

310:11:
3/4


Naval Research Reviews

SERIAL DISPLAY
DOCUMENTS
LIBRARY
DEPOSITORY

JUN 15 1992

UNIVERSITY OF ILLINOIS
AT URBANA-CHAMPAIGN





Office of Naval Research
Four / 1991
Vol XLIII



Biological Neural Networks

Digitized by Google

Original from
UNIVERSITY OF ILLINOIS AT
URBANA-CHAMPAIGN



Neuro Network Research

The Swedish Post Office printed and released these five stamps to honor winners of the Nobel Prize in Physiology and Medicine. The Office of Naval Research's (ONR) Life Sciences Directorate has had the honor of supporting a number of Nobel laureates in various fields. In the field of biological neuro network research, ONR has been a pioneer. George von Békésy, noted on the first stamp, was an ONR contractor while at Harvard in the late 1940's and 50's; he was recognized for his fundamental discoveries concerning the dynamics of the inner ear. Two articles on acoustics by Professor von Békésy appear in ONR's anniversary book, *Office of Naval Research: Forty years of Excellence in Support of Naval Science*. A second Nobel Laureate, Professor David Hubel of Harvard, is honored by the last stamp with his co-winner, Torsten Wiesel, for the neurophysiology and neuroanatomy of the visual system of higher mammals. Professor Hubel continues today his long association with the biological intelligence programs at ONR.

Naval Research Reviews

Office of Naval Research
Four / 1991
Vol XLIII

Articles

2

Interview with
Dr. Allan Bromley

10

Neural Systems For Eye Tracking

*Terrence J. Sejnowski
Stephen G. Lisberger*

23

Chemical Control of
Neural Networks

Allen I. Selverston

9

Introduction

Joel Davis

17

Exploring Biological Neural
Networks Using Realistic
Computer Simulations

James M. Bower

29

The Role of Interdisciplinary
Research Involving Neural
Systems In The Development
of Intelligent Systems

Dr. Thomas M. McKenna



Departments

16

Profiles in Science

37

Research Notes

CHIEF OF NAVAL RESEARCH
RADM William C. Miller, USN

DIRECTOR
Dr. Fred Saalfeld

CHIEF WRITER/EDITOR
William J. Lescure

SCIENTIFIC EDITORS
*Dr. Robert J. Nowak
Dr. J. Dale Bultman
Sarah Scott*

MANAGING EDITOR
Norma Gerbozy

ART DIRECTION
*Typography and Design
Desktop Printer, Arnold, MD*

About the Cover

Shown as a background texture on the front cover is a portion of the memory structure layout used in Adaptive Solutions' CNAPS-1064 very large scale integrated (VLSI) neurocomputer. This architecture is unique because it is a network model based on neurobiological research data, which is briefly described in the inset. This Lynch-Granger model of olfactory or piriform cortex is a network model developed at the University of California, Irvine, that has been derived from extensive research into olfactory processing observed in laboratory rats. Broadly speaking, the model describes a layered, feedforward, unsupervised network that performs a hierarchical clustering of training data. Feedback paths from the cortical network to the input synchronously cancel out successively dominant components of the network input. Both the "hardware" and "wetware" portions of this research were supported by the Office of Naval Research's Biological Intelligence Programs. Acknowledgement for the cover artwork goes to Massachusetts Institute of Technology Press.

Naval Research Reviews publishes articles about research conducted by the laboratories and contractors of the Office of Naval Research and describes important naval experimental activities. Manuscripts submitted for publication, correspondence concerning prospective articles, and changes of address, should be directed to Code OPARI, Office of Naval Research, Arlington, VA 22217-5000. Requests for subscriptions should be directed to the Superintendent of Documents, U.S. Government Printing Office, Washington, DC 20403. *Naval Research Reviews* is published from appropriated funds by authority of the Office of Naval Research in accordance with Navy Publications and Printing Regulations. NAVSO P-35.

Original from
UNIVERSITY OF ILLINOIS AT
URBANA-CHAMPAIGN

An Interview With Dr. D. Allan Bromley

Dr. D. Allan Bromley is The Assistant to the President for Science and Technology and Director of the Office of Science and Technology Policy (OSTP) in the Executive Office of the President. For more than two decades, Dr. Bromley has been a leader in the national and international science and science policy communities. As President of the American Association for the Advancement of Science and of the International Union of Pure and Applied Physics, he has been one of the leading spokesmen for U.S. science and for international scientific cooperation.

One of the world's leading nuclear physicists, he has carried out pioneering studies on both the structure and dynamics of nuclei and is considered the father of modern heavy ion science. He is currently on leave from his former position as Henry Ford II Professor of Physics at Yale University, where he was founder and Director of the A.W. Wright Nuclear Structure Laboratory.



Interview By Dr. Joel Davis, Office of Naval Research

Thank you very much for making time available for this interview.

Q: Dr. Bromley, how does science policy get made in the White House and how much input is received from other White House officials?

A: Well, you have to understand that I wear two hats here. First, as the Assistant to the President for Science and Technology, my responsibility is to make sure that the President and the other senior members of the administration have the best possible science and technology input to decisions where science and technology may be an important component but where, in many cases, it is by no means the dominant one. In those situations, almost as important as having the science itself is to give the individuals involved a feeling for how reliable or how certain our understanding actually is. In the past, there have been many cases where policy has been made on the basis of very shaky science because something was presented as scientific fact without any error bars. It's very important to have a relevant degree of certainty available as well.

Q: You try to provide the error bars?

A: Yes. Where does this sort of policy get made? Well, predominantly it gets made in the cabinet councils, including the Domestic Policy Council, the Economic Policy Council, the Space Council, and the Competitiveness Council. One very

good thing, from my point of view, is that, in contrast to my predecessors, President Bush has elevated my position so that I am a member or participate in all of these cabinet councils. I am present for all their discussions, and that has some very important advantages. First, it means I don't negotiate, as my predecessors did, to get into meetings. But second, in addition to providing an opportunity to inject science and technology into the discussions as they are progressing, it gives me an opportunity to recognize when something is happening in the discussion that could have a major impact on science and technology. This could pass totally unnoticed if there weren't somebody around sensitive to these matters. Both of those areas are very important.

The other "hat" I wear is that of Director of the Office of Science and Technology Policy (OSTP). There, the primary responsibility is policy for science and technology. In this capacity, I report not only to the President and to the White House but also to the Congress because OSTP was set up jointly by the Administration and the Congress in 1976. One of OSTP's responsibilities is that of insuring a healthy, vital, productive science and technology enterprise in this country. So we have, on the one hand, issues that are within science and technology, and these we try to resolve within OSTP. With the others, concerning science and technology for policy, we draw on everything available to us and move the science and technology information into the cabinet councils and then on to the President.

Obviously, one of the most important linkages for OSTP is with the Office of Management and Budget (OMB). We have, over the past couple of years, developed a very good working relationship, so that people at all levels in my office

interact directly with their opposite members in OMB and work directly with people in the agencies as well. We try to be of as much assistance as possible to OMB when they consider the budget. This is a very different situation from the past, where this office tended to get a look at the budget after OMB had assembled it. We could comment on it--but were not involved in its actual creation.

Q: So it's presently interactive rather than reactive?

A: Yes. Now there is the question of how you make policy for science and technology? First of all, you have to recognize that we are unique in the world in that our entire science and technology program is built from the bottom up rather than the top down. In most other countries there is a science and technology budget that is voted by the appropriate legislative body, and then it's a matter of dividing that budget up in various areas. Here we have a much larger number of players at the federal level; there are more than twenty federal agencies that have something to do with science and technology. The desires of the individual scientists and engineers manifest themselves in proposals, and by the time you get a proposal you have already pulled together a group of some size that believes it is a great idea. These come up through the agencies, and the agencies filter them out. One of the good things about our system is, if one agency doesn't like a good idea, for whatever reason, the chances are very good that some other agency will find it appropriate to its mission. Generally, really good ideas tend to get supported somewhere within the system. I think that diversity of support is one of the dominant secrets of why the United States has the strongest science and technology enterprise that the world has ever seen. However, it's very important that the activities of these agencies be coherent. The congress recognized this back in 1976 by creating something called the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) at the same time OSTP was created. That council worked sporadically and not terribly effectively. Very soon after I arrived here the reason became obvious; the membership in the council was at too low a level in the various agencies. The decisions reached by the Council could be disowned by more senior people during a budget crunch. Now, with strong support from the President, we have restructured FCCSET (pronounced "FIXIT"), so that it has as members the cabinet secretaries and under secretaries and the heads of the independent agencies that are pertinent to science and technology.

Decisions once made now stay made. I have been enormously impressed at the level of cooperation and interaction that has developed within the FCCSET structure. It provides a forum where ideas can be examined by various agencies and it makes sure that agency programs merge with those of others to form a coherent national program. In the end, it gives us the possibility of guaranteeing both to the Congress and to the President that we're getting the maximum amount of research and development from a given investment.

Q: I was going to ask you a question later about the Decade of the Brain but would you consider this to be an example of....

A: The Decade of the Brain is slightly different, I'll come back to that in a moment if I might. Let me first of all emphasize, however, that there is a fundamental gap in the FCCSET structure -- by law, all its members must be federal employees. There is no obvious mechanism for bringing private sector input into decision making and policy making at the higher levels of government. For that reason, I am delighted that the President has chosen to pull together a President's Council of Advisors on Science and Technology (PCAST), which reports directly to him and meets on a monthly basis. This is a unique opportunity for the President to get outside advice and counsel unfiltered by any bureaucracy. Since I have the privilege of chairing both FCCSET and PCAST, it's my responsibility to be sure that their members talk together and that FCCSET activities are calibrated and validated by the private sector.

What is the outcome of this process? Each year after sitting down with the agencies, with Office of Management and Budget (OMB), and with the full FCCSET, we select particular interdisciplinary areas that don't fit comfortably in any one agency or perhaps in any one scientific field. We select these for crosscut analyses, where we try to find out what exactly is being done in this area across the entire federal government. That effort provides us with an inventory, usually for the first time, on what is relevant that is actually going on in all the agencies. We then use that as a base on which to build a coherent program that will take that field forward. Last year, for example, the President selected the three crosscut areas that we studied as special budget initiatives. They were global change research, mathematics and science education, and high performance computing and communication. This year we are adding two more to the list that are currently in progress: materials science and technology and biotechnology. In these cases, then, we bring forward not just a heterogeneous collection of agency programs that are loosely directed toward a topic but rather a coherent, integrated, national program that focuses all the agencies' activities on the issue at hand. I have been very pleased by the willingness on the part of the people in the trenches in all the agencies to put a tremendous amount of work into this and to come up with the kind of program of which we can all be proud.

But I do not want to leave you with the feeling that FCCSET is only concerned with activities that have strong budgetary components. An example without specific budgetary activities is in the area of risk assessment and risk management. Increasingly, all the agencies in the federal government are concerned about matters of risk. It is a critical, integral part of their activity; but if you look carefully, you find that the basic principles and approaches applied are very different depending on the specific agency. We put together a group under FCCSET that is chaired by the deputy director of Envi-

ronmental Protection Agency (EPA) and has representatives from all the agencies. This group is charged with developing a consensus with respect to the principles and approaches that will be used by all the agencies in matters of risk assessment and risk management. We are not trying to come up with a manual of practice because there are substantial differences among problems from agency to agency. It would be impossible to have one sort of cookbook approach, but if we have this basic agreement on what the approaches and principles are, and we recognize that risk assessment has, in the past decade, become a real science in its own right with university activities of considerable magnitude, then I think we can do a much better job of tackling this question throughout the federal government. That's one example. Another example is a consideration across all the agencies as to how the United States should respond to the opportunities or problems that may arise in 1992 when Europe finally comes together as a integrated entity. There are both problems and opportunities here, and we want to be able to approach them from the vantage point of all the different agencies in the most effective way. That's a much too long answer, but that tells you something about what's going on.

Q: What are your current thoughts about the role of DoD agencies that support science and technology; certainly I'm speaking of The Office of Naval Research (ONR) and The Office of Naval Technology (ONT).

A: First of all, it's important to emphasize that the concept of investing in the future through science and technology is one that President Bush had clearly established even before he was elected. We are very fortunate in this Administration, and it's unique, when you stop to think of it, that John Sununu, Dick Darman, Mike Boskin, and Roger Porter are all strong supporters of science and technology. There is intrinsic support throughout the White House, and these are all people who understand science and technology. Given that, what I've been doing is to try to emphasize both science and technology more than has been done in the past -- the T in OSTP, to be specific. I have appointed, for the first time in the history of this Office, an Associate Director for Industrial Technology. I did that for two reasons. One: because I think that it is extremely important in terms of our economic competitiveness and national security. Two: because I wanted to send a very explicit message to the industrial community that the Bush administration is serious about working with it. These developments to which you refer are not changes in direction as a lot of media articles have been trying to argue. They are parts of an evolutionary process, where we've been focusing on various parts of the overall S and T program and moving them forward. For example, in the past there has never been any statement as to what U.S. technology policy really was. None. So shortly after arriving here we undertook to develop such a statement, and it was issued in September 1990. It's not the world's greatest document, but it's important because it was signed off on by

everybody in the White House and it provides a basis on which we can build.

The important thing is that there has never been any question about the appropriateness of federal support of basic research because by its very nature you cannot tell when, where, or to whom the benefits of basic research are going to flow. Therefore, no individual organization can usually justify the necessary investment. I maintain that exactly the same thing is true for generic technologies, those extending up to the competitive phase of development.

It used to be that the civilian sector was built strongly on research and development done in the military. Now, one would have to say that the military builds strongly on research and development done in the civilian sector, and in the future our national security probably will depend just as much, if not more, on economic strength as on military strength. But, by the same token, Desert Storm has made it abundantly clear that the edge we have is a technological one. It's a technological edge not only in weaponry but also in training, and that's something that people have not fully appreciated as yet but will.

It's still very clear that we, in this Administration, do not believe that it is appropriate for the federal government to select winners and losers either in terms of selecting specific industries or selecting specific technologies that have reached the stage where potential users know what the application is going to be. But we do believe that the federal government has a very real role in that period from the basic discovery -- where the federal government has in most cases already played a very strong role -- up through the point where you can assess the technology for its application in a number of areas.

We know that our competitors abroad are, in almost every case, working coherently. The government is working with the various industrial players in those countries to share the risk, share the cost, and reduce the development time of the basic technology to the point where people can say, "Yes, I can see how that would apply." Then the competitors split up and compete vigorously with one another and with the rest of the world. That's healthy. That's what we're suggesting.

Does the defense department have a role to play? Indeed it has a major role to play. If we go back to the postwar decades, the late 40's, the 50's, and the early 60's, the dominant supporter of research and development in the United States was the Defense Department. That was good for many reasons: it was an enlightened support mechanism that had been invented in the Office of Naval Research by Robert Conrad and Emanuel Piore, and it served the nation extraordinarily well. The basic principle was that you find the best people you can, as judged by peer review, you provide whatever funding you can to support them to do what they want to do, and you keep the hell out of their hair while they're doing it. Those are principles that continue to serve us well. During the Vietnam period, however, many of the bridges that had been built between the military and the universities were destroyed, and both sides

suffered. The military needs contact with bright young people and with the senior faculty in the universities. And the university needs contact with the military to have some real world problems to work on and to feel that they're part of the defense of the nation. It was pointed out a long time ago that in a democracy, the people who are charged with the security and defense of the citizenry can't feel themselves alienated from that citizenry, or you have a very dangerous situation. I'm very pleased now to see that the Defense Department is rebuilding these bridges with the universities. I expect that these are going to grow substantially.

We have a very serious problem that we have to address at the national level. Our defense industries contain an enormous national resource of know-how, technology, and personnel, but as our defense expenditures go down, as everyone agrees they will, we have to be creative to make sure that we don't lose this tremendous resource by inadvertence or by neglect. At the federal level, we must be aware that we have a responsibility to make sure that these talents are used in other areas of national importance. The defense research and development agencies are going to have a very critical role for the next few years. First of all, they are going to be involved in this reorientation of much of the research and development expertise in the country. Also, at a time of reduced international tension, we have to be very certain that we protect ourselves against being technologically blind-sided. I've argued, and will continue to argue, that although the total defense funding in the country will go down, the basic and applied research in the Defense Department is essential because this is just the time when we need that increased understanding to protect us against technological surprises.

Q: You've just recently come back from a trip to the Soviet Union, do you think that the startling, rapid changes there will influence science in US and Western Europe?

A: Very much. There will be major changes because the Soviet pattern was one in which education was controlled by the Ministry of Education, while research was controlled by the Academy of Sciences and to some degree by the State Committee on Science and Technology. As a result, there has been a definite separation between research and development and education. This has already been recognized by the Soviets as something that is going to change, and they are changing in our direction as rapidly as they can. In the Soviet Union, for a very long time--since 1917 in fact--the academicians have had a privileged life; they have had control over large sectors of the intellectual life of the country. I think that as the Soviets move to a more democratic governmental structure, control will become much more diffused and the role of the academy will no longer be as unique as it has been since 1917.

Q: Who will take that role?

A: We're already beginning to see the development of scientific entrepreneurs throughout the Soviet Union. For exam-

ple, on a recent visit I was delighted to find that one of the absolutely world class laboratories in the Soviet Union is headed by an old colleague and friend, Alexander Skrinksky. He's an academician--a young academician--but he decided many years ago that if he were going to have the kind of budget that is required to do the kind of research he wanted to do, world class research, he was going to have to get funding over and above what was available from the central system. He therefore developed factories and now makes high precision instruments and equipment that are sold around the world. In fact, right at the moment, that group is doing some very important studies for the superconducting supercollider (SSC) group in Texas. They happen to have unique facilities for it. We're going to see a lot more of this entrepreneurial development in the Soviet system. When I talked to the people who were doing geophysics, I found that they have built, on the side, development factories where they make enormous quantities of industrial diamonds. They make jewel quality emeralds, for example, that are now going out on the world market, and they pick up a very substantial fraction of their operating expenses from these entrepreneurial activities.

When you talk to Soviet Scientists these days, their highest priority areas of discussion are two in number. The first is: How do you convert from a centrally controlled to a market economy? I'm afraid that as yet the Soviets are still hoping that they can find a painless way to make the transition. I'm not convinced that there is a painless way, and until such time as they recognize that, I don't think the transition will actually be accomplished. Secondly, the Soviets have an enormous resource in their defense-related manufacturing industries and are terribly concerned about how to convert that resource into civilian use. They will look to us for as much help as we can provide: frankly we're not as certain about how to go about this as they think we are. It's a problem we all have to solve.

Given the state of the Soviet economy, obviously there will be a serious constriction in the amount of funding that can be made available to research and development for the next few years. One thing that I think is very important is the new Joint Commission on Basic Research that was established in 1989 between the US and the USSR. This Commission is going to play a very important role because there already have been a large number of individual-to-individual cooperations developing between scientists and researchers in the Soviet Union and in the US and between institutions in the Soviet Union and in the US. This marks a total change. In the past--and I've been involved since 1972--it was all from the top down. A group met in the Kremlin or here in Washington, and we decided upon the areas in which we were going to cooperate, and we picked the people who are going to be involved in the cooperation. Starting with this latest Commission, which I have the privilege of chairing from the American side with Deputy Premier Nikolai Laverov (who also happens to be Vice President of the Academy and Chairman of the State

Committee of Science and Technology) on the Soviet side, we have agreed that the commission will act to facilitate individual and institutional cooperation, and we'll do it from the bottom up. The National Science Foundation (NSF) and other American agencies are now getting proposals from joint programs that have both Soviet and American sponsorship. In parallel the Soviets have established their own fund for the support of such joint research and will also accept joint proposals for consideration.

There are some marvelous things the Soviets can bring to this cooperation. For example, there are eight kilometers of undisturbed sediment in the bottom of Lake Baikal, the world's biggest lake representing 20 percent of the fresh water on the planet. Eight kilometers of undisturbed sediment provides a permanent record of climate and the biosphere for hundreds of millions of years, and we have just begun to work with the Soviets to mine some of that record. Since the Soviets do not have the computer facilities we do, they have traditionally focused on applied mathematics, analytic mathematics, and software; and there are a great many areas where there is a very strong potential for a beneficial cooperation. We can take advantage of the algorithms and software that the Soviets have developed based on this analytic mathematics and put that together with our hardware capabilities. In yet another area, one that dates back to the early 1950's, the Soviet Union decided that making very pure materials was going to be important, and they have evolved techniques for making almost any kind of pure material. Some are more than two orders of magnitude purer than are available anywhere else in the world, and these are becoming very important as we move to micro-and nano-electronics and nano-fabrication. They've always been very strong in optics and laser work. They invented the fusion Tokamak--and there are a great many other areas where we can cooperate as equals and with mutual benefit.

Q: Many of us in the Life Sciences have been pleased to see OSTP and yourself taking a positive role in organizing the "Decade of the Brain" activities; and in fact, I remember reading a striking quote of yours, about understanding the cosmos "within" as opposed to the cosmos "without." Some very optimistic goals have been set in this program in both the medical and technological areas. Do we have a chance of doing that kind of thing in the next couple of years?

A: I will be surprised if we don't far exceed the goals. For example, we thought that we had set very ambitious goals in our program for high performance computing and communication, and we discovered that within the first four months of a five year program we had already reached the goals we'd established for the next three years. People respond. The very fact that goals are established serves to focus people's attention. It accelerates progress because it tends to attract people from areas that may at first be viewed as peripheral to the problem at hand. But, people very quickly discover that they

have many coherent contributions that can be pulled together. I'm very optimistic.

We recently had a briefing for the President's Council of Advisors on Science and Technology from Roger Porter of the National Institute of Health (NIH), and he mentioned the rate at which the technology that allows you to non-invasively study what is going on in the human brain is advancing. For example, positron emission tomography (PET) scanning and scanning with ultrasensitive magnetometers, where we can begin to focus not just on brain structure but on brain function is going to lead to breakthroughs that we can't even imagine at this point. I'm very enthused about it, and I think we're off to a good start. I had the pleasure of releasing our report on the "Decade of the Brain" at the annual meeting of the American Academy of Neurology in Boston. There was very real enthusiasm. We've now entered a time when there has been a major move forward in the scope and power of the technology. The conditions are extremely good for a real breakthrough. I'm only sorry that we weren't able to do this while Representative Silvio Conte was still with us because he was the first person who came to talk to me about these issues when I first arrived in my office. He specifically wanted to talk about what we could do in the decade of the 1990's to further understanding of the brain, and with his help, and with that of the President, we were able to get the 1990's designated as the Decade of the Brain. I'll be extremely disappointed if we don't have some major surprises that none of us could now predict.

Q: Along those lines, my particular interest in what I do at ONR is neural networks and I wonder if you are familiar with the area?

A: Oh yes! It is discipline that has interested scientists from both the physical and neural areas, and there's no question that it's going to be an exciting area as we move forward. What will be of great importance is not just these various initiatives that I've talked about but the cross fertilization between them. For example, the high performance computing and communication initiative is moving forward at tremendous speed. The capacity that we have for handling information and for presenting almost any database in graphic form takes advantage of a unique characteristic in the human brain. People tell me that about 75 percent of the acting brain is processing visual information. We have a unique ability to detect patterns and develop hypotheses when you see something in graphical form that escapes us completely when we see a great stack of computer printouts or lists of numbers. This makes possible a qualitative as well as a quantitative difference in what we can do.

We know that we're going to need the kind of scalable, massively parallel configurations. You get the kind of increasing scope and speed that is essential if we're going to be able to handle the data that are coming from the probes that we are sending up as part of our initiative in global change. In The Mission to Planet Earth program, scheduled for the latter part

of this decade, the EOS A configuration has been designed to send back information equivalent to the entire library of congress every 4.8 days. Unless we have greatly increased our sophistication and speed and information handling, we will simply be deluged with data that will have no value. We are now trying to develop the hardware and the software that can cope with these data flows. The point I'm making is simply that all of these Presidential initiatives are tied together by technology.

Q: We don't see many big science projects in the life sciences. Do you think the Human Genome project is a trend in this direction?

A: It really isn't big science. It's many individual researchers with their students working on a coordinated program.

Q: The superconducting supercollider (SSC) will require eight billion dollars, triple the estimation of a few years ago. My question to you is can the country afford it and/or is this the best way to use scarce scientific resources?

A: The country can afford it. It is to civilization today what the cathedrals were to those in the 16th century. It's the pushing outward of the frontiers of human knowledge to the farthest reaches. The SSC is designed to answer some of the most fundamental questions that we could ask of nature. Where did mass originate?

If we were doing it over again, I would have preferred to plan and build the SSC on a much more international basis. We could have come together as an international community and asked ourselves: What is the next question that needs answering to push forward our common frontiers? We could then decide what we needed in the way of facilities and instrumentation: where we were going to put them, how were we going to fund them, how we were going to build them, and how we could take advantage of the best possible international expertise available. That is the wave of the future. The SSC is probably the last large scientific project that will be handled on an individual-nation basis, where we come up with the idea and design and then invite other nations to join us. That's not going to happen in the future.

Q: Along these lines about major projects, how about the space station and it's cost relative to the microgravity research that is going to go into it?

A: As I have said repeatedly in the past, the space station is not now, and never has been a scientific project. It is fundamentally an adventure. It is the first step on the critical path that will take mankind away from Earth. And again I feel that, as the last remaining superpower on the planet, if we can't spend 0.2 percent of our GNP in taking our species off the home planet, then there's something fundamentally wrong with our priorities. But we should not make the mistake of thinking that we are doing this because we want to do some specific science. The microgravity and life science experi-

ments that have been discussed so far do not, in themselves, justify building the space station. But once we build the space station as the first step in a great adventure, then there's some wonderful science that we can do on it. The first and most important is the life science that will tell us what we need to know before we can take the next vital step into space, into permanent exploration and exploitation of space.

There is a tremendous amount we don't know. We don't know whether imposing a tenth of a g would prevent the unfortunate effects of weightlessness. Or maybe it's 0.8 g, or maybe nothing happens until you're at 0.5 g and you've got a threshold. These are the things we're going to have to find out. How can man survive in a fundamentally hostile environment at 0 g?

There are also a great many interesting things we will want to do in microgravity. We will have microgravity for several years during the construction period in between the actual shuttle flights. The station will be absolutely quiet in these intervals, and we can set up marvelous microgravity experiments. But again I would say that they, in themselves, do not justify building the Station.

I'm concerned that so many people in the scientific community have a feeling that funds for the station are fungible, that if they're not spent on the station they will be spent on science and technology of greater personal interest to these individuals. That's completely wrong. Both the station and the support of science and technology that has been a hallmark of the Bush Administration are investments in the future. We're not arguing and debating about space versus science and technology, or about big science versus little science. The debate is about whether we're going to invest in the future or whether we are going to support immediate consumerism. Are we going to build parking garages? Are we going to do various things that are of immediate significance today, or are we looking to what our children and grandchildren will do and what they will need to take them to their frontiers?

Q: The well publicized Stanford case has reopened some of the questions about overhead charges by universities doing federally funded research. Do you have any thought about this process and should it be altered, and how?

A: You may recall that back in 1986 David Packard and I wrote a report for the Reagan Administration's White House Science Council entitled On the Health of U.S. Colleges and Universities. A very large section of that report was devoted to discussion and to recommendation concerning the question of indirect costs. It's very important to emphasize at the outset that indirect costs are just as real as direct ones. There tends to be a feeling that somehow they are not real or not as real. They are. It is also true that the indirect cost situation is complex, it's difficult to understand, and very few people do take the time and effort to understand it. As a result, it has been a major source of friction and controversy between universities and the federal government and between university faculty and uni-

Original from

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Four/1991

7

versity administrators. There has been an enormous amount of misunderstanding and wasted effort involved.

We recommended a whole series of changes in 1986. The universities were not ready to accept them then, but within this past year, I sent our report back to the American Association of Universities and asked them to take another look. Strangely enough, they now think that it's a very good report. Part of the reason is an unfortunate coincidence between things like the Stanford yacht, the Baltimore case, the South Carolina case, and a series of similar unfortunate isolated incidents that came together in time. These have tended to convince people in the Congress and in the public that there has been a breakdown of ethics and standards in the universities. I don't happen to believe that that's true. But it has made it very attractive from a political sense to take some form of punitive action against the universities. I don't think that's wise either because we make a greater demand on our universities than any other country in the world. We expect them to develop new knowledge, and at the same time to train young people to use that knowledge effectively. We have to take a whole new look at the question of indirect costs. The Bush Administration is recommending that a cap be placed on indirect costs in the administrative subcategory because that's the area where there is the greatest creativity on the one hand, the greatest argument, and the greatest difficulty in understanding. The cap was set at about the weighted average currently in the United States (26 percent). There have been moves in the Congress to come up with a lower number, and there have been moves in the university community to come up with a higher one. I think that we should look on this as a holding situation; and within the next eighteen months, we will have an opportunity to resolve these problems in a more complete and final fashion. It is important to note that the Administration task force, led by OMB, made a special effort to work with a representative group of university presidents and will continue to do so. We will be involved very heavily in this task force as will many of the agencies in OSTP to really rethink these questions. In the long run the universities are probably going to require federal help to get their facilities and their infrastructure back to an acceptable level. Once that happens, however, we should arrange so that the cash flow to the universities is such as to allow them, if they manage their affairs in a reasonable fashion, to maintain themselves on a self-sustaining basis. If it becomes necessary to maintain the infrastructure, to have higher overhead rates than the current average, this will clearly require funds that would otherwise tend to go to research. There will always therefore be tension between the direct and the indirect aspects of these costs. But they're both real and they both have to be paid.

Q: What do you think has been your greatest contribution to US science policy?

A: I'm absolutely the wrong person to answer that question. If you ask me a different question, which is: What am I most

pleased with since I've been here, I can give a few answers. First of all, I've been extremely gratified by the support that I have received from the president, from all my colleagues in the White House, and from the Hill. There is a tremendous reservoir of good will in the Congress and a generally shared belief that we as a nation are underinvesting seriously in science and technology. Secondly, I have been enormously pleased and proud of the quality of the people that have agreed to join me here in OSTP. They are an absolutely outstanding group. Thirdly, I think that we have made a real step forward in getting the FCCSET mechanisms to work in pulling together, coordinating, and integrating what goes on not only among agencies and within government but also between government and the private sector. Those are the aspects that I am particularly pleased about. A tremendous number of opportunities, however, are still out there.

Introduction

Joel L. Davis

*Program Manager, Biological Intelligence
Cognitive and Neural, Sciences Division, Office of Naval
Research*

This issue of Naval Research Reviews represents the second opportunity for me to “showcase” some of the outstanding work being performed by some of our ONR investigators. The previous issue (37(4): 1985) predicted that biological neural networks were shortly to be developed (“We are at the point where realistic and powerful models of cellular “information processing” networks underlying learning and memory can be developed”). Seven years later, this issue of Naval Research Reviews presents work from three laboratories well on their way to precise functional descriptions of neural processes.

The first article represents a novel addition to the usual Naval Research Reviews articles. Dr. D. Allan Bromley is the Assistant to the President for Science and Technology and Director of the Office of Science and Technology Policy (perhaps better known as “the Presidential Science Advisor”). In this position he has an inside view and major influence on the support of science in this country. The interview took place in September, and although the intervening months have seen some policy changes, the basic description of the role of OSTP in the political process remains the same. My interviewing strategy was based on some “priming” questions coupled with a good deal of listening. Editing for this article was kept to a minimum.

Drs. Sejnowski and Lisberger describe a collaboration that is becoming a hallmark for ONR funding in the area of biological neural networks. Professor Lisberger is an acknowledged expert in visual motor control. Simply stated, Dr. Lisberger wants to answer such questions as: How does a primate stabilize the image on his retina while the head and body are moving? Solving this problem might have important implications for solving similar problems in robots and autonomous vehicles. Dr. Sejnowski’s demonstrations of neural network applications (such as his NETALK demonstration) have greatly contributed to the popularity of neural network approaches in solving real-world problems. The strong acceptance of the idea that biological systems contain design principles that can be applied to novel solutions for intractable problems underlies our ONR science support strategy.

Dr. Allen Selverston’s work centers on modeling a reprogrammable central pattern generating network. One unusual characteristic of this network is that it exists as the stomatogastric ganglion in the spiny lobster. Dr. Selverston provides good reductionistic arguments for choosing to study and understand a simpler, invertebrate system. The ganglion provides a flexible neural network with outputs (e.g. spatio-temporal patterns) that can be tuned or configured differently. Some of the reconfiguration occurs via a modulation or biasing effect by chemicals known as neuromodulators. Once again, the ultimate applications of Dr. Selverston’s work will probably be in better artificial networks and motor control devices.

Dr. James Bower, a recently tenured faculty member at Caltech, is a pioneer in biological neural network research. His work has centered on the rodent olfactory system as a model, a choice that has been shared by other major investigators. Dr. Bower’s work shares the view with the other authors in this issue that even the most detailed experimental investigations on elements in a biological system will not permit an understanding of how the elements work together as a functional system. However, computational neuroscience offers a strategy by which the elements can be combined to postulate critical properties that have not yet been experimentally observed, generating predictions that can be empirically validated, leading to more sophisticated models that generate further predictions, and so on, until the fundamental operation of the system is understood.

The last article by ONR Scientific Officer, Dr. Thomas McKenna provides an interesting contrast to a Naval Research Reviews article by Dr. Alan Meyrowitz (43(2): 1991) printed last year. Dr. McKenna provides a firm basis for the strategy adopted by the ONR Biological Intelligence Group which contrasts with the classical artificial intelligence (AI) approach of the traditional computer scientist. Dr. McKenna highlights some recent successes and makes a strong argument for the role of interdisciplinary efforts in devising intelligent, autonomous, and adaptive systems.

The 1990’s have been decreed by Congress and the President to be the Decade of the Brain. Understanding and emulating the functional capabilities of brains of all kinds remains both a fascinating intellectual and scientific enterprise with implications for the Navy and society in general. I hope this issue of Naval Research Reviews conveys some of the current excitement in the field.

Neural Systems for Eye Tracking

*Terrence J. Sejnowski
Howard Hughes Medical Institute
The Salk Institute and University of California at San Diego*

*Stephen G. Lisberger
University of California at San Francisco*

Abstract

We are just beginning to understand the principles of neural computation that make brains such exceptionally effective computational devices. Some of these principles are outlined here and illustrated with a network model of the primate eye movement systems for retinal stabilization and eye tracking.

Introduction

The brain of a house fly has less than 1 million neurons, about the same as the number of gates on a large VLSI chip. This brain is tiny – less than a mm³ in volume – but flies are among the most successful species on earth (Figure 1). Flies have good visual systems; they can navigate; they display complex courtship rituals; they are even capable of some learning¹. The power dissipated by the fly brain is less than 10 microwatts, making it an exceptionally efficient computational device, well beyond the capabilities of current technology.

Compare, for example, the operating characteristics of a fly brain with the largest and most powerful digital computers in existence, which are approaching peak speeds of 10 gigaflops (billion floating point operations per second). A

supercomputer is about the size of an automobile, costs tens of millions of dollars, consumes hundreds of kilowatts of power, is constantly attended by a team of hardware and systems specialists, and devours a steady stream of programs – an enormous human sacrifice. Supercomputers cannot see, have not been known to fly, and despite being networked have not yet mated or reproduced. They can, however, simulate wind tunnels, the weather, and even neurons, the building blocks of brains.

This dramatic comparison points out that we are missing something important. The approach taken by researchers in artificial intelligence to the problems faced by a fly is to write clever programs for the supercomputer. However, programs by themselves are not enough – the raw computational power required to solve the problem must also be available and be cost effective in terms of power dissipation and physical size. Even if such a “fly program” could be written, the available hardware for running it would be too bulky to get off the ground. Here we may have something to learn from the fly. This is the starting point for the field of neural computation, the goal of which is to understand the computational principles that are used in neural systems, which include artificial as well as biological neural systems. Some of the principles that have

emerged are summarized here. These principles are illustrated with an example of a brain system that has been worked out in great detail: eye tracking of moving targets.

Basic Principles

The memory and the processor are physically separated in the von Neumann architecture commonly used in digital computers; this separation causes a bottleneck in the flow of information between processor and memory. In contrast, many brain circuits work together in parallel to accomplish a task. The sensory representations, the neural processing, and learning mechanisms in brains are all tightly interrelated and directly tied to the neural "hardware". These tight relationships, necessary for computational efficiency, make it feasible to reverse engineer a creature's brain to understand how a complex behavior is possible².

Natural selection has favored biological systems that are fast and reliable, but not necessarily optimal. When threatened by a swat, a fly does not wait around to compute the exact trajectory of the looming hand, but immediately reacts with a

backward flip. This reaction is so fast that it defeats many predators, including humans who do not aim behind the fly. Even though a particular algorithm used by the fly is not optimal, it may be optimally matched to the neural architectures that are available, at many levels of organization³. Figure 2 summarizes the spatial structures that are found in brains, ranging from the molecular level to the systems level. Computational principles can be found at each of these levels. At the systems level, the entire brain should not be considered a single processor but a collection of special-purpose processing devices. Systems of neural networks in a fly brain are dedicated to different functions, such as visual processing, flight control, and special sensory-motor routines needed for feeding and mating. Within each such system, special-purpose neurons specialize for particular tasks, such as the giant motion-sensitive neurons that are used for orienting the fly during flight⁴. Each of these systems is quite complex; not one of them is within the reach of current supercomputers.

A principle of organization found in brains of nearly all species, from flies to humans, is the topographic map. The surface of the body, including the retina and the cochlea, is represented in the brain by sheets of neurons organized in isomorphic registry with the sensory surface. There are often distortions in a map, such as the greater representation of the foveal region of the retina compared to the periphery. Mappings preserve neighborhood relationships: nearby points in the visual field, for example, are mapped to nearby neurons in the visual cortex. This arrangement minimizes the wiring that is needed to connect neurons that must compute spatial and temporal differences. Wires are also expensive on VLSI chips that embody artificial neural networks, such as the artificial retinas and cochleas built by Carver Mead⁵. There are many maps in the brain. In the visual cortex alone over 24 distinct maps of the visual field have been identified, each one computing different aspects of the visual world (Figure 3).

The level of investigation in the nervous system that has received the most attention recently is the network level. Neural networks have become a hot area of research in part because of the promise that they might be the key to tapping the computational power of brains. However, we have already seen that there is a lot more going on in brains at other levels of investigation that is also important. Nonetheless, progress at the network level is essential. In mammalian cerebral cortex there are about 100,000 neurons in a mm^3 , each connected to thousands of other nearby neurons. This is a high degree of connectivity compared to a gate on a VLSI chip, which is typically connected to only a few others. For the brain, however, the connectivity is actually quite local: a neuron is connected to only a few percent of its neighbors, and to very few neurons that are more than 1 mm away. The current generation of neural network models only has a few thousand simplified model neurons and barely achieves the level of interconnection found in a mm^3 of cortex.

Figure 1.

*A head on view of the blowfly *Calliphora erythrocephala* that emphasizes the large compound eyes. Each contains around 5000 ommatidia, which provide the fly with all its information about the visual world⁴.*



The technology used to construct neurons and synapses is ultraminiaturized. Neurons are surrounded by a lipid membrane 10 nm thick that maintains differences in ionic concentrations and an electrical potential across the membrane. Proteins in the membrane selectively gate ionic currents according to the membrane potential. Interestingly, the biophysics that governs ionic currents in neurons is similar to the physics that governs the movement of charge carriers in semiconductors⁵. Thus, the computational primitives that are used by neurons to represent and combine information are not that different from the primitives that are available to chip designers working in the analog domain of processing. The main difference is in the efficiency of processing: a neuron dissipates around 10^{-15} joules/operation, around 10^{-8} less than the energy per operation using the current generation of silicon technology. This is why flies are so small and efficient compared with supercomputers.


Neurons cannot represent absolute values of sensory data with high accuracy because of their limited dynamic range – firing rates are typically from 1 to 100/second. Furthermore, statistical variability of the spike arrival time requires either time averaging or a population averaging to achieve even one significant figure of accuracy. This limitation favors the representation of differences, rather than absolute levels. Thus, visual neurons are more sensitive to contrast (spatial difference) and movement (temporal differences) than to absolute intensity levels, and large DC offsets are adapted out. The same principles are also used for designing analog VLSI circuits that mimic neural systems where there are variations in the biases of the transistors⁵. An example of how adaptation is used to calibrate the performance of eye tracking is given in the next section.

A Dynamical Network Model of Eye Tracking:

Time is an essential part of computing. There must be a match between the flow of information into a computing system and the time required for the computation. In contrast, the first generation of feedforward neural network models were designed for static pattern recognition. In a feedforward network, the sensory input, such as an image, is presented to the input layer and the output appears after a few steps of processing (Figure 4). Furthermore, simulations on digital machines ran very slowly. A new generation of dynamical neural network models have been developed that take time into account through the intrinsic temporal properties of the model neurons and feedback connection within the network (Figure 4). VLSI chips have been developed that greatly speed up the processing so that large networks can be operated in real time. In this section we summarize a dynamical network model that was constructed to perform target tracking. It is

Figure 2.

Schematic illustration of levels of organization in the nervous system. The spatial scales (left) at which anatomical organizations can be identified varies over many orders of magnitude. Icons (right) represent structures at distinct levels: (top) a subset of visual areas in visual cortex; (middle) a network model of how ganglion cells could be connected to simple cells in visual cortex, and (bottom) a chemical synapse³.




based on the actual circuitry found in primate oculomotor system and is an illustration of the general principles outlined above.

In primates, the oculomotor system is able to keep the eyes pointed at a target in spite of movements either of the head or of the target. Research on awake, behaving monkeys has led to an understanding of the basic design principles of the oculomotor tracking system. First, different tasks are accomplished by separate systems that rely on different sensory inputs: tracking is maintained during head turns through the use of the “vestibulo-ocular reflex” or “VOR”, which uses vestibular inputs that report angular head motion; tracking of target motion is accomplished through “pursuit” eye movement, which uses visual inputs that report the direction and speed of target motion with respect to the eye. Second, the neural networks for these separate oculomotor subsystems are largely separate in the brain. In the case of the VOR and pursuit, they interact only in certain premotor structures and, of course, in the motoneurons that control the contraction of

Figure 3.

An illustration of the principal connections between the lateral geniculate nucleus (LGN), striate cortex and extrastriate visual areas in the primate. Connections are shown for one eye only. Frontmost block is a portion of V1 containing half of a left eye column and half of the neighboring right eye column. Many connections and possible connections are not shown⁷.



the eye muscles. Finally, calibration of each oculomotor subsystem is maintained through a long-term adaptive process that senses tracking errors and adjusts the performance of the relevant neural networks so that subsequent performance is improved. In the cases of the VOR and pursuit, errors result in image motion across the retina and are reported to the oculomotor system by a part of the visual system that is specialized for the analysis of moving images.

The neural networks that mediate smooth tracking in the primate brain are found in the brain stem and cerebellum. Analysis of the connection and function of neurons in these networks has indicated that these networks are organized in the way diagrammed in Figure 5. Briefly, the VOR is mediated primarily by direct pathways that travel in the brain stem and transmit vestibular inputs to the extraocular muscles with very short delays of less than 15 ms. Pursuit is mediated by visual inputs that arise in the cerebral cortex and are relayed to the extraocular muscles through a part of the cerebellum called the flocculus. The neural networks for the VOR and pursuit are largely separate and share only the flocculus and some of the premotor neurons in the brain stem.

Our computational research has concerned two aspects of the tracking system. First, we have investigated the function of each of the pathways in Figure 5 by analyzing a computational model that had a homologous anatomical structure. The model consisted of a collection of artificial processing units

that were interconnected in a specified pattern, with the strength of the interconnections designated as weights. Each unit formed its output at time t as a non-linear, sigmoid function of the sum of the value of each of its inputs times the weight of the connection from that input.

As shown in Figure 5, the connections among units could include both forward and backward connections. Like the brain, the model in Figure 5 has highly selective and specific connections and is not fully recurrent in the way that an arbitrary neural network might be. We presented the model with a series of tracking tasks that required accurate tracking in each millisecond over a period of one second. The tasks were selected to represent the full range of performance of the primate tracking system: head moving and target stationary; target moving and head stationary; and head and target moving together. We used a gradient descent optimization procedure called "backpropagation" to adjust the parameters of the model (the weights of the connections between units and the time constants of integration of the units) until the difference between eye velocity and target velocity was minimized.

Allowing a powerful optimization procedure to adjust the internal parameters of the model until it performed its given tasks well, we could explore how different configurations of the model were able to solve the task. We discovered that the model achieved the same excellent performance as the primate oculomotor system as long as all the pathways known to exist in the brain were included in the model. However, all the pathways shown in Figure 5 were necessary: elimination of any pathway prevented the optimization algorithm from finding a set of parameters that could produce stable and accurate performance.

This part of our analysis highlighted several important points. First, it revealed that the standard methods of neural network optimization can be used successfully in a system, like the pursuit system, where the output from the model changes its input. Second, it demonstrated the critical importance of the dynamic features of the task in determining the structure and the performance of the model. Third, it emphasized the elegance of the organization of the brain for the task of tracking and raised the possibility that the design principles of the brain's oculomotor tracking system might be helpful in the design of artificial tracking systems.

The second part of our analysis concerned the way in which the biological and artificial tracking models could calibrate their own performance. In monkeys, the VOR normally generates eye rotation that is equal in amplitude and opposite in direction to a head turn: the gain of the system (defined as eye speed divided by head speed) is about 1.0, even during head turns in darkness. If a monkey wears spectacles that miniaturize or magnify the visual scene, the gain of the VOR is gradually adjusted so that it reaches values as low as 0.25 or as high as 1.8. We think of this adjustment as a simple form of learning and memory that is designed to minimize the motion of retinal images during head turns. Research in one of our

laboratories⁶ has provided strong evidence that the site of this adjustment is in the brain stem at the synapses marked with a "B" in Figure 5.

We next asked about this site(s) of learning in the model illustrated in Figure 5 and investigated why it selected those sites. Starting with the model that performed well on the basic tracking tasks of pursuit and the VOR, we reactivated the optimization procedure under conditions that required the VOR to be 0.25 times its normal amplitude and pursuit to remain normal. The model reduced the amplitude of the VOR by adjusting the weights at two sites, labeled "B" and "H" in Figure 5. Restricting the sites at which weights could be adjusted revealed that this combination of sites of learning was both necessary and sufficient to allow the model to adjust the amplitude of the VOR. Three factors were responsible for constraining the possible sites of learning in this way: 1) the structure of the model, initially dictated by the need for the system to perform well on both the VOR and pursuit; 2) the need to continue excellent performance on pursuit as the amplitude of the VOR was adjusted; 3) the demand of accurate performance throughout the one second duration of each training trial. This is a large advance over static models which placed very few constraints on the site of learning in a neural network like that found in the brain's eye tracking system.


The exact changes found at sites "B" and "H" in the model helped to explain a paradoxical finding in the literature on learning in the VOR in monkeys. In the model, learning of a smaller amplitude VOR was associated with decreases in the connection weights at both "B" and "H". The decrease at "B" is in the correct direction to cause a decrease in the VOR, but the decrease at "H" is in the wrong direction and should have caused an increase in the amplitude of the VOR. Similar changes have been found at the sites that correspond to "B" and "H" in the brain. The analysis of the model helps us to understand that this combination of changes in the brain may be necessary: changes at connection "B" to reduce the amplitude of the VOR and changes at "H" to maintain dynamic stability in both the VOR and pursuit.

Conclusions

Our simplified model of the oculomotor system has served us well in providing answers to the questions that we asked, but it has one major shortcoming that would prevent us from building a working model. The visual processing in the model sidestepped the difficult job of segmenting the moving object from the image and dealing with noisy sensory data. We are currently developing network models of visual motion areas MT and MST in primates (see Figure 3) that could be used for constructing a complete, working model of smooth eye tracking. However, the visual processing required is highly computation intensive and special purpose hardware would be needed to get the model to run in real time. We

Figure 4.

Three types of networks. (A) Feedforward network with one layer of weights connecting the input units to the output units. (B) Feedforward network with two layers of weights and one layer of hidden units between the inputs and the output units. (C) Recurrent network with reciprocal connections between units.



anticipate that analog VLSI chips will soon be available for performing the motion processing⁵. Building a working system is the ultimate test of our model, a goal that we hope to achieve with further support from Joel Davis.


Eye tracking is only one of the brain's many systems. Several other brain systems are about as well worked out and are ripe for modeling, including the sound localization system in the barn owl, and the jamming-avoidance system of electric fish². Collaborations between neurophysiologists and neural modelers, such as the one summarized here, could lead to a better understanding of these systems. The lessons learned could also be applied to solving difficult problems encountered by the Navy.

Followup

The field of neural computation is still in its infancy. A large number of able researchers from other disciplines have entered the field, which continues to expand. New journals have been founded, such as *Neural Computation* (MIT Press) and *Network: Computation in Neural Systems* (IOP Publishing), and large international meetings are held regularly. Dr. Patricia Churchland and one of us (TJS) have written a book on computational neuroscience called *Computational Brain* (MIT Press, 1992).

Figure 5.

Diagram of the basic oculomotor circuits used for the VOR and smooth pursuit (top) and the architecture for a dynamical network model of the system (bottom). The vestibular system senses head velocity and provides inputs to neurons in the brainstem (B) and the cerebellar flocculus (H). The output neurons of the flocculus make inhibitory synapses on neurons in the brainstem. The output of the brainstem projects directly to motoneurons, which drive the eye muscles, and also project back to the flocculus (E). This positive feedback loop has an "efference copy" of the motor commands. The flocculus also receives visual information from the retina, delayed for 100 milliseconds by the visual processing in the retina and visual cortex. The image velocity, which is the difference between the target velocity and the eye velocity, is used in a negative feedback loop to maintain smooth tracking of moving targets. The model (below) incorporates model neurons that have integration time constants and time delays. The arrows indicate connections between the populations of units in each cluster, indicated by small circles. The dotted line between the eye velocity and the summing junction with the target velocity indicates a linkage between the output of the model and the input⁵.



Biography


Terrence J. Sejnowski is an Investigator with the Howard Hughes Medical Institute and a Professor at The Salk Institute for Biological Studies where he directs the Computational Neurobiology Laboratory. He is also Professor of Biology and Adjunct Professor in the Departments of Physics, Neuroscience, Psychology, Cognitive Science, Electrical and Computer Engineering, and Computer Science and Engineering at the University of California, San Diego, where he is Director of the Institute for Neural Computation and the McDonnell-Pew Center for Cognitive Neuroscience. His main research interests are to understand how brains work and to develop new computer architectures based on the principles of neural computation.

Stephen G. Lisberger is a Professor of Physiology at the University of California at San Francisco and the director of the W. M. Keck Foundation Center for Integrative Neuroscience. His main research interests are to understand visual motor control and simple forms of learning through studies of the eye movements in primates. He combines experimental approaches to these problems with techniques from control theory and network modeling.

References

1. Tully, T., Boynton, S., Brandes, C., Dura, J. M., Mihalek, R., Preat, T. and Villella, A. Genetic dissection of memory formation in *Drosophila melanogaster*, Cold Spring Harbor Symposia on Quantitative Biology **55**, 203-211 (1991).
2. Konishi, M. Deciphering the Brain's Codes, Neural Computation **3**, 1-18 (1991).
3. Churchland, P. S. and Sejnowski, T. J. Perspectives on Cognitive Neuroscience, **242**, 741-745 (1988).
4. Bialek, A., Rieke, F., de Ruyter van Steveninck, R. R. and Warland, D. Reading a Neural Code, Science **252**, 1854-1857 (1991).
5. Mead, C. *Analog VLSI and Neural Systems* (Addison-Wesley, Reading, MA, 1989).
6. Lisberger, S. G. The neural basis for learning of simple motor skills, Science **242**, 728-735 (1988).
7. LeVay, S. and Nelson, S. B. in *The Neural Basis of Visual Function* (eds. J. R. Cronly-Dillon) (Macmillan Press, London, 1991).

Profiles in Science



Carver A. Meade

Carver A. Meade is the Gordon and Betty Moore Professor of Computer Science at the California Institute of Technology where he has taught for 22 years and has been a principal investigator for the Office of Naval Research for many years. Having pioneered in the fields of solid-state electronics and management of complexity in the design of very large scale integrated circuits, he has been the leading force in the development of a design methodology for VLSI. He has written with Lynn Conway, the standard text for VLSI design, *Introduction to VLSI Systems*. His recent work is

concerned with modeling neuronal structures, such as the retina and the cochlea using analog VLSI systems. His new book on this topic, *Analog VLSI and Neural Systems* was published in 1990.

Professor Meade has worked to bring about a general awareness of computation as a physical process, rather than purely a mathematical one. He has strongly advocated the importance of unifying technology and architecture into a single discipline and has shown the importance of this unity for the future of the field at large.

Exploring Biological Neural Networks Using Realistic Computer Simulations

*Dr. James M. Bower
California Institute of Technology*

Introduction

Over the last several years, there has been a dramatic increase in interest in taking a more computational approach to understanding the functional organization of biological nervous systems. This effort involves both the development of new experimental techniques to sample the complex patterns of activity generated by the nervous system (Bower et al., 1988) and the development of a more model-based theoretical approach to understanding the significance of neurobiological data (Koch and Segev, 1989; Zometzer et al., 1990). A large and growing component of the later effort involves the development of numerical computer simulations modeled on real brain structures. These so called realistic or structural computer simulations are based on the actual physiology and anatomy of real brain neurons and circuits (Figure 1) and generate simulated neural signals as outputs (Figure 2; Bower 1990). Especially when linked to the experimental investigations of neuronal behavior, computer modeling of this sort has the potential to provide powerful new insights into how brain structure reflects brain function (Bower, 1991; Hasselmo et al, 1991).

The growing interest in computational neuroscience and computer modeling of biological systems is driven by several interrelated factors. First, from the point of view of traditional neurobiology, modern experimental techniques are producing ever increasing quantities of physiological and anatomical information, yet it is ever more apparent that this data by itself does not necessarily reveal nervous system function. Perhaps this is best seen with respect to the relatively small oscillatory neural circuits that are known to control rhythmic movements in invertebrates (Selverston, 1985). The physiology and anatomy of these so called central pattern generators have been studied extensively for the last twenty five years, and a great deal of structural information has been obtained, yet it is now generally recognized that without models, the complex behavior of even these relatively simple networks will not be understood (Kleinfeld and Sompolinsky; 1988; 1989; Getting, 1989; Ryckebusch et al., 1989).

A second factor in the growing emphasis on model-based approaches to understanding nervous system function is the increasingly interdisciplinary interest in this question. The dramatic growth of research in so called neural networks within the engineering community, for example, has estab-

lished a new climate for collaboration between engineers and biologists (Lippman et al., 1991). While much of the engineering side of neural networks remains a purely engineering effort with little obvious relevance to neurobiology, a growing number of engineering students are becoming interested in the nervous system for its own sake. Similarly, students in physics and computer science are also increasingly interested in studying the complexities of nervous system function. At Caltech, this trend led us several years ago to establish a new graduate program in Computation and Neural Systems (CNS) which specifically cross-trains students in neurobiology, engineering, physics, applied mathematics, and computer science. The continuing infusion into neurobiology of students from these more quantitative disciplines will inevitably result in a rapid expansion of modeling efforts.


Finally, a third factor in the increasing development of realistic computer simulations is the actual availability of computers powerful enough to model the complexities of the nervous system. The rapid growth of the computational power of graphics workstations, in particular, has allowed neurobiologists to begin building sophisticated biologically realistic simulations in their own laboratories. For example, seven years ago, our own first simulations of the olfactory cerebral cortex took 30 minutes to generate 100 msec of data on an IBM AT modeling only 70 extremely simple neurons with a total of 5,000 interconnections (Wilson and Bower, 1986). Today, the same 30 minutes of cpu time on a SUN SPARCstation II generates 500 msec of activity in 5,000 much more complex neurons with a total of 750,000 simulated connections (Wilson and Bower, 1989). Presently, our laboratory is now using the new Intel Touchstone Delta concurrent supercomputer which is based on 512 workstation class computing nodes. Because realistic brain models quite naturally map onto concurrent computers (Nelson and Bower, 1989; 1990), this machine will allow a dramatic increase in the size and complexity of our realistic models.

Why Realistic Simulations

The primary emphasis in this article, and in our own research, is on building realistic numerical simulations of biological systems (Bower, 1990). Inevitably, such simulations are themselves quite complex and depend on the values of many detailed parameters. Accordingly, the question should, and has, been raised, as to what benefit such complex models have in understanding brain function. For example, it has been argued that a complex simulation of a biological network may, in principle, be no more understandable than the complex network itself. In our experience this is incorrect for several reasons. First, this argument assumes that the simulation already exists to begin with. In fact, building complicated biologically realistic simulations is a process that begins with the initial construction of simplified, though realistically ex-

Figure 1.

This figure schematically represents the structure of our network simulation of the olfactory cortex of the rat. A shows the spatial pattern of inputs to this structure. Axons from the olfactory bulb enter from one corner of the cortex (the lower left in this case) and then project across the cortex as shown. B shows the organization of excitatory neuronal connections within the cortex itself while C and D illustrate the pattern of inhibitory neuron connections. In the simulation, each of the components shown is represented by a mathematical expression which simulates its behavior in the real circuit. Data can also be obtained from each of these components for comparison with experimental results. Further information about the model can be found in Bower, 1990, from which this figure was used by permission.



pandable models, and then adds complexity as necessary (Bower, 1990). It turns out that this process is itself a large part of the reason to build models of this type. For example, through model manipulations, it is possible to quantify what is actually known about a particular circuit which often leads to a much better understanding of what is not known. This in turn suggests ideas for additional experiments as well as providing a context in which to interpret the data once it is obtained (Hasselmo et al., 1991). In other words, the value of the model is not so much in its eventual capabilities, but in the model building process itself. In effect then, simulations used in this way serve as tools to explore the functional significance of the components of real neural networks. Obviously, this requires that the models be realistic to begin with.

A second reason that a complex model is different from a complex biological network has to do with the ease with which each can be studied. While neurobiologists now have at their disposal many powerful experimental techniques, experimental biology is still very hard and time consuming. Simulated experiments in a computer model, however, are tremendously easier. For example, it is truly possible in a computer to “noninvasively” record neural activity, or to generate extremely well defined neuronal lesions. Of course, all such *in computo* experimental results must inevitably be checked with real biological experiments. However, when models are realistic they actually can be checked while also improving the design of biological experiments.

Strong interaction between a realistic computer simulation and experimental biology also increases the likelihood that one will discover features of the biological system that were not previously recognized. In our view this is the ultimate measure of the value of a model. Thus, for any modeling effort it is important to ask what one knows after building the model that was not known before building it. Realistic simulations, in our experience, are much more likely to meet this expectation than are more abstract models which are usually explicitly designed to demonstrate the feasibility of a preconceived functional idea. Of course, sometimes even seemingly realistic models are built for this purpose (c.f. Granger et al., 1988), although models of this type also often explicitly omit well


known, but inconvenient biological features. This is not to say that both realistic or abstract demonstration type models can not be valuable when exploring interesting ideas. We also use sometimes more abstract versions of our realistic simulations to explore emerging functional ideas (see Hasselmo et al., 1991). However, the process of building this sort of model is much less likely to result in the generation of new ideas about the structure and function of the network being considered. This is why the base of our efforts to understand neural systems remains the interaction of realistic simulations with experimental investigations.

Technical Challenges For Realistic Modeling

The increasing use of realistic models within neurobiology poses a number of interesting technical challenges and opportunities. First, from the point of view of programing, one characteristic of this type of modeling is that parameters and even model components are regularly changed or added. This means that simulation code must be easily mutable without the need to completely reorganize the structure of the software each time a change is made. Second, it is important for these modeling efforts to be in a position to take advantage of advancements in numerical methods which can dramatically

Figure 2.

This figure demonstrates the evolving pattern of neuronal activity in a simulation of olfactory cortex. The trace at the top is a simulated electroencephalographic (EEG) record from the model which resembles quite closely activity seen in the real brain. The panels at the bottom of the page represent the patterns of simulated activity in different model components at times corresponding to the EEG trace as shown. Each individual framed area shows the spatial distribution of activity across the full extent of the simulated cortex and each row of panels shows activity in one of the three types of simulated neurons. The top and bottom rows show the strength of inhibitory neuron effects at different positions within the cortex (darker areas have stronger effects) while the middle row shows which of the cortical output neurons are active. Further explanation of this figure can be found in Bower, 1990, from which it is reprinted with permission.



speed up complex simulations (Hines, 1989). Third, as realistic simulations become more complex, advancement in the field will also continue to depend on access to state of the art computer resources. Ideally then, the software base for these simulations should be largely machine independent. Fourth, realistic simulations, like the brain itself are made from very similar components linked together in different ways depending on the particular structure being simulated. This means that many different modelers can in principle use the same basic structural components in building their models. As a result, the potential exists for a considerable overlap in effort, or alternatively, for considerable cooperation. Fifth, overlapping needs also exist with respect to the analysis and display of modeling results. Given the complexity of the models, sophisticated graphical displays are critical to evaluating model behavior, and many of the same displays can be used for different models. Again this common need will result in either effort duplication or cooperation. Sixth, realistic computer simulations also potentially provide a new means for exchanging information about biological systems. In this sense, a structurally realistic simulation can also be viewed as a form of data base, which in addition to encoding information about the network being simulated, can also generate new information. This will only happen, however, if a common language develops for simulations. Finally, as discussed above, structural simulations are most effective when used hand in hand with experimental science. This means that experimental neurobiologists themselves should be in a position to build and manipulate models of their systems. While neurobiologists are adept at using experimentally sophisticated techniques, as a whole they are traditionally not experts in computer programming. Thus, ideally there should be some means to allow biologists access to simulation technology without necessarily requiring advanced programming skills.

The GENESIS Simulation System

For each of these reasons, we have spent the last six years developing a modeling environment for the realistic simulation of biological systems (Wilson et al., 1989). This modeling software called GENESIS for "General Neural Simulation System" was developed to support the simulation of neural systems ranging from complex models of single neurons to simulations of large networks made up of more abstract neuronal components. In our own lab we have used this system to model biological structures ranging from parts of neurons (Shankle et al., 1989; Lee and Bower, 1990; De Schutter and Bower, 1991), to whole neurons (Lee and Bower, 1990; De Schutter et al 1991), to small networks of neurons (Ryckebusch et al., 1989), to large scale realistic network simulations (Wilson and Bower, 1988; 1989; 1990; Bhalla et

al., 1989) to much more abstract models (Hasselmo et al., 1991). GENESIS has also provided the basis for laboratory courses in neural simulation techniques at the Marine Biological Laboratory in Woods Hole, MA (Supported by ONR) and is currently being used for instruction at several Universities around the country including Caltech.

To meet the requirement of machine independence discussed above, GENESIS and its graphical front-end XODUS are written in C and run on numerous graphics workstations under UNIX (OS 4.0 and up), and X-windows (version 11.3 and 11.4). The system has also been ported to concurrent computers including both workstation-based transputer arrays and the Intel Touchstone Delta mentioned previously. From a users point of view, the system allows modelers to build both models and graphical interfaces using a high level scripting language that takes advantage of libraries of model components and graphical objects used in previous simulations. At run time, GENESIS is highly interactive allowing all model parameters and features to be altered. The libraries themselves are user extensible which means that new components are continually being added. Both the educational and research use of the system is supported by the GENESIS users group (BABEL) which is based at Caltech. For those interested, full source for the simulator as well as four tutorial simulations and three papers that describe organization of the simulator in more detail are available via FTP from genesis.cns.caltech.edu (internet number 131.215.135.64). Any other questions about the system or its distribution can be sent to genesis@caltech.bitnet.

Conclusions

Realistic simulations of neural systems are becoming an ever more important tool in ongoing efforts to understand the functional organization of the nervous system. However, it is important to realize that, even with sophisticated computers and computer software, the challenges involved in this task are enormous. Arguably, few if any, physical systems are as structurally complex as the nervous system. As a result, even todays relatively powerful machines can be quickly overburdened as models become more and more realistic. For example, we have estimated that a modern graphics workstation based simulation of the 6 million cerebellar Purkinje cells found in the rat cerebellar cortex with each cell represented by a modest 1000 compartments, would still take 5700 years to generate 100 msec of simulated activity (Nelson et al., 1989; Wilson and Bower, 1989). In addition, even reasonably sized simulations often generate massive amounts of data containing complex spatial and temporal information for which analysis tools are not yet available. Further, to be useful this data would ideally need to be compared to data obtained from recordings of multiple single real neurons, yet biological recording techniques are currently limited to monitoring at most a few

neurons simultaneously. While we and others are working on silicon-based recording techniques that should allow tens or even a few hundred single neurons to be recorded at the same time (Bower et al., 1988), major developments in this area of experimental neurobiology are necessary. The growing use of modeling techniques to study patterns of activity in neural networks is likely to spur the development of these new experimental approaches which in turn will increase reliance on realistic computer simulations.

Biography

Dr. James M. Bower is an Associate Professor of Biology associated with the graduate program in Computation and Neural Systems (CNS) at the California Institute of Technology. His publications and research primarily concern the functional organization of the mammalian olfactory system and the mammalian cerebellum. He was a charter member of the organizing committee for the annual Neural Information Processing Systems Meeting in Denver Colorado and is Program Chairman for the First Annual Computation and Neural Systems meeting to be held in San Francisco this coming July.

References

1. Bhalla, U.S., Wilson, M.A. & Bower, J.M. 1988 Integration of computer simulations and multi-unit recording in the rat olfactory system. *Soc. Neurosci. Abstr.* 14,1188.
2. Bower, J.M. 1991 Piriform cortex and olfactory object recognition. In: J. Davis and H. Eichenbaum (Eds.) *Olfaction as a Model System for Computational Neuroscience*, MIT Press, Cambridge, MA.
3. Bower, J.M. 1990 Reverse engineering the nervous system: An anatomical, physiological, and computer based approach. In: *An Introduction to Neural and Electronic Networks*. S. Zornetzer, J. Davis, and C. Lau, editors. Academic Press, New York, NY pp. 3-24.
4. Bower, J.M., Wilson, M., Banik, J., Nelson, M. and Rasnow, B. 1988 New technical approaches to monitoring and interpreting the dynamics of real neural networks. *IEEE Eng. Med. Biol. Soc. Proceedings*. 20: 1921-1922.
5. De Schutter, E., Gundappa-Sulur, S., and Bower, J.M. 1991 A computer simulation of plateau potentials and synaptic interactions on Purkinje cell spiny dendrites. *Soc., Neurosci.* 17:1383. *European Neuroscience Association Abstr.* (in press).
6. Getting, P.A. 1989 Reconstruction of small neural networks. In: *Methods in Neuronal Modeling: From Synapses to Networks*. C. Koch and I. Segev, editors. MIT Press, Cambridge, MA., pp 397-438.
7. Granger, R., Ambros-Ingerson, J., Henry, H., and Lynch, G. (1988). Partitioning of sensory data by a cortical network. In *Neural Information Processing Systems* (D.Z. Anderson, ed.) pp. 317-337. American Institute of Physics, New York.
8. Hasselmo, M.E., Wilson, M.A., Anderson, B., and Bower, J.M. 1990 Associative function in piriform (olfactory) cortex: Computational modeling and neuropharmacology. In: *The Brain, Cold Spring Harbor Symposium*. Cold Spring Harbor Press, Cold Spring Harbor, New York, pp. 599-610.
9. Hines, M. 1989 A program for simulation of nerve equations with branching geometries. *Int J. Biomed. Comput.* 24: 55-68.
10. Kleinfeld, D. and Sompolinsky, H. (1988). Associative neural network model for the generation of temporal patterns: Theory and application to central pattern generators. *Biophysical J.* 54, 1039-1051.
11. Kleinfeld, D. and Sompolinsky, H. (1989). Associative network models for central pattern generators. In "Methods in Neuronal Modeling: From Synapses to Networks" (C. Koch and I. Segev, eds.), MIT Press, Cambridge, Massachusetts 195-246.
12. Koch, C., and Segev, I. (editors) *Methods in Neuronal Modeling: From Synapses to Networks*. MIT Press, Cambridge, MA. (1989).
13. Lee, M. and Bower, J.M. 1990 A computer modeling approach to understanding the inferior olive and its relationship to the cerebellar cortex in rats. *Advances in Neural information processing systems*. Vol. 2, D. Touretzky, editor. Morgan Kaufmann, San Mateo, CA pp. 3-24.
14. Lippman, R.P., Moody, J., and Touretzky, D. S. editors. *Advances in Neural information processing systems*. Vol. 3 Morgan Kaufmann, San Mateo, CA.
15. Nelson, M.E. and Bower, J.M. 1990 Brain maps and parallel computers. *Trends in Neuroscience* 13: 403-408.
16. Nelson, M., Furmanski, W. and Bower, J.M. 1989 Simulating neurons and neuronal networks on parallel computers. In: *Methods in Neuronal Modeling: From Synapses to Networks*. C. Koch and I. Segev, editors. MIT Press, Cambridge, MA., pp 397-438.
17. Rumelhart, D.E. and McClelland and the PDP Research Group. *Parallel Distributed Processing*. MIT Press, Cambridge, MA. (1986).
18. Ryckebusch, S., Mead, C., and Bower, J.M. 1989 Modeling small oscillating biological networks in VLSI. In: *Advances in Neural information processing systems*. D. Touretzky, editor. Morgan Kaufmann, San Mateo, CA. pp. 384-393.
19. Selverston, A. I. (1985). *Model Neural Networks and Behavior* Plenum Press, New York.
20. Shankle, W.R., Wilson, M., Stevens, D.R., Perkel, D., Bower, J., and C. W. Cotman. 1989 Computer simulation of excitatory postsynaptic currents due to varying ratios of NMDA to non-NMDA receptors. *Soc. Neurosci. Abstr.* 15, 948.

21. Wilson, M.A. and Bower, J.M. 1990 Computer simulation of oscillatory behavior in cerebral cortical networks. Advances in Neural information processing systems. Vol. 2, D. Touretzky, editor. Morgan Kaufmann, San Mateo, CA pp. 84-91.
22. Wilson, M.A., Bhalla, U.S., Uhley, J.D., and Bower, J.M. 1989 GENESIS: A system for simulating neural networks. In: Advances in Neural information processing systems. D. Touretzky, editor. Morgan Kaufmann, San Mateo, CA. 485-492.
23. Wilson, M. and Bower, J.M. 1989 The simulation of large-scale neuronal networks. In: Methods in Neuronal Modeling: From Synapses to Networks. C. Koch and I. Segev, editors. MIT Press, Cambridge, MA.. pp. 291-334.
24. Wilson, M. and Bower, J.M. 1988 A computer simulation of olfactory cortex with functional implications for storage and retrieval of olfactory information. Neural information processing systems. pp. 114-126 D. Anderson, editor. Published by AIP Press, New York, N.Y.
25. Wilson, M., Bower, J.M. and Haberly, L.B. 1986 Computer simulation of piriform cortex. Soc. Neurosci. Abstr. 12:1358.
26. Zornetzer, S.F., Davis, J.L., and Lau, C. (eds) 1990 An Introduction to Neural and Electronic Networks. Academic Press, New York, N.Y.

Chemical Control of Neural Networks

Allen I. Selverston
University of California, San Diego

Recent work on the lobster nervous system may tell us a lot about how the neural circuits in our own brain operate

The computational abilities of the brain lie in the enormous number of interconnections between its neurons. As a result of studies conducted over the last ten years, we now know that the strength of these interconnections, as well as the biophysical properties of the neurons which make them can be altered as a result of exposure to a particular class of chemical substances known as neuromodulators. Neuromodulators are produced naturally in the body and the results of their actions have been well documented. At one extreme we can describe the **behavioral** effects of neuromodulators. At the other, we are beginning to understand how neuromodulators work at the **cellular level** in terms of receptor binding, second messenger activation and phosphorylation of channels. But an enormous gap exists between these two levels of analysis largely due to the fact that for the mammalian brain there is a dearth of knowledge about how neurons are actually connected with each other to form circuits, how the synapses forming these circuits operate and how the neurons themselves operate as individual biophysical entities.

It does not seem possible to bridge the gap between these two levels of analysis in vertebrates because there are simply too many neurons and synapses involved and techniques to work out synaptic interactions between so many cells do not

exist. Invertebrates, however, have fewer, larger and reidentifiable neurons, making it possible to determine the detailed structure of local circuits electrophysiologically. In this paper I describe how one particular well-characterized neuronal circuit in the lobster can serve as a model for understanding what cellular mechanisms are involved in neuromodulation and how the modulatory effects permit the computation of new and different spatio-temporal output patterns. It will be shown that these new patterns can generate different behaviors which are of use to the animal.

An understanding of a small system such as this can have an important role in suggesting how more complex circuits such as those found in the brain may operate. But in addition to increasing our biological knowledge, the stomatogastric nervous system is a living neural network which we know works. This means that in theory the biological mechanisms can be translated to principles which can be "back-engineered" into artificial neural networks. Such artificial networks should eventually be able to replicate all of the functions of the biological system. If this goal were to be achieved, it would not only be of use in understanding how biological systems work, it would provide the basis for implementing the generation of artificial motor patterns. These patterns could be useful for robotics since they would allow a wide range of dynamic behaviors with very few "neurons."

The Stomatogastric Nervous System of Lobsters

Why work on lobsters? Quite simply, they offer a cluster of advantages which cannot be found in any other known preparation. Like many other invertebrates, the neurons are large and identifiable, meaning the same neuron can be found in different animals. The neurons also show the subthreshold activity (post-synaptic potentials) necessary for determining the circuit electrophysiologically (Figure 1). But the primary reason is that they provide models for understanding the general properties of neural networks and they have given us some insights as to how real biological neural circuits actually

operate. Such simple circuits are the only means presently available to bridge the huge gap between cellular and behavioral data. The heart of the stomatogastric circuit is the stomatogastric ganglion. Unlike the mammalian brain, the nervous system of the lobster, like other invertebrates, is distributed segmentally throughout the body in groups called ganglia. Each ganglion is a self-contained integrating system containing sensory inputs and motor outputs, and communicates with the rest of the CNS with connecting nerves.

The stomatogastric ganglion, when attached to two "higher ganglia," the commissurals, produces two highly stereotypical motor patterns. The production of such spatio-temporal patterns of nervous activity is a general feature of all nervous coding from the production of sensory inflow, to the

Figure 1.

Neural circuitry of the stomatogastric system in the spiny lobster. Most of the neurons are located in the 30-celled stomatogastric ganglion. Cell bodies for the IVN neurons are located at the base of the brain. GPR, AGR and PSR are all sensory neurons while APM, MPN and PS are identified neuromodulatory neurons.





Figure 2.

The gastric mill central pattern generator. *A*: Diagram of the nerves associated with the stomatogastric ganglion (STG). Inputs from the esophageal (EG) and circumesophageal (CG) ganglia enter the STG via the stomatogastric nerve (STN). The STN can be reversibly blocked with cold or with isotonic sucrose. *B*: Cartoon of the gastric mill circuit determined by paired intracellular recordings. *C*: Extracellular recordings from nerves carrying axons of gastric mill motor neurons during active cycling. The top six traces are motor neurons while the trace labelled R-SON shows E cell impulses. *D*: Simultaneous intracellular recordings from gastric mill neurons including Int 1, the sole interneuron in the CPG circuit.



central processing of that information, finally to the generation of appropriate patterns.

The patterns produced by the stomatogastric ganglion (STG) lead to the sequential contractions of striated muscles which operate teeth and filters within the lobster stomach. When the STG and the commissural ganglia are removed from the outer wall of the stomach and pinned out in a Petri dish filled with saline (Figure 2A), two cyclical rhythms can be produced for many hours. The STG is therefore a central pattern generator (CPG) which can operate free of sensory

input (which is left on the stomach wall), and free of all hormonal and chemical influences normally present in the animal. So, in essence, we have a small (30 cell) cluster of neurons, in operating condition with which to study the cellular basis for the production of spatio-temporal patterns. How do the cellular and synaptic properties of this small system of neurons interact cooperatively to produce these patterns and control a particular set of behaviors?

One of the two patterns produced by the STG is shown in Figure 2C. These are extracellular recordings made from the motor nerves just before they insert onto particular muscles so that their identity is assured. One sees bursts of impulses occurring at a frequency of about 15-20/min for the gastric rhythm. The cell bodies of the neurons can be identified by intracellular recording (Figure 2D) and pairwise recording has produced the circuit shown in Figure 2B.

The stomatogastric CPGs are controlled by higher centers which integrate sensory feedback with information from the brain. In addition, some 13 known neuromodulatory substances have been described in crustaceans by radioimmunoassay or immunohistochemistry (Table 1). When applied to the *in vitro* preparation these substances radically reorganize the burst patterns (Figure 3). The behavioral significance of the altered patterns can be determined by direct observation of the gastric mill in intact animals after the animal has been injected with a neuromodulator (Figure 4).


The stomatogastric system, therefore, can serve as a model with which to study the genesis and control of a specific spatio-temporally organized behavior. Our goal is to understand the mechanisms underlying these processes from the cellular to the behavioral level. How do neuroscientists know when "understanding" has been reached? One way is to replicate the system mathematically – i.e., build a model. Such a model should not only duplicate the normal behavior of the system but be able to generate predictions about the output under a variety of perturbations or conditions which can be tested experimentally.

CCK Neuromodulation of the Gastric Mill

As an example of how stomatogastric patterns can be influenced by neuromodulators, consider the effects of cholecystikinin (CCK) on the gastric mill. CCK is a mammalian gut peptide known to play a role in gastric motility and satiety. Antibodies to this small peptide have been made which are immunoreactive to a similar peptide in lobster. Using these antibodies we can show a CCK-like peptide is present in all cell bodies "upstream" from the stomatogastric ganglion which send their axons into the STG. In addition, both immunosustaining and radioimmunoimmunoassay indicate the pericardial organs, neurohemal structures similar to endocrine

Figure 3.

Effect of the neuromodulatory peptide proctolin on the gastric mill rhythm. A: Control pattern in a combined *in vitro* preparation. Notice the frequency of the bursts and the phase relationships between them. B: After perfusion of the STG with proctolin, burst parameters show a new steady state condition characterized by significant changes in burst length, frequency and phase relationships. C: & D: Block diagrams illustrating the parameter changes induced by the peptide.



glands in mammals, are extremely rich in a CCK-like peptide (CCK-LP). When purified extracts of the pericardial organs are compared with pure mammalian CCK-8 chromatographically, it is clear that although both are reactive to the same antibody, they are not the same molecule. Nevertheless, because it is reactive, the antibody can be used to answer several important questions about the neuromodulatory role of the CCK-LP in lobster.

1. By stimulating the sole input nerve to the STG, a CCK-like peptide can be released into the ganglion and collected from the saline bathing the ganglion.
2. Injection of CCK-8, even though not identical to lobster CCK, starts the gastric mill rhythm *in vitro* as the level of CCK rises in the blood.
3. The specific CCK blocker proglumide stops the gastric mill after it has been elicited by feeding or by injection. The blocker also works on CCK-induced patterns *in vitro*

Thus CCK appears to play a real role in the biology of the animal. We have looked at two other substances which can produce entirely different patterns – pilocarpine, a cholinergic muscarinic agonist, and proctolin (Figure 3), a small pentapeptide found in neurons sending axons to the STG. In all cases parametric changes in the output of the STG can produce different motor patterns leading to stable new movement

trajectories of the teeth (Figure 4). We are beginning to understand the mechanisms of how one fixed anatomical network can thus become functionally rewired as a result of the action of specific modulatory substances (Figure 5). Preliminary experiments with the second messengers C-AMP and IP_3 in both the pyloric and gastric systems suggest that they may provide the molecular mechanism of changing the cellular properties and possibly the synapses as well (Figure 6).

One additional important discovery has been the fact that not only do the neuromodulators switch the patterns to a different state, but they can “switch” neurons from one functional circuit to another. Two particularly interesting phenomena which support this concept have recently been described. The entire stomatogastric system actually consists of four separate central pattern generators: the esophageal, the cardiac sac, the gastric mill and the pyloric. One of the stomatogastric neurons, the VD cell, normally bursts with pyloric rhythm and is synaptically interconnected with other neurons in the pyloric circuit. However when the cardiac sac CPG is active as a

Figure 4.

A: Diagrammatic section of a lobster's thorax showing the stomach with the endoscope in position and two of the three teeth of the gastric mill. B-D: schematic drawings of the three teeth in typical positions as seen through the endoscope; B is the resting position.




Figure 5.

Scheme for the production of different behaviors by the same neural circuitry. The neural circuit at the top can exist in two states, off in the case of no excitatory input and on, in a basic core output mode when receiving descending input from other portions of the central nervous system. Modulators 1 and 2 can functionally reconfigure the circuit by altering synaptic strengths and cellular properties. This results in fundamental changes in the output patterns leading to two different behavioral states.




Figure 6.

Coding scheme for the action of neuromodulators at the cellular level. Peptides or amines can be released from neurons or neurohemal organs to eventually reach a neural network. Some or all of the neurons will bind the modulator starting a cascade of events which initially consist of raising the level of a second messenger and finally of phosphorylating specific ionic channels. The changes in channel function can result in alterations of input resistance, endogenous bursting, plateauing as well as other specific channel conductances.




Table 1.

Neuromodulatory substances presently known to be contained in the lobster.

Lobster Neuromodulators

Acetylcholine
GABA
Histamine
Octopamine
Serotonin
FMRF amide
Proctolin
Substance P
Red Pigment Concentrating Hormone (RPCH)
Small Cardioactive Peptide B (SCP-B)
Cholecystokinin (CCK)

result of sensory nerve stimulation, a strong excitatory synapse from this CPG to the VD cell becomes sufficiently strong to pull the VD cell out of the pyloric circuit and into the cardiac sac rhythm, i.e., the VD neuron switches from one CPG to the other (Hooper & Moulins, 1989).

In a separate set of experiments, Marder's group at Brandeis has shown that under the influence of the lobster neuropeptide red pigment containing hormone (RPCH), three neurons of the cardiac sac central pattern generator connect strongly with neurons in the gastric mill CPG in a way which produces a conjoint new "blended" rhythm (Dickinson, et al., 1990). Here the role of neuromodulator is to greatly strengthen some synaptic connections between the cardiac sac and the gastric mill CPGs which are usually too weak to be observed. Since it is well known that many muscles must participate in a wide variety of different behaviors these experiments are beginning to shed light on the neurological principles involved in the chemical control of different behaviors.

What we have discussed so far has been a qualitative description of the STG patterns and their modulation. The many interactions taking place between even the few neurons of the gastric mill and pyloric circuit can serve as a superb model system with which to develop quantitative modeling techniques which may have applicability to the formation of spatio-temporal patterns generated by more complex systems.

Application to Artificial Neural Networks

This new work exemplifies several themes emerging from studies of how small neural networks produce spatio-temporal patterns which may be useful in making artificial neural networks more flexible. First, the rhythmic output of the circuit

derives from both the membrane properties of the component neurons and the synaptic interactions among them. Second, many neurons have conditional oscillatory properties that contribute to the generation of network oscillations under some conditions but not others. Third, modulatory inputs control the expression of these conditional membrane properties, as well as the strength of certain synaptic connections within the pattern generating network.

Consequently, neural networks can provide variable outputs or can be "tuned" or "configured" in different ways. Comparing the features of many small neural networks shows that similar tasks can be performed by neural networks with different constructions and that a single neural network can produce many outputs.

The new understanding is that neuronal networks in the biological realm are not "hard-wired" but are made up of neurons whose membrane properties can be modified and connected by synapses whose strength, sign and time course can vary to produce multiple output states. This makes recent advances in computational neuroscience especially exciting since models show complicated tasks like the production of spatio-temporal patterns can be produced by groups of simple neuron-like processing elements hooked together by synaptic rules. Studies of modulation of simple networks like those found in the stomatogastric ganglion will help to understand how large biological networks generate behavior and process information. The challenge is to fully understand how small neuronal networks work, how larger networks of simple neuronal-like elements work, and how the complexity of real neurons and the richness of their connections define and constrain the possible functions of a neuronal network. We need to determine, therefore, which properties of biological networks are not shared by artificial neural networks. Insights from the study of small systems such as the stomatogastric ganglion may help in this endeavor.

References

1. Dickinson, P.S., Meccas, C. & Marder, E. Neuropeptide fusion of two motor pattern generator circuits. *Nature* 344:155-158, 1990.
2. Hooper, S.L. & Moulins, M. Switching of a neuron from one network to another by sensory-induced changes in membrane properties. *Science* 244:1587-1589, 1989.

Biography

Dr. Allen Selverston is a professor in the Biology Department of the University of California at San Diego. His publications and research are primarily in the field of small neural networks, especially of the lobster stomatogastric ganglion.

The Role of Interdisciplinary Research Involving Neural Systems in the Development of Intelligent Systems

*Dr. Thomas M. McKenna
Office of Naval Research*

A Brief History: The Divergence of Artificial Intelligence and Biology

Many of the pioneers of computer science, such as McCullough, Wiener, and Von Neumann drew upon biological inspiration for their work. The Office of Naval Research sponsored much of the early work on automata, self-organizing systems, cybernetics and neural networks, including the pioneering work of Rosenblatt and Widrow. In the seminal work of McCullough and Pitts (1943), and the Perceptron work of Rosenblatt (1962), an analogy was established between artificial processing elements and real neurons, between network connections and axons, and between connection strength and synapses. These simplified analogies formed a framework for the subsequent development of the field of artificial neural networks using rules for setting weights based on training and global network parameters. In these heady early days, there was not a very clear distinction between neural nets, automata, computation, cybernetics, and the search for machine intelligence.

Biological systems were seen as the most complex systems in the world, they were recognized as computational entities, and biological analogies from brain organization, evolution, ecological systems, and DNA coding were freely used as inspiration to drive the scientific development of intelligent machines. By the late 1960's a number of developments emerged to separate these fields, and even put them against each other. The Perceptron work ran into theoretical limitations, for as we know in hindsight, it lacked hidden layers and the means for training them. Moreover, computer technology was inadequate to directly implement many neural network ideas on a testable scale. Simultaneously, computer science solidified as an intellectual field, and within it emerged an approach to artificial intelligence as a software engineering problem, inspired in part by the emergence of cognitive psychology, which in turn was encouraged by the analogy between computing and thought. In 1969 Minsky and Papert (1969), in print and in person, attacked the already stalled neural network program, on the basis of the inability of Perceptrons to perform nonlinear classifications, and apparent problems in learning and scaling. Simultaneously, they promoted the artificial intelligence agenda for solving a host of

problems in vision and robotics. Subsequently, government agencies, in particular the Advanced Research Projects Agency proceeded to spend many hundreds of millions of dollars pursuing this line of research.

The Symbolic AI Program

Classical AI is based upon the Physical Symbol System Hypothesis, which posits that general intelligence operates through the manipulation of a system of symbols. Symbols represent the objects of the external world, and intelligence is presumed to involve our computations with, or manipulation of, these symbols or representations. As has been pointed out previously (Beer) this is only a hypothesis about human information processing, not a definition of intelligence. As an hypothesis, it must be reified by empirical investigation to determine under what circumstances it provides a sufficient explanation of intelligent behavior. AI has achieved successes within artificial worlds, like chess-playing, and solving some difficult mathematics problems, often by brute force computation. However, this approach has also led to some very expensive failures in dealing with the real world. Hundreds of millions of dollars were spent developing "machine vision" over several decades. The weakness of the approach is shown by the rapidity with which connectionist, neural net, and sensory chip approaches have eclipsed the AI approach in so called "lower" and "middle vision". It is not surprising that AI researchers have now retreated to higher levels like image "understanding". The failure of AI machine vision also led to expensive failures in the development of autonomous land vehicles, in contrast to the rapid improvements in performance achieved in these systems by incorporating neural network control. Overall the problem with AI systems has been that they tend to be narrow and brittle: they do not function well beyond limited domains, they are very sensitive to the representational choices made by the designers, and they tend to fail whenever real conditions deviate from the abstract world. These limitations are particularly evident in robotic control problems and have led to alternative, bottom up approaches such as the subsumption architecture and computational neuroscience approaches.

What is the source of the brittleness of AI solutions? All symbolic AI systems rely critically on the knowledge representation chosen, and much of the power of such systems rests on the careful crafting of these representations for each problem attacked. The rest of the system is then dependent on this representation. This leads to a lack of flexibility and generality and the inability to transfer experimental results from one AI group to another. Quite possibly, what is most needed is some *fundamental* understanding of how information is actually represented in the brain, how these representations in different sensory modalities interact, and how general inference and memory retrieval rules operate on these representations. This

question will be crucial for any real advances in artificial intelligence in the broad and deep sense. Computational neuroscience has not yet provided the answers, nor has symbolic AI. An increased understanding of animal and human intelligence, as provided by analyses of both brain mechanisms and cognitive processes, will be necessary for any significant improvement in intelligent systems. The most pragmatic approach to developing useful and robust intelligent systems is to study what are still the only systems that actually exhibit intelligence: biological systems, including humans.

With regard to machine learning, given the truth of the principle that initial knowledge is the most important element, the prime motivation to study learning, grounded in neural networks, is to explain how the initial knowledge came to be. Many symbolic machine learning programs are large constructions to begin with, and after learning, they have only converted some knowledge hidden in the architecture into explicit data structures (Lenat, 1983).

An even more serious problem for symbolic AI is its inability to scale up beyond "toy" problems. This is ironic considering the original attacks by Minsky on early neural nets over the scaling issue. AI suffers in this regard because of its rule based nature. The number of rules required for real world problems, the problem of maintaining them, the amount of precise programming required, and the brittleness of the result make the scaling up of symbolic AI systems to real world problems a highly dubious undertaking. Practically, achieving an intelligent system by AI programming is beyond the scope of current software engineering.

While it is logically true that human intelligence represents just one possible type of intelligence, and furthermore that AI may identify some other form of intelligence that may surpass human intelligence, this line of reasoning is empirically misleading. Thirty years of symbolic AI has not resulted in a single instance of software that comes close to equalling human intelligence. Intelligence is a powerful and impressive *natural phenomenon* having evolved over hundreds of millions of years. Creating intelligent systems by artificial means will not be easily achieved.

A Pragmatic Alternative: Biological Intelligence

It has been claimed that "biological correctness" is being unfairly applied as a criterion for the suitability of symbolic AI approaches to intelligent systems. Far from it, the pragmatic approach is to seek knowledge from biological systems, complex physical systems, and mathematics, among other sources of inspiration, and to incorporate them into new computer technology. One of the great advantages of biological neural systems is their adaptiveness. They can achieve goals in spite of great perturbations in the environment and develop

Interdisciplinary Approach for Neurobiologically Derived Neural Nets

PI	Brain Region	Model/Algorithm	Implementation
Edelman (Rockefeller)	Visual & S-I cortex	Group Selection Theory Darwin III	Simulation Robot prototype
Lynch & Granger (UCI)	Olfactory cortex Entorhinal cortex Olfactory bulb	Combinatorial Matrix Hierarchical Clustering Algorithm	Simulation 3 Patents VLSI silicon cortex
Cooper (Brown)	Visual cortex	RCE Neural Net BCM Theory	Nestor Learning System Intel/Nestor VLSI
Mead (Caltech)	Cochlea, Auditory system Retina	Electronic circuit equivalent Markov Random Field	VLSI cochlea VLSI sound localization VLSI retina
Berger (Pitt)	Hippocampus	Non-linear System	Transputers VLSI design
Bower (Caltech)	Olfactory bulb Olfactory cortex	Biologically Realistic Model Nets	GENESIS system
Freeman (Berkeley)	Olfactory System	K-III Nonlinear System	VLSI Classification System

appropriate responses to totally novel sensory events. It is not the case that the DNA code alone dictates how neural systems develop, thereby imposing harsh constraints on biological networks. A modern biological account of development of neural architectures and biological intelligence would distinguish at least four stages: (1) evolutionary development encoded within DNA, (2) epigenetic development (fetal), whereby physiological processes including simple local cellular interactions, and chemical gradients, interact recursively with DNA readout to produce the complex organization of tissues and organs, (3) the critical period of postnatal development in which sensory experience modifies the neuroanatomical substrate supporting the basic functional sensory and motor systems, and (4) diverse forms of learning in child and adult, in which functional connections are modified by experience through changes in the strengths of existing synapses, changes in the shapes of dendrite branches, and changes in the nonlinear biophysics of neurons, as well as changes in neural population dynamics. Obviously this account is greatly simplified, but the point is that by deriving the *biological principles* involved in experience-dependent self-organization of networks and the rules which govern synaptic plasticity during learning, we provide powerful tools for producing adaptive neural networks. This has a positive impact on this emerging technology, rather than imposing a limiting constraint. The ONR programs in biological intelligence have specifically funded the fundamental research into the development of a

host of learning rules (e.g., Hebb and anti-Hebb rules) from the study of biological systems. Such biological principles and rules are vital for the design of self-organizing, adaptive neural networks.

Neurobiological approaches also allow principled analyses of cognitive and perceptomotor processes into functional components. The growing interest in this approach is reflected in the emergence of a distinct discipline called cognitive neuroscience, as well as new ONR programs in this area.

The primacy of adaptiveness to unknown but structured environments exhibited by biological systems is also evident in new approaches to AI, namely the so called "grounded symbolism" school of AI (e.g., Brooks, 1989, 1990). This new AI methodology decomposes intelligence into a hierarchy of individual behavior-generating modules, whose coexistence and co-operation let more complex behaviors emerge. This bottom-up AI approach overlaps with the emerging advances in the computational neuroethology field base on models of simple systems such as insects (Beer, 1990), and should prove a fertile area for the development of adaptive systems, such as mobile robots, which incorporate many biological principles.

The ONR sponsored development of neurally inspired machine learning algorithms such as the backpropagation algorithm (Rumelhart & McClelland, 1986) has already led to widespread investigation of neural networks for applications such as signal classification, pattern recognition, and control. Furthermore, biological inspiration continues to provide a rich

source for the evolution of faster and more powerful artificial neural network architectures and algorithms. For example the Radial Basis Function (RBF) networks provide robust signal classification and are 2 to 4 orders of magnitude faster to train than backpropagation in multilayer perceptrons. The conceptual foundations for radial basis functions derive directly from the neuroscience concept of neuronal receptive field. The observation that neurons respond selectively to inputs over a small portion of their input space (the receptive field) led naturally to the development of RBF neural networks with tuned processing elements in the hidden layers (Moody and Darken, 1989). Hence, even within the mainstream of artificial neural networks, biological inspiration continues to provide effective new principles and architectures. Biology is not a limiting constraint to technology but a source of vital new concepts.

Neural Network Research is Diversified

To scientists and managers not directly involved in interdisciplinary research in neural systems, the neural networks field may appear monolithic (e.g. neural nets may be equated with "connectionism"). However, closer examination of the neural networks field reveals that it consists of a number of distinct subfields and related emerging disciplines. Each of these subfields differs in its approach and research agenda. These major subfields include:

Artificial neural networks - An engineering and industrial field with a pragmatic approach to using the tools of neural networks to solve problems in the real world and build electronic devices to support these applications. Performance and ease of implementation are primary drivers, and many incremental improvements and hybrid approaches have evolved combining neural networks, time-frequency feature extraction, and simple heuristics. The source of inspiration for these neural net solutions is irrelevant to this group. At present, commonly used neural networks include backpropagation in multilayer perceptrons, radial basis function nets, probabilistic neural networks, CMAC, and higher order nets. Several of these nets were developed under ONR support in the seminal period from 1957-1985. It is estimated that this neural net "industry" will achieve a worldwide size of one billion dollars by 1995 (Tom Schwartz, IJCNN meeting 1991).

Connectionism - The scientific study of neurally-inspired (but not constrained) models for cognitive and AI systems. Unlike symbolic AI, the emphasis is not on pre-defined, task-specific representations, but rather upon developing systems which discover representations.

Biological neural networks - The "reverse engineering" of the nervous system. The development and implementation of novel architectures and algorithms, for general technology

applications, based directly on neurobiology, in particular computational neuroscience. The size and funding scope of this area is minuscule relative to artificial neural networks, but ONR has been instrumental in the recognition of, and support for, this emerging area.

Computational neuroscience - The "computational" approach to neurobiology, by modeling of natural systems in great detail, from single neurons to systems level. Accurate simulations make predictions that can be empirically tested, and the goal is to achieve as much realism as possible, rather than develop new technology.

Computational neuroethology - The focus is on computer modeling of the neural control of behavior in simpler whole animals.

Complex systems - An active new area wherein physicists and mathematicians bring analytic tools from nonlinear dynamical systems to bear on problems in biology and cognition.

Obviously, neural network researchers are not a homogeneous group but consist of workers in various fields, each bringing particular mixes of interdisciplinary skills to bear on diverse problems. None of them, except for computational neuroscience, is concerned about whether a new algorithm is "biologically correct", although they have all adapted useful concepts from biology among a host of other influences.

The neurosciences continue to be a dynamic and growing enterprise, and they will continue to provide a rich source of new principles, architectures, mechanisms, and behavioral repertoires for technologists seeking new solutions for intelligent, adaptive systems. Through interdisciplinary research, we can take advantage of this resource or, through prejudice and parochialism, we can ignore it. In a highly competitive world, there will be consequences either way.

ONR Interdisciplinary Programs in Neural Networks

Our current ONR programs in neural networks support a direct and ambitious interplay of neuroscience, mathematics, computer science and electronics. The goals of our programs is to develop electronic implementations of computational and/or signal processing architectures based on models of neurobiological systems. Our current program is designed to develop intelligent systems per se. In fact, if one considers the hierarchy of levels of organization in neural systems from molecules, synapses, neurons, networks, sensory-motor subsystems, to (brain) system level, and a roughly analogous parallel hierarchy of electronic implementation from semiconductor physics, through Very Large Scale Integration (VLSI) to board level, computer, and high performance computer levels, the existing programs at ONR focus on the levels from neuron to subsystem and stops short of the system level. The highest level of integration of such a hierarchy, the human

brain, corresponds to the focus of attention of cognitive science efforts but is clearly beyond the immediate focus of our current neural network efforts. The ONR interdisciplinary approach to neural networks is outlined in Figure 1 for seven groups of investigators headed by the named principal investigator. A research group will: (a) focus on the physiological and anatomical features of a particular brain region (1st column), (b) extract from that region a basic architecture or algorithm which captures the computations performed by that brain region, and examine the theoretical properties of this network (middle column) and (c) implement this new architecture as a computer simulation or in hardware, in order to examine the performance of this net on signal processing or classification tasks (3rd column). Such collaborations require the modelling and computer simulation skills of mathematicians, computer scientists and electronics engineers. The novel features of the next generation, neurobiologically-derived neural networks, include: advantages in network size scaling (the brain has at least 10^{14} neurons), a faster learning rate, modular organization for decomposition of inputs, distributed computing, self-organization of nets, and the ability to handle dynamic signals with information in multiple time scales. These are the technology goals we expect to achieve based upon current and future research advances. In order for neural networks to handle the demanding signal processing and control problems faced by the Navy and industry, the capabilities described above are essential.

Among the many accomplishments of this program are the following developments: (1) an efficient hierarchical classifier (Lynch and Granger) based on a realistic model of the olfactory cortex which has been implemented in VLSI and which is under evaluation for sonar, handwritten character recognition, and machinery vibration diagnosis (Granger, et al. 1989, Ambros-Ingerson, et al. 1990, Leen, et al. 1991); (2) new functional forms of the Hebb algorithm governing adaptive weight changes (Sejnowski & Stanton 1989, Brown et al. 1990); (3) a biological neural network based on parietal cortex which computes 3-D head centered space by fusion of 2-D retinal inputs and eye position information (Mazzoni, et al. 1990); (4) establishment of an equivalence between the Widrow-Hoff LMS rule of adaptive filtering with the Rescorla-Wagner model of incremental animal learning and link it to neural mechanisms in the cerebellum (Thompson & Gluck, 1989, Gluck, 1991); (5) a novel neural architecture based on biology learning mechanisms in invertebrates was successfully applied to kanji character recognition (Vogl, et al. 1991); (6) a large scale model based on the interactions of multiple cortical visual areas of primate was able to exhibit perceptual grouping and figure ground separation in vision (Sporns, Tononi & Edelman, 1991); (7) demonstration of the connection between a model of visual cortical plasticity (which has been successful in phoneme recognition and 3-D object recognition) and the statistical method of projection

pursuit (Intrator & Cooper, 1991); (8) a new class of supervised learning algorithm, based on a model of the cerebellum, that was orders of magnitude faster than backpropagation (Moody, 1989); (9) a new class of analog neural circuits with local connectivity was developed for high speed image processing (Chua & Yang, 1988); and (10) a number of analog VLSI sensory chips were developed based on the cochlea, auditory brainstem, and retina (Mead, 1989). They are capable of rapid enhancement and tracking of visual targets, and speech processing. Recently these chips have incorporated more realistic features of single neurons, and have demonstrated capabilities in recognition of sequences of signals.

Where is Neural Network Technology Headed?

The current status of neural network R&D, particularly artificial neural networks for applications, can be gleaned from analyzing the presentations at national conferences or by the results of the DARPA neural network program on comparative performance. The former reveals that neural networks can produce solutions to specific industrial problems in pattern recognition, control, design, logistics, machinery monitoring, adaptive routing in power grids, financial predictions, credit risk assessment, and autonomous vehicle control. These solutions can be obtained with relatively modest investments in time and resources. Frequently, they are not total neural net solutions, the neural net is just one component in a larger system. While recent progress is encouraging, too often reports have an anecdotal quality, failures may not be reported, and comparative performance is often not specified.

By contract, the DARPA program was specifically designed to provide comparative performance measures between neural networks and current statistical or signal processing techniques on a number of classification and pattern recognition problems including: sonar signal classification, automatic target recognition, seismic signal classification, and speech recognition. These are problems where considerable engineering expertise has already been applied, yet improvements are necessary for automatic classification of real world signals. In each of these applications, neural network approaches matched or exceeded the current best non-neural techniques with only a 2 year development time. Neural network hardware implementations currently underway should lead to substantive improvements in signal classification capabilities for DoD problems and industrial needs. During the course of the DARPA program substantial refinements in neural network architectures and algorithms were made that dramatically improved the accuracy and speed of training for the particular problems addressed.

This is a dynamic field, and comparisons between a generic network, like backpropagation, with a more traditional

signal processing or AI technique is misleading since substantially better performance and orders of magnitude greater speed is likely to result from a more suitable network architecture and feature extraction procedure. In other words, the very encouraging results from the current DARPA comparative performance studies are, if anything, only a conservative estimation of the results that will obtain for the neural networks now evolving. Many of the existing comparative performance estimates have been made using neural networks implemented in software on general purpose computers. Direct implementation of neural networks in VLSI chips is producing supercomputer levels of speed, such that the inherent parallelism of neural architecture is being approximated. The past five years of engineering experience with neural networks have shown it be cost effective for many pattern recognition problems. The development time necessary for effective neural network implementation is substantially below that of alternative signal processing or AI solutions. The emerging evidence of neural network performance suggests that further funding is indeed warranted. Indeed, to ignore the potential of this technology for rapid improvements in signal processing speed, accuracy in recognition, and compact implementation for noisy, variable and information-rich environments is to invite technological surprise.

A typical current pattern classification performed with a neural network will have several stages of processing. For example, acoustic signal classification requires (a) a pre-processor stage (detection, feature extraction and optimization) using signal processing techniques such as energy detectors and time-frequency digital representations, followed by (b) an artificial neural network classifier, and possibly (c) postprocessing by a heuristic or symbolic AI system. The latter is used for arbitration if there are multiple signal paths through multiple nets, for data fusion with other signals, or for recognition of sequences of classifications. However, neural net technology is evolving for both the preprocessing and the postprocessing stages.

New neural techniques are now demonstrating good performance in feature extraction, temporal sequence recognition, and data fusion. Hence, it may soon be possible to develop complete neural network solutions for some demanding pattern recognition problems, such as scenario recognition and context sensitive recognition. The capabilities of current artificial neural networks are being extended by (a) development of dynamic networks for time sequence processing (Pollack, 1991, Giles, 1991, Todd & Loy, 1991, Granger & Lynch (personal communication), (b) structured networks for faster learning (LeCun, 1989, Grover, 1991), and (c) modular networks.

Thus, for real world problems, neural systems would likely consist of a number of neural networks rather than one homogeneous network. For example, modular, hierarchical

networks are being developed by artificial neural network research (Mjølness, 1989, Seibert & Waxman, 1991), and by biological neural network research (Granger et al.). Indeed, Waxman's neural system is capable of recognition of 3-D objects from 2-D scenes, hardly a trivial, or low-level problem.

Toward Intelligent Systems

It has been asserted that the proper application for neural networks is as a sensory or peripheral pattern recognition processor in hybrid architectures in which symbol processing takes the effective controlling role (Meyrowitz, 1991). The application of neural networks for cognitive problems has been less studied than symbolic systems for certain historical reasons, but connectionist models have already been used for purposes of language understanding (Waltz & Pollack, 1985, Cottrell, 1985, Hanson & Kegl, 1987, St. John & McClelland, 1990), logical reasoning (Ajjanagadde & Shastri, 1991, Derthick, 1990), game playing (Tesauro & Sejnowski, 1989), and autonomous vehicle control (Pomerleau, 1988). There has been connectionist support for traditional production systems (Touretzky & Hinton), including an attempt at finding a basis for the SOAR architecture (Cho et al. 1991). And recent research has also addressed what Meyrowitz (1991) called one of the "grand challenges to achieving with neural networks the computational capabilities afforded by symbol manipulation", namely the issue of dynamic representations for symbol manipulation Pollack (1990), Smolensky (1990), and Ajjanagadde and Shastri (1991).

Finally, neural network research is addressing a central problem for cognition, the acquisition of languages. Giles et al. (1991) have developed recurrent neural network learning architectures which learn formal grammars, and from which finite state automata can be extracted. Pollack (1991) has developed a recurrent network which uses the phase transitions in network learning to induce complex languages. Analysis of the relationship between the system dynamics of iterative neural networks and the generative processes (of grammars) may provide an exciting and productive new direction for both information science and cognitive neuroscience. These developments provide new opportunities for developing interfaces between neural networks and traditional computing. This is a large task, and many issues await empirical resolution. Clearly, *a priori* arguments from the AI community regarding the limitations of neural networks as the basis of intelligent systems is premature.

Neural network solutions will soon be able to take advantage of tremendous speed increases by direct implementation in VLSI hardware and the neural supercomputers under development. By stark contrast, symbolic processing has never taken advantage of new architectures, such as Cray supercomputers or the Connection Machine.

Hence, the role of neural networks in future intelligent systems remains an open research question. There have been encouraging developments in this direction, and this interdisciplinary effort will continue to benefit from detailed consideration of biological neural architectures and function.

Interdisciplinary research requires some tolerance to allow the interplay between the scientists working in different fields with different goals. It is a mistake to take the evaluative criteria of one field and apply them to another. Not all neural net research is good AI, not all connectionism is good psychology, not all computer science is good mathematics, and not all of AI is good computer science. In fact, researchers in AI tend to ignore the very issues of tractability and efficiency which are critical to the computer science perspective. In order to achieve progress on such ambitious, long range goals as developing intelligent, autonomous and adaptive systems that function in real world environments, new research accomplishments will be required of many disciplines, and these must be melded together effectively by cooperative endeavors. Yet, if productive interdisciplinary research is to proceed we must resist the temptation to critique one discipline by the limited criteria of another discipline.

Biography

Dr. Thomas M. McKenna is a Scientific Officer in the Biological Intelligence Program of the Division of Cognitive and Neural Sciences of the Office of Naval Research. He also manages portions of the DARPA neural networks program. His research interests include single neuron computation, computational modelling of sensory and motor systems, computational neuroscience, legged locomotion, and neural network applications in sonar signal processing and machinery diagnostics.

References

1. Ajjanagadde, V. & Shastri, L. Rules and variables in neural networks. *Neural Computation*, 3, 121-134.. 1991.
2. Ambros-Ingerson, J., Granger R. & Lynch, G. Simulation of paleocortex performs hierarchical clustering. *Science*, 1990, 247:1344-1348.
3. Beer, R.D. *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*, 1990, Academic Press, Boston, MA
4. Brooks, R.A. A robot that walks: Emergent behavior from a carefully evolved network, *Neural Computation*, 1989, 1(2):253-262.
5. Brooks, R.A. Elephants don't play chess. *Robotics and Autonomous Systems*, 1990, 6:3-15.
6. Brown, T.H., Kairiss, E.W. & Keenan, C.L. Hebbian synapses: Biophysical mechanisms and algorithms. *Ann. Rev. Neurosci.* 13:475-511, 1990.
7. Cho, B. Rosenbloom, P. & Dolan, C. Neuro-soar: A neural network architecture for goal-oriented behavior. In: *Proc. 13th Ann. Mtg. Cognitive Science Soc.*, Chicago, 1991, 673-677.
8. Chua, L.O. & Yang, L. Cellular neural networks: Theory. *IEEE Trans. Circuits & Systems*, 35: 1257-1272, 1988.
9. Cottrell, G.W. A connectionist approach to word sense disambiguation. TR154, Rochester: Univ. of Rochester, Computer Science Dept., 1985.
10. Derthick, M. Mundane reasoning by settling on a plausible model. *Artificial Intelligence*, 1990, 46:107-158.
11. Giles, C.L., Miller, C.B., Chen, D., Sun, G.Z., Lee, Y.C. Learning and extracting finite state automata with second-order recurrent neural networks., *Neural Computation*, in press.
12. Gluck, M.A. Stimulus generalization and representation in adaptive network models of category learning. *Psych. Sci.*, 1991, 2(1):1-6.
13. Granger, R., Ambros-Ingerson, J., Anton, P., Whitson, J., & Lynch, G. computational action and interaction of brain networks. In: *Introduction to Neural and Electronic Networks*, S. Zometzer, J. Davis, C. Lau, (Eds.), Academic Press, New York, 1989.
14. Grover, M. Artificial neural network sensor fusion and track correlation, Defense Group, Inc., technical report (DARPA & ONR), 1991.
15. Hanson, S.J. & Kegl, J. (1987) PARSNIP: A connectionist network that learns natural language grammar from exposure to natural language sentences. In: *Proceedings of the Ninth Conf. Cognitive Sci. Society*, Seattle, 106-119.
16. Intrator, N. & Cooper, L.N. objective function formulation of the BCM theory of visual cortical plasticity: Statistical connections, stability conditions, 1991, in press.
17. John, M.F. St. & McClelland, J.L. Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, 1990, 46:217-257.
18. LeCun, Y., Boser, B., Denker, J., Henderson, D., Howard, R., Hubbard, W. & Jackel, L. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1989, 1:541-551.
19. Leen, T., Webb, M. & Rehuss, S. Encoding and classification in a model of olfactory cortex. *Proc. Int. Conf. Neural Networks*, July, 1991, 2:553-559.
20. Lenat, D.B. & Brown, J.S. Why AM and Eurisko appear to work. *Proc. AAAI-83*, 236-240, 1983.
21. Mazzoni, P., Andersen, R.A., and Jordan, M.I. Ar-p learning applied to a network model of cortical area 7a, *Proc. Int. Joint Conf. Neural Networks*, June 1990, 2:373-379.

22. McCulloch, W.S. and Pitts, W.H. A logical calculus of ideas imminent in nervous activity, *Bulletin of. Mathematical Physics*, 5:115-133, 1943.
23. Mead, C. *Analog VLSI and Neural Systems*, 1989, Addison-Wesley, Reading, MA.
24. Meyrowitz, A.L. Neural networks: A computer science perspective. *Naval Research Review*, 1991, 43(2):13-18.
25. Minsky, M.L. & Papert, S.A. *Perceptrons*. 1969, MIT Press, Cambridge, MA.
26. Mjolsness, E., Gindi, G., & P. Anandan. Optimization in model matching and perceptual organization, *Neural Computation*, 1:218-229, 1989.
27. Moody, J. Fast learning in multi-resolution hierarchies. In: Touretzky, D.S., Ed., *Advances in Neural Information Processing Systems I*, Morgan Kauffmann, San Mateo, CA, 1989.
28. Moody, J. & Darken, C.J. Fast Learning in networks of locally-tuned processing elements, *Neural Computation*, 1989, 1(2):281-294.
29. Pollack, J.B. No harm intended: A review of the Perceptrons expanded edition. *J. Math. Psychol.*, 1989, 33:358-365.
30. Pollack, J.B. Recursive distributed representation. *Artificial Intelligence*, 1990, 46:77-105.
31. Pollack, J.B. The induction of dynamical organizers, *Machine Learning*, 1991, 7:227-252.
32. Pomerleau, D.A. Efficient training of artificial neural networks for autonomous navigation, *Neural Computing*, 1991, 3:88-97.
33. Rosenblatt, F. *Principle of Neurodynamics*. Spartan Books, Wash. DC, 1962.
34. Rumelhart, D. & McClelland, J. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. I & II, MIT Press, Cambridge, MS, 1986.
35. Seibert, M. & Waxman, A.M. Learning and recognizing 3D objects from multiple views in a neural system. In: *Neural networks for Human and Machine Perception*, H. Wechsler, Ed., Academic Press, New York, 1991, in press.
36. Sejnowski, T.J. & Stanton, P.K. Covariance storage in the hippocampus. In: S.F. Zornetzer, J.L. Davis, & C. lau (Eds.), *An Introduction to Neural and Electronic Networks*. Academic Press, New York, 1989.
37. Shastri, L. A connectionist approach to knowledge representation and limited inference. *Cog. Sci.* 12(3):331-392, 1988.
38. Smolensky, P. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 1990, 46:159-216.
39. Sporns, O., Tononi, G., & G.M. Edelman. Modeling perceptual grouping and figure-ground segregation by means of active reentrant connections. *Proc. Nat. Acad. Sci.*, 1991, 88:129-133.
40. Tesauro, G. & Sejnowski, T. A parallel network that learns to play backgammon. *Artificial Intelligence*, 1989, 39:357-390.
41. Thompson, R.F. & Gluck, M.A. A biological neural network analysis of learning and memory. In: S.F. Zornetzer, J.L. Davis, & C. lau (Eds.), *An Introduction to Neural and Electronic Networks*. Academic Press, New York, 1989.
42. Todd, P.M. & D.G. Loy. *Music and Connectionism*, 1991, MIT Press, Cambridge, MA.
43. Touretzky, D.S. & Hinton, G.E. A distributed connectionist productions system. *Cog. Sci.* 12(3):423-466, 1988.
44. Vogl, T.P., Blackwell, K.L., Hyman, S.D., Barbour, G.S., & Alkon, D.L. Classification of hand-written digits and japanese kanji. *Proc. Int. Joint Conf. Neural networks*, July 1991, 1:97-102.
45. Waltz, D.L. & Pollack, J.B. Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science*, 1985, 9:51-74.
46. Werbos, P.J. Backpropagation: Past and future., *Proc. Int. Conf. Neural Networks*, I, 343-353, IEEE Press, New York, July 1988.

Research Notes

Dyslexia and Vision

An important finding about parallel pathways in the primate visual system by ONR supported investigator, Dr. Margaret Livingstone (Harvard Medical School) has recently been tied to the clinical problems of dyslexia and learning disorders. The results, described in a recent issue of the Proceedings of the National Academy of Sciences (and front page of the NY Times) suggests that dyslexia may result from a brain abnormality that involves the sense of vision.

Dr. Livingstone's earlier work had described the visual system as composed of two major pathways. The magnocellular system is composed of large cells that carry out fast processes (used for seeing motion, stereoscopic vision, depth perception, low contrast). The parvocellular system operates at a slower speed and specializes in color, detailed forms and high contrast. Dyslexic patients had no trouble seeing briefly presented high-contrast stimuli. However, at low contrast, when only the fast magno system should be responsive, the dyslexics showed a greatly diminished response, suggesting that the magno system was sluggish. Finally, autopsies on five dyslexic brains showed magnocellular layers to be more disorganized and smaller than normal. These results are a good example of the serendipity that can occur in basic research. Dr. Livingstone was primarily interested in understanding visual information processing. The neural understanding of dyslexia is a fortunate byproduct.

Neural Bases of Learning and Memory

Long-Term Potentiation (LTP) is a long-lasting use-dependent increase in the efficacy of synaptic transmission in the brain that is thought to underlie certain forms of learning and memory. For the past twenty years a scientific debate has raged around the question of synaptic site and mechanism. Is LTP a result of enhanced neurotransmitter release (presynaptic)? Or, is LTP a result of increased sensitivity (e.g. more receptor sites) at the postsynaptic site.

Two new pieces of evidence would seem to favor the postsynaptic hypothesis. First, ONR-sponsored work at the University of California, Irvine has shown that neurotransmitter release can be enhanced pharmacologically without an increase in LTP. Secondly, new ONR-sponsored work at Champaign-Urbana has provided electron microscope evidence of morphological change after LTP. Specifically, LTP produces an increase in synapses on short, wide-necked dendritic spines which are thought to reduce longitudinal electrical resistance thereby making the current-conductance more linear and enhancing the synaptic "weights" in the system. Is this the end of the story? The history of the field suggests not. In fact, ONR-sponsored work at the University of Iowa using newly developed microsensors able to measure presynaptic quantal release of glutamate in vivo, is gearing up to answer some of these questions about the basic mechanisms of learning and memory.

DEPARTMENT OF THE NAVY
OFFICE OF NAVAL RESEARCH
ARLINGTON, VA 22217

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, \$300