An Interview with

BRUCE G. BUCHANAN

OH 230

Conducted by Arthur L. Norberg

on

11-12 June 1991

Pittsburgh, PA

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

1

Bruce G. Buchanan Interview
11-12 June 1991

Abstract

Buchanan describes his work in artificial intelligence, the development of the Stanford Artificial Intelligence
Laboratory and the artificial intelligence (AI) community, and the role of the Information Processing Techniques
Office (IPTO) of the Advanced Research projects Agency (later the Defense Advanced Research Projects Agency)
in AI research. Buchanan describes the work of Ed Feigenbaum, Josh Lederburg, Wes Churchman and Les Ernest at
Stanford. He discusses the changes in AI funding, including developing additional NIH funding, with the Mansfield
amendment which stipulated defense supported research should have defense applications. Buchanan concludes
with a comparison of artificial intelligence and computer science development.

DATE: 11 June 1991   INTERVIEWER: Arthur L. Norberg

LOCATION: Pittsburgh, PA

NORBERG: Professor Buchanan, can we begin by your reflecting back on the middle 1960s and your decision to go to Stanford University?

BUCHANAN: Sure, the story begins actually in 1964 when I was a graduate student in philosophy at Michigan State and I was looking for a summer job. I interviewed with people from System Development Corporation. It was an interesting kind of job, essentially policy analysis that required some mathematics. I was a mathematics major as an undergraduate. But it was strictly driven by a fiscal reason. I didn't hear from SDC, but I did get a call from Ed Feigenbaum who was at Rand and was spending summers at Rand. Apparently, SDC had sent my application over to Rand and someone thought it was bizarre enough to interest Ed, because I was interested in questions of theory formation. I was doing a dissertation on scientific discovery. I had some ideas that it shouldn't be as mysterious a process as the psychologists were making it out to be with a creative, "Ah, ha, Eureka!" experience. It was totally mysterious. So I was trying to be somewhat systematic in looking at the process of discovery, but not from a psychological point of view, more from a logical point of view. And Ed called me. I don't know if you have ever had any phone conversations with Ed, but sometimes they're a little unusual in that there are long pauses. And I was intrigued by someone who was clearly thinking so much about what I was saying. I think he was actually thinking about something else, but anyway, on the basis of essentially one long phone conversation with him I packed up and spent the summer at Rand, the summer of 1964. Ed and I were working then on a program that would learn from examples as a prototype for theory formation. And I went back to Michigan State at the end of the summer, having learned a lot about computing and been exposed to lot of ideas in AI. There was very little written, but most everything that Newell and Simon had done came out as Rand Reports and so I read all of those when I was there. There was an excellent library at Rand, and a very helpful staff. It was a remarkable set of resources for me. And Ed and I talked frequently, and Tonge was there.

NORBERG: Tong? T-O-N-G?

BUCHANAN: Tonge. Rand had done their two or maybe three summer institutes in AI the previous summers. This was not one of them, but there were, I think, three graduate student interns. Bob Balzer was one that summer. Oh, yes, Elliott, who went to Florida State, but didn't stay in AI, and myself. Anyway, at the end of the summer, I began working on my dissertation again and it changed shape, because I now saw that the ideas that Newell was talking about were in fact a mechanism by which we could define the process of discovery. So those two things came together, and the second half of my dissertation is quite different than the first half. But still, I was I guess set on a career of teaching philosophy. I asked Ed if he would write a letter of recommendation for me for teaching jobs. He said he would be delighted to, but he and Lederberg had just started a project in scientific discovery, theory formation, and wouldn't I like to consider coming out for a year? My initial reaction was, "No, I really should be teaching philosophy." But a second conversation with him convinced me that it was worth a year. It was exciting. Ed was clearly excited by it. He and I worked well together. He had read a lot in the philosophy of science and was clearly very interested. A chance to work with Lederberg was intriguing because I certainly had heard of him. I had no idea what it would be like. So I decided to give it a year and the first year... first year and a half I was there on a Stanford appointment. We lived in Berkeley, because Ed had been working with Wes Churchman at Berkeley. So I interacted with Wes quite a little bit in that first year and tried to carry on some of the ideas he had been writing about in his book on inquiring systems, and mapped that into the Lederberg, Feigenbaum thinking about data interpretation, and essentially driving back and forth, and finally then I decided to move down to Stanford.

NORBERG: Well, let me ask two questions. How did the interaction with Feigenbaum go on? What was the sort of daily contact that you had with him? What was the working style of the two of you?

BUCHANAN: It was very congenial. Ed is not given to social banter, and that suited me just fine. So the interactions were intellectually intense, and I always came away feeling that I had a lot more to think about. No

matter how much I had prepared, it wasn't... He always had the incisive question that made me feel I had blown it, but that wasn't his intent. He really was just exploring. Ed was busy with a lot of things, so I saw him less frequently than daily. The Rand experience, until my security clearance came in I was out in the trailer in the parking lot, and didn't see Ed very much at all. But when he came out it was always to talk about stuff that I had been reading, or IPL IV, which I had been working in. So my recollection was that we met a few times in the first month at Rand until my clearance came through. And at Stanford, too, he was the director of the computation center, so I was taking a lot of his time. We had weekly project meetings, and probably two or three other meetings in a week of roughly an hour.

NORBERG: Now, what was the Lederberg, Feigenbaum project in those early years? It certainly hadn't advanced to DENDRAL.

BUCHANAN: I think in Lederberg's mind it was the DENDRAL project, because Lederberg had a NASA contract for data interpretation to try to find evidence of life on Mars. And Lederberg believed that mass spectrometry would be powerful enough to provide such evidence if it were there. He realized that the current method of interpretation was catalogue look-up, and that that was certainly limited on a new planet with an earth-based catalogue. So he felt that he needed to develop some first principles for interpreting mass spectra. And that was, in fact, the whole DENDRAL project laid out with a very specific application, namely, to go to Mars, scoop up samples, look for evidence of organic compounds. But one of the critical pieces was the DENDRAL notational algorithm that Lederberg had developed earlier, and then programmed, I think it was in 1963, possibly he didn't finish that until 1964, but that was the heart of the whole DENDRAL program, and that notational algorithm was a unique naming convention for organic molecules, for any chemical molecules, and at least for the ones without cycles. It allowed a step by a systematic generation, just a simple translation of the notation into a prospector [?] generator. And that's what we were working with initially. So if you think about generating a test paradigm in AI, here was the generator. It was just being handed to us without a whole lot of work on our part. It was amazing.

NORBERG: So, what was your role in this then? In those early years, now?

6

BUCHANAN: Well, Ed and Josh felt that my thinking about scientific theory formation, and that included hypothesis testing. In fact, Ed and Josh didn't distinguish hypothesis generation and theory generation, or hypothesis testing and theory testing.

NORBERG: Did you distinguish between them?

BUCHANAN: Yes, I did, but I didn't insist on it, because I saw that if we could... and I thought that working with single hypotheses was certainly easier than collecting them together and unifying them in some sense. And we had to start there. We couldn't jump into a whole theory formation exercise yet. But my role, then, was, I guess, twofold. I was initially starting to encode a predictive theory of mass spectrometry, so that the candidate hypothesis coming out of the generator could be given to a prediction program to predict testable consequences that we could then match against the original data. So I spent the first couple of years encoding that predictive theory, learning LISP also. LISP was new to me. Georgia Sutherland was working on the LISP implementation of the generator and building books into it so it could be constrained. And Georgia and I certainly worked on each other's code, but very largely they were separate projects.

NORBERG: I interrupted you when you were saying that Josh and Ed didn't distinguish between hypothesis testing... hypothesis generation and theory generation. What was the implication of that? I interrupted you before you could get to the answer, I think.

BUCHANAN: I think operationally there was no important implication, except that I felt a little confused having grown up with a distinction. It just felt that it should be maintained. But I could put aside those questions because the day to day work was so interesting and seductive. It opened up a whole new world.

NORBERG: How?

BUCHANAN: All the stuff that we read about in the philosophy of science we were able to begin to operationalize, as we say now, and it provided a means for understanding in this context what we had to mean by simplicity, if we were going to test the simplicity of hypothesis. We had to be very precise; we couldn't just wave our hands. And the in-principle arguments that we grew up with from Carnap and others just aren't good enough for computational means. We have to have a small set of hypotheses to test. We can't be testing some power set of hundreds of billions. But that part was fascinating. But I was, I think, clearly seduced by the medium. Here was an opportunity to create something de novo and see the product in action. And I think everybody around the AI lab was feeling the same way. It was just this exciting.

NORBERG: Now, around the AI lab, I assume you mean the AI lab at Stanford, but we started out this conversation by you telling me that you were spending time with Wes Churchman...

BUCHANAN: I was.

NORBERG: ... at Berkeley. What was the split here? Were you essentially spending all your time at Berkeley during the first year, or year and a half?

BUCHANAN: It started out I was spending three days a week at Berkeley in the management science lab in the business school - Barrows Hall - and two days a week at Stanford. Then I shifted to two days a week at Berkeley, and then one day a week at Berkeley, and finally we moved to Palo Alto. The Management Science Lab certainly... There were people there who were interested in AI, but I maintained a teletype link to the Stanford PDP-6. And then I got time on a PDP-6 and later PDP-10 over in Electrical Engineering at Berkeley, and I could have all the time I wanted in the hours between midnight and 7 a.m. So I changed my schedule radically in my Berkeley days, and essentially started my day about 3, there all alone in the dark in the basement of Corey Hall. But time just flew. It was extremely exciting. What Lederberg had asked me to do was look at the mass spectra of the 22 known amino acids and try to

8

build the predictive theory such that our program could predict the major features of those spectra. And we did have something, but it was never as good as we wanted. I didn't understand then the stochastic nature of mass spectrometry well enough, and there were elements in chemistry that I believe were not well enough understood quite by anyone. And I thought they must have been understood and I was trying to do that.

NORBERG: How was this going to be used once the system was designed? Was information going to be assessed on-site on Mars?

BUCHANAN: That was Lederberg's dream. A mass spectrometer would go to Mars. The samples would be scooped up and micrograms put into a mass spectrometer. The program would analyze the data on-board, and that would direct the next experiment. Now, in the end there was a mass spectrometer that was on-board, but the signals were sent back to earth and analyzed over the successive years with catalogue look-up.

NORBERG: Yes, I do recall that part of it, but didn't realize that the origins of DENDRAL were for this purpose.

BUCHANAN: It was very clear. We were motivated, we were actually trying to work on a timetable so that this could happen. Sending a PDP 10 to Mars was not possible, and we realized, I guess, rather quickly that the physical constraints would not let us do it, but that didn't change our enthusiasm.

NORBERG: Let me take you off that topic for a moment and ask you a couple of other sort of institutionally-directed questions here. Who else was in the AI Lab at Stanford at the time?

BUCHANAN: McCarthy and Feigenbaum were co-PIs. Les Ernest was the manager. If you haven't talked to Les, I think you would find him to be a remarkable resource, because John was not particularly interested in the administrative aspects, as you would guess, and left most of that to Les, almost all of it. And Les had a remarkable ability to translate whatever people were doing in whatever field in terms that the ARPA program managers would

find had some relevance. So a lot of that was at a local level. Let's see, other major investigators... Raj Reddy had just finished, and was an assistant professor working on speech understanding. He had a group there that included Lee Earman. Jerry Feldman had just finished, and was there. Kent Colby was there working on PARRY. It's a simulation of a paranoid patient. Colby is a psychiatrist. There was a robotics group. Dick Scheinman was there. And there was the Les Ernest project too. There was work on chess. Barbara Liscov Hubermann was an early McCarthy student. There was quite a bit of systems work going on to make the PDP-6 into a research machine that was reliable. Jerry Feldman came soon... I can't remember the timing very much. And then some time around 1970 or 1971 Terry Winograd came. I believe that's about 1971 or 1972. I did an introduction to the COMTEX series. Do you remember that?

NORBERG: Yes, I have that.

BUCHANAN: Okay, and that names quite a number of those early people. It was an exciting place.

NORBERG: Yes. Now, you suggested that there were two forms of funding for this organization. I don't know how much you knew about it at the time. One of them being the DARPA money and the other being NASA funds. Is this a correct recollection? Did you know at the time where funding was coming from?

BUCHANAN: For the DENDRAL project? Yes, partly because Lederberg's Instrumentation Research Laboratory in the medical school was NASA funded. Several of the people we interacted with on instrumentation questions were there. Elliot Leventhal was the head of that laboratory. And then physically separate in the AI lab we were dealing with ARPA funding.

NORBERG: On the same project? That is, were funds being commingled for the same project?

BUCHANAN: Yes.

10

NORBERG: They were.

BUCHANAN: Yes.

NORBERG: Do you have any idea what the split might have been?

BUCHANAN: I have no idea.

NORBERG: I am not sure it's very important, but I was curious.

BUCHANAN: Well, it lent some stability, I am sure.

NORBERG: What sort of interaction did you have with DARPA people at the time?

BUCHANAN: I had almost no interaction. I have no recollection. The machine intelligence workshops that Don Amicci [?] ran and Bernard Meltzer [?] were the places where I interacted with other members of the ARPA community, and I don't even recall if there was representation from the ARPA offices there. There might well have been.

NORBERG: There are several directions we can take this at the moment, but let me go back again and ask you about the development of AI in that period. It would seem to me that by the middle 1960s that artificial intelligence as a research community in the United States was a recognizable group of people and that the set of problems that this group had identified for themselves was also well known. Is that a fair statement?

BUCHANAN: The community was, I think, pretty well recognized,relatively small. It included just MIT, Stanford and Carnegie Mellon, with a few out-liers, but essentially those.

NORBERG: Does your statement about Stanford include SRI as well?

BUCHANAN: Yes.

NORBERG: That's all one group.

BUCHANAN: Yes.

NORBERG: All right.

BUCHANAN: Although, in fact, Stanford, maybe every place, included separate little fiefdoms. Yes, I do include SRI.

NORBERG: So there were just these three groups in AI. You're talking about the size of the community.

BUCHANAN: The community being well-defined. The research problems, I think, were not so well-defined. In the large, anything that had to do with mechanizing intelligent behavior was fair game. And there were a number of people in this community working on a game plan. A number of people working on logical inference and mechanizing theorem proving, and sometimes people were doing game playing through theorem proving. So, no, I don't think the problems were terrifically well-defined. What we were doing with DENDRAL seemed to us, and I think to other members of the community in the 1960s to be somewhat different from what the rest of the community was doing. And we felt a good deal of difficulty in explaining the relevance of chemistry to AI. At one of those machine intelligence workshops, I think it was Danny Bobrow who asked, "What does this have to do with AI?" I remember

12

during a talk I gave at MI-5, I believe it was, somebody began criticizing the amount of chemistry that I felt I needed to introduce. And I will forever be grateful to John McCarthy for saying, "Just listen, will you?" John, I think, got the picture. Besides, John had no trouble at all understanding the chemistry. But no, we did feel somewhat at odds with the community.

NORBERG: Well, now, why would the subject matter be a problem for these people? After all, a number of the organizations interested in AI were dealing with different kinds of subject matters, if we think about across town here at Carnegie Mellon, they were interested in psychological issues. And the people at MIT were interested in other sorts of ways to mimic human intelligence. They were dealing with a different set of questions, I would think.

BUCHANAN: I believe the perception was that those other groups were dealing with general principles of thinking or reasoning or representation, while we were dealing with mass spectrometry. We knew that we were trying to represent the principles of mass spectrometry in such a way that we could extend the scope of the program. But there was a lot that was very specific to chemistry and what we were doing. The generator was not a general purpose, hypothesis generator. It was a generator of descriptions and chemical structures, and it wasn't good for anything else, except in the wildest stretch of the imagination.

TAPE 1/SIDE 2

NORBERG: Well, let me carry that in two directions then, if you don't mind, Bruce. First of all, could you characterize what the AI community believed they were about in the late 1960s in terms of what it is they were looking for?

BUCHANAN: There was a good deal of work on theorem proving and representing knowledge in an axiomatic system. Many of those people felt that they were on the track of general purpose mechanisms. So if we only axiomatized mass spectrometry, then we could turn it over to their general purpose AI programs and we would instantly have the intelligent behavior we needed. Now, I actually tried that. I didn't advertise it, because I felt that I

had failed in the past, but I tried to axiomatize my predictive theory of mass spectrometry. And I found that I just had too many hundreds of axioms. I was getting lost in the details, and it was not at all satisfactory. I gave up on the project, partly feeling that it was my fault, that I just wasn't clever enough. Suppes had been able to axiomatize quantum mechanics, and surely I could do mass spectrometry. And I couldn't.

NORBERG: Well, what's the difference between Suppes' approach to quantum mechanics and yours in mass spectrometry that prevented you from doing it but allowed him to do it?

BUCHANAN: I think maybe it was a level of understanding - his and mine, but also of the physics community and the mass spectrometry community, because the principles by which mass spectra were interpreted at the time, and I think may well still be interpreted, are phenomenological, very much heuristic rules. So it didn't allow axiomatization that way, I think. Anyway, Nilsson in particular felt that theorem proving was the answer to everything. Actually, the seeds of that are in McCarthy's 1958 paper on programs of common sense, in which he is proposing that a logical representation and theorem proving methods will suffice for getting programs to reason about events in the world. He hasn't given up on that at all, except he has found that general, first order logic is not powerful enough. He has been on that research track for a very long time. And the people doing psychological modeling were looking at general principles of human thought - symbol manipulation, long-term, short-term memory, matching; those general principles by which then turn something over, you have got an intelligent program. We were looking at the other side of what turns out to be the same point, namely, the very specific knowledge that you have to give to one of those engines. We had built a special purpose engine, but we were looking for general mechanisms for representing the knowledge.

NORBERG: Well, can you give a couple of examples of what the other people were doing who were interested in axiomatic principles from the beginning? What sorts of investigations were they doing that would have justified their position that you were wasting your time in chemistry?

14

BUCHANAN: There was a good deal of work in the Blocks World. It was an invention that came out of MIT that was a very well-constrained environment for these kinds of general purpose investigations. And if you have a puzzle involving block stacking, how do you solve the puzzle? You can axiomatize positions of three blocks and the movements of a hypothetical or real robot arm, and prove then that Plan A will succeed in stacking the blocks A, B, C. So that was one kind of investigation.

NORBERG: So the simpler the system the more likely the axiomatization [?] would occur.

BUCHANAN: Yes, Feigenbaum derisively called those "toy problems." Now, we know that there's immense power in simplifying a domain so that you can focus on the critical element. That's why biologists work with *e. coli*.

NORBERG: Sure, yes. Scientists in general find that to be more conducive to...

BUCHANAN: Yes, but... Lederberg did not want any of the details of mass spectrometry interpretation to be left out, because I think, first of all, he wanted a system that would work on real mass spectra. But second, I think he felt that in order to have programs contributing to science in a general way, we just couldn't afford to neglect the details in thinking about the design of those machines. And one example is in erroneous data points. A mass spectrometer is an electronic instrument, so there is noise in the system, and, yes, of course, we could always work with hypothetical, clean data, but then the system would not transfer and not scale up into real problems. So right from the start we didn't do that. We tried to take the messiness as it occurred, which gave us no end of grief, but I think it was the right decision.

There was a lot of work on the Tower of Hanoi in those years, a lot of work on cryptarithmetic, on ?   ?   ?   .
There was a lot of work on the eight [?] puzzle, so you see these small puzzles posed - well-constrained, small number of objects, absolutely guaranteed perfect information. Now, how can you solve the puzzle? Mass spectrometry was not like that.

NORBERG: But I would have thought that mass spectrometry is... I was going to say prototypical, but I think that's

probably not right... is a fairly well-defined taxonomic problem, similar to, say, chess, which also has a basic set of

rules, which, if you can identify them, you can then develop some sort of a program for simulating a chess game.

Isn't that true?

BUCHANAN: The rules were not so well understood.

NORBERG: For mass spectrometry? Even after 75 years of research in this field.

BUCHANAN: It wasn't 75 years, though.

NORBERG: Well, in terms of mass spectrometry it's 75 years, isn't it?

BUCHANAN: No, no, it was relatively new in the 1960s. The first text books in mass spectrometry were written in

the 1960s.

NORBERG: Well, what is it I am thinking of that was going on in the 1920s and 1930s in physics laboratories in terms

of studying mass spectra of materials?

BUCHANAN: That may well have been. It may be the analysis of chemical structure that was entirely a new

application.

NORBERG: I see.

BUCHANAN: And the physics are well-understood. Otherwise, they couldn't build the instruments, but those are not the same principles that you need for interpretation. You are essentially having to read the physics backward and it doesn't allow it.

NORBERG: I see. Now, what I just heard you say (I'm going to give you my interpretation now - not a statement about mass spectrometry, but about the field and the reaction of your colleagues to spending too much time in chemistry) suggests to me that what you saw then, and probably still see, if I am right, is that there were two basic communities here. There were the people doing axiomatic work to develop general principles for the investigation of intelligence. And there were a few people who were developing a system to analyze this very specific case of the spectrographic analysis of materials, and that you were not dealing, forced by the particular situation, in universal principles that could be applied to the other problems that the first group was addressing. Fair?

BUCHANAN: That's largely correct. There undoubtedly were other groups who felt they were not part of the mainstream, but we certainly did feel that we were not part of the mainstream in those early years.

NORBERG: How would you characterize the differences that have been ascribed to MIT and Carnegie Mellon's research in AI?

BUCHANAN: I am not sure how to interpret that.

NORBERG: All right, well, I am trying not to lead the witness, of course. It's my understanding that, in fact, there have been two different approaches to the study of artificial intelligence. One of them is the complex processing ideas that developed through the work of Newell and Simon and others, and the second is the sort of machine intelligence developments that are principally illustrated by looking at MIT. And that is the development of some sort of a dry machine, as Winston calls it. Dry meat for reproducing intelligence. We don't care how it's done at the substrate level; we're only interested in mimicking human intelligence at the macro level, and that that is machine-

17

based, as opposed to thinking about human intelligence, or human-based. Now, I see those as two characteristics of the AI community, certainly in the 1960s, and it goes on well beyond that, but certainly in the 1960s. And I am trying to fit this work on DENDRAL into one or the other of those, and I don't see it fitting in either case, which suggests there's a third pole in this.

BUCHANAN: I am not sure it's a third pole; maybe more of a synthesis. Remember, Feigenbaum came out of CMU, so he had a good deal of knowledge about the work on human information processing. And the Lederberg-Feigenbaum hypothesis partly underlying DENDRAL was that what a trained mass spectrometry analyst knew was what the program ought to know. So there was that transfer from human knowledge - thinking - into machine terms. On the other hand, we... maybe this was... I don't know whose influence, but both Feigenbaum and Lederberg came to believe it, that we didn't care if the program arrived at the right conclusions by the same line of reasoning. And, in fact, the Lederberg algorithm was not anything that a person could think through. People do not think that systematically about chemical structure. I believe that Feigenbaum and Lederberg had already worked out the synthesis of human thinking and machine-based reasoning before I got there. It was not ever an issue. We never felt that we were simulating a human analyst. And it was immediately obvious from the time I got there that Lederberg's algorithm was far superior to anything a person could do manually. And so, it never occurred to me that we should be simulating. On the other hand, what came to be known as knowledge engineering depended very much on understanding what it is that a trained chemist was thinking about, not necessarily the individual inference steps, but what patterns was he looking for in the data, and how did he link those patterns with partial conclusions?

NORBERG: I'll come back to that in a minute. I want to go back to your footnote that you were mentioning to me before we turned on the tape recorder in connection with the attitudes of DARPA as you now understand them, and ask you to comment in the same way, but differently, in terms of elaborating this interaction with DARPA that you didn't see in the 1960s. Is that true that you never met any DARPA people as far as you can remember at that time?

BUCHANAN: Not in the early 1960s.

NORBERG: We're talking about middle and late 1960s.

BUCHANAN: Yes, not in the 1960s.

NORBERG: Not in the 1960s, okay. Was there any discussion around the laboratory about support and about the interests of the funder?

BUCHANAN: Oh, a good deal. Les Ernest was acutely aware of the politics of dealing with any funding agency and almost all the support was from ARPA. Ed and John were going to the principle investigator's meetings, and so they were acutely aware of the attitudes within the Information Processing Office. It was called IPTO then.

NORBERG: Yes, it was. Now, this was discussed when they returned from PI meetings, or was it discussed by Ernest in the laboratory.

BUCHANAN: Let's see. Ed would convey to us in the DENDRAL project what other people were doing, and comments that he had received about the DENDRAL work. In fact, I think it was in one of those meetings that Danny Bobrow asked Ed, "What does chemistry have to do with AI?" Yes, but from my point of view those were points of information. They were not statements, not discussions steering our work.

NORBERG: When did they become discussions to steer that work?

BUCHANAN: In the 1970s, when we were asked to write the justifications of our research in terms of military applications. Then we began having discussions about what we felt ARPA's interests were and how they were selling their programs to Congress, and which of those buzz words we could use to describe data interpretation or any of the other work. We used to invent phrases that had a general purpose sound to them when the instantiation

19

was entirely within DENDRAL. We did that also to cover early funding of the MYCIN project. ARPA had zero interest in clinical medicine, and yet data interpretation, production rules, those were phrases we could use in order to justify using some ARPA money for this exploratory project.

NORBERG: Now, was anybody else funding MYCIN, as well?

BUCHANAN: No, not in the early days. It was entirely bootlegged. After we had a few years of research under our belts, we got NSF funding. We had some funding from a division within NIH, which is now called Health Care Technology... I have forgotten the name of the office. It was BHSSR at the time. But no, we had to work hard in the 1970s to justify using ARPA money for MYCIN. We didn't mention infectious disease or any of the substance of the knowledge base.

NORBERG: Well, it wasn't really necessary, was it, to get funding from IPTO? I would think they would be interested in the principles you were trying to explore.

BUCHANAN: We felt it even more strongly that it was necessary not to mention medicine. But again, as an exploratory project, it wasn't entirely clear to us what we were exploring. We didn't have nearly as clear a picture of the knowledge-based programming paradigm that we now have and can project back into the past.

NORBERG: Did your interaction with DARPA people increase when there was more concern about the justification - yours; I'm speaking of you, not the group. I want to know who you knew and who you didn't know, and so on.

BUCHANAN: Very largely Ed was doing almost all of the political administrative work between the DENDRAL project, or later the heuristic programming project and ARPA, and all of the funding agencies. It was, I now have come to recognize, a real gift to those of us who were mired in the programming details. So, yes, I met with ARPA

visitors, but almost always just at Ed's invitation, since he was meeting with them, there may be questions they would have that I could help answer. But also, it was just by way of helping me understand the whole process.

NORBERG: During these meetings, then, with any DARPA visitors who happened to show up, did you have to avoid discussion of the subject matter? Was medicine not mentioned? And if so, did you people rehearse these things before hand?

BUCHANAN: Yes, we did. We knew that there were lines in the research that we wanted to emphasize and specific words and pieces of work that we should not mention - that they would raise flags. It was just generally better to avoid that kind of controversy. A very minor controversy that came out was in a set of three reviews of Weisenbaum's book. I don't know if you saw that Stanford technical report.

NORBERG: Is this the 1972 one on *Computer Power and Human Reasoning*?

BUCHANAN: Yes. McCarthy and Lederberg and I had all separately written reviews of that book, and I put them together into a Stanford Technical Report, thinking it would be generally interesting. And I acknowledged ARPA funding, but learned later that that had created some flak for Les Ernest. ARPA did not want to be associated with those ethical questions. For ARPA to be seen as having funded the reviews of that book, apparently the ARPA office felt that that had been politically unwise.

NORBERG: Well, this would have to have been around 1973 or 1974 when those reviews appeared, and Licklider was back running the IPTO office at the time.

BUCHANAN: Okay, I met with Licklider and Feigenbaum once or twice at Stanford. I didn't hear this criticism directly.

NORBERG: Sure. Well, I am just trying to fit it into the office at the time. Who would be doing the...?

BUCHANAN: Well, I can tell you exactly when that report was issued.

[INTERRUPTION]

BUCHANAN: 1976.

NORBERG: November, 1976. I see, okay; and caused some difficulty.

BUCHANAN: Minor flak, but it was one of those phrases or topics that was better left unmentioned. So anytime we were being visited by an ARPA program manager, we were aware of sensitive topics.

NORBERG: How did this reflect itself in the development of proposals then? Were the writing of new proposals different in, say, the early 1970s than they had been in the late 1960s? Let me ask the question differently so we can elicit a straightforward answer here. When did you first become involved in the writing of proposals? If at all, maybe you didn't.

BUCHANAN: Oh, I did. Ed and Josh were taking care of ARPA proposals in the 1960s. They would get folded into the general Stanford AI lab proposal that Les Ernest would fabricate with the boilerplate on it. As part of every proposal, we would be asked to describe past work, of course, and future work was at least partially an extension of past work. I don't have a specific recollection. It's inconceivable how Ed and Josh would not have asked Georgia and me to write some paragraphs. It would have been a very normal thing, but I don't recall.

NORBERG: All right, so you really don't know anything about the construction of proposals in the 1970s. The reason I ask that is because in 1974... before 1974, I have seen some e-mail messages from Licklider to the community.

It's not clear to me how broadly that community was defined at the time. I'm not even sure it went beyond the AI people. But in any case, he sent an e-mail message around April 1974 indicating that there were some serious questions being asked back in the home office about the funding, and what should be funded and what not, and how to elaborate more milestones for AI research and so on. This is just about the time that Heilmeier was coming in as head of DARPA. And I have seen other statements of Licklider around the same time criticizing other proposals. Let me give you an example. Al Newell wrote an outline of a proposal in this same period, which he sent by e-mail to DARPA. Back came not by e-mail. but a hardcopy, a 27 page criticism of the outline from Licklider, suggesting all of the things that need to be dropped from the outline, not even discussed in the proposal, which confirms what you were just saying about, "We don't want to discuss medicine," and so on, and in addition, suggesting what ought to be in the proposal and how it ought to be framed. At the same time, alerting the community that funding would no longer be done on a multi-year basis, but would probably be from year to year from that point on, or at least for the foreseeable future, which turned out to be the case from that time on, as a result of that. Do you remember any of this going on, any discussion in the laboratory at Stanford about the funding situation in the mid-1970s?

BUCHANAN: Yes, with respect to those topics that were generally more interesting, and therefore we ought to try to ? into. I think, the phrase was, fusion research, multi-sensor integration. Those were phrases that ARPA knew they were trying to sell. Anything we could say about that would provide a link. Partly for these reasons Ed started the HASP project, which he has talked to you about, I am sure. It was mostly sensor integration, and it was a military application. It was interpretation of sonar signals from Navy sensors in the Pacific Ocean to try to interpret the movements of the Russian Navy. Well, I think there are many reasons why Ed was interested in this, but one of the most important was that without that close link to a military problem and a military application our ARPA funding was in real jeopardy. Ed and Josh also broadened funding for the heuristic programming project, as the DENDRAL project evolved into. And we got NIH funding about 1971.

NORBERG: So that's earlier than the period I am talking about.

BUCHANAN: Well, I think they saw the handwriting on the wall. They saw that biomedical applications were harder and harder to sell; that there was just a lot of chemistry. We needed funding to develop for a lot of clinical medicine. And...

TAPE 2/SIDE 1

BUCHANAN: ... mentioned, and one of those in particular was in our work on tutoring systems within MYCIN. We had ARPA funding, and we justified it as multiple uses of a knowledge base for clinical diagnosis or tutoring medical students. But we never mentioned those terms. It was multiple uses. That would have been in the 1975, '76 proposal, because the work actually culminated in a thesis in 1979.

NORBERG: Okay, we...

[INTERRUPTION]

BUCHANAN: Well, let's see. Their proposal to NIH to work on biomedical computing and problem solving was in fact what funded the NIH resource, the SUMEX computing resource. I believe the first year of that funding was 1971; that's verifiable. But the proposal writing would have started in the late 1960s, or the thinking and planning in order for a 1971 start date. And also then, Ed and Josh encouraged me to apply for an NIH career development award. And that was awarded in 1971, and was five years of guaranteed funding for my salary. So it partly removed me from the necessity to worry too much about ARPA funding for myself.

NORBERG: I see. I didn't realize it was a five-year award. I knew you had it, but I didn't realize it was five years.

BUCHANAN: It was another one of those marvelous gifts, and I think the existence of those career developments awards is another story, but they provided a lot of just this kind of enlightened funding.

24

NORBERG: The suggestion here that I am taking from your remarks then is that Lederberg and Feigenbaum saw that they needed to have other support for this research as early as, let's say, 1968 as an arbitrary date here. Now, that is about the time that Roberts is taking over as the head of the IPTO office. His predecessor was still there, I think, in 1968; that would be Taylor. Now, Roberts has expressed some criteria that he used to certify funding - to... (what's the word I want here?)... in his decisions about funding... expressed some criteria about how he did that. And one of the things that he was not interested in, as he told it to me, were things that were just small improvements over what was going on already. But he didn't get to the other end of that (I didn't think to ask it either). He didn't get to the other end of that of trying to tell me about which things he was not willing to fund. I just didn't even think to ask that question; I don't know why now, but I didn't. I wonder if it's possible that these people in the office expressed some lack of interest in what was going on in the DENDRAL project.

BUCHANAN: Not that I am aware of, but it's entirely possible. Ed is politically astute, and he would have picked up instantly on that one.

NORBERG: Okay, I will send him an e-mail message and see what he responds to that... to the question without saying where I heard this. I'll ask him what gave rise to the approach to NIH. Now, one last question here has to do with NIH funding generally. Were you aware of any programs in NIH for stimulating computing? Never mind research in computer science, but stimulating computing as an aid to the biomedical sciences.

BUCHANAN: The Division of Research Resources under Bill Raub [?], who just finished his term as acting director of the NIH, but Raub was the program manager for DRR, and he became convinced that a computing resource... a resource for symbolic computation in the biomedical sciences was needed. Part of every resource - mass spectrometry, computing, whatever - is core research, and it took some argument, mostly by Lederberg and Feigenbaum that computing was worthy of resource status, because DRR had prior to that funded only large instrumentation grants. But core research was part of the criteria, and our core research on techniques for symbolic

computation then became that part of the proposal. It could never have happened if Raub had not been an enlightened administrator.

NORBERG: Yes, but they were funding large instrumentation facilities before this?

BUCHANAN: Yes. For instance, a mass spectrometry facility at Cornell; things like that.

NORBERG: I see, where the emphasis was on the use of the material, not on computing.

BUCHANAN: Oh, that's right. And so, setting up a computer as a resource was a new idea. It required some selling, but once Raub believed it, he was firmly behind the required selling to peer review.

NORBERG: All right, let's pick up from there tomorrow morning.

BUCHANAN: All right.

DATE: 12 June 1991

NORBERG: I want to go back to where we ended up when you mentioned that you had received a career development grant from NIH and this was a five year award, and therefore you and no one else had to worry about your salary for those years. What else did it bring with it? Were there research funds attached to the award?

BUCHANAN: No, as I recall it was just salary. And it didn't change the nature of what I was doing, but it was some affirmation that the meta-DENDRAL work that we were just starting was worth doing. That was the program we were writing to learn the rules that DENDRAL used to interpret mass spectra. And so it was peer reviewed, partly on the Stanford environment and partly on the work proposed. And they put some money there for that work and my salary. But I was still involved with the ARPA contract and ARPA-related work. That didn't change.

NORBERG: Was it during this five-year period, or maybe just after, that you were appointed to the faculty at Stanford?

BUCHANAN: Let's see. The move from research associate to research computer scientist to adjunct professor wasn't entirely clear to me. I guess the adjunct professor was the faculty appointment. That was about 1975.

NORBERG: So it was just about at the end of that period.

BUCHANAN: Yes.

NORBERG: Okay, now, when did MYCIN come along?

BUCHANAN: MYSIN started about 1972. Ted Shortliff was a medical student, who was taking some time out for a Ph.D. in an interdisciplinary area that he defined as medical informatics. And we had been running a journal club on computers and medicine, and Ted was a participant there, and that's I guess how I got to know him. And he was also in a class that I was helping - forced - teach in statistics. Ted was talking with Stan Cohen about projects. I was looking for a different project. We had been working on DENDRAL for over five years. Meta-DENDRAL had started. We thought that clinical medicine was ripe for the technology, and Ted believed it and took the ball and ran with it.

NORBERG: Now, how would that compare... I'm thinking now of the origins of the project. How would the origins of the project compare with the origins of DENDRAL, where Lederberg came up with the algorithm, which then people explored for the next five or six years? Was this the same sort of thing?

BUCHANAN: No, I think it's quite the opposite. DENDRAL was conceived by a very creative person, who had a grand vision. MYCIN was conceived by a student and junior researcher, namely me, and we were working from the bottom up to see just what was possible. The idea started to be a program that would review actions already taken in a medical center, to pull out those that were questionable and ought to be reviewed by a peer review committee. And we decided that critiquing was not the right psychological model for a computer program with physicians' actions. But prospective systems to keep people out of trouble might be more acceptable. I'm not sure whose idea it was - Ted's or Stan's, but that was a transformation that only happened after we started the work; not part of the grand vision.

NORBERG: What I am trying to do is to compare notions about MYCIN in the early 1970s with notions about DENDRAL in the middle 1960s to see whether or not the field of AI had changed its objectives and that this was a response to that. That is, people were no longer looking for grand principles that would govern the development of what I would call knowledge-based systems. They were now being much more applications-oriented, in a sense.

BUCHANAN: I don't think at that time that a major shift had started, although there were several groups working on computer-aided medical diagnosis. There was a group here in Pittsburgh. There was a group at Rutgers. I am not sure where else at the moment. But these people were interested in medical applications, but still it was a very small community. And the major thread of AI work was still on the grand principles. Ed refers to a paper that Papert and Goldstein put together in which they talk about a paradigm shift. Did he mention that to you?

NORBERG: No.

BUCHANAN: I have a reference for it, but they were the ones from MIT who were looking at what was going on at Stanford and a few other places and noticing that the major emphasis had shifted sometime in the mid-1970s from grand principles of human thought to knowledge-based systems. So that paper may be worth mentioning. It was the first... We thought there was something different. We were not giving it the label yet that it has now come to have.

NORBERG: Could you characterize for me then what you see as the major events in AI from 1965 to 1985?

BUCHANAN: Oh, boy...

NORBERG: That's a tough assignment, I realize, but I'm trying to see what your reaction is to the question actually, rather than to the specifics.

BUCHANAN: I guess I would have to think about categories of events, because hardware development was extremely important to us. And the development of LISP personal workstations made a very big difference in the way we would think about program construction and use. I guess from my perspective, with acceptance of the knowledge-based paradigm as a major organizing principle for AI programs is the major development. The thing that at least, from my point of view, has made the biggest difference. And that takes a number of forms. It is not a commitment to a rule-based system, but it is a commitment to the development of a separate, explicit declarative knowledge base that can be manipulated from the outside manually with editors or with other programs such as a learning program, or a knowledge-based editing assistant. And that explicit knowledge base then gets read and interpreted by a general purpose inference engine. And that inference engine is not in itself particularly powerful. It's a journeyman, sylogistic inference system.

NORBERG: Now, when did this occur? What are the aspects of that development?

BUCHANAN: Well, I have found threads of that in the DENDRAL work, and some of it is relabeling what we were doing, but it was there. In DENDRAL we called it table-driven programming, because we realized that encoding chemical knowledge in LISP procedures made it very cumbersome to change. And we were changing things daily, as the ? exchange, as we needed to broaden the scope of the system, and so on. Something as simple as the masses of chemical elements, just putting them in the table is an obvious software engineering principle. But we

29

weren't doing it initially. Those masses were hard-coded into procedures, and each time that was referenced... we had to notice each time if we were going to change them at all. For instance, add another element ?   ?   . So there are threads there, but... And also, in the introduction to the MYCIN book Ted and I talk about the work that we had done in DENDRAL, that I had done on production system encoding of the mass spectrometry rules used in the predictor. And initially that was just a set of procedures. After Newell gave a talk at Stanford about 1968, Feigenbaum and I talked a long time about using production rules for this mass spectrometry knowledge. He was using them for purposes of psychological modeling, and it appeared to be a representation that we could use. And so Ed gave me the task of rewriting a system as a production system with the mass spectrometry knowledge held in those explicit declarative conditional sentences, which I did. And it made my life very much easier then in what we call now knowledge engineering. The on-going dialogue almost daily between me and a post-doc chemist, Allen Duffield working with...

NORBERG: Allen who?

BUCHANAN: Allen Duffield.

NORBERG: Duffield.

BUCHANAN: A super person - very patient with me. Allen would articulate how it was he was interpreting some data and what the program ought to do in order to get it right. And then I would try to put what I was hearing into the program. So there are the threads of knowledge and engineering, too. But we didn't put all of that together really until the MYCIN project. And after some initial exploration with other representations and variations on the problem formation, we decided that production rules would be an appropriate vehicle for all chemical knowledge that was needed for infectious disease diagnosis, and we ought to try to be systematic in separating the medical knowledge from the logical rules of inference. We tried to be right from the start, which was very different from what we had done in DENDRAL. We had been making some claims that we could substitute something else for that medical

knowledge, and the inference procedures would be able to run over them for a diagnostic task, just as they did in medicine. And it was Bill Vinaly [?], a student of mine who had been working on MYCIN, who actually made good on that claim, first in a class project and then he cleaned everything up and made a full package out of e-MYCIN.

NORBERG: Now, what's the date we are talking about here, roughly?

BUCHANAN: Vinaly [?] finished that and published his dissertation in 1979, and he probably worked for four or five years on that.

NORBERG: Who else was doing this kind of work? You made a reference to Newell, but Newell was working on a different set of problems at the time.

BUCHANAN: Newell was still working on small puzzles that you posed to human subjects in order to get them to think out loud, and he was looking at those protocols on cryptarithmetic problems, for instance, but was using production rules as a way of encoding what he thought human subjects were doing when they were manipulating the puzzles. Probably the piece of work that is closest and intent and perhaps even architecture was Joel Moses' work on symbolic integration in the MACSYMA package at MIT. And the MIT folks claim that MACSYMA is the first expert system. MACSYMA and DENDRAL developed more or less simultaneously, and at my level there was no contact. I didn't even know about MACSYMA. Ed certainly did. He was much better read than I was, for one thing, but also, it would have been discussed at the ARPA principle investigators meeting. But I don't think we drew a parallel here. We didn't think of that work and ours as similar until much later. The other work in medical diagnosis bears some similarity in that it required a good deal of knowledge about medicine. But people at Pittsburgh and Rutgers and elsewhere were not systematically trying to separate the medical knowledge from inferential procedures.

NORBERG: Oh, yes, in looking at the Amarel work I think I can understand that comment. Is the suggestion that the outlook of the different groups working in AI, and I am thinking principally of the three big groups - MIT, CMU and

Stanford, is it fair to say that there's a similarity between what went on at CMU and what you, Feigenbaum, and Lederberg and Shortliffe were doing? But it's different than what was going on at MIT, other than Moses obviously, at MIT and other groups at Stanford, who were working on the machine-oriented side of AI. What I am trying to do is separate AI into essentially two camps.

BUCHANAN: It's not geographical though, because McCarthy's group and the SRI group under Nilsson were doing more MIT-like work, predicate calculus and its extensions for theorem proving. And there was a good deal of robotics at those three places. Less robotics at CMU in those days. And, you know, in the DENDRAL project, we had no interest at all in robotics; we saw it as a diversion. I am not sure it is just two things going on then, because the Carnegie people were interested in production systems for psychological modeling. And the nature of their production systems was very different from the kinds of things we were trying to develop, which could be understood and edited by a chemist. So from their point of view we were taking something that was very simple and clean syntactically and muddying it up with a lot of semantics of chemistry. From our point of view, what they were dealing with were structures that operated at a micro level with movement of individual symbols in and out of long-term memory, or short-term memory.

NORBERG: Who's *they*?

BUCHANAN: They: Newell was certainly the intellectual leader of that group; John McDermott was one of the major players in that group; Don Waterman, who had done a dissertation with Ed, had gone to CMU and was working with Newell on production systems. But at a meeting in roughly 1980 McDermott presented a paper on production rules, production systems, and 17 different ways of thinking about conflict resolution, which is the choice of which rule to invoke next. And Feigenbaum's comment, at the end of McDermott's talk was, it sounded like McDermott was coming from Mars. That what they were doing had nothing to do with the manipulation of knowledge. And McDermott later said that that had a profound effect on him - that he thought he was doing mainstream AI. And in many respects he was, but whatever Newell does is almost necessarily mainstream AI.

McDermott, I think, saw the difference between the micro level needed for psychological simulation and the macro level needed for knowledge representation. And he then became a convert. You know, the origin of the name R1, don't you?

NORBERG: No.

BUCHANAN: Don't you know...?

NORBERG: Yes, I know what R1 is.

BUCHANAN: Well, McDermott named it R1 because he said, "Yesterday I couldn't even spell knowledge engineer; now I...

NORBERG: ... R1. [laugh] I see. Your comment about Newell raises an interesting question. I assume it was said factiously that whatever Newell is doing is by definition mainstream AI. I am not sure there is a mainstream AI.

BUCHANAN: Good point.

NORBERG: Until maybe the 1980s. I am willing to concede that today there is a mainstream and there are some black and white areas.

BUCHANAN: No, I think you may be as right now as in the 1980s.

NORBERG: You see, this was involved in my question of the two camps, as well, looking to see whether there are, say, two main threads. And the implication of your comment about Newell is that there was one main thread and the rest of the people are out-liers.

BUCHANAN: No, I didn't mean it that way. I meant that Newell whatever was doing it was irresponsible for anyone to ignore. It was going to be important somehow. And Newell has been on the track of psychological modeling in human information processing using AI as a way of understanding that. He has had many diversions, but that has been a consistent line of thought for him. But he is willing... I think I have heard him define AI as the science of thinking in humans and machines - the science of intelligence. And there certainly needs to be a lot of interplay between these two threads. Nobody has ever denied that it's necessary to know about both things. The question is whether machine intelligence takes the same form. I don't think Newell is committed to saying that, but his research program has been to understand human thought. Now, I am still bothered with the two camps, because we didn't... We were dividing the world up more according to what people worked on. There were people doing robotics; because there were people doing speech understanding; people doing game playing, theorem proving.

NORBERG: Well, compare that for me then with, say, medical people, since that's an area that you're highly competent in. The people doing anatomy, physiology, neuro... whatever it is now...

BUCHANAN: Surgery.

NORBERG: Surgery, and so on. Is there an analogy to be made between the various areas you just mentioned for AI... yes, robotics, knowledge-based engineering systems and so on, to the specialties in medicine? Or is that too extreme of an attempt?

BUCHANAN: No, I think that's fair. Years ago, McCarthy talked about the two main themes in AI research being the representation of knowledge and the use of knowledge - inference. And if you take that as the core "theory," then the specialties build on that core, just as in medicine, in ways that partly intersect with each other, but partly have their own identity. So the robotics people have to be very much more concerned with...

34

NORBERG: To make sure that we got all of that comment, you were saying that the robotics people had to be much more concerned with mechanical engineering.

BUCHANAN: But unless they're concerned with representation and use of knowledge involved... knowledge of whatever assembly line they're working on they don't have a complete system. Their priorities are different from ours. But, just as we needed a good deal of knowledge about organic molecules and mass spectrometry, we still had to be concerned with representation of that knowledge and how it would be used. General level strategies for, say diagnosis of problems... We began to see that MYCIN had one strategy; the internist system here at Pitt had a somewhat different strategy. But we could talk about a strategy for diagnosis in medicine, electronics, mechanical equipment. And the strategy was still roughly constant.

NORBERG: And was the strategy the same for people who were building robots, for people who were developing, say, LISP machines?

BUCHANAN: Let's see now. Again, they were not talking so much about diagnosis of problems then, but when they did start talking about it, yes, a good deal of commonality at that strategy level. Speech understanding also has its own idiosyncracies, but the speech signals themselves are data just like mass spectrometry data. So there's a data interpretation strategy that various people proposed for speech understanding. And one of the most influential of those was the blackboard architecture that was developed by Reddy and others at CMU for speech understanding, but it has a good deal of use in knowledge-based systems for data interpretation elsewhere. I am trying to give you the idea that, yes, there were specialties, but they were feeding back into the core as well as building on it.

NORBERG: Let me take you in a slightly different direction and then I will sort of circle around I hope back into these comments to make sure I understand this fully. What would you say has been the relationship between AI and other

areas of computer science? And again, let me tell you why I ask the question. Some people, like Marvin Minsky, would like to claim that AI is, in his words, "the avant-garde area of computer science." And therefore, most... he makes this conclusion... and therefore, most of the important developments in AI... in computer science, such as time-sharing, are driven by AI work. Now, I am not asking you to confirm or deny that. What I am asking you though is, what do you perceive has been the relationship between AI and other areas of computer science, if any?

BUCHANAN: Let me answer that in two ways. Sociologically, the people doing AI have been quite largely separate from the people doing other parts of computer science, and have been separated from computer science at many institutions in various artificial ways. I am not sure of the reasons for that. AI is, in many persons' eyes, less rigorous. Now, it may be less mathematical. There are fewer theorems to prove, but I don't think that implies less rigor. Apart from sociology, in terms of the content I oversimplify things often in an introductory lecture to lay people about what AI is. I use a two-by-two matrix where on one axis we're thinking about the type of information that is being processed, and on the other axis the type of processing. And the type of information may be numeric or symbolic. The type of processing may be algorithmic or heuristic. The 1940s emphasis was on numeric algorithms for computing trajectories and such. IBM moved that into business data processing with payrolls and employee records, but it was still algorithmic processing. The box with heuristic processing on numeric data - I think the best examples would be large scale numerical simulations that are just unwieldy if you try to deal with the simulation algorithmically. You need to introduce some heuristics, some simplifying assumptions to get things to run in a timely fashion. And then AI is that fourth quadrant, whose emphasis is symbolic information and heuristic inference. But of course those are not sharp lines. But then it means that AI has a piece of computer science that it emphasizes, that the other disciplines are quite willing to acknowledge they do not emphasize. I think the major difference is in the algorithms versus heuristics.

NORBERG: What do you think has been the contributions of AI research to other areas of computer science?

BUCHANAN: Well, I have heard Marvin saying things like that, and I think to some extent he's right, because...

36

NORBERG: You think he's right about what?

BUCHANAN: He's right that AI work drives some development in other parts of computer science. I think we would

be hard-put to say all such innovations and developments were driven that way. But as we are trying to tackle larger

and harder problems we seem to run out of resources faster than other people. So something like time-sharing was a

way to get more resources here. I don't know if you would make that case about personal work stations; some people

might. The ALTO machine was partly developed for symbolic computation. It was strong in a whole line of Xerox D

machines, certainly. [long pause] I think it's extreme if we try to claim that...

NORBERG: Okay, let's not claim that then. Let's drop down to a more realistic position, and try to cite things that

have been influential. Two things come to mind when I think about AI and the sociological history of AI, and that is

the seeming continual need of the AI people to justify themselves to other computer scientists. And the way that

seems to show up is that other computer scientists don't think the AI people do anything that's really useful. You

said rigorous, but there are other people who would say "really useful" in that context. So the question is why did AI

people in many instances feel they had to continually justify themselves to other areas of computer science? That's

the question in my mind. And the second part of the question has to do with what I see as a very disparate set of

problems being investigated that may or may not be a response to this criticism from other areas of computer science.

 Now, you described ten minutes ago that while many of these research problems that people were working on - you,

McDermott, Newell, Papert, and so on - many of the research problems, while they looked different because of the

subject matter and the approach are indeed feeding back into the core.

BUCHANAN: Yes.

NORBERG: I don't think that's obvious to everybody, and when I started this task it wasn't obvious to me. I am beginning to see it more clearly as time passes and as I do more reading and try to fit things together. But it's not clear that everybody else does.

BUCHANAN: Pushing Minsky's hypothesis that core AI feeds back into computer science generally. Let's see, sociologically, I'm not sure what you say is entirely true, that there is a need for a continual justification. I think that's largely stopped, if it was there at all. Among the people who were most confident about their computer science skills - McCarthy, Minsky... I don't know, Newell, Reddy - I think it's not an argument and you don't find them trying to justify what they do. They just do it. I think that's true. And I suspect that the kinds of justification comes from people who, perhaps like myself, have come into AI from a different discipline and, lacking training in all of computer science, were less comfortable than the Newells and the Reddys about the algorithmic numerical sides. So that may account for some of it. AI certainly has attracted people from a lot of different disciplines. I don't know the extent to which the work out of which AI came could be distinguished from core computer science at all - McCullough, Pitts, Rosenblat, Wiener - those people who were very interested in human thought but didn't think of themselves as working on anything but computing machines. The fact that machines operate on arbitrary symbols is exploited in AI, but it was certainly well-known as the von Neumann principle. So I don't think von Neumann would have divided the world up into that two-by-two matrix, for instance. That's overly simple.

NORBERG: Well, there are criticisms of AI and they are continuous in the history of AI, so how are we going to deal with that? Take the simple example of NSF taking a very long time before they supported AI work. It's really only in the 1980s that they began supporting AI. Before that, according to Minsky, it was very difficult to get the AI projects through a peer review process. And when you did get it through the amount of money was so small that it wasn't worth while in the face of ARPA support. Now, do you think that that is a correct analysis of the situation?

BUCHANAN: Yes, I think there's a good deal of inertia always in the system. And I think when there's a somewhat flamboyant group working on rather ill-defined problems with a good deal of enthusiasm and having fun colleagues

are going to be suspicious, and I think that's largely what happened. AI is fun, and what these other people do can't

possibly be as much fun. So I think there's a suspicion and jealousy. And so, if they are controlling the purse

strings, as peer groups do, they will be reluctant to let the new kids in.

NORBERG: Okay, I am going to ask another question with two parts to it also, and we'll go back to the personal

again. While you were at Stanford, how much contact did you have with the rest of the AI community? Was it

through people passing through Palo Alto, such as the Newell lecture you heard, or did you attend meetings

elsewhere, or did you visit elsewhere, and if so, where?

BUCHANAN: People came through Palo Alto a lot. The fact that the DENDRAL project was at the AI lab meant that

we saw a lot of people passing through as McCarthy's guests, or there to see Reddy. And Feigenbaum and

Lederberg certainly had their own group of people they wanted guests to meet. But we went to a lot of meetings, and

I think the most influential in the 1960s were the Machine Intelligence Workshops out of Britain - Mickie-Mettzer

workshops. They were international, but I would guess that 2/3 of the participants were from the U.S. There were

maybe 50 people at those workshops, but over the course of about two successive meetings you met a very large

fraction of the people who were working in AI. Now, since I was junior, I guess I largely tagged along with

Feigenbaum and that gave me access to conversations with all of the principals. But I didn't feel I had Minsky's ear.

There was no reason really for him to pay attention to me at that time -maybe not now. But nevertheless, there were

discussions in which everyone participated. We were presenting the DENDRAL work, as I mentioned yesterday, and

some people were listening anyway. But those volumes also then became a principle set of research resources for

everybody. There weren't any textbooks yet. In fact, in 1964, when Ed asked me to come out to Rand there was only

one book in AI, and that was *Computers and Thought*.

NORBERG: Yes, Feldman and Feigenbaum.

BUCHANAN: So, the answer to your question is, both. We had a lot of visitors and we went to several meetings.
Ed went to more than I did.

NORBERG: How about contact with other areas of computer science?

BUCHANAN: Far less for me personally and I think for the project and the whole AI lab. At Stanford the AI lab was
two and a half to three miles behind the campus and many of the students who were working there would attend
classes sporadically, but spend all of the rest of their time at the AI lab. So they were not in contact with the rest of
computer science, nor were we. When I started attending Stanford faculty meetings, there were several faculty I
really didn't know at all. I had never talked to. I mentioned a conversation with George Forsythe when he was
chairman, but he was a numerical analyst. But I had very few conversations with Forsythe about computer science or
AI. I saw him socially. He and Ed were friends. Very little contact.

NORBERG: Let me slightly shift again. I think I'm still on the same topic, but you may not think so. As you were
coming off the development award in the middle 1970s was there any discussion around the laboratory about funding
issues that you remember? Reactions to changes at DARPA, for example.

BUCHANAN: Well, I think those discussions never stopped, Arthur, and I don't recall any specifically pointing to a
problem with getting new funding.

NORBERG: But how about getting the laboratory funded? It's the Heilmeier years that I'm fishing for here, and what
happened in 1973, '74, '75 that caused a shift in the justifications needed at DARPA for funding in computing. It
wasn't just AI, although AI became a particular target for justification by that group -this is not the same justification
I was talking about before, the justifying to ones colleagues, but justification to the Department of Defense that this
work should be supported. There seems to have been a shift in the 1970s, and I am trying to see whether or not you
were aware of that shift, and whether there were discussions in the lab.

BUCHANAN: I was aware of it, partly... Those were the times at which we changed the name of project from the

DENDRAL project to the Heuristic Programming project, partly to send a message that we were not just doing

chemistry. Heuristic programming had some meaning in the computer science community and to ARPA, and also,

realistically, to reflect the fact that there were additional projects. Ed had started the HASP project, as I mentioned,

specifically to address the Heilmeier question. And it was with a good deal of forethought about which area, and just

how to attack it. But it was a very deliberate decision on his part. And then we could bring a little bit of chemistry in,

if we used different words, and a little bit of medicine, using different words, with a major emphasis on signal

understanding of military sonar signals. That was important.

NORBERG: But did you people feel it was important to be in AI? Why couldn't it have been done elsewhere? Was

this really a concern that you people had, or was this just to justify to DARPA in order to get funding for the

laboratory?

BUCHANAN: I don't even understand the first part. I mean, I guess we believed that what we were doing could be

applied to additional problems and it was good to demonstrate it. But I think left to our own devices we would have

demonstrated it with biomedical problems.

NORBERG: That's a fair answer and I think it's probably right. It's probably right. I would like you to react to two

things for me, if you will. Coming from the introduction of Barr and Feigenbaum's handbook (I don't know now

whether it's the first volume, or if... It's Volume II), in which they talk about the DENDRAL system. You can just read

it for yourself. It's the indented paragraph there that I am interested in to see whether this is similar to the story you

told yesterday or slightly different.

BUCHANAN: Yes, I have no quarrel with that.

NORBERG: But, you see, when I was talking to you yesterday I was trying to suggest that the foundations of DENDRAL were rather specific to produce results in a given area of AI, rather than the theory formation issue that you were interested in. And... I am not saying this in the right way. This is not the way I meant it.

BUCHANAN: In nearly completed form...

NORBERG: No, I am not arguing with you about it. What I am trying to say is that in the literature that I have read the emphasis seems to be on this; not on theory production, not on theory formation, not on theorem proving, and so on. Those are things that don't seem to appear in the literature now when people are trying to say that DENDRAL was one of the first expert systems. That early, let me call it, large vision seems to have disappeared from the later literature.

BUCHANAN: In some of the very first papers on DENDRAL, one that Lederberg and Feigenbaum wrote before I even got there, they're talking about mechanizing scientific inference. And again, this kind of blurring of hypothesis formation and theory formation. And both came under the name "induction." And the report that Watson and Feigenbaum put together talked about machine induction. And those are the kinds of results in chemistry, results of reasoning by induction, that they're talking about there. I think it's still all consistent. But when you use a grand phrase like mechanizing scientific theory formation, it's harder for people to take you seriously than when you demonstrate that you can get the answer to a data interpretation problem in this   ?   , although we all along wanted the project to develop along the lines that it finally did with meta-DENDRAL. But we didn't start meta-DENDRAL until about 1971, because the other pieces were not yet in place. We needed a lot of machinery just to manipulate chemical graphs effectively. And we didn't know it at the time, but we couldn't have started meta-DENDRAL without this explicit representation of the production rules that were driving the DENDRAL interpretation, because meta-DENDRAL was writing those in that same syntax, and then turning them over for use. I think that was a key element, and we just saw that there was an awful lot we had to understand about data interpretation before we could think about the theory formation. Lederberg was truly amazing in all these years, and once he formulated a

42

problem precisely enough in his own mind to see his way through to a solution with technology that was either in hand or could be put in place, he moved on to the next problem. He considered it a solved subproblem. So as soon as meta-DENDRAL was laid out in his mind as a large, heuristic search program, then he was off on the next project, leaving us with more than five years of work to do to make good on his vision.

NORBERG: I see.

BUCHANAN: It was amazingly stimulating and frustrating for us, because we were so slow by his standards.

NORBERG: Okay, one last question and that is, in the later years, middle 1970s, do you remember visits to the laboratory by DARPA personnel? Did you ever interact with any of these people?

BUCHANAN: Bob Kahn, several times. I forget when Kahn came.

NORBERG: Well, Kahn was a program manager from 1972 on, so he could have been there several times.

BUCHANAN: Ed and Bob Kahn had established good rapport. I am uncertain of the dates, but we had numerous visits from Kahn, and I am sure Ed had lots more. There was a time in the mid-1970s when Ed was being recruited to be program manager, or whatever...

NORBERG: Director.

BUCHANAN: ... Director, I guess, and would only do it on the condition that ARPA move to the West Coast. [laughter]

NORBERG: Well, the reason I ask that is the second comment that is in this report that one of my students did, that I am interested in your reaction to, because this looks like an extreme position to me.

BUCHANAN: [Reads section.] Okay, my contact with Licklider came through meetings - workshops, conferences. So I don't recall meeting with him directly. I haven't seen this comment from Ed, nor have I heard it, but there's no need to doubt it. Ed has a very good memory.

NORBERG: But these are tight budget times, 1974 and 1975, and to have someone come along and say, "You're really underfunded; now I am going to give you twice as much," seems like an odd comment, when there was rather significant competition for resources.

BUCHANAN: Well, let's see. What would be the date?

NORBERG: It has to be 1974 or 1975, because that's when Licklider was in his second term as Director He was gone by the end of 1975.

BUCHANAN: How early could it have been? No earlier than 1974?

NORBERG: Right.

BUCHANAN: Because we had results coming in earlier than that. What was the end of his first term?

NORBERG: 1964.

BUCHANAN: Oh, that's too early.

NORBERG: Yes.

BUCHANAN: And Ed would certainly remember that it was Licklider.

NORBERG: Yes, because it wouldn't have been Roberts. It's not Roberts' style. He wouldn't tell you how to attack your budget. Okay, I was just looking to see whether you...

BUCHANAN: No, I don't have any recollection of that.

NORBERG: Okay, this is the sort of thing if it were me in my institute I would just say, "This guy just told me to double the funds, people! Get your projects ready. We're going to do some new work."

BUCHANAN: We were expanding partly by diversifying. We had NIH money coming in. I wasn't managing the budget at this time. I started taking over the budgeting duties later.

END OF INTERVIEW