anderlini, Luca. "Communication, Computability and Common Interest Games." Santa Fe Institute Preprint 89-002, 1989.

Arthur, W. Brian. "'Silicon Valley' Locational Clusters: When Do Increasing Returns Imply Monopoly?" Santa Fe Institute Preprint 89-007, 1989.

Bagley, R. J., J. D. Farmer, N. H. Packard, A. S. Perelson, and I. M. Stadnyk. "Modeling Adaptive Biological Systems Biosystems." To appear in *Biocybernetics*, 1989.

Barr, Edith S., Matthew Koebbe, Gottfried Mayer-Kress, and Toshiro K. Ohsumi. "Similarities between the Dynamics Underlying Rossler's Band and a Suspension of a Quadratic Entire Cremona Transformation." Submitted to *A Chaotic Hierarchy*, edited by G. Baier and M. Klein. Singapore: World Scientific, 1990.

Belen, A. d'A., and J. D. Farmer. "Artificial Life: The Coming Evolution." To appear in the proceedings of the symposium commemorating Murray Gell-Mann's 60th Birthday (1989).

Cline, J., S. S. Pinsky, and L. M. Simmons, Jr. "Renormalization of the Delta Expansion for $\lambda(\phi^{2+2\delta})_2$." *Phys. Rev.* D39 (1989): 3043

Dressler, U., J. D. Farmer, and T. Taylor. "Higher Order Lyapunov Exponents." Unpublished, 1989.

Ecke, R. E., J. D. Farmer, and D. K. Umberger. "Scaling of the Arnold Tongues." *Nonlinearity* 2 (1989): 175–196.

Farmer, J. D., and J. J. Sidorowich. "Exploiting Chaos to Predict the Future and Reduce Noise." In *Evolution, Learning and*

(continued)

Cognition, edited by Y.C. Lee. New York: World Scientific, 1988.

Farmer, J. D., and J. J. Sidorowich. "Predicting Chaotic Dynamics." In *Dynamic Patterns in Complex Systems*, edited by J. A. S. Kelso, A. J. Mandell, and M. F. Shlesinger. Singapore: World Scientific, 1988.

Farmer, J. D., and J. J. Sidorowich. "Optimal Shadowing and Noise Reduction." Unpublished, 1989.

Forrest, S., and J. H. Miller. "Emergent Behaviors of Classifier Systems." To be published in the Proceedings of the Ninth Annual CNLS Conference on Emergent Computation, 1989.

Forrest, S., and J. H. Miller. "The Dynamics of Classifier Systems: Empirical Results." Los Alamos National Laboratory Technical Report LA-UR-89-760, 1989.

Goodwin, B. C., and S. A. Kauffman. "Bifurcation, Harmonics and the Four-Color Wheel Model of Drosophila Development." *Theoretical Models for Cell-to-Cell Signaling*, edited by A. Goldbeter. London: Academic Press, 1989.

Kauffman, S. A. "Adaptation on Rugged Fitness Landscapes." In *Lectures in the Sciences of Complexity*, edited by D. Stein. SFI Studies in the Sciences of Complexity, Lect. Vol. 1. Redwood City, CA: Addison-Wesley, 1989.

Kauffman, S. A. "Principles of Adaptation in Complex Systems." In *Lectures in the Sciences of Complexity*, edited by D. Stein. SFI Studies in the Sciences of Complexity, Lect. Vol. 1. Redwood City, CA: Addison-Wesley, 1989.

Kauffman, S. A. "Requirements for Evolvability in Complex Systems: Orderly Dynamics and Frozen Components." In *Complexity, Entropy and the Physics of Information*, edited by W. H. Zurek. SFI Studies in the Sciences of Complexity, Proc. Vol. VIII. Redwood City, CA: Addison Wesley, 1990, in press.

Kauffman, S. A. "Cambrian Explosion and Permian Quiescence: Implications of Rugged Fitness Landscapes." In *Evol. Ecol.* (1989), in press.

Kauffman, S. A., and E. Weinberger. "The N-K Model of Rugged Fitness Landscapes and Its Application to Maturation of the Immune Response." *J. Theor. Biol.* (1989), in press.

Li, Wentian. "Mutual Information Functions of Natural Language Texts." Santa Fe Institute Preprint 89-008, 1989.

Macken, C. A., and A. S. Perelson. "Protein Evolution on Rugged Landscapes." *Proc. Natl. Acad. Sci. USA* 86 (1989): 6191–6195.

Marimon, R., E. McGrattan, and T. J. Sargent. "Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents." Santa Fe Institute Preprint 89-004, 1989.

Mayer-Kress, G. "A Nonlinear Dynamical Systems Approach to International Security." To be included in the proceedings of the annual AAAS conference, San Francisco, January 1989.

Mayer-Kress, G., and A. Huebler. "Time Evolution of Local Complexity Measures and Aperiodic Perturbations of Nonlinear Dynamical Systems." In *Quantitative Measures of Complex Dynamical Systems*, edited by N. B. Abraham. New York: Plenum, 1989.

Mayer-Kress, G., and D. Campbell. "Challenges and Limitations of Mathematical Modeling of Environmental Issues." Proceedings of the International Conference on Technology-Based Confidence Building, CNSS, Los Alamos National Laboratory, Santa Fe, hosted by L. Rosen, July 1989.

Mayer-Kress, G. "Monitoring Chaos of Cardiac Rhythms." Proceedings of the 6th Annual Industry Conference and Exhibition, *Biotech USA*, October, 1989.

Miller, J. "The Coevolution of Automata in the Repeated Prisoner's Dilemma." Santa Fe Institute Preprint 89-003, 1989.

Miller, J., and S. Forrest. "The Dynamical Behavior of Classifier Systems." Santa Fe Institute Preprint 89-005, 1989.

Perelson, A. S. "Modelling the Interaction of HIV with the Immune System." In *Mathematical and Statistical Approaches to AIDS Epidemiology*, edited by C. Castillo-Chavez. Lecture Notes in Biomathematics. New York: Springer-Verlag, 1989.

Perelson, A. S. "Immune Network Theory." *Immunol. Rev.* (1989): 5–36.

Perelson, A. S. "Immune Networks: A Topological View." In *Cell-to-Cell Signalling: From Experiments to Theoretical Models*, edited by A. Goldbeter. New York: Academic Press, 1989, 261–272.

Rust, J., R. Palmer, and J. Miller. "A Double Auction Market for Computerized Traders." Santa Fe Institute Preprint 89-001, 1989.

Rust, J., R. Palmer, and J. Miller. "Participant's Manual" (for the Double Auction Tournament), Santa Fe Institute, 1989.

Scheinkman, J. A. "Nonlinearities in Economic Dynamics." Santa Fe Institute Preprint 89-006, 1989.

Segel, L. A., and A. S. Perelson. "Shape Space: An Approach to the Evaluation of Cross-Reactivity Effects, Stability and Controllability in the Immune System." *Immunol. Let.* 22 (1989): 91–99.

Segel, L. A., and A. S. Perelson. "Some Reflections on Memory in Shape Space." In *Theories of Immune Networks*, edited by H. Atlan and I. R. Cohen. Berlin: Springer-Verlag, 1989.

Segel, L. A., and A. S. Perelson. "Shape Space Analysis of Immune Networks." In *Cell-to-Cell Signalling: From Experiments to Theoretical Models*, edited by A. Goldbeter. New York: Academic Press, 1989, 273–283.

Zurek, W. H. "Dynamics of Computation: Algorithmic Complexity and Information Metrics." *Nature* 341 (Sept., 1989).

Zurek, W. H. "Algorithmic Randomness and Physical Entropy." *Phys. Rev. A* (Oct., 1989).

—RKBV,GF

1989 Contributors to the Santa Fe Institute

The Santa Fe Institute and the programs that it sponsors or co-sponsors have been supported in 1989 by the following:

Academic Institutions

Brandeis University
St. John's College
The University of Arizona
The University of California
The University of Illinois
The University of Maryland
The University of New Mexico
The University of Texas at Austin

Foundations

The Bay Foundation
The Arie & Ida Crown Memorial
The Harry Frank Guggenheim Foundation
The John Fitzgerald Kennedy Library, Boston
The Esther A. & Joseph Klingenstein Fund
Richard Lounsbery Foundation
The John D. & Catherine T. MacArthur Foundation
Research Corporation
The Russell Sage Foundation

Laboratories

Los Alamos National Laboratory

Government Agencies

Air Force Office of Scientific Research
Department of Energy
National Science Foundation
NATO
Office of Naval Research

Corporations

The Advest Group, Inc.
Aetna Life Insurance Company
Batterymarch Financial Management
Booz-Allen & Hamilton, Inc.
Citicorp/Citibank
Cray Research, Inc.
Digital Equipment Corporation
Fiduciary Trust International
Heinz Co. Foundation
IBM
David E. Shaw & Co.
Morgan Stanley & Co.

Individuals

Lewis & Margaret Agnew, Los Alamos
Jean Anderson, Tesuque
P. W. Anderson, Hopewell, NJ
J. B. Ault, Santa Fe
Richard & Helen Brandt, Santa Fe
Jack D. Campbell, Santa Fe
Bengt & Julia Carlson, Tesuque
Albert & Molly Clogston
Stirling & Rosemary Colgate, Los Alamos
George Cowan, Los Alamos
Paul Craig, Davis, CA
Irwin & Florence Cromwell, Santa Fe
Carl Djerassi, Stanford, CA
Jerry Donato, Santa Fe
William Dougherty, Santa Fe
Sam & Valerie Dunnam, Austin, TX
Marcus Feldman, Stanford, CA
Donald & Margaret Hagerman, Los Alamos
Lawrence Huntington, New York City
Stuart & Elizabeth Kauffman, Santa Fe
L. D. Percival & Elizabeth King, Santa Fe
George Kozmetsky, Austin, TX
Donald Kummer, St. Louis
Tom Lang, Albuquerque
Robert & Nancy Maynard, Aspen, CO
Faith Meem, Santa Fe
Seth & Margaret Montgomery, Santa Fe
Robert & Ann Noyce, Austin, TX
Thomas Pick, Chicago
Dr. & Mrs. David Pines, Tesuque
Nan Robinson, New York City
Ross Robinson, Santa Fe
Assis Kumar Sakar, Philadelphia
Arthur & Libby Spiegel, Albuquerque
J. I. Staley, Wichita Falls, TX
Eugene & Gail Stark, Los Alamos
Thomas & Janet Thompson, Las Cruces
Anthony & Ireene Turkevich, Chicago
Charles Tyng, Santa Fe
Philip Voegelin, Santa Fe
Hans von Briesen, Jr., Santa Fe
Henry & Mary Wider, Albuquerque
Susan Wider, Albuquerque

—SW,MA

Staff News

After directing the Economics Research Program for the last year and a half, W. Brian Arthur has returned to Stanford University. Arthur has been with SFI continuously since June, 1988, expertly guiding the transformation of a loose network of researchers into a full-scale, on-site research program. Arthur remains closely associated with the program as part of its Steering Committee and as a member of the 1990 Summer Study group. His vision and hard work during the past year and a half is greatly appreciated, and we look forward to his return to Santa Fe.

Robin Justice has joined the Institute staff as our Computer Systems Manager, replacing Stephen Pope. Justice has extensive experience as a computer systems analyst, most recently with Los Alamos Technical Associates. At SFI he'll be responsible for systems coordination as well as working with staff and researchers on specific projects including programing, database management, and software generation.

Stephen Pope has resigned as the SFI Computer Systems Manager to accept a position as a staff member in the Advanced Computing Laboratory at Los Alamos National Laboratory. We will miss Stephen's expertise and energy in the Research Wing, but are happy that he will continue to visit as a Member of the SFI and as a contributor to the Time-Series Forecasting project.