

An Interview with  
EDWARD FEIGENBAUM

OH 157

Conducted by William Aspray

on

3 March 1989

Palo Alto, CA

Charles Babbage Institute  
Center for the History of Information Processing  
University of Minnesota, Minneapolis

Copyright, Charles Babbage Institute

Edward Feigenbaum Interview  
3 March 1989

Abstract

Feigenbaum begins the interview with a description of his initial recruitment by ARPA in 1964 to work on a time-sharing system at Berkeley and his subsequent move to Stanford in 1965 to continue to do ARPA-sponsored research in artificial intelligence. The bulk of the interview is concerned with his work on AI at Stanford from 1965 to the early 1970s and his impression of the general working relationship between the IPT Office at ARPA and the researchers at Stanford. He discusses how this relationship changed over time under the various IPT directorships and the resulting impact it had on their AI research. The interview also includes a general comparison of ARPA with other funding sources available to AI researchers, particularly in terms of their respective funding amounts, criteria for allocation, and management style. This interview was recorded as part of a research project on the influence of the Defense Advanced Research Projects Agency (DARPA) on the development of computer science in the United States.

## EDWARD FEIGENBAUM INTERVIEW

DATE: 3 March 1989

INTERVIEWER: William Aspray

LOCATION: Palo Alto, CA

ASPRAY: I'd like your comments about how you first became involved in getting DARPA support (ARPA at that time), and what your interactions were with the program officers at the time.

FEIGENBAUM: The first thing to say is that for me the period of coverage that you're interested in starts at 1965, so you're talking about 1965 through the early 1970s. That's quite a while ago. We're talking about 20 years ago, so details dim in the mind. I'll do as best as I can, and salient personalities will come up, but I certainly won't be able to remember all the individuals and my relationships with them. In any case, it all started out with J.C.R. Licklider, whom I didn't know, but only knew through his writings in psychology as a psychoacoustician. Licklider came out to Berkeley, where I was teaching at the time. I was at Berkeley from 1960 through late 1964. Licklider came out to Berkeley. Licklider approached me about setting up a DARPA supported artificial intelligence project.

ASPRAY: At Berkeley?

FEIGENBAUM: At Berkeley. The other person he came out to see to stimulate a DARPA project was Professor Harry Huskey, to stimulate a project in timesharing. So he came with those two things on a trip to Berkeley. Now the question is, where did he get hold of my name? And how was I known to him? I was a young assistant professor barely visible above the surface. The answer is through graduate school contacts. Obviously, Herb Simon and Al Newell were not invisible to Licklider; they were probably good friends of Licklider through the psychology channels. Licklider was asking around and wanting to put together a program. One way you put together a good program is to find the good people. He probably asked Newell and Simon (I've never checked this out with them) "Did you have any good students? Is anyone doing great work around the country?" They probably said, "Feigenbaum at Berkeley," and Licklider came out to check that out and promote that idea. Ditto with Huskey, who was well known in computer systems. Licklider knew McCarthy quite well through the old MIT and BBN

connections. Licklider came out to see if he could get McCarthy in on the principal investigator team that we now call the "community." Immediately after Licklider set us up with some funding at Berkeley, Harry Huskey went on sabbatical to India for a year, which was a kind of crushing blow. You don't really start up a new project and then go off to India for a year. So another professor... a brand new guy that had just been hired at Berkeley, David Evans -- later to become very famous as a founder of Evans and Sutherland -- took over Huskey's role for that year. Since Huskey and I were listed as co-principal investigators of the Berkeley grant, Evans got me involved with the design of the timesharing system, which eventually became commercialized as the SDS 940 timesharing system. So, together, Dave Evans and I picked the SDS machine. Of course, Dave was a great hardware engineer and I didn't know anything about that, but I got a lot of education from Dave during that process.

ASPRAY: Did you actually participate in this?

FEIGENBAUM: Yes, a lot. However, I disappeared from the scene at the beginning of the fall quarter at Berkeley in 1964. So here's Huskey on sabbatical and here's Feigenbaum leaving the whole thing to go to Stanford. The rest of the SDS project that became the 940 proceeded after I left, so I really didn't have much to do with it -- and nothing after that time.

ASPRAY: Can I ask a question about the timesharing system? Was this an attempt to get centers working on timesharing that could independently develop ideas? There was already the active group at MIT.

FEIGENBAUM: That was correct. Exactly right. Licklider saw the MIT group as the large machine group. He was going to set up a mid-size machine group, which was going to be us at Berkeley. We were a competitive group to MIT; competitive in the sense of working on a different class of machines. That definitely was in Licklider's mind. He was quite explicit about that. Then I came down to Stanford. I had no trouble whatsoever. I don't remember the slightest bit of trouble in getting my funding either started or moved from Berkeley to here. It was continuous without losing a step. I moved to Stanford January 1, 1965, and it is approximately at that time when Ivan Sutherland took over from Licklider as head of the office. It was quite a shock to us all because Ivan was so young. For him to

be office director seemed anomalous, but Ivan was brilliant and reasonably well known to us -- I wouldn't say a friend, because I didn't know him that well, but he was well known to us through his work on Sketchpad. It was easy to deal with Ivan. He understood our needs. He appreciated the research having come out of that kind of research background. McCarthy opted to create a very substantial laboratory at the time, and Ivan gave it his blessing. My part of the effort was quite small at the time. It was the original Dendral work that was the grandfather of all expert systems. That was the first expert system, although we didn't call it that at the time.

ASPRAY: Small in the sense that you had not large machine requirements? Or not very many students and programmers working for you? Or all of those?

FEIGENBAUM: We didn't have many programmers and students. We had a substantial machine requirement, but then McCarthy was getting a PDP6 as part of his laboratory, so it wasn't going to be a worry about how to get enough computer time to do the job.

ASPRAY: Okay. How much did you see Licklider and Sutherland along the way? Did they come to visit you regularly? Were they participants in the writing of your grants when you had to go back for additional money to DARPA?

FEIGENBAUM: My vague recollection after 20 years is that we hardly saw them at all. It's also my vague recollection that although we would not submit a proposal to ARPA without having passed it by those guys first for comment -- reworking, criticism -- in no sense did they help to write it. It was all our stuff. They might help to shape it so that it went through the ARPA front office better.

ASPRAY: But the technical content was...

FEIGENBAUM: ... was all ours -- 100%. When Fano gave his speech about Project MAC at the 25th anniversary celebration of Project MAC in October, he mentioned that the initial relationship between ARPA and Project MAC

was a handshake between himself and Licklider. The kind of proposals we were sending into ARPA were more than a handshake, but little more than a handshake. They were short statements of what we thought was important to do, and, roughly, how we were going to do it. That was good enough for Licklider. And that was good enough for Ivan too. This is in dramatic contrast to the way it is today. The way it is today it's almost like you were proposing a super-collider -- massive documentation, milestones and deliverables, and all the arcane stuff that goes along with careful and prudent management of government funds.

ASPRAY: This may not be relevant because you only had the Dendral program going, but I had the sense that a number of these early grants that the IPTO office made were umbrella grants where the money was there to work on a whole range of problems, and you didn't have to carefully identify exactly which ones and what percentages were being devoted to each of them, but the money was there for the decision of the group on-site.

FEIGENBAUM: Well, mine was small, and it was particularly aimed at certain particular questions that I was posing. However, having said that, the spirit of the money was just exactly as you described. The spirit of the money was when Ed wanted to change, there was some good reason for doing that, so no one was going to question that he did it. Or if he decided halfway through a year that suddenly the hottest problem had become a machine learning problem, which he never described in his proposal, that was fine with them. It was better described as "center of excellence" money than it was as research contract money, although they never called it that -- not here. At Carnegie-Mellon they called it a center of excellence, I think.

ASPRAY: Okay. From what Licklider said, I understand that it was also a term that had been used by Licklider very early on when he was thinking about how to organize the office.

FEIGENBAUM: Now, you posed the question of whether funding from ARPA was any different in kind from the funding that came from other places. Today it's commonplace to observe that for the normal investigator-instigated proposal that is peer reviewed in a normal program, the National Science Foundation funding comes maximally to \$150,000 or \$200,000 a year. That's today. At that time, that number wasn't 150,000 to 200,000; that number was about

50,000 at the time -- 50,000 to 75,000 was that same range -- the most you could expect. A typical grant might have been 25,000 or 35,000. That number, at the time, like the current number, is peanuts compared with what you need to do a threshold level job in systems-oriented AI research. I'm not talking about theoretical AI research that you might do with pencil and paper, but where you're building big, experimental systems. That money was much too small, and the most it could be used for was dealing with a very small project, like maybe one research scientist that I would hire, one student, and a fraction of my salary. That's a little project. Or it could serve as seed money for getting prepared for launching a bigger initiative that would be funded by DARPA. We've had NSF money on and off. But it is just far below the level of funding that I could conceivably want. It has always been political in two senses. One is that the peer review process is highly political. When I served on the NSF Computer Science Advisory Board and we were the review panel for what the peer reviewers and the administrators had done, you could see this in computer science. Computer science not being a discipline of well established cultures and traditions, you would find wide dispersion of opinion, so that the average ranking of a computer science grant, let's say on a scale from one to five -- would be around three, with some of the peer reviewers saying it was excellent and some saying it was terrible. It would average out to about three, whereas in other disciplines where things are much more clear-cut you'd find that the ones being granted peaked in around the 1.2 level, and there were a whole bunch of 4s and 5s which was the lousy stuff. They weren't all meshed in around three. So that's political from that point of view. Political from the point of view of pleasing the congressmen, we used to get the grants presented to us at the meetings in *alphabetical order by states!* ARPA was not that way at all. ARPA was focused on excellence. Now, people outside the ARPA community of investigators would say that that's an insider's view, that in fact it was highly political; it was just that there was an "in" group and an "out" group. If you were a student of Newell, or Simon, or Minsky, or Fano, you were in. That was very political. That was political in the British sense of being a minister by virtue of where you went to school, and your accent. So in that sense you could say that ARPA was political. Now, as far as other sources of funds go, from the early 1950s, one of the most benign and forward-looking sources of funds was ONR. During the period you're talking about there was a wonderful ONR computer science executive, Marvin Denicoff, who supported a number of investigators, some inside the ARPA community and some outside the ARPA community, but always with small amounts of money. So I never even went to Marvin. I never needed an amount of money that small.

In Licklider's second term, when the first results were coming in from Dendral and they looked really good, Licklider came to me and said, "Ed, this is ridiculous. You're asking for an amount of money a factor of two too small. You want to grow. Double your budget." Now, I'm a very conservative and incremental person. I always thought my budget should grow at about inflation rate. I like to keep my head below the wall and not try to get shot at by asking for too much money or in any way making funding waves, and here is Licklider telling me to double my budget request, which I did and we grew. And later on, we effectively doubled again. We're talking about, say 200 to 400 to 800 thousand. I think those were the steps, although the 400 to 800 didn't happen all at once. It happened over a stretch of time. Now what were the other candidate funding places? Air Force Office of Scientific Research. I never made any contact with them.

ASPRAY: Do you know if anyone in the AI community had support from them?

FEIGENBAUM: Yes. I think so. I think Nils Nilsson had support from them when he was at SRI. I think there was a wide variety of people who had support from them, but as an ARPA investigator I never did.

ASPRAY: Okay. NIH, I suppose, is another.

FEIGENBAUM: NIH, yes, of course. NIH over the years has been a bigger supporter of our work than ARPA, believe it or not. Well, the key player in that institution is the Division of Research Resources (DRR), now defunct, which was an institute-like division of NIH that was chartered to fund expensive computer resources, expensive instrumentation resources, and expensive animal facilities and clinical care facilities. DRR isn't related to an organ or a disease. It is related to big ticket items. A particular program in there was a program on biotechnology. Its most recent incarnation was called BTRP (Biotechnology Resources Program). I guess it was called BRP at the time. It was headed by a superb person, the best government official I had ever met up until that time. His name was Bill Raub, now the Acting Director of NIH. He was imaginative, forward-looking, and highly supportive of shared computing resources. Joshua Lederberg at the medical school and I (as director of Stanford's computer center) put in

a grant application for a shared resource for medical computing in the medical school. We built a timesharing system called ACME on IBM equipment, supported by Bill Raub's BRP.

ASPRAY: When was this?

FEIGENBAUM: This was in the 1966-67 time frame. Subsequently, Lederberg and I transitioned that activity from a medical school timesharing system to a new timesharing system on a PDP10 -- the PDP6 era had just ended; we were transitioning into the PDP10 era to support a community of investigators doing artificial intelligence research in the areas of medicine and biology.

ASPRAY: All at Stanford.

FEIGENBAUM: No. Half at Stanford, and half around the nation. The way that we would share this resource around the nation was through a new-fangled contraption called a computer network. ARPA had just developed one. Our AI lab was on that network. I knew all about those networks. I could get Licklider's approval in his second term to hook this NIH-supported machine to the ARPA network, because we were also ARPA supported, and NIH reciprocated by supporting ARPA-sponsored work with their computing resources, so Licklider gave the permission to hook this facility to the ARPANET, and it became what is now a very famous computer facility called SUMEX-AIM. SUMEX stands for Stanford University Medical Experimental Facility. AIM stands for AI in Medicine, which was the name of the national community. We had a two-part management structure that involved an executive committee here at Stanford and a national executive committee. The chairman of the national executive committee throughout its entire life was Don Lindberg, "Mr. Computer Science and Medicine" from the University of Missouri, now director of the National Library of Medicine. Through the years, NIH put a lot of money into SUMEX AIM, its core research, and separately funded work for the Dendral project. By this time, the Dendral project had transitioned into the principal investigatorship of Carl Djerassi, the famous chemist. So there we were at a time in which money was fairly scarce and given in small amounts, collectively getting easily a million dollars a year for those two efforts from NIH. That has kept going. We're currently still funded by NIH on the SUMEX-AIM project until summer of

1991, I think, at the level of a million and a half a year. NIH has been a vigorous supporter over the years. The contrast between the two is that no matter how supportive and benign the administrators are, NIH has a deeply rooted peer review system. They have a lot of trouble doing things that the peer reviewers don't want them to do. So NIH support, while it has been generous, has always been a struggle, a lot of work. When we would have to renew our NIH money, we would just salt away a half a year of our life to get that NIH money renewed. DARPA money was never like that. DARPA money was a matter of making sure that the person in charge of the office, like Licklider, or Larry Roberts, or Bob Kahn, knew what you were doing, and that the program managers down below were fairly routinely informed of what you were doing either by electronic mail, or by visits. Then the next proposal generally was handled routinely. There were two periods in the time frame you're asking about in which there were major exceptions to that. One was what are euphemistically called the "Heilmeier years" (1973-75) when George Heilmeier, now Chief Technical Officer of Texas Instruments, was Director of DARPA. During his tenure in office he was quite harsh on AI. He subsequently became AI's most vigorous supporter in the industrial world. He thinks of himself as "Mr. AI in industry," and he really deserves that title. The story on that is told in the Texas Instruments chapter of my new book with McCorduck and Nii, *The Rise of the Expert Company*. We interviewed Heilmeier. We tell that whole story. But the Heilmeier years were rough in defending AI and getting proposals through it. Still, I didn't suffer any budget cuts that I can remember, or if there were budget cuts they were small. The other time was around 1970 or 1971, when Steve Lukasik, then Director of ARPA, began to ask the question, "What has AI done for me lately?" We had to really get our act together and answer that question in written form for Lukasik.

ASPRAY: What kind of answer were you able to give?

FEIGENBAUM: Well, the answer was quite broad and involved a vigorous defense by many brilliant people of a range of things from cognitive science all the way to computer speech understanding -- vision and so on. But actually, the big win was having Dendral, to show that the dream of a really intelligent machine was possible. There was a program that was performing at world class levels of problem solving competence on problems that only Ph.D.s solve -- these mass spectral analysis problems. All we had to do was tell the Dendral story, and that was enough of a promise of good things to come that that convinced Lukasik. So, what people really said was, "Dendral saved AI in

the first wave of criticism, and MYCIN, which was the medical diagnosis system we did here, saved AI on the second wave of criticism." Of course, that is a very Stanford-centered view of events and others will cite other work that "saved AI" in those difficult times.

TAPE 1/SIDE 2

FEIGENBAUM: Let me mention, in connection with this, that during that period when Lukasik was pushing on AI, the critical person there handling the problem was a man named Steve Crocker -- program manager for AI. I believe Steve had just followed one of our former students, Cordell Green. Cordell had gone to ARPA, worked there for a while, then left to come back to Stanford. Steve, I think, took over from Cordell. I remember just endless interactions with Steve Crocker about a big, fat document -- a "What is AI?" document -- that we were putting together for Lukasik. That document is possibly in my archives at the Stanford Library.

ASPRAY: Good.

FEIGENBAUM: It's in its computer printout form. Of course, we were using text processing a long time before other people, and it's in that form. We would ship drafts back and forth over the ARPANET, get lots of people to contribute sections. And I was kind of the academic side pulling things together. He was the ARPA side helping to facilitate things. By that time we had gone through a transition period called Bob Taylor. In spite of what you might gather from current rewritings of history that have to do with the fact that Bob Taylor became quite well known for his work at Xerox PARC, at that time he was considered an anomaly in the office; namely, someone was being promoted to the directorship of IPTO that didn't have a doctorate and didn't have any background in computer science. Imagine, from Licklider to Ivan Sutherland, and then all of a sudden, Bob Taylor from NASA -- a management background from NASA, what you would call not a fully credentialed person. Now, the story of how bruised Bob Taylor felt about that is told in a book called *Fumbling the Future -- How Xerox Invented the Personal Computer and Ignored it*. It's somewhat a book about Bob Taylor. There's a story in there about Bob's tour of duty at ARPA. It was an anomaly. The anomaly was solved and the community of investigators was calmed by virtue of

the fact that a "part-time consultant" called Chief Scientist was brought in, named Larry Roberts. Well, we all knew Larry. Larry was the obvious successor to Ivan Sutherland. The two of them were graduate students at MIT at an earlier time. So the fact that Larry was called Chief Scientist and Bob was called Director of the office made things okay.

ASPRAY: As I remember that, it was a four year period that Taylor was the director -- 1965 to 1969.

FEIGENBAUM: No, that can't be right. It was Licklider who stepped down in 1965. Ivan took over in 1965.

ASPRAY: 1966, that's right, because Ivan left in 1966 to go to Harvard.

FEIGENBAUM: Taylor's tenure was 1966-67, a two year period. Bob was very competent, very good, very benign, listened a lot. For me, Bob made a crucial, important decision. I made the case to him that the Dendral project needed to be separate from McCarthy's AI laboratory for various reasons -- some logistical, some political. He said, "Fine." He just separated the money into two separate contracts that allowed McCarthy to be one PI and me to be another PI. It allowed me to be my own person here at Stanford. So I'm forever grateful to Bob Taylor for that.

ASPRAY: Do you want to say a little more about Taylor or the end of Taylor's time? And then maybe we could take this up again some other time.

FEIGENBAUM: No, I just wanted to say that it was during the Larry Roberts time that Lukasik was putting all this pressure on. It wasn't terribly bad pressure. Lukasik, I thought, was a marvelous director of DARPA. He was seriously evaluating whether AI was one of these normal things that you support and then you drop, or whether AI was something which he would support for a very long time. At the end of all of this Lukasik gave a speech to the Principal Investigator's Conference in Los Angeles. I remember it so vividly. He stood up there and he drew the "curve" of the normal DARPA project which goes to zero at the end of five years. Then he drew a horizontal line across the board and he said, "This is AI. This is my long-term investment." So that was really a turning point.

ASPRAY: This would have been the early 1970s?

FEIGENBAUM: Well, 1972 is what comes to mind. Then, the other thing that occurs to me is when I mention principal investigator's conference, all of us agreed that that was the best conference we went to all year long. The ARPA PI's got together in a room for a few days and talked science, and helped the ARPA people plan their next year or two of projects. The latest ARPA principal investigator's meeting held in Dallas last October had 300 people there. And that's an impossibly large group. We had a ballroom to hold all those people. That was a very unfortunate kind of principal investigator's meeting. But the old ones, where we had Allen Newell, and Alan Perlis, and John McCarthy, Marvin Minsky, myself, and a few people from the networking area, Dave Evans and a few other people from the graphics area, others from the time-sharing area; those were terrific -- 20 or 30 people; absolutely stellar, wonderful meetings. Some people from SRI... I guess Doug Engelbart used to come and talk about his work. (NLS, an editor in which screen-oriented work was coupled with a mouse for the first time.) Anyway, those were great meetings.

ASPRAY: When you mentioned Lukasik and the timing, what comes to my mind immediately is the Mansfield amendment. Did that cause some problems for the AI research -- directly or indirectly in terms of is this really defense-related research, or is this just basic research that the Defense Department is supporting?

FEIGENBAUM: The following things happened, short term: for the first time we started to worry about military relevance of our work. That is, to start putting together words to indicate the military relevance. What we all knew about the military could be put in a thimble. So all the stuff we crafted in this regard was silly. Second step: the DARPA people looked at it and said, "That's silly. We can't write that stuff. This is ridiculous. We'll have to take over that job for you." So they started to craft all the language, which they would then slap on to our proposals. We would never see that language. They would write all that stuff. The only people who ever saw it were the students who would later dig it up under Freedom of Information Act. It wasn't in the Stanford version of the proposals. It was only in the stuff that went up to the DARPA front office. Somehow, they'd manage to get a hold of it and then

they'd bring it on campus and say, "See, McCarthy is working on such and such, and such and such." McCarthy would say, "What do you mean? I never heard of that?" (laugh) Okay. That was number two. Number three: the Mansfield amendment dates the start of the bureaucratization of ARPA. It's the end of Eden and the beginning of the real world. ARPA up until the Mansfield amendment was Eden. It simply was trouble-free paradise for working, getting your ideas approved, getting funding for them. Then the Mansfield amendment was only one tiny little nail, but after that there have been endless numbers of nails that have made ARPA very bureaucratized, slow, much more difficult to deal with.

END OF INTERVIEW