An interview with Terry A. Winograd

Vendelø, Morten Thanning

Document Version
Final published version

Publication date:
2002

License
CC BY-NC-ND

Citation for published version (APA):
Vendelø, M. T. (2002). *An interview with Terry A. Winograd.*

Link to publication in CBS Research Portal

An Interview with Terry A. Winograd

Morten Thanning Vendelø[1]

Copenhagen Business School

Department of Informatics

Howitzvej 60, 4.

DK-2000 Frederiksberg

Denmark

Phone: +45 38 15 24 08

Fax: +45 38 15 24 01

e-mail: **mtv@cbs.dk**

http://www.cbs.dk/staff/mtv/

An Interview with Terry A. Winograd

I. Introductory Note

Terry A. Winograd (born 1946) began his academic career within the field of artificial intelligence. His early research on natural language understanding by computers was a milestone in artificial intelligence. Later he moved to the field of human computer interaction, and within this field he has done extensive research and writing on design of human-computer interaction. Foremost, focusing on the theoretical background and conceptual models for human-computer interaction design.

Terry Winograd received his B.A. in Mathematics from The Colorado College in 1966. He studied Linguistics at University College, London in 1966-1967, and earned his Ph.D. in Applied Mathematics at Massachusetts Institute of Technology (MIT) in 1970. From 1970 he was an instructor and assistant professor of Electrical Engineering at MIT, before coming to Stanford University in 1973, where he is now a professor of Computer Science. At Stanford, he directs the Project on People, Computers, and Design, and the teaching and research program on Human-Computer Interaction Design. He is one of the principal investigators in the Stanford Digital Libraries project, and the Interactive Workspaces Project. He is also a consultant to Interval Research Corporation, and serves on a number of journal editorial boards, including the Journal of Human Computer Interaction, Computer Supported Cooperative Work, and Personal Technologies.

Prof. Winograd is a longtime advocate for socially responsible computing and is a founding member and past national president (1987-1990) of the Computer Professionals for Social Responsibility, and is on the National Advisory Board of the Association for Software Design. His publications include 'Understanding Natural Language' (1972), 'Understanding Computers and Cognition: A New Foundation for Design' (with Fernando Flores, 1986)[2], and 'Bringing Design to Software' (with John Bennett, Laura De Young, and Bradley Hartfield, 1996, which brings together the perspectives of a number of leading proponents of software design.

II. Terry Winograd and Artificial Intelligence

Prof. Winograd's initial breakthrough in the field of artificial intelligence was with the SHRDLU program, which he wrote at the M.I.T. Artificial Intelligence Laboratory in 1968-1970. SHRDLU is described in his dissertation, issued as MIT AI technical report 235, February 1971 with the title: Procedures as a Representation for data in a Computer Program for Understanding Natural Language. It was published as a full issue of the Journal of Cognitive Psychology, vol. 3, no. 1 (1972), Understanding Natural Language (Academic Press 1972).

SHRDLU understands natural language. It carries on a simple dialog (via teletype) with a user, about a small world of objects (the Blocks World) shown on a display screen. SHRDLU has complete knowledge of the internally represented block world consisting of colored cubes, pyramids and boxes on a flat surface. The program simulates a robot, it accepts commands in English with regard to the block world, carries out the command, and explains how it did it and why certain actions were performed. In addition, it has the ability to learn about new tasks.

SHRDLU uses important ideas about human syntactic semantic and problem solving activities and about their interactions in understanding natural language discourse. Understanding of English requires an integrated study of syntax semantics and inference. Winograd felt that the best way to experiment with complex models of language was to write a program, which can actually understand language within some domain. In this case with a robot which has a hand and eye and the ability to manipulate toy blocks. The program attempts;

1) to be a usable language understanding system.
2) to gain a better understanding of what language is and how to put it together.
3) to understand what intelligence is and how it can be put into computers.

---

[2] In this book they take a critical look at work in artificial intelligence and suggest new directions for the design of computer systems and their integration into human activity.

SHRDLU is head and shoulders above contemporary systems when it comes to intelligent conversation. Although its domain of discourse is restricted to a tabletop world of colored objects SHRDLU really understands this world in terms of the relation between semantics and the physical properties of the blocks and the tabletop. It consists of subsystems that parse interpret and construct sentences, carry out dictionary searches and semantic analyses and makes logical deductions.

SHRDLU uses Halliday's systemic grammar, which emphasizes the limited and highly structured choices made in producing syntactic structure abstracting the features that are important for conveying meaning. The parser is special and interprets the recognition grammars. Meaning is covered by the development of a formalism for concepts within a language user's model of the world representing objects events and relationships. Semantics is represented by a system, which is developed to work in conjunction with the parser, a dictionary and the problem solving programs. It considers not only meaning, but also context.

III. Purpose of Interview

The present interview is one of a number of interviews with people who are or have been involved in AI research. The purpose of these interviews is to learn about their views of the AI-field and the work of other AI scientists.
As it appears from above Prof. Winograd was heavily involved in AI research earlier in his scientific career, but he decided to leave the field of AI. Elsewhere he describes his conversations with Hubert L. Dreyfus as influential to his own "complete shift of research direction, away from artificial intelligence towards a phenomenologically informed perspective on human-computer interaction" (Winograd, 2000). Using this as the point of departure it was assumed that he would provide interesting answers to questions such:

- What attracted him to AI in the first place?
- Being a "second generation" or may be even "third generation" AI-scientist, how does he think that his view on AI differs from the view held by "first generation" AI-scientists?
- What is he view on the work by other AI scientists?

- Why did he choose to leave the field of AI?

And thereby, contribute to our understanding of the history of AI.

V. References

Winograd, T. (2000) Foreword. In: M. Wrathall and J. Malpas (eds.) Heidegger, Coping and Cognitive Science – Essays in the honor of Hubert L. Dreyfus – volume 2. Cambridge, MA: MIT Press.
(Available at: http://www.idiom.com/~gdreyfus/70Celebration/Foreword2.html)

VI. The Interview[3]

The following conversation took place between Professor Terry A. Winograd (TW) and Morten Thanning Vendelø (MV) on February 29. 2000, in the Department of Computer Science, Stanford University, California.

MV: Could you tell a little about your educational background. I know that you got a BA in Mathematics from The Colorado College, then went to University College in London and studied linguistics there, and then went to MIT for your Ph.D. but is there a storyline in choosing that path for your education?

TW: The storyline is that I went to a liberal arts college that had no engineering and very little mathematics. So my mathematics major didn't mean that I did a lot of mathematics. I was interested in language, and they didn't have a linguistics course, but I convinced the anthropology teacher to give me a readings course in linguistics. Also, I became acquainted with computers, not through coursework but outside class during my last couple of years in college. In addition, I saw some writings by Marvin Minsky describing artificial intelligence and thought it was pretty fascinating. So I applied to MIT because of these computer interests, and I also applied for a Fulbright Grant to study linguistics abroad. That was basically opportunistic. People said: "It is wonderful experience to go to Europe for a year on a Fulbright, you should do it." So the driving force was going to Europe, not the specific content. At that time my belief (which I now realize ironically was false) was that there was no interesting computing work going on outside the United States. If I had known what was going on in computing in England I might have done something different. The only language I had studied in college was Russian, and at that time going to Russia on a Fulbright Grant was not an opportunity. So I was limited to English speaking countries and since I was interested in language I decided to do a year of study in linguistics. So I simply went through the catalogues of the British universities looking for what seemed to be the most interesting linguistics program from the point of view of what I wanted to do. I found the program at University College, London, with a

[3] This transcription of the interview has been with Terry Winograd for review.

professor Halliday and I applied to it although I did not have a very coherent intellectual plan. I came from Colorado, a small state, and since Fulbright had a quota by state I had an advantage in getting the grant. So that's how I got the Fulbright and studied linguistics for a year in London. If I had not got the grant, I would have gone directly to MIT and not done the linguistics work. Because once you were at MIT, you couldn't do computer science and linguistics at the same time.

When I would meet Chomsky's students at a party and say: "I work in the AI Lab." they would turn around and walk away. That was it. There was no communication, but active hostility between Minsky and Chomsky. So when I came back from the year in London to start at MIT I already had a lot of interest in language and a year of background in linguistics. That is how I ended up doing a language project.

MV: You said that you became interested in computers. What was it about computers that made you interested in them and in AI?

TW: When I was a junior in college a professor in medicine who had been at a larger medical center ended up moving to this small hospital in Colorado Springs. He had previously hired programmers to work on his research. He had a very early computer, which he used for doing calculations on radiation therapy. So he sent a note over to the math department saying: "Do you have any math students that can help me with my computer?" So my first exposure to computers was a personal computer. There was a room and there was this desk-sized thing (a CDC 160) and I was the only person using the computer. I had a great time so that was how I got into computing.

MV: But how did you become interested in AI?

TW: The question was, which field was combining computers and language, because they were the two things I was interested in.

MV: You said that at MIT there were no connections between computer science and linguistics?

TW: Minsky and Chomsky were both very strong personalities. You could go to one or the other but you couldn't be in the middle. That is how MIT is, it has strong boundaries between its departments.

MV: But you were allowed to do computers and linguistics?

TW: Minsky wanted me to do linguistics and prove that his students could do better linguistics than Chomsky's. I was not the first. Dave Waltz, who eventually went on to do work in vision, was doing a research project on linguistics when I got there, as was Gene Charniak. And in the prior generation Daniel Bobrow and Bertram Raphael had both done language projects. So it was a major stream of work in the AI Lab.

MV: What was it that made you make this connection between AI and language, why did you choose this connection between computers and language?

TW: I had the question: "How could you make computers use language?" And by definition the answer is AI, because computers don't do language by themselves. Fortran programs didn't do language, except for Joe Weizenbaum's program, but that was not a standard Fortran program. My direction wasn't because of any philosophical commitment to artificial intelligence. It was because of an interest in understanding how language worked, using computers as the tool.

MV: How would you describe the artificial intelligence environment at MIT at that point in time? If you talk to some people who went to Carnegie Tech in the late 1950s and 1960s they describe it as an intellectual supernova?

TW: It was a very exciting time because people were getting their hands on machines that previously either were available only for serious and highly specialized scientific work. We could try things out that nobody had ever tried before. So people came up with all sorts of interesting things, some of which turned out to be very important and some of which didn't. The excitement is like being in a gold rush or exploring new territory. There are all these nuggets lying on the

ground and you can pick them up to see what you can do with them. Also, there was a very strong sense of that we were building the future.

I would not call it an intellectual supernova in that there was no emphasis on deep intellectual thought. The quest was to build stuff and see what it did, and then build more stuff and make new things happen. But it wasn't intellectually grounded. Carnegie was much more grounded in Simon's and Newell's theoretical interests and AI was driven by their theory. Whereas at MIT it was more a hacking approach if you can call it that. Those people highly respected in the lab were not cognitive researchers, but virtuoso programmers -- hackers.

MV: So this linking at Carnegie between economics, AI and cognitive science did not exist at all at MIT?

TW: No, it was looked down upon as a kind of waste of effort. "We are building wonderful new stuff, why worry about how people think about economics?"

MV: You completed your Ph.D. at MIT and then you went to Stanford. How would you describe the environment at Stanford compared to MIT?

TW: They were pretty similar. John McCarthy had developed AI Lab here. It was a bit more independent because it had its own building up in the hills -- it wasn't even sharing a building with the rest of computer science. It was funded in the same way as the MIT Lab, which is by large umbrella programs. This meant that individual projects had a lot of latitude to do whatever they felt like and research didn't have to be justified on a project-by-project basis. There were a lot of machines and a lot of work on robotics. As for its emphasis, it was pretty much a sister environment to MIT, and thereby distinct from Carnegie, with its more cognitively grounded environment. Stanford put more emphasis on using formal logic because John McCarthy likes formal logic and Marvin Minsky didn't.

MV: Did you have any contact with Ed Feigenbaum about AI when you arrived here?

TW: It depends on what you mean by contact. He was not in the AI lab. I was up on the hill with McCarthy's lab and Feigenbaum was somewhere on campus with

his expert system projects. He was there when they hired me, and he was on the faulty committee and so on, but in terms of day-to-day contact we had very little.

MV: Was it because you had these two different views on AI?

TW: Again it was two different camps. There was Feigenbaum students and McCarthy students, and Feigenbaum faculty and McCarthy faculty. And you really didn't cross those boundaries. It wasn't as hostile as between Minsky and Chomsky, but it was still a very clear divide.

MV: If we look at these two different ways of doing AI, and were to do an evaluation of them today, How would you evaluate them in terms of their progress / contribution to AI?

TW: It is good question. I think that Feigenbaum was much more eager and willing to say: "I want practical and commercial applications and I don't care how deep or interesting the theory is." If you look at his so-called theoretical principles, they are very shallow. The focus was on practical use. If you look at the fifth generation work, it is fairly mundane from an intellectual point of view, but they actually tried to make it commercially relevant. On the other hand, McCarthy is really a pure mathematician. His interests have nothing to do with making money or applying AI. They have to do with coming up with deeper theories. I happen not to agree with McCarthy's theoretical leanings, but I think that it is good that he had them and was trying to drive the program from a conceptual point of view, instead of a kind of opportunistic way.

MV: You said that you didn't interact with the linguists at MIT. But did the work by, for instance, Noam Chomsky inspire you?

TW: It did indirectly, but my direct source of inspiration came from the year in London where I studied with Halliday. If you look at what Halliday was doing in the context of the larger picture of generative grammar, he was trying to adapt his theories, which came more from a social perspective, to a generative form. In my opinion, the merger never worked. If you look at the subsequent work of

Halliday in linguistics, the attempts to do generative grammar have been much less successful than the return to systemic grammar's roots as a socio-linguistic analysis. But certainly nevertheless the training was important, I took a course in transformational grammar during that year, so I had learned it. Chomsky's basic insight was that synchronistic language can be described fairly well with a generative rule-based system. That was at the heart not only of his work but also of the work by anybody who was trying to use computers, because if you cannot put language into a rule-based system, then you cannot program it. So in that sense I think it was very much along the lines of Chomsky. The differences are at the next level where Chomsky posited a structure of transformations, which was not computationally implemented while those in AI were focused on finding appropriate algorithms. So I was on the algorithm side and not the formal grammar side.

MV: Returning to your time at MIT. Would you say that your view of AI is similar to that of Marvin Minsky's? How would you position yourself in relation to him?

TW: We differ in a couple of major ways. One is his basic faith that the intelligence embodied in organisms, for instance people, is very similar in nature to the programs we write in AI. Not necessarily in detail but in basic nature. I disagree with that fundamental assumption about symbolic processing, which I critiqued in the book I wrote with Fernando Flores. The other key differences are Minsky's view on the role of human values and the role of machines. If you ask: "What is Minsky's religion?" Then it is some abstract notion of the progress of intelligence. The values he was pursuing had to do with a drive towards higher intelligence as a value in itself. I have always been a much more political, social oriented kind of person for whom other human values take priority.

MV: In my first e-mail I mentioned that you might be what we can call a second or may be even a third generation AI scientist. How would describe yourself in contrast to the first or may be the second generation?

TW: People like Minsky and Newell and so on were the immigrants. None of them started their intellectual life in artificial intelligence, as it didn't exist. So they

went in and made it happen. Then the second generation, which I would include myself in, followed in their labs. At MIT there were Tom Evans, Daniel Bobrow, Bertram Raphael and Adolfo Guzman and so on, and then in a later round there were Gerald Sussman, Carl Hewitt, Dave Waltz, me, Eugene Charniak, Pat Winston, and others. But there wasn't any major transition between those two. I think it was more gradually increasing machinery and more sophistication, but it was pretty much the same spirit. We did not have to develop the context. The context was there, the machines were there, LISP was there and so on. So we could just take a problem and apply those things.

MV: And then your own idea for your block-moving program SHRDLU, how did it come about?

TW: Basically the block-moving problem was chosen because I wanted to do something more concrete. Bill Woods at Harvard was using airline reservation systems as a language domain, but I didn't find them very interesting. Minsky had the idea that I should do something about stories for children, because he thought they were much more simple than stories for adults. I was not so sure about that and wanted to do something else. There were others at the lab building robot hand-eye systems that actually moved blocks around on a tabletop. So I don't know if it was me or Minsky, but we arrived with the idea of doing a system that conversed about block moving. But there was never any real robot arm connected to my work.

MV: It was common for AI scientists from Carnegie Tech to go to RAND and many people perceive RAND as very important to the development of AI. Did you ever go to RAND?

TW: No and I cannot remember anyone from RAND coming to MIT, and I cannot remember that anyone from MIT went to RAND. I remember it as a pure Carnegie connection.

MV: In the middle of the 1980s you wrote and published a book with Fernando Flores, where you articulate a more critical view of AI. How did you meet Fernando Flores?

TW: I went to a meeting where I met a Chilean scientist, Francisco Varela, who said: "How is Fernando Flores?" I said: "Fernando who?" And he said: "Fernando Flores. He is at Stanford." And I said: "Oh well Stanford is a big place and I don't know where he is." And he said: "No, he is in the Computer Science Department at Stanford," and I said: "No, this cannot be, I am in the Computer Science Department I go to all the faculty meetings and I have never heard of Fernando Flores." He said: "I know Fernando Flores is in the Computer Science Department." So when I got back I looked him up, and it turned out that he was indeed there. After being in the government of Salvador Allende that was overthrown by a coup, Fernando had been imprisoned for three years. The San Francisco Chapter of Amnesty International took up him as a "prisoner of conscience," and one of the conditions for his release that the Chilean government put up was that he had to have a job waiting for him outside of the country. Two professors from our department, Bob Floyd and George Dantzig, had a position they could support from various grants that were close enough to what he had done before going into government that they could justify hiring him. So without having a specific demand on his work they created a research associate position for him at Stanford so that he could have a job, so that they could get him out of prison. But he had just been in prison in Chile for three years and had been flown to California, where he was getting oriented and reconnecting with his family, so he was not spending his time in the Computer Science Department.

I looked him up, and he is a tremendously intellectual guy, probably the most intellectual, in some deep sense, that I have ever known. He is also practical, but he is the kind of thinker who can read six philosophy books before breakfast. He has this incredible mind always thinking and always looking for more, so he wanted to find out who at the department were interested in talking to him. We started talking in a casual way, then he handed me a book on philosophy of science and said "You should read this." I read it, and we started talking about it,

and we decided to write a paper about it that turned into a monograph that turned into a book. It was a gradual process of finding him interesting, and finding the stuff we were talking about intellectually stimulating. He only kept his job at Stanford for a short time, and then went to Berkeley as a Ph.D.-student, because he wanted to have a Ph.D. In Chile he had gone into politics before he had finished his Ph.D. I was officially on his committee there, but it was never a question who was leading and who was following.

MV: You are also one of the founders of the organization called: Computer Professionals for Social Responsibility. Did it have any connection with your decision to leave the field of AI?

TW: No, the two issues are really separate. My interest in CPSR grew out of nuclear war issues. In those years any organization that had social responsibility in its name was basically trying to prevent nuclear holocaust. The American government seemed to be headed down this track, and every group said: How can we help? First there was Physicians for Social Responsibility, then Educators for Social Responsibility, Architects for Social Responsibility, Psychologists for Social Responsibility, and so on. It was clear that there were a lot of computing issues involved in warfare and when the Star Wars program (SDI) came along it was very clear that it was based on assumptions about computing that were not valid. We picked up on this as the key issue. The only place where it touched on the areas of AI was to the degree that military funding and military applications were based on practical applications of AI. But as to whether the theory of AI in general was right or wrong, a good idea or a bad idea, was just never in the purview of that kind of politically oriented activity. CPSR was not philosophically oriented but politically oriented.

MV: And then after AI you moved on to HCI, what let you in that direction?

TW: I think that I was always doing it. I just didn't realize it. When I was writing systems to use language there were two somewhat independent motivations. One is to model and understand human language, and the other is to make computers easier to use. What I realized was that making computers easier to use, was not

the same as making them use ordinary language. There are many other issues involved in how to get people and computers to interact well. For a lot of the cognitive studies, you can turn them around. Instead of thinking of how to model people with the computer you use cognitive analyses of people, to better fit the computer to them. – how to interface with them, rather than how to duplicate them. Much of what we have learned in AI is very relevant to HCI.

There was a period of about five years where I had no label for my work. Because I was out of AI and I was doing a sort of philosophical writing, but I wasn't a philosopher. I realized that really HCI was more compatible if I wanted to be in the university environment. The kind of cognitive-philosophical reflection that I was doing with Flores was not a viable program for graduate students in a Computer Science Department. There needed to be something much more concrete that the concepts could be applied to. If you read the book with Flores you see a kind of open promissory note at the end, which says all of this theory should guide you in the design of systems. So that is really how I got off into Human Computer Interaction -- asking how to make some of these more philosophical considerations guide us in the design of systems that work with people.

MV: Another issue related to this is that a lot of people tend to view AI as the ultimate science in the sense that the big question we have here is: What is the nature of human intelligence? And therefore, it is one of the biggest questions that you can work on as a scientist, and most else come second. Is this an opinion that you share?

TW: I think that my sense of it, and this is in a broad sense my opinion for scientific research, is that for a scientific topic to be good to work on there have to be at least three things that come together. One is resources -- funding. One is that it is an interesting problem, and the third is that you are at the right point in intellectual and technical history to have some leverage on the problem. Scientists a hundred or two hundred years ago could have thought that it was a fascinating problem to understand how the brain works, and they might have built models of it. But they wouldn't have gotten very far because they didn't have the background knowledge. My sense of AI is that we are not there yet. If I

were a young person going into science today, I might very well go into biology or neuroscience, because I think there are all sorts of things that we yet don't know about real brains and nervous systems. We need to develop better fundamental principles of how information systems and biologic information systems work, without trying to pretend that they are what we have in silicon, just sloppier. After another 10, 20 or 50 years of great research in that area it will be the moment when somebody can say: "Ah, we can put this together." My own hunch about where we are is that we are still missing the basic knowledge that will make AI a productive science.

MV: So you don't share these more optimistic views on AI?

TW: Well, what I just said is that it will happen. I just think that we are much further away than many AI proponents think we are. I certainly don't believe that the basic science that we need is there, and that it is just a matter of putting together the mechanisms. It is going to take quite some time before we get there, and it is going to come out of the biologic science and not out of more computer development at this point.

MV: So more work is needed, but in different areas than many AI-scientists believe. Do you share the reductionistic view of AI that some AI-scientists pursue?

TW: There are two different levels. I share the physicalist view. If you say: "Is there anything going on in my mind which cannot be explained by the motion of molecules in my brain?" I don't believe it. I believe that thinking is purely a mechanical process. If by reductionistic you mean, when I am thinking about a dog, will that be correspond to activating some kind of computer-like symbol for the concept of dog, then I don't believe it. So it is a physical reductionism but not a symbolic reductionism, if you want to distinguish among those two.

MV: Given that the needed biological knowledge is in place, how far do you think we can go with artificial intelligence?

TW: Starting from scratch and building up a mind will not be the way it goes. It will be more likely to incorporate hybrids that combine real biological systems with some amount of artificial stuff. We will be able to engineer changes to real biological systems genetically or chemically. So the question: "Could you build intelligence?" will not be the interesting question. It will be more a question of extending and adapting intelligence. I don't think that there are ultimate limits in the sense that we can never achieve a certain kind of intelligence. I don't know where they are, and I have no particular reason to believe that they are not hundred or thousand years away. But in principle the brain is a mechanical system, just a very complex one.

MV: If you look 20 years or so ahead, do you then have any hopes for artificial intelligence. In terms of what would be a good outcome of the work being done in the field?

TW: I think that what will happen in twenty years is a continuation of what we see now. Consider speech understanding. When I was a graduate student there were heated arguments as to whether a computer needed full logical understanding to understand speech. It is easy to come up with words that sound alike and whole phrases that sound alike but mean different things, and you have to understand this. What has happened is, that as computers were able to process more and more data and compute higher level statistics and more analysis, we have got programs for speech dictation that achieve accuracies up to 90% without any logical understanding. You have programs listening to you, which don't have any logical understanding of what you are saying and they do it pretty well. I think that we are going to see those boundaries pushed more in various applications such as visual recognition. It is not necessarily going to be because of new insights into how people do these tasks, which would be scientifically interesting. Instead it is going to be because you can just throw enough processing power at it to do an acceptable job. So I think that a lot of things which people in the early days though of as proofs of deep AI, including chess as an the obvious example, will just be handled by not deep AI but by a lot of processing power. We have almost seen success at chess that way. Computers haven't quite been beating Kasparov consistently yet, but they are doing well. AI

success will be in a number of applications, none of which will have the flavor that the original AI people though about, which is: "This machine is thinking like a person." Instead, it will use extensive computation and incorporate statistical methods. When I was a student in the AI Lab nobody there even studied statistics, nobody even mentioned statistics. Today non-symbolic techniques are critical in most branches of AI. As for any major imminent breakthroughs in AI, I don't see them, but usually you don't see a breakthrough before it happens.

MV: OK, this was my last question, so thank you for your time.

VII. Selected works by Terry A. Winograd

VII.i Books

Winograd, T. with J. Bennett, L. De Young and B. Hartfield (eds.) (1996) <u>Bringing Design to Software</u>. Menlo Park, CA: Addison-Wesley

Adler, P., and Winograd, T. (eds.) (1992) <u>Usability: Turning Technologies into Tools</u>. New York, NY: Oxford University Press.

Friedman, B., and Winograd, T. (eds.) (1989) <u>Computing and Social Responsibility: A Collection of Course Syllabi</u>. Palo Alto, CA: Computer Professionals for Social Responsibility.

Winograd, T., and Flores, F. (1986) <u>Understanding Computers and Cognition – A New Foundation for Design</u>. Menlo Park, CA: Addison-Wesley.

Winograd, T. (1972) <u>Understanding Natural Language</u>. New York, NY: Academic Press.

VII.ii Journal Articles

Oviatt, S. L., Cohen, P. R., Wu, L.,Vergo, J., Duncan, L., Suhm, B., Bers, J., Holzman, T., Winograd, T., Landay, J., Larson, J., and Ferro, D. (2000) Designing the User Interface for Multimodal Speech and Gesture Applications: State-of-the-Art Systems and Research Directions. Human Computer Interaction, vol. 15, no. 4, pp. 263-322.

Winograd, T. (1995) Environments for Software Design. Communications of the ACM, vol. 38, no. 6, pp. 65-74.

Winograd, T. (1994) Designing the Designer. Human-Computer Interaction, vol. 9, no. 1, pp. 128-132.

Flores, F., Graves, M., Hartfield, B., and Winograd, T. (1988) Computer Systems and the Design of Organizational Interaction. ACM Transactions on Office Information Systems, vol. 6, no. 2, pp. 153-172.

Winograd, T., and Flores, F. (1987) Response to Reviews of Understanding Computers and Cognition. Artificial Intelligence, vol. 31, pp. 250-261.

Winograd, T. (1980) Extended Inference Modes in Reasoning by Computer Systems. Artificial Intelligence, vol. 13, no. 1, pp. 5- 26.

Winograd, T. (1980) What Does It Mean to Understand Language? Cognitive Science, vol. 4, no. 3, pp. 209-242.

VII.iii Book Chapters, Reviews, and Popular Articles

Winograd, T. (1995) Heidegger and the Design of Computer Systems. In: A. Feenberg and A. Hannay (eds.) Technology and the Politics of Knowledge. Bloomington, IN: Indiana University Press, pp. 108-127.

Winograd, T. (1994) The Norbert Wiener Award for Social and Professional Responsibility. Cybernetica, vol. 37, no. 3/4, pp. 387-392.

Winograd, T. (1991) Thinking Machines: Can There Be? Are We? In: J. Sheehan and M. Sosna (eds.) The Boundaries of Humanity: Humans, Animals, Machines. Berkeley, CA: University of California Press, pp. 198-223.

Davis, R. (ed.), Dreyfus, S., and Winograd, T. (1989) Expert Systems: How Far Can They Go? AI Magazine, Spring, pp. 61-67.

Winograd, T. (1984) Some Thoughts on Military Funding. CPSR Newsletter, vol. 2, no. 2, pp. 1-3.

Winograd, T. (1976) Artificial Intelligence and Language Comprehension. Lead article in a report of the same title by the National Institute of Education, February, pp. 1-26.

Winograd, T. (1976) Computer Memories--A Metaphor for Human Memory. In: C. Cofer (ed.) The Structure of Human Memory. Freeman, pp. 133-161.

Winograd, T. (1974) Artificial Intelligence - When Will Computers Understand People? Psychology Today, vol. 7, no. 12, pp. 73-79.

Winograd, T. (1973) A Process Model of Language Understanding. In: Schank and Colby (eds.) Computer Models of Thought and Language. Freeman, pp. 152-186.

Winograd, T. (1973) Language and the Nature of Intelligence. In: G. J. Dalenoort (ed.) Process Models for Psychology. Rotterdam, NL: Rotterdam University Press, pp. 249-285.