Interview of Edward Feigenbaum
2012 Fellow

Interviewed by:
Jon Plutte

Edited by:
Dag Spicer

Recorded: March 20, 2012
Mountain View, California

CHM Reference number: X6409.2012

Plutte: This is Jon Plutte with Ed Feigenbaum at the Computer History Museum. It's March 20, 2012. Thank you for joining us. I'm going to start off with some questions about-- I'd actually like to start-- it's a little bit reversed, but what got you interested in artificial intelligence? How did you start your career?

Feigenbaum: Stop for a minute. Let me just get that answer organized. It's actually a great historical story, which I think Kirsten has heard again and again.

Plutte: So, if you could look at me, it'd be great. Go ahead.

Feigenbaum: I first got in-- <coughs> excuse me. I first got interested in computers as a kid, when I didn't know anything about the kind of computers that we have today, but my father was an accountant and he would do work at night, and he'd bring home large mechanical calculators of the kind that we have here in the Computer History Museum. And, being a nerdy kid, I became a master of those calculators and I loved to show them to my friends and I showed off-- I mean, those friends were great at football and tennis. I was great at calculators, but then I got to Carnegie Tech, what is now called Carnegie Mellon University, and during my undergraduate time there as an engineering student, I happened to come across one of the true geniuses of the 20th century, Herbert Simon, who was in the process of co-inventing what we now call Artificial Intelligence, along with Alan Newell there and of course John McCarthy and Marvin Minsky at MIT. I was taking a class as a senior from Herb Simon called "Mathematical Models in the Social Sciences," and he walked in after the Christmas holiday and he said, "Over the Christmas break, Alan Newell and I invented a thinking machine." Well, that really blew our minds, a six-person seminar. We could ask questions: "What do you mean by a machine? What do you mean by a thinking machine?" And then he handed us a manual of the IBM 701. [a mainframe computer] Carnegie Tech didn't have a computer at that time, and I took it home and read it from cover to cover. I remember staying up all night and by morning I was born again. I was going to go into this field, and I got a job at IBM in the summer, learned how to program. I went back to Carnegie to be one of Herb Simon's graduate students and then did my thesis in a part of artificial intelligence that we now call computer simulation of human cognitive processes, the psychology end of AI. Later on, I switched to the engineering side of AI.

Plutte: What is the goal of artificial intelligence research?

Feigenbaum: In the early days, artificial intelligence research broke up into largely two subfields with somewhat different goals, although they worked together. One was the goal of modeling the information processes of human thought, giving precise definition to exactly what was going on in the mind during cognition and perception, and when I say modeling, I mean that we were trying to model as one would think that physicists model, where we wanted precise, verifiable models. Indeed, I first began to work on that for my Ph.D. thesis with Herb Simon and then continued on for several years after that working in that area.

The other part of artificial intelligence is the engineering side of the discipline that wants to build the smartest possible machines, whether or not those programs-- by “machines” I mean “software”-- whether or not that software is a good model of human thought. The spirit is the same as with aviation. You want to get a jet plane that flies the highest and the fastest, not necessarily [one that] flaps its wings. In recent years, we’ve had an augmentation of those goals. Around 1980, what became a major goal of artificial intelligence was the goal of machine learning, not just the performance of cognitive acts as well as people but learning from sensory signals, as well as people do and even better. And, finally, artificial intelligence moved on to the goal of augmenting human intellect, not simply matching it or replacing it, but augmenting it as an amplifier of human endeavor, you might say continuing the evolution of the mind, except continuing it in software and silicon.

<crew talk>

Plutte: Just to follow up on that a little bit, there’s a lot of people who’re afraid that artificial intelligence is trying to replace human thinking, what do you say to somebody who believes that?

Feigenbaum: Will artificial intelligence replace human thinking? Well, first of all, it’s unlikely in any reasonably short term, because the amount of commonsense memory and knowledge of the world that humans have so far outstrips what we’ve been able to accomplish with artificial intelligence that at best these programs will be adjuncts to human thought, but in certain areas, artificial intelligence programs will be replacing human thinking, in areas where human thinking is either too slow or has certain defects.

People are not very good, for example, at combinatorics. Anything involving combinatorial manipulation, they’re unsystematic. So, it would be good to have a machine replace a person. People don’t exploit knowledge very systematically. They grab knowledge out of their memory in a rather haphazard and unstructured way, so programs being systematic in certain areas can far outperform the abilities of human beings. Now, we’re already in a situation in which some artificial intelligence programs are replacing people in critical situations where we may feel uncomfortable: for example, AI in robotics on the battlefield, AI in managing one or more drones in drone strikes. So, AI is capable of doing tremendous good. It’s capable of making superb-- sorry. It will be capable of making superb medical diagnoses from vast amounts of evidence in the medical literature, but it will also be able to produce complex strategic and tactical plans in military warfare, as well as equally successful military defense strategies.

Plutte: Can you explain for me the difference between inference and knowledge-based approaches to AI?

Feigenbaum: Let me..

Plutte: I’m sorry. One second.

Feigenbaum: Hang on.

Plutte: Are you okay?

Feigenbaum: Let me go back to that other question. If people are concerned about robots replacing human labor, they have a very good reason to be concerned. Robotics is advancing extremely fast, and much of the endeavor that we're asking robots to do is not the most intellectual kind of endeavor. It's usually an endeavor that involves manipulation of hands and arms and objects in space, and that can be done very well by machines. It can be done rapidly, precisely with heavy objects, light objects, precision tooling and so on. So, we will, just as we have seen jobs move because of globalization, we will also see jobs move even within the United States because of robotization, and there's really nothing we can do to stop that. That's a technological given, the way we're heading, but we can prepare for it. Now you asked another question.

Plutte: I have a follow-up, but I'll do that in a while. Explain the difference between the inference and knowledge-based approach to AI.

Feigenbaum: Inference is about reasoning. Inference is something that you learn in a college class on logic, which you may have taken in the philosophy department or you may have taken it in the math department, where it's called mathematical logic. It's all about the intricate step-by-step allowed rules of reasoning that give you guaranteed, proven and accurate results, essentially mathematical truth, results that accord with mathematical truth.

The knowledge base of an artificial intelligence program consists of the symbols that represent what it is you are reasoning about. Those are called the domains. The world consists of enormous numbers of domains of knowledge. Almost everything is its own domain of knowledge and has specialized properties and attributes and values of those attributes and collection of things that we call situation-action rules. If I'm in this situation, then I take that action; or I [can] believe this other set of symbols that can be used in the inference step.

Now, it turns out that by experiment we discovered [by] approximately 1968 what I originally called the "knowledge is power" hypothesis, but it was so validated that it became known as the "knowledge principle" -- that the quality of the behavior of these programs that we write is almost entirely a function of what's in the knowledge base and not what's in the inference engine, so all that stuff that you learned in that math course about 20th-century mathematical logic or even 19th-century from Boole onward? Well, turns out all you really need is a little Aristotle and maybe augmented-- not maybe, certainly augmented with Bayes' theorem for probabilistic reasoning. So if you have Aristotle and Bayes in the box, you have enough reasoning, but you can never have enough knowledge of a domain, because knowledge always improves the result.

Plutte: And this may be too broad a question. How does that specifically relate to your work? You talked about that a little bit, so how did you use these two concepts to advance your work?

Feigenbaum: In the early 1960s I decided to move from the psychology side of artificial intelligence, the computer simulation of human thought, to the engineering side. It was just an inner feeling I had that I really did want to work on the problem of the smartest possible machine, what the great statistician I.J. Good in 1956 called the "ultra-intelligent computer." Well, that's a very big concept and you have to work on a piece of it, and the piece that I chose was not puzzle-solving, not deduction, which is most of what AI had been working on, but induction, the induction of hypotheses from data or even more grandly the induction of theories from hypotheses and data. And even that's too big a problem to work on, so I chose to model that as scientists would do it, scientific induction, because in my view, scientists were professionals--they were experts--in induction. That was their business, developing hypotheses and theory from data.

Well, it happened that I was extremely fortunate to recruit as a colleague a legendary scientist at Stanford, Joshua Lederberg, a Nobel prizewinner in genetics, and he was interested in exactly the same problem, coming at it from the point of view of a scientist who wanted to see what computers could do, a reflective scientist who wanted to see what computers could do to automate scientific work. So, he and I and a small group of young scientists developed a program called the DENDRAL program, which formed hypotheses in a small area of physical chemistry called mass spectrometry. We were fortunate there to recruit another genius, Carl Djerassi, head of the Stanford Mass Spectrometry Lab but also the inventor of the birth control pill, so we had two geniuses working on this project, and that gave rise to the [expert system]. Then from that we developed applications in medicine and engineering and even military applications.

Now, that gave rise to a big picture of what it is we were doing. We were using some inference processes but we were mostly extracting knowledge from the heads of experts in an area in the quest to model the knowledge of those experts and put it to work. Almost all of the work came in the extraction and modeling of that knowledge. The extraction part was called knowledge acquisition. We hoped it would eventually be done by machine. In fact, we did a major piece of work in trying to develop a piece of the theory of mass spectrometry automatically, and it happened. We actually were able to publish a brand-new piece of physical chemistry in the *Journal of the American Chemical Society* done by a program, although we didn't give the program any credit; we just put it in a footnote. So, that's where we got the idea of the knowledge principle, that almost all the power of the program came from what it knew. The reasoning processes were extraordinarily simple. We didn't have to dive into 20th-century mathematical logic and pull out fancy things.

Plutte: Can you talk about the early enthusiasm for thinking machines and how has this thinking evolved?

Feigenbaum: Hold on a second. I want to see how I started my answer for that. Let's talk about the early enthusiasm for artificial intelligence and the later stages of disillusionment and re-enthusiasm about artificial intelligence.

It's often said that we overestimate what can be done in the short term and we underestimate what can be done in the long term. The Computer History Museum contains wonderful examples of that in almost every exhibit. Well, the idea of modeling human thought was a breakthrough, a paradigm breakthrough in the early '50s pioneered by, of all people, Alan Turing, whose [birth] centenary we are honoring this year. And it was carried on by people like Newell and Simon and others of the time, Herb Gelernter and some students of Marvin Minsky into some spectacular early successes. For example, the very first artificial intelligence program to follow on from Turing's great 1950 vision was the program that I had mentioned, which was- I should scratch that- was the follow-on to Turing's great vision of the 1950s was the very first artificial intelligence program, the Logic Theorist program, whose task it was to prove theorems in Chapter 2 of the great monument of 20th-century mathematical logic, Whitehead and Russell's "Principia Mathematica," and it proved all the theorems in Chapter 2, but one of the proofs, Theorem 2.85, was done in two or three lines compared to two pages that it took Whitehead and Russell to prove that theorem.

Simon sent the proof to Russell and Russell was mightily impressed in written correspondence with Simon. So, there were spectacular successes. Herb Gelernter's program the next year, 1957, proved theorems in high school plane geometry and took the New York State Regents exam in plane geometry. It was the only taker of the exam in-- sorry, in 1958. It was the only taker of the exam in 1958 that got a perfect score. These early successes led to a phase of overprediction, which seemed actually to be validated by the arrival of expert systems on the scene. Here were programs that were performing at or better than the level of experts, of Ph.D.s in mass spectrometry, for example, or experts in medical diagnosis.

Eventually, this line of work hit a wall. It hit the wall of brittleness, when you go outside the area of specialization, the domain. The programs don't work right, because-- remember the knowledge principle. They don't know enough about the rest of the world, so they break. But people have vast amount of common sense, with which they can paper over these breaks in the reasoning process. People use language and exhibit fluid language skills, which computer programs were not able to do, especially at that time. So, there came the inevitable period of disillusionment in the public's eye and in the eye of the media, but quietly AI kept making progress and recently has achieved very spectacular results, again, that have captured the public imagination. I'm thinking of the autonomous driving vehicles of the Stanford and Carnegie Mellon people who won the DARPA Grand Challenge prizes for driving over the desert and in urban settings. I'm thinking of IBM's Watson program, one of the most spectacular performances ever of an AI program. I'm thinking of the Siri program, which [was] developed out of a DARPA-sponsored research program integrated by SRI and used by Apple in its most recent release of its iPhone. These are dramatic steps forward in the use of AI. So now the overprediction is starting again, and we're probably overpredicting what's going to happen in 2030 and we're probably underpredicting what's going to happen in 2050.

Plutte: I'm just intrigued by the concept of human common sense papering over the missing pieces or something like that. Could you just talk a little bit about your thoughts on what that is, and is that something we can replicate, or is somebody working on replicating that, or is that a side question? I don't have an answer off of that.

Feigenbaum: The big breakthrough in the use of commonsense knowledge in artificial intelligence came simply with the realization that it had to be done. It wasn't a simple job. It was an excruciating, painful, long-term job of codifying millions and millions of statements about ordinary human activity, and there was no one who wanted that job except one person, a former student of mine, Douglas Lenat, who did it at a research organization in Texas, and when that research organization folded he set up a corporation called Cycorp to continue to do that work, which he is doing at the present time. Cyc knows five million things currently of ordinary commonsense knowledge, not only about day-to-day activities but also commonsense knowledge about different areas of work, like, for example, if you're a doctor there are some things that you know that are specializations of your field, but there's other things you know which are simply common sense about people and patients and how they act. And Lenat and his sponsors are just- his sponsors are paying him to codify vast amounts of commonsense knowledge, which can be used to bridge between those parts of the knowledge base that are well developed and those parts for which there is either nothing or only ill-developed parts.

Plutte: Very interesting.

Feigenbaum: I might say that Lenat's work is in contrast to much work in artificial intelligence which is on the theoretical side of representing commonsense knowledge, so here's an important thing to know: that the first AI conference was attended by people who had experimentally developed programs to achieve significant goals. I mentioned the Logic Theorist program of Newell and Simon, but also by John McCarthy, whose paper in 1958 was a significant landmark. It was called "Programs with Common Sense." And McCarthy's theories over the years developed the theoretical basis for commonsense reasoning, but it didn't produce concrete programmed artifacts to do that. It was Lenat and people using the experimental method that developed the actual programs that took us significantly down the road toward realizing that dream. Let me also say that artificial intelligence is mostly an experimental discipline. We're not as mature as physics in which we can support both a vibrant experimental organization and a vibrant theoretical organization. We are mostly experimental. Theory is not a strong part of the work in artificial intelligence.

Plutte: How would you describe the study and development of expert systems to the average person?

Feigenbaum: An expert system is an artificial intelligence program that is aimed at doing intelligent thought. It's an artificial intelligence program that seeks to perform intelligent tasks in specific domains of expertise, where we notice and respect expertise. Almost all fields of human endeavor are fields in which bodies of expertise exist. The expert system builder wants to get that expert knowledge out of the heads of the practitioners and model it as symbolic structures that we put in the memory of a computer and we turn over to the inference processes to solve problems. So, an expert system is really a model of an expert's way of thinking and an expert's knowledge of a field. The people who do that kind of work, the work of extracting knowledge from experts in a very careful and systematic way, and then representing it for a machine are called knowledge engineers. And they are the people who build expert systems.

Plutte: Can you give me a very common example of an expert system t—one that we encounter daily? Can you talk about it a little bit?

Feigenbaum: Let me give you a couple of examples of expert systems that we might use [perhaps even] every day. When Apple's Siri came out, it was -- in fact, even when Siri was in an early stage of alpha testing, you would notice that Siri knew some things very well. And the rest of the world it didn't know much about. It knew about restaurants and it knew about hotels. And it knew about travel planning and things like that. But if you asked it anything else, it would just say, "Sorry, I really don't know about that." And now it says so in a very pleasant and most cheerful female voice.

Later, when Siri was incorporated into the Apple iPhone, people noticed that [they] changed its areas of knowledge. Apple decided it ought to know other things besides or maybe instead of the hotels. Maybe hotels were not the source of the greatest interest to phone users. But what was of interest was how to get the iPhone to do certain things for you. That was a real interest. So Siri was able to do that, send a message, send a message to a certain person about this. That's one example of an expert system driven by a beautiful, natural language front end where you could actually see the shift in knowledge bases over time. Another example comes from Microsoft Research. It's not widely in use but it could be. It's part of the Bing system for mapping and highway travel. The system developed by the artificial intelligence lab at Microsoft Research includes an immense amount of knowledge about specific-- 72 specific areas of the United States. Historical traffic flows, so that if anything happens on the road that you're on, particularly a freeway, it will immediately tell you where to get off that road and route you through city streets that, for that time of day, and both weather conditions and all kinds of other contextual things, would be the best for you and how to get back on the freeway on what exit ramp. Now, there's an expert system that you can just call up from let's say your iPad in the car and watch it happen. It'll tell you what to do.

Plutte: What is the purpose of developing an intelligent machine?

Feigenbaum: One might as well ask what's the purpose of developing a computer in the first place. Computers were developed to do complex, high-speed arithmetic. They were developed during World War II and then in the post World War II period, particularly by Von Neumann working at the Institute for

Advanced Study at Princeton, to do what we later came to know as H-bomb calculations. Well, artificial intelligence programs were born with the idea of doing symbolic computation, the symbolic manipulations that we call thinking for complex problems. And maybe we could accomplish it at levels of difficulty that people couldn't handle. We also wanted to develop artificial intelligence programs because we believed that thinking was no more than information processing in a complex program called the mind, which ran on a complex machine called the brain. And by modeling those information processes, we'd be able to get a better handle than psychology ever had before on [how] things worked in the mind. And in fact, it caused a revolution in psychology. It's called the cognitive science revolution. One might ask what was the purpose of genetic mutations from the chimpanzee to the human brain. The human brain is the most highly evolved mechanism in all of evolution. With the human brain, we have been able to invent the world. We are the masters of the artificial world. We created all the artifacts. We are now becoming the master of the natural world. And yet, there is very little difference between a chimpanzee's genes and human genes. Now, what was the purpose? Man's brain is not the -- humankind's brain is not the end of the evolution of the brain. And we're in the process of evolving thinking capacity not biologically but artifactually with AI programs.

Plutte: And I know we sort of touched on this in a lot of different ways. But the direct question is how would this be helpful to humanity?

Feigenbaum: Well, humankind is a tool builder, has always been a tool builder. This was the latest -- computers and AI and other activities of the information revolution are the latest tools that humankind has invented. And as we know, the tools can be invented for better or for worse. The tools of -- one of the inflection points of human civilization was the Industrial Revolution, the invention of, for example, the steam engine and things associated with it. And yet, the Industrial Revolution led directly to World War I and World War II, vicious and ugly uses of industrialization. It's the same with artificial intelligence. It could be used to now a little bit and in the future very much to do exquisite and complex medical diagnoses of a kind that we never had available for us especially in the light of enormous amounts of new information that can be brought to bear on a medical diagnosis, for example, genetic information and the interaction between genes and environment. Or it could be used to plan immensely complex military operations like the world has never seen or complex military defenses like the world has never seen. It's our choice.

Plutte: And again, I know we talked about this a little bit but it'd be interesting to ask this question in the context of everything we've talked up to this point. So what made you interested in expert systems? It seems like it can relate to that previous answer.

Feigenbaum: Hang on, let me go back to the-- right now, mankind is faced with a tremendous opportunity, the opportunity to use hundreds of billions of pages of information available on the internet. But we lack the cognitive skills and we lack the time. We lack the persistence and the energy to examine [it all]. It's totally out of our capability to examine hundreds of billions of pages. But suppose we could use an artificial intelligence program for information extraction, for information ingestion, for information summarization that we could-- suppose there were a very smart program behind a Google search or a Bing search. Just think of the power that that would be to all of us every day. We wouldn't just have to ask questions that involve keywords. We could ask questions that involve concepts. That would be of tremendous use to mankind. Sorry, now ask the question again. I don't think your question made much sense but ask it again.

Plutte: Yeah. No, that's why I was going to go to skip to the next one because I think you were starting to touch onto this, which-- and I think you may have just talked about this but what the next big challenge for AI research is going to be.

Feigenbaum: <coughs> Excuse me.

Plutte: Do you have any water?

<crew talk>

Plutte: The other question was about your interest in expert systems. It seems like you could kind of tie those pieces together. You could just leave this on your chair if you'd like.

Feigenbaum: I think the other question you asked is better.

Plutte: The one I just asked. Okay, if you chose to do that.

Feigenbaum: Because I think it's-- the other one is redundant. I've said it five times in different ways.

Plutte: Okay. What are the challenges for artificial intelligence going forward?

Feigenbaum: Obviously, there are application challenges, moving step-by-step toward this ultra intelligent machine. In the early history of AI, the first 25 years, the research activity was mostly top down driven by high level symbolic reasoning activity, exemplified by chess playing, proving theorems and logic, expert systems. That was the kind of activity that dominated. In the last 25 years, the field shifted around completely or almost completely to a bottom-up view, recognizing patterns and learning things from watching low level signals and putting them together into higher lever hypotheses. We call the whole spectrum the "symbol to signal spectrum." And for a while we were working only on the symbol part of the spectrum. And for the last 25 years, we've been working on the signal part of the spectrum. But these two are beginning to merge. It's becoming important for example, for the autonomous driving vehicles to not only interpret what their radar and lidar signals mean but also to have a higher level of understanding about the situation in which the car is placed, the streets, the people, the traffic lights, the stop signs, the rain. So the next 10 years will see a major merging of these two streams of activity into a stream called integration. People are now building integrated systems, encompassing both, high level reasoning and lower level signal processing or signal understanding. You can see it in robotics, for example. It's not enough to have a robot that just walks around the floor and knows how to manage its mobility but doesn't know anything about what the world wants it to do. It would be great to have robots with knowledge of the world and some reasoning ability, to be able to accomplish real tasks. And that's the job of the next 10 years, integrated systems. I'm really a big fan of integrated systems because it is a way of guaranteeing that AI remains on the track of the experimental method. As I mentioned, AI has learned almost everything it's learned by the experimental method. And building integrated systems is a way to ensure that the experimental method will continue. It's system building and that's experimental.

Plutte: Okay. What do you think is the next big challenge for the technology industry in general?

Feigenbaum: My view of the biggest challenge that all of us in information technology face in the next 10 years is communication technology. We have been making major strides in every part of information technology except communication technology. The applications people, the software people, the computer processing people are way ahead of the communications people. I would like to see fiber to every house. I happen to be lucky enough to have, courtesy of a Google experiment at Stanford, fiber coming to my home. It's an absolute game changer. It changes the nature of my use of the Cloud, of the Internet. We're in a desperate situation as it regards communication to mobile devices. We have these absolutely silly things about limiting the bandwidth to individual devices for 10's of dollars a month for a trivial amount of data, 150 megabytes of data or 250 megabytes of data. This is totally ridiculous. And it's partly the communication-technology industry people but also partly the inertia in the proper use of the spectrum that the government for so many years has controlled that is used by so many different parts of our society, including television. We have to overcome that inertia. We have to have some more breakthroughs. There will be no Cloud revolution without a revolution in communications technology.

Plutte: This is a slight digression. But we've been talking about here how in many ways, third world countries are ahead of us because they never developed the infrastructure that we have and they have to develop new like cell phone infrastructure. Do you have any thoughts on that? Maybe they're leapfrogging over where we are to the next big thing.

Feigenbaum: When you have a major infrastructure and a mature corporate environment, you automatically have tremendous amounts of inertia. It happens at all times. It happened when radio towers were going to replace copper as a way of transmitting phone calls around the United States. But companies had tremendous investments in the copper and the trenches in which the copper was laid. So that put inertia into the process. They were the -- companies like AT&T thought of themselves as switching companies. So when [pioneers of] the early Internet, particularly Bob Kahn and Vint Cerf, went around and talked to people like AT&T, the answer was: “Communication is switching. It’s not packets.” There’s always inertia. When you’re starting afresh in a new territory like Africa or after some major disaster, you can overcome that inertia simply by leapfrogging it. And that gives a certain advantage. On the other hand, the highly developed corporate structures and highly developed technical environments give us other advantages in capital and scope.

Plutte: All right, thanks. What advice would you give to a young person just graduating from high school now?

Feigenbaum: Let me talk first about the young people in high school who are heading for college. My first piece of advice is when you get to college, spread your wings as widely as possible. You’re there to become a citizen of the world and a citizen of the world of knowledge. There’s plenty of time for specialization.

Secondly, look around you and find among the otherwise excellent faculty the few geniuses that are on that faculty and align yourself with those geniuses. Find a way to get those geniuses to allow you to do research on some of their projects. You cannot imagine what you will learn from those geniuses. The unarticulated knowledge that these people have of how to ask the right questions and how to cut through all the brush of a complex search for an answer and how to slice through it in a simple and clean way, that’s knowledge that is not knowable any other way. It’ll take you years if ever for to develop that yourself.

For people who are not going on to college, let me just say that in every field of endeavor, in every person, there is an expert. There is expertise in every person and in every field. My favorite auto mechanic is really an expert at what he does. That’s why I go to him. So learn as much as you can. Spend what AI people once called 10 years of apprenticeship, learning the necessary knowledge to be an expert. Recently, some writers and psychologists have called it -- have converted the 10 years into 10,000 hours. Spend the 10,000 hours to become an apprentice, to learn what you have to learn to be an expert. In a knowledge-based society, expertise is highly valued. And finally, my advice to both of those groups is knowledge has a very short half-life. We are moving very rapidly. What you know will become obsolete very fast. So you have to be a continual learner. You have to dedicate yourself to education all the time, continuously for your whole life until you are no more.

Plutte: Great. One final question. What does it mean to you to become a CHM fellow?

Feigenbaum: Now, that’s where I need you to hold it up there.

Plutte: I have it right here.

Feigenbaum: The mission of the Computer History Museum is to preserve and present for posterity the artifacts and stories of the information age.

Plutte: I'm sorry, that was me.

<crew talk>

Plutte: All right, let's try it again. Here you go.

Feigenbaum: The mission of the Computer History Museum is to preserve and present for posterity the artifacts and stories of the information age. I'm very gratified that my story was selected-- got to start over again.

Feigenbaum: The mission statement of the Computer History Museum starts out, "to preserve and present for posterity the artifacts and stories of the information age." I'm very gratified that my story was chosen to be recognized as part of that history. I'm also gratified and, to be honest, when I got the news, quite surprised that my story was chosen because [of] the collection of people in the information revolution who made enormous contributions. That collection is so large. It encompasses so many threads of the information revolution. And so, I'm very gratified that I was selected out of that very large and worth[y] group. And finally, I was gratified-- I am humble-- a great deal of humility when I look at the people who have already been elected to the Hall of Fellows. Those people provided the concepts, the ideas, the artifacts, which form the basis for my contribution, that I was able to fit into. And it's really a great honor for me to be able to belong to their club.

Plutte: Okay. Well, thank you very much. That was great. Is there anything else you'd like to add? Do you have a question?

Feigenbaum: I'm going to say something about expert systems being accelerators of human mental activity but I don't like the world accelerators. If I tell you that -- like I have a big AI program hanging off my belt and it's my power tool. I don't want to call it an accelerator. What's a good word for that?

Plutte: I mean it could be a powerful tool. It could be…

Feigenbaum: I don't want to use the word tool because I've used it too much. I don't want to call it a rocket engine either because that's cute but it's not…

Plutte: After burner.

Feigenbaum: Huh?

<overlapping conversation>

Plutte: You could just go simple and say it’s something that will help us a lot, you know, something that will help us achieve more, you know, rather than doing a…

Feigenbaum: Amplifier.

Plutte: Amplifier, okay, go for it.

Feigenbaum: Okay. When we think of expert systems as amplifiers of human thought, we have to realize expert systems have always been used in the service of human activity in some area. They’ve never been stand-alone expert systems. And in fact, because of that, expert systems almost by definition came to be known as those AI programs that know how to explain themselves. That’s extremely important because humans have to come to trust the line of reasoning of the expert system. You can’t just have a black box announce this is what you are to do. [Rather, this is what you ought to do because of this fact and this hypothesis and this reasoning step. You have to explain what it is you’re doing. And that allows expert systems to be amplifiers of human thought. One of the greatest challenges of the next generation is how to enhance that amplification, humans cooperating with machines. Think of it as co-evolution of people and computers and AI software.

Plutte: That’s all I have. That was great, thank you.

END OF INTERVIEW

