THE EVOLUTION, GAMES AND LEARNING CONFERENCE: REPORT

LOS ALAMOS CENTER FOR NONLINEAR STUDIES
William Bricken
May 1985

The Center for Nonlinear Studies is "... a resource serving the needs of the entire nonlinear community." It sponsors a yearly multidisciplinary conference (reported below), and research into:

dynamical systems

"... only a few degrees of freedom are needed to describe EXACTLY the observed chaotic behavior of a system having A PRIORI an infinite number of degrees of freedom."

reduced dimensional materials

"... a spatially extended excitation modifies the medium surrounding it in such a way as to "dig a hole" for itself and become strongly localized."

scientific computing

"... it is typically necessary to use computers in an exploratory mode, simply to gain some intuition about systems that cannot be treated analytically."

The conference included two dozen technical presentations (each an hour in length), strong computer demonstrations of games and cellular automata, and about 200 participants.

My comments on the presentations follow. Murray Gell-Mann presided over the opening session.

DAY 1

Gian-Carlo Rota (MIT philosopher mathematician):

A story in memoriam of Stan Ulam. Basically, we never perceive objects, only functions, roles in context. The punch line was to formalize "as", as in "I see the pages of paper AS a book." No hints about how.

Irving Good (Virginia Polytech philosopher biologist):

Progress is not a necessary concept to evolution. Evolution is a little more probable than creation, so let's believe Darwin, even if he did turn into a Lamarkian as he matured. Someone in the audience accused the speaker of a

vicious attack on the Lord's creation, and the question period ended abruptly.

James Valentine (UCSB oceanographic paleontologist):

Evolution is not stable in time. The fossil record (from the ocean) shows rapid bursts of diversity at all hierarchical levels followed by selective smaller extinctions. Diversity radiates outward and is not a sequential process. Evolutionary intermediates (like non implemented ideas) are local, nondiversified, and rapidly changing. About 1/3 of all niches are OPEN when times are stable. Extinction keeps niches open and change chugging.

Norm Packard (Princeton young mathematician):

Let's view evolution as a slow trajectory of changing state spaces in the space of state spaces. Each space has its own fast dynamic. (Gell-Mann interrupts: you can do that merely by increasing the number of parameters in a single state space.) Thus corrected, Norm showed us some nice pictures of colorful cellular automata.

Bernardo Huberman (Xerox PARC computer scientist):

Excellent research about the laws governing the relations between parts and wholes. To get effective decoupling between hierarchical levels, let the large scale structures act as static constraints while integrating over the lesser structures. Adaptable algorithms and architectures accept variation in structure without variation in organization. They have an architecture of densely interconnected decision nodes that iteratively converges to uniquely recognize input structures.

Stephen Wolfram (Princeton cellular automata guru and young mathematician):

Cellular automata are discrete approximations of partial differential equations (or iterated continuous mappings on the Cantor set). But some classes are so complex that the easiest way to predict the next state is to calculate it. The implication: evolution is computationally irreducible; the best way to predict what comes next is to wait and see. For complex systems, prediction and behavior are equivalent. We can only predict general features, we can only specify general goals; leave the details to emergence. But we can, via a control meta-structure, guide the details. The fixed-points of cellular automata form a natural language for control of computation. We find them by random search since their attractor basins are fractal. Actually, Stephen is out of ideas at this point. He is also guilty of one-dimensional perspective.

End of Day 1 Comments

Score:
Computer Science: 1
Mathematics: 0
Biology, the umpire, is blind

Idea of the Day:
The real competition in evolution is between existence and extinction.

Close Second:
For complex systems, predicting behavior means observing behavior.

Raspberry of the Day:
one-dimensional cellular automata (beautiful but dumb)

Cliffhanger of the Day:
Will experimental mathematics seduce the formalists? Thus far in our story, the lets-take-a-look approach of trendy new mathematicians has flashed its sexiest curves. The old men of the continuous school have been entertained by the naivety of youth. Tomorrow, the games.

DAY 2

Hans Berliner (CMU, AI in backgammon):

Old polynomial twiddling in evaluation functions. Some folks just can't mature past numbers.

Morton Davis (City College NY, ?):

Why did this guy make a fool of himself? NEVER present an idea that is thirty years old without consulting the references. (For the curious, it was shallow evaluation, shallowly thought out, shallowly presented.)

Peter Frey (Northwestern psychologist):

Classical (a la Samuels) feature weighting in the board game Othello, with surprising results. First, the experts are wrong: the program rewrote the strategy of the game and won the World Championships. Second, never play a losing strategy twice (obvious, but seldom implemented). Excellent research in a domain that is better suited for machine play than for humans. Note that knowledge engineering not machine learning decides features. The machine can decide among the selection of features. Non-linear evaluation is mandatory.

John Maynard-Smith (Univ. of Sussex, evolutionary biologist, DUDE):

A fundamental contributor to mathematical population biology. Others at the conference deferred to his wisdom (A typical comment: "Since John is in the audience, I will be very careful about what I say."). In the game theory of interaction, all asymmetries must be perceivable. That is, mental preparation is irrelevant. If you glue a small lead weight on the back of some spiders, they will win the showdown for the territory because they move as if they were bigger. Major concept: Evolutionary Stable Strategies (ESS), an operational definition of fitness. Characteristics of an ESS: needs many trials to stabilize transients, never completely disappears, recent trials contribute more, probability of occurrence determined by external payoffs (sounds like Pavlovian reinforcement, but the animal chooses it autonomously). Animal behavior is far more sophisticated and complex than our simplistic models.

Peter Schuster (Univ. of Vienna, biological mathematician):

A student of the Eigen school (see Hypercycles, Eigen's book). Superb mathematical modeling of autocatalytic processes. The Europeans excel in theoretical work. Selection and mutation are the same concept. Replication processes is the key to understanding evolution. Selection is error-free replication. Mutation is the error term in this process. First order autocatalysis must be error-free. Higher order autocatalysis must be constrained: this is the source of the no-free-lunch concept (i.e. the environment is finite). By the way, multiple reentry of LOSP imaginary variables models poly-autocatalysis as complex indication waves. Schuster studies simple replicating systems (template induced RNA synthesis). Necessary components: a template or initialization, strand matching or unification, and strand separation or drawing conclusions. Autocatalytic reactions are robust, they wash out random effects. That is, evolution is driven by self-reference not by mutation.

Stuart Kauffman (U. of Pennsylvania evolutionary biologist):

Autocatalytic reactions are the inevitable origin of life. Cleavage and joining of simple molecules is all that is necessary to form a supercritical set which begins massive catalysis. The problem is to find a differentiation mechanism, otherwise its universal Gaia (well that phrase is redundant, Gaia is universal; differentiation is unnecessary.). Provocative ideas, slightly askew.

End of Day 2 Comments

Score for the Day:
Mathematical Biology: 3
Computer Science: 1

Idea of the Day:
Learning and evolution are self-referential, not random, processes.

Runner-Up:
If the domain is right, the machine can teach the humans (eg: Othello rewrite).

Let's Talk About This More Award:
Kauffman's inevitable autocatalysis.

Raspberry of the Day:
Search the literature before you open your mouth, Morton.

Wisdom of the Sages:
Human simplification is anti-descriptive.

Wisdom of the Sages, Part II:
If you do not have a self-referential foundation, you cannot know what you think you know.

Growing Suspicion:
The better the school, the better the speaker.

Special Session on Computer Games and Simulations

Bruce Wilcox (GO):

An American codification of the inscrutable tactics of the world's simplest complex game. Excellent for beginners, misses the whole point for experts. Still, the computer is a powerful analyst of local phenomenon. I learned a lot. On the IBM-PC: interface is poor, play is moderately good, theory is excellent. Bruce is negligent in not writing a book.

Hans Berliner (BACKGAMMON):

The World Champion because of lucky rolls. Old but entertaining. No interface, play is excellent, theory is good.

Peter Frey (OTHELLO):

A delightful experience. Fast, smarter than the reference materials, very educational. On some 64K machine: interface is poor, play is superb, theory is ground-breaking.

Robert Hyatt (CHESS):

A well-studied topic, this work is becoming a good product. Still hard to beat, still idiosyncratic, still evolving. Interface is excellent, play is very good, theory is established.

David Slate (CHECKERS):

Another world champion. On the IBM-PC: interface is poor, play is excellent, theory is old but good.

Stephan Wolfram (one-dimensional cellular automata):

A comprehensive research tool; the user can specify thousands of rules and initializations. Very fast, very pretty, but there may not be a lot more. The information is so dense that it must be "analyzed" visually. Social faux-pas of the conference: Stephen tried to SELL cellular automata post-cards.

Chris Langley (automata, LIFE, genetic algorithms):

A consummate hacker and a beautiful system. Excellent interface, very fast processing, complex theory. He demonstrated, in living color, self-reproducing automata, foraging ant-automata, and the best LIFE processor I've ever seen (10^5 pixels, 1/2 second update). Works for Holland at Michigan.

Summary of Machine Demonstrations

Humans are no longer in competition for games. Even very simple learning algorithms outstrip our tricks. The relative rate of evolution of ability is about 20 to one in favor of machines over humans.

Automata are startlingly beautiful, approaching a visual aesthetic.

The quality of the interface dictated the assessment of the program.

DAY 3

John Conway (Cambridge, pure mathematician and DUDE):

How LIFE was invented. Most random sets of rules will produce a Universal computing machine (Turing equivalent), the problem is finding the subset of the space of configurations that is Universal, and proving it so. Emergent phenomena are not necessarily describable in a condensed form. (Prediction is behavior, again.) Modern mathematics is an experimental discipline guided by a juggler's intuition.

L. Valient (Harvard, psychological computer scientist):

Assumptive rehash of generalization of Predicates. Fifteen years out of date. Learning is the computational deduction of a recognition algorithm. Granularity is the main issue. I/O testing is never sufficient for learning. Pigeons overgeneralize.

John Denker (Bell Labs, game-author physicist computer scientist)

My vote for the most provocative research. Neural networks come of age. 10^5 electronic neurons, each consisting of a resistor synapse and dendrite-axon connections, all fully interconnected. This architecture can be configured by adjusting the resistance of each element into a system that models gradient fields. Learning occurs by local modification of the gradient topology; fixed-points form the stable system responses, and thus the memory of what is learned. Extremely robust and fault tolerant; holographic distribution of memory (ie: the system still works even if half of the electronic neurons are broken). Known matrix algebra techniques drive the optimal formation of gradient fields. Think in terms of eigen-behaviors: fixed-points of dynamic fields as evolutionary stable strategies. Through-out the conference, John kept reminding speakers that neural networks have evolved way beyond perceptrons.

Puye Cliff-Dwellers (circa 1300 A.D., damn Cortez):

The afternoon was for sightseeing. Vibes of ancient ceremonies intermingle with atomic fusion. The Bomb and the Canyon are both secret.

End of Day 3 Comments

Score:
Pure Mathematics: infinity
Psychology: 0

Idea of the Day:
Universality is common; recognition is rare.

Most Pregnant Research:
Neural networks

Raspberry of the Day:
Psychologists trying to do Computer Science.

Wisdom of the Sages:
Play now, formalize later, but formalize eventually.

DAY 4

Stephan Brams (NYU, mathematical political scientist):

Game theory can be modified by probabilistic payoffs to give stable strategic results. Specifically, if we threaten the Russians with probabilistic retaliation conditioned inversely on the extent of their threat, a stable cooperation can evolve. The inversion comes from the incredulity of big threats. Good mathematics based on silly premises (a common mode); who says the superpowers are playing the same game? Brams is better known for the game-theoretic analysis of God: omniscience is over-rated because, in some games, it doesn't help One to win. I need a new category similar to science fiction: mathematic fiction. The story (or proof) is credible, but the premises are outlandish.

Stanley Reiter (Northwestern, mathematical economist):

Heavy mathematical modeling of the goods consumption mechanisms. Memory is equivalent to communication; just allow the message to carry its antecedents, which it does anyway for comprehensibility. There is a direct trade-off between message density and complexity of process: more communication simplifies processes. The conclusion is that free speech simplifies government, and inversely, Communist regimes foster unsolvably complex management problems. There is something here for system architecture. I fell asleep during the details.

Michael Conrad (Wayne State, biologist computer scientist):

Evolutionary processes are self-organizing. An attempt at the algorithmic specification of evolution. This work started well and then degenerated into complexity. The simple model produced strange behavior, and rather than finding the original source of mis-modeling, Conrad added much more complexity (patches) until it was impossible to follow. Then he had his graduate students add still more "features". A good example of bad research.

J. Deneubourg (U. Brussels, ant biologist):

Individual ants learn, and display idiosyncratic behaviors. Behavior of species and colonies differ markedly. Ants know how to amplify individual learning to group activity. When you look closely, a single ant is not a automaton, it is fairly focused in its behavior. New hope for simple machines learning complex patterns.

Jack Cowan (U. Chicago, neuropsychologist):

This guy spoke like a DUDE. All animals have a topological map from retina to cortex. The wiring is very plastic and can regenerate. By surgical intervention (like slicing an eye in half and regenerating the optical fibers), we can map the connective process. Generally the topology and the boundaries of the mapping is consistent over a wide variety of mutilations and interventions. The technical details of this way of exploring brains are complex, but can be modeled by very general dynamic equations. Experimental surgeons have tried ingenious experiments, and strange phenomena occur (eg: interleaved multiple maps from eyes added onto backs). Somehow the neural connections find the right pathways to preserve topology.

John Holland (U. Michigan, computer scientist AI):

An AI architecture that learns. Three nested systems: the inner most is the standard rule-based expert system with its goals coded as meta-rules; on top of that is the simple learner, driven by the utility (usage) of the existing rules; and on that is the discovery learner, the genetic algorithm. The simple learner addresses the credit assignment problem. All rules fire in parallel. When the environment changes so that it rewards (via a payoff structure) the system, the active rules share the payoff. Payoff propagates through the rules by a commodity model: if R1 calls R2, when R2 is successful, it pays R1. The actual mechanism is more complex, and is regulated by passing messages between rules. All messages are coded as bit arrays, one bit per feature (so for example, 101 might mean YES trunk, NO wings, YES grey). A # is used as a don't care bit; it is the means of generalization. Firing of rules actually means matching bit sequences. So

a rule looks like: 10##0#1//####101, with the double slash separating condition from action. Actions are messages, and the rule pays to post its message, which might trigger either payoff (feature recognition) or other rules. The genetic algorithm crosses rules randomly, just like genetic cross-over mutation. Weak rules are weeded out because they don't get called and consequently don't get paid. Strong rules compete for payoff and those that survive (get lots of payoff) are effectively learned, or created. So wealth via usage is equivalent to biological fitness. Strong rules give birth to children which increase the chance of mutation of that rule. Extinction of weak rules leaves blocks of rules that act like rule sets. We are talking 10^5 iterations for things to happen, but they are serious enough to have built a chip. This work created a lot of enthusiasm, and Holland's followers are devoted. On close inspection, the idea looks clunky; the enthusiasm might be from lack of AI knowledge. However it is the first to actually achieve evolutionary modeling.

End of Day 4 Comments

Score for the Day:
Frog mutilators: 1
Concept mutilators: 1

Idea of the Day:
Each ant is an individual.

Second-place:
Memory and communication are indistinguishable.

Get Em Gossiping Award:
Genetic algorithms let expert systems evolve.

Raspberry of the Day:
If you don't understand something, have your graduate students make it more complex.

Gross Me Out Award:
Take half an eye, attach it to the wrong half of a brain, and watch the frog jump backward toward its food.

DAY 5

Marcus Feldman (Stanford, biologist and semi-DUDE):

All you computer scientists modeling biology don't know any biology. The genetic algorithm, for instance, is haploid, so it models some rare fungi. All animals are diploid in their chromosomes. Further, "random" mutations

are under genetic control. That is, there are genes that specify the fragility of the chromosome. The whole biological story is much more complex than you suspect. I don't remember the technical details.

Terence Sejnowski (Johns Hopkins, computer scientist):

It will take about a million CRAYs to make the ALV work. In the brain, memory and CPU are coextensive. Perceptron architecture has evolved into Boltzman machine architecture by introducing intermediate (hidden) units between input and output. Have a fully connected array of weighting elements that learn patterns by probabilistic updating of weights. Performance levels out at 95% after 70000 trials. Noise in the system is good in that it helps the system get out of local minima. A good idea that looks like it will die out.

James Anderson (Brown, computer scientist):

I skipped this one.

End of Day 5 Comments

Score for the Day:
      Computer Scientists: 0
      Biologists: 2

Idea of the Day:
      Genes control their own mutations.

Computational Explosion Award:
      Progress in Biology is occurring faster than our understanding or modeling of it.

SUMMARY

What do they call a dumb biologist? A physicist.

That is, self-referential domains dominate reality. By simplifying out the fundamental process, by separating the observer from the observed, we guarantee an understandable and totally inaccurate model. Even pure mathematics is becoming an exploration of self.