

An Interview with
PATRICK H. WINSTON

OH 196

Conducted by Arthur L. Norberg

on

18 April 1990
2 May 1990

Cambridge, MA

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

Patrick H. Winston Interview
18 April 1990
2 May 1990

Abstract

Winston focuses on his work in computer science and artificial intelligence at the Massachusetts Institute of Technology (MIT) and funding of projects through the Defense Advanced Research Projects Agency (DARPA). Winston discusses: computer science and artificial intelligence research, the work of Marvin Minsky and Seymour Papert, the Laboratory for Computer Science and the Artificial Intelligence Laboratory at MIT, his own work in the AI Lab, the programming language FRL, the changes in DARPA support over time, and the influence of DARPA support on project design.

PATRICK H. WINSTON INTERVIEW

DATE: 18 April 1990

INTERVIEWER: Arthur L. Norberg

LOCATION: Cambridge, MA

NORBERG: I'd like to start back in the late 1960s when you were doing your Ph.D work here at MIT and ask you to tell me a little bit about how you selected the problem that you worked on and what the support was for that project, as you knew it at the time.

WINSTON: I'm not sure what you mean by support.

NORBERG: Well, was it direct project support, or was it just the general support that DARPA was providing to the AI Lab, or did it come from somewhere else?

WINSTON: This is back in an era when DARPA had a considerably different model of university support than it has today. This is a time when the notion of a center of excellence was part of the idea as distinguished from deliberately not part of the idea.

NORBERG: What do you mean by "deliberately not part of the idea"?

WINSTON: Well, in recent years DARPA has not funded entire laboratories as centers of excellence but has, instead, supported specific projects within laboratories. So whereas in the beginning this entire laboratory was almost one hundred percent supported by a single DARPA contract, it is now by contrast supported by many different ARPA contracts and many other sources of funding as well.

NORBERG: We will come back to that later. I just wanted to make sure I understood what your comment meant. So back to the late 1960s.

WINSTON: Back to the late 1960s. I came to the Artificial Intelligence Laboratory because I was interested, ever

since I can remember, in the question of intelligence and what makes people intelligent, and how people might become more intelligent, and how machines might become intelligent. And I moved from one field to another looking for something that would help me along the path of understanding or gaining insight to those questions. So as a graduate student, I studied information theory, and neural biology, and all sorts of fields of that kind before I discovered that there was such a thing as artificial intelligence and that there was a center of artificial intelligence activity here at MIT. So I became involved in the laboratory as a consequence of my interest plus the interest that Marvin Minsky took in a term paper I wrote for this subject. On the basis of that term paper, he invited me to work in the Artificial Intelligence Laboratory for the summer.

NORBERG: Do you remember which summer that was between bachelors or master's or...

WINSTON: I think it must have been between master's. I finished a master's degree in operations research. So at this point I was hoping to get involved in the artificial intelligence laboratory, not quite knowing how to do it. It was a community in the sense that... how to explain it? There was a lot of esprit de corps. It was a bit tight knit. For a basically shy graduate student, it was difficult to figure out how to penetrate it. The people in the Artificial Intelligence Laboratory seemed to know everything there was to know about everything, especially computers. So for anyone with the least inclination to be shy, it wasn't clear how to get involved. The only faculty were Marvin Minsky and Seymour Papert at that time, if McCarthy having already gone to Stanford. Minsky and McCarthy at that time were famous people surrounded not only by... They were eccentric in various ways. It wasn't easy to talk to them; it wasn't easy to get an appointment. If you did, they tended to be distracted. So at that time, I was wondering how I might become involved. So I took Marvin Minsky's subject, wrote a term paper he liked. He invited me to work here at the laboratory during the summer, and the rest went on from there. I was, during the subsequent year, here at the laboratory as a research assistant. Tom Binford, who is now at Stanford, was working on problems of vision. I was interested in semantic nets and wondering how I might do a thesis in the general area of semantic nets, a thought that seems rather naive in retrospect.

NORBERG: Why?

WINSTON: Well, putting it in today's language, I was operating as a person intrigued by mechanism rather than concerned with the task, concerned with the problem and the need to bring a solution to the problem. Like many students in the intervening years, I was obsessed by this new representational idea and looking for a place to apply it, rather than the other way around - facing the problem and bringing the right tools to bear on it. In any event I was speculating on the use of semantic nets in a blocks world context and a very obscure footnote in my thesis proposal said, "Well, gee, if I can describe block structures as semantic nets, then perhaps I could gain something from the differences between two descriptions of different scenes and their semantic net form and that might be a basis for learning something." And maybe I gave an example or something. But of course that footnote turned out in the end to be the key idea in my thesis, and the semantic net stuff was put in its proper place. It was a tool and a mechanism and not the idea. So that is how I got involved, and I proceeded to work on that over the course of the next two years.

NORBERG: Who else was in the lab at the time, besides Minsky and Papert and yourself?

WINSTON: Well, Gerry Sussman was very much a part of it. Although he is younger than I am, he was already deeply involved in the AI lab for a year or two before I got here. Other graduate students of note were Eugene Charniak, and Dave Waltz, Bill Henneman. These were all people who were, I think, here when I arrived. Then there were hackers of course, the hackers being the most scariest group of all, because they were like a submarine community. (Laugh) They were very tight knit and very binary in their outlook on people and so on. And already very famous; they included people like Tom Knight, and Jack Holloway and Richard Greening. Bill Gosper as well.

NORBERG: It seems to me there are about two hundred in the lab now, aren't there?

WINSTON: That's right.

NORBERG: And you just made a list of about eight.

WINSTON: I think at the time I became director of the laboratory there were on the order of forty or fifty people.

NORBERG: What year was that?

WINSTON: That was around 1972.

NORBERG: The director of the laboratory in 1972?

WINSTON: Acting director, yes.

NORBERG: That's the one piece I couldn't find out about you, by the way.

WINSTON: What's that?

NORBERG: When you became director of the laboratory.

WINSTON: Oh, the date is a little obscure because I became acting director and was acting director for an undetermined period of time, and then I became actual director, as Marvin says.

NORBERG: Let me go back to the first two years, though, before you took over the lab, because I'm interested in how the lab ran as well. What was the association between the faculty and the students? If Minsky was so difficult to talk to if you didn't know him, what was he like once you got to know him? What was it like inside the laboratory?

WINSTON: Well, I should refine my notion of how difficult to approach... Minsky has never been difficult to approach. He just seemed that way because anybody sufficiently famous will seem difficult to approach to a shy graduate student. I'm still awed by Minsky. In a sense, nothing's changed. I'm sure you've read Feynman's various

books. His remarks about the monster brains he had a great deal of awe of when he was a graduate student struck a very resonant chord with me when I read them.

You ask, what was the place like. The first thing to say, I guess, is that it wasn't one place. It changed a lot with time, and it changed a lot depending on exactly what your geographic location was. Everyone had a mission, everybody worked hard, everybody stayed up all night. The phenomenon was of course encouraged by the fact that in those days computer time was a precious resource. We might work all night so we could get more of it than we would if we worked all day. Some people stayed up late; they went out to dinner as a consequence of that. There was a sort of 'round-the-clock community. There was a sense of, I think, great excitement because we didn't know how hard anything was. We were all young and ambitious and full of expectation that progress would be extremely rapid. So it was a time of great excitement. It was a time when I think the coupling between what any individual did and the mission of the sponsor was a lot looser than it is today.

NORBERG: Could you be more specific about that? Do you remember in 1970 or 1971 having any interaction with people from DARPA?

WINSTON: As a student?

NORBERG: Or immediately after that.

WINSTON: Sure. As a student my first interaction with DARPA people was at a meeting at Allerton House at the University of Illinois, which was a great occasion for me because I met other graduate students who subsequently became quite famous. For example, Alan Kay was there talking about this lunatic idea he had for a dynabook.

NORBERG: He's still talking about it.

WINSTON: It is pretty much here by now, but it seemed impossibly distant then. At that time I first met Steve

Crocker, who was at that time a graduate student here at MIT. I don't remember exactly when that meeting was, but it must have been fairly early in my graduate career because I didn't really know very many people. Steve Crocker was a senior graduate student at the time, and not well known to me. So I kind of met him for the first time, and shortly thereafter he went to ARPA, I guess. So my memory of that meeting is a little fuzzy. But I think Cordell Green was at ARPA at the time. So he was a DARPA personality that I met while still a graduate student. For the most part, though, DARPA was somewhere behind Minsky and Papert and Russell Moscar (?), who was their administrative officer at the time. It was there. I kind of knew vaguely about it, but it wasn't something that we were involved with on a day-to-day basis. And that was probably a good thing, but as in all things there are pros and cons. We were somewhat sheltered from the realities of funding and whatnot in those days. And that had the defect of making us less aware of what that world was like. On the positive side, we operated in an environment in which we were driven by ideas, not by applications and not by a current sponsor and its demands. So I think that was probably a good thing. I think it was a very creative time and it was a time when we could do very long-range things. We could work on the blocks world and make what, in retrospect, are still lasting achievements I think, without excessive concern about the coupling between our blocks world research and the eventual military applications.

NORBERG: What do you remember about your first interaction, then, with the DARPA people? Not at that meeting, but after you became acting director.

WINSTON: Well, this was a long time ago. I think probably my most vivid memory is, from time to time, ARPA will come up in force with the director of the office and several lieutenants and then we would put on a half-day, all-day show. It's never routine, but of course the first time it is especially scary.

NORBERG: Yes, I've been through it. (Laugh)

WINSTON: I think I remember Steve Crocker coming up in advance, as some kind of advance man on my first presentation, helping us to prepare ourselves for the real event. And I think Dave Russell must have been... Is that right? Was Dave Russell office director at that point?

NORBERG: He might have been.

WINSTON: Early 1970s?

NORBERG: Yes, he might have been.

WINSTON: And of course we had the event, and I thought we did a pretty good job. I always end up thinking we do a pretty good job, because we have some pretty good people. (Laugh)

NORBERG: What were the consequences of the meeting, do you remember that?

WINSTON: No. I think the consequences of these meetings are always that we passed. That has always been my sense. I don't think it is the nature of all the people who have ever worked at ARPA to heap praise upon what they've seen. So I've always taken the view that at these presentations there is considerable downside exposure and not much upside opportunity in the sense that, if you drew a curve of performance versus future consequences, I think that you could muff it and that could have dire consequences for you. So you try to be sure you are on the right side of that step function.

NORBERG: When you took over the lab as acting director now in 1972, do you remember why that was the case? Why were you chosen? What were the reasons that were given for either Minsky or Papert deciding not to do it anymore? Was any of that public?

WINSTON: What do you mean by public?

NORBERG: Well, I asked Minsky this question, why he gave up the directorship of the laboratory, and he gave me an answer. And I'm wondering whether that is a public answer or not, and whether that was well known, and why

you? After all you only had a degree for a couple of years at that point. And while I accept that most of the people in the field were very young, other than Minsky and Papert and McCarthy and a few more like that, could this be seen as an unusual event?

WINSTON: Well, it was certainly seen by a lot of people at MIT as unusual. The provost couldn't believe it. (Laugh.) But we had more or less presented the provost with a fait accompli, so eventually he got used to the idea. It wasn't out of character for Minsky and Papert for two reasons. One is that this place has always been very egalitarian. There has never been any big emphasis on... There has never been much visibility of "These are the professors, these are the students, and these are the staff." There is no stratification like that. In those days we were all so young, you couldn't tell the difference without a program. Minsky and Papert, I think, encouraged that. They were interested in ideas, they didn't give a damn about degrees or anything else. So one reason it wasn't strange was because of this egalitarian outlook. The other reason it wasn't strange was because it had a history of doing that. Gerry Sussman was in charge of the famous summer vision project of 1967, or so, as an undergraduate. I don't think Marvin ever felt that he had any great interest in running projects, and he was always into the delegating the day-to-day responsibility for such projects to somebody else. And I think he just did it to people to whom he for some reason or other trusted most. Russell Moscar (?) had a great deal of... He was very much involved in running the laboratory when Minsky and Papert were directing the laboratory. For some reason or other, I suppose I was in the right place at the right time. There weren't a lot of other people; there weren't any other obvious choices, so they...

NORBERG: Why not?

WINSTON: Well, I was almost the senior... Let's see, I guess I was the first faculty member after Minsky and Papert.

NORBERG: What were you expected to do as director?

WINSTON: What was I expected to do?

NORBERG: Was there a job description?

WINSTON: No, I made it up myself, which is, I think, what people are supposed to do when they get a new job. First thing I did when I became director is I went around and talked to a dozen laboratory directors and deans and other high officials at MIT and asked them how to make a great laboratory. I was astonished to see that, as far as I could tell, only one of them had ever thought about the question, and that was Jay Forrester. I had trouble getting it out of him though, because when I sat down to talk to him he presented me with tea and cookies, and after I explained what I was there for, he went to great pains to explain why artificial intelligence was a bad idea. Then I probed a little bit and said something like, "Well I suppose in any case it wasn't hard to have a great laboratory because everybody must be highly motivated by the Whirlwind computer and the thought of putting together the first one." And Jay looked at me like I was the king of the fools and said (laugh), "The laboratory wasn't about the computer at all; it was about the protection of the United States against air attack from the Soviet Union." That remark always stuck with me. The lesson I learned was the way to have a great place is to have a passionate mission. I don't think we've ever had quite the same kind of mission as that, but we've all had the mission of understanding intelligence and changing the world as a consequence. So it's a pretty good substitute for protecting the United States against air attack.

NORBERG: Had the mission been articulated before you took over?

WINSTON: I don't know... it's not on a coat of arms or anything. I think it is a mission that everybody felt. It was never a mission that people talk about a lot. It sounds almost religious when we talk that way. But I think it's there.

NORBERG: But there must be a sort of a socialization process that goes on inside the laboratory to convey this sense of mission to the new people who come in.

WINSTON: Yes, I think so.

NORBERG: And in some places that is usually done with some sort of a document or an orientation, or a working

with a particular person and they...

WINSTON: My view is that it is always determined by the person on the top. I've been amazed at how much vast organizations, Fortune big 100 companies, how companies of that size can still have their own characteristic ambience as determined by what everybody perceives their leader to be like. The leader may not actually be like that, but people shape themselves in the mold or select themselves for that company as a consequence of what they perceive to be the characteristics of the leader. In the case of the laboratory, Minsky was the leader, and I think he has the sense of wanting to make a contribution as fundamental as Darwin and Freud. We all see that and convey it in turn to our students. We really want to make a fundamental...

NORBERG: Let me go back to 1970; we're getting closer to the answer that I want to hear, or at least the kind of answer I'd like to hear. Let's go back to a question I asked you earlier, which we only investigated partially, and that is, what was the interaction inside the laboratory then? For example, there is a large space out there, outside your office which looks like many people could congregate there at any given time in the day, maybe at coffee hour in the morning or something.

WINSTON: That was one of Marvin and Seymour's last great ideas, while they were still directors.

NORBERG: Did anything like that happen then during those years? And how did people interact with Minsky and Papert, as graduate students, as young faculty, to talk about these ideas, this mission that you are talking about here?

WINSTON: It is a little hard to answer that question for two reasons. One is it was a long time ago, and one is that I have to separate my own particular circumstances from those of others in the laboratory. When I was a graduate student, there was what seemed at that time to be a large staff and there were graduate students. Those groups definitely formed two different nuclei - the staff and the graduate students. They were for the most part on separate floors, the staff were working with the time-sharing system on the ninth floor to a large extent. Graduate students for

the most part had offices on the eighth floor. I tend to be a sort of shy type myself, so as a graduate student, I didn't have very much interaction with Minsky. A number of staff people were much more comfortable with blazing into Marvin's office and striking up a conversation. So they had a different kind of interaction than I did. I recall seeing Minsky only three times in the course of my Ph.D. dissertation, or something like that.

NORBERG: I see. That's revealing.

WINSTON: Once to give him a proposal, once to see if I was doing anything useful, and once to hand him my thesis. (Laugh) So Minsky has never been the sort of person who would carefully set up a weekly appointment with students.

NORBERG: Who did you interact with then, just the other graduate students on this problem you were working on?

WINSTON: Oh, yes. Graduate students tended to interact a bit. Some of them more than others. I tended to do most of that work on my thesis by myself. Other graduate students influenced each other quite a lot. It's a little further along, I guess the 1972 time frame, but Eugene Freuder worked quite a lot with Dave Waltz on his thesis and made a number of valuable contributions to Dave Waltz's thesis. Bertholt Horn and I talked quite a lot - not so much on each other's theses, but just about what was interesting.

TAPE 1/SIDE 2

NORBERG: When you became acting director, did you see any need to change that style of operation in the lab?

WINSTON: Before we get into that, maybe we ought to go back and talk a little bit more about thesis interaction. Now, I didn't see Marvin very often, but when I did see him, it was incredibly valuable. He always seemed to be very uncanny, and other graduate students that I talked with at that time told me they had similar experiences with him. They would begin to explain their idea and Marvin would misunderstand the idea to be a better idea, which wasn't the

idea that the student actually had. So the student ended up going away much enriched.

NORBERG: Do you think that might have been deliberate?

WINSTON: I don't know, perhaps. I don't think so, actually. I think Marvin is just so damn smart that as he is guessing what your idea is, his guess is often better than what the idea actually is. When I submitted the thesis, it was interesting. Marvin and Seymour were both members of my committee with Marvin being the supervisor, and they had very different styles in reading the document. Seymour made very global comments, like "You really ought to have another chapter here" or something like that, which distressed me enormously because I wanted desperately to be done. (Laugh) Marvin on the other hand would say, "Well, gee why don't you say it this way" and there would be ten lines of stuff. And I would look at it and say, "In fact, why don't I say it that way?" So there are probably lines in my thesis that are near verbatim transcripts of Marvin's suggested rhetoric. So they were both very valuable in critiquing the thesis, although with very different styles.

Coming to your question of what did I do when I became acting director... well, I made a lot of mistakes, of course. But also I did one or two sensible things. I think what I became good at was finding the intersection between DARPA's interests and local interests, where local interests are driven mostly by what is scientifically exciting. To me I've always felt that the work done in a laboratory should in so far as possible lie at the intersection of three circles, one circle being what the sponsors are interested in, one circle being what is scientifically exciting, and one circle being what has the potential for changing the world. I think I exhibited a talent for finding those intersections and looking at the work that we were doing in the laboratory and finding the connections that were there waiting to be found between what people were doing and what the sponsors were interested in.

NORBERG: Can you give me some specific examples, and try to make them time specific as well?

WINSTON: That's a little harder. (Pause) I tend to remember some of the things I failed to get ARPA interested in better than I remember the things I managed to get ARPA interested in (laugh).

NORBERG: Well, maybe one of those would help you to remember some of the things you did.

WINSTON: Yes. Well, I'm not sure about the timing on this, but it was probably around 1975 to 1978. I was trying to interest ARPA, Dave Russell in particular, in the work that Gerry Sussman was doing on circuit understanding. And a very peculiar thing happened. So much of the work that was done in the early 1970s was done on toy problems in toy worlds. I thought that "Gee, if we brought this a little bit closer to the real world, things would be a lot easier."

NORBERG: A lot easier in what way?

WINSTON: Well, I thought DARPA would be more interested if it wasn't the blocks world. So when Gerry Sussman started to work on analog circuits and understanding them, writing programs that could describe their behavior and whatnot in human terms, I thought that was a great thing. But then Dave Russell looked at it and said "Oh, that's not important because when electronic stuff breaks you just replace it with another board." So I was quite shocked by that because I discovered for the first time that if you take the work into a domain where people know something, it can occasionally be a lot worse than if you leave it in a domain where they don't. (Laugh) You see the problem was that while it was in the blocks world, the only way it could be appreciated was for its scientific content, whereas by doing the same work in a world of electronic circuits, then it raised the question whether the applications direction makes sense or not. So it complicates the picture. In that instance, some very good scientific work was not funded by DARPA as a consequence of the DARPA person in charge not getting beyond the immediate application possibilities to the question of scientific underpinning. It turned out, of course, a few years later Sussman turned his attention to digital circuits which was a matter DARPA was vastly interested in. There were two new people, and they had quite a different reaction. But it was essentially the same theme.

NORBERG: Why did you feel you needed to have some sort of an application program of interest to DARPA? If they were going along and funding this laboratory adequately for all those years, why not just let the scientific interest continue to drive the program and let them continue to support that?

WINSTON: Well, during the 1970s there was what we all lovingly refer to as the Heilmeier era. And there was tremendous pressure to produce stuff that looked like it had a short applications horizon. One never knows if the apocryphal stories are true or not, but I think I heard an apocryphal story in those days about George looking at some surface-scratching piece of work and saying, "Well, let's wait six months, give it to users, and see how they like it." So I definitely remember that those Heilmeier days of the mid-1970s were days in which you had to find a way of explaining the work in applications terms. It wasn't enough to study reasoning; you had to talk about how it might be applied to ship maintenance or something.

NORBERG: That seems stretching it quite far to me.

WINSTON: What?

NORBERG: To move from the laboratory, scientific experimental situation to maintenance of ships - in the middle 1970s - maybe it's not so extreme today.

WINSTON: Well, we were talking about some heuristic scheduling of things and whatnot. I'm not sure about the timing on this, but Ira Goldstein who was working on a lot of stuff during, I guess, the middle to late 1970s, perhaps more to the late 1970s, I suppose, but still the Heilmeier era. When did George leave then? I'm trying to get these dates adjusted.

NORBERG: I want to say 1978, but it could have been 1980.

WINSTON: Sort of around that time frame Ira Goldstein was doing some work on office scheduling and things, and we mapped that into ship maintenance scenarios, with no intent to mislead. Everybody knew what we were doing in one domain and describing it in both domains. That was interesting work because it is probably a good example of stuff that started out in a sort of artificial world but eventually emerged in important ways in other forms. If you'd like

I could digress into that.

NORBERG: Go ahead, as long as it's on your mind now, I'd like to explore that.

WINSTON: Okay. Goldstein, and you'd have to check the AI memo dates to actually nail down the time, he erected a very broad program that included some natural language and some scheduling and some work on frames. This was just after Marvin's "Frames" paper. He implemented a system, which I think he called FRL, Frame Representation Language. And that eventually led in two or three directions. The two most conspicuous were as follows: one of Ira's graduate students went to Teknowledge and was either the key developer or one of the two key developers in the KEE system. So that work that started here in a sort of toy office scheduling world eventually ended up being the basis for, I guess by most accounts, the enormously successful KEE expert system. It went off in another direction because some people at MITRE had a copy of FRL and got interested in scheduling aircraft sorties. They developed a prototype for some sort of short-range sortie plan or something. As I recall the claims of the day, they reduced that kind of planning exercise from twenty-four hours to two as a consequence of the prototype that they had built. And that has more recently become part of something I think it called the Templar System. A lot of these tracks tend to get overgrown, by the time the thing actually hits the military point area.

NORBERG: Yes, and they are very difficult to uncover.

WINSTON: Yes. But it is interesting for me, who has seen it all go through those various stages, to see that this systems work to attack aircraft deployment actually goes all the way back to Goldstein's activities in the mid- to late 1970s here.

Well, back to Heilmeier's day. I finally figured out what George was talking about only a few years ago when I was a member of DARPA's advisory group for information processing. One day I heard him talking about the problem with the ISTO office when he was director was that it was doing all this great stuff, but it wasn't explained to him how they were going to win the war. In other words, a lot of the stuff that the office was doing was telling about how well this

will save money, and this will do that, and this will make it easier to do something else, but what George was looking for was war winners. Something that would really make a difference in the case of actual context. And when he said that, suddenly it all made sense. Before that it looked really strange, because here is a guy who was feared by the artificial intelligence laboratories while he was in charge, who when he left DARPA, carried the AI crusade on himself at TI. It seemed peculiarly out of character until you realize that he must have always been a fan of the technology; he just wasn't getting the community to offer up any war winning ideas.

Those were days also when I think Colonel Russell ran the office in a sort of military style. It was orders from above that got translated into orders for below. So I think we felt that Dave was transmitting orders down but not carrying ideas up.

NORBERG: Who was handling AI then? I want to say... not Green, but...

WINSTON: Steve Crocker was for a while.

NORBERG: There was somebody who followed Crocker, though, before Simpson came in 1984.

WINSTON: Oh, there were a variety of people. There was Larry Druffel, who was involved partly. There was a Carlstrom.

NORBERG: Carlstrom is the name I was looking for. Carlstrom claims that he had to fight very hard to keep the AI message before management in DARPA in that period.

WINSTON: Yes, I think that's right.

NORBERG: Carrying the torch when no one else would.

WINSTON: That's right. And we all had a great deal of respect and admiration for him for doing it. He was clearly a visionary in the sense that he realized there was something very good going on here. It was very clear that he was working really hard to deliver that message.

NORBERG: Did you take any steps to try to change the funding of the laboratory as a result of this attitude on Heilmeier's part?

WINSTON: One is always looking for sponsors, and we looked for other sponsors, but I very quickly began to realize that ARPA was a unique resource in terms of sponsorship. And I very quickly concluded that it was a very good thing that the country has a variety of ARPA research sponsoring agencies with different styles. I always felt that DARPA could do things that the National Science Foundation couldn't. It was very clear that the differing missions of DARPA and the National Science Foundation led to different kinds of sciences. NSF after all is run by peer review, which means that it tends to favor established sciences, whereas DARPA has always been run by individual princes of science, people with great vision or a good intuition who could do things that couldn't be done if things were slowed down by an elaborate peer review process. They can start new fields if they recognize a smart, young guy like Minsky was back then and support him.

NORBERG: Well, I'm taking two messages from that analysis. And one of them is that it's all right if you're on the side of the good and the true with respect to DARPA, if your new field is the one they wish to support. But if it isn't then it doesn't matter how good your field is; there's nothing to come from that. Whereas on the other hand with NSF, while the risk is greater through the peer review process, the rewards might be spread around over a number of new fields, and therefore you get something.

WINSTON: Well, yes, they're also spread around over geography. Yes, I mean with DARPA you've got to take the down with the up. If DARPA decides it doesn't like what you are doing or think it's relevant, then you lose. And you can't complain, because it is the same mechanism that led you to be funded handsomely in the first place. So if DARPA were suddenly to up and decide that artificial intelligence is no longer relevant to its mission, then that's it.

I'd complain, of course, but I wouldn't feel bad about losing.

NORBERG: You wouldn't? Why not?

WINSTON: Well, I'd feel terrible about losing. (Laugh) I'm trying to say something sort of complicated. If the guy who is running DARPA makes a decision and it runs against you, I think you have to do your duty and complain, of course. But I would be, on another level, philosophical about it. I think I would say, "Well, that is that person's prerogative. ARPA has to run that way. ARPA can't do what it has done and achieve what it has achieved by running in any way that is different than that." So, yes, did we seek other sponsors? Yes, sure. Who did we talk to? The National Science Foundation; we talked to NASA. And what I found was that all those possible funding agencies were operating on a different scale from ARPA in terms of its funding for this kind of research. So the same effort that would produce fifty thousand dollars from one would produce half a million from DARPA, because their budgets were different, their scale of thinking was different, their station in the country's research hierarchy was different. So we sought out other sponsors, but we remained realistic about what benefits we would derive from that.

NORBERG: The second message that I heard was that you were willing to massage the labs program in various ways in order to continue to meet the DARPA mission. Is that a fair statement?

WINSTON: It is approximately right, but I think that word may be the wrong word because it has connotations.

NORBERG: Sure.

WINSTON: I chose my words deliberately. I was seeking to find intersections between what the laboratory was doing and what DARPA either was interested in or could be persuaded to be interested in. So in some cases it was a matter of pointing out the potential applications of a piece of work. And in some cases there were attempts to interest ARPA in new areas. Again, for some reason I remember the failures more than the successes. For example, at one point I remember sitting in Colonel Russell's office suggesting that it might be a good idea for ARPA to have a

program of leg locomotion, since a lot of the world can't be accessed by wheeled and tracked vehicles. And he said, "No" very directly. So that was that.

NORBERG: Was that because they only take able-bodied people into the military? (Laugh)

WINSTON: I don't know exactly why it was, but I couldn't get him interested in that. Years later, or course, ARPA did have a very influential and important program on leg locomotion. So you say "massage," no, that wasn't quite right; "looked for that intersection," that is correct.

NORBERG: I see another interesting phenomenon going on at MIT at that same time. This is the period in which the laboratory for computer science changes management also, in 1974. And Dertouzos comes aboard as the director. What sort of an interaction did you have with him as directors of two important laboratories here?

WINSTON: Well, let's see. I'm going to defer that question because there's another thought on my mind that I think might be worth getting onto your tape. And that is that I think, looking back on the 1970s... several thoughts come to mind, all disconnected from one another. One of the big mistakes I think we as a community made in the 1970s was to suppose that the applications benefit of artificial intelligence would be to save money by replacing people. And looking at the world from a commercial point of view today that appears naive, because for the most part people aren't interested in saving money by replacing people. People aren't interested in saving money in general. What people are interested in is making new money. The reason for that is plain. There is only a finite amount of money you can save by replacing people, but you can always dream of unlimited new revenues if your idea is on the revenue side rather than on the expense side. So I think today I tend to think about the applications more in terms of making people smarter rather than in terms of replacing people. If you are thinking about this from a robotics perspective, the chore isn't to replace people on a production line; that doesn't do any good anyway. The cost of production is often a small fraction of the total cost of the manufacturing enterprise. The real benefit in something like robotics is to make it possible to build products that you couldn't make before. The real benefit of an AI system of today is that it opens up new opportunities rather than saving money.

[INTERRUPTION]

Well, let's see, we were turning to the relationship between the Artificial Intelligence Laboratory and what used to be Project MAC. Some of this of course goes back to before my time. After all, the Artificial Intelligence Laboratory had its roots in RLE and then when Project MAC was formed it went from RLE to be a part of this new thing. And then in about 1970 it split apart from it, mostly over a squabble about how ARPA money was being divided up, or something. And then there were several Project MAC/LCS directors in the early 1970s. I guess it went from Fano to Licklider to Ed Fredkin and finally to Dertouzos.

NORBERG: What I was thinking of when I asked the question was whether there was any coordinated efforts on the part of both of you to encourage DARPA in particular directions.

WINSTON: Not coordinated. Both laboratories had, at the very highest level of abstraction, the same mission, to ensure that DARPA continued to be enthusiastic about supporting computer science in general and our particular programs in particular. But not coordinated.

NORBERG: Because Dertouzos claims that he had some particular ideas in mind, and indeed there is good evidence to show that he tried to get DARPA interested, and in some cases he was quite successful and in other cases not. He spent a good deal of time, apparently, at this. At the same time he was trying to change the funding of LCS in order to make it less dependent on DARPA. I'm just wondering whether there was any carry over there, whether there was any discussion between the two of you at any time. Maybe offhand discussions, it doesn't necessarily have to be anything formal.

WINSTON: No. Not as much as perhaps there could have been or should have been. The two laboratories have always had a different style and a different mission. To some people that is invisible; to other people it is very visible. We inherit from the good old days a much more informal and unstratified style than LCS has had. We also view our mission to be different. Their mission is to do computer science research, and our mission is to understand

intelligence. That means that computer science is only one of the many disciplines that people in our laboratory interact with. We always felt it was important to maintain that separate identity as a consequence. So if you look at our faculty during the early 1980s, there was a big swing toward robotics, as a consequence of some new funding. That in turn led us to have a considerable fraction of our faculty from the Brain and Cognitive Science Department. And we have mechanical engineering faculty now in the laboratory as well. So I think we view ourselves as having a somewhat different mission and necessarily having a separate identity as a consequence. In the mid 1970s there were discussions, attempts to merge us back with LCS. And we seem to have resisted that.

NORBERG: Did that in any way give rise to a competitive attitude about funding from DARPA?

WINSTON: I think so. I think there was a sense of competitiveness associated with it.

NORBERG: If they get the money, you don't, and visa versa, that sort of thing?

WINSTON: Well... I don't think it was ever that. Competitive doesn't necessarily mean that you think you are going after the same money. I think the competitiveness was of a different form.

DATE: 2 May 1990

TAPE 2/SIDE 1

NORBERG: When we finished last time, we ended the tape with just the beginning of a discussion of AI as a sub-field of computer science. And your last comment as I remember it was that AI people felt that they had to prove that AI is an important place for doing work and an important place within computer science. How was such proving done? Can you elaborate on that for me? What sort of justifications were provided for the field as you understood them, and so on?

WINSTON: Well, I should start taking my own notes because this leaves a couple of topics that I want to be sure to be get down to.

NORBERG: If you want to take those up first, go ahead.

WINSTON: No, that's okay.

NORBERG: Whatever is on your mind at the moment.

WINSTON: Well, maybe I will take one or two of those side issues up before we get back to that. First of all we discussed last time my perception of my role vis a vis interfacing between ARPA's needs and our faculty's and staff's scientific interests. And one of the things I tried to do in those days was in fact to convince people at DARPA that artificial intelligence should be viewed by them as a field rather than as a project. And that was an important thing because DARPA has always had the idea that they would get in and out of projects in something like a five-year time frame. And we certainly didn't want them to get out of AI, because AI should be viewed as a field like computer science is a field or physics is a field, with projects to be run inside of it on a five-year rolling basis, but not it itself being viewed as a five year thing to be gotten into and out of. So that addresses the question of AI as a field.

Then there's another thing that comes to mind when you talk about AI being a sub field of computer science. I don't view it quite that way. I think AI should be viewed as a field that intersects with computer science but that also intersects with other fields. If you look at our present faculty, you'll find that forty percent of them are in the Brain and Cognitive Science Department. They are computational psychologists. I don't know if that is an official term, but that's what they are. Other people who find their home in an artificial intelligence laboratory would be mechanical engineers interested in the mechanical side of robotics, or computational linguists who had interacted a great deal with the linguistics community, which is after all the community that holds the custodianship of large amounts of knowledge about language. So I've always felt that it is constraining to think of AI as a sub-field of computer science, in part because we reach out in different directions and in part because that is a natural consequence of a

different mission. Computer science has the mission, I suppose, of making computers do ever more wonderful things - make them faster, make them this and make them that, and understanding the complexity and all the theory of that. To me artificial intelligence has the different mission of understanding intelligence from a computational point of view. So we have a lot of heritage that comes from computer science. We're different from other attempts in understanding intelligence because of our roots in computer science, but I think it is uncomfortable for me to think of artificial intelligence as a strict sub-field of computer science.

NORBERG: Is that a widely held view do you think among people in AI?

WINSTON: Widely, I suppose, means more than 50 percent. Probably not. I think it is a view that is held by a substantial minority. It would certainly be on that I think would be held by about 40 percent of our faculty members in this laboratory who come from Brain and Cognitive Science.

NORBERG: Well, how then did this point of view develop, do you think?

WINSTON: This point of view?

NORBERG: Yes, your point of view. Not in your individual case, but among this say 40 or 50 percent.

WINSTON: I think it is not a top-level view. It is corollary to a top-level view. The top-level view is "What are you for? Are you for advancing the state of the art of computers, or are your passions focused on this question of what is intelligence?" And understanding what intelligence is requires understanding more than what you would understand if you just understood computing. You have to understand the way the world is. It is helpful to understand a little bit of animal intelligence. Not because we are interested in simulating animal intelligence, but because our own intelligence in particular is an example of an intelligence, and the way we seem to have been built and factored the various necessary tasks is demonstrably one way of doing it. It is helpful to understand one way if you are trying to understand ways in general.

NORBERG: Getting back to the original question that I asked you this morning. Given that there are at least two different views of the relationship, let's say, between computer science on the one hand, and AI on the other, and maybe AI and other fields, how do you think, historically now, the field has attempted to justify itself, the AI field, how has it attempted to justify itself both to other people on the periphery of AI and to the funders?

WINSTON: Well those, I think, are perhaps different questions. Let me try to address them one at a time. I think the way the field has worked with its sponsors has always been to cite potential applications and to stand on previous applications. And the shift between those two is of course a function of the maturity of the field. In the early 1970s, we were almost exclusively talking about potential. And here in the 1990s we've got plenty of successful applications to brag about. And in between it has been a mixture of those in proportion to how much had been done. I think up until the early 1980s, talk of actual applications was mostly talk. And even until 1982, or whenever your study stops, we were mostly talking about a handful of successful applications. Up until about 1982, you could still count on the fingers of your hands everything that looked like it might be construed as beyond a prototype in terms of its successes in application. Now that we've had exponential growth or something, now there are thousands so that establishing credibility is no longer a problem. It is interesting to note that at the time there are thousands of successes, at the same time there haven't been very many successful companies. But that is a consequence of poor business judgement rather than the technology. I'll say one word about that, just because I probably shouldn't leave that thought dangling. I think that many companies that started up in the early 1980s during the euphoria period of AI failed because they were filled with the euphoria of being an AI company. And you can't sell technology; you can only sell solutions. And those companies didn't have solutions to sell; they had technology to sell. And that only works in periods of euphoria. Once the euphoria has worn off, then people are interested only in getting their problems solved, and many of those companies were not accustomed to that idea, they were not equipped for that idea, they didn't have the right personnel for that idea. As a consequence they didn't do well. Nevertheless, if you look inside of many big companies--Dupont is a conspicuous example; the logistics work of the United States Air Force is another conspicuous example of what appears to be a tremendous success--you find that there the work is driven by the problem rather than by the technology, and that is where the big successes are.

NORBERG: Can I ask you just one incidental question then? Can you name a couple of big successes for me?

WINSTON: Yes... big successes. Well, they go all the way back to X-con I guess, which is the big success that launched the AI business. Dupont is conspicuous because it has three or four hundred successes by their own count, in terms of little tiny expert systems that help on small problems in the chemical industry. I mentioned the Air Force thing. All these things tend to be argued here and there, but there are claims that the Air Force has saved itself half a billion dollars, by now more I suppose, in the area of purchasing through the use of simple expert systems. The American Express system gets wide press these days for being another conspicuous success. It is a success of a slightly different generation from the rule-based expert system successes, and to my mind it marks the beginning of a sort of second generation of success. So there are lots of successes, and most of which are in Feigenbaum's book. Feigenbaum has done a pretty good job of collecting together some representative examples. I wouldn't want to suggest that his book is complete by any means, but the samples are representative. Now, let's see we were...

NORBERG: The first half of the question had to do with justification among colleagues.

WINSTON: Justification. Well, justification is always been "Let's understand what the problems are, and talk about how the technology can be deployed to deal with those problems, and let's buttress our arguments by citing past successes." I think that is how we always tried to work with our sponsors. And we talked a little bit last time about my role in helping to make that marriage between the technology and the problems, and I have thought about that a little bit since. I think to a large extent a lot of what I was able to do in the 1970s was provide basic coaching skills in presentation, so that people with a very complicated technical story to tell and who would be inclined to tell a very technical story and leave the rest to be guessed or to be inferred. I would always encourage them to find a way of articulating how their work would be relevant to a DARPA goal and getting that out in the first sentence or two. Showing one's hand early, I guess, is the message that I continually harped on. Showing one's hand early, making sure that the work was not presented in such a way so as to inadvertently make it sound like it was an endless enterprise. I learned very early on that phrases like "continue to work on" are not well received by sponsors,

because it sounds like you'll be continuing to work on it for the next two hundred years. (Laugh) So I encouraged people to think instead in terms of the kinds of terms the DARPA people told me to think in terms of.

NORBERG: Which were?

WINSTON: Milestones that have words in them like "report on," "test," "bring up to a measurable level of capability," "intermediate measurables," "We'll try this algorithm on one picture, we'll try it on twenty pictures, we'll try it on two hundred pictures." Not just "We'll continue to work on this algorithm." I mentioned last time that George was driven nuts by big proposals with no milestones. Well, we tried to understand what George wanted and others and to make sure that we were comfortable with what they wanted and that we provided them with the stuff in the form that they needed it. So I think I had a talent that was useful in those days in that direction of presentation skills, making sure the message got out early in a presentation, and that the presentations were kept focused on that intersection of our scientific interests and DARPA's applied interests.

NORBERG: Where do you think you achieved those skills?

WINSTON: I think I gradually created those skills as a consequence of just interest in teaching and an interest in verbalizing heuristics for good communication through lectures. I often, especially in my early teaching days, would sit with other people interested in these things in the back of a room and talk not only of the material that was being presented but the style in which it was being presented and our perceptions of what was being done well and what was being done not so well. These days I give an annual lecture at MIT on lecturing, which is a presentation of a compendium of all the various heuristics I learned over a period of time. See this is a natural thing for an AI guy, to look at what many people regard as a skill and try to understand its structure and develop a *armamentarium* of heuristics for improving one's capacity in this skill area. I am making more things explicit so they can be talked about and improved on.

Well, let's see maybe we should now turn to the question of AI proving itself to people in other fields. I think if we're

talking about the 1970s, we are talking about something quite different from what we are talking about today. In the 1970s, there was all sorts of speculation about whether computers could be intelligent and whatnot, just as there continues to be today. The problem is that in those days there were only a few successes to be pointed to and critics could always fall back on the argument that "You can't do this, because you haven't done it yet." I mean, no one ever said that, but if you look at their tangled and tortuous arguments and took them apart, it seemed that that was at the core of many of them. And it is hard to deal with that. We hadn't done it yet. We hadn't done anything yet. So I think it was natural that in those days we developed a kind of insularity in the sense that we said, "Okay, the hell with you guys. We'll go off and do this, and we'll get in touch when we've solved the problem." This was never made explicit by anybody either, but that's my sense of how the community was feeling, if the community can be said to have feelings. So we didn't talk to linguists, and we didn't talk to psychophysicists, and we didn't talk to even the mechanical engineers. We just went off and reinvented a whole lot of stuff as a consequence of that isolation. But I think it was a necessary period of isolation because it was a time when we were very fragile and we needed just to roll up our sleeves and get to work and separate ourselves from endless debates about whether we could do it and just go out and do it. On the periphery, a conspicuous example of that is time-sharing. I guess I've kind of lost track of who actually invented time-sharing. I know John McCarthy must have been heavily involved in it. Maybe it's arguable that AI was an influence or the demands of AI were an influence on the development of computing, even in that context. But in any case, in the early 1970s here at MIT, we were doing amazing things with ITS. The hackers were: Richard Greenblatt, Tom Knight, and Jack Holloway. Knight put up what I believe was the first Bitnet display on ITS. And this was at a time when a lot of the faculty members in other parts of computer science had no idea that this sort of stuff could be done or had already happened up here in the relatively isolated Artificial Intelligence Laboratory.

Well, what's different? I think things began to change a little bit, or shall we say change a lot, as we gained more confidence in what we'd accomplished. And by the late 1970s, we began to see people working with people in other fields perhaps for the first time, in this laboratory at least. I suppose Newell and Simon have always viewed themselves as psychologists of a sort, but in this laboratory we were pretty isolated until the latter part of the 1970s. At that time, Dave Marr was a big influence here and slowly the psychology department of that day began to get

accustomed to and enthusiastic about David's work, and he became a faculty member in the Psychology Department. And his work is a testament to the usefulness of understanding what has been learned in another field and looking at it from a computational point of view. His genius was to understand the utility of that. So he invented stereo theories that were invented and rejected as a consequence of their biological implausibility, and then he went on to other stereo theories that were more consonant with biological observations. He regrettably died too soon, and I think as a consequence of that his legacy will be his methodology more than his actual scientific results, because he was only part way through that process when he was taken ill with leukemia. He was a big influence. He was a big influence on his own students and on the other faculty, like me for example, who admired what he was doing. That was not universal either, of course. I think Marvin Minsky, to this day, thinks that that approach to vision did more harm than good. Marvin's impatient and wants it done and felt that those guys were moving too slowly and were not thinking... Well, Marvin, would have to characterize this himself because I can't speak for him very well, because on this particular point I think I disagree. But to attempt it, I think Marvin feels that those guys were focusing too much on low-level early vision questions and not enough on higher-level, more cognitive things that in the end I think Marvin feels are more important to visual understanding.

NORBERG: But you say you think you disagree with that position?

WINSTON: Yes, I think I disagree with that. I think it is very important to understand those details of early vision as a means toward understanding the whole of the visual understanding. Let me elaborate on that. I was never very enthusiastic about the early work on speech at Carnegie, because it was always billed as successful because of the blackboard architecture. And my view was that if you didn't get the early phonemes processed well, it didn't matter what architecture you had on top, you'd never succeed. But on the other hand, if you got them processed well, and once again it didn't matter which architecture you had on top, because the final process can be handled by almost anything. By analogy I think you can't do much in terms of understanding an image unless you understand what you need to know about the early vision processes well enough to support the higher-level reasoning processes. So yes, I guess I believe more than Marvin does in the importance of understanding things down at that grungy level of detail where you are struggling to deal with edges and things.

NORBERG: Is it possible that these two positions will converge?

WINSTON: Yes, absolutely, I think they're converging now. As we speak, there's a connection machine upstairs that might be humming away on all sorts of vision processes with a view toward understanding their integration. And that's possible now because of all that work that was done on edges and shading and stuff in earlier times. So integration is very much a theme of today, and Shimon Ullman, one of our faculty, takes it to an even higher level and tries to understand how early vision procedures might be driven by higher-level requirements. He calls this a theory of visual routines. He is partly motivated by psycho-physical observation, namely that some aspects of a scene seem to be computed almost instantaneously and others seem to require a serial scan across the image. He would argue the things that seemed to require a serial scan are driven from above.

Well let's see, where were we? Justifying ourselves to other fields. I think my view of what was happening is that we were often criticized by Dreyfuss and others as not doing anything worthwhile, by people in other fields saying "Oh, there's nothing to it." So our determination was to separate ourselves from that criticism and go off and do it and prove ourselves by accomplishment. Once we had survived our fledgling stage, then I think we felt more comfortable working with people in a wide field because our identity had been established. It was no longer a danger that AI would be subsumed by some other field and disappear as a discipline. It was on the map so we could begin to formalize this without risking our own demise, around the late 1970s. It was about that time, you know, that Stanford guys like Feigenbaum, and Shortliffe, and Davis were working with physicians. Here Gerry Sussman was working with real electrical circuit designers to understand how they do it. So it was not only a time when people began to collaborate with other people interested in intelligence, but with other people interested in other real people and real application demands as well.

NORBERG: Didn't that idea start somewhat earlier than the late 1970s though? I seem to recall, was it Shortliffe and Davis, or Shortliffe and Duda? I can't remember the specific piece now. In reviewing rule-based expert systems in *Science* magazine, six or seven years ago, commenting in the opening paragraphs of the piece that in the early 1970s

AI researchers gave up the idea of trying to find general principles of intelligence and began working more specifically on very narrow areas in order to, as you would it seems to me say now, "prove themselves" or prove that the field can accomplish something.

WINSTON: That's a slightly different twist, I think.

NORBERG: On my part, or on their part, if I've represented them correctly?

WINSTON: Let me tell you what I think, and we can figure out if we're saying the same thing. I'm saying that in about the mid-1970s people, at least here and I think at Stanford too, began to ask themselves about real applications and began to be found in company with real experts in those applications areas. You are talking about a slightly different thing, namely the time when people gave up searching for a holy grail, a universal mechanism, and began to feel that maybe this enterprise of intelligence would require lots a mechanisms, not just one universal mechanism. Now I think 1970 sort of marks the beginning of new thinking about that. If you look at the GPS and that sort of stuff, I think that was late 1960s. And it was at that time that Newell and especially Simon were making predictions about how smart the machines would be in ten years, not because they were foolish, as John McCarthy once pointed out. This is my recollection of what John McCarthy said once. McCarthy was saying that was not a foolish remark that Simon made; it was an attempt by a responsible scientist to be responsible to society about something he thought really could happen. And what he thought really could happen was that the general problem solver might be deployed on the problem of making itself smarter to produce a sort of chain reaction that would lead to an intelligence of phenomenal capability, which would in fact, had it happened, have had tremendous impact on society by now. So there was poor Simon trying to be responsible in a time when everybody was saying that scientists should be responsible, only to get elevated by his own petard, because something that very well could have happened didn't actually turn out to happen. To this day, Newell I think believes in universal engines. I think he believes that SOAR is a universal theory of intelligence. But I think in the early 1970s, around 1970, some people began to think there was more to it than just one mechanism. Now we have that diversion complete in the form of Marvin's Society of Mind theory, which is a complete antithesis to the SOAR attitude that there is one universal mechanism. So when I said

mid-1970s I think I was talking more in terms of collaborations with real scientists. When I talk about early-1970s I think we're talking about a shift in how people in AI were thinking about how AI would evolve.

NORBERG: Do you think the shift then produced the interaction with these scientists?

WINSTON: Yes, sure. It's that shift in scientific thinking as well as a desire to get on with the production of real demonstrable application successes.

I have another thing on my notes that I meant to be sure to bring up. That is the last time we talked about the tremendous missionary work done by David Carlstrom in ARPA in some of the more difficult times. I meant to mention, and I don't think I did, that Marvin Denicoff at ONR was a tremendous influence on us during that difficult period. In Marvin Denicoff we felt we had a visionary ally who would always be there in our defense. So we always felt that although he wasn't at ARPA, he was a person who would be telling our story at every opportunity.

TAPE 2/SIDE 2

WINSTON: Denicoff not only had a cheering effect around here, he was also constantly dreaming up new programs that ARPA might undertake that would serve as umbrellas for the research that we wanted to do. He, for example, started a program that was labeled shipboard maintenance, or something like that, that was broad enough to support some of our work in robotics as well as some of our work on scheduling and multi-person coordination. At another time he dreamed up and launched a program in very large data bases. Those things had varying success in terms of their influence on ARPA thinking and their staying power even within ONR, but they nevertheless provided seed money for things that subsequently grew into major efforts.

NORBERG: Can you be specific in one example?

WINSTON: Yes. The robotics work that Marvin started. My recollection of the dates is going to be a little fuzzy

here because this was a while back, but Marvin started a program in shipboard maintenance and that supported some work in robotics, and that seed funding enabled us in turn to obtain major funding from the Systems Development Foundation here at MIT, and that funding in turn enabled us to build up this incredible strength in the robotics area. Then once the Systems Development Foundation had spent all of its money, there was a very solid foundation for ARPA support in the area. So the trail is a little tenuous, but to me it all goes back Marvin Denicoff's support.

NORBERG: Sure. Can we invert that process? Can I ask you did ARPA support any areas that ended up doing the same sort of thing? That is the Institute shifted to support from another organization and ended up with ONR or somebody else?

WINSTON: Let's see, it's always hard to scan for examples of that sort. You're asking for an example in which DARPA provided some sort of seed funding and went somewhere else?

NORBERG: Yes, if that was ever the case. Maybe it was not.

WINSTON: Well, sure, perhaps the most conspicuous case is ARPANET, which was an entirely DARPA funded deal that eventually became a universal system of tremendous consequence. I think there are other things that have happened like that, but the problem in tracking them down is that if ARPA does a big demonstration say, and then nobody buys into it, but a few years downstream somebody does something that they wouldn't have done had that thing not happened. Sometimes it's ARPA again picking up the ball. I mean, speech is a good example of that. A big ARPA project canceled, many years later the phoenix rises from the ashes. There have been various Navy test beds for AI over a period of time, natural language testbeds in, I guess, the mid-1970s, testbeds for anti-submarine warfare, or work in that area. A more contemporary example would be the FRESH testbed in Hawaii for maintaining knowledge of Navy ships and their interactions with support for decision making for that sort of thing.

Let's go back to that Navy testbed, the Lifer (?) thing, the work that was done by SRI. I don't know what ever happened to that testbed, but that work sure became the substrate on top of which commercial natural language

systems, many of them, have been built. The work that Feigenbaum and those guys did on anti-submarine warfare with the HASP system was stuff that provided the literature foundation for current work on expert systems for that sort of thing. I'm a little out of touch with that, but a few years ago I was doing a study for NRAC, the Navy Research Advisory Committee, of which I was a member, on the question of AI and anti-submarine warfare. And you could see that the stuff they had done at HASP was before its time, but nevertheless laid some of the groundwork for new work that was being done for the next generation of submarines. I guess what I'm saying is that when you look for these threads, you see the tracks often get well covered up. Things move with the movement of an individual. I mentioned last time, I think, another one along these lines where the work that Ira Goldstein was doing on frames in the context of office automation. He moved to MITRE where they did a deal on aircraft sorties, which in turn led to the development of the Templar system. I don't know what state that's in. Last time I looked they were just issuing RFPs. Those are examples of sort of tenuous but discoverable things.

NORBERG: Let me ask you for a piece of advice then as long as we're on this subject, because this has direct implications for what I'm trying to do in developing some sort of a case study in the history of AI. And that is if you were in my shoes, knowing how difficult it is to uncover many of these connections, what do you think would be the best examples for me to try to uncover and address in order to answer two questions. One of them would be simply the sort of milestones in the field itself, in AI, that are significant say in robotics, or in vision, or in rule-based systems or knowledge engineering. Areas like that rather than the whole field. I can't do the whole field. That's the first question: how does one select the milestones and find the evidence? And secondly, how does one show the influence of DARPA in those areas? It's not to say that they had to do it all, but how does one uncover their influence?

WINSTON: Well, I think I'll have to think about those questions. Those are difficult questions. I think the way to attack those is to attack bits and pieces of them. It's impossible to gain a total solution. I'll try to nibble away at those questions as we proceed. As you spoke, another thing that I think DARPA has had tremendous influence in for which unfortunately I think is too contemporary for your study is the development of the Common LISP standard. I think that DARPA's encouragement of that evolution has been essential. I don't think it would have happened

without DARPA's encouragement. I hope DARPA does more in that area to promote the commercial and defense use of the LISP program language. But it is absolutely clear that without DARPA it wouldn't have happened. What would you have instead? You would have fifty miserable local dialects of LISP and absolutely no portability, not even a remote hope of commercial use of the language, which is of tremendous importance because of its productivity relative to all other popular languages, all other languages period.

I think the way you have to look for these tracks is to talk to the people who carry the knowledge with them. I think Richard Brown, who was at MITRE for a number of years and is no longer there, was the person who carried Ira Goldstein's ideas to MITRE and then married them to this problem of sortie planning. That in turn led to other directions.

NORBERG: Do you know where Brown is now?

WINSTON: I'm sure I could find him. I will find him, but I don't think he'd be a person to pop out of my Rolodex. When we're through with the tape, I'll initiate a search. In the case of that DARPA testbed, what happened? Well, let's see. Gary Hendricks was very much involved in LIFER (?), was very much involved in that testbed. That testbed was a foundation for shaping his ideas on natural language interfaces and that led to a very popular PC product providing a natural language database for a lot of people. So another thing to do would be to talk to Gary Hendricks about his views on how that thread worked.

NORBERG: It would be nice to trace these things back starting with the end result and the applications to brag about as you referred to them earlier.

WINSTON: Yes. I mentioned last time that you'll find Ira's work emerged in another way through the KEE knowledge engineering environment. And it happened as a consequence that that influence was carried by an individual - one of Ira's students named Greg Clements. I think he left MIT before getting a Ph.D. and went to work for Intellicorp, I guess. As a consequence of that you'll find a lot of FRL ideas in KEE, which is now the substrate for a large number

of successful expert systems. So there is a kind of broadcast type of connectivity between original work here and applications there.

NORBERG: Did you say FRL?

WINSTON: Yes. Frame Representation Language. There was a lot of confusion about that. Because at the time Xerox had a language they called KRL, Knowledge Representation Language. FRL was a sort of lean and mean approach relative to KRL. By the way, very often you find this sort of dichotomy in AI. It is most conspicuous in the robotics area. We have always had two types of arms. One type is an arm you do research with, and the other is the type you'd do research on. And FRL, by analogy, was the type of knowledge representation language that you'd do research with, and KRL at Xerox is more of a language that you'd do research on. And the ones that you do research on are always broken and fragile and never work when anybody else wants to use them. That's because they are an object of research in their own right.

NORBERG: Could you make the same criticism of some of the rule-based expert systems that were developed, too?

WINSTON: It wouldn't be criticism, in fact it would be opposite of criticism. It's the way it ought to be. It's foolhardy to try to do research in robotics with an arm that somebody is hacking on a nightly basis for his thesis. You're asking if there are analogies in the expert systems world?

NORBERG: Yes. Because I've been reading a recent criticism about MYCIN and DENDRAL and the fact that they are not very useful systems, they are really research systems, and wondering what the meaning of that criticism is as someone from outside the field who is just looking at this and thinks those are two of the best examples of development in the 1970s.

WINSTON: Well, I think that could still be true. The question is whether they have the kind of robustness you need in order to think of them as broadly applicable, perhaps even commercial, tools. If you are building a new gasoline

engine, you may be the only engineer that can get the damn thing to run, and that's what those things were. They were patched and hacked and repatched and rehacked, and not documented very well, and not supported with commercial principles of software management or anything. So it is not surprising that they might be awkward to use. The commercialization of a piece of software requires that you have a kind of documentation, and QA, and even interface thinking that you don't want or need when you are trying to do something for the first time. So I don't think there is any surprise there. Those systems were tremendous scientific achievements, but surely the software wasn't in any shape for direct commercial exploitation.

NORBERG: Do you know whether they directly influenced some of the commercial systems that are around now?

WINSTON: MYCIN and DENDRAL?

NORBERG: Yes.

WINSTON: Surely MYCIN did. This is East Coast and that is West Coast, so I track that stuff only very remotely. But there was EMYCIN, and EMYCIN surely was another influence on KEE no doubt and certainly had an influence on Teknowledge. I couldn't cite chapter and verse on that one.

NORBERG: That's all right.

WINSTON: I've got another example of that kind of spinoff. For a while Gold Hill Computers was doing quite well.

NORBERG: Which computers?

WINSTON: Gold Hill. You say your study only goes up to 1982?

NORBERG: That's just an arbitrary position.

WINSTON: This may be useful for background. There was a local company trying to get itself into business, and they started off doing a PC-based LISP and they wanted to build an expert systems shell, and they were trying to think how to do that. They finally decided that the best way to do that would be to do what they knew how to do best. So the expert systems shell reflected a great deal of what was in Gerry Barber's Ph.D. thesis done here. You know, another example we haven't talked about, I guess, is all the hardware spinoffs from this technology, and certainly DARPA has been... There the tracks are very clear. The jury is still not in on whether the whole LISP machine thing was a good idea. And if it was a bad idea, whether it was a bad idea for technical or for non-technical reasons. But after all, DARPA supported the LISP machine work here at MIT and that led to several licenses, one of which was to Texas Instruments. A lot of machines were sold and a lot of things were developed as a consequence of that development. Those companies are not doing very well in the LISP machine business these days. But there is quite another thing that needs to be observed, and that is the ideas in those machines are still a great influence to the evolution of computer hardware.

NORBERG: In what sense?

WINSTON: Well, for one thing they give the rest of the computer community something to argue about. But if you look at the LISP machine one of the big innovations... not innovations. I guess I better back up a couple of paces and do this a little more right. I often ask the LISP machine guys if we could patent that stuff. And they say, "No, there's nothing in here that hasn't been done before. The miracle is in the artful combination of things." A few new twists, but for the most part it was the artful combination of existing technology into a new concept driven by this idea of a personal workstation that was important. If you didn't get all these things right, like the display and the memory organization and the processors, terminals and whatnot, you might as well not have any of it. One of the things that was regarded as an important component was the idea of hardware data types, because LISP doesn't like to have declarations and therefore you need something down at the hardware level to know whether you are talking about a floating point number or not. These days, with RISC architectures and whatnot, there has been a kind of temporary movement away from that. But as you see things evolving toward the future--I'm told by my hardware

friends--you are beginning to see things like hardware data types emerge in the MIPS chip in one form or another. You are beginning to see a convergence of some of the ideas that were promoted in the special purpose engines and the ideas that have been promoted in the RISC line coming back together again. Where were we? I think I lost the thread here.

NORBERG: No. We're still talking about examples and the influence of DARPA.

WINSTON: Oh, yes. And then the other one I wanted to be sure to mention was the Connection Machine. There is a machine that I think had tremendous influence on what people are thinking. It's a successful company. Lots of machines have been sold, lots of applications have been developed, many of them in the clandestine community I'm told, it's rumored, things having to do with information retrieval from large text databases and whatnot. That was all absolutely positively a consequence of DARPA funding without which it would have never happened. So in the hardware area... I mean hardware has always been close to the hearts of AI people. Somehow a lot of hardware innovation comes out of AI laboratories and I've always regarded that as a consequence of the challenges that are posed by the enterprise of AI, and the attraction that those challenges have to people who are interested in developing fantastic hardware. That's why the hackers came to the AI lab because of the challenges that they found here and the egalitarian, free-wheeling, "let's do it, roll up our sleeves and go" type atmosphere that we had then and have tried to nurture since.

NORBERG: How about your own company?

WINSTON: Well, I haven't mentioned that because it is much more contemporary than your study.

NORBERG: That's all right. You see we may be able to trace some of the DARPA work influences on that company too, which would be earlier.

WINSTON: Let's see. We combined ourselves into a company for a variety of reasons. One because of the esprit de

corps that a small organization always has. One because we thought we could make some money. But it was a combination of having fun and making money. When we started the company we also felt that there were lessons we had learned from other failures that would be important to us in optimizing our prospects for success, most important of which we've already touched on - namely you need to form a company around somebody's problem, not just taking it from around the technology itself. So as we developed our marketing style explanation of ourselves, we always talked about ourselves as a company that knew how to blend established and emerging technologies. We always talked of AI with the word enabling attached to it. We viewed AI as an enabling technology, not as the thing itself. We always talked about ourselves as providing systems that solved problems in which AI was only a small fraction of the total lines of code, but nevertheless an important fraction because it made the whole system possible - perhaps. We always tell our customers that if there is no AI in the solution that's fine with us, because our purpose is to solve the problem, not to hype or promote a particular technology. So what did we find ourselves doing? Well, we found ourselves thinking that an important application area for AI is a in real-time deployment of resources. And half by luck and half by design, we stumbled into the area of airport and airline management as an area in which expensive resources need to be deployed in a sort of right now way. You can't wait for a 3090 to crank away for three hours on a linear program in order to decide where to park a late arriving aircraft at a gate. So we began to think in terms of systems that would help people do that. Help because it was clear very early on that you needed to work in what I would call a mutual blunder stopping mode. The system prevents people from making errors and people override the system when they know something the system doesn't. To give you an example, when there is a rainstorm somewhere the whole monthly schedule goes out the window, tremendous stress, lots of things happening all at once. It's very easy for people to make mistakes that cost real money. You might break the speed gate on an airplane if you park it at a gate with an awkward protrusion. You certainly don't want to have an airplane try to blow itself back from a gate where the windows are fragile. The baggage guys are going to complain that the connecting flights aren't close together. The maintenance guys are going to complain that the Boeing aircraft are on one side of the concourse. There are a lot of things that have to be kept track of, and that is where the machine can stop people from making blunders. On the other side, if Gorbachev arrives at Logan Airport, that's not something any expert system builder has ever thought of before, and that will probably require human intervention. So they operate together in a mutually blunder stopping way. And very important, I think for the future of commercial AI is that the

effort is not to replace people; it is to make people smarter. As we think about ourselves and our future as a company, I think we are convinced that the upside for the systems that we are in a position to build is in the direction of making people smarter and hence enabling new revenues as distinguished from eliminating people and merely cutting expenses. So we try to think on the revenue side rather than the expense side. Did we discuss this last time? This is a point that is often missed in talking about the utility of robotics. People will say something like, "Well, the factory labor is only twenty percent of the cost, so if you took it to zero, you still couldn't compete." I think that is the wrong way of looking at what robotics can do because...

NORBERG: We did not talk about this last time.

WINSTON: I'll say a word about it then. Because what is forgotten when people say that is that eliminating the labor isn't the important thing. It is creating the flexibility, the ability to turn around on a dime, the ability to get a new product out the door every six weeks, or every six months, or something instead of years of tooling. The ability to build a better product or a more uniform product is where you really have an impact on manufactured products with the introduction of robotics. It is not only the elimination of people; you may not even eliminate any people. It is in the improvement of product and turn-around capability that impacts the revenue side as distinct from the expense side that people should be thinking about. That is I think the way the Japanese think about it. In any case, it is the way people in my company think about it.

NORBERG: Two questions about the company. First of all what is the company's name?

WINSTON: Ascent Technology. That is a remarkable coincidence; we weren't in the airline business when we started up. Oh, I have to get back to this now. Okay, so what DARPA-inspired technology did we use? Almost none, but I'll get back to that. Almost none is unfair. What I should be saying is that almost nothing really new, because the nature of the problem is such that AI technology from very early on is what was needed for this problem. I spoke too hastily. It does have some rules in it, and rule-based systems were after all nurtured by DARPA. It does have some constraint propagation in it. And after all, constraint propagation was a computing metaphor that

emerged at least three times and in every case supported by DARPA.

NORBERG: When you say emerged three times, what do you mean?

WINSTON: Well, people argue about what piece of work constraint propagation should be most associated with. To my way of thinking, the thing that really put constraint propagation on the map was Dave Waltz's thesis on line drawings. The thing that the historians will trace it back to is some work by Cordell Green in formal reasoning. The people who aren't interested in anything until it reaches the point where you can see an application at the end of the tunnel think more in terms of maybe the work of Sussman and his associates in the electronics arena. How do we use it? We use it as a generalized... you know constraint propagation is sort of related to spread sheets and stuff. We use it sort of in that mode. What we do is we say "Well, if a plane is late leaving this airport and if it comes back to this airport, then you can calculate the soonest it can possibly come back and you may discover that its assigned gate slot is empty, which enables you to solve somebody else's problem with that empty gate slot." So it's just a little, not particularly sophisticated, but essential use of the notion of constraint propagations. So it is a legitimate claim.

What else? Well, it turns out that you have to do a certain amount of computer archeology on airline systems in order to interface to them. They were written twenty-five years ago or so; they are not very well documented. You can't get inside them because they are written in a language that nobody speaks anymore, so the only way to get the information off these mainframe systems is to have your computer pretend that it is a human at a terminal. So the stuff that comes back of course is meant for display on a screen. It is not particularly well documented, and it has been patched since it was documented. Maybe the documentation has been lost. So you really have to do a sort of archeology to figure out what is coming back. Sometimes what comes back is only semi-structured. The message may include a highly structured statement about estimated time of arrival along with a birthday greeting from one gate controller at an airport to a friend of his at another airport. So we had to use a little bit of DARPA sponsored style natural language technology in order to parse and handle that message traffic.

If you look even a little harder, you will see that to our immense surprise, we make use of the blackboard metaphor in a very peculiar way. We write all of our important variables values into an Oracle database where it finally became instantly available to any other program that wants to use them. So the Oracle database becomes very much like a blackboard with respect to the whole system. So here I am, I started off saying we didn't use any DARPA sponsored technology and now I'm reciting a long list. I guess the reason for that is that we're not taking something that was developed here five minutes ago and carting it off and using it commercially. There is a time delay.

Another thing to be said is that here we are having gone into this business several years ago with no thought of any DARPA orientation, only to discover that quite accidentally DARPA has become very enthusiastic about military transportation as a focus. So now we come full circle back to the capability that the company has to support that DARPA effort, not through research because this is a company that solves problems, not a company that does research, but rather through our capability of providing a substrate into which all sorts of research can be poured and demonstration out of which all sorts early, fast prototype systems can be demonstrated.

NORBERG: Would you say, though, that the major customer for the company from the military point of view, from the Department of Defense point of view, would be the services rather than DARPA now, if we talked about major transportation problems?

WINSTON: Eventually, sure. But all that needs some DARPA nurturing before it can reach the point where a very robust, deployable system can be developed.

NORBERG: Why is that? If you are already using this for commercial airline systems, why wouldn't it be just transferrable into military transportation problems?

WINSTON: Because the problems are only analogous, not exactly the same. You know, in a commercial setting, you have a different set of rules and a different set of objectives.

END OF INTERVIEW