Interview with Bart Kosko, Portland, Oregon, July 1993.

ER: Why don't we begin with your date of birth and where you were born?

BK: I was born February 7, 1960, Kansas City, Kansas, in a region called Strawberry Hill, which is a Slavic community.

My father was Russian. At the time he was a building contractor. My mother was Serbian, a housewife. I grew up in a Serbo-Croatian culture. My first musical instrument, for example, was a mandolin, an approximation of a brac, which is a Slavic instrument. My cousins were in The Tambouritzans, a balalaika type orchestra but with Yugoslav instruments. I thought everybody spoke a little Serbo-Croatian or Russian. I ate that kind of food, lived that sort of life style, knew about the feud between the Croats and the Serbs.

ER: There was no academic interest in the family?

BK: No academic interest. I come from a long line of peasants. My grandparents all came over from the old country in Eastern Europe. They came over poor and all ended up in a Slavic region. Slavs, Poles, Russians, Dalmatians, Yugoslavs, Czechs, the whole group, in Strawberry Hill.

ER: So there was a rather intense family scene?

BK: Exactly, on that side of the family, on the Yugoslav side of the family, it was very much family oriented, very positive. The Kosko side, the Russian side, was a little more aggressive. Everyone's fighting, always getting into wars, those sorts of things.

ER: Did you have brothers and sisters?

BK: I have one brother three years older. All during high school he was the intellectual. I was always off doing other things, more in the arts. I was an outdoors kid at the time. He got the scholarships first and then went to school. We all went through the hippie era.

ER: I'm curious what your earlier childhood was like.

BK: Well let me say it was very interesting. I grew up very early. I grew up as part of a Kansas City street gang. My brother was the youngest member of it and he was three years older. I grew up very early. I had my first sexual experience with a girl when I was four. A deep dark secret. I didn't find out that my brother and his friends had had their experiences at the same time until I was 25. So we were part of a group of bad boys in the street gangs in Kansas City. Most of those kids grew up to become complete hoodlums. I don't know if they're in jail or not. When I was in second grade, for whatever reasons, my father thought it was a good idea to move out into the country to a farm. That's where we lived for a fair amount of time. That was the best part of my life, out on the farm, wide open, minimal government.

I've always had a problem with the government. When I was three the government took our house, through eminent domain, to build freeway 635 for the Kansas City airport. That was a shocking thing to tell a little kid. Your house is going to be torn down and destroyed. So when I was seven, eight years old, I found myself on the farm, hunting and fishing and animals. It was bucolic. It was just wonderful. Wide open. My father was a very open minded man.

That all ended for me when I was ten when the house burned down. A few months after that my father died. I began living with other people then. My brother and I split up. I still stayed largely in the farm community. We had another farm after that very briefly, in the same area. Through high school I stayed in the city of Lansing, Kansas, outside of Levenworth in the northeastern corner. My high school never had more than 400 people in it.

ER: Were you with your mother during this period?

BK: Sometimes with her, often split up, living with other people.

ER: Relatives?

BK: More often friends. I had different friends during high school. I went through various phases at this point and went through the hippie phase very early, when I was young, and got deep into drugs when I was twelve. That began for me in part as an herbalist under the influence of Euell Gibbons, "Stalking the Wild Asparagus."

It was Kansas and so you start smoking some local pot and then trying different kinds of herbs. My brother was deep into chemistry at this point. I remember we read a copy of Aldous Huxley's "Doors of Perception." I was the time around 12, 13 years old.

We wanted to try mescaline, and did. I had some very mind opening experiences when I was young, with LSD. I was deep into the music, rock music. Then I had a bad trip, I think when I was fourteen. A very bad acid trip, a paranoid trip, and I got completely out of it and was turned off to the whole culture, including rock music.

For whatever reason I got deep into classical music at that time. I had some training in the mandolin and I switched that over immediately to a violin. I got a copy of book on orchestration, and I began to study that. By the time I started high school I was writing my first little violin concerto. I got some supervision at the local college, from a music professor there, and got deep into music.

So at a farm school in Kansas I saw myself as the next Beethoven, and started writing a lot of works, small works, big works. Within a year, I was getting performed by members of the Kansas City Philharmonic. And so during high school, where I did all of the fucking and fighting that you do in high school, I'd come home at night, and every night before I'd go to bed I'd make sure that I wrote some music. I got into the habit of the discipline of creativity. I would not go to sleep until I had written at least a few bars of some project. By the time I was 16, I was reasonably good. By the time I was 17 I'd won the Young Composer's contest in the country. Now I started getting scholarships. The best one was from USC. And out to USC I went.

ER: So music was always a kind of central organizing principle in your life.

BK: I began with music. The big intellectual event for me happened when in my senior year in high school, I learned physics and I lost my faith in God. This was a big crisis for me. Physics pushed God right out of my head and replaced it with physics. There was an alternative description of the universe. I was deep into Newtonian mechanics. I had a professor, rather an instructor, Bill Geier. Since Lansing, Kansas is next to Fort Levenworth, there were an awful lot of people cycling through the military school there, high caliber instructors who teach pro bono and step in to teach at the local farm schools. This fellow did that for me. He came in to teach our physics class. When I stood in graduation line and I had my robe on he walked up to me and said, "Here is your graduation gift."

It was May, 1978, and it was a copy of that year's Pulitzer Prize winner by Carl Sagan, "The Dragons of Eden," which is speculations about the origins of human intelligence. This was the first book I'd ever read on the brain.

So now I'd lost God, and I was deep into physics, and got this book and read it cover to cover in a day or two. He handed it to me, I think, then, rather than earlier, because this was the sort of town that was very much bible belt, at least in the school. This was considered a very radical book. It got me thinking, and was the first time I ever understood the idea of neural networks and synapses in the brain.

Most importantly, it had a mechanistic view of mind. I'd long since become a materialist in my philosophical point of view. And here at last were some speculations on how that might actually be brought about. That was my first acquaintance with neural networks. I'm happy to say that many years later when I wrote my first textbook on neural networks, I wrote thanks to him, and he was nice enough to write a kind response in return.

When I went away to USC under a complete music scholarship, I was already working on my first symphony. I had a contract to record. I got it on the basis of an orchestral overture to the Count of Monte Christo. When I came to USC I wanted to be the next Richard Wagner. I wanted to make films and score them and write them. I wanted to be the artistic superman of the day. I had no concept of how any of these things worked. I didn't know that USC, for example, is located in the ghetto. I thought of LA as filled with movie stars and all those sorts of things a farm boy in Kansas would think about Los Angeles.

So I got out to USC on a full scholarship to the top music school on the West Coast. The first thing I did was take placement exams. These were for undergraduates, and I passed them. Understand at the time I was very advanced in my musical work, had several copyrights, and I was orchestrating my first symphony. So they had me take the PhD exams and I passed those.

They didn't quite know what to do with me because the music that they taught was atonal. And what I like is very tonal. It would tend to place more like the late 19th century in the sense of harmonic romanticism and the orchestra. There was no way in hell that they were interested in doing that. So we had a big fight, and the way it was resolved is that they got me into the graduate course in film composing, which was what I wanted to do anyway, and let me keep the scholarship for a year, provided I got the hell out of music. So I did.

So all of a sudden, before classes had even resumed, this guiding principle of my life, this goal, the music, all of a sudden was cut out from beneath me. Now I was drifting. But my hobbies had been philosophy and science. I'd done a lot of reading. The four years of undergraduate school led to two degrees, one in philosophy, one in economics. Those in turn have led to the two fields I work in now, fuzzy logic and neural networks. The philosophy led to fuzzy logic and the economics led to neural networks.

ER: You said that you had read widely in philosophy and science before you got to USC. You referred earlier to Aldous Huxley and "Doors of Perception." Was that what first made you start thinking about philosophy?

BK: I think you think about philosophy. The two big questions were goodness and Godness. "Does God exist?" and "Is this right?" These were the questions of ethics. Is the universe just? Is that particular action just? It had nothing to do with psychedelics.

My brother periodically came back. He was at Northwestern at the time, and he had just taken the latest course in philosphy or math, and we'd debate this sort of question. That was the age of "Zen and the Art of Motorcycle Maintenance." You'd sit around smoking pot or eating a magic mushroom and talking about metaphysical questions. Although we lived on a farm, we had some terrific drug parties out there. Thinking back, it was one of the best parts of my life, but I saw my relationship to the government was very much like my relationship to God. It was one I questioned and one I lost in some sense.

Life in Kansas, now that I look back on it, was very good. Wide open. A chance to do my own thing in my own way, make a lot of mistakes, pursue a lot of paths. The whole drug culture was a part of that.

ER: What made you choose economics?

BK: That's an interesting question. Because if you're really going to be a fanatic about philosophy you have to have a political philosophy. So the first thing a young man in philosophy encounters is Marxism. It seems the most radical. The first book I read on it I recall was "Das Kapital." I read that before I read the "Communist Manifesto." Having lived on a farm, and worked so much with the supply and demand process I wasn't impressed with the labor theory of value. I was impressed though with the revolutionary spirit as I think most young people are. That's what brings them to Marxism. To man the barricades, not the labor theory of value.

So at the same time I found that the collectivist type arguments really weren't for me, I was looking for a social philosophy. And this is where I stumbled upon libertarianism. This idea has two goals: Maximal personal and economic freedom.

Now that I look back upon it I think that that's where I got the idea for the fuzzy cube because there's a fuzzy square used to define your political position. It has two axes. One axis is from 0 to 100 percent for political liberties, and the other axis for 0 to 100 percent for economic liberties. You slice up the square into four pieces. One quadrant that says low political freedom and low economic freedom is the populist position, an Archie Bunker position, a government control position. Another quadrant that says higher economic freedom but low personal freedom is the conservative position, the Rush Limbaugh position. The other quadrant, diagonally opposite from that, is high personal freedom and low economic freedom. That's the modern position, the bourgeois, the petty socialist, the modern liberal, or Bill Clinton, for example.

That leaves the last quadrant, high economic and high personal freedom. And that's the libertarian position. I saw that drawing and that's me up there in the far corner. The real question is, what's the optimal size of government?

I started taking courses in microeconomics. I thought minimum wage was a good idea, why not raise it by 50 percent, why not raise it by a thousand percent? Minimum wage, rent control. It all sounded good to me at the time. I was poor. I was working on work-study all my time at USC. But these notions collapsed with a supply and demand analysis. So I found myself increasingly in that upper box arguing for a government that was limited.

At the time I thought the limit could in theory go to zero and maybe some day would. That was really the essence of Marxism. The real vision of Marxism was the state would wither away. In fact, it tended not to wither, but the ultimate Marxist state was a very libertarian looking thing, with complete wealth, complete freedom from the state and so forth. It was the means of achieving it that was debated.

So I thought, well, basically I was in agreement with that, that we all wanted to end up in that upper quadrant in that final corner of complete freedom. At the time I was getting deeper into the libertarian philosophy, I was also being forced to register for the draft. That tends to sharpen one's position on these issues. I was in the first draft pool. The cut off line began for those born January 1, 1960. I was born February 7, 1960. That made me think a lot about the idea for the free market or volunteer draft.

The case for a volunteer military seemed very clear. In case after case after case I saw that the alternatives were really, were you going to have the government provide a monopoly or have competition for a given service. It had nothing to do with the nature of the service itself, whether it was protection by way of the courts, or police protection, or in the end even national defence. The effects of philosophy and economics for me were to erode my belief in the state. Like a lot of people I lost my belief in God with science. With more science and philosophy I lost my belief in the state.

Now what happened is that part of my work at USC was in political philosophy. My instructor there was John Hospers, who was the first Libertarian presidential candidate. He was at USC and wrote a famous book called "Libertarianism", the first book on it. Right away I became a campus libertarian. My friends and I took charge of the speakers' committee. I remember, for example, we made Ted Kennedy speak outside in the rain one day. We brought in speakers and kept out others. I remember we brought in G. Gordon Liddy.

The other thing that I did is begin writing essays. At this time, when I was 19, I tried to become a professional writer. I started writing basically porn stories under a pseudonym. It was very hard work. I took courses in the graduate program for writing at USC to train myself for writing. I started making some extra income and I learned the discipline of writing papers. I started writing essays. Essays about the draft, the abolition of victimless crimes, the nature of liberty. I would run those in the Daily Trojan [the student newspaper]. Then I sent copies to my friends and comrades at other universities. At KU, Northwestern, at different schools. They would run the essays under their own name and they would respond to the questions themselves. That was their obligation. I remember at one point I was hitting an audience of over 100,000 people. That was a very powerful feeling when you're 20 years old.

When I got into philosophy a little more deeply, I started from political philosophy. The modern philosophy is the philosophy of science, and the works of Quine and the logical positivists. Every statement was either true or false. It was meaningful if and only if it was an empirical or testable statement or a logically trivial statement. In other words either a statement of math or science. Questions of ethics and questions may have personal meaning to you, but had no cognitive content.

My heroes were the members of the Vienna circle of logical positivism in the 1930s, Rudolf Carnap and Quine, who came and visited us at USC in the Philosophy Department. I began taking courses in symbolic logic. So I had lost my faith in God, and I turned to science. I was a fanatic to find some kind of foundation to stand on. I very soon realized that the language of science is math and that the structure of math is logic, and the essence of logic is basically these Aristotelian assumptions of the black and the white. I remember, for example, running across the books of Ayn Rand and her big propagandistic novel called "Atlas Shrugged." The three acts of the drama are labeled the three so called laws of thought of Aristotle, "A equals A", "A or not A", or "Not the case that A or not A is a contradiction."

I began to question these laws of Aristotlė. This was to me the next big changing point of my life. I got deeper and deeper into logic and was taking graduate courses in symbolic logic while still an undergraduate. Suddenly I had a crisis. I couldn't find a single statement of the world, about the world, descriptive world, the world of factual truth, to which logic applied. In other words, I couldn't find a single statement that was either 100 percent true or 100 percent false. The grass is green, the sky is blue, the dirt is brown. Any of these statements were matters of degree. But by logical law, they had to be true or not true. They had the same status as the statement 2 equals 2, or 2 equals 3. This was a great crisis for me. I saw a mismatch between a grey world and a black and white science.

Then I stumbled on multivalued logic, as it was called then. And that to me was a big breakthrough. I thought this was a central issue, trying to get the language to match fact. That led very quickly to my fuzzy research. Now at the same time, in economics, I got deeper into the study of free markets. This was the beginning of the Reagan revolution, so the free market was beginning to get very popular. But there are certain theorems that a lot of people don't know about. One of them is called the Coase Theorem. For this Ronald Coase was given the Nobel prize in economics a few years ago. It says that if transactions' costs, in exchange, are zero or small, and if property rights are well defined, then the market outcome is Pareto optimum, or is efficient. Pareto optimum means it's a kind of global equilibrium where it's impossible to make somebody better off without making someone else worse off. It's almost like an ideal Marxian state. So you can view, for example, the economy as a big game, an exchange game. If you hit a state of Pareto optimality you'll never move out of it. The Coase Theorem was bouncing around the halls both of the econ school at USC and its law school.

The other theory that was very hot at the time, that brought me back to biology, was the theory of sociobiology. This was one of these grand sweeping world view theories. The selfish human. We're just gene machines and the chicken is the egg's way of making more eggs. We're DNA's way of making more DNA. I began to look into this business. I read E.O. Wilson's book on sociobiology. That was a lot of work. Population biology, ecology, there are many things packed into that book. I started writing my first technical paper, on the marijuana market. I viewed that market as a sociobiological type game looking for what's called an ESS, or evolutionarily stable strategy, another notion of global equilibrium. That's what is called a Nash equilibrium in game theory. Like the Pareto optimality in the Coase Theorem.

Increasingly I was thinking about the social systems that emerged through so-called invisible hand mechanisms, for example, language. No one invented languages, they just sort of evolved. Most social institutions. General market outcomes. Supply equals demand equilibria. The concept began to fill my head that the more agents enter the game, the quicker and in some sense the better the equilibrium you reach. So between sociobiology and the new mathematical economics, and my political enthusiasm for free markets

I was driven deeper into the mathematics.

I was very good at logic. When I'd taken calculus earlier I was basically self taught and never had much enthusiasm for it. I hadn't run with it as I'd run with music, for example. So when I was 20, right before Christmas break, one day I went to the book store and bought for $1.50 an old calculus text. I took it with me over Christmas break and worked every problem in each chapter. I thought it was easy and trivial and simple. I kept doing this and doing it like an exercise workout, with discipline. Pretty soon it was like when I'd worked with physics or with musical theory. I got very deep into it.

So when classes resumed, and the new semester came around, I enrolled in an advanced calculus class and got the only A in the course. That began a new line of training for me in mathematics. I pursued it with religious zeal.

I had to learn topology. So I got the Schaum's Outlines books and did some background work and got my first notion of topology, very general notions of connectedness and compactness and smoothness. I studied the transformations of one system into another, so that I was able to handle the works of Gerard Debrue, who won the Nobel prize in economics in 1983. And in particular the key book here for me was our bible called "A Theory of Value." Now the entire ecomomy is reduced to a fixed point of the system, Brouwer fixed point. You can picture it in terms of, again, a unit square if you view the positive diagonal as a locus of fixed points where x equals f(x). Then it's impossible to draw a curve, a smooth curve, from the left to the right without hitting that diagonal at least once. So the theorem says that a continuous function of a compact convex set into itself always has a fixed point. Debrue had this clever idea of showing the economy in terms of market clearing functions and setting them up in terms of the compact convex set of price vectors, which is a simplex. Each price vector is like a probability vector, the components all add up to one. It maps into itself. Debrue showed there had to be one such price vector that decentralized the economy. That, in effect, was Adam Smith's invisible hand. I thought well gee, this is the ultimate proof of laissez faire capitalism, it falls out of the Brouwer fixed point theorem and the Kakatuni fixed point theorem that extends it.

I got very deep into that and then ran into a neural theorist, the works of Morris Hirsch, Mo Hirsch. His book, "Differential Topology", I couldn't really handle it, and I still have an awful hard time with it, but increasingly the works in economics dealt with what are called generic systems, structure of equilibria, properties that hold almost everywhere. The more agents you had, the more likely that you might see something like you see with neural systems, that is, exponentially fast convergence. It was this idea, systems equilibrating, that got me into neural networks. And at the same time I was working with the ESS, the evolutionarily stable strategies of game theory. The idea was to extend that idea to CSS's or culturally stable strategies. I did that with my marijuana paper, which it took me many years to get published.

Another thing that helps concretize my ideas is trying to write it down and get it published. The idea of the invisible hand, the convergence of systems, Debrue's work, Arrow's dictatorship theorem, the Coase theorem, all these things, convinced me that broad social structures could be construed in economic terms. Somewhere, someday, I don't remember how or when, I thought of the brain in the same way. Why couldn't the brain act like a big economy?

It was in that context that I ran across an obscure paper by Stephen Grossberg, very hard for me to read, about competition as an organizing principle for biological systems. Not just at the broad level of the Darwinian slug out, but even down at the level of the structure of the brain. I have to give Grossberg credit for bringing me into the neural field. In some sense it began with Carl Sagan, but it was with Grossberg that I saw the mathematics that he had taken along the right lines. He actually has some theorems about tying behavior, agents, in effect, to brains.

So then I wrote my own paper called "Equilibrium in Local Marijuana Games", which was a game played between growers, rip offs, and narcs. I wrote that paper after watching a Ted Koppel program on ABC Nightline, on the growing of marijuana. There were narcs trying to raid the patches, but they found that as the narc raids went up, the rip offs went up as well, that was the game played between growers and rip offs. So modeling that phenomena, I had my first mathematical result. I derived a global equilibrium for the entire grass game. Given any combination of agents, any mix of growers, rip offs, and narcs, you'd always converge, and exponentially quickly, independent of initial conditions, to a stable outcome.

I had just graduated from USC now, and had no money. I was accepted into graduate school in math at UCSD. During that summer I remember I checked out over 200 books at the USC library because I now spoke the mathematical language. I read books on traffic theory and books on engineering and as much as I possibly could on population biology. I laid out my first novel on sociobiology. I was writing a little fiction along the line. I did publish a little story.

I tried desparately to get my marijuana paper published. The only place I thought I could do that was in "High Times." They accepted it but rather than paying me for the article they offered me advertising space. That was really not the sort of thing that I do so it didn't get published at that time. It took many years. It didn't appear until 1991 and I changed it along the way but the essence remains, the result that I derived when I was an undergraduate.

In my last semester at USC, when I was taking various math courses, I wanted to understand general relativity. It's just an intellectual goal. To do that you need to understand differential geometry, the calculus of curvature, calculus on manifolds. They offered a course on it that semester with a fellow named Mark Kac. He was one of the editors of "The Annals of Probability."

I remember Mark walked into the class, a big man, big powerful presence, white hair, bright red tie. And he said, "I'm here to tell you the truth, and only the truth, but not the whole truth, because that would scare you." I was very taken with the man. I began pursuing differential geometry much more than I otherwise would have and I became his protege. He was new at USC and I was just about to leave USC. He wanted to see as many people as possible go into pure math. Within a month or two he was writing letters of recommendation for me and that's how I got into UCSD, in the math department.

I was offered a scholarship in the econ department, the economics program at USC, but I wanted to pursue math. I was really into the idea that math was the language of science and this is the new religion. I just couldn't learn enough of it. So through the help of Mark and others, I made it to UCSD. I remember, we'd have many discussions about the nature of probablity. He was a rabid probabilist and also was called an operationalist. That led me to believe that what exists was what you could measure. He had worked with Richard Feynmann and other on the foundations of quantum mechanics. If you asked him, "Is the moon there if you're not looking?", he would say, "No."

I thought now here's a guy I have great respect for, saying something that is clearly false, clearly outrageous, to a realist, and yet there's a reason why he's doing that. I'm very suspicious of probability, just from its general problems, and also because probability looked a lot like God. It was used to explain a lot of things and you could never catch it in the act. You could never find the real thing.

I remember asking Mark about multivalued logic and he pooh poohed it. We talked about what was called the measure theory of probability instead. So when I was now 22, and at graduate school at UCSD in mathematics, a TA for calculus and so forth, I began to pursue measure theory, or the formal theory of probability. My economic development went on hold.

At the time I was very poor. I had no money and a minimal scholarship in the math department. This was at the height of the recession at the time, in 1982, '83. I raced through and got a Masters very quickly, in one year, and got a job. I sent out 70 resumes, I recall.

The only job I could get was at General Dynamics. I had mixed feelings about this. I mean here was a libertarian about to go to work for the largest defense contractor in the world. It was very seductive. They were paying me $30,000, and I had never had anything like that. I had maybe earned $5,000 a year, if even that. So I got $30,000 to be, in effect, an artificial intelligence consultant, a mathematical internal consultant. It was very sexy. I could continue PhD work gradually on the side. I had to take it. I'd always been poor, and I wanted to have my first home and all those sorts of things.

So before I knew it, I'm at General Dynamics. I actually started July 11, 1983. I began the security interviews and all the things that you have to do to work there. Again I felt the omnipresence of the state that had plagued me since I was age 3. But the one thing that was very good at General Dynamics was they had a library at the Convair division where I was in San Diego, and unlimited resources. In fact, they changed the copying machine policy because of me, because I copied so many articles.

The first thing I got into there was fuzzy logic. I had stumbled across that term, fuzzy. Before I'd always heard the field described at multivalued. I read the works of Lofti Zadeh and checked out all the books I could find on it and had ordered for me other books on the subject. I went to seminars and just devoured that field and related ones. I looked a lot more at artificial intelligence, which is interesting because of the problems it dealt with, but it's not satisfying. There wasn't the mathematical basis there.

It was also at this time, at General Dynamics, that I began looking at networks. I wanted to apply the economic notions I'd worked out to military planning. Somewhere along the line I ran across the words "dendritic tree." And so it was time to start looking at neural networks. I got some primers on neurobiology. Before long, I ran across Stephen Grossberg's book, "Studies of Mind and Brain", which I think had just come out. I had a hell of a time reading that. In fact, I still have a hell of a time reading that book. It's my favorite book in the field, but very deep.

For me these two fields that were to play such a heavy role in my life began to come together. The first research type thing I did in '83, was to come up with something called a fuzzy cognitive map. Now the term "cognitive map" has been used by many people in psychology and also in political science as a way to relate causal events. The idea is pretty obvious, just to allow fuzzy causal events as well. So if this node goes up then that connected node goes down to some degree. The nodes themselves can stand for fuzzy sets, like the strength of a government, or political activism, or these kinds of very abstract notions to which all events belong to some degree.

My application was driven by the problem of how you put values on a target. I was working on some smart weapons at the time, largely the Tomahawk system. There was the problem of shooting several Tomahawks at different targets. It is a relative evaluation problem. It's the problem of the target value of a bridge. It's worth a lot before the tanks go over, it's worth very little once they've gone over.

So we looked at expert systems, decision trees, and they didn't really handle the problem. Then the next thing that happened to me was the cognitive maps. I wrote my first paper using them. The analysis suggested that they really should have feedback. But the minute you put feedback, there were closed loops and busted tree structures. Then the doors of AI shut and you couldn't do inference. The question was, "Well what could you do with it?" There was just one idea. The hell with graph search, why not just let the thing spin around, and see if, like an economy it would cool down and it would equilibrate.

Really, this was a neural network I was playing with. It was the year after the Hopfield paper came out. I read that, and other papers. I started to see the neural connection, though I was still thinking of the cognitive map in terms of causal prediction, which is the essence of philosophy. You know, Hume said, causality is an illusion, and when you say A causes B really it just means if A then B. Causality was a constant conjunction of events. And I thought, well, that's interesting. I'd like to see if we could have an adaptive causal structure. We could change the causal links in a cognitive map.

The problem I had with the Hume idea, which was a correlation idea, really a Hebbian idea, was that it grew spurious causal links. The idea that came to me was that of John Stwart Mill who said that causality is a concommitant variation of events. To me the simplest way to deal with variation is as a change, as a derivative, and as a product you just multiply. This is where I came up with my first neural contribution, what I later called the differential Hebbian law, or the differential synapse. But it was originally designed for causal prediction on a cognitive map with no intention of having any neural interpretations. Much later I saw the work of Harry Klopf, who did apply this to neural networks, in slightly different form.

I began to study how a cognitive map might behave, and to see its matrix structure. I read a book edited by Jim Anderson and Geoffrey Hinton and saw that you could reduce a lot of neural networks to linear algebra, followed by some non-linear operations. I worked that out for cognitive maps.

At about the same time, I was proving my first theorems on the foundations of fuzzy set theory. By now I was caught up in the debate between fuzziness and probability. Most people just said fuzziness was probability in disguise, and I wanted to see whether that was true. I expected it might be, or it might be the other way around, since I sat in the fuzzy world. At this point I had met Lofti Zadeh and many of the founders of the field. Lofti took me on, in effect, as a long distance graduate student, and steered me into electrical engineering, which is why I moved from UCSD to UCI and UC Berkeley. I could find no interest at UCSD in this at all.

In any event, my goal was to recast a theory from the foundations. So I thought, well, it works like this. In the most general case, you have a set of objects, and all possible subsets of those objects, what's called the power set. If you have n objects, you have 2 to the n subsets. I wanted to work with that set of all possible fuzzy subsets, which is infinite, even if n is finite. That's what's called the big sigma algebra. It's all in the algebraic structure of the sets. For me the big advance was when I saw that that sigma algebra had the structure of a unit hypercube. A brain state in a box, in effect. Maybe that was a triggering event or maybe it was the Libertarian cube. I don't know what it was. I remember thinking that a Rubik's Cube has eight corners, just as a set of three objects has a power set with eight objects in it. Any point inside a Rubik's cube corresponded to a fuzzy set of three elements, where each element belonged to some degree.

About this time, 1984, I ran into Robert Hecht-Nielsen. He, like me, made up the neural underground in San Diego, the neural fuzzy underground. He had designed a processor called a fuzzy associative memory. It really wasn't fuzzy, but it did have some outcomes where output values could take on a spectrum of values. Robert and I got along, We never fitted in at UCSD, we were never members of the PDP group, and when we tried to get in later, we were always persona non grata. So we were on our own.

Robert was at TRW then, at the AI lab. Very soon I moved to a smaller company called VERAC, the V stands for nothing, but the rest is Engineering Research Analysis Corporation. I quickly became manager of adaptive systems. Also I ran a local neural network interest group. We'd invite various people to speak, for example, David Rumelhart. It was all an underground movement.

Now to step aside here from the technical issues, there was also a political development. Something happened in 1985. I thought at the time the most important conference series was the ICAI and AAAI series of artificial intelligence conferences. That's where I met, physically met, Lofti Zadeh. I saw him being pummeled on the panel on uncertainty. I watched my hero Lofti Zadeh try to sell fuzzy logic to the AI experts in the section called "Management of Uncertainty in Expert Systems." They didn't buy it and, in fact, they joked about it. At the same time I would scan the "Proceedings" and hardly ever saw a neural paper. Robert and I were both separately trying to get papers published, in my case neural and fuzzy, in his case just neural. It never worked.

I remember the 1985 AI conference which was at UCLA, very close to home. This was at the height of the AI movement. I'd done some work with the Symbolics company, developer of the LISP machine. I'd developed a program for the government called ADBM, an Adaptive Distributed Ballistic Management system, which tried to organize the SDIO system [Strategic Defense Iniative Office, often called Star Wars] like a big capitalist system in the sky. The folks at Symbolics Graphics had worked out a videotape that was very popular, and brought me, the libertarian, to the headquarters of the Strategic Defense Initiative to pursue money. I began to see from that point that the way we pursued mission planning in the military was purely socialistic, as well

For me a turning point was in 1985, at the AI conference. This was the peak AI conference in terms of attendance, and certainly in terms of enthusiasm. It was held in our backyard at UCLA and my fuzzy friends were there and my neural friends were there. Robert Hecht-Nielsen and I went up to LA. Our papers were routinely rejected, of course. But what we saw there made a deep impression. Just the panache of the AI community. The two competing vendors of LISP machines, LMI and Symbolics both had separate limousine services take you to the Beverly Wilshire, another very large downtown hotel. There was all the food and drink that you wanted.

There was a big Malibu party from Symbolics. It was very posh, the structure of the tutorials, the plenary talks, the way the venture capitalists were running around trying to fund the field. And somewhere in there the idea began to emerge, it was sort of at a subliminal level still, why not do something like this for the neural or fuzzy fields?

That same year, after the AI conference, in the summer of 1985, the president of VERAC, my company, happened to be the IEEE San Diego chairman. He said, "How would you like to be Chairman next year, in 1986. The only requirements are that you say, I do, and that you be an IEEE member."

I wasn't a member. So he wrote me the letters and I became a member, and I said, "I do."

So when 1986 rolled around, I had a vote on the executive committee of the San Diego IEEE section. At the same time I was teaching a course at UCSD at night on fuzzy theory, in which I was developing a lot of my fuzzy ideas to the point where I was giving out homework problems on them. Often I would create a new theoretical idea, like the idea of fuzzy entropy, and within a week, I'd already assigned homework problems on it to the students, and for myself to explore. This was in January 1986.

Then I did something. The local section of the IEEE had about $30,000 in the treasury. They felt that each year they were eating that up by about $5,000, and they were slowly going broke. They needed something to raise revenues. I had the idea, well why don't we get a bunch of these AI vendors together, and we'll throw a mini conference. We'll call it the AI Slugout. We'll have LMI compete with Symbolics, and bring in all the smaller vendors that are selling AI machines. These included Sun Microsystems, Silicon Graphics, and other companies that have AI packages, the AI software tools people, and so forth. So at UCSD I rented Mandeville Auditorium for $1,000, courtesy of the IEEE, and brought together several of these vendors.

In April, 1986 we had the AI Slugout. It was on a Sunday afternoon. I'll never forget it. The curtain hadn't been drawn, it was to start at 2:00. I went behind the scenes and there were at least 20 people, most of whom were major vendors. The two biggest ones refused to show up, Symbolics and LMI, but the smaller vendors showed up and we had a show. I peeked ouside the curtain and there were exactly six people in the audience. There were far more vendors than attendees. So I told the vendors that the program had been delayed about half an hour. I ran outside with my friends and we began to grab anybody walking by and tell them they gotta come see this artificial intelligence conference and gathering. We got maybe 15, 20 people. That got a critical mass going. Other people trickling in. I think we put up a small sign and at one point it reached almost as many as a hundred people.

It was judged a success by the local IEEE section but more importantly, I had the vendors contribute, I believe $200 apiece. The net result was we not only paid for Mandeville Auditorium, but we also made a nominal profit. In the eyes of the local IEEE section this was a big stamp of credibility for me.

We'd been talking, always talking, about perhaps having a conference someday as a major revenue enhancer. They weren't sure about what. They were talking about having maybe something in manufacturing, where there already was a conference, and trying to co-sponsor it. In the first week I believe of June, 1986, "Business Week" had this famous article on neural networks. In it was a picture of Robert Hecht-Nielsen bending over his neural machine at TRW, and John Hopfield, standing with his arms folded in front of a picture of a neural network. Everybody was talking about it. Robert had thrown the first neural short course in the fall, attended largely by neural researchers, and had scheduled another one for later in that summer. I was about to begin a course at UCSD on the subject.

I sat at the monthly IEEE meeting, I had brought a copy of "Business Week" with me and said, "Have you all seen this?" I showed the pictures of Robert and John Hopfield, I said. "Point of fact, we have a very large neural effort here in San Diego, the PDP group."

At this point the PDP books were not out but we'd all seen advance copies of it and talked to them. Every week, "EE Times" had been running articles. I had some of those with me. And I said, "Why don't we throw a neural network conference?"

I put forth the motion, got a second and won approval to explore the issue. I called Robert Hecht-Neilsen and bounced the idea off him. He seemed enthused. I think we had lunch some time after that to talk about it. The next big step was when I brought Robert and his golden mouth to the next board meeting of the San Diego IEEE section. Robert and his magnetic personality convinced them that not only could we do it, but we could do it with panache just like the AI guys did.

That was the plan. We would structure this just like a big glossy 1985 AI conference. We would have the tutorials, the plenaries. We'd do it right in a big hotel. We'd have a party, a banquet, the whole thing. And most of all, we'd use the mailing list of the AI community. We would schedule the conference two to three weeks before the big AI conference. We'd beat them at their own game. Sure enough, the AI community sold us their mailing list, and we kept copies of the brochures and proceedings from previous conferences, and completely copycatted what we viewed as our competition.

Now the problem, of course, is that a local section, a city section of the IEEE has no authority to throw an international conference, not even a region wide conference, and yet we'd gone ahead and done it. The IEEE is a volunteer organization. If somebody has the initiative to do something, it's very hard to stop them.

By the fall of 1986 we had already printed a preliminary brochure. We went around to the many feuding factions in neural networks. The feuds were really beginning to heat up now. The PDP books were out. There were different camps and we felt the only thing to do was have a level playing field. Offer, capitalist that I am, a very lucrative tutorial fee, but conditional. The deal was, if the conference didn't make money, and nobody thought it would make money at this point, the tutorial presenters got nothing. If it did make money, they could make as much as $15,000. So there was a risk in the package, something that later caused us a lot of problems with the IEEE.

Many of the leaders of the neural field who we asked to participate and to give tutorials literally hated each other's guts but now had a joint, common self interest in this conference. We wanted something on a very broad scale in San Diego, something that looked a lot like the AI conferences. The idea I learned from Marxist politics is this that you don't just criticize a field, you don't just shoot holes in somebody's book, what you do is you build a second book. And then you shoot holes in the first book. That was the idea. The problem was that we didn't have nearly the money that you need to fund this big conference. In terms of marketing, it was all boot strapping. We had to commit to a hotel, the entire Sheraton. The only way we could get it was commit to the entire thing, to fill it up. That meant we had a legal liability of over a half million dollars, maybe as much as a million dollars. The San Diego section, with $30,000 in the bank had just put itself on the line for over a half million dollars.

We began to line up more and more of the scientific talent. I was technically the general chair, but we called it the organizing chair. We had Stephen Grossberg be honorary chair, but de facto general chair. We brought in all the major players in the neural field. We had various schedules, and cut off dates, and appointments. By late January, when my daughter was born, January 30, 1987, about that time, we were at the first cut off date. And we had done some linear prediction of attendees. We thought we'd have several hundred people enrolled by February 1. I think there were less than 20.

Now came the calls to close down the conference, including from the local IEEE, San Diego. People were getting scared. Worse, people had begun to hear about the conference in the IEEE. Who the hell were these people in San Diego who have put the IEEE on the hook for a half million dollars? And the reputation of the IEEE? Who the hell is a 26 year old puke named Bart Kosko? How many conferences has he managed? Who's Robert Hecht-Nielsen? On and on and on and on.

In time there were petitions circulated to close the conference. This occured as late as a month before we actually threw the conference in June of 1987. It was very much a classic success story. The great bulk of attendees registered in the last two or three months. Then we knew we had it, a smashing success. The rest is history. We had the conference, about 2000 attendees. And the San Diego section, which had broken a lot of rules of the IEEE, had become the richest section in the history of the IEEE.

All of a sudden, the game changed. And of course there was the risk package for the tutors. I made a lot of money, everybody made a lot of money. We were heavily criticized.

I also want to give credit. Right before the conference was held, in the spring of 1987, there was an emergency meeting about what to do about this conference. The executive director of IEEE, Eric Hertz, who happened to be a former head of the San Diego section, flew in and met with myself and my other PhD advisor, Professor Alan Stubberud from UC Irvine, who had been the chief scientist for the Air Force and was about to become director of Region Six of the IEEE. So they met and Merrill Buckley was there, who was about to become president of the IEEE. Merrill wanted to close us down and Eric Hertz stood up for us. He put his job on the line to Buckley, who was, I think, on the Governing Board of the entire IEEE, and said, "Do you want to fire me or not?" And Buckley backed down and the conference survived.

But something I'll never forget, when it was all over, when Robert and I were talking, with Stubberud he told us the golden law of bureaucracy: "It's easier to get forgiveness than permission." He was going to let us go this time, but don't ever pull something like this again.

And so it was a success. After the conference in June, and the formation of the INNS [International Neural Network Society], the field had congealed to a sort of stable state as we now know it. At that point, there came the question, what would be the next conference? The reason the IEEE had done it, and had in the end endorsed it, was because they knew there was going to be a lot of money when they saw the final attendance. It became a big dog fight. Different societies in the IEEE wanted a piece of the action. The IEEE is a massive bureaucracy, and it in effect took over the conference.

At the same time, Steve Grossberg's INNS set up the first INNS annual meeting in Boston. It looked like something of a competition would emerge. So the real interesting drama was how there ever came to be an ICNN '88, because the IEEE had formally cancelled it. During the summer of 1987, the IEEE wanted to proceed with the next annual conference. It was a question of who would run it. There were many fights about this. The one thing they agreed upon is that Bart Kosko and Robert Hecht-Neilsen would have nothing to do with it. We could live with that and that was fine. The problem was we felt they were about to kill the goose that laid the golden egg. They were going to let the conference series die. That was unacceptable.

So in September, October that year Robert and I asked IEEE what the status was, and they said, "Well we've killed it. No one can reach a consensus."

This was totally unacceptable so Robert and I decided to do it again. Supposedly there's an Al Capone saying, "Do it first, do it yourself, and keep doing it." I appeared as program chair for 1988 and Robert as general chair. We had our say. It wasn't just ego, although there was certainly a lot of that. It was more than that. What happened is, the IEEE had shut down the conference. Robert and

I said, "To hell with this. We're going to go to the 1987 NIPS conference and sign up as many people as we can for the next ICNN conference in '88."

We asked Teuvo Kohonen to be the honorary chair. He gladly did it. We set out with the same structures we had before, went down the list, filled out the slots, found the people we'd like to invite. By fall the IEEE had circulated letters, and different members had circulated letters, to every major society, asking them not to support Hecht-Nielsen and Kosko in this endeavor. There would be no ICNN '88.

Robert and I really took a big risk here, because if we failed we would be the Milli Vanilli of neural networks. So we went around and Robert at this point, of course, was head of his own company, the future IBM of neural networks. [Now HNC Software, Inc.] We signed up all the major neural researchers, just as we'd done before. I think Steve Grossberg didn't want to do it because he thought it would interfere with the INNS annual meeting but nevertheless we signed up all the major people. We would have had serious egg on our face if it didn't come to pass.

Once we had a slate, with all the intellectual fire power to do it the next year, and we had the hotel reserved because we had the reputation from the previous conference, we had the management structure, what happened is, my PhD advisor, my good friend and mentor, Alan Stubberud was about to become Region Six director of the IEEE. He was also now a boss at the National Science Foundation, which carries a lot of weight in the academic community.

There was a great screaming match held in his office at NSF, between Robert Hecht-Nielson arguing for the conference and the current president of the IEEE, Troy Nagle, who opposed it. They hammered out a basic deal, in which the local San Diego section got some cut, a future conference series and the IEEE got theirs. This was the basis of what continued as the ICNN series. In the following July, I guess it was, we had the conference. I was program chair and Robert was general chair. We gambled and won.

ER: I think you've given us a very complete history.

BK: So that's how the conference series began. To get back to technical developments, for me one of the big achievements of my career is the thing called the BAM, or the bi-directional associative memory. That was in 1985, back in those times when I was very much taken with the idea of global stability, Hopfield style networks, the ball rolling into the energy well. At the same time I thought the neatest idea in neural networks was Grossberg's adaptive resonance theory, that you learn only if you resonate. But to me the weakness of the adaptive resonance paradigm was that global stability was not part of it. You might be searching through an awful lot of grandmother cells before you resonated finally and learned the pattern.

At the same time I was fooling around with fuzzy associative memories. And these were fuzzy matrices that mapped fuzzy vectors or points in the unit hypercube into other points in the hypercube. I remember pushing a vector through a matrix and getting out a different sized vector on the other side and then pushing that back through the other way. When I pushed the output vector back through, I had to use the transpose or flip the matrix over and then a fixed point developed there. That was an interesting property.

I wondered what whould happen if you did this with a regular matrix, with linear algebra operations, and did the usual threshholding. So you push vector A through the matrix, and out pops B, and now transpose the matrix, push vector B back through and out pops A. I did it and, lo and behold, it always stabilized. It seemed for any matrix that was always the case. So this was one of these moments, these epiphany moments of scientific discovery, where now you have a theorem to prove. I quickly proved a simple version for discrete, additive BAM, showed the global stability, that is, any matrix always stabilizes.

I got very excited by this. I also did it to an auto-associative matrix, a Hopfield net, but in the Hopfield case you had to update one neuron at a time. In the BAM, you update the entire vector at a time.

I did many extensions of it. If this result was that robust, you should be able to change the weights slowly and then you have what I thought was something more like the real adaptive resonance theory, in the sense that you have both the neurons changing and synapses changing. You could extend the idea to what I called an ABAM, or Adaptive BAM and have a Hebbian learning law. The system always converged to fixed points. If you use the competitive learning law, it would still work. Now you have something that's very similar to the ART model. I kept extending results and finally got into a random domain. As long as you have a system perturbed by a noise of finite variance the system will always cool down. In effect the learning was structurally stable. It was robust.

That was one line of research. Along the way, I found a learning law that I think is the most important idea. That was the idea of combining the competitive learning with differential Hebbian learning. In other words, using not just a signal, but the signal velocity. The question had always been, how could a neuron or synapse compute a derivative? It's a very complicated calculation and very unstable numerically. But if you have pulses, as you have with real neurons, then it falls out very simply. The derivative is just the pulse minus the expected signal value. In other words you can estimate the derivative at any moment by whether there's a pulse. If there is, it's a positive derivative, if there's not, it's a negative derivative. That's a biologically plausible mechanism.

By this point also I'd worked out a lot of theorems on the foundation of fuzzy sets, all housed in the geometry of a hypercube. It was my approach to try to see things in math, because if it's a real structure you can always picture it. One inspiration was the brain state in the box neural model where the allowable state space of a neural system is bounded and is in effect equivalent to a unit hypercube. There ought to be a neural connection. The next level was to go from a cube to two cubes and mappings between the cubes. This was the idea of a fuzzy associative memory. If that mapping changed with time, then you had an adaptive fuzzy associative memory, or, if you like, a neural fuzzy associative memory. So each fuzzy rule defined one of these little mappings from one cube to another. If the air was cool, then it turned down the air conditioner a little. If the air was cold, then it turned it down a lot. Each one of those rules is a mapping. The fuzzy system has all those rules firing in parallel, converting inputs to outputs.

I began to pursue this, building it from cubes to multiple cubes. By this point, late '80s, Hal White had proven that feed forward neural networks could approximate any function if you used enough neurons. I was convinced you could do the same thing with fuzzy systems. I found a very simple proof. If you view a rule as a patch in the state space geometry, mapping from the input, trying to estimate the function that maps from input to output, that's just a curve that would go through some high dimensional space. You could cover the graph of that curve with patches and average the overlapping patches and that gives you back a fuzzy system. In my dissertation I'd worked with this averaging process. I'd developed my fuzzy integral. The dissertation was called "The Foundations of Fuzzy Estimation Theory" and I did it with Stubberud and Zadeh.

I went back to that work and extended it to this problem of function approximation. I was teaching the subject at USC. I had a class where I introduced, libertarian that I am, a capitalist mechanism. I had $1000 prize money. Whoever develops the best neural fuzzy project wins the prize, gets some local press. I got a neural computer dating service, and the usual applications to robotics, and lunar landers and a wide variety of fuzzy applications. Some were bawdy, some very clever, some were done on video tapes. Each time I'd teach a class, I would do a video countdown from the previous class. It was a place to try out new ideas.

I was convinced that I saw a quick proof that you could approximate any function with a fuzzy system. Unlike the neural proof, this is a little more constructive, because the patches were rules. We had geometrized a piece of knowledge as a rule and you could estimate those patches with neural networks. The rules could not just find the first patch, but tune them. In time we found a one two punch of unsupervised learning to estimate, and supervised to tune, was the best combination. We tried it out for a class, and applied it to Widrow's truck backer upper. Widrow showed that you could back up a truck and trailer in a parking lot into a loading dock with a neural system. We showed you could do that also with a small set of fuzzy rules and then showed you can convert any neural system into an epsilon equivalent fuzzy system, that had a similar input output characteristics, but with the fuzzy system you can open the black box, and you had a set of structured rules.

For me that brought together these lines of research that began with Carl Sagan's "Dragons of Eden", the works of Debrue, the economics, the libertarianism, the philosophy, all come together for me in what I call the FAT theorem, the fuzzy approximation theorem. It says you can always approximate any continuous or measurable function on a compact set to any degree of accuracy with a finite set of fuzzy rules. These rules can be very general. I now prefer to pick these as ellipsoids, or regions of attractability in mathematical analysis.

ER: Maybe we should talk a little bit about how you got into doing your recent book, "Fuzzy Thinking". It's for the general reader. It tries to describe these ideas, make them more accessible and has almost a buddhist point of view.

BK: I was, as I said, a writer all along, and always did freelance writing, a lot of political writing, very little of which paid. So in pursuing the libertarian writing and my writing of fiction, most of which is done under a pseudonym, I got into the discipline of writing every day. I write every day, I exercise every day. Get up in the morning, I pay my quota to myself. Write a certain quota of words, and exercise and then I go to the University to do whatever it is that I do.

That way I turned out my first book, a very large book, "Neural Networks and Fuzzy Systems", with a lot of my work and the work of other people. I edited a volume called "Neural Networks and Signal Processing", and now there is a third book coming out called, "Fuzzy Engineering" part of a three book package from Prentice-Hall.

Along the way I also began writing essays, philosophical essays, on the philosophy of neural networks and fuzzy systems, in the tradition, I thought, of the old philosophers, who would learn as much science as they could, stand at the periphery of science, and then speculate. That had always been the nature of metaphysics, and ethics and the like in the past.

I was asked to write a a series for "AI Expert", an AI magazine, in 1989. The managing editor called me up, and asked me if I was interested in having a regular monthly or bi-monthly column. We decided we'd call it "Meditations" and it could be about anything I wanted. I sent him several topics, and we agreed that the first essay would be an article entitled, "In Defense of God" because after having lost my faith in God, that personal God, I gradually regained belief in God of sorts. I wrote that essay and sent it in, and of course, there was a change of editor, and a very anti neural-fuzzy editor took over and, by God, there's no way Kosko's going to have his own column. Again I had the situation where somebody had to go to bat for me. The editor had to send the article for review to other people for an outside opinion. They did publish the essay in defense of God. But that was it. I would get no series.

But in time that essay secured me a contract for the popular book, "Fuzzy Thinking." A scifi reader read it, and some other people, and it led to a profile in the "LA Times Magazine". Different editors and agents read that article and around and around it went. They wanted a book talking about God and speculating.

Here was the argument about God. The strongest arguments against God had always been the fact that God could never be defined. That is, you can't say what he is. You can't say that he is. The one thing that the neural networks taught me is what they call recognition without definition. That is, you could recognize the pattern of a face without having the ability to define it.

Surely dogs and animals have this property. They can't articulate anything.

This idea is a negative one, that the inability to define God is in itself not a sufficient reason not to believe in him. There are other reasons you may not believe in God. The fact that your prayers don't work is the most popular. Or AIDs spreads, or the worst get on top and stay there, or science is enough, or whatever the reason happens to be. But just because you can't define God, that itself is not sufficient. That was point number one.

The second thing that both disturbed me and delighted me was that science seems to track math, but doesn't have to. The classic example is, Maxwell puts forth some equations for electricity and magnetism, you manipulate the equations and then, poof, out pops light. As a mathematical prediction you get a wave equation, and sure enough, we found that to be verified in experience. The same way, with the general relativity equations, what immediately falls out is a wave equation, hence the prediction of gravitational waves, or gravitons.

For some reason science tracks math, and it doesn't have to, logically. That impresses me. It seems to me that God, the power, whatever you want to call it, is the he, she, or it that wrote the math. The idea of the math maker. This is the one we take orders from. I think if this continues, you know, we take orders from a Pythagorean Theorem, these big theorems, in a thousand, a million, a trillion years from now, if that continues, I think the idea of the blue print in the sky will be a little clearer. We tend to recognize that pattern even though we can't define it.

So anyway, that was the content of the essay, the last line of which was, there may be no God, but the math maker in science is his prophet. That idea of speculating at the periphery of science intrigued some editors.

I had other essays. I'd written essays on the Buddha. I saw a sort of historical duality here, between Aristotle and the Buddha, between "A or not A", versus "A and not A". To Artistotle, the pink rose is red or not and the Buddha says it's both red and not red, to some degree. So around that historical boxing match, I cast a book, "Fuzzy Thinking", whose thesis is, everything is a matter of degree. Looking at the world through a set of grey glasses, and driven, motivated by this problem that Einstein and many others articulated, the mismatch between precise black and white math and science and a grey world.

ER: If you were advising someone, as you probably do as a professor, who is getting started in the neural network field, what is it that you tell them?

BK: If they're a graduate student, I encourage them to take as much mathematics as they can for training. But what sells a field is a vision. What sold Marxism was a vision of the state withering away and complete freedom. Some people are sold by the idea of becoming the next Steven Jobs of the field. A major in information science, neural networks, fuzzy systems may just be the means to do that. Other people are fascinated by the idea of really truly understanding the brain, or maybe building the Commander Data of the future.

I like to ask researchers, well, where do you get your ideas. The only answer I've ever seen that makes any sense is, you vary your input if you want to vary your output. Do lots of things. If you've gotta take drugs, take drugs. Take long walks, meditate, watch a lot of movies, learn a new language, read different books, argue the other side of the debate, anything you possibly can to vary your stimuli. And then you have to, as they say, keep the ass in the seat. You actually have to sit down and write. Do it in a disciplined way. So I think if people have that, have a certain minimal training in mathematics, the problem will take care of itself. Because neural networks are inherently interesting. And I've believe that will continue well into the next century.

ER: Where are neural nets now? What do you think the future looks like?

BK: I'm very skeptical of subsidized science. If I look at the fields of AI, neural networks, and fuzzy logic, I see an inverse relationship between government funding and commercial products. I heard estimates, that in the past 39 years, something like $100 billion dollars went into AI. And we all know there's not a product to show for it.

Neural networks, there's been at least $100 million or more, and finally, some products have come out. We called the first conferences as a gamble. We thought there would be products in three, four, or five years, get the the enthusiasm going and something would come out. Finally some things did. Pap smear recognizers, bomb detectors, lots of process controllers and so forth, it's taken a long time.

And for the fuzzy systems, there was in effect zero research investment and now billions of dollars in product.

Let me just tell you a little story. In 1987, after the success of the first neural conference, we were at the first NIPS Conference in Denver. Robert and I were soliciting, signing up people, to talk and to chair the sessions of ICNN of '88. One of those people was Carver Mead. There was a meeting about what should be done to get more money from the government. That was the question. Only one person had a dissenting view, and it was Carver Mead. He put forth a Gresham's Law. Gresham's law in economics says that bad money chases out good. His view was that bad researchers chased out the good ones. We ought not get any money at all from DARPA. I was very impressed by that idea and there's a lot of wisdom to it. When you pick winners you tend to end up supporting dinosaurs.

I'm a bit suspicious of government funding and worry about appealing to government agencies. This is something we did of course. The first conferences had government panels and each government agency got up and talked about the money they could give out and on and on. I think the best thing that could happen to neural networks, happened. That was the collapse of the aerospace industry, the end of the cold war and that set of training wheels that the field had begun to support itself with.

ER: Do you see the future as in some way bound up with fuzzy systems?

BK: I think that's inevitable. First off, there is the interplay between the two systems, the intertranslatable neural fuzzy systems. And I realize, and it breaks my neural heart, that I want to use the neural systems now just to tune a fuzzy system. But you can use either one as approximators. There is this problem with the neural system, that when you learn something new, you may forget what you've already learned. In the fuzzy case you can open that black box, study the rules, and see how they're changing. A neural system sees abstract patterns in the data, and those patterns are fuzzy sets, a concept like "cool air", for example, or the setting of the motor speed to "a little" or "a lot". In the next step, it begins to reason, or associate those patterns, into fuzzy rules. But the system itself, that turns inputs to outputs, is a fuzzy system. So I think, at least from an applications point of view, for many years the way to go is a neural fuzzy system.

I think an area of future application is neural fuzzy systems in the small, at the nano level. The so called nanobot. Viral swarms of little computers that can recognize the abstract pattern of a cancer cell or an AIDS virus and eat it and convert it into healthy nutrients for the other cells and, in time, to repair the cells one cell at a time. If you can fix up a smashed up car a part at a time, why can't you resurrect the body a cell at a time? In the bigger picture, we can reduce death to a problem of molecular engineering.

Maybe the neural fuzzy nanobots will help bring us back. I intend to find out. As of now I'm about the 350th person signed up to be cryonically suspended upon death. I've gone for whole body suspension. Most of my colleagues have gone for head only, the idea being that if you could resurrect the brain, its synaptic structure, from nano devices you could also regrow the body from the head stump too, from the information in the DNA. It was really my final conversion to materialism through neural networks that drew me to this. That I am my synapses. If I can resurrect those, repair those, and fill in missing links with some clever averaging algorithm, I may come back. I would not bet my life on cryonics but I am more than happy to bet my death on it.