

ON BEING HUMAN IN THE COMPUTER AGE

Michael A. Arbib

COINS Technical Report 82-17

(October 1982)

This technical report is the text of a lecture delivered on April 15, 1982 as part of the Chancellor's Lecture Series at the University of Massachusetts at Amherst.

CHANCELLOR'S LECTURE

Delivered by Michael A. Arbib

April 15, 1982

Introduction by Chancellor Henry Koffler:

Ladies and Gentlemen, it is a great pleasure for me to welcome you this evening to the final event in the 1981-1982 Chancellor's Lecture Series which completes the third full year in which I have been privileged to introduce a number of the most distinguished scholars and speakers on this faculty. This has been the eighth year of this special series created to do public honor to the most deserving of our faculty members, to bring together interested members of the wider scholarly community, and to transcend the usual boundaries of the academic disciplines which each of us holds most dear. I am convinced that these Lectures have accomplished these aims well, of publicly sharing and celebrating the best aspects of higher education. It is my hope that this lecture series will continue as a major showcase of this university's excellence.

Tonight's speaker is Michael Arbib, Professor of Computer and Information Science and one of the world's leading scholars and theorists in the fields of brain theory, cybernetics and artificial intelligence. Professor Arbib attended the University of Sydney in Australia and received his Ph.D. in Mathematics from M. I. T. for a thesis on probability theory. After visiting appointments at Imperial College in London and at the University of New South Wales in Sydney, he taught for five years at Stanford where he developed a cybernetic approach to cognition which was to develop later into his theory of schemas. In 1970 Professor Arbib joined the faculty of this university, and as Chairman from 1970 to 1975 he had primary responsibility for converting the then-existing Master's level Computer Science Program into the full-fledged Department of Computer and Information Science of today in which he is now a Professor.

He also made major contributions to the creation at this University of both the Center for Systems Neuroscience and the Cognitive Science Program. In addition to the fields already noted, Professor Arbib's interests include the algebraic theory of computation and control, and linguistics as well as the social and philosophical implications of all these areas of study. He has given numerous invited lectures here and abroad, has published

extensively in scholarly journals, and is the author or co-author of many books including *The Metaphorical Brain*, which received the American Society for Information Science award as the best information sciences book in 1973. Other books are *Brains, Machines and Mathematics*, *Computers and the Cybernetic Society*, and *The Design of Well-Structured and Correct Programs*.

Professor Arbib has been selected to deliver the world-renowned Gifford Lectures (which is second only to the Chancellor's Lecture!) in natural theology at the University of Edinburgh in November of 1983 with Professor Mary Hesse of Cambridge University. His topic on that occasion will be *The Construction of Reality*. I want you to know he is in very good company -- previous lecturers who have given the Gifford Lectures have been Henri Bergson and Nils Bohr, Arnold Toynbee, Reinhold Niebuhr, Alfred North Whitehead, and quite a few others. I think we can be proud of having Michael represent himself and the scholarly world of this University in 1983 in Edinburgh.

Tonight he will address the topic *On Being Human in the Computer Age*. It is with immense personal pleasure and institutional pride that I introduce to you one of my most distinguished colleagues of this institution, a true renaissance man, Professor Michael Arbib.

The Computer Age

Ladies and Gentlemen, it's a great honor to be here tonight, an honor tinged with sadness because this is the last Chancellor's Lecture at which Henry will be the Chancellor, so I think I speak for all of us when I wish both Henry and Phyllis every good wish and every success in Arizona, and I'd also like to wish all of us every success in finding an administrative team who will look after us as well as we have been looked after in the last few years. Now for the talk.

Goethe said, and it was something that my thesis advisor repeated to me when he wanted me to concentrate on my thesis, "If you would master the infinite, take the finite and master it from all sides." When I chose the topic On Being Human in the Computer Age, it was with the aim of really giving you two talks: one on trying to assess the extent to which progress in artificial intelligence and brain theory and cognitive science could tell us to what extent we are machines, and the second on trying to see the way in which computerization and automation were changing the workplace and changing the human condition. But I chose a rather grandiose title, On Being Human in the Computer Age, and so I am afraid that there are times when I may get a little carried away tonight and speak about things of which I know little. But if the curse of the modern academy is overspecialization, perhaps a little hutzpa will do no harm.

When I say "the computer age", I'm very much aware that this is not just the computer age. When we try to describe this century of ours, there are many other terms besides 'computer' that come to mind. This is the century of horror, the century of Hitler and Stalin, of genocide and the gulag. It's a century of staggering anachronisms, when Iranians who would return their state to the fundamental Shi'ite religion of the 8th century use cassette recorders to spread the word of their Ayatollah, and when perhaps the group most adept at the art of TV networking and computer mailing in this country comprises those who would have us adopt the most fundamentalist form of Christianity. It's also the century in which we have become more conscious than ever before that we live in this world and not just in a village or a town or a nation. I was

most struck by this in 1969 when the Appollo 11 astronauts returned from the moon, and splashed down in the Pacific Ocean 3,000 miles from the United States. The TV announcer said, "The astronauts are home!" For the first time in human history, just to be on the surface of this planet was to be home. Yet at the same time that we come to think of ourselves in this global perspective, we find that our world is threatened by overpopulation, by pollution, by the depletion of natural resources, and this very consciousness of one world is shattered by the resurgence of nationalism as each nation tries to grab for what it can get. So all this will be in the background as we talk on that finite topic, the impact of computers at the end of the 20th century:

When I talk about the impact of computers, I'm going to be really thinking of the computer in two different ways: the computer as a tool, and the computer as a metaphor. In speaking of the computer as a tool, I'm going to stress two particular aspects tonight: the increase in automation, changing the nature of work in our society as more and more tasks become accessible to the computer; and the way in which the computer is changing our notion of literacy, as what it means to read and write and think gets changed by the computers that we can interact with and by the rise of the computer as a personal device which each of us, at least in the well-developed societies, can have in our own homes.

When I talk about the computer as metaphor, I'm going to be looking at artificial intelligence, the attempt to program computers to do things that you would swear were intelligent until you knew that a computer had been successfully programmed to achieve them. The question I will ask is: to what extent, as we come to better

understand the mind of the machine, can we come to see a continuity with our own minds? Is there an impermeable barrier between man and machine, or are we highly sophisticated mechanisms; and if so, what does that mean?

The Image of the Computer

But before I get into these questions, let me just step back a little and discuss with you the image of the computer. For each of us the computer means something different, and I would like to ensure that we have certain agreements about what is being discussed. For those of us over 40, the computer is probably the symbol of conformity -- we remember when computers required information to be put on paper tapes or punched cards, and we remember the old jokes about "I am a human, do not fold, spindle or mutilate me." But if you're under 20, the computer is probably something quite different -- what you play video games on, -- and for you a computer is fun, and a friend, and probably a lot better than your parents. We now see the computer going from the tool of the elite, as it were, to something that is quite universal, and I think that we will see a transition to computer literacy.

There was an article a few years ago in the New York Times about how much we take literacy for granted. It gave interviews with various people who managed to survive in New York without being literate: the desperate ploys when in a restaurant of saying "What looks good to you on the menu?" and then following suit; or of trying to get the instructions to a friend's house that didn't require one to read the street signs. The ability to read and write

has become part of our subconscious; we no longer think about it. Yet before the invention of the printing press, to be literate was to be an exceptional person, to be part of a small elite. I was struck, when reading the biography by John Gardner of Geoffrey Chaucer, that in the middle of the 14th century Chaucer was very fortunate to have a tutor who was exceptionally rich and possessed all of 50 books. Yet each of us takes for granted that we can accumulate hundreds of books of our own, just as we take for granted the use of reading to find our way around our world.

Just as we have become used to street signs and menus, so will we become used to the computer; and, in the words of my colleague Connie Wogrin, we will come to find the computer as "invisible" as reading and writing. We have long been used to sharing our thoughts with distant friends by writing letters to them; it's not all that long ago that people found it frightening to communicate by telephone, and now we all take this to be a matter of course; and now we're beginning to use electronic computer-mediated mail. The computer is going to change the way in which we think and solve problems. There will be no need to be very good at mental arithmetic when we can use the computer to keep track of the arithmetic details and let us concentrate on the problem-solving. The ability to write down our thoughts has long relieved us of the skills of rote memory that were common in preliterate societies. With computers we will develop yet different memory skills to learn about retrieval, and about how to best organize information within our computer network so that we can find what we want.

There is a professor in the School of Education at this University who believes that the 3 R's are the tool of the elite and looks forward to the day when Ph.D.s will not be required to be literate any more. But I suspect that he is mistaken. There will still be the need to articulate carefully what we have to say, to edit it, to accumulate it, so that we can best express our thoughts. To what extent those skills will require us to write in longhand vs. typing into our console and editing, I don't know. With the development of computers we will see not so much a decay of skills that we have, but rather their enhancement. But there is still the fear that the computers are warping our lives in ways that we don't really understand.

E. G. Schumacher, author of "Small is Beautiful", distinguished machines which serve men from machines which require men to serve them. Unfortunately, the distinction is a very hard one to hold. What is the telephone? When you want to contact a distant friend, it's a machine that serves man. When you are happily ensconced in the shower and the phone rings, it's a machine that forces men to serve it. Similarly for the car, and similarly for the computer. We cannot yet think through well enough what it is we really want for ourselves. We are seeing increasing automation taking over not only the blue collar jobs but also the white collar jobs. In some ways this is good, because suddenly each of us can afford the services which before only the elite could afford: we can have check accounts, we can have dishwashers, all because we have mechanical slaves. The only trouble is that there are a lot of people who made their living doing the drudgery that those machines are taking over. How, then, do we build an adaptive dynamic society

in which people are not defined by some particular social role and then thrown aside when society discovers that that particular skill is no longer needed? The problems that we face, while they may be exacerbated or brought into focus by the computer, are much broader than that. They are human problems, but human problems taking a different shape at this stage in our history.

On Being Human

There has been much work in philosophy and religion and literature and poetry and drama and art to try and address the issue of what is to be human, and I shall certainly not try to recapitulate all of it right now!

Those of us working in artificial intelligence have tended to stress things like problem solving, game playing, question answering, ability to use some aspects of a natural language, and elements of vision in our study of 'intelligence'. We are encouraged by our increasing ability to write computer programs that we think provide some insight into the human mind. But we must not pretend that those dimensions exhaust what it is to be human. True, Aristotle and many since him defined man (and woman) to be the rational animal. But there's a lot more to being a human than being rational: there are the dimensions of emotion, love, compassion, of having a family, being part of a community, belonging to a society. Unfortunately, I won't say much about these aspects.

These are the positive sides of being human. A title like "On Being Human in the Computer Age" has a utopian ring to it -- Let's all be better humans through love of our machines! But there are other things about being human: there's evil, hatred, jealousy, and war. In fact, if we were to bet on two things that will continue right through the computer age, they might well be war and hatred. But we will be a little more utopian, to see some way in which we can steer our way through the shoals of the remaining 6,470 days of the 20th century. [This is how machines change you: I happen to have a calculator which allows one to subtract one date from another -- if it wasn't for that machine, I would never have known that there were 6,470 days left in the 20th century. This is perhaps a reason to despair about the overprevalence of computers in today's environment...]

Although much of what I say will be focussed on one of those aspects of our humanity most affected by current developments in computers, I'm still trying, in those words of Goethe, to get that handle on the infinite. I do not think that being human is some sort of unary thing that we already know. I do not believe that there is already an ethical or religious system in place that exhausts what it is to be human, and that our look at the computer age can fit neatly into that system. I think that there is no ultimate reality in being human, and that if you believe there is an ultimate reality, you are doomed to delusion or to despair. Rather, I see our reality as a contingent one: we do the best we can. And that best that we can do is conditioned both by our intuitive feeling as humans, as members of our society, and by our rational analysis of social, cognitive, and psychological "forces".

Man is a Machine

I really believe that "Man is a Machine". A friend in the Theatre Department was going to make me a breastplate with transistors and flashing lights, and I was going to dramatically rip open my shirt at this moment to justify my claim that I was a machine, but ...

Humans are very resistant to new ideas, yet after we become used to them we forget that there was ever any problem in accepting them. Copernicus advocated the idea, one that had lain dormant for about 2,000 years, that the earth actually moved. This was obviously a stupid idea -- the earth is stationary, it's the sun that rises, not the earth that turns, as you notice every morning that you get up early enough. When Copernicus first came out with this idea, although it created a certain stir, the Church wasn't very worried. But the religion of that time, Christianity as reflected in Dante's universe, had the tiers of hell within the earth, and then -- in spheres around the stationary Earth -- the planets, the stars, and beyond that God's heaven. But if the earth was just another orb moving through the heavens, and if the stars could extend to infinity, where were heaven and hell? So by the time Gallileo came on the scene, his theories were ripe matter for the Inquisition.

We have assimilated all that. Although we still say the sun rises and probably mean it, we understand that this is just a way of speaking and that, yes, it is the earth that rotates as it moves in its elliptical orbit about the sun. It would be hard to imagine that a legislature in this country would enact a law saying that

equal time should be given to the view that Dante's universe is an acceptable scientific theory. It's strange, isn't it, that Darwin's theory has not fared so well, so that it is still a matter of heated debate in certain state legislatures as to whether or not we really did descend from (horror) apes. In fact, if one looks at history, it is probably bad enough that we descended from some of those humans that were around a little while ago...

What I want to argue is that our study of minds, brains and computers is bringing us to a similar conceptual revolution. Freud showed us that much of our apparently free behavior has deep roots within our subconscious. Freud started as a neurologist, and he began to develop a mechanistic view for understanding what might be apparently irrational in our behavior. We are now coming to better and better understand the mechanisms of mind, the mechanisms of the brain. I would like to think (perhaps wrongly) that 100 years from now it will seem as absurd that there were people back in those benighted days of the 20th century who doubted that mind was a mechanistic phenomenon as we now find it strange that people should doubt the concept of evolution or that the earth does indeed move.

I should offer a caveat here about what I take a scientist to be -- I assure you I'm not quite as dogmatic as I sound. There's one image of science as progressing in terms of "normal science", some theory that explains just about everything, until too many facts accumulate that cannot fit into this framework. Then for a while there is turmoil, until these new facts are assimilated in a new theory. This is the Kuhnian model of scientific revolutions, moving from one paradigm to another. In fact, my experience -- reflected in a book like Feyerabend's "Against Method" -- is that

science is much more pluralistic, even perhaps anarchistic. Not that scientists go around with bombs, throwing them into each others' meetings, but anarchistic in the sense that science does not progress by the Academy saying "this is the true theory," but rather as a result of many people trying different theories. Now let us see the dichotomy in views of the relation between mind and brain. Some scientists, perhaps starting from a certain religious world view, are struck by how huge the gap is between those aspects of mind that we can now understand in terms of computer programs or neural networks -- certain aspects of vision, memory, and motor control -- and the full richness of being human to which we have already alluded. There are others of us who note that it took billions of years for an amoeba to evolve to a human, and that without the fossil record and a bit of good theory, you wouldn't really believe that such an evolution could happen; and then argue that, similarly, there is no discontinuity that we can yet see that would distinguish machines in their evolution from humans.

In our Department of Computer and Information Science at this University, and in our interactions with our colleagues in the Cognitive Science Program in Psychology and Linguistics and Philosophy and elsewhere, we have begun to explore a number of ways in which the work in computer modelling and brain modelling can give us an insight into cognitive mental activities. We have robots that can pick up eggs without breaking them. We understand how it is that a frog can detour around a barrier to get at a worm. We know how a computer, if asked questions about a relatively simple domain of knowledge, but asked in English rather than a programming language, could come up with the correct answer. So we see this

progress: we are well beyond the amoeba, perhaps we have already reached the worm. Some would say the worm is a long way from the human. Others of us say that we haven't been stumped yet, so let's keep going.

In this latter spirit, I would suggest that the notion of mechanism can help us understand our humanity. I said before that our understanding was pluralistic: if we are going to talk about our emotions, then for much of the discussion our normal everyday discourse will be fine; but if we are looking for an effective drug therapy, then the mechanistic underpinnings become crucial. Again, for many human problems, the language of sociology, of people in interaction in society, becomes crucial. So when I say that man is a machine, I am not trying to restrict us to some narrow form of mechanism, but rather to suggest that for a rational analysis of the human condition, an evolving notion of mechanism is what is required.

Such conceptual evolution has already occurred in Physics. It's almost 300 years since Newton published the Principia Mathematica, and at that stage we began to see the Universe as one vast machine. But in the 20th century, the definition of physical machine has undergone its revolutions -- the phenomena of relativity and the quantum mechanical change from determinism to probability amplitudes have changed our notion of mechanism. I don't see any reason why our increasing attempts to build better computers and to better understand ourselves won't in the same way synergize to grow yet new notions of machine.

To conclude this part of my talk, let me say dogmatically (I say "dogmatically" because Bruce Aune, in his Chancellor's Lecture last year, argued this view in detail) that there are notions of free will and social responsibility that are in no way incompatible with what I take to be the mechanistic view. What does it mean to have free will? I think it means that we can act in accord with our own complexities; not that we do something that is totally unrelated to our entire preceding life. As Bruce Aune noted, we would be shocked if we could not begin to understand the roots of our actions. Freud has helped bridge that gap by showing us how to take the occasionally inexplicable and try to dredge up forgotten (or repressed) experiences to better know how we behave. Some people have thought that quantum mechanics takes us from determinism to free will, but if we always jump at the toss of a die rather than thinking through what it is that we really want to do, we don't have free will. And so one can continue in this way, of trying to think through the complexities of our behavior, trying to see how this behavior grows out of our own self, seeing how that self grows out of both our biological background and our social experience to understand human responsibility within a mechanistic view of the human mind. I think we will be able to match this with a sufficiently complex view of mechanism.

The Computer in the Workplace

I have tried to suggest that the computer metaphor is not a static one. I don't say that we think in binary code, or that we cannot understand something unless it is fed through one ear on paper tape. Rather, we are beginning to build a vocabulary which

goes far beyond the normal vocabulary of mentalism, just as Freud had already begun to enrich our vocabulary for thinking about the mind.

With that I want to turn from the computer as metaphor and to think about the computer as a tool, and I want in particular to examine the nature of work. Some people, in defining freedom, have suggested that freedom from toil, freedom from work, was the goal of human life. But I think that, for most of us, work is part of the definition of our life so long as that work is "meaningful". Some people work because they know they will starve if they don't work, and that gives work meaning enough. Others of us, fat cats of the professoriate, like to think of our work not as a means of subsistence, but rather as a means of our development, etc. The point is that I think it really is part of being human not only to be a member of a family, to have friends, to be part of a society, but also to have some work that is meaningful for us, and to have some reasonable standard of physical well-being.

Marx, or at least the young Marx of the Grundrisse, observed that in feudal society people had fixed roles. They didn't have to question who they were: society was static, and God had decreed each person's status -- whether bondsman or lord, he had his duties. When the machine came, many of those roles which had hitherto been fixed for humans became jobs that could be done by the machine. And this had effects both good and bad: it was liberating -- suddenly people were not locked for life in the role that the fall of the hereditary dice had cast for them, but they could begin to explore new roles. But it also brought alienation because suddenly they were no longer possessed of a worth that all society recognized as a

farmer or a craftsman, but they were on the labor market and they were worth what they could get in competition against the machine. Suddenly people could no longer count on that sense of meaning.

In the 20th century the pace of change has picked up. It takes perhaps ten years for some forms of employment to become outmoded and outdated. What we have to ask ourselves as we try to be human in the computer age is perhaps not so specifically tied to computers or automation, but rather "How are we going to cope with massive change?" I claimed that we are (in some sense) machines, but are we machines to be thrown out on the dust heap once we no longer play our economically useful role, or are we rather special machines like a fine clock that is so beautifully engineered that you will keep it even when it stops keeping accurate time?

It is easy to look at the temper of our age in terms of the prevalence of machines, just as there are some people who say "The weather has been so bad lately because of the nuclear bomb," or "We wouldn't have had that snow if it hadn't been for the Falkland Islands." People are often possessed of a very simple view of causality, trying to find just one cause for each thing that ails them. It seems to me that if we look at the history of modern society we can see people displaced from their jobs because of machines, but we can also see people displaced from their jobs because of a certain economic system. Depending on who you are, you can see technology or capitalism or communism as the culprit. One of the things that we learn in cybernetics is the notion of mutual causality, that you can't just isolate one item and say "Here is where the chain of cause and effect starts, here is the cause and here is the effect, and that's it." So it is that we must come to

understand a complex interacting system in which the presence of the computer, and the presence of automation, is changing the idea of what people can achieve as a whole. There are things that we can do in the way of predicting the weather, of mass-producing goods which were hitherto available only for the wealthy and of letting everybody have a hifi and a stereo, that would not have been possible without the machine. But at the same time, many people who had a job which gave some meaning to their life are now without it. Some people show little sympathy, stating that if these people no longer have an economic role, they should go out and do something which "shows their initiative." However, I believe that we must redesign our social structures so that the benefits of technology 'trickle down' to all these people. But, of course, it won't just 'trickle down' -- it requires some coherent vision of social justice.

I cannot offer such a vision tonight. But what I want to say very strongly is that we are in a situation where the nature of the workplace means that some of us, with good connections or good education or just good luck, can benefit tremendously from the new opportunities of the machine, to enjoy privileges that only a few could have enjoyed in the past. Remember that tutor of Geoffrey Chaucer's who was rich indeed to own 50 books back in 1354. Tragically, we have seen an increasing level of unemployment amongst minorities, and especially amongst black youths. This is not something that responsible citizens can allow to continue. We have to understand, then, how we can provide opportunities for these people to enable them to make their contribution to society. In no small part, that has something to do with education.

Reconstructing Society

We have been talking about these problems as if they were really quite new, but they are not. The introduction of cotton spinning machines in Italy in the 1300's caused a great dislocation of the industrial base in medieval Europe. If one reads about the Industrial Revolution of 150 to 200 years ago, one is struck at how many of their societal problems are ones that, if only one changes the names of the machines, we recognize today. The 19th century came up with a great invention -- the labor union. Men recognized that the new machinery was concentrating economic power in the hands of the few. That wasn't new, because economic power had been in the hands of a hereditary aristocracy, but now it came into the hands of a mercantile aristocracy. And so what workers did was to band together to say that "You think that you should not pay us more than an equivalent machine would have to be paid in terms of power and maintenance. But we are humans. We need more than that. We cannot live on these wages." And so this great adaptive social invention, a response to the first round of mechanization in modern times, was the labor union. However, I fear that at the end of the 20th century the union is becoming maladaptive, because unions have become identified with large bodies of workers who themselves are identified in terms of some particular technology or trade. As that technology or trade changes drastically, then too often the work of the union is to turn back the clock to say "You must keep using this obsolete, outdated technology which gives our people employment."

I think we have to see a move towards a more adaptive society in which we still have special interest groups, advocates for the welfare of groups of workers who are skilled or unskilled, who are skilled for the future or only skilled for the past. But these groups must come to define their obligation to people not in terms of the demarcation dispute of holding to a static framework, but by helping people understand what is going on, so that they can adapt to change. Here, interestingly enough, is where I think the computer may be developing to a stage at which it can help us with these new social needs. Many of our concerns with the social or industrial machine are concerns with bigness, concerns that power has been concentrated in the hands of large companies or big government, with workers organizing in large groups to meet this. I would argue that we can now move into a pluralistic world where, because we have computer networks, because people can have access to large data bases no matter where they live, with sophisticated expert systems, knowledge networks and visual processing devices, we can begin to have people defining their skills and their trades in terms of much greater diversification. The effective unit of organization can be relatively small and thus restore more of the human scale.

I do not see technology as determinate, in the sense that technology must make us all better people, or technology is such that it must inevitably ruin our social world. In the hands of the KGB a data bank can be an invaluable tool of oppression. In a democracy, access to home computers hooked into a computer network of data banks can be part of creating a new level of informed citizenry. But even though computers can be used for good and for

bad. I still think we must face up to the fact that in this age of increasing automation (increasing, I remind you -- automation has been with us since the first irrigation ditch, the first water wheel), technology contributes to many social changes which are threatening to disenfranchise many of us, to deprive many people of their livelihood.

But again, I don't think we can see this in a uni-causal way, of the machine causing the social malaise. We might say that greed, need and power drive our economic structures, and it is the economic demands of big business that make a cheap machine more desirable than a well-fed person. And the computer extends the reach of this automation, so when we talk about being human in the computer age, we're not simply saying "Let's program our computers to do what we want them to do," nor are we throwing up our hands in despair and saying that the computer will dehumanize us all. We have to think in a multi-causal way about the way in which, as we come to better understand ourselves, as we come to better understand the tools that machines give us, as we better understand our interactions within society, we can use this knowledge to begin to address the problems of a changing society.

Using that knowledge means that we have to be far more self-conscious. Here's a current example: a few years ago there was an energy crisis. Petroleum was running out, and we were going to have to find alternate sources of energy. But then the price of gas went down 10 cents a gallon, and people forgot the energy crisis. Well, we cannot solve our pressing social problems with that sort of instant amnesia. We are going to have to learn to assimilate the complexities of our world, and (perhaps this isn't a

bad thing to say at a Chancellor's Lecture) we are going to have to be a uni-versity again. We are going to actually have to talk to people in other departments. We are going to have to not only talk about our little specializations, but to try to understand whether they mean anything in the larger context of society's problems.

We tend to think in very narrow terms, whether we're humanists or technologists. We tend to define a problematique in terms of something that we can publish in a paper within a year; we have all too little time to step back and really wrestle with these larger issues and try to understand how what we know can make sense outside our specialty. The Chancellor mentioned that the title of my Gifford Lectures was "The Construction of Reality". I don't want to develop the theme of those lectures tonight, but I do want to suggest that the range of topics we have discussed reminds us that we have a reality to construct, a social reality. Even though I've talked about man as a machine, even though I've talked about deterministic models helping us understand ourselves, I have also said that within that framework we can understand free will -- we are responsible machines. And as responsible machines who care for our children we must try to bring all our rationality to bear to ensure that as this computer age continues, hopefully long after the threat of nuclear war is a memory of a strange and aberrant past, we can still be truly human.

Thank you very much.