# Oral History of Peter Friedland

Interviewed by:
David Grier

Recorded May 16, 2018
Mountain View, CA

CHM Reference number: X8635.2018

# Table of Contents

# Peter Friedland

## Conducted by Software Industry Special Interest Group

**Abstract:** In this May 2018 interview with David Grier, Peter Friedland discusses his early education and upbringing in Allentown, Pennsylvania. After earning an electrical engineering degree from Princeton University, Friedland went on to earn a PhD from Stanford University, where he participated in the DENDRAL and MOLGEN Projects. Friedland then spent nearly 7 years working as a research scientist at Stanford. During this time, he was one of the co-founders of the early AI and expert systems companies Intelligenetics/IntelliCorp, and Teknowledge. In 1987, Friedland joined NASA Ames Research Center, where he oversaw NASA basic and applied AI research for eight years. The discussion of his tenure at NASA includes descriptions of several notable advances in AI applications, including autonomous robotics systems, unsupervised data classification, PI-in-a-Box, and expert scheduling systems. The remainder of his career includes co-founding the knowledge management company Intraspect, a second stint at NASA, and a short period at FX Palo Alto. Lastly, Friedland describes his work as a consultant for the Air Force Office of Scientific Research, where he continues to advance applied defense research projects and initiatives.

**David Grier:** Today is May 16, 2018. My name is David Alan Grier, and I'm conducting an oral history interview of Peter E. Friedland. This is part of the ongoing oral history program of the Software Industry Special Interest Group, which is part of the Computer History Museum in Mountain View, California. We are here at the Computer History Museum. This interview is being videoed and recorded and will be transcribed, edited, and posted on the Computer Museum History website.

**Early Education and High School**

**Grier:** Are you ready to begin, sir?

**Peter Friedland:** Ready.

**Grier:** July 17, 1952, Brooklyn, New York. Tell us a little about your family and the Brooklyn that you were born into.

**Friedland:** The Brooklyn I was born into, let's see. Well, I lived on Ocean Parkway in Flatbush. This was the Brooklyn of great delicatessens and the Brooklyn Dodgers. The old Trolley Dodgers. My father was a CPA, and my mother a teacher. We moved to Allentown, Pennsylvania when I was six. I grew up in Allentown, although most weekends we drove back to New York to see our relatives.

**Grier:** Okay. You went to high school in Allentown. Was there anything about that high school education that interested you in technology or computers or anything else?

**Friedland:** Well, it turns out Allentown had an extremely enlightened public education program. Starting in third grade, they had special programs for kids who were identified as being good in school, in math, science, English, and the like, called Opportunity School that lasted through sixth grade. There were honor classes in junior high and high school. The public-school education in Allentown was really excellent.

Also, Allentown was heavily involved in activities like science fairs. One reason was because the publisher of the *Allentown Morning Call*, the newspaper, was on the national board of Science Service, which was running things like the Westinghouse Science Talent Search at the time. The science fair in Allentown, the Lehigh Valley Science Fair, was a really big event.

**Grier:** Did you participate?

**Friedland:** Yes, I got first place in chemistry in my junior and senior years and was reserve champion in my senior year, beaten out by another high school student who built a proton accelerator in high school. It was hard to compete with that with just some chemistry programs.

But to judge the quality of the Allentown high schools, three kids from Allentown, in my senior year, including myself, were among the 40 winners/finalists of the Westinghouse Science Talent Search. Our town went crazy over it and had a testimonial dinner for the three of us, among other things. So, it was a little bit crazy. It actually resulted that our local congressman and Mack Trucks flew the three of us to the Apollo 13 launch, and we sat in the VIP visitor stands next to John Stennis and Spiro Agnew.

**Grier:** Wow.

**Friedland:** And Wernher Von Braun. We actually saw the Apollo 13 launch.

**Grier:** May 1972?

**Friedland:** It must have been May 1970, and watching a Saturn V launch was really amazing from close up.

**Grier:** I imagine.

  

**Friedland:** You felt it more than heard it. Anyway, our local school district was, as I said, extremely good. We actually were second only to Bronx High School of Science that year in winners of the Westinghouse Science Talent Search.

## Early Interest in Chemistry and Computing

**Grier:** What led you to chemistry?

**Friedland:** A couple things. I'd always been interested in sciences, built electronic kits like Heathkits. I don't know if you're old enough to remember Heathkits, but in fact, I watched the Apollo moon landing, Apollo 11, on a Heathkit color TV I had built myself.

I'd always been interested in that, and I really got serious about chemistry when I went to an NSF [National Science Foundation] summer science program. It's unfortunate there aren't as many of these now. Back then, NSF sponsored a whole bunch of summer high school science enrichment programs.

I went to one at the Mount Hermon school in New Hampshire, where college professors taught chemistry. In fact, I got interested in certain types of inorganic chemistry at that program. Also, it was actually my first experience with timesharing computing because they had a terminal for the Dartmouth Basic System, the first real practical timesharing system. This was the summer of 1967.

When I came back from that, I also was very fortunate that I got involved working with a college professor at Muhlenberg College in Allentown. Last name was Staley. I forget his first name now, but he was  my mentor in chemistry. That's actually where I did my project for the Westinghouse Science Talent Search.

So, there was sort of a dual thing with interest in computers and chemistry. My high school actually had an IBM 1620, which you have here downstairs in the Computer Museum. That was my first real computer. There was a computer club at my high school that made a lot of money by running a computer matching dance.

**Grier:** At Dartmouth, you obviously would be learning to program in Basic.

**Friedland:** Well, I was at Mount Herman school learning how to program in Basic.

**Grier:** In Basic. On the 1620?

**Friedland:** Fortran. Well, before we got a Fortran compiler and a disk, we actually were programming in SPS, Symbolic Programming System. It was before we got one of the hard disks, which you also have downstairs here. You put your program into the punch card reader, first into a pre-compiler to check for errors and then into a compiler. You actually got back a deck of object code, which you put back into the computer. Then in my senior year, we got a disk, a big disk drive, so you no longer had to reenter your object code into the computer. That was my first serious exposure to computers and also to serious chemistry.

## Summer Job as a Scientific Programmer for Air Products and Chemical Corporation

**Grier:** On your resume you list working as a Scientific Programmer at Air Products and Chemical Corporation. Was that through working with your mentor?

**Friedland:** Well, no, the mentor was at Muhlenberg College, during the school time. I had a summer job for two years as a programmer at Air Products and Chemicals. People get the wrong impression of Allentown from the Billy Joel song, which is really about Bethlehem, which was a steel town, not Allentown. Allentown was a much more mixed industry town. There was Bell Labs, and Air Products was one of the biggest employers. I got a summer job at Air Products as a programmer for two summers.

**Grier:** What sort of programs did you write for them?

**Friedland:** It was working mostly on programs that modeled the condition of various gases at very cold temperatures. There's a thing I still remember called the virial equation of state that allows you to model pressure and temperature, etc., as gases are compressed into liquid form. I rewrote some of those programs and modified some of them, and just did general programming work mostly in Fortran.

**Grier:** Mostly in Fortran. Closely supervised, or did they just sort of say, "Here's the problem, do it"?

**Friedland:** A combination of both. Certainly, closely supervised originally. Then when I turned out to be a good programmer, less so. But I worked closely with the head of the department. I think his name was Steve Israel. Anyway, he was a very good mentor for learning more advanced programming.

## Undergraduate Work at Princeton

**Grier:** In September 1970, you go to Princeton.

**Friedland:** Yes.

**Grier:** You applied, and your Westinghouse scholarship helped draw attention to you.

  

**Friedland:** Yes, I think it pretty much guaranteed entry. Being one of the 40 winners of the Westinghouse Science Talent Search that year pretty much guaranteed you got into whatever college you wanted.

Three kids from my high school, not the same ones as the Westinghouse winners, got into Harvard and MIT [Massachusetts Institute of Technology] that year. We all turned down Harvard, which made the local Harvard rep almost die. He never had three kids get into Harvard from Allentown before, and we all turned Harvard down. Two went to MIT. Two of my friends went to MIT, and I went to Princeton.

One of the things that attracted me to Princeton was I also wanted to, as an undergraduate, be exposed to not just science and engineering, but all forms of education. Actually, during my high school days, I'd also been a pretty good writer. I had an article in the *Atlantic Monthly*. There was an *Atlantic Monthly* writing contest.

**Grier:** What was the article on?

**Friedland:** It was an essay on the book *Too Late the Phalarope* by the South African writer. What was his name? The most famous South African writer, you must know.

**Grier:** It's *Cry the Beloved Country*.

**Friedland:** Yes.

**Grier:** Alan Paton.

**Friedland:** Paton, right. Yes, very good. You're a historian and writer. You should know these things. I should remember these things.

**Grier:** We all should remember.

**Friedland:** The other thing that really attracted me to Princeton was I'd been reading in a journal that a new professor had come to Princeton named Todd Wipke, who was interested in the mix of computer science and chemistry. I had written him during the summer.

One thing about Princeton that's very nice is that, even though it's obviously one of the great universities, undergraduates really come first. There's a lot more undergraduates. It's a relatively small school, about 4,000 undergrads and about 1,000 grads at the time. Undergraduates, it is not unusual for them to be heavily involved in research programs from day one, as opposed to a place like Stanford, where while undergraduates sometimes do get involved, that's much more of a rarity. You know, it's basically graduate students doing work with professors on research.

I got involved with Todd's project from day one being there. I had a job as a research assistant as opposed to working in the cafeteria or something. The project was a theme that actually continues to today, which was using computer science, using AI, to help design synthetic pathways for chemical reactions. How do you design a reaction, a set of reactions, to produce a given desired molecule, whether it be a drug or any other desired molecule? What's the pathway to doing that? Also, he, jointly with another professor, Bob Langridge, who is a computer graphics chemistry guy, had won an NSF award for a DEC PDP-10 computer. It was a great system to program on. I ended up doing my senior thesis on that.

At Princeton, everyone does a senior major research project, a senior thesis. In fact, probably the most famous senior thesis that came out of Princeton was Wendy Kopp's one that was the plan for Teach for America; that started as a senior thesis at Princeton. My daughter ended up working for them for a while.

But anyway, I was a member of this research team for four years. It was a good introduction both to AI and to computer chemistry.

**Grier:** You very much went to Princeton knowing that chemistry was going to be what you were going to be studying.

**Friedland:** Well, I was interested in both chemistry and computer science. I'd always been interested in the combination of computer science with solving hard problems in science and engineering and medicine.

**Coursework at Princeton**

**Grier:** What did you take in terms of computer science while you were there?

**Friedland:** Well, my advisor —it turned out he was Peter Denning. I don't know if you know Peter.

**Grier:** Yes, I know Peter.

**Friedland:** When I went to Princeton, there was no computer science department. This was in 1970. The computer scientists were in the EE [Electrical Engineering] department, so I took courses there. Peter was my advisor.

  

I took a course from him essentially on operating systems and hardware, on things like paging and all that sort of stuff. There was a course on pattern recognition from I think Theo Pavlidis, which was the closest to an AI course there. An AI course was started, but I forget who taught it in my senior year.

There was a course in numerical analysis. That was the hardest course I ever took. There were five people in it, and one person got an A. Two people got F's out of the five. One got a D+. It was my only C at college. I got a C+, and that was probably the proudest I was ever of a grade. A guy named Forman Acton taught the course. The book was *Computer Methods that Work*. So, that was sort of my first introduction to heuristics. But I was also very lucky.

I think probably my favorite courses at Princeton were several in history and philosophy of science, taught by Thomas Kuhn who of course, wrote about paradigm shifts, and C.C. Gillespie, the edge of objectivity, and Mike Mahoney.

Actually, I think my favorite little bit of discovery at Princeton was a paper I wrote for one of the history and philosophy of science courses. I delved in the library into the *Journal of the British Academy*, and I think I found the earliest connection between electricity and magnetism. It was like a letter to the editor, as it were, in I think the 1400s or 1500s of a ship captain who was asking the learned academy to explain how come when a lightning bolt hit near the ship, his compass went crazy. I wrote a paper on that. I was very proud of I think discovering one of the first mentions of that connection between electricity and magnetism.

I took a course on modern literature from—good lord, I just forgot his name—the guy who actually was Hemingway's official biographer. Being able to take courses like that was pretty amazing.

**Grier:** You did a fairly broad group of things, yet you still went directly toward computer science.

**Friedland:** Yes. I actually got dual degrees in chemistry and electrical engineering.

**Grier:** Okay, but you chose electrical engineering, more specifically computers.

**Friedland:** Well, it is because that's where computer science was. I was also interested in electrical engineering. As I said, I like putting together Heathkits and the like, but it was a pretty broad background. I got a physics course from a guy who turned out to be famous also later, P.J. Peebles, and learned about the basic introduction to electromagnetism in physics. [turns out P. J. Peebles won the Nobel Prize in physics a year after we did this interview]

## Choosing Stanford for Graduate School

  

**Grier:** But you're coming to the end of it, and you have choices. You clearly could have gotten a very well-paying job. You could have gone to graduate school in chemistry or chemical engineering. You could have gone toward being a professor. And you go to Stanford for computer science.

**Friedland:** Yes.

**Grier:** What were you considering?

**Friedland:** Well, I had decided that I wanted to get a PhD in computer science. At that time, I was starting to get really interested in AI. The top three computer science places were Stanford, MIT, and Carnegie Mellon. I applied to all of them. I got into all three of them. I decided to go to Stanford despite the fact that I'd never been west of St. Louis. I think I was leaning toward that anyway. I got a call from Ed Feigenbaum telling me I really should go to Stanford because, at the time, it was snowing everywhere else and he was playing tennis.

**Grier:** So, pitching the weather.

**Friedland:** The weather in California certainly turned out to be a draw. Having now been here for 44 years and lost all immunity to humidity, I can't stand going to places like Florida anymore.

**Grier:** But it was artificial intelligence that pulled you here.

**Friedland:** It was that, and it was particularly the connection between AI and other sciences. Professor Feigenbaum had started what was called the Heuristic Programming Project. It was the connection between AI and really interesting applications in the sciences, medicine, and engineering that made Stanford particularly attractive.

## Graduate Research at Stanford

**Grier:** During your work at Stanford, you worked on a couple programs: the CONGEN program for constrained generation of chemical structures and the MOLGEN project.

**Friedland:** Yes. The first big expert or knowledge-based system ever built was a program called DENDRAL, which was a collaboration between computer scientists like Ed Feigenbaum and very famous chemists like Carl Djerassi, who of course was famous for a lot of reasons and died recently, which was very sad. But Djerassi, besides being an incredibly good chemist, the discoverer of the birth control pill, and the founder of Syntex, was also a playwright and had several well-produced plays on the national level. He was an expert. They jointly, and their students, designed a system, which was able to outperform expert human chemists in taking mass spectra and predicting what was the structure of an organic compound. CONGEN was the constrained generator part of DENDRAL that produced potential candidate molecules based on the mass spectra.

The first program I worked on there when I was getting my PhD was doing some work on CONGEN. I wasn't a major contributor, but I wrote some of the code for that. I was very lucky to be at Stanford just at the time when molecular biology had moved from being a very data poor science to a very data rich one.

Until about 1974–1975, determining the base sequence of DNA, the string of the letters A, G, C, and T that makes up our genetic code, and it's the language of life, being able to determine say 10 base pairs was an effort of many, many man-months, a laborious, extremely laborious chemical process. But almost simultaneously in 1974–1975, two teams—one in the U.S., which was Wally Gilbert and Allan Maxam, and one in Britain, Sanger a professor, Coulson his post-doc—developed two different techniques that enabled you to read hundreds of base pairs in a short period of time. That meant that instead of having very small amounts of data, molecular biologists had potentially thousands and thousands of base pairs of data.

At that time, one of Ed Feigenbaum's major collaborators, who had introduced him to Carl Djerassi, was Josh Lederberg, who became a Nobel Laureate in biochemistry. He left Stanford in I think 1979, or 1978, to become president of Rockefeller University. Ed had been asking Josh, "Is your field ready for collaboration with computer science?" And Josh had said, "Not yet," until then when he said, "Yes."

So, I was looking for a PhD thesis project, and I was lucky enough to stumble at the right time into the interaction with molecular biology. My thesis was about attempting to model the process by which human scientists design experiments in molecular biology.

**Grier:** I just have to throw this in because this is another part. During your period was the early days of Arpanet.

**Friedland:** Yes.

**Grier:** Did that shape your research environment at all?

**Friedland:** I would say not in the early days. What shaped the research environment then was Stanford had a DEC-10, and later a KL-20, a DEC timesharing machine that came as a result of an NIH [National Institutes of Health] grant for artificial intelligence in medicine. Stanford was a major part of that community that included scientists across the U.S.

  

Having email connectivity to many of those scientists through the network was important, but the most important thing was having easy access to the timesharing system where we did our work. Later, about halfway through my PhD project, everyone moved sort of from that framework for doing the actual computing work to specialized Lisp machines. At Stanford, it was mainly machines from Xerox, what were called the D machines. Dolphin, Dorado. There was another D machine; I forgot the name of it. We sort of had our individual D machines.

**Grier:** For your work rather? Okay, this was before 1979.

**Friedland:** Oh, yes. When did we have D machines? It must have been starting in about 1978 or so.

**Grier:** So, toward the end of your dissertation.

**Friedland:** Toward the end of my dissertation.

## PhD Dissertation Research

**Grier:** Okay, going on to sort of a theme that comes out of yesterday. You were working artificial intelligence but were very clearly identified with a specific problem, or if we use the business term, vertical.

**Friedland:** Yes.

**Grier:** How important was the fact that you were something of an expert in the chemistry involved? How did that shape your work in artificial intelligence? Did it make it easier?

**Friedland:** Well, I mean I did not have a background in biochemistry, so I did a lot of reading and learning. Mostly, getting time from the great biochemists and geneticists at Stanford and their PhD students and post-docs was what was extremely important. I learned a lot about biochemistry then.

**Grier:** Did that shape the computing?

**Friedland:** Well, the way in which the molecular biologists thought is what shaped the computing, learning how the human experts designed experiments. I mean, it was very much the approach of not modeling things at a neural behavior; we weren't building neuromorphic computing systems. Attempting to model in the computer the reasoning processes that truly expert scientists use in designing experiments is what shaped our AI approach to things.

We were not building formal logic models of human thinking. We were building sometimes rule-based, sometimes what's called frame-based, representations of human processes.

Probably the major innovation I developed as part of my thesis work was the notion that… Let me go back just a little bit. Designing experiments is a planning problem. Most prior AI planning work had been in very simple domains, usually the domain called the blocks world where you're starting with three or four or five blocks and designing the plan to move them from one configuration to another. That's what's called a combinatoric problem because there are many possible states of those blocks. The problem was, how do you navigate this space, when do you backtrack, and how do you design the most efficient plan to go from state A to state B, from a starting state to a final state?

Now, the discovery that I made was that human experts almost never do that. They almost never do backtracking. In designing an experiment: if a step in a biochemistry experiment makes it too hot for the next step in your plan, you don't go back and redesign the first step. You make it colder in most cases. If it makes it too acidic, you do something to make it more basic. You do a repair step. My view is that planning in the science fields I was studying at least was not particularly combinatoric but was very knowledge rich. The problem was teaching the planning system enough knowledge to know how to fix things along the way.

The other discovery I had was that, in this field, scientists rarely designed things *de novo*. They took what other people had done and modified them. For example, one of the major types of experiments that was being done in biochemistry was cloning experiments. Actually, there were three things that revolutionized biochemistry. All of them came out at the same time period. One was the method for doing sequencing, which resulted in lots and lots of data.

The second was the discovery of what was called "restriction enzymes." Prior to that, when you were cutting DNA into manageable pieces; it was a totally random process. You threw in something that denatured the DNA and cut it up into random pieces. Then, if you were lucky, you got pieces about the right size. A discovery, and I forget who discovered it, was that there were certain enzymes which specifically look for patterns of four, five, or six bases in a row, the base pairs of A, G, C, and T, adenine, guanine, cytosine, and thymine, and would cut only when it found exactly that pattern of four, five, or six bases. Now, you had a tool for snipping DNA at precise places into precise pieces.

The third thing was the innovation of Stan Cohen and Boyer—I forget his first name—but Stan Cohen is a famous Stanford geneticist and Boyer a famous Cal geneticist. Sorry, not Cal, a UCSF [University of California, San Francisco] geneticist who developed a technique called "cloning," which enabled for the first time, if you had a piece of DNA you wanted to make a lot of copies of, you could analyze it, doing what was called "cloning experiments."

  

How do you do cloning? They developed a technique for doing that, and they wrote a very short little book we got a copy of which explained the steps to do cloning. That resulted in thousands of cloning experiments worldwide basically following the plan of seven or eight steps but modifying each step according to the needs of their own experiment.

In fact, I ended up designing a system called Skeletal Planning, which took this skeletal plan, this framework for a plan, and figured out how to instantiate it for a whole host of different experiments. Instead of there being combinatoric backtracking, like in a typical planning problem, there was almost none of that. The work there was the knowledge in knowing how to specialize each step in that plan and when to insert other little intermediate steps to correct things. That was the view for how scientists designed experiments.

**Grier:** This goes against the common vision of planning. What does this provide as insight into the problems of artificial intelligence? It certainly simplifies it.

**Friedland:** Well, we in the Heuristic Program Project, later named the Knowledge Systems Laboratory, were very interested in how humans design stuff. Ed Feigenbaum had come from the Newell and Simon legacy at Carnegie Mellon. That's where he got his PhD. In Simon's famous book, *The Science of the Artificial*, he describes that one reason for studying human cognitive behavior is to learn a way to model it in a computer.

Other PhD students, Ted Shortliffe, who designed the MYCIN system, an MD himself, modeled the system after what he viewed were the rules that infectious disease experts used, consciously or unconsciously, to decide what antibiotics to use. The DENDRAL project was modeled after how folks like Djerassi thought about problems in mass spectral analysis. Other students were doing projects in other areas of science and engineering. There was one student who was modeling how lawyers think about developing arguments in court.

In my case, we were learning about, at least the best we could, how human scientists thought about and designed experiments.

## Knowledge Capture

**Grier:** Just talk for a moment about knowledge capture, since that was a theme of yesterday.

**Grier:** You were involved in knowledge capture, and that led you to develop a model of human experiment design. You're a knowledge engineer. Tell us about knowledge capture. What did you learn about knowledge capture in this process?

  

**Friedland:** Yes. I suppose I proudly would call myself a knowledge engineer. The first thing is that it's a pleasure to talk to domain experts and to get them thinking about how they do reasoning. It's not always obvious at first to them how they think about problems, but the people I worked with actually really enjoyed thinking deeper about how they actually did their work in science.

Knowledge engineering is really a process of engaging experts and getting them to think with you rigorously about how they solve problems. Folks like Josh Lederberg and Doug Brutlag, another biochemistry professor at Stanford, and Larry Kedes, a professor of medicine at Stanford, and their post-docs, I think it gave them more insight into how they were solving hard problems. We tried very hard to figure out accurately how they were thinking. That sometime took a while to go from, "Well, I just had a flash of insight," to what that really meant. What did you have to think about? What was the knowledge you used to make that discovery?

There was a later project where I supervised a PhD student, Pete Karp, who's now head of biochemistry computing at SRI. We studied a scientist called Charlie Yanofsky, who probably would have gotten a Nobel Prize, except he died too soon. He had a remarkable insight into solving genetic regulation in the tryptophan gene; it was a problem that was really great for studying because all the information was there that any trained biochemist could have come up with the insight he did. He did only because he actually was able to think in two dimensions at once to imagine the ribosome both as a machine for producing proteins, but also as taking up space. It turned out nature had evolved a genetic regulatory mechanism that took advantage of that. When he had this flash of insight, which was caused by him being able to think in two dimensions at once, it turned out to be correct.

That was an amazing discovery. I really, really enjoyed the time spent with those experts trying to figure out how they think and also trying to figure out what knowledge specifically they used in solving their problems, which became the bulk of the knowledge that we encoded within our systems.

## Postdoctoral Work at Stanford

**Grier:** Okay, you finished your PhD in 1979. And you stay in the computer science department.

**Friedland:** Yes. I stayed in there as a senior research associate because there was a lot more work to do in the MOLGEN project, turning what I had done into a system that could be duplicated for other problems in molecular biology. I had a really good PhD student called Yumi Iwasaki who refined the skeletal plan concept into a system that could be repeated on many different systems. And Pete Karp, who I mentioned, modeled the thinking of this other famous biologist, Charlie Yanofsky. I stayed there in the computer science department until 1987.

**Grier:** Right. You could've gone elsewhere, if you would have wished to have been a professor of computer science.

**Friedland:** It turned out I didn't really want to be a professor. I really enjoyed doing the research more than the teaching. I enjoyed working with one or two students at a time, not teaching big classes, so I wanted to stay in research of that type.

**Grier:** Okay. A little bit of an aside, but it's relevant to people who research these sorts of things. You're with the university. The interaction with SRI, did you have any? What was it like?

**Friedland:** I did because SRI had a really good AI group of people: Nils Nilsson, Peter Hart, Dick Duda, Earl Sacerdoti, lots of folks. They were also doing work in planning and had a whole community of work doing that.

I was mostly research assistant at Stanford for my stipend as part of our PhD project, although I did one little thing as a teaching assistant for Nils in his course on AI, his introductory course on AI. That was my one experience of teaching large groups of people. I also became the editor for Nils' book on artificial intelligence.

I frequently visited the SRI folks because they were our colleagues. Many of them had joint appointments at Stanford. Nils was at both places at once. So, SRI was part of the same group of people doing this work.

**Grier:** Right. But it was just sort of a world apart, and you were friendly and interacted on some things.

**Friedland:** They were research colleagues. We got together for joint seminars. I asked people like Nils for advice on my thesis. I'm trying to remember; but I believe Nils was on my thesis committee.

## Intelligenetics and Other Spin-Offs

**Grier:** During the same timeframe that you're doing this, you're also with IntelliCorp. You list three, IntelliCorp, Intelligenetics, and Teknowledge, in Palo Alto, California.

**Friedland:** Yes, as I was just about finished with my PhD thesis in 1979… To back up a bit, from our point of view getting three or four hours a week of deep conversation with all of these wonderful scientists was incredible. One of the things we did to pay them back was write some pretty simple computer programs.

  

These programs helped with  DNA and protein sequence analysis, although they didn't have too much to do with AI research. Essentially, finding patterns in the strings of letters representing DNA base pairs or protein amino acids. Now that they had this data, there were looking for stuff. What were interesting patterns of the data? What were the portions of a DNA strand that actually encoded for proteins? And what were other sequences that were basically control signals for when to stop and start doing that? I ended up writing some programs for that, and they turned out to be useful.

The folks they were useful for talked about them to their colleagues who also were getting these data. All of a sudden, we had lots and lots of molecular biologists around the country wanting to use these programs. They were still running on our timesharing system, on the SUMEX-AIM computer. We created a guest account called Genet that provided access eventually to thousands of outside scientists. And we got credit for that in lots of research papers that people wrote.

It got to the point that two things happened. One, it got to a point where the director of SUMEX-AIM, Tom Rindfleisch, politely asked us, "You know, you're using two-thirds of the cycles on SUMEX-AIM. Could you please get off and find your own computer?"

Also, there became sort of an issue. We had no problem with releasing the stuff for free to academic scientists, but there was a growing molecular biology industry, and companies were being formed. Scientists at Stanford and lots of other places were wearing dual hats at the companies they started at their universities. While we thought it was quite reasonable since we were funded by NSF and NIH to let the academics use cycles on SUMEX, it didn't seem right to let the companies freely use this stuff.

Basically, we got a ruling that they weren't allowed to. It led to a few hard feelings because scientists at other universities who were wearing dual hats said, "Well, how come you let me use it for my research work at Dartmouth or someplace, but I can't use it at the company I started?" We said, "That creates a conflict with the federal government with providing free access on those machines to corporate ventures."

We realized those two things at once. We realized, first of all, we didn't want to abandon the community, so we were going to have to get our own computer. We also saw this growing commercial market. Plus, there now was starting to be this growing commercial interest in expert systems. I think it was in early 1979 that all of us in the HPP, the Heuristic Program Project, were asking, "Should we all get together and form a company?"

Ed and I were sort of the strongest in saying," Yes, we should do this." We offered all of our colleagues in HPP the opportunity to join us. Most said, "No, we're not that interested at this time." But our two main biological collaborators at the time, Larry Kedes and Doug Brutlag, were. We said let's form a company that both does the bioinformatics type of work as well as explores uses of AI in science and other applications. The goal was, to take the software we developed at Stanford, which now we had permission to license from Stanford because the NSF had decided that instead of software being developed as part of university grants that it was public domain; that actually it was owned by the university with the proviso that the university really do an honest job of making it commercial and available to the world as a whole. This included the software we had developed for DNA and protein sequence analysis, asw well as the beginnings of what we thought was an AI product for molecular biologists to help them design experiments. In addition, we had general tools for building AI systems that we thought could be used for many other applications.

We started this company, I remember, sitting around in Doug's house in Ladera, a community near here, and thinking about names. To be honest, I don't remember whether it was me or Larry who first thought of Intelligenetics. We created a company called Intelligenetics, and being in Silicon Valley, we did the rounds of venture capitalists.

Got a lawyer. Lawyers in Silicon Valley serve several purposes, but probably the most important one from my point of view is they become your business advisors. Most law firms will basically waive your fees until you raise venture capital money as part of being part of your advisory team. We found a great lawyer Greg Gallo. At the time, the company was Weir, Fletcher, and Friedenrich. That firm went through about eight different name changes over the next 20 years, but Greg Gallo was a senior partner there over that entire time. He adopted us, and we adopted him. In fact, he became the lawyer for all the companies I founded. So Intelligenetics got started.

Interestingly enough, about two years later the group of people at Stanford who hadn't wanted to form the company decided," Yes, it's time to form one." About 14 or 15 people from our Stanford group and then a couple of our other strong colleagues like Rick Hayes- Roth, who was at Rand, and a couple of other people who were colleagues and friends, some from SRI, formed a company called Teknowledge, which was designed to be a training and tool building company for AI systems. I was still at Stanford but spent a certain percentage of my time at Intelligenetics and later a smaller percentage at Teknowledge both as a board member and a consultant helping design things.

One of the first things we did at Intelligenetics was apply to the NIH division, which had funded the SUMEX-AIM computer, for a replacement system specifically for first serving the academic molecular biology community. That was called BIONET. We wrote a proposal for BIONET. It got awarded. We had a five-year grant, and we ran that for the community and for literally thousands of projects across the country and the rest of the world.

Five years later, we applied for renewal, but at that time, there was a lot of rumbling in the community about, do we really want this run by a commercial company even though we were running it as a nonprofit just taking expenses for the BIONET facility charging minimal fees to pay for costs? I remember one group of scientists didn't have the $1,000 or so to pay the year for BIONET use for academics and actually had a BIONET bake sale to raise the money.

**Network Connectivity Before the Internet**

**Grier:** There's interesting because this shows a tech company breaking away from a university.

**Friedland:** Yes, which became increasingly common at Stanford.

**Grier:** Which becomes increasingly common. You talked about first how you created accounts on the original SUMEX-AIM machine for them to use. But this is still in the age of Arpanet. That's the only connection they have. You're starting with just the communities of ARPA.

**Friedland:** By that time, most universities had nodes on the Arpanet.

**Grier:** Most research universities.

**Friedland:** Most research universities did. There were couple of ways to logon to SUMEX-AIM. Certainly, you could do it through the Arpanet. Actually, most people logged on in the early days through their 1300, 1400 baud modem. There was a dial-up number through Telnet.

In fact, I take that back. Very few actually used an Arpanet connection. That was mainly limited still to the major research universities. Most users of the Genet account on SUMEX-AIM logged on through their friendly headphone coupler modems.

**Grier:** I’m not sure what this is, maybe referring to the TI Silent 700 terminal?

**Friedland:** That was the thing, most of the times they were using things like… What was the company that was making the smart terminals? Data something, I forget. Actually, Heathkit Zenith produced a smart terminal for a lower price. We actually had a building party among grad students. We offered every grad student in the computer science department, one year, a free Heathkit Zenith smart terminal if they built it. We had a building party where I taught them how to build Heathkits over a weekend.

**Structure of Intelligenetics**

**Grier:** Going, again, to the separation process. You have it set up. It's working where you set it up basically for the use of academics. It's clearly expanding that use.

  

**Friedland:** Yes.

**Grier:** You set up a separate computer facility.

**Friedland:** We moved from SUMEX-AIM to BIONET.

**Grier:** Where was BIONET located?

**Friedland:** Where was it located physically?

**Grier:** Not on campus at that point?

**Friedland:** No. It was not on campus. Was it physically at our Intelligenetics location? God, I don't even remember. I think we rented separate space for the computer facility.

**Grier:** Okay. But it's a separate computer.

**Friedland:** It's a separate computer. It's a separate DEC KL-20 computer.

**Grier:** That is, in effect, selling services back to the university?

**Friedland:** It wasn't really selling services. We were required to do cost recovery and use the money from NIH to buy the computer and to use for separate research projects that Intelligenetics did. You could apply to do separate research projects on computing and biochemistry on the system, but the sort of production use for DNA sequence analysis and the like was done on a cost recovery basis. It was a very small amount. It was to pay for maintenance and operations in the system. It was like $100 a month or something.

**Grier:** Okay. So, this was very much a not-for-profit activity.

**Friedland:** BIONET was definitely a not-for-profit.

**Grier:** Not only not-for-profit, it was subsidized by your NIH grant.

**Friedland:** Well, the NIH grant paid for the computer. It paid specifically for operations. It paid for some of the staff cost. We were required to do at least five or ten research projects under the grant for advancing the art of computer-aided biochemistry. People applied for those. Some people at Intelligenetics could apply for that, but an independent board of advisors decided who got the research grants to do computer-aided molecular biology.

  

**Intelligenetics Growth**

**Grier:** Okay. Again, sort of working on this arc a little bit. You're listed as founder and vice president?

**Friedland:** I was a founder and vice president of Intelligenetics.

**Grier:** Vice president of what? Or just vice president. They handed you the title: here be a vice president.

**Friedland:** Yes. Vice president and member of the board.

**Grier:** And this ran through 1986? Or was that your participation and it ran longer?

**Friedland:** Well, Intelligenetics the company was interesting. The company went public in 1983 mostly based on the expanding trend of computer-aided molecular biology. The biology part of the company caused it to go public, but starting around 1984, 1985 the expert systems trend in the world, the hype in the world became the dominant force. Our second offering was almost entirely fueled by that growth. The molecular biology part dominated the company in the first several years, but slowly and then rapidly the expert systems tool and applications part began to become transcendent.

**Grier:** You have an IPO in 1983, four years after you begin operations.

**Friedland:** Yes.

**Grier:** You have enough revenue that the bankers feel good about all of this and felt you could sell this?

**Friedland:** In the words of Thom Weisel, head of Montgomery Securities, "Revenues is a bad thing." Remember the motto of that period of time in the Internet world was the power of growth, so profit was a bad thing. Revenue wasn't bad. Profit was bad. Revenue was supposed to be used to fuel growth. You got measured on the projected growth of the company as opposed to individual revenues and God forbid profit.

We did have substantial revenues. I think by 1983 we had… Oh, I didn't bring the detailed list with me, but by 1983, it was about $2 million a year.

**Grier:** People will be looking it up who use this thing anyway.

  

**Friedland:** About $2 million a year in revenue. In 1984 it was about $5 million. It jumped to like $10 million right after our second public offering.

**Grier:** Here's the question. Some of that money was cost recovery, and some of that was from your commercial use.

**Friedland:** Well, talking about revenue. I'm not counting the cost recovery.

**Grier:** That's entirely from the commercial side?

**Friedland:** Entirely from the commercial side.

## Formation of IntelliCorp

**Grier:** Okay. That IPO was in 1983. This lists your termination as 1986. Do you move on from this? Do you stay on the board?

**Friedland:** My termination? In January 1987 I became part of NASA, so my time available for participation became extremely limited. We had our second IPO in 1986. A year earlier, because of the fact that the computer science side, the AI side, was growing, we decided to rename the company IntelliCorp. Intelligenetics became a sort of division, a branch of IntelliCorp.

Then in 1986, Amico, who was a significant customer of IntelliCorp, but mostly on the Intelligenetics side, offered to buy a 60% share in Intelligenetics. In other words, take the Intelligenetics division, physically move it away and make it a joint venture between IntelliCorp and Amico, with Amico owning 60%. So, it moved half a mile down El Camino from the IntelliCorp facility. At this time, the offices were at Rengsdorff and El Camino, in an office complex there. It moved down to across San Antonio, across Grant Road. The other side of Grant Road at Mountain View. IntelliCorp was in Los Altos, and we moved Intelligenetics to Mountain View and still on El Camino and ran it as a separate company.

After I joined NASA in 1987, I really didn't have much involvement with IntelliCorp because part of that might have been a conflict of what I was doing with the federal government. I was still on the advisory board for Intelligenetics, which was ruled not a conflict. I was still a stockholder.

**Grier:** You just had no active participation or limited active participation.

**Friedland:** In the IntelliCorp side, I still had friends there, but I had no active participation in IntelliCorp after I joined NASA in 1987.

**Grier:** The other company of this era was Teknowledge. You were primarily a consultant on that?

**Friedland:** Yes. I was one of 20 cofounders. What I did for Teknowledge occasionally in that period until 1987 was participate in site visits to potential customers where we would learn about their problems and describe how Teknowledge could provide expert systems solutions to their problems.

**Joining NASA**

**Grier:** Okay. But you've already hinted at the next stage of your career, which is NASA. That starts in January 1987.

**Friedland:** Yes.

**Grier:** How did that job come about?

**Friedland:** Well, in 1986 or beyond I was thinking about what I wanted to do next. We had sort of come to a point in MOL GEN where we could declare success and be happy with its spin-offs and things. It turned out that my old advisor at Princeton, Peter Denning, was now head of an organization called RIACS, Research Institute for Advanced Computer Science, which was on the NASA Ames campus and was sort of their in-house advanced computer science shop. In late 1986, the space station program was being planned.

I learned some interesting things about the way things operate. It turned out the most powerful person as far as NASA was concerned, at that time, was not the NASA administrator, but a person with the rather unassuming title of chief clerk of the House appropriations subcommittee for HUD [Housing and Urban Development] and independent agencies. This is a guy named Dick Malloy who basically designed the legislation involving space policy. One of the things he got into the legislation was that the space station program was to devote up to 10% of total budget to research and development in, autonomous systems and robotics. That meant that NASA now believed that it really needed to develop a substantial AI and robotics program.

Peter Denning and I had kept in touch. He called me and said, "NASA wants to create this program." There were a few people there who knew a little bit about AI. There were a few people sort of around all of the different NASA centers who knew a little bit about AI, but no serious research program.

At the time the center director was Bill Ballhaus, who is tied for being the best boss I ever had. Bill was one of the few people who was both a really good CEO and COO [Chief Operating Officer] at the same time. What I mean is he had the vision and the way to describe it as CEO, but he also both knew how to run things. He knew how to hire a deputy who was really good at running things. I talked to Ballhaus and what they wanted to build. I forget who told me the most important question you should ever ask someone like that is, how long are you going to be around?

**Grier:** How long was he going to be around?

**Working for the Federal Government**

**Friedland:** Well, he said he certainly had no plans of leaving and he really wanted to see this developed. He was someone who had come up through the ranks as an engineer who had pioneered the concept of computation through flight, which means Ames had had and when I went there still had the largest collection of wind tunnels in the world. Wind tunnels are used for the design of aircraft, but he pioneered work on computational fluid dynamics [CFD] and the way it interacted with testing and wind tunnels, so you could basically do computer design as a major complement to the wind tunnel design and eliminate the need, in many respects, for wind tunnel testing of certain types.

So, he was an engineer turned executive with being appointed Ames center director.

He basically said, “I will give you the resources to make this happen. If you do sensible things that sort of violate the rules, I'll have your back.” That meant things like, “There's a certain number of slots I can give you for civil servants, and I'll get you all I can. But if that's not enough, you can hire contractors, and for all practical purposes treat them like civil servants.” You know, things that were normally verboten and probably did violate the regulations like sharing offices, and for all practical purposes, senior contractors having civil servants underneath them in teams, we did. It was probably technically against formal regulations. It certainly wasn't illegal. But Bill said, “As long as you get done what's good for NASA, I'll have your back.”

We also had a couple of people at NASA headquarters who led the funding who became part of our team who really felt the same way. A guy by the name of Mel Montemerlo was my partner in crime from 1987 to 1995 while I was at NASA. His boss Lee Holcomb also was a tremendous asset. During that period from 1987 until about 1994 it was a pleasure to work at NASA.

One thing I learned afterward, with respect to the Air Force, is that the federal government can either be a good place or a terrible place to do R&D. It's a good place, not just R&D, to get missions done if the infrastructure believes they’re part of the team, if they’re behind you. If procurement and legal and contracting and personnel and the like—except procurement and contracting are about the same thing—are behind you and they believe they’re part of the mission, it's a pleasure to work there. You can get things done.

  

If those organizations believe their job is to protect the organization from crazy scientists as opposed to help them make things get done, it can be a real pain to work there. At the time, all those people, all of those departments reported to the center director, Ballhaus, and if he needed things done in legal, he could call up the head of legal or who reported to him and say, “Hey, I need you to review that contract tomorrow, not two weeks from now. Get it done.” He could tell personnel, “Get those people hired in a month instead of a year.” It made things really easy to get things done.

**One NASA Operation**

**Friedland:** The other thing we were able to do was treat it as a sort of a “one NASA project.”

**Grier:** Meaning what?

**Friedland:** Well, NASA, like any large organization, has internal rivalries. There are nine NASA centers. You know, people say you work at NASA and they say, “Oh, you're in Houston.” Well, no. Actually, NASA headquarters is in Washington. There are nine NASA centers that do various things. There are four manned spaceflight centers: Johnson in Houston; Marshall in Huntsville, Alabama; Kennedy in Florida; and Stennis in Mississippi just across the river from New Orleans. There are two space science centers: JPL [Jet Propulsion Laboratory] and Goddard. There are three-and-a-half research centers. I say half because the fourth one sometimes is part of Ames and sometimes isn't, which is Langley in Virginia. Lewis is now called Glenn in Cleveland. Ames is in Mountain View, and Dryden is now called Armstrong, which is the NASA part of Edwards Air Force Base. Each of those centers is a partner, but also a rival of the others. Rivals for permission to build buildings, and rivals for money.

At the time when I got there, there were sort of these nascent AI organizations at many of the other centers. and certain people at those centers were jealous of Ames getting the responsibility. When I first got there, the steering committee for assigning money to this was very political. One of the things I said to Lee Holcomb was I'd like to replace the people on the steering committee with people at the centers who are really in it for designing and building the whole NASA program. He said, “Yes, go ahead.”

That was actually a bit of a fight to get some people who were more political than technical off that group, but headquarters backed us, and we ended up with the group of people who really did something that was very unusual at NASA. We assigned money based on merit rather than by center allocation. We developed a process for essentially receiving grant proposals from all of the centers for the work that we were going to do under this whole AI program. I cochaired that with Mel Montemerlo. We had representatives from every center on it. Took a little bit of prodding to get it to be a one NASA operation, but we really did.

  

**NASA Applied Research**

**Grier:** Can we go from that to the actual kind of projects that you were doing?

**Friedland:** Yes. Well, a little bit more detail that's necessary to understand. It ended up being two research laboratories: one at Ames, the largest one and a smaller, but still very significant one at JPL. Applications projects were at all of the centers. At Ames the goal was to hire very, very good research scientists. Lead scientists who were good enough to be professors at Stanford or MIT or CMU, but also wore two hats that they would do the basic research but that wouldn't be the end of the stuff when they published. They would do whatever was necessary to take what they had done as basic research and turn it into NASA applications, which we did.

Over the period of within two or three years, we had not only dozens of research publications but also over a dozen fielded applications in NASA's main missions in space science and human space flight and in aeronautics. That part of the period lasted from 1987 until about 1994 and it was sort of the golden age of doing this at NASA.

Not that we didn't have occasional problems. The hardest thing I ever had to do in the federal government was fire somebody. Believe me, you don't want to ever fire somebody in the government, but I had to. That ended up going through two years of stuff ending up at the Supreme Court. Literally the Supreme Court, although they rejected the case. They were the final arbiter that was allowed to accept or reject the federal appeals case, and they did it. It went through all including the Federal Court of Appeals for the Washington circuit, the first circuit court which takes all government internal final appeals. In fact, it's where most Supreme Court justices come from.

We had a program that grew in size from an original approximately $10 million program to a NASA-wide program of well over $100 million.

**Grier:** Okay. $100 million program headquartered here.

**Friedland:** At Ames. Yes.

**Grier:** You were the director of it here.

**Friedland:** But comanaged by a group that a headquarters person and I co-led. I was the director of the Ames Laboratory.

  

**Constraint-Based Scheduling Problem**

**Grier:** During this period, can you get into some of the projects and the work that you were doing?

**Friedland:** Sure. A lot of very interesting projects. One that Monte Zweben, who was in the group yesterday, developed. Monte became my deputy and also ran a project in constraint-based scheduling. He glommed upon what was probably NASA's hardest scheduling problem as their domain to work in, which was the problem of scheduling what was called the processing of the space shuttle orbiters after they came home. That was a process that originally was supposed to take 12 days.

I found out about this because the flow sheets that Monte showed me had 12 days on each of them. I said, “Why is it 12?” He said, “Oh yeah, that was the original plan to turn around orbiters in 12 days. Fly about 10 times as many missions as it actually flew.”

It actually ended up taking about two months to process each orbiter. This was the process done at Kennedy Space Center at what was called the orbiter processing facilities. There were two of them so you could process two orbiters at a time. Essentially, fixing up any damage, replacing tiles, putting more of the fuels and other supplies (liquid supplies, gaseous supplies) that you had to do. It was a complicated problem because there were a lot of timing constraints, like when you were doing tiles, they used a very toxic adhesive to replace the tiles on the shuttle. You couldn't have anyone else who wasn't in hazmat suits doing that. It was a very delicate process.

It was a process that cost the agency about $1 million a day mostly in labor costs to what was then Lockheed Space Operations Corporation, all overseen by a single guy called a flow manager. Monte became very friendly with two of the flow managers, the one for Endeavour and one for Columbia. At the time when we got there, the scheduling was done literally by long sheets of paper on a wall. People cut and pasted things. Their total budget per shuttle allocated for paper and X-Acto knives and glue was $12,000 per mission.

Monte built a just-in-time scheduling and rescheduling system that enabled you to find correct solutions at any time that you had to reschedule. But they could stop at any time with the guaranteed correct solution. It might not be the optimal one, but some of these were timing dependent. It ended up replacing this totally manual system with a new computer graphical and automated system for those two shuttles. The cost savings to the agency, mostly by reducing the need for weekend overtime, was well over $1 million per shuttle flight. We ended up getting the largest ever Space Act Award, monetary award, for service to NASA from that project. That was a really interesting project, advances in the constraint-based scheduling, lots of good publications, and a book out of it, but also the application to NASA.

  

**PI-in-a-Box**

**Friedland:** Another very interesting, much more an expert system project, was called PI-in-a-Box. When you're a space scientist in a laboratory at MIT, Rice, Stanford, or somewhere and wanting to do an experiment on the space shuttle, they're planned years in advance. If you're real lucky, you get space on the shuttle to do the experiment. Almost always, it's an astronaut scientist who carries out the experiment—what are called the mission specialists are trained astronauts. There are mission specialists, and there's pilots, pilot commanders, who are astronauts who are usually former Air Force pilots or Navy aviators. The mission specialists are scientists who become astronauts.

In fact, when I first came to Ames I shared an office with Ellen Ochoa, who ran a photonics laboratory. She really wanted to be an astronaut and eventually did succeed, becoming the first flautist in space. You're allowed to take one personal item, and she was a symphony grade flautist. She took her flute. Later she became head of the astronaut office and is now the Johnson Center Director.

So, those people are trained on your experiment, but they're not like they were your post-doc doing your experiment. If it broke, in most cases, that was it. A power supply failed or something, and this thing you had planned seven or eight years for was done They weren't capable of on-the-fly data analysis the way you are on the ground. If you're doing an experiment in science, and something interesting happens…

**Grier:** You can adjust the experiment.

**Friedland:** Yes, it might be the thing: "Wow, that's the discovery." An anomaly might be eureka. You couldn't do that: analyzing the data and then re-planning on the fly. So, we wanted to do those three things for the experiment. We collaborated with a very famous space scientist, Larry Young from MIT, who studies vestibular physiology in space or why humans throw up in zero gravity. He had several experiments on the shuttle using a device called the "rotating dome," which was this dome you put your head into where spots of different colors rotated. Apparently, you can measure the eye motion and other things to help you understand how the human vestibular system changes under zero G.

Larry actually tried to become an astronaut, and he was the world's oldest astronaut candidate. He made it to the training phase but not to actually flying.

We tried to replicate much of Larry's knowledge in this PI-in-a-Box experiment, which we actually flew on space. It turned out the astronaut who used it was Shannon Lucid, who became famous because she had, for a while, the most hours logged because she was on Mir, a NASA astronaut on the Russian space station, for a long period of time.

  

It worked. It was also the first flight of a true laptop on the Space Shuttle.

## Modifications Necessary for Space Flight

**Friedland:** We used the Apple PowerBook for the experiment, and we learned what it takes to fly something on the space shuttle. You couldn't use the batteries on a laptop because they were worried that, at the time, the nickel-cadmium batteries could explode. We had to connect it to the shuttle's supply. And we had to get a space-rated power cable.

I remember the project manager at the time, Silvano Colombano, who once a week came into my office and told me why we weren't going to fly; one time he came in and said, "We're not going to fly." I said, "Why?" He said, "Well, we need a space-rated power cable." Okay, what's the big deal about getting a ten-foot power cable? It's two wires. What's the big deal? Well, it has to be space-rated. Well, okay, who does that? Well, there's a lab at MIT that's funded by the Johnson Space Center that has a two-year backlog.

This was before the Internet, so my Internet at the time for NASA was my collection of NASA phone books for each center. I look in there and look up the head of the space flight projects branch at JSC. I call up that person and get the branch chief. He said, "Yes?" I said, "Well, you know I'm another NASA branch chief. What can you do to help me?" He said, "Well, the reason we have such a backlog is we can only afford one contractor at the lab to do it." I said, "Well, suppose I go 50/50 with you for the next year on another contractor. Will you put me first in line?" He said, "You could do that?" I said, "Yes, this is a couple million-dollar project. I'm happy to save you $35K." So, I called Mel Montemerlo and I said, "Hey Mel, is it okay if I sent $35K to Johnson so we get a cable for the PI-in-the-Box?" He said, "Of course."

**Grier:** So, it's a $35,000 cable, but more came out of it than that.

**Friedland:** Well, it was $12K for the cable but $35K for the contractor co-shared with JSC to make them do it. So, we flew. We flew PI-in-a-Box. That was the first expert system ever in space.

## Unsupervised Data Classification

**Friedland:** Another really significant project was on the data-analysis side. One of our best scientists was a guy by the name of Peter Cheeseman, who was a Bayesian reasoning expert. He designed a Bayesian classification system called AutoClass. When people talk about machine learning these days, almost entirely they're talking about supervised, neural-net-based classification systems. It's an old idea. The idea is at least 35 years old. The reason it works so dramatically well today is because of two things that are totally apart from AI. One is the availability of huge amounts of accessible training data for almost anything. If you want to train your system, say, to classify dogs and cats, there's millions and millions of pictures of dogs and cats on the Internet. You can have people on Amazon Turk label them dogs or cats, and then you can give that training data to your neural net.

And the other thing: that takes lots and lots of computing power. People discovered that GPUs that were originally designed as graphical processing units for gaming computers and the like are wonderful.

**Grier:** For this problem.

**Friedland:** Running those deep-learning neural nets.

**Grier:** Right.

**Friedland:** Huge amounts of specialized computing power, huge amounts of data make supervised learning systems do dramatic things. Now, the goal we had was to do unsupervised learning. You get data from a space mission, an astronomical mission. It's studying things in the infrared, or the UV, or the visible. You don't have them classified because that's the job. Which blob is a galaxy? Which blob is a nebula? Which blob is a quasar?

Now, human scientists do that off the photographic data. Peter's system, in theory, found where the classes were. It said, "Okay, there seem to be four distinct classes of astronomical objects." Peter's system looked at the infrared astronomical satellite (IRAS) data that was our testbed, lots and lots of data, image type data, in the IR [infra-red] and said there's this many classes, and here's what's in each class.

The astronomers looked at that and said, "Yes, that's better than we did in pulling things out. This class is clearly a galaxy. This class is clearly a nebula. That's clearly quasars. Let's replace our former catalog, which was about 2 inches thick of astronomical objects, with the one you guys gave us." Human beings still said what the meaning was. The system didn't say this is a galaxy. It said this is a class of things, and there are this many classes of things. It worked remarkably well. The proof was, as I said, I have this catalog that the astronomers adopted as the classification catalog for the infrared astronomical satellite.

Those are three examples. I can probably give you a dozen more, but those are three different classes of examples. In each case, working with human experts was key. In the case of the scheduling system, working with Eric Clanton for Endeavour and then Wayne Bingham for Columbia—who were the flow managers was important. In the case of IRAS, working with the infrared astronomers. In the case of PI-in-the-Box, working really closely with Larry Young is what mattered.

  

**Leaving NASA**

**Grier:** We need to move on to Intraspect. Talk to me about the transition from NASA to Intraspect.

**Friedland:** Well, starting in mid-1994, the NASA administrator, who'd been a good guy, was replaced by Dan Goldin, who could best be described as NASA's Rasputin. He was a madman. NASA, in its best of times, is normally a meritocracy, lots of incredibly smart people there. You really can't get too much of an ego when you're associating with people like Jeff Cuzzi, who was one of our collaborators. He's the chief project scientist on the Cassini mission. His nickname is Lord of the Rings. If you're around people like that, just like being at Stanford around people like Josh Lederberg, you can't get too much of a swelled head. In the best of times there, NASA's a meritocracy.

When Goldin came in, it turned out the primary criterion for his lieutenants were people who would tell him how wonderful he was, reminiscent of certain other things. NASA went from being a meritocracy to being very political. For reasons that I could go into it—it'd be about an hour to go into, but I won't—he fired the Ames center director. Replaced him with someone who was, by all standards, totally incompetent to be there. In fact, a very strong belief, which is probably true, is that he picked someone who would give him an excuse to close Ames.

The people who were not good enough to prosper in a meritocracy came out of the woods. It was becoming a very unpleasant place to work. It was obvious that I could no longer really do my job appropriately. I was told, "Good news, we don't have to work with people at other centers anymore. Headquarters will decide who gets what." It was pretty clear that we'd had a great run, but the structure we had built couldn't survive as people would leave. So, I left.

**Knowledge Management at Intraspect**

**Grier:** So, you left. Intraspect, what was the thought for that?

**Friedland:** Well, I'd done a short tenure in a company called EIC, which was one of the early Internet-enabled companies, and had met Tom Gruber, who I'd known earlier. He was a later PhD student of Ed Feigenbaum's.

**Grier:** Right.

**Friedland:** And I got to know Craig Wier, who had been a DARPA project manager. We were all at this little company trying to develop web-based applications. It turned out that it wasn't the right setting for us because the company wanted to stay small and do little projects and we had what we thought was a big idea.

So, we left. The big idea was what came to be called knowledge management, which was making a corporate memory for an organization so that you could reuse rather than reinvent. In large organizations, if you have a problem to solve, the odds are very, very high that somebody else somewhere in your organization has solved a similar problem already, if it's a large enough organization. But in most companies, trying to find out who that person was and how it was done and why it was done is nearly impossible. I mean, companies like Citicorp have hundreds of thousands of employees.

Tom was one of the very first people to combine email with the web. He had built a system at Stanford originally that produced a web-based mail system. He viewed the integration of unstructured data, meaning everything from Word documents, PowerPoint, email, web conversations, blogs, video, everything that isn't in a relational database, integrating that into a corporate memory system so you could find it, subscribe to it, and find it in the context of prior use. We formed a company originally called Colloquy.

We later changed the name to Intraspect, mainly because Tom hated the name Colloquy. He originally liked it, and then he decided he hated it. Since he was our CTO, we agreed to change the name. It was a better name. We all agreed.

We started a company, got venture funding, did the rounds with the idea, and began Intraspect. Started, I guess, in 1996. It grew. We attracted increasing amounts of venture capital, found a bunch of corporate partners, developed the system through many levels of development, and went through a couple of CEOs as we talked about yesterday. Hard to find CEOs who are good enough that satisfy all your venture capitalists at the same time.

I was CEO for a while until we got our third major round of funding, and we were growing enough that we realized that I should find somebody, when it was becoming a sales oriented company, as we developed this as a product, a shrink-wrapped product, the Intraspect Knowledge Management System. We started winning awards from the major analyst groups and so on and developed a market. The goal was to find somebody who was much more knowledgeable about sales and marketing than I was.

We found a CEO, Jim Pflaging, and it was time to go public. Unfortunately, we missed our timing by about three months. The company had grown to the size where the valuation according to the bright young MBAs who came over from Morgan Stanley, who was our underwriter, told us it was going to be a half-billion-dollar valuation. And then the market went… We had written the S-1, about to go public in about May of 2000, when the technology bubble went boom.

**Grier:** Started sliding down.

  

**Friedland:** More than sliding, it went boom. I mean we're talking three months after the bright young MBAs told us our god-given valuation was $500 million, the same bright young MBAs explained our god-given valuation was $35 million. Our board, which at that time had people like Michael Boskin, the famous economist, on it, said to Morgan Stanley, "We really should wait six months until the market recovers."

That six months turned out to be about 10 years. Unable to raise funds in the public market, although we had grown to a high expense level. We were in this building just north of SFO [San Francisco International Airport] in Brisbane off 101 [freeway] in the top floor of a new building that had been built, a 10-story building. The board wanted signage so we could attract the best sales guys, so $1 million dollars for signage on the building.

**Grier:** A million?

**Friedland:** Yes, I voted against it. I was the only person voting against it on our board at the time. When it later became time to cut back because the availability of funds was hard, what it meant was it was harder to sell into the big organizations that were our market because they wanted guarantees that the company they were buying from would be around because, not only were they buying a major software package, they were buying services for it. Only a few companies were big enough for companies to trust, even though we were a far superior product to something like Microsoft SharePoint, which was a competitor, but not really; they knew buying a far inferior product for more money from a company that would be around was a lot safer.

So, new sales were drying up. We earned far more revenue from existing customers on our support side than on our product side. Eventually, the VCs decided they were going to try to get some of their money back and sold the company for $35 million.

## <u>Return to NASA Ames</u>

**Friedland:** That was in late 2002, at which time, a very interesting development happened. I had gotten a call from the new Ames center director. A new NASA administrator had been appointed. Goldin had been replaced. The new guy saw what was happening at Ames and replaced that center director with actually a good person. That person got the history of what had happened, called me up and said, "Look, I want to apologize to you for everything that was done at Ames. Would you consider coming back as our chief technologist?" Which I did.

For somewhat bizarre reasons, during the time I was at Ames from 2003 to 2006, I was actually the highest paid employee at NASA. At the time, the NASA administrator was Sean O'Keefe, who was the most political and least technical administrator NASA ever had. He was a protégé of Vice President Dick Cheney. He basically left the technologists alone, which was good. He had people underneath him who were good.

At the time, you may remember, there was a presidential scorecard for federal agencies. You could be red, green, or yellow in about eight categories like personnel, finance, project planning, something like that. Nobody was green on everything. One of his goals was to take NASA from a yellow to a green on personnel. At the time, there was a little-known federal thing called excepted science and technology positions. There were only four of them in the government at the time because no one really knew about the process. But if you could make an argument that a science or technology person was essential to the agency, you could get a position where you were waived all the normal salary requirements. You could be paid up to the vice president's level.

So, I went back, and the center director Scott Hubbard said, "Okay, Sean wants us to do what we can to use this position." I said to Scott, "Okay. Can you really make an argument that without me, a computer scientist, that NASA would fall apart? Just give me the SES slot that we talked about." He said, "Well, no, Sean really wants this. Since he's the one who's going to approve it at the NASA level, and he just came from being deputy director of OMB, the Office of Management and Budget, which has the final approval, it'll probably get approved." I said, "Okay, fine. If you want to pay me more than your salary, go ahead." This was in, I think, around September. So, I said, "But if this isn't done by Christmas, let's go back to the default state."

So, it goes along. We hear that it's been approved by Sean. It's been approved by Office of Personnel Management, but it's sitting at OMB. A week before Christmas, the head of personnel calls me up and says, "You're not going to believe this. It was actually approved." It was called the ST-EX position.

**Grier:** You were the fifth person.

**Friedland:** So, I was the highest paid person at NASA for a short period of time. I stayed on. I was chief technologist. We did a lot of interesting things in developing some of the new sciences at NASA in nanotechnology and the like.

I ended up leaving NASA three years later because another new NASA administrator fired Scott and said there won't be any more basic research at NASA because we can't afford it, because we have to get the shuttle replacement done. Once again, NASA went from being a good place to work to being a bad place to work if you wanted to develop long-term science and technology.

**Grier:** At that point, just going over the three years, you're overseeing all research that Ames is looking at, correct?

**Friedland:** Related to that question, NASA itself had gone from a new way of managing its money. Previously, money had been in distinct categories. Money for salaries was different from money for travel, different from money for research grants. You could never cross the money. It went to a system where all money was in one pot, and your job was to manage it.

Along with that, each center had the right to put aside a certain amount of money for internal R&D, IRAD as it was called. I managed that budget, which was sort of a free pot of money to encourage people to do innovative, usually interdisciplinary things.

The first IRAD project I ever funded was sort of interesting. We got together the people who had been developing thermal materials like heat shields with the people who'd been doing nanotechnology to see if we could develop a material for a shield that would be both a heat shield and a radiation shield at the same time. I managed that. I was the center director's chief advisor on science and technology, a member of the executive council, and head of new business development at the center.

Under Scott, that was great. I decided to leave NASA when he was fired.

## FX Palo Alto

**Grier:** So, then you moved on to FX Palo Alto.

**Friedland:** FX Palo Alto is the research laboratory for Fuji Xerox.

**Grier:** That makes it the successor of Xerox PARC?

**Friedland:** No, not the successor. It was on the same campus as PARC but was a completely separate laboratory as part of the Japanese company Fuji Xerox. Fuji Xerox is not part of Xerox. Fuji Xerox is a joint venture between Fuji Film and Xerox that is owned like 70/30 between Fuji Film and Xerox. Xerox actually divested itself of it entirely later on.

I was hired to turn the laboratory into a Silicon Valley laboratory—in other words, be like a Google with lots of people arguing in the halls at all time, everyone there nights and weekends, all the usual stuff. It turned out, after a year of trying to do that, that wasn't really what the company wanted despite their saying so. They wanted more a vanity laboratory to say they had one in Silicon Valley. I realized that, and even though they were paying me a lot, life was short, and there are better things to do.

## Air Force Consulting

**Friedland:** I went from there to what I've been doing for the last 10 years.

**Grier:** Which is your own consulting firm.

**Friedland:** Well, my own consulting firm, but almost entirely for the Air Force. I had done a little bit of consulting work for NASA and for University of California, but the real fit came about because Ed Feigenbaum had been chief scientist of the Air Force as one of his tenures. One of the things he had done was help the Air Force set up its Asia-Pacific office.

The Air Force Office of Scientific Research funds about half a billion dollars of basic science worldwide. Unlike the NSF [National Science Foundation], the Air Force looks for basic research wherever it can find it in almost every country of the world, except the few you can probably guess. I really liked that idea of global science. Ed had been the chief advisor to AFOSR [Air Force Office of Scientific Research] on a bunch of these things. He was going to retire from that, and he asked me if I would be interested to do it. He introduced me to Brendan Godfrey, who was the head of AFOSR and to Ken Goretta, who was the head of the Asia office, and we hit it off.

Originally, I spent most of my time sort of as a talent scout in the Asia-Pacific region, which means leading tours of the various universities and research centers to find good places for the Air Force to invest. I also ended up helping to develop the AFOSR-wide autonomous system programs, which had sort of fallen in disrepair from the programs of about 20 years ago. I helped recreate them and became an advisor to the head of the organization and the chief scientist besides doing these scouting trips which now include Europe and South America.

One of the things that the Air Force has is a personnel problem in getting people with the right technical skills to be program officers in the foreign offices; it is very difficult to hire civilians for these jobs, so most of the program officers are PhD lieutenant colonels. But the system that assigns them to jobs only very loosely looks at their qualifications in any one technical discipline. It's like, "Who's available?"

I ended up becoming the mentor to three successive program officers for information sciences in the Asia-Pacific region and one in the South American region who weren't computer scientists. I had to teach them about AI and autonomous systems and the like. They were great. They really wanted to learn. We became great partners and ended up doing two, three, four trips a year to the various regions.

Another thing was that the chief scientist about four years ago, Chuck Matson, and I both believed that AFOSR wasn't taking into account enough what are really the long-term needs of the Air Force, that doing basic research at a mission agency, like NASA or the Air Force, was different from doing basic research at the NSF. You still want to do fundamental research. The research itself doesn't have to be tied to a specific military application. But if a four-star general comes by and asks, "Why are you doing this crazy stuff?" the program officer ought to be able to say, "If this crazy stuff succeeds, it will help enable the following class of things in the future."

  

The program officers need to know enough about the place they work. While AFOSR Program Officers never tell their university grantees, “You have to work in a military domain,” they (the Program Officers) must have the notion of why they’re doing this work in planning, doing this work in nanomaterials. It may not succeed, and that’s okay, but if it succeeds, why was it valuable to the organization?

In fact, there’s a really good book called *Third Generation Research*. Most people who read it get the wrong thing from it. It tells you how to build all these bubble graphs and things. Like if you make these bubble graphs, magically, you’ll be doing good research. But in there is one really important concept for the CTO of any large organization: Be able to understand why if long-term research that may not succeed does succeed it was worth your investment. It’s like the play, “The Producers”; the worst thing is if you invest *n* million dollars in a research project, it succeeds, and it wasn’t worth the *n* million dollars.

Chuck asked me to do a short-term study that involved talking to people in charge of real Air Force missions. What plausibly were their long-term capability needs? Tie that to technology, and tie the technology to basic research, so we had sort of a plan for doing that. It worked out well. He really liked the result of that.

Then we embarked on a one and half year more comprehensive study to do exactly that. That would also have the side effect of engaging enough people within the Air Force operations and research establishment that they would buy into the final document we produced. We finished that work a few months ago, and it’s in its final form now. But that was a very, very interesting project to do for the Air Force.

I’ve been essentially working half-time/full-time for the Air Force under various mechanisms. Every time a mechanism for hiring me works, after it’s done, somebody decides it was improper. So, my book about the Air Force is going to be subtitled, “If it ain’t broke, break it.” This isn’t just true for me. I mean, the Air Force is in a bad state for much of the organizational infrastructure. It’s like what NASA became when headquarters decided the bright idea was to have legal and financial and procurement and personnel not report to the center director, but report up the chain several levels, which meant now their job was to protect the guy several levels up from any embarrassment, as opposed to getting the mission done. Unfortunately, the Air Force in peace time is in that state. I’m not hoping for a war, but I do know in war time it reverts to get the mission done. Getting things done is aggravating because things you would think should take a week or two to do end up taking eight or nine months. That’s the frustration part about the Air Force.

The good part of the Air Force is there are a lot of awfully smart people and really interesting things to accomplish. You get a lot of satisfaction from watching something proceed up that chain.

  

I'll just give you one example of that: Nine years ago, on my very first scouting trip to Australia, I found a group called NICTA, which is an academic-government-commercial set of labs. National Information and Communications Technology Australia, it stood for. It's been renamed recently to Data61. But a group of researchers there were doing amazingly good work in complete, provable validation of operating systems kernels. They're hacker-proof if done, if the research succeeds.

The research we funded went really well, won several major Australian awards for best computer science research. We turned it into applied work. It went from AFOSR to one of the Air Force research laboratories funding them to make it a full-scale system. DARPA adopted it as part of their project for an autonomous helicopter, making it the kernel of the operating system on this autonomous helicopter, making it provably hacker-proof. The DARPA demonstration project flew, and it was great to see the evolution of that from very basic research through applied stuff to a flying helicopter in a relatively short period of time. That's the kind of thing that really excites me when we do that.

## Final Thoughts

**Grier:** Okay. This is the point where you get to add what you would like that I didn't lead you into. What subject would you feel needs to be part of this that I have not touched on?

**Friedland:** Well, let's see. I captained Stanford's national winning college bowl team. We beat Yale in the finals.

**Grier:** Excellent.

**Friedland:** I don't know. I think we've covered most of the issues. I think that the thing that I think I'm most interested in, that I really like and have liked it since being engaged in it at Princeton, is the notion of interdisciplinary work. Where is the intersection between computer science and cognitive science and things?

One of the programs I helped start at AFOSR I think is actually the very best example of that. One of the things that the military, DoD, had put at the very top of their list that we talked about a little yesterday is the notion of effective human-machine teams where the machine agents, whether they're robots or analyst computer systems or whatever, make the transition from tool to partner.

Achieving that transition is inherently a multidisciplinary process. We have a program in human-machine trust and effectiveness at AFOSR that funds grants in 15 different countries because that's inherently a multicultural issue. The way in which face-oriented societies are different than the way in which honor-oriented societies are different from the way in which dignity-oriented societies like ours, merit-based societies react. At our last program review of that, there were 15 countries represented and 12 different scientific disciplines—philosophers, computer scientists, cognitive scientists, experimental psychologists, anthropologists, sociologists, a couple others. That was just great.

Among other things I enjoyed there is I get to learn one new word per program review. The latest, my favorite one, was from the anthropologists; I learned the word hypergyny.

**Grier:** Hypergyny?

**Friedland:** You know what that is?

**Grier:** No, I don't.

**Friedland:** Hypergyny means marrying up. So, attaining social status or wealth or whatever by marrying someone upward from your status.

**Grier:** Here's the deal: in every couple, there's always one that does that.

**Friedland:** Well, ideally, both think that. I learned that from one of our anthropologists at UCLA [University of California, Los Angeles] at our last program review. That has been something that I think I've been totally fascinated by throughout my career, the notion. It's really fun to work with people outside your own narrow discipline.

When I go to an AI conference, I tend to avoid what I think are the really boring machine learning talks on, say, half a percent improvement on the back-propagation algorithm. I tend to go to those talks where I see people talking about... Oh, there was a great talk from some of the Cornell folks on work they had done with materials scientists on automated discovery of new materials. Those are the kind of things I really like helping fund and hearing about.

**Grier:** Well, thank you very much for this.

**Friedland:** Oh, sorry, I only had one other thing. I actually met my wife through the MOLGEN project. It turns out, she was a cancer biologist, one of the first group of cancer biologists that graduated from Stanford. She was part of the MOLGEN Project as a domain expert. I met her at graduate student housing, and we ended up being partners, both in life, and for a while, in research.

**Grier:** There you go. Thank you.

  

**Friedland:** Thank you.

END OF THE INTERVIEW

  