An Interview with

NILS NILSSON

OH 155

Conducted by William Aspray

on

1 March 1989

Palo Alto, CA

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

Nils Nilsson Interview
1 March 1989

Abstract

Nilsson begins the interview with a brief historical overview of DARPA-sponsored AI research at SRI, including his own work in robotics, research on the Computer Based Consultant, and related research on natural language and speech understanding. He notes the impact of the Mansfield amendment on DARPA funding for these projects at SRI.

The major portion of the interview is concerned specifically with his work in robotics during the period 1966-1971. He describes the significance and relationship of this work to the larger field of AI, particularly the intellectual problems it addressed and the enabling technologies it helped develop.

In the last section of the interview he gives a general impression of changes over time (from the early 1960s to the early 1970s) in funding trends and research emphases at DARPA. He concludes with a short list of contributions to AI research that came out of DARPA-sponsored work during this period.

NILS NILSSON INTERVIEW

DATE: 1 March 1989 INTERVIEWER: William Aspray

LOCATION: Palo Alto, CA

NILSSON: There was a lot of work on a project at SRI that led to the development of a robot called "Shakey." I was a project leader of that for a while. [There was] lots of communication back and forth with ARPA. There are many, many things that have to do with ARPA's administration and how they're set up to support projects, and also many things that have to do with the intellectual development of AI as a field that came from that particular project, Shakey the robot. That project had sort of a life cycle, a history. It ultimately led to the production of this robot, and some filming of it, and some experiments that we did with it. That project ultimately ceased to be supported by them (DARPA) [because of] a number of changes at ARPA and change in what they wanted to support, . They knew they still wanted to support work at SRI in the general field of artificial intelligence, but instead of it being something where they provided us money and we provided the vision and guidance -- "it's over here; this is where we want to go" -- they felt, "Okay, it's time. We don't want to go that way anymore. That was good work. We don't want to do that; we want to do this." So we had the attitude, as do a lot of researchers, I think, that we're not necessarily in the business as "guns for hire," that you tell us your problem, we'll go solve it. We're in the business to do science. We know where we want to go. [We] look for sponsors who want to go in our direction to support it -- and to the extent that ARPA wants to do that, fine. If they don't want to do it, we'll find somebody else. More or less, except that if they pull a million dollar project out from under you, you bend a little bit. So we tried to change their view of what they thought they wanted into something that was perhaps a little closer to what we were doing, that would enable us to do much the same sort of thing we had been doing with Shakey the robot, but not exactly directed at robots -- something a little bit different. We can go into more detail on what that project was, but it involved using much of the same technology, software and everything else, except not for a robot, for what we called a computer-based consultant. We called it the CBC Project. That project had to do with building a computer system that could give advice to an apprentice technician about what that apprentice technician should do next in order to achieve a certain task: repair some equipment or something of that sort. From our point of view, instead of using the motors on Shakey the robot we were using the muscles of the apprentice technician. We were still thinking of ourselves as doing robotics. We were doing the reasoning, planning, and figuring out what should be done, and so on. The technician was recording his observations to our system. He was on the one hand a sensor, telling us what things were like. On the other hand, we would do the reasoning and use him as a high-level operative: if you don't know how to do this then here's how to do it -- here's more details. If you do know how to do this, do it. Tell us when it's done. They didn't quite know what to make of that at DARPA, and they let us continue with that for a little while.

ASPRAY: What is the time period?

NILSSON: The initial work on Shakey was proposed in about 1966, I think, and that led ultimately to DARPA funding in maybe 1966 or 1967. The project probably went through this transition perhaps in about 1971, 1972. We'd have to look at exact dates, but somewhere in there.

ASPRAY: And during that period, what percentage of the Shakey research was being funded by DARPA?

NILSSON: All of it. We had some additional research funds at SRI at that time -- small funds from the Office of Naval Research. At SRI there was a certain tendency to put all this money into a pot, and [just do] what we're doing. Of course, you have to untangle it for the reports. But our attitude was that we were heading in a certain direction. We got some money from ARPA. We got some money from ONR, and so on. But, 90% probably was from ARPA. There were some changes at ARPA. The person who was supporting the Shakey the robot work... Let's see, I've forgotten exactly who was who during this time period, but at the very beginning of the Shakey the robot work it was Ivan Sutherland.

ASPRAY: Right. Then it would have been Bob Taylor.

NILSSON: Then Bob Taylor was happy with it. Then Larry Roberts was more or less happy with it. And then near the end of the Larry Roberts period they decided they didn't want to continue the Shakey the robot. And there was Dave Russell.

ASPRAY: Who was an interim director, as I understand.

NILSSON: He was kind of interim. They don't ordinarily let the military people... They try to get some technical person from the technical community, although that changed too. My perception was there were two feelings about this robotics stuff from the standpoint of DARPA and the military people. One feeling was [that] it's too far out. Maybe there will be robots that will help us on the battlefield, or will help us in supply depots, or warehouses in 2030, or past the year 2000. Maybe we should be working on these things, but it's too far out. We have to mobilize the technical community to deal more with some specific problems we have in this decade. But on the other hand I think they felt [that] maybe it was also a little too scary to be supporting. It was somewhat contradictory, because if it was so far away how could it be so scary? But, you know, "robots". We're telling our congressmen we're building these robots that are going to do this, that and the other thing, and they want *people* to be in control. So they wanted us to orient the project much more toward what they called Command and Control Applications, in which a commander, given all the information would be aided by various kinds of decision aids, getting reports from the battlefield or whatever, ultimately would be making decisions, but would have certain suggestions presented to him.

ASPRAY: How would something that's strictly robotics fit into this Command and Control work?

NILSSON: Well, we weren't too interested in the Command and Control stuff. David Russell started saying things about Command and Control and we thought, "Fine, but we're not too interested." It was Licklider, and Licklider's staff (?), that decided we aren't going to do the robotics anymore. Licklider wasn't too sure what he was going to have us do. Russell, who was a deputy or in the office at that time wanted to move us more toward Command and Control. But we weren't too interested in that, so we suggested this Computer Based Consultant idea to Licklider as a way, from our point of view, to continue what we were doing but to appear to be responsive to their needs for somewhat closer-in things. We pointed out that the military had huge bills [for] repair and maintenance, and it was very, very difficult to train technicians to repair complicated equipment; and these kinds of reasoning systems would be very important in upgrading the general level of competence in shops. So we thought [this] was closer than 2020. It might not be exactly deliverable in two or three years, but then, after all, it's called the *Advanced* Research Projects

Agency. They ought to be looking at this sort of stuff. We thought it was just great. We were happy to go along with bending our work, at least to that extent, and getting rid of the robot. Then there were some reports written. Some interesting technical stuff came out of that. Perhaps later we can talk about the stream of technical ideas that emerged from all that work, from Shakey, and the Computer Based Consultant. At the same time, or right about then, ARPA also decided to get involved in speech understanding. There was this big study to do speech understanding research and proposals. SRI fielded one of the teams. It was CMU, and Bolt, Beranek and Newman, and a combination of SRI and the Systems Development Corporation. So we had a group in our artificial intelligence research center doing speech understanding work.

ASPRAY: This would be early 1970s now?

NILSSON: This is the early 1970s. This is during the transition, moving from robotics into the Computer Based Consultant -- 1971, 1972. I'm a little sketchy on dates. At SRI, we weren't doing much on low level speech with a microphone and acoustic processing. But we were doing more on the parsing, and once you've got it into words, how to make sense out of it. So we had a natural language understanding group. Just to make the Computer-Based Consultant a bit more jazzy, we decided to use some of it as one of the application areas of our speech understanding, or natural language work, because if we were going to be communicating with this apprentice, we ought to do it in English. So, we had in mind that the apprentice would be able to ask various questions of the system, "Where is the tool to remove a flywheel?" and "How do I take off this pump?", etc. The system would then calculate what to do, and be able to generate natural language sentences to the apprentice. We even hooked up a little speech synthesizer, and it would say things like, "Remove the pump." That was all kind of fun. As a matter of fact, it led to some very interesting work in natural language having to do with discourse strategy, because as you know, in natural language processing work, things get segmented into phonetics, syntax, semantics, and what might be called pragmatics. You're taking into account the whole discourse and the context. The context is very important for resolving ambiguities in English, [i.e.] pronouns, etc. We discovered various instances in which a particular pronoun referred not to the immediately preceding noun, or anything that happened in the last half hour, but to something at the very beginning of the task. You could block the task out into block structures -- big tasks, super tasks, subtasks, sub-sub-subtasks, in such a way that depending upon which block you were in, it gave key information as to how to resolve pronominal references. Barbara Gross did some work on discourse strategy. Lots of recordings were taken of actual people and experts discussing in English how to fix equipment. We had an expert who knew about how to take apart an air compressor, and an apprentice who had never seen an air compressor. They would talk back and forth. Barbara and her colleagues recorded all this. That informed not only our work in natural language processing and discourse strategy, but even some of our work on planning. That would have previously been robotics, but now was planning about how you generate a plan to put an air compressor together or take it apart. So that was rumbling on; we thought things were just fine. But then there was increasing pressure from DARPA. They couldn't see exactly how they were going to use this. At that time an application that was more logistic and oriented toward warehouses and so forth wasn't as neat, in their opinion, as something that was like a bomb. They wanted something on the front lines, something that exploded. So Command and Control was pushed forward again. Finally, Dave Russell said, "We're cutting off the Computer Based Consultant project." That was probably after Licklider had left. "You guys aren't going to work on it anymore." Just done. "We like the stuff you guys are doing in natural language, and perhaps some of that natural language work could be applied to helping in the Command and Control area, where a commander could type in, or have somebody type in how many ships are in the Mediterranean. There would be ability to understand that and get at data bases." So the emphasis in problem solving and automatic planning stopped. We carried on [in it] at a very low level, but we had to get funds from other places and essentially it stopped. Most of the work then shifted into the natural language work at SRI, and some techniques involving building reasoning systems that could reason to help answer questions after the natural language is translated.

ASPRAY: Were there really any alternative sources for large amounts of money for this research?

NILSSON: No, it was just DARPA. We might have been able to get $100,000, but we were running on this Computer Based Consultant project at probably a million bucks a year or more , having six to ten people, having to support a lot of computers, building special equipment for it. So DARPA was the only game in town on that one. I think that the combination and the interaction of DARPA politics and what they saw their mission as being was also occurring during the time of the Mansfield amendment. They were reacting to that. That combination, with what the researchers were trying to do, really did change the emphasis of the field -- not only at SRI, but at some other places too, because there was some other robotics work going on at about that time [that was] supported by ARPA. The Stanford AI Lab, under John McCarthy, had some work in robotics and vision and so forth. That worked, although you will have to get the details from McCarthy, who was certainly influenced. At about the same time the so-called Lighthill Report came out in Britain.

ASPRAY: Right, that devastating report.

NILSSON: He was against some of the stuff that the AI people were doing in robotics. At that time there was some very good robotics work and work in planning done by a man named Austin Tate at Edinburgh, that was closely related to some of the work we were doing at SRI. You might want to talk to him too; you'll get a feeling of how people overseas perceived us. That is just now getting restarted again over at SRI, here at Stanford, MIT, various places. It would be interesting to see, coming from completely outside of ARPA, how long it took after ARPA cut off funds in that area for the field, to regroup and then get started on it again. It's still an intellectual problem... Now there is funding again for it. Even ARPA is providing some of it. But the tragedy of the whole situation was that if that had gone on we would have been much farther ahead now.

ASPRAY: Can you succinctly characterize the intellectual problem?

NILSSON: Yes, I'll try. Let's see, is there anything else connected with all that? SRI did then move into this Command Control subject -- some interesting stuff happened from that too. The Shakey project gave us an umbrella to work on a cluster of problems. One of them had to do with how a system represents the world that it's in and the actions of which it's capable so that given a particular goal that the system has to achieve (that goal could be to accomplish something in the world, to configure the world in some particular way) and given the fact the system has some idea of its own capabilities, of what their effects are, how does it string together a plan of action automatically, knowing that if it does this action, a certain effect will happen, [and if it does another action] another effect will happen. How does it string together this chain? How does it write itself a program of things to do in order to achieve its goal? People presumably do a lot of this. We, of course, act out of instinct and do a lot of things automatically. But occasionally we think, and we decide: "If I do this, and then I do this, then such and such will be the case; but I shouldn't forget that I have to do this thing first, because I have to remember [for example] to take the car keys if I'm going out to the car, because once I get there if I don't have the car keys I can't start the car, so then I've got to..." We make that problem easier for ourselves by always having our car keys in our pocket. Sometimes you have two cars, and you don't have both car keys -- you have to remember it, right? So people occasionally think about this stuff. One of the main intellectual problems was how can we get a robot system to do that? After having done it, there's the other problem of, given the plan, how to execute it? It's one thing to have a plan. It's another thing to carry out the steps in it and to monitor whether or not the things you predicted were going to be the case after each step, because if they aren't, maybe you're off the track. So then there's a question of replanning, and how to recognize that you're off the track, how to replan, and start over. You're off the track, so just start the problem all over again. You're in a certain situation, [and you need to] sense as much as you can to see what situation you're in, find out what the goal is, and make a whole new plan all over again; you do that at every step. But that's somewhat wasteful, because in many cases the new plan will have a lot of shared features with the old one; it's just a matter of getting back on the track. So that was one intellectual problem, dealing with all that. The second intellectual problem was a problem of integrating the reasoning system that's doing this reasoning with something that actually operates in the world. It is one thing to reason about what you should do; it's another thing to actually do it and to allow the sensors to affect the information you have that you reason with, and then deciding to do, and then actually turning on motors. I think Lewis Carrol has said some things about that. So, [the problem is] actually building such a system [and] making it all work together. At the same time there was a lot of work on vision in artificial intelligence. We had a television camera on Shakey the robot, and visual processing routines that could extract information from the image about what the world was like, where doorways were, where big objects were. We didn't want the project mainly to be a research project on vision. We wanted to use whatever vision abilities other people had produced at that time, which weren't all that great then, but they were adequate to help Shakey move around. We had to cheat a little bit and paint the baseboards dark compared to the light walls, so it could see where the floor and walls joined -- stuff like that. So the first [intellectual problem was] planning and executing plans, and the second one was integrating a

system that has sensors, action, and reasoning altogether. We didn't appreciate the depth of that problem then as much as we do now. This has emerged as a central theme of intellectual ferment in AI in the last two or three years. It would be worthwhile to look into that in some detail. I think we would have faced those problems and been farther ahead had we continued on in 1971. But now, in 1989, they're very key. The key thing that comes up there is that in reasoning systems, if you're really going to reason about what to do, there are some computational complexity difficulties. It might take an unbounded amount of time to come up with a plan. So there are real questions of how much planning do you do. Yet, in the real world you have some real time constraints: you have to perform certain actions by a certain time. So, there's a trade-off between how much reasoning you do and how much knee-jerk reaction you do instead of reasoning. Integrating those
two things in some kind of elegant, seamless architecture, which reacts when it needs to react, but plans when it can plan, is a challenge. It has led to a lot of activity lately. There are some people who have despaired of reasoning. They are just building pretty much reactive systems. Rob Brooks at MIT is one person. He's building some very nice little robots that he thinks of more as creatures, not reasoning systems. He's not trying to get back to humans right away. Can we even build a turtle, or a bee, or something? Nothing to do with reasoning; they're just reacting. He feels even that is hard enough. And he thinks of himself as being in AI. So he's devoting most of his attention to that. There are some other people at MIT, some people here. I've also been doing some work on the reaction side of the spectrum. There are other people who point out that, "Well, it's one thing to react, but if you don't think a little bit, you might have some problems." So that's the second intellectual problem: combining reasoning with action. We were just beginning to face that in Shakey. Now [there is] the third one, which we didn't really deal with with Shakey the robot, but we're getting around to it. As a matter of fact, some of us wrote a paper called, "New Directions in Robot Problem Solving" in 1972, Richard Fikes, and Peter Hart and I, that also served as the proposal to ARPA to continue new work in robotics in 1972, which they rejected. We had some ideas of things we wanted to do, and now we're facing them. That's the third intellectual problem: that it's one thing to build a robot that works in a world where nothing happens, except whatever the robot does. Time is of no matter there. It doesn't matter that it takes an infinite amount of time to solve a problem, because time isn't moving. Nothing's going to happen, unless you, yourself, want to get the job done fast. There are no other robots; there are no trains coming. So the other question is, if you're going to build a robot that exists in a dynamic world where other things are happening, then

there is this time pressure. So not only does a robot have to be able to predict the consequences of its own action, it has to have a model of the dynamics of the world and what is going to happen. It has to be able to do one of two things: it has to be able to accommodate itself to the dynamics, knowing the world is going to happen as it happens, [or] it has to interleave whatever its plan is with the dynamics of its world. So that means it has to be able to reason about time -- not only reason about objects, but reason about how time is moving and what's going to happen at various times. That's a great challenge. There's been a lot of work in AI on that. [It is] one thing to accommodate itself; [it is] another thing to interfere and say, "Well, unless I do something, the world's going to, turn out this way. I can actually interfere in what's going to happen in the world and it won't turn out that way if I do such and such." -- very interesting philosophical problems about that, you see, because you immediately get to this whole point about free will and determinism. If the robot predicts that it's going to interfere and the world won't turn out that way, it has to not get so confused about that that it decides, "As long as it isn't going to turn out that way I don't have to do this." It has to think of itself as being part of the world that it's predicting. That might cause certain kinds of problems, just as it sometimes causes people [certain problems]. So, some very interesting questions came up there. Now, the next point that we didn't conceive of wanting to work on in the Shakey days, but many people are working on now, is that not only are things happening in the world, dynamics, but there's a question of, "How do you model those things that are happening?" It's one thing to say that there's a block on top of another block, to model a static world. How to model a dynamic world? Well, there are always differential equations. Physicists have been modeling dynamics for a long time. But in a certain sense that's too detailed. There are a lot of people in artificial intelligence research who are working on what they call "qualitative physics". They don't care exactly about the differential equations. They merely know that if you take a basin of water and empty it, it will be empty. You don't have to do the fluid dynamics. John McCarthy talks about if a speaker is sitting some place and somebody comes by with a cup of coffee and spills it and it lands on the table, he'll jump back. The hydrodynamic calculation is to figure out what the flow would be; he just sort of knows he's going to get wet. So how do you do that? There is a lot of work in AI on "qualitative physics" and modeling these hypotheses. But now there is another thing. Suppose in addition to routine, interesting physical phenomena that are happening in the world, there are other robots. Of course, you could model them in terms of their transistors, and the circuits in them, and the effects of those circuits, and so on, just as the qualitative physics people might. But if they are very complicated robots, that's going to be a little hard to do. You might, instead, want to ascribe to them ideas that they make plans, and that they have goals and beliefs, and if this robot believes something and has a certain goal, it's going to try to act to achieve that goal. And it probably will. I can predict that something is going to happen, because that robot is just as smart as I am. So, how do many robots deal with each other? From the one standpoint, how do you merely accommodate your actions to the actions of the other robot? Or from another standpoint how do you interfere with them? Maybe you can communicate with that other robot, tell it something that it doesn't know. If it accepts what you tell it, instead of it planning to do one thing it will plan to do something else or plan to achieve its goal in some other way, or maybe you can even change its goals. So there is work in what I call intelligent, communicating agents, where there are communication anomalies. The reason for saying all [this] is there is a whole stream of intellectual problems that I think were well on their way to getting started in 1971 or 1972. Some of them couldn't have been pursued, because there has been a lot of interesting work that has happened independently in the late 1970s and early 1980s. But now that it has happened, it allows us to continue.

ASPRAY: What are the enabling results that have intervened in this period?

NILSSON: Two or three things. One enabling [result] is a much better ability to deal with representing another agent's knowledge. There is work on representing what it means for another agent to know something, and how that agent might reason with that knowledge. People like Bob Moore, who is now at SRI, Joe Halperin at IBM, Kurt Konolige, and others have worked out very elaborate representational schemes for reasoning with knowledgeable agents. That's absolutely required in order to build a robot that can act effectively in the company of other robots, because it has to be able to be reason about what they're doing. That's one thing. The second thing is that we just began, in the Shakey the robot project, to deal with this very important idea in AI called the "frame problem". That has attracted a lot of attention of philosophers too, by the way -- [there is] an interesting little book by Pylyshyn; a collection of articles by AI people and philosophers. We proposed a system called "Strips" for dealing with that problem in Shakey the robot. However, it's a very, very deep problem -- deeper than we thought at the time.

TAPE 1/SIDE 2

NILSSON: John McCarthy, probably in the late 1970s, was one of the early people in AI to recognize it as a fundamental problem. He proposed a technique for dealing with it called "circumscription". Since then there have been a whole host of people in AI who have looked at the whole problem of what is called "non-monotonic reasoning". Circumscription is one approach. Ray Reiter, who is now at Toronto, has proposed some techniques called default theories. These techniques now can be incorporated in planning systems to help us deal better with the frame problem. It doesn't really solve it, but [we are] much more sophisticated than we were in 1971. There is a time in science where it's too early to attack a problem. The stuff has to bubble a bit more and you have to be able to have the foundation on which to stand. There are certain arguments that could be made that it was just as well the robotics work stopped in 1971, because you had to do this other stuff first.

ASPRAY: All of the examples you have given of enabling technologies have been from the AI field itself. What about from computational complexity, or from underlying power of machinery to do the reasoning tasks?

NILSSON: From the standpoint of the intellectual work, I think the machinery that was present at the time would have been sufficient. It presented various difficulties, like the computer couldn't live on the robot. We had to communicate by radio. That makes the problem a little more expensive, a little more difficult. It doesn't make it impossible to deal with conceptually. It probably makes it impossible to do any practical applications, but many of us weren't in it for that reason. I don't think there were any new results in computational complexity that bore on our problems. Any of the results that might have come up there were more liable to [lead to] misinterpretation than to be helpful, because they told you what can't be done. A lot of people in computer science would say, "Well, AI's intractable, because the problems are all exponential, or NP-hard." My point of view about that is that that might be be. Posed in that fashion, they may be intractable. [But] people do them. We have to find a way to pose the problem in which whatever the solutions people get are tractable solutions, because, obviously, it works. It has certainly been helpful that the enabling technology on computers has been getting more powerful. Now we can conceive of putting as much computing power on board a robot that in 1970 we had in a big room. So Flakey the robot (son of Shaky at SRI now, a nice little robot) has a Sun microprocessor, hard disk, and everything on board. They still have more powerful computers off board, which is always going to be the case. You'll want to connect to the most powerful thing around, and the most powerful thing around probably isn't portable. Ultimately they will be able to connect it to a Connection Machine, or to a Cray, or something. You're always going to have the temptation to use radio links or something to get more power. But it doesn't make the experiment easier. And of course, vision. The vision people are now almost at the point of producing special-purpose hardware to do a lot of the very early processing. [There is] a lot of very good work [being done] at MIT: Carver Mead's work, for example, in his little company that he's got on analog VLSI, on simulated retinas and so forth. All that is going to be very helpful. I have no doubt that people could say, "Look you guys, it's still too early to work on robots. Wait until the years 2000, or 2010. Look at all this nifty stuff you're going to have. One of the reasons that McCarthy and others even thought about the problems that circumscription is supposed to be an answer to is [because] people [were] experimenting with trying to build planning systems and were running into the frame problem. If we hadn't done that, that problem might have just not even been exposed for some years. So you do what you can. They go hand in hand.

ASPRAY: So, we've essentially gotten to 1971 in this overview of the ARPA story. What happens with ARPA funding in AI from 1971 on through the next few years?

NILSSON: I suppose if you drew a plot of it, it probably never really went down much. I suspect it kind of went up. There may have been little wiggles, but in general the thing was going up. I don't know that I can see any general trends toward it being more applied or less applied over long periods of time.

ASPRAY: Did you continue to get ARPA money?

NILSSON: Yes. I kind of quit ARPA in a huff. I said, "Okay, if you guys aren't going to support a Computer Based Consultant project (which I was the project leader of), and are going to push it all into Command and Control, you don't have Nilsson. Go find somebody else." So I continued doing stuff that I wanted to do. But SRI continued to get ARPA money, probably even more than we had gotten before. We were just doing different kind of stuff.

ASPRAY: What sorts of projects were going on?

NILSSON: Well, they were largely directed at natural language interfaces to databases and vision. There was a large group at SRI doing work on vision, and that probably expanded. I just didn't do any ARPA work. I got started with some ONR work and got interested in this intelligent communicating agents -- several robots -- and started some new things there. If there is a trend [in ARPA funding] that dates all the way from Licklider's early days at ARPA through Ivan Sutherland at least, to the present, there was a change. You have to be very careful looking at these things because I have a very subjective viewpoint. In the early days, it seemed to me that ARPA's main mission was [that] information science is very important; we don't know exactly what should be done, but let's go find the very best people we can in the country and just dump money on them. They'll do something. We can give them broad, general guidelines, but you go find the Minskys, the McCarthys, the Feigenbaums, the Allen Newells, the Simons at leading places. Don't worry too much about whether the money is spread fairly, or equitably, or is Arkansas getting enough, or whatever. You just find leading places, set them up, build a laboratory for them, give them the best equipment, give them money, try to get them to get the best students, and stand back and wait for the goodies. I think goodies came -- timesharing, and LISP happened, and all the things that happened in AI -- expert systems. When Feigenbaum's work on expert systems got done, it wasn't because ARPA had a particular direction: "Let there be expert systems." Feigenbaum was self-directed in what he wanted to do, and as a matter of fact, it wasn't in any military problems. The main emphasis came from chemistry, biology, and medicine, which in those days ARPA let him do. But then, it seemed to me, that as they got more and more successful, ARPA thought more and more that they knew what they wanted. They would put out [demands]: "Buy one each unit of research directed at a following thing." I don't know how successful that has been. To a certain extent, the speech understanding project was a little bit that way, because they got a group together to decide what they wanted, and they put out very specific goals and worked toward those goals. I think that was by and large the reason it was successful. But when you get military people deciding there ought to be command and control, there ought to be this, or ought to be that, what happens is not so much that they get exactly what they want, but that money goes to somebody, and those people, in the course of what they're doing, end up doing some good stuff, which might, in fact, be rather unpredictable from the standpoint of the planner figuring out what it was he thought he was buying. So it sounds self-serving, but I've always tried to argue with ARPA that they did better in the old days, when they knew that the problem was to support leading laboratories, to keep the infrastructure going, and to make sure that the very best people were free to follow their instincts on where things ought to go. That was a very difficult message to get across objectively for someone who was looking for that kind of money. But I think there's been a lot less of that at ARPA over the years. So if there is a trend, it was a trend away from that and more toward specific programs and projects that these people would cook up. Bob Kahn decided, well, we needed more money in AI. You just can't go to Congress and say, "Give me some more money in AI and I'll go dump it on MIT and Stanford and CMU." You have to have some mission. So he cooked up this deal about the Strategic Computing Program. I was involved in helping to plan some of that. You had to cook up these things about specific applications. You were going to have a battle management, a pilot's associate; you were going to have a this; you were going to have a that. I think for a lot of people in the field, from their viewpoint, this was just an exercise you had to go through to get the budget up. There would be some trickle-down. We would have to go do this, and somebody is going to have to go build a vehicle and have it run around on roads and things -- Martin Marietta or somebody. But there would be some trickle-down to the universities that would be doing some good work as a result of all that. I think we've been a little frustrated. I don't know that the trickle-down has been all that great. I think that ARPA still thinks in terms of some particular program name. You've got to sell it on the basis of some product that you're going to deliver.

ASPRAY: It may just be the fact of life in government funding.

NILSSON: I think it probably is. Maybe the anomaly was that Licklider and Ivan Sutherland got away with the budget that they had, that they could just go dump on. A similar thing happened at Xerox for a little while. When Xerox PARC got set up they apparently had some money so that they could go and get the very best people in the world, bring them to Xerox, have them do whatever they wanted. Now Xerox didn't capitalize on that, but a lot of good stuff came out of that. They didn't get very much direction from the top at Xerox. They hired Danny Bobrow, Alan Kay, and a whole bunch of people, and they just had a ball. Out of that, ultimately, came the MacIntosh, all kinds of stuff -- a whole industry. Xerox didn't capitalize on it -- too bad for Xerox, but at least good stuff happened. But that was probably an anomaly also. It sort of ended when Xerox tried to direct it more. The more they directed it the weaker the lab became, and finally Bob Taylor took his group and left Xerox, essentially over how much direction he was getting from the central Xerox headquarters and he set up at DEC. Maybe the point is that these little periods where you do have that much freedom are the quirks, and [you should] take advantage of them when you can; and the rest of the time you just have to see what you can do with the system.

ASPRAY: If you were to try to characterize in a list the contributions in AI that came out of ARPA-sponsored work up until 1971, 1972, what sorts of things would you list?

NILSSON: I guess I would want to refer to some notes to make sure I wasn't missing some stuff. There was a good amount of work in early stage robotics -- the so-called hand-eye work, both at MIT and at Stanford. Many of the dissertations, for example, that came out of MIT at that stage -- people like Terry Winograd, his SHURDLU program, and Drew McDermott, Sussmann and some planning work -- that was all supported by ARPA. Similarly there were things going on at Stanford. Raj Reddy got a dissertation in the late 1960s here for some speech understanding work that he did. There were students of John McCarthy and Feigenbaum. Cordell Green was a student of mine. All that was supported by ARPA. These were more intellectual than practical. Some of them later led to the foundations of expert systems and speech understanding work. But that foundational work was important and had to be done, and wouldn't have been done without ARPA's support. So I think the thing to make sure of is [that] you look at what the intellectual and the practical accomplishments were at the major AI laboratories in those days -- SRI, CMU, MIT and Stanford. I'll say that 90% of that stuff that came out of those laboratories was supported by ARPA. One thing to have available is that each of those laboratories was asked to write a history of the labs over the years for the Comtex Technical Note Series. I think these have all been published in the *AI Magazine*, or at least we can get them for you. I wrote one which was the history of the SRI Lab up until about 1980, when we delivered all of our SRI technical notes to Comtex for their microfilm.

ASPRAY: We have copies of the Comtex materials in our library, but I haven't looked for these histories -- in fact, I didn't know about them.

NILSSON: No, there are these histories. I wrote one, and I think Pat Winston or maybe Pat Winston and Minsky did one for MIT. There was one for CMU, probably by Raj Reddy, and Stanford. Bruce Buchanan [might have done] it for Stanford. Those histories would be very worthwhile, because so much of it is ARPA. Some of the people talked a little bit about the difficulties and joys of working with ARPA in their histories. Some of them mainly stressed the intellectual history and the technical history. Those would be very good sources.

END OF INTERVIEW