Oral History of Brad Allen

Interviewed by:
Avron Barr

Recorded May 16, 2018
Mountain View, CA

CHM Reference number: X8641.2018

Table of Contents

Avron Barr: Today is May 16, 2018. My name is Avron Barr. I'm conducting an oral history of Bradley Allen. 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're here at the Computer History Museum today, and this interview is being videoed and recorded and will be transcribed, edited, and posted on the Computer History Museum website. Let's get started.

High School and First Exposure to Programming

Barr: Tell us a little bit about your life today. Where are you, what are you doing, and what's life like?

Brad Allen: I am currently Chief Architect at Elsevier, which is a publisher of scientific and medical information. I am living in Los Angeles—in Manhattan Beach, California, to be specific—and really enjoying being part of a company that's trying to push forward what we're trying to accomplish with science and medicine and make the world a better place. Working as the chief architect means that you end up having to do a lot of things covering a lot of different topic areas and challenges. It's a good life right now.

Barr: Good. Glad to hear that. Anyway, I'm going to take you to Elsevier later, but I want to go all the way back and see how you got here.

Allen: Sure.

Barr: Where did you start out?

Allen: I grew up in the northeast of the United States, in a number of different places. We lived in Pennsylvania, New York, and Ohio.

I went to high school outside of Philadelphia and started my career as a computer programmer, if you want to put it that way, back when I was 14. Some folks who were at the Burroughs Corporation, which at that time had a significant presence in Paoli, Pennsylvania, had an Explorer troop that they convened to bring kids in and introduce them to computing technology. I started working on some of the old Burroughs equipment, and then at the high school we had access to some timesharing machines. I just started programming and really loved it and wanted to basically go in that direction as a career.

Barr: You had a lot of exposure to programming in high school.

Allen: Yes. It was very early days, but it was a good opportunity to really get into it. I had been exposed to some of the concepts and [was] really fascinated by some of the history of logic, the

original work by people like [Alan] Turing and [Kurt] Gödel and so forth and just [was] kind of captured by the idea of symbolic logic.

Barr: Even back then, even in high school?

Allen: Yes.

Barr: Did you have courses in high school?

Allen: No, not so much, but there was a math teacher at the school who had access the timesharing system that we were able to work with, so there was exposure there, for sure.

Influence of Stewart Brand's *II Cybernetic Frontiers*

Barr: Was it just computing, or did you have other academic kinds of interests back then?

Allen: I was interested in general in history and science and a lot of those kinds of areas, just generally curious, but computing is the thing that really captured my interest. There was a moment back in high school when I encountered a pivotal document for me, which was Stewart Brand's *II Cybernetic Frontiers*. I remember very clearly going to a bookstore in King of Prussia and buying that book. The first half of it was a portrait of Gregory Bateson, the cyberneticist.

Barr: Yes.

Allen: The second half was the article that Stewart Brand wrote for *Rolling Stone* on life at Xerox PARC and at the Stanford AI Lab. That was a picture of everything that I wanted to go and do. In the back of that article in that book, Stewart Brand makes the mention about where do you go to college for computer science? In there, it says the best place to go is Carnegie Mellon.

Barr: That was it? That's why you went?

Allen: That's why I went. That's why I ended up in Pittsburgh at CMU [Carnegie Mellon University].

Barr: When was that?

Allen: That was 1976, my first semester there.

Undergraduate Studies at Carnegie Mellon University

Barr: What did you study, and why?

Allen: Well, at the time, the computer science department there was a graduate program only, and those of us who'd stumbled into CMU looking to become computer programmers at the time ended up in the applied math program. There was an applied math degree with a computer science specialization; that was the track that they had at the time.

At that point in time, it was just a tremendous group of people there. It was a jumping off point for a lot of things including the stuff we'll be getting into later. The undergraduate people, of which I was one, had a great deal of exposure to what was going on with the graduate courses. It was very easy to get involved with a lot of the courses that were being taught to the folks who were coming into the Ph.D. program. It was just a wonderful environment.

Barr: Your degree was in applied math, but you were exposed to a lot of computer science courses too?

Allen: Yes.

Barr: As well as projects.

Allen: Exactly.

Barr: It does sound like a good environment. What was life like at CMU? What kind of things were you doing in terms of studies and coding and life?

Allen: Well, CMU is a place that is one where you do best if you know exactly what it is you're trying to achieve. At least it had that flavor at the time. It's very, very intense, focused work; you drove yourself as hard as you could to be able to deliver in that sort of environment.

At the same time, it was a blast because we were getting exposed to things. There were very, very few places at that time even among the elite universities where you could get access, for example, to Alto machines, to some of the very early DEC equipment that became so pivotal later on. It was a lot of fun because it was a very challenging environment, but at the same time just a bunch of people spreading their wings and looking for things that they could make a contribution.

AI Work at CMU

Barr: It was a very interesting time. I'm wondering, if you think back, were there particular people that were especially influential during your years there?

Allen: The folks that were engaged in trying to understand where to go with artificial intelligence applications were first and foremost for me. I got involved through Jaime Carbonell's Introduction to AI course. Jaime taught that and at the end of that course they gave opportunities and said, "Here's a project that we'd like people from the class to work on that is involved with some of the research that's being engaged in."

Mark Fox was a graduate student at the time but [was] very rapidly getting involved in the commercial exploitation of artificial intelligence. He had a project that he suggested that I went and contributed to and became involved with. Then that led to the work with Mark at the Intelligent Systems Laboratory.

So, Mark, Jaime, and other people in that environment like John McDermott who were involved in the early blush of expert systems were really important people that I engaged with at the time. There were the people who were just there that you could interact with.

I took a Combinatorics course from Jon Bentley, and that was in the beginning of a lot of the stuff that he later put into publications like *Programming Pearls* and the like. He was a fascinating teacher. That was a great, great class.

Peter Andrews taught logic and worked very heavily on higher-order logics and their automation. I took a number of courses there.

Garrel Pottinger, who was in the Philosophy Department actually as a logician, was somebody I spent a lot of time with. That harkening back to that early interest in logic and automated reasoning. Those were people who had an impact and an influence on me at the time.

Barr: You mentioned Mark's project that you got involved in. You were a programmer on the project?

Allen: Yes.

Barr: What was the project about?

Allen: We did a number of things in Mark's lab but the most significant one at the time was a project for Westinghouse, planning production and manufacturing of turbine machines. It was a lot of the stuff that he and Monte Zweben, for example, got into significantly as time went on around the question: How do you improve the state of job shop scheduling? How do you make it more efficient? How do you make it not [necessarily] optimal, but deal with the complexity of the task by applying heuristics and notions of search through a state space that you could define based on the constraints associated with the job shop floor? Working on those kinds of programs, some things around simulation of job shops and the like.

CMU's Intelligent Systems Lab

Allen: This was around the time of the founding of the Robotics Institute, which Mark's Intelligent Systems Lab was part of, so we were all engaged with projects that demonstrated the application of artificial intelligence to those kinds of industrial problems. Westinghouse was a founding sponsor of the Robotics Institute. That was a big source of those kinds of problems that we were tackling and the areas that we were getting into.

Barr: Do you remember anything about what tools you were using. There weren't many tools back then. What were you programming in? What was that like?

Allen: We were using Lisp, and if I'm not mistaken, it was Franz Lisp in an early version, in that vein of development on VAX 11/780s primarily. It was a great environment and one that was a tremendous amount of fun as well because we were sitting there on one of the main CMU Arpanet nodes at the time. We had our bulletin boards up on Usenet and internally. We had access to a lot of interesting and innovative software at the time that was just internal to CMU, but you could see ripples through in later days: the ZOG net, which was Alan Newell's attempt to build an early form of hypertext, and the usual kinds of Unix entertainment sort of things that were up there. We were programming away in Franz on 11/780s and later got into using the [Three Rivers] PERQs, but it was primarily DEC equipment at the time with Lisp as the main programming environment.

Barr: You were probably one of the first groups that had access to the Internet as a tool for your research and for your lives. Can you remember whether that was influential, the fact that you could connect in an instant with anybody else doing research in AI?

Allen: Yes. There weren't a lot of people in general to deal with. It made a huge impact in our ability to talk within that community and primarily within the CMU community, but it wasn't so much the kind of thing we do today on social networks, interacting with people globally with very, very little effort. It still took an effort to reach out.

It was still more of a personal relationship thing to be dealing with people at Berkeley or at ISI [Information Sciences Institute] at USC [University of Southern California], and so forth, but we used email very fluently and ran into a lot of the things that have challenged people ever since in terms of etiquette. I remember very clearly reading the threads where Scott Fahlman initially proposed the use of what ultimately has become the emoji. You know, the whole thing where people got bent out of shape at each other over a misunderstanding expressed on a bulletin board, so Scott came on and said, "You know what? We need a thing that allows us to tell people when you're kidding. Here's this thing with a colon and a dash and a parenthesis."

It was very early days, but that way to communicate and get information across back and forth was something we moved into very naturally. It was completely consistent with the experience that we have today, but it was a very early look at that. Back in 1978, there were not a lot of people who were interacting with each other over email in that way.

Barr: How long were you at CMU?

Allen: Seven years all told from in the door to leaving for California. When I joined Mark's lab, I became an employee of the university. That was a tremendous thing at the time because that took the pressure off of tuition. You could pace yourself. Living in Pittsburgh at the time was very economical. It was a great life, and it was hard to leave because it was such an amazing environment at the time.

All told, calendar-wise in my time as an undergraduate and time as an employee at Robotics, it was seven years.

Barr: Did you graduate before you started?

Allen: No. I started working I think as a sophomore, in the second semester of my sophomore year as an employee of the university, and that was great.

Barr: Yes.

Allen: Getting paid and getting a degree and working in this incredible place.

Barr: Having real work to do as well as, yes. I can see; it sounds really good. Was it a natural transition? Were other students following the same path?

Allen: Yes. I mean, there were a lot of folks that took that path where they weren't in the graduate program, they were undergraduates, but they showed ability.

At the time, and still to this day, a lot of money flowed into that environment that created a demand for people who could take on the programming tasks who weren't otherwise engaged in either getting a Ph.D. or teaching classes or being a professor. There are people who are there to this day who have basically risen in the CMU organization to positions of responsibility who I knew there as an undergraduate. We were all undergraduates working at the time.

Working at CMU’s Robotics Institute

Barr: Looking back, what do you remember about the work you did at the Robotics Institute, and how did that influence you?

Allen: A lot of it was focused on trying to take these very hard practical problems in industry and take this technology which allowed us to express information symbolically to be able to reason about that, to be able to apply search algorithms to come up with solutions that satisfied a number of constraints. Assembling all of those things but delivering them in an industrial environment was a key aspect of what it was we were trying to achieve there at Robotics.

It was different then. There was a lot of theory going on at the time. Dana Scott was there at the time, as an example of that level of work. Alan Newell and Herb Simon were still active. A lot of work on operating systems and so forth, but the work that we were engaged in was, how do you get what came to be called expert systems technology in a position to be able to deliver value in an enterprise environment? That was my first exposure, certainly, to the challenges that you engage with when you actually try to deal with it. How do you integrate with existing systems? What are the organizational dynamics that either enhance or frustrate technology transfer and adoption? You know, it was all an exposure to those sorts of things very early on.

Barr: When you came to CMU, had you had a focus on AI? Was that already in your mind as a place you wanted to go or a possibility?

Allen: Well, the thing that I had walked in there with was this very romantic notion of the power of logic and the tradition, which I still gain a lot of strength from, if that's the right term for it, of this path from the early 20th century, the time around 1930, the Golden Era of mathematical logic and that notion of automated reasoning being brought forward to be able to create if you want to call them thinking machines. Using that terminology of the time, the romance of that was the thing that I was interested in. In fact, I remember Simon and Newell's general problem-solving book as something that I got before I actually came to Pittsburgh.

Barr: Really?

Allen: Yes. Now I can't say that I necessarily understood every nuance of that, but that was the thing that I had in mind that that attracted me to that environment.

Barr: Who did you work with most closely at the Robotics Institute?

Allen: Well, with Mark, certainly. Jaime Carbonell and folks who were working with him on some very early kinds of natural language processing tools that came into play. Where we physically worked, we were co-located with a number of folks on the robotics end, Hans Moravec and his lab, David Bourne.

We ran into a lot of folks in the hallway worrying about making robots move around without hitting things and that sort of stuff. Very early days on that. That's about 1982 to 1984. So, there were a lot of folks. Jim Crowley—I believe if I'm not mispronouncing his name—was working on robotics problems at the time.

I have a very clear recollection of Marc Raibert's earliest work on locomotion. You know, every day I would go down, and at the time it was called Wean Hall. It's now Newell-Simon Hall, if I'm not mistaken. He had his Leg Lab, so sitting in the middle of this concrete and cinder block room was this mechanical pogo stick hanging from the ceiling that would jump up and down. And if you look at the videos now that Boston Dynamics, for example, shows of the things that they're building there's a direct lineage from that Leg Lab and that work that Marc was doing to the kinds of amazing stuff that those folks are doing today. So those were the people in the environment at Robotics.

Barr: As you look back, this was your extended undergraduate career.

Allen: Yes.

Research Spirit at CMU in the 1980s

Barr: It was unusual, I think, in the degree to which you were involved in research and especially programming. Did that work? Do you think that was unusual, and did that work as a way to get involved in artificial intelligence?

Allen: Well, it worked for me brilliantly. There were other folks that were there with me at the time. Mark Wright, Gary Strohm, Paul Haley, all of whom I subsequently did a fair amount of work within the expert systems days. We all came along that progression, and it was a tremendous introduction to real problems at the bleeding edge in an environment that was extremely supportive, not competitive or judgmental in a way that some organizations are. It was just, “We're all in here trying to figure out how to use this cool stuff to do cool things.” I see a lot of antecedent in that to some of the things people are trying to accomplish now with getting younger folks to become coders and be introduced to computers and so forth. It was just it was a very natural thing. I think it was an alignment of purposes. The folks who had raised money to fund research needed people to actually implement things, and those of us who were implementing them got a great deal of value out of that.

Barr: There was a lot of stuff going on in AI at other places at that time. Partly because of the Internet and email and partly because of just going to conferences together, there was a lot of cross-seeding of things.

Allen: Yes.

Barr: As you look back to those years at CMU, what was really influential? What were the highlights of the intellectual progress during your years there?

Allen: The realization that people were beginning to build systems that actually touched real-world problems, beginning to spill out into practical applications, was a big thing. Up until that point we were working, or people were working, on toy world problems or games. In that environment a lot was going on with work that continues to this day: on solving, or not solving in the true technical sense, but getting computers to play effective games like Othello. Paul Rosenbloom did that work at the time that I was around CMU. Hans Berliner was there.

Those were all interesting AI challenge problems, but the breakthrough into being able to build things and deliver them to business was pretty impactful, at least from the perspective that I had. And there was a lot of excitement in and around that. A lot of that was generated by things that were goals and initiatives driven out of the federal government or the military to get to this vision of autonomous systems that were able to work in a variety of challenging environments. That created a lot of work in places like MIT [Massachusetts Institute of Technology] and out here on the West Coast, where you were building environments to work in that way which is the birth of the Lisp machine and Lisp as an industrial programming environment. So, a lot of that was going on. That whole issue of “We're about to take this stuff out of the laboratory and bring it into practical use” was a big thing at the time.

Barr: Yes. It shaped a lot of what went on. Do you have a favorite project from that period?

Allen: I think working on the system that Mark particularly pioneered, the ISIS system; I forget the acronym exactly, but it was the scheduling system that brought forward some concepts that were pioneered originally in some of the speech understanding systems. Things like beam search as an algorithm, and looking for ways to be able to make it easy to come up with a satisficing solution for a job shop scheduling problem by applying that kind of technique, going in and implementing that system, building some of the visualization of the exploration of the state space with some very primitive kinds of graphics capabilities that we had at the time. That was the project that crystalized not only for Mark where he was going in terms of his research, but again this notion that we can build some sophisticated and useful systems for industry. I particularly enjoyed my contribution to that project.

Computing Curriculum at CMU

Barr: Well, I have one more question about CMU, and it's more or less about the academic path that you took there. You were in applied math as a department, but clearly, you were exposed to the computer science curriculum and the AI curriculum.

Allen: Yes.

Barr: Anything stand out there as particularly good or particularly problematic in the way that you came into the field?

Allen: Well, it was funny because the Mathematics Department, which at the time was headed by George Moore, was focused on problems around differential equations and that sort of thing. It was a bunch of old “diff eq” guys, and the idea of having to deal with a horde of aspiring hackers was not exactly what they were focused on. But that didn't turn into any kind of barrier or impediment.

The people who wanted to gain access to what was happening in the computer science part of the curriculum had no problem doing so because as a natural part of the track, we were meeting the professors, we were meeting the TAs [teaching assistants] at the time. People like Rich Korf, now at UCLA, was my TA for the Introduction to Software Engineering course, which at the time was taught by Elaine Kant and Guy Steele—just tremendous folks. Very quickly, because of that track, you got that exposure. If you had any desire to get involved at all, you were welcomed into that environment.

I ended up taking something on the order of nine semesters of logic, including four graduate seminar kinds of things where I was given the opportunity to explore a wide raft of different topics and work on things like decidability problems and the like. I had an opportunity to present in graduate seminars, summarizing some of the recent results in that area.

There's somebody I should mention: Garrel Pottinger, who I brought up earlier, was just tremendously impactful for me. And Garrett Birkhoff, was an incredible mathematician. He was the son of George Birkhoff himself, who was one of the leading lights of mathematics in early 20th century America, and [Garrett] was a buddy of John von Neumann. At the time, he was semi-retired, but he was teaching at CMU a course in abstract algebra.

Having the opportunity to engage with people like that, with that history... He was tremendously supportive of just the interest in the work. It was a great environment. There was really no impediment to somebody who was motivated to go out and gain that knowledge. You were there in an environment where you could get it.

Leaving CMU for the West Coast

Barr: That leads to the next question. Why'd you leave?

Allen: It was time. I think the other thing to be said about my undergraduate career was I was not necessarily the best of students from a traditional perspective. I gained a lot out of it, but it wasn't clear that there was a real graduate school path for me at that time. After I had been

there for a number of years at Robotics, it was clearly time to figure out what the next step would be.

That happened in 1983, just about the time that the real venture interest in expert systems companies began to blossom. The opportunity presented itself based on the fact that a lot of money was starting to go into that sector. Some of the first wave of expert systems companies started to happen, and they needed people who were knowledge engineers. Outside of the communities at Stanford, and to some extent at MIT, there were no people who could look people in the eye and say, "Yes, I've built that kind of expert system for an industrial application." You know, they couldn't say that, but I could and some of the folks that I worked with at CMU could as well.

All of a sudden, an opportunity presented itself in the context of that the emergence of the first generation of expert systems companies.

Barr: What did it look like as you were ready to leave? Did you look around at possible places to go, or did somebody come and get you?

Allen: We were headhunted. But at the same time, myself, Mark Wright, Paul Haley, a number of others in the environment at CMU had been thinking about what would it take to start a company. What would a company look like? We had serious discussions about the questions: Could you build a company around a language? What would that look like and so forth?

We had a number of conversations with people like Rick Hayes-Roth when he came through, just exploring that idea. In the absence of external people looking for people to come join them, we probably would have ended up doing something on our own in that vein.

But people came calling who had just raised rounds and were going to build out the companies that are going to solve this problem for industry. That was an easy way to get started into it.

Barr: One of those early ventures was at CMU.

Allen: Yes.

Barr: Carnegie Group.

Allen: Yes.

Barr: Is that after you left, or was that something you considered?

Allen: I considered it. I got an offer; very clearly, I remember the letter from John McDermott. I said I would come and then within the week I got an offer I couldn't refuse from Inference Corporation. It's not something I'm necessarily terribly proud of and Mark Fox has mentioned that he's still mad at me about that. But I knew everybody who was involved in the initial work and founding and getting Carnegie Group off the ground. I remember very clearly the conversations associated with that while Mark and the rest of those folks were trying to form that. That was tempting, but at that time I had been seven years in Pittsburgh and the West Coast beckoned for all sorts of different reasons.

Barr: Do you mind if I ask you what some of those reasons were?

Allen: Sure. One is the very simple attraction of being in California. It's laughable now even to think about it, but at one point in time while I was at CMU, Tim Leary came through and was doing his standup philosopher thing at the time. He gave a talk, and [in] part of that talk [he] said something to the effect that Western civilization is going faster and faster in a westward direction and it's just going to hit the California coast and go into space. That stuck with me: I really maybe want to end up in California.

I had had an opportunity to come out and visit a couple of times. Mark Fox and I came out to a workshop on agents that USC ISI ran out in Idyllwild. It was my first introduction to Los Angeles, and it was just an amazing, beautiful place. It was very attractive from that point of view. I remember coming out and interviewing for Inference Corporation and ending up at LAX and walking out <laughs> having just gotten off a plane from a very dreary and snowy Pittsburgh and there it was: the clear blue California sky. So, there was that.

I had family who had recently located to Los Angeles, and that was part of it. But it was the sense of adventure, and California being a place to go. A lot of people that I knew had gone to ISI as part of a very close relationship I think that had been engendered between CMU and ISI at the time. There were people out here saying, "Come on in, the water's fine." It was a very natural step from that perspective.

Barr: The computer science community was still pretty small then, so everybody knew everybody pretty much.

Allen: Yes. It was an interesting environment here in Los Angeles. At the time, Bob Balzer was the lead guy at USC ISI, and he collected quite an interesting group of people. They were a significant cornerstone of what happened with the emergence of the Internet and so forth. There were some folks doing work at the main campus at USC and UCLA [University of California. Los Angeles]. Not much in the commercial environment. There was a scattering of tech companies, but you're referring essentially to the computer science community as a whole, right? We were all just basically a few scattered nodes on the original Arpanet, and people knew each other.

Move to Inference

Barr: When you got to Inference, what did you start working on and why?

Allen: Well it was an interesting environment. Chuck Williams and Alex Jacobson had been working for a number of years already.

Barr: Really?

Allen: Yes. Alex came out of Hughes Research, and Chuck was out of the USC ISI group. They had hooked up to build a natural-language understanding system, and Chuck was building that on top of Lisp. One of the challenges they had was, "We need an environment to make this run fast and be a thing that will work beautifully and deploy it."

They managed to get involved in the original founding of Lisp Machines, Inc. with Rich Greenblatt and were engaged in the series of events that happened between LMI and Symbolics at the time. They also did some work with Stephen Wolfram on a very early version of what ultimately became Mathematica, which at the time was called SMP. They had been working for a number of years I think starting in 1978, and Alex had managed to attract an investment from Venrock partner Tony Sun at the time and was part of this initial set of investments that the VC [venture capitalist] community made in artificial intelligence companies.

Chuck, who had been wrangling this by himself and a few others, was in a position where he needed to move quickly in terms of getting people, and it was a struggle because Los Angeles was not Palo Alto. It was not near any of the traditional AI centers. Chuck was working with a headhunter at the time. Chuck was just working himself to death and got very sick and run down and was on the phone emoting to this headhunter about, "I'm just not getting any traction," and so forth. The headhunter said, "Well, just give them offers they can't refuse." That was the thing that basically hooked me to come out and start with them.

When I got to Inference, there was the work that Chuck had done on the LMI machines. He had built a system that ultimately was the system that became ART, which had some interesting and innovative elements, principally around what in the subsequent years came to be known as assumption-based truth maintenance. It was the ability to have a way to be able to manage sets of assertions and reason in those sets of assertions independently so that you could build a way of doing hypothetical reasoning in a very structured and disciplined way. It was brilliant work, and it was very efficient and fast. It was Chuck's baby, and he had a lot personally invested in that.

Inference Automated Reasoning Tool (ART)

Allen: When I came from CMU, the way that I showed up attracted a lot of interest from some of the other folks that were there: Mark Wright and Paul Haley. It was like "This is a great opportunity," so they came out in short order. We came there from CMU with a perspective that was borne out of the Digital Equipment work on XCON, using production systems and the work that we had been doing in Mark Fox's lab, which was very much in this production systems model of building rule-based reasoners. A lot of the focus was on speed and the ability to do that. Some of the work that Mark Wright and I had done very early on in the 1982, 1983 timeframe on hybridizing rule-based systems and what we called schema systems, which were essentially frame-based or object-based knowledge representation systems, we brought all that out there and worked with Chuck to mold that together into a cohesive whole. It was essentially this big toolkit of things that you could apply to a problem to see if you could make some progress on that.

We showed up in the spring of 1984. I got there in February 1984. I think Mark and Paul ended up there a few months later, and the goal was to build something that we could demonstrate at the AAAI Conference in Austin in August.

Barr: In August, wow.

Allen: We ended up doing the classic eating and sleeping in the office thing for that period of time and just moved heaven and earth to build what we demonstrated in Texas as ART, the Automated Reasoning Tool. We brought a lot of ideas together and worked very, very quickly in a collaborative way to synthesize this whole, and it was an interesting time.

Barr: How many of you were there besides the four you've mentioned? Was it a big group or a small group?

Allen: There were maybe about a dozen people, all told, some of whom were there before we showed up. A couple folks from that era. Rich Schroeppel [who] is, I believe, currently at University of Arizona, was there and was an amazing guy to work with because of his just incredible ability as a mathematician particularly as it related to encryption. He was a coauthor on the MIT HACKMEM document. We had another fellow from Aerospace Corporation who joined us, Charles Kalme; his claim to fame in our world was he was the last American to defeat Bobby Fischer in a game of chess.

There was a whole range of individuals there that were very, very talented, but we were all told about 12, 15 people. The core group was Chuck, who was dealing with the truth maintenance system; Paul, who implemented the rule system; Mark, who implemented the schema system; and I did my best on the interactive development environment that was part of the shell. That was the team.

Commercial AI in the 1980s

Barr: As you recall, from Inference, looking around at this AI world that was just developing, this commercial AI one, what struck you? What were you thinking as you looked at your competitors and at the marketplace about how Inference would proceed?

Allen: Well the thing about Inference was, and I think through its existence, until it found its niche in the 1990s, was we always had a chip on our shoulder because we weren't one of the companies out of Stanford or so forth. There was this competitive feeling. It's like, "We're going to show those guys up at Stanford what we can do."

I think there was a mutual feeling of competitiveness amongst all those different companies, so it was an environment where once we released ART, the demand from the federal government or people who were being financed by the federal government for the software was enormous. It was of the order of, "We're going to buy every version of everything that we can get our hands on just to figure out what to do with it and how it works." So, it's just this tremendous early adopter demand.

We released this stuff in August, or launched it is probably the better way to say it, and immediately we were doing about $1 million a month in revenue, just on the basis of selling this, selling training associated with it, starting to do some of the early expert system projects under the aegis of Inference and so forth. That experience was being recapitulated pretty much everywhere amongst those early companies, so there was this feeling of, "Wow. There's something here" and "We're riding a wave." Just tremendous excitement associated with that.

Barr: Somebody once described that wave as every Fortune 500 R&D lab buying one of everything.

Allen: Yes. Right, that was a signal from the marketplace that I think, certainly in retrospect, we took the wrong way.

Barr: We all did.

Allen: But it lasted for a while, lasted long enough for us to garner a significant investment from Ford. It led to some real growth in the company, and it led to a lot of interesting projects. The early work we did with the NASA Johnson Space Center, the work we did with American Express and the Authorizer's Assistant, these were significant systems that we were able to build in the subsequent two years. The period between 1984 and 1986 were some of the first successes that the expert systems people claim in systems that were automating decision-making [in] one way or another. That was being recapitulated across the industry.

The other thing too, I think, that's worth noting is people were trying to differentiate on the basis of features of languages and tool capabilities that were pretty arcane, certainly in hindsight. At the time we thought, "Well, you can't really be serious unless you have this particular feature," or "Here's a new way of doing stuff that just is simply not achievable in this other way."

A lot of these systems brought together a whole mass of techniques, some of which were well integrated and some of which were not. They were presented in this jumble, and it was left up to the people who were implementing the systems to pick and choose what was effective and what was not.

That process of discovery about what was material in terms of achieving results was mixed into that whole thing. As people piled into the industry to take advantage of the market that had appeared, there was a lot of confusion about what you needed and a lot of different ways of talking about architecting systems that created a lot of confusion and ideological rigidity in terms of how people position[ed] themselves. You know, they're production system people or they're model-based reasoning people later and the like. They were all differences that in the long run weren't necessarily as meaningful as people thought they would be.

ART Implementation and Applications

Barr: As you looked at the competitors, did you think that ART in particular was an advantage? ART was sort of your ticket, your product.

Allen: Sure.

Barr: I'm just trying to get a feeling for how that went for you.

Allen: Well, I think we made a lot of the fact that we had taken a lot of the ideas on the efficient implementation of production-based or rule-based systems and built a nice implementation of that. Part of what we were saying is that you can reason quickly in a forward-chaining direction. Part of it was the ability to reason hypothetically that came from Chuck's work, which was pretty unique. None of the other expert system shells did that, and that was a big differentiator for us. There were a bunch of other features that you could add into that, but those are the two that were the hallmarks of what we would position as advantages when we went in to compete against people like IntelliCorp and Teknowledge, and the like.

I think in hindsight, and I know from speaking with Doug Lenat that he would say there was a lot of sweet technology that turned out not necessarily to be either easy to use or, again, something that made a difference in terms of people's ability to implement actual applications. But what we had in mind is we had a really fast rule engine, and we had the ability to do hypothetical reasoning.

Barr: I gather you were involved both in designing and building ART and in the applications of it, some of the applications anyway. Thinking back, what strikes you as the most interesting work that you did as in the application side?

Allen: Well, for the first year and a half there, I was pretty tied up with the tool itself. There were others that we brought into the environment: Mark Malitz who worked on the first system we did for NASA, and Laurent Piketty, who was involved with the Authorizer's Assistant for American Express. We billed out those folks, brought them in. They're the ones who were the feet on the ground making the expert system stuff happen.

I came in later as the other folks took over the evolution of the tool. Did some work with NASA on building systems to help manage software libraries and did software parts composition and things like that. Did some work for DARPA-funded projects on expert systems for detection and recognition of Test Ban Treaty violations and things like that. It was a smattering of a lot of different kinds of simple expert systems work, but a lot of what stood out for me were those first initial systems. Just the satisfaction of having engaged with organizations like NASA and American Express or American Airlines and built these things out and delivered them, and working with the teams to support them, because the tool itself was something that was more or less used as is.

A lot of the challenges that we had with those systems were around integration. For American Express we had to get that system, which was at the time running on Symbolics machines, integrated with the mainframe-based systems that their authorizers were using to be able to make credit decisions. There was a lot of baling wire and chewing gum and serial ports involved to make that work. Just being there and supporting those guys was a big part of what we were doing as the central team.

Enterprise AI Systems

Barr: As the years went on, things changed, Inference changed and ART changed, and the marketplace changed. What do you think was, as you look back from those early years of everybody needed AI and there was a plethora of projects to choose from? How did that change over the subsequent five or seven years?

Allen: Part of it was an evolution of what was the acceptable set of techniques and tools that you needed to build an expert system. The things that we built were very expensive, and they ran on expensive hardware at the time that was not the kind of hardware that anybody in a traditional IT organization had any experience with before. You know, this is predating the PC, and to some extent, a lot of our systems were birthed out of the minicomputer environment.

Barr: They predated Sun workstations too, for that matter.

Allen: The VAX. They did. I mean, very early. They were just starting to emerge, but they hadn't achieved significant traction in traditional enterprise use. Traditional enterprise was dominated by the mainframe. As we worked our way through the initial bubble of business that came from the early adopters that were being driven by the government, we started to run into the scars on my back sort of IT guys who were saying, "Well, how am I going to do this? You're not touching my mainframe." That was a significant wall to hit.

There was a lot of industrial R&D and people who were willing to work with the newer tools and the newer platforms to do certain things, but over time it became clear and certainly by 1988, 1987, 1988, 1989 the initial wave of enthusiasm had died off because people were looking for something easier and cheaper. By easier I mean we were still dominated by the notion of the knowledge engineer as an expert who was going to be the person to come in and know how to use the tools to be able to solve this decision-making problem.

That was a very difficult thing to do because people were coming up to speed [and] constantly changing and innovating slew of ways of building expert systems as well as trying to understand how to do the knowledge elicitation task and trying to get things pulled together in a way that would then be easy to transfer into an organization. That was a problem that was difficult to solve. When you add it to the pressure to try to understand how to translate the technology into a way that mainframe enterprise computing would be able to absorb, a lot of the original companies hit a wall.

ART-IM

Allen: The thing that was interesting in the Inference experience was one of our initial flagship customers: the NASA Johnson Space Center. The different centers had different tool sets that they had championed over time. They said, "Look. We want to build these kinds of systems, but we don't want to have this hardware, and we don't want to pay your license fee," so they basically reverse engineered the rule language in ART and released their own implementation, effectively as open source.

At the time they were a very important customer for us and we didn't want to go to battle over that, so we let that happen, without very much comment. Then ultimately took the code base for that system that those folks down at Johnson had built and made that into a new version of ART. Called that ART-IM. We were able to go to a lower price point to be able to run on more traditional hardware and so forth.

Barr: But its functionality was only a subset of the original ART?

Allen: Right. The hypothetical reasoning that was gone, other aspects of it were limited. It wasn't the rich, interactive environment that you had with Lisp machines and so forth, but it was

good enough to build rule-based expert systems and do that reasonably well, and a lot of other people followed the same trajectory.

Turning Point in the AI and Knowledge Engineering Industry

Allen: When Neuron Data came out with a combination of very simple, very focused, beautifully done on the Mac at a price point that was several orders of magnitude lower than the initial stakes of investment for somebody just getting into it in 1984, that was a huge turning point for the traditional expert systems companies. We all had to go in that direction. It was three things. The difficulty of doing technology transfer into a traditional computing environment. The collapse of the price point, so to speak, and then the remaining challenges in articulating how somebody could use the technology effectively and getting that into a broader and broader set of people as common programming practice. That created an environment where in 1988 and 1989 people were looking around and saying, “How are we going to progress the companies? How are we going to grow and build business moving forward?”

Barr: That wasn’t just Inference. That was across AI companies.

Allen: Right. If you look back at places like Syntelligence and so forth, those people were starting to get to a place where it wasn’t about building a tool and then solving all problems. It was about maybe building a tool but focusing on a given application and building differentiation there. In retrospect, that’s a challenge that every software company faces these days. That has become much more professionalized, I would say.

You know, the notions of lean startup practice and MVPs and product market fit, that’s all a way of talking about the challenge that we faced. We’d come out with an idea. We had a certain amount of success initially. That turned out to be illusory with regard to a sustainable growing market, and companies were forced to try to understand how to pivot to things where they were delivering solutions that made sense to people.

That was the challenge that Inference was faced with and at IntelliCorp, Teknowledge and all the first generation of companies and some of the later players too. A lot of them, although they were building simpler tools that ran on easier to ingest technology stacks. If they had come at it from perspective of, “We’ll have a general tool and then other people will be able to apply it,” [they] were facing that same challenge.

Barr: Yes. As you say, everybody was having this problem with the original conception of expert systems technology. How did Inference deal with it?

Allen: Inference went into a period of introspection and exploration of different alternatives, and I took it upon myself to try to understand how to tackle the problem from the point of view of,

"How do you make knowledge engineering itself much, much easier? How do you build expertise into a system without a lot of the things that we had run into?" The kinds of things that we talked about in the early expert systems days were about, you interview the expert and you get a bunch of rules and then you go and do that, just put it in the engine. That was great for certain very small, narrow things that weren't very complicated.

One of the things that we had fallen into the trap of, because of our specific technological path was rule-based production systems —firing off of activations based on the state of the database and so forth —turned out to be really hard to understand how to control. You would get all sorts of unintended things happening, which if you took a much simpler approach to the things that emerged later with machine learning and decision trees and the like, you didn't have to worry about that. But at the time we were going, "It requires expertise to understand how to wrangle rules and put it all in."

Transition to Customer-Support Systems

Allen: So, I wanted to understand how to do that easier and better. At the time, work that was coming out of, for example, people associated with Roger Schank at Yale, Kris Hammond at [Northwestern] University in Chicago, and so forth. They were beginning to explore a technology called case-based reasoning, which was more around basically describing a problem and a solution. Then you make the problem you're facing something that triggers the retrieval of previously experienced problem solution pairs, and you bring a solution to the decision-maker as a recommendation. As we look back at it from the perspective of our times, it's a search problem, but search engines were not a thing at the time. That wasn't a natural solution.

What I said was there were two things that were happening. One was that sounds like an intrinsically better way to build a system for helping people make decisions, and we had recognized the fact that there was an opportunity in customer support organizations to help codify that kind of problem-solving knowledge and deliver it to people at that point—again, pre-web—who were manning telephone banks to be able to answer questions being posed from customers.

That led to the creation of a product called CBR Express, which was focused on that. It's like building case bases to help telephone support people initially and then ultimately people were working on products and services and supporting them in the field. That was the thing that we went off and explored, and that grew. We started to market that. We brought that out in the 1990, 1991 timeframe, and that showed real promise.

At that point in time, Alex Jacobson had been moved aside by the board, and they had brought in Pete Tierney, who was an Oracle veteran, to helm the company. Pete was supportive of that. There was a whole part of the company that was still very much wrapped around enhancing and developing the ART expert systems shell, which at that point in time had grown into a product

called ART Enterprise. But I think that Pete and the board saw the real opportunity was being manifested in the context of delivering these customer-support tools on the top of systems like CBR Express.

Transition into PC and Networked Systems

Allen: Simultaneously, companies like Clarify and Vantive up here in the Bay Area had begun to emerge focused on the problem of customer relationship management. This was a significant component of that, so we began to feel our way through to a use case in an application where the technology wasn't that hard for people to learn how to use. They can integrate it fairly easily, because at this point in time we're starting to work on top of PCs and networks within enterprises that were free of the constraints of the mainframe world. It began to grow, and that was the pivot that Inference began at that time.

Barr: As that split evolved, it had implications for the company; it had implications for you too.

Allen: Yes.

Barr: How did your own work change over the course of the evolution of CBR Express?

Allen: Part of it was crystalizing the notion, evangelizing it within the company, doing the early work and getting the early version out, helping build the team that brought that to market. Then it was beginning to take a look at the other kinds of opportunities that became clear, which were in this retrieval engine space.

But the company's direction began to diverge from my own personal interests around that time because I had been there at this point from 1984. This was starting to get into 1993, 1994. There was an interesting period where we launched CBR Express, and it caught the eye of folks up at Microsoft. They said, "We're not so much interested in CBR Express for our help desk, but we are interested in embedding it into our future operating systems work."

I spent about a year and a half working very closely with folks up at Redmond in the context of the Cairo operating system project, which was one of the versions of NT that didn't necessarily come to fruition the way that they had hoped. But a lot of that was involved in trying to say, "Okay. Well, here's a thing that you could bring to bear that would allow you to be able to pose questions and get solutions that you could embed into an operating system." You know, have that be part of the functionality of what was intended to be the Cairo version of NT.

Barr: Is this the help system for Cairo?

Allen: Yes. Not so much the help system, but what it ended up morphing into was an approach to doing information retrieval as part of the operating system. All of that ended up on the cutting-room floor at Redmond, but there was a long period of time where we were working very, very closely, spending time up in Redmond, engaging with folks like Paul Moritz and Jim Allchin and ultimately having a meeting with Bill [Gates] in the board room and that kind of a thing.

We ended up actually cutting a deal with them for exclusive use of the technology. That was a big milestone for us because it was a real validation of what it was we were trying to accomplish. It ultimately led nowhere in terms of what it was that they shipped, but it had a bit of an impact because I think it opened up the eyes of folks up there in Redmond with regard to retrieval and its use, not just for customer support but as a capability in and of itself in software.

A lot of those ideas were beginning to evolve and formalize at the time. But around 1994, we started to see some of the first glimmerings of the web as we came to know and love it. I was very fascinated with hypertext technology, hypermedia technology, and had been following it up until that point, until Tim Berners-Lee had the brilliant insight that you don't want to force people to maintain the consistency of linking backward and forward: just let people go and write it and link it and let the network go from there. When the web appeared, it became clear that this was a tremendous way to be able to deliver that decision-making functionality.

Alex Jacobson, who had become chairman, was kicking around for something to do and I began talking to him and we started talking about, "This is a thing we'd like to proceed and do." The main company was about to spin out ART Enterprise as its own company, Brightware, and focus on CBR Express. Essentially the board gave us a hunting license to go off and explore to see what we would do with next. That's where we began to lay a path out of Inference and into the next stage of things.

Limbex Spin-Off

Barr: So, a big transition at Inference at the time. As you say, the ART Enterprise group split off to a company. CBR Express had its own life. What was life at Inference like? Was it trouble?

Allen: By that time, we had spent close to 15 years in business. After the initial boom in the mid-1980s, we'd gone through a few phases of waxing and waning —you know, some layoffs, some growth, and so forth. People were veterans, and they had been there a fair piece of time.

I think people were still engaged and the like, but it was a much more mature company than it was in the mid-1980s and the boom days. People were encouraged by the fact that we had found a thing to pivot to that had real opportunity. The sales organization was being productive in that regard. Pete saw a lot of opportunity there.

The ART side of the company was, I don't know... It's not fair to characterize it as tilting at windmills, but there was a feeling of frustration that they hadn't broken through with what they saw as the promise of those kinds of technologies and promises, the concepts. But it was an okay environment at that point in time.

Barr: You were about ready to start your next company around that time, with Alex Jacobson, I gather?

Allen: Yes.

Barr: What was the point of transition? Did the Inference board decide that it was time for you to leave, or when did you decide?

Allen: This was in 1994. We had been having these conversations, and we were pretty open about it. Pete and the board were fairly supportive. They gave us, as I said, a hunting license to go and see if there was an opportunity there. Alex and I went out and talked to a lot of people.

In the process of talking to people we ran into Howard Morgan, who at the time was a board member of Quarterdeck. Quarterdeck had just brought on a CEO by the name of Gaston Bastiens, who had come from Philips originally. He had been the head of the Newton project at Apple and had ended up at Quarterdeck and was focused on taking Quarterdeck and turning it into a company that was an Internet company.

At the time, the notion of an Internet company was something that was relatively undefined, and Quarterdeck thought of Internet software as things that you basically sold just like antivirus software. You put them in boxes and shrink wrapped them, and you had end caps at stores like Fry's and so forth. They were looking for product.

What happened was the things that we were talking about in the concepts that were evolving were around this question: How do you make retrieval better on the web? At the time the search engines of the web were things like WebCrawler, Lycos, and the like, which were all interesting but not operating at the level of capability and scale that ultimately in a few years Google would. I came up with an idea: Wouldn't it be great to build a desktop assistant to help you manage your search queries and repeat them and show you new information and have your search assistant environment? We proposed that, and Quarterdeck said, "We'd love to fund that."

We did a deal where Quarterdeck put up the money, and Inference took a piece of the company that Alex and I founded to do that. Quarterdeck funded it. Again, we went off on a five-month sprint to get not this time to AAAI show but to COMDEX with a product and that was the birth of a product called WebCompass.

Barr: At a company called Limbex.

Allen: Limbex. It was completely incubated by Quarterdeck. We were working very closely with the Quarterdeck people in terms of the launch and the packaging and positioning and marketing. We took it to Comdex in Las Vegas.

WebCompass Development and Sale

Barr: Who came from the group?

Allen: From Inference?

Barr: Yes.

Allen: It was just me and Alex. I hired a few folks over the Internet at the time: Brian Ulicny, who's now at Thomson Reuters, had just come out of MIT; Jay Nelson, who is a tremendous professional working in the LA area; Kristina Lerman, who's a professor at USC ISI now, and Linda Rudell-Betts was a librarian because we thought that was an important piece of capability.

We all pulled together and built this thing called WebCompass. Took it to Comdex and in just a confluence of weirdness ended up getting *Byte Magazine*'s Best of Show Award. It was not necessarily a tremendous predictor of future success, but it was a great honor.

It was being of the moment, something different but of the Internet and intriguing, and it had one interesting ultimate impact, although it ended up not necessarily being in sync with where the Internet went and evolved in terms of where that capability of search assistance and the learning sat, whether it was on the server or on the PC and the like.

But Jon Udell, who was a columnist at *Byte* and was one of the people involved in the decision-making process associated with the award, and I had a conversation. One of the things I remarked to him was that one way you can think of building systems on the web was using a concept which originally had been pioneered by some folks at CMU and out here on the West Coast of a remote procedure calls. When you make a request of a website, you can view that as a function call and the results coming back is the response to the function call. You can actually build and compose systems that way.

Jon noted that and he mentioned that to Tim O'Reilly apparently—I learned about this years ago, years later in a blog post that Jon had put up. Tim said, "That's a great idea." Ultimately, his whole championing of Web 2.0 was birthed out of that initial insight. I can't really lay claim to that concept, but it was the first instantiation of how you build things on the web which compose systems in a way where you're using HTTP calls to be able to glue stuff together.

WebCompass was fun, but it was very quick. It was like 15 months from the start to the point where we had an exit with a purchase from Quarterdeck. When we got to COMDEX and we won the award, Gaston turned right around and said, “We have to own you guys,” but that conversation took a long time to come to fruition for the usual reasons around tax events and the like. Ultimately, that was very successful. It was very, very quick.

Founding and Exiting TriVida

Barr: Was there any AI there? I mean you’ve been working in AI for so long. Did you bring some with you?

Allen: Not so much in that system. The one that we did immediately after that, though, at TriVida [did]. Once we had had the success with Quarterdeck and the exit, Alex and I turned around and said, “Okay, let's do another one.”

Barr: Did you consider yourself a serial entrepreneur by then? Was it part of your consciousness that you were going to be starting companies?

Allen: By the time I got through the second one, yes. But at that point in time it just seemed like the natural thing to do. Right? Plenty of ideas, everybody has them.

That was a system which was coming back to things that I had learned about machine learning in the 1990s and, in particular, technologies which today people would talk about as neural networks but were basically single-layer linear functions that are trained in a way where they would learn online. This is out of work by people like Littlestone and Warmuth and online learning with algorithms like Winnow and so forth to be precise about that. The idea was let's wrap one of those kinds of engines in a web server, expose an API to that, and then allow people to be able to tackle problems using machine learning.

We had learned the lesson at Inference that we needed to have an application to go after. After some initial exploration, that system was turned into one of, as I understand it, the first applications that focused on cross-sell/upsell personalization on websites. We ran these servers that would learn from data that people would provide us by instrumenting their websites with cookies so we could track user behavior in terms of what pages were being looked at together, build the models, make the predictions, return those as responses, and integrate that into them.

This is all in the 1997, 1998 timeframe. Had some initial experiences with people like Barnes & Noble and the like. It was an example of a personalization software as a service very early on. Ended up in an exit in 2001 with a company called Be Free back east.

Barr: How many people were at TriVida?

Allen: I think at TriVida we got to about 40 before we exited. The exit was in February 2000, which was a propitious moment because that was the very moment the initial web market peaked. We got out through the open window with the sash hitting our heels as we got out the other side. It was a pretty crazy time. Immediately after that, we were part of Be Free, which was an Internet marketing company.

Barr: I don't remember that. Where were they out of?

Allen: They were out of Framingham, Massachusetts. Sam and Tom Gerace, a pair of brothers had built that company. They've since morphed into a number of companies that came out here to the West Coast, essentially selling affiliate marketing services and the like. They were assembling a suite of technologies to cover that space.

We were part of that group for a couple of years until the market began to pinch and they had to reduce headcount and cost. They said, “We can't have an LA office anymore,” so I and our technical people were out the door.

We went down to [the] Typhoon restaurant in Santa Monica and had a drink. I looked around at my best two technical guys, Cormac Twomey and John Jensen, and said, “Let's do this again.” That was the point where I was probably characterizing myself as a serial entrepreneur. In retrospect, that was a decision I should have made more prudently.

Founding and Selling Siderean Software

Barr: Why?

Allen: Because the third company that I did, Siderean Software, exploited a lot of what I had been seeing as opportunities from the earlier expert systems days and this idea of making retrieval better and more effective. It brought in a lot of technology from the semantic web community into building a semantic search platform. That just was the experience that, I think, many serial entrepreneurs go through where you think you know what you're doing, but ultimately, you find out you were simply luckier than most as opposed to better than most. That was a very chastening experience.

Siderean did a lot of good work at building [out] that concept. It pioneered along with other companies like Endeca some of the early ideas associated with faceted search and navigation. It did some early work on customizing what had become standard platforms for search like Solr and so forth, but we struggled to find the right fit for the market. We ran into a lot of issues associated with, again, technology transfer and the burden placed on people trying to absorb the technology from the semantic web underpinnings that we used to organize and architect the system.

It was a company that had its run for about seven years with venture-backed folks in the Los Angeles area, Clearstone Ventures, Red Rock Ventures up here, InnoCal down south. Ultimately, we had to wind that down in 2008 when the recession hit. There was just no more appetite, and the investors just declined to continue the grind to try to find an exit with that, so we wound that down. That was a chastening experience.

Barr: I can imagine. You weren't involved in the winding down of Inference?

Allen: No. By that time, I had moved on, and they were acquired in 2000 by eGain, which was another company in that customer relationship management space. But by that point in time I had moved on long since.

Barr: This must have been a kind of traumatic personal experience to go through after having such success with a couple of startups.

Allen: Yes. When that kind of hubris hits you, you really are that invested. Seven years is a long time. It's hard to walk away. People discuss this more broadly: there is a responsibility that as a founder you feel to your investors to stay the course. The people who do that well know when to say, “Okay, that's it.” But a lot of people feel the pressure to say, “Well, we’re just going to continue to go through until you can't do that any further.” That's the difference between soft landings and hard landings, and the shut down in 2008 was not a soft landing by any means. Of course, that wasn't a period in time when everyone else was having a fine time of it. It was a very challenging environment for everybody. But, no, it was a significant lesson.

Barr: It was a hard landing even though the company eventually got sold?

Allen: Yes. There was a sale of assets and that happened a long time after everything had disbanded, so I regret that. There were some interesting things that we did at Siderean that people have been bringing forward in other applications and in other ways from different sources. This notion of how to exploit knowledge to make search that much more effective I think is one that is still a problem that people are going to gain real benefit from working.

Everyone has become accustomed to the tradition of typing three or four words in the search engine box and getting a bunch of results back and weeding through that. Google and it’s like have made that a universal global experience, but people still struggle in many, many different cases when it's an answer they're looking for that they need to be able to take an effective decision and an action.

That's something that harkens back to the expert systems days, which was about, “Tell me what I need to do in this situation based on your expertise.” At a certain point, discovery toward the purpose of making an effective decision, there's a lot of overlap there. What we did at Limbex,

some of the earlier systems at Inference, and what we were trying to do at Siderean shares that common theme. How do you wrangle information to get it at the right place at the right time for somebody to know what it is they need to do?

Risks and Rewards of Startups

Barr: I'm sure that's part of your current life too, but I thought I might ask you as you're now in your life about to transition to a big company how you feel about your entrepreneurial days. Is there another startup in your future?

Allen: I don't know. It's something that I think about not every day but quite often. I do know, based on experience, how difficult that is and the responsibility one has as a founder. If I were in a position for somebody to trust me with that responsibility, again, I would welcome that, but I know the costs.

As one goes on in one's career and one's life it's a significant question about whether those costs are worth incurring for not just your quality of life but the quality of life of everyone else you're involved with and your family and friends and so forth. Going that direction, again, is not a decision you take lightly. It's a very difficult thing to do when your risk/reward tradeoffs are set to a different place than they were in your 20's or 30's or even 40's.

But the thrill of the chase, the thrill of taking a raw idea and shaping it and positioning it and bringing it to market and evangelizing it and getting traction and seeing people get benefit from it, that's a tremendous thing. You can do that in a large organization as well, but the romance that many people bought into with regard to doing that as a startup is always something that's a siren call that beckons.

Barr: Speaking of romance, you've now been in LA almost your whole life. Do you still love LA?

Allen: I do. It's a tremendous city. Over the time that I've been there, I think it has blossomed and matured. In the early years that I was there it could have its rough edges. We went through some challenging times with the riots and so forth. I live in a community that is privileged relative to the entire LA basin, but I can drive five minutes and be in a different place along those lines. That's true of many urban environments, but the tapestry of cultures and cultural experiences and so forth in LA is just tremendous. And the beach isn't bad either.

I have enjoyed it, and it's been wonderful to see Los Angeles have a Renaissance, if you will, with technology as a vital sector. It's always going to be an entertainment town, but what's happened with Silicon Beach and the like… I've been involved with one of the incubators down there: Amplify. I've done mentoring and advising to a number of startups in that environment. It's

been gratifying to see that part of my life get more of a footprint in LA. But LA, in general, has just been a tremendous place to live. I'm very happy there.

Joining Elsevier Labs

Barr: After being a serial entrepreneur, what were you thinking that led you to take a job at a staid publishing company like Elsevier?

Allen: Well, the salary was one. Money was an important thing. But what was interesting was that Elsevier had been a prospect of Siderean. I had spent a lot of time with executives there, understood their problems and their challenges. When I wound down the company after a brief stint working with folks at Symantec, they reached out to me and said, “You know, Brad, we could use your skills and experience in-house to try to affect some of the stuff that we’re trying to drive forward here.” I said, “Okay, I'll try that.” They gave me a tremendous opportunity to remain in LA, which was important to me from a family perspective at that point.

I came into the Elsevier Labs group at the time and shortly thereafter the head of Elsevier Labs came to me and said, “I'm going to be moving on myself. I’d love you to take the position,” so I assumed that responsibility within the group.

Then a number of years later Dan Olley came in to be CTO [of Elsevier] and he said, “I want you to step up to the chief architect’s role and help me wrangle the things that we need to wrangle to get the company where it needs to be as a technology firm.”

Scientific Publishing and the Future of Research Communications

Barr: Can you tell us a little bit about that, about what you are actually doing there? I don't want to get into anything confidential.

Allen: There's a lot going on with scientific publishing and what researchers are trying to do in terms of the future of research communications. There's a lot of controversy in terms of the ecosystem and the economics associated with it. Elsevier has a role that goes back 140 years. We’re talking about a company that is probably at least two sigma out from the normal lifespan of a company. This is a company that has been around for a long time. There's a lot of tradition.

Not long before I got there, it was effectively a manufacturing company. It would commission publications. It would go through the process of putting those things into beautiful shape so that they can sit on library shelves. They would print them out, put them between pieces of cardboard, and ship them off to your university library. The transition that the web represents for organizations, not just traditional manufacturing organizations but traditional media companies, has been a tremendous challenge.

The company before I got there had already negotiated the first step of that transition from being a print publisher to being an online publisher and had done so very successfully to the point where they had managed to do so without changing their business model. But the sorts of things in the long term that are an impact of the web and the way in which information can be replicated and transmitted and so forth, how scientists are beginning to want to work, are things the company is going to need to adapt to.

In the time that I've been there, I've been privileged to be part of the group of people who have, together with the rest of our senior management team, brought forth a consciousness that that change is the thing that needs to be pervasive throughout the company. The transition from a traditional publisher into a provider of information solutions, and principally ones that are based on the data of science and healthcare, that's been a huge culture change and a huge change in perspective that didn't exist when I started there. Participating in that has been an interesting thing.

The whole thing about having an immensely successful company that needs to change course and figuring out how to do that without destroying the company in the process is the thing that we honor and idolize in Silicon Valley; it's the startup story. Right? But there's this other story about how you take a tremendous legacy asset like a company of that sort and change it in a way where it can survive and profit another 140 years. That's been an interesting process to be a part of.

Barr: Yes, hard problem. There are more failures than successes with those kinds of problems.

Allen: As there are with startups, but it's been fascinating and very illuminating. It's an organization that is orders of magnitude larger than the ones that I was dealing with as an entrepreneur and before that working at Inference and at CMU. It's been a fun learning curve, a learning curve I probably should have gone through a little bit earlier in my career, but one that I've enjoyed nonetheless.

Barr: Good. It's a big switch, a big challenge I should think.

Allen: Yes.

Early 21st Century AI Renaissance

Barr: As we're starting to wrap up, I want to try to draw some themes. Where's AI in your life now?

Allen: AI is something I deal with, investing in, evangelizing, and understanding how to use effectively every day. As chief architect, a lot of the things I'm working on are about building the

data and knowledge underpinnings into the company so that it can build the class of applications that can leverage that technology to provide value to our customers, helping them make decisions either as researchers or medical professionals or nursing students, which is another constituency for us.

I am responsible not just for the enterprise architecture function but for the Labs' teams still. I've brought folks into the Labs team, Ron Daniel and Paul Groth, who I've known in previous incarnations and who have a lot of roots in the semantic web community and standards. Together we've been working with an increasingly growing group of data scientists and natural-language processing professionals and semantic web experts to build that platform and to build the technology capabilities to realize the promise of those technologies. It's very much front and center for not only myself but the company and senior management as a whole because we see that as part and parcel of how we are going to be delivering value in the future.

Barr: I thought I'd ask how you felt about the AI renaissance that we are in the middle of? You've been in this for a long time. You've seen a lot of different aspects of AI technology. What's going on with AI? It's all over the front page these days.

Allen: Yes. Part of it is amusing to the extent that you see the same sorts of things happening that we saw back in the mid-1980s: abuse of terminology and astronomical expectations with regard to the capability of systems and the like. Some things that are parallel are actually pretty amusing. As a challenging rising nation state, we have China as opposed to Japan. We have outrageous starting salaries for young researchers in industry, yet again. We have specialized hardware that people have to figure out how to integrate.

But there are significant differences. I think first and foremost the impact of a globally networked community of software practitioners and information sharers has just been amazing. In the context of Elsevier looking at the rate at which ideas are explored, disseminated, and acted on and replicated is just amazing—how quickly that's happening now, where I can open up something that somebody has just put up on a preprint service, go to GitHub and get the code and work with it hands-on and do that on a weekly basis. I think that has put a much larger group of people in a position to work with the technology, deliver it, and build it into systems in a way that we just had no ability to do back in the 1980s. That's a huge difference.

I feel that I'm personally taking advantage of that. It feels like a rebirth in a personal way as much as from the point of view of the industry or the community as a whole because it feels like once again I'm participating in that chain of progression from the early ambitions of the pioneers of the 1950s to what it is realizing.

I had this experience a few years back: The ACM had [a conference celebrating] the 100th anniversary of Turing's birth. It was co-located with the Turing Award for that year, which was won by Judea Pearl for his work on probabilistic graphical models. They had assembled as

many of the Turing Award winners as they could muster and had a number of different panels. They had one panel that had folks like Raj Reddy, and Ed Feigenbaum—I think, he was there—and other people talking about all that's been accomplished.

Take Raj, for example. Raj was there in 1971, 1972 with speech understanding and recognition. That work flowed through in a continuous chain of development to the systems that we have now, realizing speech recognition capability on phones, speaker independent, all of that. That was the thing that he had in mind. It took 40 years to deliver that. They were up there with these exhausted, but mildly pleased looks on their faces like, "Yeah, it actually happened."

I think that some of the promise and some of the ambition we have is still going to take decades to get through, but we're getting real practical results out of real work that's been done in these areas of deep learning and trying to build systems on top of that. Those sorts of algorithms are leveraging not only the stuff that came out of the past, the work in symbolic AI and so forth, but are enabling new ways to be able to build more and more capable systems. It feels great that I still am in a position to be able to leverage and work on those kinds of applications. I'm in a more traditional corporate context now but in an area where that is just as important to the company's success as it was when I was trying to do that sort of thing in a startup.

I think a lot of this work that's going on right now is great. A lot of it will still hit the same sorts of walls. There's only so much you can do with [machine] perception in terms of exhibiting real human judgment. Right? I think a lot of the insight from folks like Rod Brooks that created a huge revolution in building effective robotics systems needs to be brought into and applied to the kinds of decision-support systems that we're building. That will allow us to get through some barriers, but we'll still run into the sorts of things where people want things that are about people-level decision-making that we just are not going to be able to capture with machines.

In the end, I think, the old debate between artificial intelligence or intelligence augmentation is a valid one. If you pressed me on it, I would have to say I'm with the intelligence augmentation crew because the thing that will transform and continue to transform civilization is going to be the partnership of people and machines working together in ways that each of them is doing what they do best as opposed to general artificial intelligence as a goal, per se. But it's a great time, and it's been fun to watch. I feel that in my own small way I'm still participating in that movement.

Final Thoughts

Barr: It is an exciting time. What are your expectations about things you're going to see in the next decade from where you sit that seem really important?

Allen: A lot of the time, particularly as it relates to Elsevier, has been thinking about, “How has science transformed?” One thing I think that's certainly the case is things that have taken off and grown wildly as a result of the introduction of computer technology have really been around communications, changing how science is done in a way that leverages in a real sense what's happening with communication as opposed to trying to replicate the paper-based way in which we did that.

I think that that's going to come down to a situation where you'll be able to have people start real experimental programs in the hard sciences, and the hard science is done the way that software entrepreneurs up here start software companies. They'll sit down in a café somewhere. They'll get on the local Wi-Fi, and they’ll bring up some services. They'll say, “I want to do this and this and this.” It will invoke a whole range of services that are either front ends for contract research organizations or automated laboratory test equipment and that kind of thing. Go off and do that. Bring all of the experimental results back, do some initial organization and automation and so forth, and allow people to do laboratory work without going through the considerable efforts both in terms of the process of becoming a blooded researcher in a university as well as just the sheer labor involved in putting something together—leveraging that sort of thing. It’ll be accelerating science and having the machine [work] in the ways that were pointed to by Doug Lenat’s AM system or Pat Langley's work on automating discovery where the machine becomes a partner in that process, a virtual marching robot army of lab assistants that would empower scientists in trying to drive those sorts of things forward.

So much is bound up about how we manage to continue the progress associated with science and technology and research in what will be undoubtedly, in my mind, a much more challenging environment, with challenges like climate change, the geopolitical environment, and so forth. It's going to be rough to continue a lot of the progress that we've seen over the past decades unless we’re able to figure out how to do it leveraging these kinds of tools that we have.

That goes to this point I was making earlier about intelligence augmentation. It's about getting people in a better position to be able to communicate effectively, not only between themselves but with machines, to do things that they simply wouldn't be able to do individually or do very easily.

Barr: Well, I've got more questions, but I think that's a good place to stop. I think we've covered a lot of ground. Fascinating stuff.

As you were speaking just now, I realized there's still stuff from back in your CMU days that has never really been applied, AI ideas that will probably find their way into future products and economies. But it's really amazing how much of the work that you did back then has just mapped right into things the world needed.

Allen: I like to think so. Ideas are cheap in a certain sense. A lot of people have them. The good ones keep coming back no matter who's involved. It's been fun from that perspective, trying to wrangle these things from back then into results today.

Barr: Good. Thank you very much. This was a pleasure. I'm sure a lot of people would enjoy this video.

Allen: Thank you.

END OF THE INTERVIEW

