Future of Health Technology Symposium

Presentation by:
Marvin Minsky, PhD
Author, *The Society of Mind* and *The Emotion Machine*
Donner Professor, Media Arts and Sciences, MIT

This podcast is presented by The Centers for Disease Control and Prevention.
CDC – safer, healthier, people.

Marvin Minsky:
Well, there are lots of problems to solve, and I'm going to talk about just one of them, but here's a big list of things that bother people like global warming and epidemics and energy sources and education and terrorism and other things. And no one knows what to do about any of those. What we need is a government that has some technical people to find technical solutions to what are treated as political problems. But we've heard few references to one big problem that's sort of like global warming in that it should strike us or grow rapidly over the next 50 years, and I should have brought my laser pointer. Is there one?

Okay, one of the big disasters, along with putting carbon dioxide into the atmosphere is improving health so that people live longer, and I saw an estimate that the increase in longevity for the past half century has been about one year for every four, so people born in 2000 will live 12 years longer than people born in 1950. But I expect it will be quite a lot worse than that; by worse I mean better, and that's because as we understand more genes, we'll probably understand more about aging and dementia. We don't mind aging as long as we don't get demented and crippled and so forth. There ought to be a lot of progress there. No one knows much about aging, although there are some very good scientists who work on it in other animals. It's hard to do it with humans or alligators because they live so long. Mice are convenient and so forth because you can find out a lot about longevity in a small number of years. But one theory is that aging depends only on a small number of genes, and the argument for that is very simple; chimpanzees live twice as long as other mammals their size or more, and nobody knows why the primates are relatively long lived. And humans appear to live about twice as long as chimpanzees and we're only five million years from chimpanzees in evolutionary times unless you live in Kansas, or wherever it is. And as you know the genome of chimpanzees and humans differ by a very small amount. So it might be that most of the difference is the effect of less than a hundred genes, and presumably the improvement is not in making beneficial genes better, but in modifying genes that didn't work so well, making slight improvements in them. But that's all speculation. But anyway that means that in the next 10 or 20 years, we might understand the function of maybe half of those genes and make people live 200 years instead of 100. That's an advantage except for economics because the population of the earth is going to reach about ten billion in the fairly near future, but there's not enough food to feed 20 billion, so other countries will have to emulate China and reduce the number of children to one per family on the average. And if they don't, other countries will complain and various ways to enforce them will appear as the shortage of food advances around the globe,

accelerated by the conversion of things you could otherwise eat into gasoline for cars; the ethanol move is considered disastrous by many people because it takes almost as much energy to make ethanol as it saves because the amount of arable land on the planet, there's not a lot of extra land now to grow stuff, and as global warming proceeds, there will be a lot less because many arable areas will get droughts. Anyway I don't want to waste your time with all that gloomy stuff.

What it does mean is that there will be only one child per retired person pretty soon. So everybody who's 140 years old, there will be one 70 year old to do all the house work and industrial production and so forth. So it's a sort of funny crisis, but as far as I can see – well anyway, this is all to say that we're going to need robots and very intelligent ones just because there will be no one else to do the work. We have to get the productivity up. Social Security is in crisis, and many talked about we could solve that if people were more thrifty and saved up. But savings won't help, investments won't help, because there's no way that that money in banks could produce more food, see what I mean? It's not an economic problem in the ordinary sense because as soon as there is no labor force, of course automation will fill half the gap because of the non-intelligent robots who get better and better. But if you look in an automobile class, most of the fabrication of parts and most of the assembly of parts is done by industrial robots. Generally they're very clumsy, very little dexterity, but you don't need dexterity in mass production things because the trick in the factor is never let go of the part. Therefore, you put it into a jig or a cradle or something on the belt and the next robot can grab the thing without any feedback, and that's more or less how things work. So there hasn't been much advance in the smartness of robots since about 1980. The field that AI made spectacular progress between 1950 when computers first appeared, and 1980, and then it slowed down because what happened was that each specialist said, these robots aren't smart enough, and one specialist would say, maybe they need better logic; let's make a logic engine that can do more. And another specialist said maybe it's because they're not learning well from experience. Let's provide them with extremely sensitive statistical methods so that they can look at a lot of events and find the significant factors and learn that way. And other people say, no that's too limited, let's make something that imitates evolution, and a field called Generic Programming was developed, which was very dramatically successful for a few years. What it did was to make a computer program and give it a task, and if it didn't work, you'd make a thousand copies of that computer program with mutations, usually not very smart ones, and hope one of those thousand programs would work better. And if it did, then you'd copy it and make variations and so forth.

So Genetic Programming marched along and as it turns out, it was able to solve certain problems amazingly well. A researcher named Douglas Leonard in Austin applied a very smart version of genetic programming to some problems in designing computer circuits and created some new designs for generally simple things that nobody had discovered before. So there are few victories like that. But generally if you give a genetic algorithm some other kind of problem, it won't solve it, and if you give a neural network some particular problem, it won't solve it. If you look at the neural network literature which is enormous, there must be 50,000 people working in that area, and

every week somebody publishes a paper saying, look it solved this problem. You never see a paper saying, look it couldn't solve this problem. In genetic algorithms you never see a paper say it totally failed to make progress here. You never see a paper by logicians which say – well sometimes you do, they're different, but there's something wrong with my field in other words. In Psychics, if you have a theory and Slack or Serrin or some great laboratory runs an experiment and the theory isn't confirmed, that's a great thing, especially if it's someone else's theory. So Physics progresses by doing critical experiments and success is routine, and failure is the dramatic thing that wins you the Nobel Prize, like when… won the Nobel Prize for the violation of symmetry, parity. A very sensitive experiment, and all theories of physics had predicted these two experiments would give the same result because they were just mirror images but slightly different. But the theory doesn't work you in except in my field. But that's because they don't have support for basic research. I've got a lot of slides, but I'll summarize them because there's too many.

So what are we going to do about this shortage of labor? Who's going to do the housework, who is going to install the solar cells when they get efficient enough to be cost effective? Because tacking them on a hundred million roofs of homes is going to take a lot of work, and climbing up and cleaning them every few months. But the worst one of all right; looking from the point of view of CDC and NIH, and anybody who is actually alive is what are we going to do if our health fails and we can't afford thirty thousand dollars a day for the latest treatment. So how do your reduce the cost of future healthcare? Number one, and the cheapest of all, and almost completely out of practice because you can't make any money on it is preventing disease. I won't say anymore about it except it's a problem.

Another problem is that the health industry is very profitable, and if anybody makes a profit, somebody else has to pay for it. And I'd like to see – I know Hulin has a vast number of important patents, and I'd love to see them applied to industry perhaps, but I'd like to see the strength of the patent system in the medical area to become illegal or something. Of course the most obnoxious one was the wave of patenting jeans, which weren't invented but just found by systematic exploration. But anyway, and then everyone knows that the administrative costs of health care are enormous. I know many young people in my generation, lots of people went into medicine. Now they are afraid to because although they like the idea of helping people and discovering things and improving treatments, they don't like the idea of spending almost half their time on the telephone trying to get a hospital admission, reimbursement forms and some of them don't, they just hire three or four full-time assistants to fill out the forms for their practice. But that's an enormous cost.

Okay, then there's the global database, how do we evaluate drugs? People say, well we've got to patent the drugs because it costs us three hundred million dollars, and we've got to find a way around that. There are lots of simple things we could do like pay people to volunteer for studies, and of course you have to pay them a lot now because if they got a chronic illness from taking a drug that didn't work, it would cost a million dollars. But if we had universal health coverage, then the three hundred million dollars

you saved the pharmaceutical companies from the testing would produce maybe fifty million dollars in compensating these poor people. So lots of issues like that. If you have a global database, if people are required to make their database public, encrypt it as much as you want, and encryption is good, but if we had the medical history of everybody and more monitoring of what they ate and what they did, recording their vital signs and so forth, then when somebody got sick, you'd ask can I find ten thousand other people who had a similar illness by the ….. and their medical histories, that would cost about fifty cents if you use the sort of hardware that Google uses to find out such and such a film and such and such an actor.

So people talk about privacy issues, and we need a publicity campaign that says we think we can protect your privacy, and if we had a universal mandatory medical database, some of your rights might be violated once in a while, and we'd try to punish the violators, but you'll live five years longer; which is worth more to you? And what drives me up the wall is looking at the senators and issues like stem cell research. And senators are pretty old, and most of them are going to get heart disease or whatever your favorite disease is, within twenty years, and if they promoted stem cell research and other allegedly controversial things, they'd have the substantial chance of getting out of that. So you can think of them as heroes; they're willing to die for their beliefs. But they haven't been educated to realize that that's what they're doing, and if they knew I'd bet they'd change their votes. Enough politics.

Why don't we have smart machines yet? I'll show you a few slides showing how smart the artificial intelligent programs were in the 1960s, not so well known, a remarkable period, and why it stopped. And actually I just told you why it stopped; it stopped because around 1980, there were lots of AI programs that did different things pretty well, but there are lots of things that none of them could do at all. And almost every specialist in the field said, I think I know the solution, its better statistics, or its better reinforcement methods for neural networks or its better logic, or whatever. And a lot of them worked on low level mechanical robots. Nothing was learned from Roomba, for example, it's the most successful financial robotic thing in history. There are many millions of Roombas selling for several hundred dollars each. What other robots are there? Anyway, the point is they are not very smart. And why did they stop getting smart?

There are several reasons why AI research hasn't made much progress; the principal one is the idea that I want to find the single secret that will solve the problem, because that's what happened in physics. Physics sort of got to a good start in the time of Archimedes, which is three or four hundred B.C., and he almost discovered calculus, and it stopped for reasons I don't want to mention except it may have had to do with religion, and we ought to find a cure for that. I don't think picking, I don't think holding up passengers at airports is going to solve the problem. In fact air travel is already too safe to be cost affective. You think about it. The number of lives lost in waiting on line is larger than the number that would be lost if you had three or four terrorist-inspired crashes per year, which we never did except that one day several years ago. Now that

they've locked the cockpit door, I don't see why they had to take my little pliers away. If the airplane broke, nobody would be able to fix it.

It seems to me that the bug in the field of artificial intelligence has been each person's idea that there's one way to do this thing, whatever it may be; natural selection, logic, making mechanical robots that learn from experience with the real world and so forth. These are all popular. The robotics thing is great for high school. I love the first program developed by Dean Caman where thousand of high schools around the world have students competing to make robots that play a certain primitive form of basketball. But nothing's learned from this in general; it's good experience for kids. I don't see that it should be done by college students or graduate students because as far as I can see nothing has been learned from the experience of building physical robots. If you build a simulated robot in a simulated world, then when it doesn't work you can find out why. In the physical world it might be because there was too much noise or the sunlight came through a window or one of the students stepped on a cable and make it loose, which is the most common. What's the most common cause in the da Vinci type things? When I ran a laboratory, I discovered we had cables lying on the floor, and there's two kind of human beings; I'd like to – I don't know the origin of them, but some of them step over cables, and others stop with their heel on it and sort of do this, and I don't know of any psychologist has ever studied this phenomenon. And neither of them know they are doing it.

Why are humans so much more resourceful and reliable than robots, than intelligent robots? Stupid robots are very reliable. And the answer is that it's not like Newton and Maxwell. See, physics progressed by people looking at complicated phenomenon and saying, maybe there's just a few simple rules. And the great thing about the Galileo Newton collaboration, although I think Galileo died the year Newton was born was Galileo sort of recognized the formula for gravity, and Newton picked up where Archimedes left off. I'm ignoring what happened in Asia, there were also some developments in science in other countries. The Chinese have a good reputation for that but I haven't heard of anything substantial beyond fireworks and a few other things. But that's just my ignorance.

So why does physics work so well? Because Newton found three simple laws that explained almost all of mechanics; incredible. Not only that, but gravity. But that didn't explain electricity or magnetism. But then Maxwell invented four little laws that are slightly more complicated, and that explained electricity and magnetism. And Einstein, bless his heart, invented one law that replaced all four of Maxwell's because he discovered that magnetism is just electricity in motion or something. I'm sure you all know that theory, but I don't remember it clearly.

So psychologists, except for a couple people like Hull have tried to find laws of psychology, and we know a lot of laws like if you ask someone to memorize a list, a long list, then they remember the last one on the list best. That's called the recency effect. And they also remember the beginning of the list second best and I don't know the name for that law. There are hundreds of such laws, and none of them are true, but

they're often true. It's very complicated. And in general psychologists have never found anything like Newton's laws, but they kept looking for simple explanations, so the whole theory developed by Fred Skinner in the 1930s was that you learn from experience by being rewarded. To this day there are tens of thousands of psychologists and hundred of thousands of teachers who have learned from them that learning proceeds through reinforcement of success, whereas the most important learning proceeds through the analysis of failures, and that's hardly ever discussed. In fact there was a review of my new book by – done on the Amazon page by Robert Restak, who says Minsky has this wild speculation that we learn more from failure than from success. I'm really annoyed by that, but he's a Neurologist.

Okay, so what's the answer? We don't know by simple rules, because our brain is made of four hundred different kinds of computers. If you look at the index to a big neurology book, you'll find that there are all these brain centers, and you look back in the text and it says, the hippocampus is involved with memory. It doesn't say how it works or what it does. And you go all through the book, and except for the visual system which is fairly well understood, but not thoroughly, it says when you remove this part of the brain, the person loses such and such abilities. Ablation experiments, you take out a little part of the brain and maybe you can describe the resulting behavior by saying something is missing, but the something missing is usually described as an ordinary common sense word. We don't have words for what the brain really does, and that's what I'm going to talk about in a couple of minutes.

So what I think is that we shouldn't imitate the physicist, and look for the basic principal of learning or a vision or pattern, blah, blah, but we should ask for each phenomenon you see in a person's behavior, what are the ten principal ways the person does that, because I think evolution does the opposite of what physics does; evolution is unscrupulous. Anything that works a little better than something else gets selected and we accumulate these, and we end up with a lot of systems, none of which work very well. However, if you have 10 of them and some way of deciding which one is best right now, then you win, and so here's a nice slide which shows 12 different ways to tell how far something is. Like if I wanted this, I could just reach for it. But if I wanted that, I would have to take a step. One of them was one meter away, and one was two meters away. How do I know how far something is? Well there's 12 ways. Texture gradients. The nicest one is knowing. See there's a glass and I know its two meters away, because I know a glass is about the size of my hand, by outstretched hand. But when I outstretch my hand, that glass is half the size of my outstretched hand, so it must be twice as far and I know how far my arm is. This is just common sense. Some distance judgments are made by knowing about the object. Some are about the laws of optics, some are from the fact that you have two eyes. If things are moving at the same speed, then the closer one looks like its moving faster and blah, blah, blah. So there are just ways, and I'm sure you can think of a few more that we do one function, and that's why we hardly ever walk off the edge of a cliff. How many of you have walked off the edge of a cliff? Actually another part of your brain says don't even get near the edge of the cliff.

Anyway I have this book with a lot of theories in it and you can get it by going to my web page. In fact there's a previous draft of the entire book on my web page, and the only difference between the draft that's free and the one that you buy is about four thousand almost insignificant changes; and in a few cases making a theory a little bit better. So if you ask Google for Minsky, I might come up first. Or if not you'll get Richard Minsky, who claims to be the world's best book binder. I got a lot of my early ideas from Isaac so put him there. There's one thing I wanted to say before I waste all your time. There are very few people that are doing research on what I'm going to talk about which is, the diversity of the intellectual activities, and the importance of common sense knowledge. This is a strange phenomenon, but although there are tens of thousands of people who are applying artificial intelligence to industrial problems, you know, people making systems that help organize your production schedule. Although there are tens of thousands of people applying these different ideas like low level robotics, or statistics or logic, there are very few people doing full time research on the kinds of problems I'm talking about and that's because I think the United States is getting into very serious trouble.

When I was a young researcher, the universities were expanding, and there were many great basic research labs out in the industry. I got a job one summer at Bell Labs working with Claude Shannon, who was a great influence. And when I got to Bell Labs they said, you know we do basic research here. We have a sister company called Western Electric who makes gadgets; makes the telephone and the switches and everything having to do with the telephone company. And at Western Electric, the standard is every product should last for 40 years. So in our living room we have a black old dial phone, which must be from 1950, and it still works. Whereas if you have a $10 phone, it probably won't work in 40 years, but who knows. Anyway, so they said, you should only work on problems that will take about 40 years. Now if you get a job in the industry they'll say, we want a progress report every three months, that's a farsighted company. So when I was a beginner, there were Bell Labs, which had the longest horizon, because it was a monopoly, it didn't have any competition, so it could do the right thing. General Electric and Westinghouse had robots. Where was Ralph Moche. I think he was at West General Electric, and he made a four legged walking robot; it weighed three or four tons; it was actually a remote controlled one. You'd like down in its body and move your arms and legs and this enormous horse could walk across the room. So he was allowed to do it but he was afraid that other people couldn't control it without training. Looks like an elephant and he would walk over and he kicked a jeep across the room just by moving his foot. Wonderful gadget.

Anyway, the point is that the universities were expanding in the 60s, from the 50s to the 80s, but now they're contracting. And the professors are obeying the live longer every four years rule so they're not retiring. So if you have professors living longer and universities not expanding, then assistant professors can't get basic research jobs anywhere. And I think my first 20 PhD's all became professors except Tom Evans, who decided to be independent. And now only three of my last 20 PhD's have become professors. No place in America for basic research; you have to go to a company. NIH gave career fellowships which were five or ten years back in the 60s, and I bet there are

none of those. Anybody know? How long you give someone a research grant for? So that's a problem. But as I said, I think the big problem is that we need not particular methods for finding out how far away something is, what the brain also has is when some of these work and some of then don't, how do you know which ones to believe? And that's a tough question, because each of those methods have limitations and somehow you learn which of your methods have limitations.

So here is my very simple general theory of what we have to do to build a smart machine and the idea is that the machine has two parts; at any moment you can consider yourself to be facing a number of problems. Some of them are solved by very low level methods, like if I were thirsty I would keep talking while I'm going over here and getting a glass, a bottle of water, and wondering if I'm strong enough to get the cap off. So that's a problem, but I'm also trying to decide what to say next and whether to look at the audience or not. It's a very important decision to make because if somebody smiles, then you're liable to stick to that topic instead of – because we're all just people and we're influenced by what other people think, even when we shouldn't be.

So what kinds of problems are there? Well I've never seen a good classification, although there are a couple of references in the Emotion Machine book about some attempts to classify these. There are visual problems and language problems and all the sorts of things that Howard Gardner talks about in his couple of books on multiple intelligences. In the popular view some people are smart, and some people are dumb, and there's a thing called IQ, which is a numerical measure of how smart you are, and that theory is almost worthless except that it has a lot of economic value because if you pick people with a higher IQ, then there's a better chance that they'll understand what you said. But as Howard Gardner - what's in the IQ test? The IQ test is a fusion of about ten things. How good are you at mathematics? How good are you at language? How good are you at social things, I forget what's in it. And Howard Gardner's theory is that every person has about eight on nine, well I think it's eight dimensions of intelligence and somebody may be very good at music or symbolic things, and another person may be very good at graphical or verbal things and so forth.

So when any person faces a problem they have to decide is this the kind of problem that I'm good at solving? Should I give it high priority, should I waste my time on it? And that's what goes in the left hand box. The mental critics that size up the present situation and then the ways to think, which is should you think about it geometrically or tactically, or linguistically, or logically or arithmetically or what. And every one of us has those capabilities and we each are good at them to the extent that we can solve larger or smaller classes of problems with each of them. I'm good at some kinds of music, and sometimes when I have a problem, I represent it in terms of several different processes because when I was a baby, I got good at counter point, so I suspect I have just some non standard way of thinking that isn't musical, but it uses the fact that I practiced very hard trying to think of two different tunes at the same time. Probably spent a year on that without knowing it.

So the important thing then is what are ways to think? One way is to – if a problem seems familiar, then maybe you don't have to solve it because if you think about it a while you'll remember another problem that's very similar and you'll remember how you solved that one, and that's called reasoning by analogy. No computer programs yet do that very well, although there is a subfield of artificial intelligence called case based reasoning, and that is a significant amount of research on that; some at the university in Atlanta. So anyway the idea is you have these critics which try to say what kind of problems we face and what kind of way to think. The critic has learned which ways of solving problems are good. So if a problem seems difficult and nothing else works, maybe you can split it in five parts. How will I get to the airport? Well I have to take a taxi or I have to find my car and start the car, first I have to go out of the house. So for almost any routine job you already know the script which consists of five or ten little jobs. You don't have to think much about it. The last one of the list is if you can't solve a problem find somebody who can and pay them or beg them into working on it for you.

In 1961 a student at MIT wrote a program that was able to do as well as MIT freshmen in the calculus course for integration problems. He couldn't solve physics problems and things like that because he didn't know enough. But it was quite remarkable because mathematicians did not have a theory of how to do this. So integration was done by trial and error, and there were books full of different formulas for integrating. And after Slagel's success with this program with just a few dozen rules managed to solve a large class of problems. So mathematicians began to work on it and by the end of the 1960s a couple of mathematicians found the general solution to integrate any symbolic function that has an integral. So that led to a lot of progress. In the middle 1960s Tom Evans, who I mentioned wrote a program that did as well as high school students on geometric analogy tests. Daniel Bobroe wrote a program that could solve word problems in high school algebra and in some ways this was the most impressive of that because nobody knew much about how to make computers understand language. In the case of algebra, even though every sentence is in fact some equation and he found about a hundred tricks so that it could not always but often figure out the right equation. But even today no computer has the kind of knowledge that every four or five year old has like you can pull things with a string, but not push them, or it annoys people when you interrupt them, or it's hard to stay awake when you're bored; of course that didn't happen today. No one else can tell what you're thinking. One of our kids once came up about the age of three and said, can you tell what I'm thinking? I said no. She said, oh I thought so, and went off very happily. So everybody knows that but you won't find it written anywhere because everybody knows it, and the editor would take it out if you put it in.

Every child who is looking at something is thinking about it in several different ways. So the little girl is saying, what would happen if I pulled out that bottom block? Should I help him or injure him? That's social. How would he react? Emotional intelligence. I forgot where I left the arch-shaped block. That's self reflective cognitive psychology. Can I reach it from here? I talked about that. How would it feel to grab three blocks at once? Some parts of your mind are able to reconstruct or synthesize something like how would it feel. Maybe you'd feel three points here, here and here. So that whenever you see

something you are thinking about it in many different ways and this table is pretty close to Howard Gardner's table of different intelligences. And the trouble is that no artificial intelligence program today does anything like this, and that's why they haven't made much progress in all these years. Instead each specialist has been trying to do it, say I've got the answer, I'm like Faraday or Peltia or Newton or someone and I have the solution to intelligence. Jeff Hawkins has an interesting book called, On Intelligence, which is immensely popular with people who admire millionaires. And he has an idea that you remember certain kinds of sequences. It's probably an excellent idea and one of ten that you need in order to make a resourceful machine, but it's typical of somebody says, I've got the secret.

Freud is my hero because he had the idea that you think in many ways and they don't always agree and conflict. Psychologists in cognitive psychology never even mentioned him. So I think I'll – here's a list of ways to think, and I just mentioned some of them. What are the kinds of critics? Well some critics are – I reached for that but I couldn't grab it, something is wrong with my range finding. That's a sort of prosaic physical thing. But another thing is, I tried the same idea four times, and it didn't work. Maybe I should change my management of my own thoughts. So you see, thinking occurs on many different levels and sometimes the problem is that you're not controlling your high level mental search for methods properly, so you have very high level critics and low level critics. Anyway, the book is full of stuff like that.

There's a thing about Microsoft Word which is that if you have something which finding a piece of text and you left find Window open, then the text won't show pictures. Now there are four hundred people who work on Microsoft Office at Microsoft, and that's why bugs like this remain year after year, because nobody knows what the others are doing.

Okay, this is more details about things inside that book. I think I'll stop and see if there's time for a couple of questions.

To sum it up, we need intelligent machines. We're having trouble getting them. I think there are some high level theories that are in place that might give us machines with general intelligence and common sense. I think my new book tells some ideas about how to do it. There's a guy in Birmingham, England named Aaron Sloman who has similar ideas. There's a guy in Texas named Douglas Lenatt who might be on the track. However none of us can get support for basic research in this area because the United States has gone intellectually broke and we don't have long term support for basic research really in scarcely any field. I think it's a crisis that nobody understands. In the 1950s and 60s, and 70s, this was the leading country in the world. Now I suspect China soon will be unless we get some intelligent people into Congress and White House. There's lots of money for lots of things but research is incredibly cheap, and we're not supporting it. Thank you very much.