



**AI: Expert Systems Pioneer Meeting
Session 1: Purpose, Structure, and Introductions**

Moderators:
Burt Grad
David C. Brock

Recorded May 14, 2018
Mountain View, CA

CHM Reference number: X8652.2019

© 2018 Computer History Museum

Table of Contents

WELCOME AND THANK YOU.....	5
MEETING OVERVIEW	5
ORGANIZER INTRODUCTIONS	7
PARTICIPANT INTRODUCTIONS.....	8

Artificial Intelligence Pioneer Meeting:
Session 1: Purpose and Introduction

Participants:

<u>Name</u>	<u>Affiliation</u>
David Brock	Moderator, CHM, Director Center for Software History
Burton Grad	Moderator, SI SIG
Brad Allen:	Inference
Avron Barr:	Aldo
Denny Brown	Teknowledge
Ed Feigenbaum	Intelligenetics/Intellicorp; Teknowledge
Mark Fox	Carnegie Mellon
Peter Friedland	Intellicorp
Paul Harmon	Newsletter
Peter Hart	Syntelligence
Gary Hendrix	Symantec
Fritz Kunze	Franz, Inc
Doug Lenat	MCC/CyCorp
Brian McCune	Advanced Decision Systems
Peter Norvig	Google
Alain Rapaport	Neuron Data
Herb Schorr	IBM/ISI
Reid Smith	Schlumberger
Monte Zweben	Red Pepper/Blue Martini
Stephanie Dick	Historian, University of Pennsylvania (attended remotely using the Beam Robot)

Colin Garvey	Historian, Rensselaer Polytechnic Institute
David Grier:	Historian, George Washington University
Thomas Haigh:	Historian, University of Wisconsin at Milwaukee
David Hemmendinger	Historian, Associate Editor IEEE Annals
Hansen Hsu	Historian, CHM, Center for the History of Software
Dan'l Lewin:	CHM, Executive Director
Paul McJones	CHM, Special projects
Len Shustek	CHM, Chairman
Ed LaHay	SI SIG, Meeting Project Manager

Welcome and Thank Yous

Burt Grad: It is May 14, 2018, and we are at the Computer History Museum in Mountain View, California. As I've already told you, we appreciate your coming here, and this meeting is being cosponsored by the Software Industry Special Interest Group. That's a volunteer group that I helped to found and run for about 100 years. <laughter> Mike Humphries and Ann Hardy are the current chairs of that. I'm emeritus at this point. Is that the right word, when you're old enough?

This meeting is being cosponsored by the Center for Software History of which David Brock is the Director. I want to thank Ed LaHay sitting back there for making the arrangements and logistics arrangements with the Museum and the young lady you saw out there, Poppy Haralson, who's just done yeoman work making sure we have all the papers and stuff we need for the meeting. We also have invited a few computer historians. They'll introduce themselves here, and they'll help direct the meeting and ask questions to make sure we're covering the topic.

The purpose of these meetings is to get information for historical purposes for current historians or historians 50 years from now. Maybe expert systems will become a major topic like the Gold Rush was at one point in time or the trains, things like that. And of course, I want to give a special thanks to the Computer History Museum. They're not only hosting the meeting for us, but they've provided all the administrative and operational assistance as well as the video and the audio stuff with Albert back there. They have arranged to have the Museum open tomorrow afternoon for those of you who haven't been here before and want to take another look. You'll be able to see some of the wonderful exhibits that are around here. How many of you have been to the Museum before? Turn it the other way around. Put your hands down. How many of you haven't been to the Museum before? Okay, for the six of you we'll keep it open.

Meeting Overview

Grad: This morning, we'll be talking about periods before most of you were working. It was in my time frame but not yours. This is the 1950s; we're using that as the beginning period for the AI [artificial intelligence] that we're checking and studying. I want to then, as I introduce that topic, tell about what you know personally, either from people you worked with or things you experienced if there was any of those. We're not getting it out of a book. What you read in a book is what you read in a book. That's already recorded. We're looking for anything new or something that you might have that might not be recorded. Then we'll finish off the rest of the morning with 1965 to 1975, when Ed and many of you got involved in this area and helped to build this whole thing. I get all choked up thinking about it.

The afternoon is the one that I'm most interested in. It's the business side. I have too many questions to ask, as Mark Fox pointed out to me. There's no way to cover all that with as many

companies as are represented here, but we'll find a way, and if necessary, our schedule is flexible. We'll carry some over to tomorrow morning if we need to. The companies are important—how you worked, what you did—we'll make sure we do that.

Tomorrow, we'll look more at how they were applied, what the applications were. Again, not so much the ones that are in the books like the one that, you know, Ed and the others put together, but things that you may remember, stories of things you worked on or things your clients worked on that are not in the books, something special and what was different about them. What was difficult, what tools you used, those are the kind of things that we want to explore.

The afternoon is a little more complex. We're talking about what happened: the demise of the industry. There are still some expert systems businesses around, but not much. Instead of growing like some of the areas did, a few multimillionaires, a couple of multibillionaires, most of you guys didn't quite get to that level. You seem to all be prosperous looking, though, so somehow you did okay.

Now, we're ending our history in 1990, in the 1990s and 2000. If we drift over a bit, it doesn't bother me, but a transition took place there from what you guys have told me. This is when all the companies died, but all of a sudden new stuff started to come in. They gave it different names. They tried different approaches, different technology. Talk a little bit about what was happening and why, from your knowledge again.

Now you may open your folders. I promised you that you could.

Now, do remember, rustling of papers and stuff in general, it's okay but will be picked up. If you do side talk, you talk to someone next to you like Mark is doing now, that will also get picked up and will prevent other people who are talking or listening. David is going to be the... what would we call you, interlocutor? No.

David Brock: Ringmaster?

Grad: Ringmaster it is; he has a whip. When someone's talking, they get to finish. Only David or I get to stop them. That's the rule. If you have questions you want to ask, he'll make a note and you'll get called on pretty much in sequence to ask the questions.

Brock: Yes, if you could just get my attention. I'll just keep a running list of everybody's name and the order. Then when people are done speaking, I'll just say who's next in the queue. That seems to have worked well in the past for these settings.

Grad: Most of you know each other, right? This is a record for purpose, for historic purposes. So, things you say, "Well, everybody knows that." Well, all these people know it. Think of me. I

don't know a damn thing, and you're educating me and you're getting a record on file so that people who look at this later on will get it. Just so you know, every session is being videoed in there. They'll all be transcribed. They'll all be edited, and they'll all be posted on the Computer History Museum website. But again, if something, a topic gets interesting, we'll pursue it in more depth, and we'll spend a little bit more time, but we'll come back to it later on.

Organizer Introductions

Grad: Any basic questions before we start around the room and have you introduce yourselves? So far, I'm in control, and that doesn't happen very much. What I will ask each of you to do is briefly, taking two, three minutes, explain your connections to expert systems and AI. You'll have lots of time later to talk in more detail, but at this time I just wanted to give a brief introduction so those of us who don't know you get that connection—what companies you worked with and that kind of thing. I'm going to start with David Brock.

Brock: Oh, great. Well, thank you. Thank you very much, Burt, and thank you everyone for coming. I'm a historian of technology, and I joined the Computer History Museum at the beginning of 2016. I came to play a role in the Computer History Museum software initiative that is manifest in a couple different ways. One, our new exhibition opened around that time called "Make Software: Change the World!" which is downstairs, and which will be open for you to see tomorrow. Also, we have a very active educational program, "Design, Code, Build," focused toward giving students experiences both of making software and lessons about the history of software. And the part of that software initiative that I have joined is the Center for Software History, which has a curatorial team and a collections team and together we're working at collecting, preserving, and interpreting the history of software. In my role with the Center for Software History I've enjoyed working with the Software Industry Special Interest Group: Burt, Mike, Ann. Also, with a set of extraordinary volunteers to the Museum; Paul McJones is one of them, who's here with us today.

As I'm sure your correspondence about this workshop has shown you, I will be playing a supporting role today to Burt's passions and enthusiasms for history, but these workshops, as some of the historians in the room can attest, have proved very valuable for people working in the history of computing. It's a real contribution that you all are making today, so thank you.

Grad: Thank you. Now I'll introduce myself. I'm Burt Grad, and I was one of the cofounders of the Software Industry Special Interest Group back in 2000 with Luanne Johnson, and we've been collecting software history since then. This is the 14th of our pioneer meetings, covering a whole range of subjects which you can look up.

Also, I was at GE in the 1950s and IBM for 20 years after that and then ran a consulting practice for software companies for another 30 years until my clients retired me. They all got older. Isn't that the strangest thing that <laughs> happens sometimes with these people? And I actually am

continuing. I work with one of my sons as a financial consultant too at a life insurance agency, and I'll be glad to give you his card.

It has been fun to do these kinds of things. There's also the *Annals of the History of Computing*, which many of you know or not? How many of you know the *Annals of the History of Computing*? Not as many as I would have hoped. I'll show you some copies. Based upon some of these meetings, we have done a whole special issue devoted to that subject. We're doing one now on desktop publishing. We had that meeting last year, so we're doing two special issues on that, and David Hemmendinger is the Associate Editor in Chief. Is that the right title? That's good? Close enough? Whether we'll do something on this will depend upon you guys. If somebody gets motivated and says, "Ooh, I'd love to do a special issue," we'll be glad to help. I ain't doing it.

Just so you understand, I do have a small connection to what you call expert systems. At GE in the late 1950s, I was designing systems, talking to the people who design products to try and automate the design. We worked on and used a logic tool called "decision tables," which we created at that time. I never knew about you people, and as far as I know, none of you ever knew about me. It eventually went nowhere, but we had fun doing it and it was quite useful.

Participant Introductions

Grad: Now we'll go around. Reid, starting with you.

Reid Smith: Okay. I'm Reid Smith. I think my point in being here is to give you a little insight into Schlumberger's history in AI, expert systems in particular. I started as a graduate student working with Ed and Bruce McKenna at Stanford. First heard about Schlumberger when a couple of people came to the Heuristic Programming Project and talked about it. I had no clue. Schlumberger's a very large oil field services company. I worked in several sectors. I was VP and Managing Director for three of them in Cambridge, England, in Palo Alto, and in Austin, Texas. The early years were done in Ridgefield, Connecticut, and I'm not going to go into detail at this juncture, but I'd be happy to talk about any of that.

I think Schlumberger is interesting in a variety of ways. One way is to look, as a simple example, at the number of AAAI fellows who are ex-Schlumberger people. There's about 13 of them. Peter, you may find some others that I don't remember. But that's one way. Another way to look at it is in terms of the number of companies that were founded by Schlumberger people either while they were in Schlumberger or later after they left Schlumberger. I think I'll stop with that. Thank you.

Grad: Thank you. Avron.

Avron Barr: Avron Barr. I've been an independent consultant for a lot of years now, but I started off pretty much when I came to Stanford as a programmer in 1972 for Pat Suppes doing computer-assisted instruction. I was co-located with some of the people from the AI Lab. A couple years later, in the spring of 1975, I took my first AI course with Professor Feigenbaum and through that got involved in a writing project that got me involved in as a graduate student in Professor Feigenbaum's lab. I spent the better part of six years pretty much looking into every AI project that I could find.

When that was done, I got involved in starting one of the early expert systems companies, Teknowledge. From there, I was able to start a career as an independent consultant, and that's what I've been doing ever since. I have possibly uniquely worked for most of the expert systems companies over the years and certainly have become familiar with a lot of the long trails of expert systems technology over the course of my career.

Grad: Thank you. Brad.

Brad Allen: Hi. My name is Brad Allen. Let's see, a point of departure for me in the expert systems world was starting from my years as an undergraduate at Carnegie Mellon University. I became a research programmer at Mark Fox's Intelligent Systems Laboratory, when it was just newly founded as the Robotics Institute. In the early 1980s game, you know, some of the earliest practical experience in fielding expert systems in an enterprise environment, and the time that the initial wave of enthusiasm for expert systems and the commercialization started to come up in 1984, I found myself being in the position of the very desirable entity of somebody with direct practical experience in building expert systems. I went out and talked to all of the different expert systems companies at that point and ended up at Inference Corporation in Los Angeles at the time. I became one of the development team that built ART, one of the early expert systems shells and then spent a few years after that really working on trying to understand how to apply that with a number of companies: Ford, American Express, American Airlines, NASA at the time.

Basically, I had the experience there of kind of going through the entire cycle of tremendous enthusiasm through to disillusionment in the early 1990s. I helped the company at the end of that period kind of pivot into the consumer or the customer relationship management space for the introduction of a product called CBR Express, which was one of the first applications of case based reasoning to that problem. That's pretty much what I'm bringing to the table.

Grad: Sounds great. And Denny, not Dennis. I put the wrong name. Denny Brown.

Denny Brown: I'm Denny Brown. I came to Stanford in 1968, and Ed Feigenbaum was my advisor from day one. After a long and undistinguished graduate career, Ed hired me as Associate Chairman in the Computer Science Department when he took the Chairman's job. I was able to teach undergraduate courses in computer science during this time and then did the

administrative work running the department. I was never much for academic research, but I loved the personal interactions that came from teaching the courses and herding the faculty and students of the Computer Science Department.

Among the things that Ed and I did during the tenure from 1976 to 1981 was to create a master's degree program specializing in AI. MSAI [Master of Science in artificial intelligence] was a two-year program that included Fundamental CS, core AI courses, and a project-oriented dissertation. In 1981, Ed and a bunch of other people in this room, 20 folks, founded Teknowledge, and many of the graduates of that MSAI program ended up working at that. Certainly, after Ed turned over the Department Chairman's job to Gene Golub, I left academia to join Teknowledge as employee two, three, four, or five depending on how you count.

At the beginning, and this is early 1982, Tek's main focus was education and training. We had a one-day executive briefing, a one-week overview of expert systems, a five-week intensive that was essentially the core curriculum from the MSAI program, and a six-month onsite mentored training where it was essentially the MSAI project. Participants from all over the world would come for six months to Palo Alto and go through the five-week training and then work on their projects and, roughly speaking, do the MSAI without the academic credentials and writing and that stuff that went on with that. So, my first role at Tek was not surprisingly to manage the education and training function. The courses were initially built by the founders, and my job was to make them more pragmatic and to make them deliverable by other people besides the founders, because most of them kept their day jobs doing what they were doing in academia. We delivered multiple instances of training in the first couple of years.

Grad: Denny, we're going to be covering companies in detail tomorrow or later today. What other things did you do?

Brown: As my role at Tek evolved and I ended up managing, helping manage the Custom Application Group, we built big custom applications for people like General Motors, Proctor & Gamble, NCR, Motorola, and others.

Grad: So that was your responsibility?

Brown: I'll have more to say about Tek and Coherent Thought when we get there. Thank you.

Grad: Thank you. David Grier will join us as historian. He will join us a little later. I'll have him introduce himself when he comes. Peter, it's your turn.

Peter Friedland: Peter Friedland. My first experience with expert systems, although I didn't know they were called that for about four years later, was as an undergraduate at Princeton in 1970. I was lucky enough to get involved for the next four years on a project called SECS,

Simulation and Evaluation of Chemical Syntheses, which was a system that Professor Todd Wipke was developing basically to apply AI-based knowledge representation and search technology to a problem of finding organic synthesis pathways. That is still a very interesting problem for AI people all over the place. Every so often I read an article about how the latest AI techniques are being applied to that problem.

In 1974, I came to Stanford as a Ph.D. student. What convinced me was getting a phone call from Ed Feigenbaum when I was deciding which graduate school to go to telling me about what they were developing, but more importantly, explaining about the weather in California versus Pittsburgh. I followed Ed's advice and, which was extremely good advice for a lot of reasons. I got to Stanford in 1974, which from my point of view is an extremely lucky time because it was a time not just when the technology that all of us have been involved in in expert systems development, but also when molecular biology hit the point where it transitioned from being a qualitative science to a more quantitative one.

I was lucky enough to be involved with developing of the first AI systems for molecular biology. That was called the MOLGEN project, which involved not just computer scientists like Ed and some of the people in this room, but also extremely famous and renowned molecular biologists like Joshua Lederberg, the Nobel Laureate, Doug Brutlag, Larry Kedes, others. Well, a lot of Nobel Laureates at Stanford donated their time in exchange for us building some not AI tools but very useful tools to help with the explosion of data that was currently coming in molecular biology. That led to not just the Ph.D. thesis but also the formation of what was certainly one of the earliest both bioinformatics and AI companies and probably the first ever to combine the two: Intelligenetics, which we'll talk more about later. Basically, it combined software technology both for helping molecular biologists sift through what had gone from a tiny little bit of data to a huge river of data and also developing follow-ons from my Ph.D. thesis in understanding how molecular biologists design experiments. That led to the formation of Intelligenetics in 1979.

I was also part of the group of 20 that has been mentioned several times that formed Teknowledge. Left Stanford in 1987 after supervising students and working on the project, later to form what eventually became the Intelligent Systems Laboratory at NASA and coordinated both the program at NASA Ames, which eventually became a very large AI laboratory, and the NASA-wide project. Was lucky enough to find people like Monte Zweben as a young kid who actually called me up and said, "You want me for your laboratory." He was right. He actually said, "You want me, and you want me to be your deputy."

Grad: Nice.

Friedland: I said, "Oh, who are you?" He came out, we talked and he convinced me. That turned out to be one of the smartest decisions I ever made. That led to that at NASA. Right now, I'm an advisor to the US Air Force on science and technology, particularly AI technology for the Air Force.

Grad: Thank you, Peter. Paul.

Paul Harmon: I spent the 1970s as a behavioral psychologist working in training. The idea of cognitive psychology and things associated with cognitive were very far away from anything I knew about. I developed a lot of programmed instruction, and then I branched out and I moved to California in the beginning of the 1970s. I focused primarily on helping businesses develop training materials of various kinds.

I don't know how, but at some point, Denny Brown called me up in the beginning of the 1980s and said, "Teknowledge is doing training, and we need someone to give us some advice on it. Would you come in, sit through a course, and give us some suggestions?" I did that, and I did it about, I think, five times maybe, one course after another. Each time I'd write out all my suggestions. Each time they'd do some of it.

At some point, they said, "We're not going to do this initial training anymore. We've done it enough, so thank you, anyway." I thought, "Great. This is a fascinating subject. I'm totally interested in the idea of expert systems. I'm going to go away and write a book that describes this from my perspective. I've watched them do it over and over again, and they're not talking to business people; they're talking way over their heads. I can write a book that's at a different level and communicate a lot more."

I spent the next year, I guess, writing a book, *Expert Systems*, which then presumably had started to sell. That established me as somebody who knew about expert systems, which is questionable, but then you've all dealt with the press; you know what we're like. <laughter>

I then proceeded to have somebody approach me and say, "Would you write a newsletter about this?" and I agreed to do that. Then for the next 10 years, I wrote a newsletter called *Expert Systems: AI for Business*. The same thing. I love being a newsletter person, incidentally. You'd never imagine how many people will invite you and then tell you all kinds of things that they probably wouldn't tell somebody else, but they'll tell you because you're a newspaper person.

My view of expert systems has always been from the outside. It's been watching things happen. It's been watching companies do things. In the early 1990s when expert systems did start to fade, I transitioned to a new field. Well, I transitioned back to what I'd done originally, which was helping people in business solve certain problems. I wrote a book on business process, which is also selling well, which is now my current business, training people in how to improve their business process. I mention that only because many of the companies I ran into in business process were business rules companies, business expert systems companies that transitioned to becoming rules companies. I met them all again as they were doing something else, in any case.

Grad: Thank you, Paul. Herb.

Herb Schorr: Hi. I'm Herb Schorr. Since this thing also says, "AI," I thought I'd start by saying that as a postgraduate student at Cambridge University, I was one of the inventors of what I think still the fastest garbage collection algorithm that exists. <laughter> I went to IBM and eventually was Director of their Computer Science Division and I set up the AI program. IBM had a big bugaboo of not doing anything like AI. I told them, "Well, we're just doing this because it throws off interesting software that worked." We had an AI Division and we did things on Prolog and other things of that sort, the fifth generation. I was the lucky man who got to tell the Japanese that IBM would not participate in the fifth-generation projects with Alain Rappaport, who's somewhere in here. There you are. You've got gray hair now. <laughter>

We set up the Applications of AI Conference, which I was told today is still ongoing. I left Yorktown and went to a job in development, where I was in charge of new business ventures for the IBM Corporation, one of which was artificial intelligence products, which focused on expert systems. But I had an unusual job: I had charge of development and finance and marketing, so I got to meet an awful lot of people. Very unusual, but they decided they had to do something. Then I left and I went to ISI [Information Systems Institute] in Southern California in the University [USC]. Basically, I set up the AI Group there, which still continues strong. I will stop at that.

Grad: Thank you, Herb. Peter.

Peter Norvig: Hi. Peter Norvig. I came to the party a little bit later. I didn't get to grad school at Berkeley until 1980 and worked on natural-language processing, not on expert systems, but using the core technology of semantic nets, starting with marker passing and spreading activation. Then as Horvitz and Heckerman and Dean Cherniak started to introduce probabilistic inference, working with that as well.

I was at Berkeley as a grad student and later as a research faculty. 1990 was one of the AI winter areas. I figured out that my grant wasn't going to be renewed, so I said, "Well, if the funding agency doesn't care about the work I'm doing, maybe I'll stop it and I'll spend all my time writing a book about Lisp instead." So, I did that and I included the implementation of E-MYCIN in the Lisp book, so there's a connection with expert systems there.

I then went off to Sun Microsystems working on what we then called information retrieval using semantic networks. We now call it search, so that helped me when I eventually found my way to Google. In between I found time to write a textbook on AI. Stuart Russell and I are now doing another revision, so if anybody has comments, now's the time to tell me. <laughter> I found my way for three years to NASA and inherited the group that Peter and Monte helped build.

Grad: What are you currently doing?

Norvig: I'm currently at Google. I'm Director for Research and actually spending a lot of time with startup companies, helping them understand how to do AI.

Grad: Thank you very much. Doug.

Doug Lenat: Doug Lenat. I've been doing expert systems research and development for many years before and after it was called that. Before, in particular, many of you will know I was a student of Cordell Green's at Stanford, 1972, 1973, 1974, working on how to build program writing programs, automatic programs, synthesis programs. We took in effect an expert systems approach to that.

By the end of that period of time, I realized what needed to be done to break the bottleneck in that, which was not have such a narrow focus on what the performance was going to be, what the task was going to be. I was fortunate enough to find Ed Feigenbaum at Stanford and became his student for my final year and a half or so as a graduate student. Did the AM program, which again, most people don't think of it this way, but it effectively was an expert system to do automatic discoveries, plausible conjectures in math. After that, decided on what its bottleneck was. That led me to Carnegie Mellon and then back to Stanford, where I was a faculty member up through the 1984 time period.

At that point, I realized that we needed to do a project to break that bottleneck that would take a couple thousand person years of effort. You know, if you do the math, you'll find that that would have taken 400 or 500 years if I'd stayed a professor.

Moving to the wilds of Austin, Texas to the just forming Microelectronics and Computer Technology Consortium, which was America's answer to the fifth-generation computing effort, enabled me to have that extra order of magnitude. It's only taken 35 years for me to get where we are. Even though most people don't think of the CyCorp, which is the company that spun out of MCC 25 years ago that I'm the CEO of, don't think of that as an expert systems company, effectively, that's what we're doing. I'm sure I'll have more to tell you about how it is that we've stayed in business all these years when none of you have. We are the last surviving large expert system company.

Grad: Gary Hendrix. Stop laughing. You can tell he's amused by the last statement

Gary Hendrix: All right. I'm Gary Hendrix. I was a natural-language processing guy. Started at Stanford Research's Artificial Intelligence Center in 1974. Thank you very much for hiring me, Peter. I really didn't have too much to do with the expert systems, so it's sort of an anomaly that

I got invited to this conference. But I feel very honored to be surrounded by all you expert systems luminaries.

Grad: Let me ask you a question, Gary. Symantec when you started it was your first company. I did an oral history of you at one point in time many, many years ago.

Hendrix: Okay. I was there at SRI, and there were a bunch of us that were interested in figuring out how we could turn artificial intelligence into commercial activities. I think we were sort of the first ones that really thought that you could make a company out of this stuff. There were a lot of good things that were happening. It was expert systems, for sure, but the natural-language stuff was starting to work really well for accessing big databases, and machine vision was coming along pretty well. A lot of things were beginning to work.

We started meeting at my house on Tuesday nights and eating popcorn and drinking beer and talking about how we might be able to make something of this. It didn't go too well because Charlie Rosen, our ringleader, was always smoking and eventually my wife threw us out of the house. We had to go over to Charlie's house, and Charlie had a very well stocked wine cellar. Well-lubricated by that, we got Machine Intelligence Corporation off the ground. And I don't know if Machine Intelligence Corporation was the very first AI company, but if it wasn't, it was pretty close. It was right there at the beginning.

Grad: What year? 1977 maybe?

Hendrix: I think it was 1978. The reason why I say that is because we got kicked out of the house when Jeannie became pregnant. because she was afraid that the smoke was going to ruin the baby's health and the baby was born in late 1978. Machine Intelligence was born at about the same time, so that was the way that worked.

Machine intelligence needed to focus on something that would, you know... We thought we could do it all, but the realities kicked in <laughs> pretty quickly. The decision was made to focus on machine vision in support of industrial automation, so basically doing robot eyeballs. That was what Charlie Rosen was interested in, and he was the one who left SRI to work full time at this new company, a very brave thing. The rest of us all wanted our day jobs.

Grad: But Symantec took another path.

Hendrix: Symantec was a spin-off of Machine Intelligence Corporation.

Grad: Ah. I didn't know that.

Hendrix: Machine Intelligence did the vision and Symantec, aided by a little grant, a \$25,000 grant from the National Science Foundation, which was big money at the time, started work on natural-language processing for microcomputers. Something to have there that you could sell with all these new homebrew computers that were coming out and the new PC and the small Apple computers that were available at the time. Symantec's contribution was to try to boil down these algorithms that had been running only on very expensive computers and make them run on these computers that you could buy for a couple of thousand dollars.

I think Symantec was probably the first company that had a product, the product was Q&A, was a flat file database and a word processor integrated. It also had the ability using artificial intelligence to ask questions of that database using ordinary English. I think that was the very first product that was widely distributed that had an artificial intelligence component to it. We sold tens of thousands of copies of that, and it got used all over the country and the world.

Grad: We'll come back to that when we talk about the companies. I wanted to make sure that you had some credentials for being here, and it wasn't just an accident that we invited you. That was the point I was making.

Hendrix: I don't know if that's credentials or not.

Grad: It's sufficient.

Hendrix: At any rate, Symantec went off and figured out that it was easier to make money not in artificial intelligence but in computer security of various kinds. <laughs>

Grad: Monte.

Monte Zweben: Well, I'm very honored to be here. I feel like I'm in the second generation of expert systems. I started my career at Carnegie Mellon. Took my first AI class with Jeff Hinton and Steve Minton. As I was mentioning earlier to Gary, I reimplemented Gary's LIFER, his natural-language system, and got a great grade on it from Jeff and Steve. I said, "Boy, this AI thing is pretty cool. I'm going to go into this." That's how I got into this career, so it's nice to be sitting next to you.

My career has been always about taking the state of the art and research it and applying it to real-world business problems. I started at Carnegie Mellon, left and built my first expert systems at Westinghouse. Actually, to pay my way through Carnegie Mellon, I built expert systems in the productivity and quality center at Westinghouse. Got exposed to Mark's work, which you'll hear about in a little while. Did that and then got a job coming out of Carnegie Mellon and running an AI lab at Mita Corporation. That was in Boston. That was a great experience for me because got to take what I learned from the Carnegie Mellon community and see the MIT perspective on AI.

It was very different. I got to take classes with Marvin Minsky after having taken classes with Alan Newell and Herb Simon. It was really great to see that contrast.

There, we built constraint-based systems, and it was a different way of looking at expert systems. We built constraint-based systems to plan military missions. We built it for space shuttle missions, and that's when I got exposed to NASA. Then I met a scientist at NASA who's not here, but he was a really great scientist, worked for Pete Friedland and me, by the name of Peter Cheeseman. He really wanted to focus on doing it as reasoning about uncertainty, and he wanted to do planning research and he didn't want to manage this lab. He said to me, "Hey, why don't you manage the lab, Monte, and I'll just do the research." I said, "Sure." I told all my friends and colleagues that I was going to manage the Artificial Intelligence Lab at NASA until I got that call that Peter mentioned about before, when he told me he had that job and I thought I had that job. On the fly I said to Peter, "Well, if I don't have that job, I want to be your deputy."

That's exactly how it went down and Peter gave me the job. I appreciate that, and that really was a great trigger for my career. Peter was a great mentor. We built a wonderful laboratory at NASA that still to this day has impact on artificial intelligence. I worked on my own research there. I specifically worked on planning and scheduling projects and machine learning projects and specifically applied that to the space shuttle program and helped plan to schedule the repair and refurbishment of space shuttles. That's what led to my first company. My first company transitioned that out into the private sector and built a supply chain out there.

Grad: What was the name of the company?

Zweben: That was called Red Pepper Software. It used the same techniques to plan and schedule manufacturing operations. That became the heart of a very large company's manufacturing product called PeopleSoft. PeopleSoft purchased us. I got the itch to do something outside of PeopleSoft while managing a large lab there. I really wanted to do something again as a startup, and I built what I think was one of the first machine learning companies called Blue Martini Software. That was at the height of the e-commerce boom where we were directly taking large scale e-commerce applications and marketing applications and directly enabling basic business people to use AI and machine learning to learn. That was Blue Martini. Now I'm the CEO of a company called Splice Machine, building a data platform for artificial intelligence. Thank you very much for inviting me and I appreciate being here.

Grad: Thank you, Monte. Next.

Mark Fox: Hi. I'm Mark Fox. I guess I trace myself back to 1967 when in my high school physics class, grade 10, I was introduced to a Digital Equipment machine, a GE Mark II or something machine that I programmed in Basic and an APL machine all at the same time. In grade 10, I started programming and decided I want to do computer science. Went to the University of

Toronto, which is my hometown, and studied computer science there. In 1974, took a course on artificial intelligence.

At that point I got bitten and I asked the faculty there where I should go. They said Carnegie Mellon. Why Carnegie Mellon and not MIT or Stanford? I'm not sure, but that's what they suggested, so I applied there.

Went there in 1975. Felt like a kid in a candy store. People who know the CMU culture know that grad students are sort of treated like faculty. You sort of go off and do whatever you want. It was an amazing place to be. I think I was bad luck for certain faculty. Rick Hainsworth was my first advisor, and then he left. Doug Lenat was my second advisor, and then he left. Finally, Herb Simon took me on, and for some reason, he kept me.

Then in 1979, I hadn't finished my Ph.D., and Raj Reddy, who was a mentor of mine, asked me to start the Intelligence Systems Lab of the Robotics Institute, which had yet to be started also. The Intelligence Systems Lab was meant to be focused on white collar robotics or basically AI as we tend to practice it. I started that up in 1980, and one of our first applications turned out to be scheduling, where we took constrained approaches to scheduling. Brad Allen was part of that project in the early days. And Monte and I have commiserated about all that stuff over the years.

Zweben: And we wrote a book.

Fox: And we edited a book. So the starting of the Carnegie Group, which is one of the reasons why I am here, grew out of the kickoff of the Robotics Institute in December 1980. I remember it very well because John Lennon had been shot, and I was getting ready for the kickoff of the Institute and programming in my office. At the demonstrations there, a venture capitalist who invested in the early days in Digital Equipment saw what I was doing in the area of knowledge base simulation and saw what Jaime Carbonell was doing in the area of natural-language understanding and said, "I want to bring you two guys together so that we have a natural-language interface to a simulation system." I thought it was a great idea except I had one problem. I had yet to start my Ph.D., so I was always in year five of my Ph.D. I had a topic, but I was having fun doing everything else, so hadn't really done it. I said, "Thank you, but no thank you. I got a Ph.D. to finish." But I started talking to Jaime. Then we brought in Raj and then John McDermott and we started... You were meeting on Tuesdays?

Lenat: Yes.

Fox: It was Sunday mornings with Raj. Everybody knows Raj Reddy, and there was no such thing as a day of rest in Raj's life. Digital got wind of it. John McDermott had very close relationships with his success with R1 Xcon. They came to us, and I'm not going to do the story

now, but I'm just going to say that one piece. They said, "Here, why don't you start a company? Here's a million dollars in contracts, a million dollars in equipment. Start." That was the genesis of the Carnegie Group at that time. I held various roles in it: CEO, VP of Engineering, back to CEO, etc.

About the same time as the kickoff of the Robotics Institute, I was doing some consulting with Westinghouse, and they introduced me to this problem of steam turbine and generator diagnosis, which their math department had spent three years applying phasing techniques to diagnose three problems with three sensors. The only problem is a turbine and generator have 1,000 sensors on it, so I designed a system. I never thought of this as being important at that time because I was also doing everything else. It was just sort of a side thing, designing an expert system that did much more than what the math department did, which went into production use as a PDS DMA Gen 8. Talked to Reid one day a couple years ago about having done that years ago, not knowing if it was still in existence, discovering that that system which has over 15,000 rules in it is still in operation by Siemens monitoring over 2,000 steam turbines and generators around the world. It may be the longest lasting expert system in production use today.

Zweben: That's awesome.

Fox: Where am I today? There's lots of things. I started another company called Novator Systems in 1983 for online retail software services. One of the first, I think, in the world to do that. That's another story. I returned to the University of Toronto in 1991 as a Chaired Professor, and today I'm a Distinguished Professor of Urban Systems Engineering and a Professor of Industrial Engineering and Computer Science.

Grad: Thank you, Mark. David.

David Hemmendinger: I'm David Hemmendinger. I taught computer science first at Wright State at Ohio and then at Union College in Schenectady. Retired eight years ago. Actually, my only connection with expert systems is having studied them during a course in AI in 1981 when we studied MYCIN and E-MYCIN. After that, I taught courses at Prolog for several years in the 1980s.

I'm here as an historian. I've been interested in history of science and mathematics during a previous life teaching philosophy. About 20-odd years ago I became a co-editor of *The Encyclopedia of Computer Science* with Tony Ralston. That got me into fifth gen computing, which I've been studying since then. I've written several articles for *Annals*. Now as Burt said, I'm the associate editor in chief of the *Annals of the History of Computing*. At the moment, Burt and I are working as guest editors for special issues on desktop publishing, which grew out of a similar pioneer meeting a year ago.

Grad: David, here, is a friend of ours, and we'll see where we go. Peter, you're next.

Peter Hart: I'm Peter Hart. I entered AI in the mid-1960s through the pattern recognition and scene analysis route, like many others did. My first big project was working on the world's first robot with artificial intelligence we called Shakey, which nowadays lives downstairs right here in the Museum. My experience with expert systems began in the early 1970s, when I led the development of a project called Prospector, which worked in the area of mineral exploration and was successful in identifying where a hidden ore deposit was. In 1983, I cofounded a company called Syntelligence, which has some points of uniqueness with respect to both business model and knowledge representation scheme; that I guess we'll get into later in this meeting. In the 1990s, I led a research center and was the CEO and chairman of a subsidiary of a Japanese company, Weco, right here in Menlo Park, where among other things, I led the development of an expert system for diagnosis that—Gary would be interested in this— combined diagnostic nets with what we called query-free information retrieval, which I thought was a nice tagline.

That was probably the last honest job I had. After that, I think, they've made me do things like be chairman and CEO of companies and things of that nature. Nowadays, I mostly travel around the world taking pictures of birds and riding a bike.

Grad: Are the birds artificially intelligent?

Hart: They are not.

Grad: They're truly intelligent. Is that what you're saying?

Hart: You want a scientific tidbit on that: it's recently been found that the density of neurons in birds' neural systems is far, far greater than the density of neurons in mammals. That's the reason that birds have such remarkable intelligence in such a teeny pea-sized neural system.

Grad: Thank you for that piece of information. Alain.

Alain Rappaport: So, Alain Rappaport. I guess I get into AI in a very different way, through some work that a friend and I were doing. I was in medical school, and my friend was at École Polytechnique in France. We were very fascinated by some debates that were going on at the time around whether Piaget or Chomsky were the right approach. We got really fascinated by the contrast in the structuralist approach of Piaget. Our friends and parents were a bit worried by the time we were spending on that theory. We decided to go for something that at the time existed in France called the Scientific Prize given by Philips Corporation every year for some kind of extracurricular work. In the U.S. it was called Westinghouse. Westinghouse was giving a prize for years.

Grad: Yes, Westinghouse.

Rappaport: Yes. We go to be laureates of this and started working on making sure that it was an interesting approach, which in the end it was really an important mission for us. We were looking at how to learn in a single shot, single event. One thing led to another. I met Raj Reddy. It was just a very impressive meeting when we had this first one and then continued to have conversations and got invited to go to CMU as a visiting scientist. I knew Mark. An intellectually unbelievable place to do what we want to do, which was really to scale the theory. Originally when we found some data about the sequence of the cytochrome p450, we did our first demonstrations on a TRS-80, which was our machine learning platform. That was a bit small. It was fascinating when we arrived in an environment with DEC-20s and Lisp machines, and we could start formalizing things.

At that point, I had my M.D. and a Ph.D. in molecular neuropharmacology. We used that kind of background to build expert systems that would be very powerful in the sense that you could design your control structure. I mean, put those in a field. Also, the Macintosh appeared, and with the Mac, we understood that the future of expert systems would be in expressing your knowledge very easily, visualization, instant visualization. With that, we started the first company.

Grad: Yes. That's what I was going to say. Did you start the company?

Rappaport: Neuron Data. I don't even know exactly where to start. It was such a fascinating experience.

Grad: Well, we'll come to the company tomorrow or later today.

Rappaport: Yes, okay, then. The second company I did was a little bit as Peter was saying, something called AI became search or vice versa. I did a search company in medicine and health care for both professionals and consumers. We were bought by Microsoft in 1987. Not 1987, in 2007. That became Bing Health, which was the approach there, to have vertical areas.

Grad: We'll cover that.

Rappaport: I'm still in search with another company that I cofounded called Nudgit. We'll talk about that as well.

Grad: Yes, we'll cover those. We're running out of time here. Thomas Haigh.

Thomas Haigh: I'm associate professor at the University of Wisconsin Milwaukee in the History Department and visiting professor in the History of Computing at Siegen University in Germany.

My first two degrees are in computer science at Manchester University. Actually, within that, AI was my main focus area. I did five courses in it, and I heard of some of you as an undergrad. I did my equivalent of a senior thesis on card-playing contract bridge programming in Prolog. At that point, obviously, there was no future in AI, so instead I went into history, which was done splendidly with my Ph.D. in History and Sociology of Science at the University of Pennsylvania.

Since then, I've been very active in the history of computing community. I've written about databases and word processing and management information systems in business, a book on Eniac, and various other things. I'm going to mention just two things I do. First, if you get *Communications of the ACM*—maybe you keep it in your bathroom and, you know, read through it sometimes—you may have come across some of my historical reflection columns. I do a couple of those a year. Actually, I was just thinking about the next one I'm going to write as people were talking. We have NASA and Google. So, in something called Google, you don't know what the moon shot is. Talking about the actual history of the space program and style of innovation versus the style of innovation that tech companies do now. I'm going to probably try and take my knowledge as a historian of technology and use it to frame things in a way that will be interesting to a broader computing community. I also edit the Turing Award winners' website and work with the ACM History Committee on a few other things.

Something I'm doing right now is with Paul Ceruzzi, who's a senior curator at the Air and Space Museum. He has a book, a history of modern computing that came out 20 years ago. I'm doing a new edition of it with him. It's going to be all changed around and very substantially revised. We're grappling with this question: How do you get the whole of the history of computing from Eniac forward into one book? Because computing becomes so many things, right? It's tricky. One of the questions that came out yesterday when I was discussing it with some other historians is, what do we do with AI? There's no AI chapter. Now it's not the history of computer science. If it were, AI would definitely get a chapter. It's not the history of computing research. It's the history of computing in terms of hardware, software, algorithms, practices, as it's actually done in the world.

From my own background, I know AI is a big interesting thing. I like it. But if I'm talking about computing, as it has the impact on the world, where does it go? As you're talking about, expert systems came and they went. Do they even need to be in the book at all? That's what I've been thinking about. Do they get a paragraph somewhere? If they get a paragraph somewhere, what should it say? That's the question I'm going to have in my mind as I listen to all this discussion. The MIT hackers pop up. They're in the timesharing kind of chapter, and of course there's the AI Lab. We mention the incompatible timesharing system. We've got video games. I think AI's going to be there a little bit. The place I know it's definitely coming is at the end, like Siri, all these massive things. I guess it grows out of case-based reasoning, but it's not called that anymore with all those massive data things, all the stuff that Google's doing. It's definitely there and self-driving cars and the last chapter. What I've been thinking about is, what are the places earlier in that story where we can say AI in general and expert systems in particular actually

made a difference enough that when I'm trying to get the whole of the history of computing into one book they should pop up?

Grad: Where does it fit in? Thank you. Tom has been an incredible asset to us in the last 15, 20 years. I think you've probably been to almost every one of our 14 meetings.

Haigh: Yes.

Grad: And has always contributed. Wrote a number of articles for the *Annals* and has just been an incredible help to us for a very long time, and I appreciate it.

Haigh: I'm just a Burt Grad groupie.

Grad: Brian, it's your turn.

Brian McCune: Hi. I'm Brian McCune. I got my start in AI from a crazy Italian professor at Oregon State University called Emilio Gagliardi. He was teaching from a great book on pattern recognition by Nils Nilsson. That led me to Stanford and the Stanford AI Lab and Cordell Green's automatic programming, or what we eventually called program synthesis or program refinement. I worked with Doug Lenat for a few years, and then we did something called the PSI System. What it really tried to do is create a topic for nine different Ph.D. theses.

It did a really good job of that. You could describe an algorithm using natural language or traces or predicate logic, and it would all get filtered down into an executable variable program. Then it would get optimized into optimal code, and at each step along the way, there was a Ph.D. thesis. This system actually ran one time in the middle of the night on an ISI computer; it was incredible.

Then I graduated and started the first of eight companies. I'm not going to bore you with all eight. I'll manage with three. The first company was called Advanced Information and Decision Systems, which we did not pronounce AIDS, ever. This was in 1979 and the AIDS epidemic came along, so we got rid of the "information."

Grad: Rid of the "I."

McCune: To stamp out AIDS. We became Advanced Decision System instead of Advanced Information and Decision Systems. That company developed expert systems for two groups. One was DARPA and the other was essentially the rest of the Defense Department. I had two main areas that I worked in. One was signal analysis and the other was information, what we called information retrieval. It's now called search. I'll say more about those later, but we

deployed a number of operational systems into the Defense Department. Along the way we also worked for Eastman Kodak, Bank of America, and Wall Street and deployed systems there.

In the text area I invented what I think was a very simple text-retrieval system in 1982 called Rubrik. In 1988, we spun that technology out into a company called Verity, which then went public in 1995. It's an actual AI company that went public, of all things, although by then it was not called AI. I don't know if we get credit for that or not. ADS was sold. We went from 1979 to 1991, and we then became a part of Booz Allen and Hamilton in 1991.

At that point, I became the Chief Technologist at Booz Allen and got involved in all sorts of different software projects worldwide. Most of them, unfortunately, very low tech, so I didn't stay there very long. After that, I was a consultant for about 20 years and retired just a couple of years ago.

Grad: Brian, thank you very much. Ed, can you speak? Ed has laryngitis, so I don't know if that's good news or bad news.

Ed Feigenbaum: I'll try to keep it short. Even though I'm going to make it short, I want to start out by thanking Burt. He said at the beginning that somehow this was my idea. Well, it was his idea and a tremendous amount of work that he did to get this show on the road. It's just amazing.

Grad: Thank you.

Feigenbaum: Let me just divide my story into 10-year periods, and I'll just point out a highlight or two in each of the periods. In the 1950s, I was an undergraduate, early 1950s, at Carnegie Institute of Technology, now Carnegie Mellon University. I was intrigued by some of the work going on in the social sciences at the Graduate School of Industrial Administration, so I took a course in my senior year with Herb Simon. It was a course called Mathematical Models in the Social Sciences and that was the highlight, that was the big bang. Herb Simon walked into the class after Christmas—we were on the semester system—and he said, "Over Christmas vacation, AI Newell and I invited a thinking machine." There were no computers at Carnegie. What does he mean by all that? He explained that it was the logic theorist. He handed out IBM 701 manuals to all of us six students taking this seminar. Reading that manual got me hooked; I knew what I was going to do now. I was going to do these thinking machines, and I was going to learn how to program these things. —

I'm going to jump ahead to the 1960s for a moment to explain why I did what I did in the late 1950s. I was teaching at Berkeley with another graduate student of Herb Simon's, Julian Feldman, and we were teaching the first artificial intelligence courses that were taught at Berkeley ever. For the students to get material to read it was really difficult. I don't think

Xerography had even arrived yet at Berkeley. So we decided to do an anthology of the top stuff, 1962, called *Computers and Thought*, which is still in print. You can still buy a brand new one in print. That book is divided into two parts. One is sort of the engineering side.

Grad: We really just have a few minutes here.

Feigenbaum: Yes, Burt. That's right. I won't tell you about that. Anyway, to make a long story short, in the 1960s, what people were focused on in the first half of the 1960s in the problem-solving parts of AI was not the natural-language part but the problem solving and not in perception and vision. People were focused on generality in problem solving. In the later acronym, KRR, knowledge, representation, and reasoning, people were focused on the generality of the reasoning—Newell and Simon with their TPS General Problem solver and McCarthy with his various theorem providing approaches.

When I arrived at Stanford, I already had set up an alliance with Joshua Lederberg in the Genetics Department who—a fantastic mind—was interested in the same thing I was interested in: namely, the modeling of the processes of scientific reasoning, particularly induction. We started a project that Lederberg had suggested having to do with inferring hypotheses to explain that mass spectra in organic molecules. The highlight of that was in 1968. Looking back over three years' worth of experiments, it was clear that the power of that system, which then for certain classes of molecules, had become world class. The power was in the knowledge that we had evoked from the heads of the mass spectrometrists and their post-docs and graduate students, and not the complexity of the reasoning process. The mantra there was, “In the knowledge lies the power.” It's the “knowledge is power” thesis, principle in AI. The laboratory picked up the theme of extracting expertise from experts and building modest reasoning systems around that. One of them, a key one was a medical diagnosis: MYCIN. There were about 12 other projects that were going on in the Heuristic Programming Project at the time.

Grad: Ed. Just a second. In our next two sessions, we'll be covering this. You started a number of things. Why don't you just quickly go and tell what they were.

Feigenbaum: Then the projects were beginning to pay off. That is, they were really successful. We were the focus of attention in the artificial intelligence world and in the application world. For example, Schlumberger noticed it and arranged a contractual relationship with our lab. Because so many people were interested in this, it looked like it was going to be a viable company, there could be viable companies started out of this work. We decided to take orthogonal dimensions to the starting of the companies. One was focused on a particular area, namely the molecular biology computations. That was started as Intelligenetics. The other was going to do the teaching that Denny Brown mentioned, and it was going to do application development. It was going to sell software to the world of people who wanted to build expert systems, and that was Teknowledge. Intelligenetics prospered for a few years, but it was eventually sold to Standard Oil of Indiana, the research lab.

I got involved with in the early part of the 1980s with the fifth-generation project, wrote a book about it that was quite widely read. Then, actually at Peter Hart's suggestion, put together an anthology, or not an anthology so much as a survey. It was the era of in search of excellence; we went in search of excellent results in the commercial and industrial world from expert systems. That book became *The Rise of the Expert Company*.

In the 1990s, I did something completely different and went off to Washington to be Chief Scientist to the Air Force because the Secretary of the Air Force thought that AI was going to be a really hot topic. I mean, previous Chief Scientists had been physicists and laser folks and people like that, but she chose a computer scientist focused on AI. I'm going to stop.

Grad: Yes, thank you, Ed. It's just you have so much to cover that it can take the whole morning, but we'll come back. Fritz.

Fritz Kunze: My name is Fritz Kunze. I was born in 1952. In 1984, I was a graduate student at Berkeley in Math in the Mathematics Department. I got very interested in a program called Maxima, which spun out of a project back at MIT. I started a company called Frond. We sold something called Franz, and we sold something called Franz Lisp. Many of the expert system companies in this place bought our product.

Four years ago, I had a big stroke. I could still talk. My memory's intact, but as the day wears on my speech may slur as if I'm drinking. I'm sorry for that.

Grad: We're not providing any alcohol. Thank you very much, Fritz. I want to finish up the introductions, so we'll go quickly around the room with the other people who weren't covered before. Then we will start in on what's called session two, which has to do with what happened between 1955 and 1965. Alright. Let me finish with the introductions. We'll go around the outside. Ed LaHay. You have your microphone there?

Ed LaHay: Ed LaHay. I'm probably the least knowledgeable person on AI, but I helped Burt set up the conference. If you've got any administrative things, just come and see me and I'll sort it out. My background is in software. I wrote custom software and product software. For the last large number of years, I've been project managing software. You'll see me walking around with a camera and Alma Jones, who'll be photographing all this, and I'll probably post the photographs on the website. If you want to get a copy of some of the photographs, let me know and I'll copy them for you. You may notice that there are some blank areas on your attendance sheet. I will be pacing after you if you've got any blank spots to fill in things like contacts and so on.

Grad: Ed's been with us in the Software Industry SIG literally since the beginning as a friend of mine from IBM days back in the 1960s. Been in a number of companies, project manager, and

very successful in a number of areas. So, thank you for continuing to help us, Ed. Colin. You can just take the mic if you want to go there.

Colin Garvey: I'm Colin Garvey. I'm a Ph.D. candidate at RPI, Rensselaer Polytechnic Institute, and in training to be a historian of AI. I'm very interested in expert systems and particularly in the relation to kind of the internationalization of AI. In the 1970s and 1980s, a few of you have mentioned, the fifth-generation computing project with Japan involved, and I've been writing and researching a lot about that, so I'd really be interested to hear all of your reflections on that period of time. I'm also kind of looking at the current situation with machine learning and deep learning. Now there's a lot of talk about competition with China, so I'm kind of interested in maybe what kind of parallels there were to the AI boom of the 1980s with Japan and what that could kind of teach us about what's going on maybe with China now. The politics and kind of the history angles are very interesting for me. I'm really happy to be here, and I've really enjoyed hearing everyone's story already, so thanks a lot.

Grad: Colin is here thanks to Ed Feigenbaum, who made contact with one of his compadres—whatever you want to call them—at RPI who arranged for the funding for Colin to come out to be able to join us. He's doing an interesting document about his view of the history. That'll be in your folders when we give them to you when we leave tomorrow.

Len Shustek: I'm Len Shustek, cofounder now 22 years ago of what has become the Computer History Museum and still Chairman of the Board. My background is in computer science. I got a Ph.D. from Stanford in the 1970s, and in the late 1970s, I was on the faculty at Carnegie Mellon for a while. I learned what Peter Friedland knew instinctively, which is that the weather is better here than it is in Pittsburgh, so I came back to the Bay Area, and over a period of 16 or 17 years, did a couple of startups in networking. I then went back to teaching, so I went back to teaching at Stanford, Computer Science, and had the idea, which was crazy, to start a museum. That was in 1996 and then built this: The Computer History Museum. Burt has asked me at lunch tomorrow to speak more about how it came to be and what we're doing, and I will do that. Thank you.

Grad: Len has been such an incredible force. He's made the Museum stay together all this time. With the help of his wife, actually, Donna Dubinsky, helped to get the money to do some of these magnificent exhibits that you see here. Bought the building. He had a vision statement that he wrote before the thing started. I don't know how many of you ever seen it. But gradually, I was like, "Let's do software now." Later: "I need exhibits. I need people coming here. We need some money flow." Finally, a couple years ago now, he was able to finally realize that through the Shustek Center, which is collecting all the nonhardware and bringing that together and helping to focus us with the Center for Software History and these things. Len has been a delightful person to work with for all these years. Let's see, who's next? Let's look at the robot.

Stephanie Dick: Hi, everyone. I'm so sorry that I can't be joining you in person. I'm Stephanie Dick. I'm an Assistant Professor in the History and Sociology of Science Department at the University of Pennsylvania. Currently, I'm a scholar in residence at the Max Plank Institute for the History of Science in Berlin, so it's just about 8:00 pm for me today. I'll stay with you as late as I can. That's the reason I can't be with you in person, though I would very much like to be.

I finished my Ph.D. at Harvard University in 2015 in the History of Science Department. Following that, I was a Junior Fellow at the Harvard Society Fellows for three years before starting at Penn. All of my work is about the history of mathematics and computing in the 19th and 20th century United States. In particular, I work on the urban history of automated theorem proving and automated reasoning, which some of you all will know as much or more about than I ever could, so I'm really excited to learn from all of you over the next couple of days, although I will be doing it in a strange cyborg form. Perhaps that's appropriate for a gathering of this kind. So, wonderful to be with you as much as I can, and I'll try not to be distracting on the robot from the back. Thank you.

Grad: I told Stephanie that I was a robot at one of the meetings before, but I covered the camera. I didn't want to see anybody when eating the donuts or scratching my ears or doing whatever. But she seems much more appropriate than I was. Hansen.

Hansen Hsu: I'm Hansen Hsu. I'm a curator here in the Software History, and I work with David Brock. I did my undergrad in EUCS at Berkeley, spent six years in industry as a software engineer at Apple. In 2005, I decided to return to school, go to graduate school to study history and sociology of computing. First, I studied at Berk and then Cornell where I finished my Ph.D. in Science and Technology Studies in 2015. Here I am today, and I work with David. I do a lot of oral histories. I'll be doing Brian's oral history on Wednesday. I'm happy to learn all about AI from all of you. Thank you.

Grad: Thank you, Hansen. Paul.

Paul McJones: I'm Paul McJones. I also got a bachelor's degree at Berkeley, 1971. I had the good fortune to take a class from Peter Deutsch back in those days on AI programming languages. Peter Deutsch, Herb Schorr, and Bill Waite did a marketing algorithm that Herb had mentioned earlier. My career was software, R&D, systems software—that's not really relevant. My one other thing I think that is a little bit relevant is I've been helping the Museum collect software source code, so I ask all of you to think, are there artifacts, source code from graduate school days or maybe early company days, who knows, that you could donate to the Computer History Museum so that future scholars and students could study what you actually did? That could be the tools, it could be an expert system with the data, with the knowledge in it, or whatever.

Grad: That wasn't very painful. <laughs> We can now move ahead with the program. Anything else that we think is vital to add before we move ahead on to the subject matter? Yes, Peter.

Friedland: Just one quick thing. I just remembered as Ed was talking. We actually had in the Heuristic Programming Project, which led to the Knowledge Systems Laboratory, through about a year a philosophers and historians spending time with us by the name of Lindley Darden. I haven't seen her in years. I don't know, have you talked to Lindley in many years? She actually wrote several books about the history of biotechnology, early biotechnology, and spent a year watching us on a visiting professor or whatever thing. You might see if you can find her, because she'll have a lot of her historian's perspective knowledge.

McJones: Could you say the last name again?

Friedland: It was Lindley Darden. I think she was a professor at the time. Of course, this was back in the 1980s that she spent time there. I think she was a professor at the University of Maryland.

McJones: Thank you.

Friedland: If you can find her, she probably has a lot of historical notes from those days of developing expert systems.

McJones: Thank you, yes.

Grad: Just as an odd point before you go - remember, this is a beginning story. This is what took place up to a point in time. AI has become... everybody knows what the initials stand for now. There's a follow-on story here with a lot to be done, and the Museum's been very interested in this area for the future. This is not an end of the story; this is simply providing that background and history. Tom?

Haigh: Oh, yes, I just remembered: in the existing edition of *The History of Modern Computing*, Paul Ceruzzi had written the introduction. Maybe people would expect to see something on the history of AI, but he thought the history of AI would be best written by machines. <laughter> That was how he got out of it in the first edition, and I'm going to see if we can maybe move a little bit beyond that.

Grad: When was that, 1995 or so?

Haigh: I think it appeared in 1998.

Grad: Just exactly 20 years ago, yes.

END OF THE INTERVIEW