Interview with Carver Mead, San Francisco, March, 1994.

ER: I'd like to ask how you grew up and what your parents did.

CM: I was born in 1934 in Bakersfield, May 1st, 1934. The reason it was Bakersfield was because my dad worked for the California Edison company back up in the Kern river valley. Bakersfield was the nearest place there was a hospital so that was where I was born.

I grew up up in the Sierra Nevada mountains in a place called Bee Creek, east of Fresno, about 30 miles back in the woods. I went to a little school with 20 kids total for the first eight grades, one teacher. A wonderful fun time.

ER: What did your father do?

CM: He was the guy in charge of the local power plant. I grew up around the power plant. He used to bring home electrical stuff, batteries and so on. I got hooked on electricity at a very early age and I've been hooked on it ever since. It's never gone away. I never had any desire to do anything else.

When I was in sixth grade a guy moved into camp. This was back in the mountains with a little group of houses around the power plant that we called the camp. The guy who moved into camp was a radio ham. This was just as WWII was ending. He got to teaching me about this radio stuff and I was just blown away. So I'd save all my hard earned money and we'd go down to the big town of Fresno once a month. I'd go down to the surplus store and hunt for bargains.

You could buy the most incredible pile of electronics for a dollar. I collected all this stuff so I could experiment with it. One thing led to another.

ER: Did you have brothers or sisters?

CM: I was the only child.

ER: Was your mother at home?

CM: She was at home. Once in a while she'd be secretary for the local school board or something.

ER: Did you build a ham radio station?

CM: I built all my own stuff. I couldn't afford to buy it. So you substituted labor for buying it. That was much better because you learned all about the stuff. I built receivers up until I could afford one of those little RME 69s, remember them?

JA: I sure do. I had one of those.

CM: Did you? Well that was the only thing I could afford.

JA: With the great big dial in the front?

CM: That was the one. Those were fun days. We didn't have test instruments really. I finally bought a Heathkit vacuum tube volt meter. That was as close as I could get to a real test instrument. You had to debug stuff blind as hell. You had no idea what was going on. That was excruciating, but it would also give you a sense for how to proceed when you didn't have but the tiniest bit of information coming back from the thing and no real instrumentation to find out. It reminds you of neural networks a whole lot, actually.

It was wonderful training. It gives you this confidence that if you muck around for a while you'll figure it out, even though you have no really sound fundamental basis for it. You just have hunches and you develop a feeling for it and after a while it gets to working and the more you work with it the better it works. And then you can't remember why it was hard back in the beginning. It's an awful lot like research in that sense of being out there where you don't know what the hell's going on. There is no obvious reason to be able to succeed and you just keep at it until you start developing a feel for it. It really tunes up your intuition.

ER: What was high school like?

CM: I moved away from home when I was fourteen. I went to the big town of Fresno where my grandmother lived and I lived with her. I went to high school there. It was great because I didn't have to put energy into fighting with my parents. I got my commercial radio license. I ended up getting a job with the local radio station and the local two-way communication place. That was about all you could do in Fresno because that's all the electronics there was. But that was enough.

I rubbed against people that knew more than I did and one of them said, "Well," he says, "you should take electrical engineering at college."

Of course in Fresno, there was Fresno State College and there wasn't any electrical engineering there. The closest you got was the Physics Department. There was one guy there who knew some electronics. I used to go over there when I was in high school. My buddies were all in college and I'd go over there and sit in on their classes. We didn't have any of these AP courses [Advanced Placement]. You just made it up as you went along.

I knew this one guy up at the radio station who said, "You should either go to Caltech or to Stanford."

I'd heard of Stanford. Never heard of Caltech. I applied to both places and I got admitted. I visited up here and I visited down there. I struggled for a while. I finally went to Caltech because it was smaller. I've been there ever since.

I took a Bachelors and Masters and a PhD at Caltech. I had no idea I wanted to do a PhD. When I was doing my Masters I stopped by one of the faculty's office and we got to chatting. He said, "Well, you know, a Ph.D's not much different from what you've been doing except you make up your own problems instead of doing problems somebody else made up."

"Hey," I said, "That's what I've always been good at."

So I thought I could succeed in a Ph.D program. It never had occurred to me before. Absolutely never had occurred to me. I had interviewed for jobs at places like Minneapolis Honeywell. I was thinking of going to work in a very standard industry job and then this guy said what he said, and I went, "Hmmm."

ER: Was he an advisor or just somebody you knew?

CM: Just one of the guys on the faculty. I liked the way he taught his course, so I'd go up there once in a while and chat. So then I thought what the hell, so I applied.

I was one of these late bloomers. I had a miserable academic record from my undergraduate career and I was just getting to be OK as I got into my Masters program. So they made up a special exam to try to weed me out. This exam was what we now call a mini oral. We've since institutionalized it. We do it at the end of the Masters year because it's such a graceful place to let somebody go because that way they don't look like they flunked their oral. It gets the faculty off the hook. I was the first one they gave it to because they figured this guy's got no business being a PhD. I'd done well enough on the Masters that there wasn't any obvious way to keep me out. So they figured that an oral exam was going to make me fall all over myself.

It turned out that, of course, oral exams are more about understanding than they are about being brilliant mathematically. I've always been better at the intuitive side of figuring stuff out than grinding out long solutions to big complicated things and so I did really well. They were blown away. They didn't think that that was going to happen so I got to stay.

ER: And what problems did you choose?

CM: It was interesting. John Linvill at Stanford, ironically enough, the place I didn't go, had just come out with this thing he called the Linvill model for transistors. You basically did a lumped approximation for the linear behavior of the thermodynamic quantities in the transistor, and then put the exponentials into the interface between that and the voltage switches, which is of course a brilliant way to look at the whole thing. I looked at it and I said, "God, even I can cope with that."

Before that time there had been the Ebers-Moll equations which were just these godawful non-linear things. The non-linearities were all mixed in with the other stuff so you couldn't see the essential relationships because it wasn't factored in any nice way.

Linvill had factored in such a nice way that you could see what to do and you could see where everything came from. So I said, "God, this is what I need to do transistor switching times and storage times." I went and worked all that out, on my own. It was enough for a thesis.

In retrospect, it's kind of trivial. But at the time it was neat because I'd figured it out myself. You know when you turn a transistor off it's a while before the minority carriers clear out and the thing goes off. You can predict all that. I had some fun with it. But mostly it was important to me because I'd figured it out myself. Nobody gave me the problem. It wasn't just one more homework problem. I had to figure out what the problem was.

ER: And you went directly from the doctorate into teaching?

CM: When I was a first year graduate student, Dave Meadowbrook had just come down from Stanford. He was teaching a transistor course. I took his course as a first year grad student and I did well in it. The next year he was going to go off and write a book, but somebody had to teach the course. He came around and he said "How would you like to teach the course?"

Here I was, a second year grad student, supposed to teach a first year graduate course. I had been teaching undergraduate courses all along since I'd been a senior. I'd do lab TAs and stuff because I liked teaching. This course was major because there were all these bright guys that were smarter than I was. But I figured it out, gradually. It really stretches the hell out of you when you've got to explain this stuff. You've got to really really know it all the way down.

That experience was what got me in love with teaching. When I got done with it all it turned out they didn't have anybody else around doing this stuff. They were out recruiting young faculty and in those days you couldn't find anybody in transistor physics because everybody had been slurped up by other universities. So they said, "How would you like to stay?" There wasn't anybody there to compete with so I had it to myself.

I went off doing tunneling and transport processes through thin insulators for about ten years. Then Max Delbruck came around and said, "Hey, these guys are saying that nerve membranes work like transistors. Is that right?"

And I said, "I don't know, let me look at it."

He gave me a bunch of papers. People were copying Shockley's stuff and putting in ions instead of electrons and it was all complete hogwash.

So I said to Max, "It's complete hogwash."

And he said, "Well, let's figure it out then."

Max was a gruff old guy. We started this little subterranean group. That was when work just started on biolayers. [Artificial monomolecular films that have certain similarities to cell membranes.]

We did a bunch of biolayer stuff. I guess the only contribution we made there was that we were able to show that the channels were ohmic.

That was my only real biology. We did that with Max and a couple of students, one engineering post-doc and one grad student in biology. We called it our subterranean group because the biolayers would break if you shook them at all. We had one vibration isolated dark room way down in sub-basement.

JA: No real animals?

CM: No, I have a hard time with surgeries.

ER: At this time were you aware of any of the the perceptron work or Widrow's work?

CM: I knew about the perceptron. I'd heard about Bernie's stuff but just vaguely. I wasn't paying any attention to it at all. That was the time I was starting to get interested in VLSI. I had these long talks with Max about how when you work with biological substances you realize there's nothing precise. Everything has to be adaptive. I was getting the idea how that had to work and how it couldn't be that there were any precise parameters. How in the hell can you build a system like this? It has to find its own zero, find its own center, and tune itself up.

At the same time I was getting intrigued by large scale MOS technology because I could see that if you're going to make big systems that's the way you had to do it. That won.

In those days digital systems were the easiest ones to conceptualize so I did them for ten years. I figured out how you made things that would organize big systems and get them to work. I had in the back of my mind that what you really ought to do is to use all the physics of a device to do the computation, not just the on or off property. But you couldn't even imagine doing that unless you had ways of designing really complex things.

JA: You were always heading towards big analog circuits?

CM: They used all the physics. I was from the device side so I wanted to use the physics. It was fascinating to me that you could get wonderful exponentials over orders and orders of magnitude. I had done a lot of that, taught courses on it. It always seemed to me a terrible waste to turn these devices into switches, but that was the only kind of system we knew how to imagine. It still isn't obvious.

ER: I think of the VLSI work as culminating with the Mead and Conway book. [Introduction to VLSI Systems, Addison-Wesley, 1980. A classic book on VLSI design.] I was wondering how that partnership came about?

CM: I had been doing VLSI. There was a talk around Caltech about having a computer science program. We didn't have one at the time. We were interviewing people for the job. They were all very traditional computer science. Nobody was really thrilled at having one more. We talked to everybody and they'd say, "Yeah you've got to have a guy on operating systems or this guy on languages." It didn't feel like the future exactly.

Then Ivan Sutherland came by and got all excited about VLSI. We ended up hiring him and the two of us started the computer science operation at Caltech. Ivan had been introducing me to a lot of people he knew like Dave Evans back in Utah and his brother, Bert Sutherland, who was running one of the labs at PARC [Xerox Palo Alto Research Center]. It turned out to be lab where Alan Kay was doing the SmartTalk project and right next door was Bob Sproul doing wonderful stuff that turned into Postscript eventually through a very circuitous route.

I went to give a talk up there and Bert said, "Why don't you come and consult with us?" So he stuck me next to Lynn Conway. She was at PARC. She worked for Bert. After my talk she came up and said, "You know, Carver, you should really write a book about this."

Then Lynn went off and the next week when I came back she had every book on integrated circuits that was available on her desk. She had looked through all of them and she said, "There isn't anything like it. Let's do it." So that was it. But that's typical. She did all the research, figured it all out.

ER: And how long did it take to write it?

CM: Two years. In a year, we had a note version and we distributed that to a number of universities that wanted to teach the course. We got feedback and finalized it the next year. That was great fun working with people, getting the courses started and giving them material, and getting MOSIS [the national chip fabrication facility] started so people could actually get chips made. That was a great partnership. We had really good times. God, it was stressful, trying to start this whole major thing and it was fun.

ER: You referred to it earlier as a ten year detour. When did you start to get back on your main road?

CM: It was '82. I had been one of the people that had supported the move to get John Hopfield to come to Cal Tech because I'd known him from solid state physics. I'd stopped by at Princeton a few times and got chatting with him. So when they started a move to get Hopfield to come to Caltech I said, "That's great. I like him."

When he got to Caltech we started talking and he said, "Let's do a course."

So we did a three-way course with Hopfield and I and Richard Feynmann that we called the "Physics of Computation." We did that course for three years. Three more different stories you have never heard. You would never imagine the stories we told had anything to do with each other. We enjoyed it immensely because we could start making connections to each other's viewpoint. That was a three year period where we were each evolving our own view of what it all meant. Then we went off and made three courses.

That period was when I learned about what is now traditional neural net stuff. I didn't know there had been any except for the perceptron which everyone knew about. I didn't know there had been ongoing work, I was frustrated because this work was all a very simplistic view of what was obviously a much more continuous, much more rich, adaptive thing. So then I decided, "Hey, I've got to figure out how to make these adaptive circuits, because that's the only way this is ever going to work."

I figured that in five years you could learn how to make adaptive circuits. I've been at it ten years and we still don't how to do any but the beginnings of it. It's been much longer than I had anticipated. Meanwhile, it's interesting that simulations are getting more adaptive. Many of them are getting more like what you can actually build.

ER: I know that many of your students have gone into business and that you yourself have been involved with various companies through the years. Maybe you could talk just a little bit about how you've juxtaposed yourself to the commercial environment through the years.

CM: That was all started by Gordon Moore. When I was a first year faculty, Gordon Moore came by and he had one of these old briefcases that used to open up at the top. He came by with this briefcase and he sat down in my office and he said, "Hi, I'm Gordon Moore from Fairchild and I'm a Caltech alum and I just stopped by to see what you're doing."

I told him about this device stuff I was doing and he said, "Would you like some transistors?"

I said, "Yeah," because I was teaching labs and transistors were expensive if they were any good, and if they weren't expensive they were no good, and either way you were damned. So he opened up his briefcase, there was an old shirt and two old socks in there, and he looked up kind of sheepishly and he said, "I travel light."

He pushed down his dirty clothes and he pulled out these manilla envelopes, you know, 8 1/2 x 11. Two of them. And one was full of 2N706's and one was full of 2N697s. I'd never seen so many transistors in my life. And he said, "Here's some rejects. I don't guarantee what they are, but some of them will be good." And he said, "Come up and visit and maybe you'd like to consult for us."

So two weeks later I went up to Fairchild. It was a little tiny place. There were probably 20 people in the whole company. I gave a talk about what I was doing and then we all sat around a conference table. The whole group was there. We sat around and talked for an hour and had lunch. It was very laid back. They were building these transistors in this little lab downstairs.

I started consulting for them. I'd come up every week and spend a day. I got hooked because I always learned something and I could get samples of different kinds of devices and I found out the things they didn't understand so I could go back and work on them. That's kept up till the present day. It's been a source of an enormous wealth of research projects, the research part that nobody has time to follow up when you're in an industrial track. You've got to get a product out.

I've kept doing that. When the companies get too big so I couldn't do that any more I'd find a smaller one. A lot of my former students have started companies asked me to get involved with them. I come up there every week. I've been doing this since the early '70's, over twenty years. I keep track of what's going on. It's been wonderful, having the best of both worlds.

JA: Do you have any interest in things like artificial intelligence?

CM: I never got involved. I guess I couldn't see a way to make a contribution there is really the honest truth.

JA: Was there discussion about AI at Caltech in that era?

CM: A little bit. And of course at Xerox PARC there was a lot of it because PARC was crawling with AI people. But I never quite found how it fit with anything I know about. I'd have chats with those people and try to understand what they were doing but I never quite got it. But with neural networks I felt right away that this was an analog thing. It's really a simulation of an analog property and that that made a lot of sense.

JA: The work you've done on analog VLSI has been largely in the direction of sensory processing.

CM: There's a reason for that. I started wanting to do learning systems. The very first chip we did was a learning chip. It was a feed forward net with outer product learning. It only had two bits of weight storage. That wasn't enough but we didn't know that. We had no idea how many bits you needed for weight storage. It did its thing and it incremented the weights and decremented the weights and so forth. You could never get it to do anything very interesting. Neither could anybody else. We didn't realize that if we had five levels it probably would have been interesting.

But in the process of trying to make it do something, I realized that if this was going to be interesting we needed real time stuff coming in, and if we were going to have real time stuff we weren't going to get it out of a television camera. I ought to be doing something about the information coming in. That's been a ten year detour. I really think of it as a way to get data that's worth learning. I'm sad to report that we still don't have data that's worth learning. It's been a much bigger detour than I had ever imagined to get sensory preprocessing to where it's in reasonable form to do anything. It's still not there.

JA: Most of the brain is preprocessing. That's really the most interesting problem.

CM: I think it is mostly. There's also a motor equivalent of that which I don't even have a word for.

JA: I heard you give a talk a few years ago where you talked about the cognitive iceberg, which I always thought was a beautiful image. [Mental life is described by analogy with an iceberg, with a little bit of cognition above the waterline and a huge amount of sensory preprocessing below the waterline where most of the work gets done.]

CM: I think it was Wundt who came up with that image. When I first saw it, I thought, "That's my idea." It was a beautiful insight. I really believe it's true. We get these full formed concepts and percepts that come floating up from below.

It's hard to do that, you know. I'm still trying to do it. I can't make an object. I can in trivial cases, but in the real world I can't make an object yet and I've been working very hard. I want an object so I can learn stuff that is interesting. I haven't made phonemes yet or anything even remotely resembling a phoneme. I'm very frustrated because I feel like I should have made more progress. But this tells me that it's a lot harder than I thought.

ER: I was reading a section from a new book. that says we must follow the rule of microcosmic prophet, Carver Mead. The book quotes you as saying, "Listen to the technology and find out what it is telling us."

CM: Yes, I did say that.

ER: I wonder what you think neural nets are telling us.

CM: Adaptation. It's the whole game. I really believe that. It's not just because of the technology I work in, it's because of the nature of the real world. It's too variable to do anything absolute. You don't make a voice recognizer by looking at absolute frequency maps. You don't make vision systems by looking at intensities of pixels. You have to develop a higher level of abstraction than that and you do that by comparing things and adapting to things. The nervous system figured that out a long time ago. But, boy, is that hard because nobody tells you what to adapt to what.

It's a problem, really, about the interface between the real world and computers. If you talk to anybody in computer vision, they'll tell you, yeah, you can make everything work, as long as you control the lighting really well. So they spend all their time going around controlling the lighting really well, and they get it all just right and then you can go and do computer vision. It looks great, but nobody told you that they spent two hours getting the lights just right and if you change the light it doesn't work anymore. You don't have that problem with your own vision system.

To me, that's that. My own belief is that what started out as a dumb kind of adaptation has gradually turned into learning down through the evolutionary time. This is wild speculation, but I just believe, deep down in my gut that sooner or later that's what ended up with learning.

ER: It's a straight line between adapting and more complicated learning?

CM: A straight line on a long scale, yeah. What people do is take a given problem and set everything up for that. The brain couldn't do that. It had to work with whatever came in. So that means it had to adapt much more than our lore has it.

If new stuff comes along, it figures out what to do with it and just keeps going. I think the learning is really long-term adaptation. It's a way of adapting to things on a longer and longer time scale and eventually that gets to be learning. I'm not clear in my own mind if there's a boundary between this rudimentary finding a level and getting rid of what I think of as the false entropy in the inputs, you know, getting rid of the junk that doesn't matter and pulling out the important things.

JA: There's not that many levels in the visual system to do it, either.

CM: And then all of a sudden here you are recognizing somebody's face. When you see your grandmother, there's a population that lights up that's different than when you see your grandfather.

JA: I remember a talk you gave in Washington. You had a lot of pictures illustrating the low selectivity of natural receptors. You pointed out that sensory receptors are mostly low Q. ['Q' in electronics describes selectivity of response. A low Q receptor would respond to a wide range of different stimuli, a high Q receptor responds to a very narrow range.]

CM: Right. That was the thing that blew me away the most in the beginning. In color vision, I realized that you see yellows better than you see any other color because two broadly tuned filter functions are crossing at that point. That was a new idea to me, that you'd encode things using the crossover, rather than by the selectivity of the receptors. Of course, if you're doing adaptive processing, it's the only way you can do it because all you can do is compare things. There are no absolute levels of anything. If you have a crossover you can compare and you get nice things so that made sense.

ER: You were talking about things that have surprised you. What have been the major surprises as you've pursued you work through the years?

CM: It was amazing to me how hard it is to have a new idea. I've always found that whenever I really understood something, somebody else had already figured it out a long time ago. Like this thing about the cognitive iceberg. It's obvious once you've seen it, and then of course, somebody's figured it out a long time ago.

Adaptive circuits have been much harder to build robustly than I had any idea. I'd say that of ten circuits that get invented around our place, one will survive. For every winner there's ten that you think are going to be just as good and there's something that isn't adaptive enough. You don't see it right away, and then later you realize, oh, it's because it wasn't really a relative circuit, it wasn't working on the natural scale of the thing. It was working on something absolute somewhere and I just didn't see it. And then it isn't robust anymore. The difficulty getting to the solutions that are really clean and robust has been amazing.

ER: Are you surprised in terms of the acceptance of your work?

CM: It's taken longer than I would have guessed to get people interested. I think part of that is that the major group of people who know about analog stuff are EE's and they're still a little skeptical about the neural stuff. So I think as the neural stuff becomes more mainstream then we'll get more EE's getting interesting in the analog way of doing the neural. I think there is still a lot of skepticism among the hard core engineering community about whether this art form really is going to turn into anything or if it is still fluff.

That's my constituency, the hard core engineers. There was downright hostility in the beginning when I started talking about neural things. They'd bristle. Then when I started a little softer sell about adaptive analog and how people started out making digital filters and then they finally found out that they had to be adaptive filters. But if you're going to be adaptive, you don't need the precision, and so then you can be analog. Engineers understand that argument. But if you start from the neural perspective they are still skeptical. I think it's turned from hostility into skepticism at this point.

ER: Could you comment a little bit on other peoples's work in the field that you think is important?

CM: I've learned a lot from everybody that I've run into. I'm shameless about adopting ideas from people. I'm not sure I can even trace the important insights to any one person. I've learned a huge amount from Geoff Hinton. I think I learned more talking to Geoff than anyone because he thinks so differently from me. Just struggling to cross that different way of conceptualizing I learn a tremendous amount.

I've learned a lot from, for me, a most intense learning experience. We have this thing called The Helmholtz Club. I don't know if you've heard of it, it's Francis Crick and Terry Sejnowski and Ramachandran and myself and John Allman and a group of other people. We meet once a month and we invite two people. We spend a whole day, from noon on. We get two two-hour presentations and then discussion. There's maybe 20 people there. I never miss one. I have somebody take my class because it lands right on top of my class. I do it anyway because it's just too important to miss.

The reason it's named after Helmholtz is its about the intersection of biology and harder core physics and engineering. And it's been an exposure for me to many different ways of thinking.

We had one talk on eye movements and we had a guy talk about sleep and all the things that happen in sleep. I didn't know anything about sleep. It was wonderful hearing all that stuff from a perspective that's trying to think of a model. The last one we had was on the amygdala. We've had about three sessions now on attention.

JA: Attention is incredibly important behaviorally and it's something people almost never build into a neural net.

CM: It seems to me that's a form of adaptation that's almost totally missing from our official models. It is extremely important because it gets the resources of the brain focused on things that belong together not because they come through one sense, but because they're important for some reason to the animal. It's a way of aggregating the stuff that's in memory and the stuff that's patched together from all the senses and exploration and all pulled together into a context.

ER: What do you tell newcomers starting out in this field? How do you guide them or direct them? What do you tell them is important?

CM: When someone feels like they want to get involved, I try to find out what instinct base they have built for themselves. People with different backgrounds have different instincts about things. People that get really good at something develop a set of instincts around what makes sense and what doesn't. They can sort really rapidly through ideas. This field hasn't got to the point where you can just grind it out. You've got to have some instinct about what's important or you just get lost.

I try see a way in which that instinct base can be effective. Then I will usually try to steer them to someone where it feels to me like the things they have good sense for fit with the way that person works. I'll say you should go and talk to Terry Sejnowski or you should talk to Geoff Hinton or you should talk to John Allman, whoever it feels to me like would resonate with that person's native sense of where things are.

If people think really differently from me, it's hard for me. If I can't see a way to bridge into what they're doing, better they work with someone that has that bridge made already. That's happened quite a lot with people that have come to me and I say, "Yeah, I think you fit really well here." And then I'll make an introduction because we're still not able to make big bridges yet in this field.

There's still pockets of stuff. You know that there's connections but you can't always see them. Most of the time it's just hard work. We still haven't got the big blinding flashes yet.

JA: Occasional sparks.

CM: We do get occasional sparks. That's always fun. You believe there's a big lightning strike out there, someplace.

ER: Where do you see the field going?

CM: I've grossly underestimated the effort it's taken to get real time data to the point where it was useful in learning systems. I would have told you five years ago that by now we would have real time stuff that people would want to learn with. In other words, data that was good enough, with a good enough representation to feed into a learning system. We're not to that point yet. I still believe we're going to do it because it's just too important. It's in the real time stuff that the richness of the network shows up. Networks will be doing things that you just can't do with AI. I don't know how long it's going to take. I'm getting gun shy of making predictions.

We've got to be able to do it in real time with real time stimuli and all adapted into some form where a lot of the invariance gets built in on the fly. That's my own personal goal. I want to provide the front end so people can build the learning system on the back end.

JA: Do you think we're still on an exponential learning curve in this area?

CM: To me, I guess exponentials happen when one thing you do makes it easier to do the next thing. The places you get the big exponents are when you start getting crossings between areas of work so that something that happens here has effects elsewhere.

It feels to me kind of linear right now. We're riding on a computational paradigm that's getting better all the time but in terms of the actual knowledge in the field it feels linear still. It doesn't feel like the insights from one way of looking at it play instantly into the others. That's what you have to have to make real exponential growth.

I think we still have the exciting period to go through in this field. It's beginning to. It feels like it's starting to tell us how it wants to be. We've been telling it how we want it to be and it's just now starting to tell us the shape it wants to go into. When that happens, that will be the big exponential.

Part of what keeps me from being able to take immediate advantage of the rest of the community's work and vice versa is that the stuff I build is so hard to simulate. People don't have the computational resources to to apply it to things they are doing. There's so much energy put into getting networks to run well in the digital paradigm that all these issues of adaptation, how to have things be self-normalizing, how to have the levels find themselves, which are for me the really hard problems, those get swept under the rug.

ER: What about the commercial future of the field?

CM: It's beginning. There's two issues there. One is neural networks working on real problems. There's a lot of that now. Most of them are run on standard computers as simulations and they're done in applications where it's not real time in any sense, like predicting the stock market. You don't need to have millisecond response and you can wait til the market opens the next morning to get your prediction updated.

In the real time area, first you use DSPs, [digital signal processing chips] and then once you get a huge bunch of DSPs then you need them to be in a little box. Then, as a last resort, you'll think about analog.

Analog is not the driving end of commercial neural networks. It can't be. There's too much momentum, too much knowledge, too much lore behind the software side and the digital side. Eventually some of that lore trickles down to applications where there are power and space critical. Battery powered things. At that point somebody's going to be willing to invest in development of an analog system. Then we're going to gradually be able to grow the analog lore up to where it can start to be a real thing. Right now it's a specialty item. By no means is it a commercial neural network thing. That's being done by the guys doing simulation software packages. And thank God for that because you need to develop a base of lore and credibility.

You really have to have done a lot before you'll ever make an analog device because it takes so long to get it right. As time goes on we'll get better at it. But at this point it is really important that we have the digital stuff out there commercially. Otherwise we wouldn't, as analog people, be able to keep the neural net paradigm alive long enough to learn the analog stuff well enough. As a field, we would be sunk if we didn't have the software that's running right now.

ER: Do you think that it's important that people are starting to make the transition from software simulations to making chips?

CM: It's a first step, right? It's beginning to make the point that there is a real thing here. It's very important that it gradually is seen as a thing that has its own sturcture, its own paradigms, its own way of being excellent, its own metrics or performance and value. What was scary for a while was that there were all of announcements of stuff that was pretty clearly garbage and couldn't even compete with standard chips that were out there for other reasons. That gives the field a bad name, when people don't know that they could go and buy off the shelf something that would work better than what they just built. That was an awkward period. I think we're past that now. We're getting serious players that are going in and doing things in a serious way.

It's fortunate that there's enough awareness about the way things are done in a neural paradigm, enough different sets of requirements, that you can still do it in software but much more optimally by building custom architectures.

ER: Your company [Synaptics, Inc.] is involved mostly with recognition problems. Do you think this is the foremost area where neural networks can make a commercial contribution?

CM: It's an area where it's very hard to do anything without a neural net. The AI people tried for years. In a week with a neural network you can do better than people have done fooling around with AI programs. I'm particularly interested in the perception end of it because I'm working from that end.

With real world data, where you don't know what the information is. Neural network paradigms can pull out that information and make it useful. That is very very clear. And it wasn't clear five years ago. Then it was a gleam in all of our eyes, but at that time it wasn't clear to anybody who was objective about it that it was going to be better than sitting down with a smart guy and writing a program. Now it's clear that there's no contest at all. That's been a big change.

You probably remember some of the early talks I gave which were, "Hey let's not over hype this thing because it could turn into another crash." We don't need another 20 year famine like we had after the perceptron. I think we're well past that now. I think, my own feeling is, it's not going to go away no matter what.

JA: It wasn't really a crash. Maybe a flat spot.

CM: I think we got through it. I think the group has been really statesmanlike bringing in the old guys. It's like saying no, this isn't a community where we're pushing people out. Let's figure out how it goes together, what the insights are from where ever they're coming from. Let's have a community that's open to new ideas and new ways of looking, or old ways of looking, or whatever. Let's pull together and find what makes these big collective systems work.

To me that's been the thing that got us through. There was a tendency for a while for people to break into groups. It feels to me like the statemanship won out over that and that that's really what caused the community to survive.

ER: Do you have any sense of what's going to happen in terms of the government's role? Traditionally this field has been funded by the more defense oriented needs of government and now defense needs are changing.

CM: I can give you the jaded view, or I can give you the optimistic view. Maybe I ought to give you both. The jaded view is that the agencies will go after the hot new fields regardless of what their nominal role is. DARPA has been a great example of an agency that has figured out how to position itself to look like it's doing the current politically right thing to do. They've been skillful enough at it so that they haven't been bashed like a lot of other agencies have.

I don't think the basic forces in Washington for who does what to whom are going to change very much. The fly wheel's too big. It's just the names will change around a little bit. The players will move around. They'll keep talking to each other and they'll keep politicking the way they always have.

The problem of having something just defense related is it was always a sham, anyway. Because with research, you don't know what's really going to happen. So you have to make up something. The more you have to make up, the less relevance there is and the less feedback there is between what works and what gets worked on. We've seen that played out over the last 20 years to the point where now there's very little feedback from what works to what's worked on. You you see huge budgets going into things that have no relevance to anything that's going to ever be useful to anyone. You see things that are really very useful that don't get any money.

In the military sector it's 20 years before anything gets fielded, and everybody forgets who did it anyway. There's no feedback. In the commercial sector, there's a lot more feedback because it's pretty clear what works and what doesn't. You're much more apt to get some sensible form of feedback.

I think because of the fact that neural networks do useful things, there will be a net positive influence from actually looking at what they do rather than making up stories about what they might do. I mean it took us 20 years to get completely disconnected from reality. It may take another 20 years of constant work to get back connected with reality.

JA: I asked about AI earlier. Do you have any interests in things like cognitive science?

CM: There are a whole bunch of cognitive people that come to Helmholtz. We always try to pair them with a physiologist to try to get that bridge. I don't consider that AI. The whole cognitive psychology, perceptual psychology side, I consider that central to everything.

JA: I'm glad to hear you say that.

CM: How would you work on a perception problem if you didn't know how you perceive?

JA: I think you have to know what the output is doing as well as what the input is doing. But there's a lot of people that don't think that, unfortunately.

CM: I think it's central. I think as a practical matter, we've actually gotten more from perceptual psychology than we have from physiology. Fortunately, physiologists are starting to use psychologically important stimuli.

JA: That's been a sub-theme for a long time, but it hasn't gotten more respectable or popular until fairly recently.

CM: Now it's kind of expected that you'll do that because otherwise how the hell do you know what will happen? I find that the most important trend in neuroscience in recent years. Of course there have been the occasional people that did it anyway, but it was an oddball thing, rather than mainstream. But now it's mainstream.

JA: It was always like real scientists count photons.

CM: There was this terrible reductionist thing for years which was just deadly. If it wasn't reducible to that level of basic physics, then it wasn't real. And of course nothing could be further from the truth. I remember Feynmann once gave a talk. Somebody asked him something about chemistry. And he said, "There's a reason physicists are so successful with what they do. And that is they study the hydrogen atom and the helium ion and then they stop."

The chemists have to deal with real molecules with more than one electron and more than one atom. Of course there's no way to do that with physics. So then you have to start making approximations. You have to start making constructs at a higher level. Like the atomic bond. Well, what the hell's that? Yes, it has some quantum mechanics in it, but it's not something that you can just solve.

The more complicated the system, the higher level and the more approximate will be the concepts that you use. The good physicists knew that. They knew that the reason they were so successful was they didn't tackle the real problems. They made up a problem that was simple enough you could actually do it and then announced they were successful and left.

In our business there's no hope whatsoever of having that kind of reductionism work. We're more like the covalent and the ionic bonds. They're approximate ideas but without them you couldn't make any progress at all.

But it is interesting that, having said it from that side, we can say it from the other side. There was a time when psychologists would argue vehemently over the meaning of some term. The philosophers were the same way, and they would yell and scream at each other and when you look back at that, those diatribes, you say, what the hell were they thinking about? It's like they were trying to be precise about something that was a very fuzzy concept anyway. It had very little to do with anything you could actually observe. Maybe what they should have been doing was more experimental. It feels bad arguing about the meaning of some word which doesn't really have a meaning yet because you couldn't operationalize it.

How many angels are dancing on the head of a pin? Maybe it was Aristotle or somebody who felt that you ought to be able to prove everything. It's only been 20 years that we've known that that was an oxymoron, you know, with Goedel. And nonstandard analysis was only worked out 15 years ago, so that it's OK to use infinitesimals again. Maybe Leibnitz had something after all. It was weird, the place we got into because of our Western upbringing that said you've got to be able to prove everything.