

Los Angeles, California, USA
1:13:53
.mp4

Michael Arbib

An interview conducted by
Peter Asaro and Selma Šabanovic

[2 June 2011]

Q: You can start by just telling us where you were born and where you grew up and where you started school.

Michael Arbib: Well, I was born in England and at the age of seven moved to New Zealand and at the age of nine moved to Sydney in Australia, which is where I went to high school and my undergraduate work.

Q: Okay.

Michael Arbib: Do you want to keep speaking?

Q: Sure, yeah.

Michael Arbib: Okay, well, I should first clarify that I don't really consider myself as a roboticist. My career has overlapped robotics from time to time and I think perhaps influenced it from time to time, but I wouldn't consider myself one of the robotics people, as it were. So I really got started in moving away from pure mathematics, which was my first love in college in this direction by reading Norbert Wiener's *Cybernetics* which showed me that mathematics could be applied in a different way from classical Newtonian stuff but thinking about parallels between brains and machines. And about the same time I came upon Grey Walter's book, which introduced me to his tortoises, and so that was probably the first real robot I knew outside reading science fiction like *I, Robot* by Asimov. But as a mathematician I was very interested in automata theory and mathematical theory of neural networks and things of that kind. And then in my final year there was a neurophysiologist working at Sydney University on the visual cortex of a cat named Bill Levick, and he wanted to learn more mathematics, and so he taught me about the brain and let me see what they did to cats and I helped him with his mathematics. And it was he who introduced me to a paper called, "What the Frog's Eye Tells the Frog's Brain," which had come out only a little bit earlier in which it built on some of what I'd been reading on the basics of the cybernetics book, for example, the "Formal Theory of Neural Networks," by McCulloch and Pitts, and it turned out that Ledford and Maturana had worked with McCulloch and Pitts and had taken not the classic '43 paper but the '47 paper on how we know universals, which showed how pattern recognition might operate over a layered neural network, and they tried to test that in the frog. And in the end they did find an interesting layered network, although it wasn't the one predicted by the theory.

And so that got me interested in the frog as a place in which one could bridge between the actual biological data and the more formal theory that my mathematics had gotten me to. It also taught me that McCulloch and Pitts were at MIT until that time. Since I'd been reading their classic papers I thought they were still in Chicago. And also as part of my automata theory

I read a book called *Automata Studies* which was edited by Shannon and McCarthy who were on the faculty at MIT and had a very interesting paper by Marvin Minsky before he became virulently anti-neural networks. And so all these people, Wiener, McCulloch and Pitts, Shannon, McCarthy, Minsky and so on were at MIT, so that became the natural place to go for my Ph.D. work. I started working with Wiener because I wanted to know about cybernetics, but I also had a research assistantship with McCulloch, and so I was involved with that group. But then it turned out that Wiener wasn't really into cybernetics anymore. His big thing was statistical mechanics, so I was his Ph.D. student for a while, but when he left for his sabbatical and we didn't get much correspondence I switched to another mathematician working on statistical mechanics, and in the end I decided that wasn't a very interesting type of mathematics. So I ended up with a Ph.D. in probability theory, stochastic processes that has nothing to do with robots or brain in terms of my own career. But midway through the Ph.D. I went back to Australia and gave a series of lectures at the University of New South Wales, and that became my first book, *Brains, Machines, and Mathematics*. So that really set me up for my future.

And the one other thing I would add about MIT was that – and it just shows how much serendipity affects one's career. A friend from Sydney wrote and said he'd read this book on control theory, and the authors Newton, Gould, and Kaiser were at MIT. Would I go and talk to them? And until that time I had known nothing about control theory beyond what was in Wiener's cybernetics book. So I went in and talked to them and they said, "Well, since you're a mathematician you'll be interested in these two people working out at Lincoln Labs, Michael Ethanessiotis [ph?] and Peter Falb." And so I went to see them and it turned out they were developing the theory of Rudolph Kalman who invented the Kalman filter. And so after my Ph.D. I spent four months driving all the way around the United States visiting people and national parks and so on confusing the hell out of them because they all assumed if I was visiting them I was looking for a job, and then they'd start talking about a job and I'd say, "No, I just wanted to see what you were up to." Anyway, one of the people I visited was Kalman, who was at Maryland at a research center, but he was just then moving to Stanford so he asked me to go there as a postdoc, which I agreed to do. So after about six months traveling around Europe – no, actually six months in London working on neural _____ and then about four months traveling around Europe and the Soviet Union I went back to Australia and then finally got to Stanford. And so during that time I probably had my first real contact with robotics research because just down the road from Stanford was SRI at that time, still called Stanford Research International before Stanford disassociated itself and then Neals Neelson [ph?] and his group were working on Shaky the Robot and the strip system for robot planning. So I think that's where I got first involved with real work on robotics.

But my own research was divided between automata theory and mathematical theory of neural networks and then getting more involved with real brains in real animals, and in particular building on the work with "What the Frog's Eye Tells the Frog's Brain" I worked with a student named Rich Duday [ph?] and we asked, "What does the frog's eye tell the frog," so the idea is, okay. Here are some signals in the brain the original paper said that might be correlated with

prey and might be correlated with predator, but how does the animal decide if there are several prey, for example, which one to snap at? So we came up with a model, which I think might've been the first winner-take-all network. You can check with everybody else and find out whether there were earlier dates than the late sixties. And then with other students we modeled mammalian systems with _____, looked at cortical networks for depth perception with Curt Boyles [ph?], looked at the cerebellum and its relation to that. That's when we started interacting with David Marr. And then I developed something called schema theory where the idea was that just as in computer science doing everything in machine code is death. You need higher level procedures, so I felt that the brain needed not just a structural account in terms of networks and neural columns and so on but a more functional level of parallel interaction distributed in corrections. Those were the schemas, and the focus there was to really think about vision, not in terms of just recognizing what was going on in the world but in terms of how that recognition could extract the parameters needed to guide action. So what I was doing then has been characterized as computational neural ethology, ethology, the study of animal behavior, computer models of the brain, mechanisms underlying that. So in a sense I was learning about robotics, not contributing to it but developing a framework that was very much relevant to robotics as, if you will, embodied computers, which had both sensors and effectors. So that gets me through Stanford. Are there any questions you want to raise before I jump into the next phase?

Q: Well, so you were conscious of the applications to robotics. How much of your work was driven by an interest in being able to translate these neural models to an artificial system, and how much was just trying to model it <inaudible>?

Michael Arbib: I think that the robotics people I was interacting with at that time were either looking at the visual front end, what computer vision algorithms might influence the thing, or in terms of extracting a model of the environment, or they were looking at the high level planning in a very abstract space. So I think that my work was probably more complementary than tightly geared to it at the time. I wanted to know really how you have a system that is extracting the information it needs to control its behavior inspired by the brain talking to the roboticist but not looking for a platform on which to actually implement it.

Q: And in terms of the cybernetics scene at MIT when you were there was there a growing interest in trying to develop robotic systems there or to continue Grey Walters' work in any sense, or were they sort of working on different –?

Michael Arbib: I left MIT a long time ago, <laughs> 1963, so at that time under Minsky's leadership there was quite a lot of stuff in what you might call symbolic AI. Minsky and Papert were writing their book on perceptrons, which many people took as the death nail of neural networks, but I think just introduced complexity theory into the study of neural networks. We

had some early work on computer vision, but I don't think there was what you might call live robots at that time. I mean I've always claimed since that I was a robot created by Marvin Minsky, but –

<laughter>

Michael Arbib: – whether that's true or not I don't know. I don't know whether you'd know this but a few years ago a group in England published a paper on unfortunately a very simple robot, but it was called ARBIB, and the good thing about it was not the robot itself but the acronym because it turns out that I am an autonomous robot based on information from biology, so what you're looking at now is not actually a human, but an artifact that Marvin Minsky created in the early 1960s.

Q: But based on biology.

Michael Arbib: Definitely inspired by biology, yeah.

Q: Did you ever encounter Shannon's rats or any of the other cybernetic robots of the early days? You saw the tortoise as <inaudible>?

Michael Arbib: No, I never saw them. Yeah, I knew Shannon moderately well. I once went rowing in the pond behind his house and capsized, and he cycled out to rescue me because he had a bicycle on pontoons that he could take out onto the pond. But I talked to him more in terms of his mathematical theory of communication, so I knew about the rat study in relation to as you say Grey Walters' robot and other efforts, but I don't remember that as something that we were actively engaged in discussing, sorry.

Q: Was it just at that point the mathematical models and the neural models were more interesting? It seems that cybernetics had that sort of moment when they built a lot of robots and then they stopped for a decade or so, and then they started rebuilding them with the more symbolic AI except for maybe the Hawkins Beast, which was another cybernetic one.

Michael Arbib: Yeah. No, I'm just saying that there were obviously a lot going on at MIT so I can't claim to be giving you an accurate picture of the whole spectrum but certainly in terms of the people I interacted with there were the formal theory of neural networks and reliability theory over in McCulloch's group. There was the move into symbolic AI by the interest in perceptrons in Minsky's group and then a lot of control theory, communication theory, mathematics, so I was not actively engaged in robotics as such at that time. So as I say I think the first serious

conversations I had on robotics were with Neals Neilson when I went to Stanford where it was at the level of planning, so in some sense it was complementary to my interest in what are the neural networks whereby a computer model of the brain can recognize objects in terms of interacting with them, navigating around them and things of that kind.

Q: And what was the extent of your interaction with the roboticists there? Did you mostly have conversations or did you work with them on any of the projects at SRI?

Michael Arbib: Conversations.

Q: <laughs> What was the first project that you worked on that involved a robotic component?

Michael Arbib: That's the next stage, the University of Massachusetts. So in 1970 I left Stanford. I was recruited to start the Ph.D. program and chair the new department of computer and information science, and so in setting that up I explicitly added cybernetics to systems and theory as the three tenants, as it were, of the department. And so I think the turning point came at a 1979 conference on the brain, which was arranged by amongst others **David Ingold [ph?]**, who was one of the people I'd been interacting with who did frog experiments. And at this meeting there were two very influential papers about human and monkey brains. **Mark Genereaux [ph?]** from France had been studying hand movements in humans and looking at the way that the pre-shaping of the hand is coordinated with the movement of the arm. And so that led me to think through the issue of how to take the schema theory that had been developed more in relation to my visual control of frog behavior and think about how it could be applied to the visual control of hand movements. And so I put out a review paper that included my initial thoughts about that and fortunately or unfortunately it gave the impression that I knew about hand control, and so I got invited to an international meeting in Australia and then I had to produce, so I had two very good Ph.D. students, Thea Iberall and Damian Lyons, who worked with me to develop a theory of hand control.

And that work finally led into robotics in two directions. One is that it raised the question of, well, can we apply this to control of a robot hand? So we got hold of what was then one of the first dexterous hands produced, the Salisbury Hand, and started applying our algorithms to that. And Damian in addition to working with me on modeling the human hand took schema theory and developed a programming language called RS for robot schemas, and that's seen modest application in the robotics community, I would say, not massive but some. And then we established something called the Laboratory for Perceptual Robotics, and I think this was somewhat novel at that time in that a lot of concentration had gone into the things of path planning for robots, and then of course in industry by that time there were already a lot of robots, but they were doing stereotyped operation. So the notion of really thinking through on the basis of my animal studies how to close the loop of perception and action, I think that was still

relatively novel. And so some of the things that came out of that, I had a couple of students, John Prager and Daryl Lawton who worked on optic flow so that as you moved – here I'd been influenced by the work of a guy named Dave Lee from Edinburgh University who had been looking at the role of optic flow in animal behavior, so we said, okay. Can we go from his observations to actually thinking about algorithms for extracting optic flow in a way that could be useful whether it's a brain doing it or a robot doing it? And then a fellow named Ken Overton, who unfortunately left academia to work at General Electric and has never been heard from since, but he was quite the operator, and so he developed tactile senses that we applied to the Salisbury Hand, so we spent time picking up eggs with the Salisbury Hand and things of this kind. And then there was one other student, Steve Begich [ph?], I think, who continued the work on tactile senses.

So I would say that was my main phase of being as it were within robotics rather than having ideas that interacted with robotics. The Lab for Perceptual Robotics is still going strong. The current director is a guy named Rod Grupen, and in addition to having a lot more robots than we ever had he has developed these ideas of perceptual and motor schemas that were part of the schema theory. Another way we developed schemas at UMass, University of Massachusetts Amherst was that we had two people, Ed Riseman and Al Hanson, who had a computer vision group, and they developed some of my ideas in schema theory and I interacted with them to think about how to recognize complex scenes so that you would have different schemas specialized in recognizing different objects in the scene but passing messages to each other so that the different schema instances could compete and cooperate to come up with the interpretation of that scheme, and that architecture was related to the hearsay architecture for speech understanding that had been developed in the early to mid-seventies at CMU, and we hired Victor Lesser who had been one of the members of that group to join us, so that got cooperative computation, distributed artificial intelligence, call it what you will, up and running at UMass. And then the other thing, which has had an influence both in neural networks and in robotics and also in machine learning is that I met a guy named Harry Klopff who was working at the air force research center outside Cambridge, Bedford, Massachusetts, and he put out a little book called *The Hedonistic Neuron* and his idea was to think about brain learning by asking, "What's in it for the neuron?" So his idea was how could you have – the neuron was in some sense adapting, changing its synapses, to satisfy its own needs. And so I got some money from him to do research on seeing if we could come up with a real theory of that.

And I had two colleagues working on brain modeling at the time, Nico Spinelli, Bill Kilmer, and while I was on sabbatical after I'd got the money they hired a postdoc to come in and work on the project. His name was Andy Barto, and when Harry Klopff gave us the money one condition was that, "I know this very bright student at Stanford just getting his bachelors degree. You have to hire him as a research assistant on-on this grant," and that guy was Rick Sutton, and as they say the rest is history. Barto and Sutton came up with the mathematical theory of reinforcement learning that has been incredibly successful and influential. Oh, and then the one other person I should mention from my UMass days is Ron Arkin, who became very

intrigued not with the hand work but with the frog work in terms of the models we'd devised on how the frog recognizes its prey, how it avoids its predators, how it can detour around obstacles and so on, and so he took that and implemented that as a controller for a robot and that became the basis for his very successful career in robotics that he's pursued for many years at Georgia Tech. So, yeah, I'd say my peak involvement with the history of robotics is during that period at UMass. And then I moved here in 1986 and I was recruited in part by George Bekey, who was then chair of computer science here, and he had a dexterous hand as well, perhaps not quite as beautiful as the Salisbury Hand in that it had less independent finger control, but it was the USC/Belgrade Hand. And we brought over Thea Iberall as a postdoc to work with us on that collaboration, and so there was quite a bit of work between the three of us then on dexterous hand control in relation to studies of human hand control.

And then we had several Ph.D. students work with us. Tony Lewis did a thesis work, which we jointly supervised on evolutionary algorithms for locomoting robots, hexapod robots and so on. But I think it became clear after a while that most of the students we were getting – oh, and just one other student in relation to this. Andy Fagg did his Ph.D. with me on what's going on in the monkey brain in relating parietal and frontal cortex and sort of taking visual information and converting it into hand control, but he was also actively involved with Bekey's group so that there was a bridge there from the brain modeling we were doing at the time into the robot control, and then Tony Lewis was looking at the evolution of neural networks. And that actually lay in some sense a basis for a later thesis at the University of Edinburgh by Alkie Icepert [ph?] on how to evolve a fish as a lamprey's spinal generator for swimming into a salamander's spinal network that could do both swimming and locomotion on land, and he later came here with his then wife to do a postdoc, but that's a different thing. So there were some successes, I would say, from that collaboration, but it also became clear that we were not making as much progress as we should, in part because both of us were doing a lot of different things, in part because the students tended to be perhaps too purely computer science students, so we had a few lucky students like Lewis and Fagg who did enjoy tinkering, but we had others who didn't really know what to do if the robot stopped.

And so at that stage we were successful in recruiting a couple of very bright then young people to come in as assistant professors and basically handed over robotics to them. So Stefan Schaal in terms of the hand control, and now that's moved on to a lot of other issues where you're really thinking about how various learning algorithms can be applied to really look at the dynamic control of complex manipulators, skeletal systems and so on. And then Maya Matarich who came to us from Rodney Brooks at MIT who came in more to look at the social robots where you don't really need to know much about the guts of the robot. It's got a simple control. It'll scoot around, but how do you work on those controllers to make them interact? And then there was a good Ph.D. student of George Bekey's named Guarav Sukhatme, so around 1999, 2000 we promoted him from Ph.D. to postdoc to assistant professor, and he has become very dynamic since. So with Matarich, Schaal, Sukhatme and some interaction with Laurent Itti, who

I hired to work on the bridge between computer vision and brain vision, we really have a very strong robotics group here.

But I would say by now I'm just a brain guy and leaving the working with the real machines to them. My own interests were captured, I would say, in the late 1990s. We got some money to collaborate with an international group of people in Japan, France, and Italy on the neural mechanisms controlling hand movements on the basis of visual input, and Andy Fagg's thesis was a part of that. And then this group discovered something called mirror neurons, Giacomo Rizzolatti and his colleagues in Parma, which bridged between firing both to code the action the monkey is doing and to code an action the monkey is seeing someone else do. And then we here at USC with a brain imaging expert called Scott Grafton imaged the human brain, so we went there looking at individual neurons coding the action two ways because at least other chunks of brain that light up more for interacting with objects and seeing someone else interact with objects rather than just looking at the object alone. And we found that that seemed to be in Broca's area, which is traditionally thought of as a language area. So ever since I've been continuing modeling the mechanic brain but also developing a theory of language evolution that suggests that the mirror neurons might be at the heart of an increasingly elaborate set of circuitry that allows the human to learn and use language, and that is my life.

Q: <laughs> Have you seen the mirror neuron model being used in robotic learning systems?

Michael Arbib: Erhan Oztop, who was the first student to model mirror neurons with me after his Ph.D. went to ATR in Japan to work with Mitsuo Kawato, who's probably the leading person on the brain robot interface in Japan, and so they have done some work. I would say their mirror neurons in the robotic case are not as interesting as the brain mirror neurons. But there was also a collaboration between them and Daniel Wolpert in London on mental state inference where they would look at how you would have internal models for different actions and how you could deploy them in different combinations on the basis of the mirror neurons. We'll be running a workshop here at the end of July where we're bridging between models of action and action recognition, neural linguistic models and the neuron for that can support that. And one of the people we'll be inviting is Yiannis Demiris from London Imperial College, who is working on imitation in robots because in fact some of the modeling he has been doing overlaps with what Oztop and I did in thinking about the relevance of mirror neurons for larger systems that could do imitation. The mirror neurons by themselves just say, "I'll recognize an action I already know," so you need a bigger system to say, "Oh, there's something I'd like to do but I don't know how to do it. How can I add it to my repertoire?"

Q: Who are some of the other partners in this international collaboration? You mentioned Parma and ATR.

Michael Arbib: Yes. Okay, so I mentioned that my interest in hands started back in 1979 with the work of Mark Genereaux, so he was the French anchor for this effort, Rizzolatti in Parma, Hideo Sakata who looked at the parietal cortex. He was basically recording from the monkey brain to look at how vision was coded in a way that was relevant to action, and Rizzolatti's group was more looking at the frontal cortex to say, "Okay. If I get that visual input what do I do with it?" And then, Gennaro was more in terms of, "Okay, those ideas are great. Let's see how I can apply them in analyzing human behavior." And then we had another modeler at that, who's since gone on to do other things, but Michael Jordan, who did his PhD with Rumelhart [ph] in UCSD, where I met him because I was on sabbatical there while he was finishing his PhD. And then he, for a while, was very much engaged in motor control. And then he basically handed over to two of his PhD students, Dan Walpert [ph] I've already mentioned, and Gara Murney [ph]. And then he went into graphical models and machine learning, and is now at Berkley, I think, where he's done very well. So that was the-the initial group. And then we brought in other people to do some brain imaging in later years.

Q: And so Kawato became involved with all this just later on when [redacted] –?

Michael Arbib: Now Ka-Kawato is a different path.

Q: Okay. <laughs>

Michael Arbib: So back at UMass, I had something – I had founded something called the Center for Systems Neuroscience with funding from the Sloan Foundation. And in the first year of that, one of the fellows who came to spend a year with us was Shun-ichi Amari, who has gone on to become the most influential person in the mathematical theory of neural networks in Japan. I mean, that was many years ago, so – he came in 1977, I think. So Amari and I organized a meeting in Kyoto in 1982 on Competition and Cooperation in Neural Nets that, in part, built on the work we'd done together in Amers [ph], but also introduced a lot of other people. And one of the bright young things at that meeting was Mitsuo Kawato, who'd just finished his PhD. So as a result of that, I got to know Mitsuo quite well, and visited him in ATR quite a few times. And so, that led to the fact that Stefan Schall, who had been doing a postdoc for several years with Kawato – when he was looking for a university position. It sort of was on the basis of that longstanding arrangement with Kawato that I came to be able to rec – help recruit him here to-to USC. But also on that basis, several of my PhD students went on to do postdocs with Kawato, so the – N-Nicolas Schweighofer, who did a PhD with me on the cerebellum worked with Kawato for several years, and he's actually now on the USC faculty, but over in biokinesiology and physical therapy, looking at how to apply his modeling to rehabilitation. How does the-the neural net – how do the neural networks of the brain change after a stroke? How can we help the patient recover their functionality? Erhan Oztop then followed in those same footsteps, and then the third person, I think, from my group who went after his PhD was Jacob Spoelstra, another person

who'd worked on modelling the cerebellum. So that's-that's a longstanding relationship. And so, it's included discussions of the work he does on robotics, but I would say the-the serious interaction has probably been more on general modeling of control of hand movements, without worrying too much whether it's a-a monkey brain, a human brain, or a-a robot brain that's-that's doing the controlling, whereas-as you know, Stefan Schall really got involved in the-the detailed algorithms for control of the Sarcos, a robot – very beautiful humanoid robot.

Q: What do you think are some of the most significant neural models, or mathematical models, that have been influential in the evolution of robotics? So, common filters, specific forms of neural networks, what's your personal impression of – what's been some of the breakthroughs?

Michael Arbib: Well, modeling the cerebellum has certainly had quite an impact. There was the-the Cam model, which was directly influenced by that. Jim Albus, who did a PhD thesis, published in 1970 on the cerebellum, not only came up with that model, but was very influential at the National Institute of Standards, whatever the – it's NIST these days, I forget the acronym. So, thinking about the cerebellum, and then Michael Paulin, a New Zealander, was the one who really pushed using the Kalman filter as a model of the cerebellum, so that ties in with that brain influence work. But on the other hand, the Kalman filter is such a general mechanism that it's certainly been well applied. I mean, clearly, I would think optic flow algorithms, in terms of robot navigation, depth algorithms, more generally, whether they're static, stereo, or dynamic, optic flow has certainly been important. The work of Schall is typical of those people who are using complex, non-linear learning models to-to really get at the problem of, "If I have a robot that is very flexible, I can't afford to use a-a sloppy algorithm the way I can with a-a robot with a few degrees of freedom and-and very rigid joints," and so on.

So I think that's a – another area. And then, where people have continued work on planning, and that's probably an area where neural networks, to date, have not been as important as more traditional AI models. From my point of view, another thing that's very important is thinking about distributive control, so you're not just looking at the – a serial flow, in this case, you're really thinking about how to distribute control to different parts of the system. And it's interesting, there, that if we look at today's motor cars, they already have vast networks of-of computers on board, doing everything from optic flow to avoid collisions, to self-parking of the car, at the other end. So the interesting point there, I suppose, is that the field of robotics has moved, perhaps, out of its specialized niche into computer science, more generally, as we think about embedding computers in systems which have sophisticated sensors, and so on. So that the-the skills that go into robotics are not just for animaloid or humanoid systems, but for very diverse vehicles, as it were, which-which can be platforms for integrating sensation and action.

Q: In terms of optic flow, do you find the work of JJ Gibson to be informative or influential?

Michael Arbib: Well, JJ Gibson is where we started. Well, JJ Gibson, in terms of the general concept of optic flow, and then part of our work has been influenced by the general notion of affordances. David Li, in terms of looking at actual use of optic flow in controlling the-the dive of the gannet towards the water to catch the fish, and the-the judgment of pace length by people doing the long jump, and things of that kind. So those-those two set the stage on which our particular quest for, “How do we come up with brain-like parallel algorithms for extracting optic flow?” took place. The debate with that man – and there actually was a debate when I was at U-Mass. We invited Gibson to come and debate me, because he had this idea of direct perception, the notion that, somehow, the brain just picked up affordances, including optic flow. A-and so where the debate was drawn was, basically, I was saying, “Yes, I agree with everything you’re saying about – we pick up this information, but look, it’s not direct. And-and we have to understand the neural networks that-that carry out these performances.” So-so, conceptually, Gibson was tremendously important. Computationally, he wasn’t in the running, and so we-we built on his work in that way.

Q: What was his reaction to you?

Michael Arbib: He was a nice guy.

<laughter>

Michael Arbib: No, no. I mean it was just what – we-we had a-a debate, it was very well attended by the-the psychologists, as well as the computer scientists, and I think we-we both enjoyed the interaction, and I don’t think either of us budged an inch, but I think the audience benefited from getting a-a deeper understanding of what Gibson was about by seeing how he responded to my questions. But perhaps getting, also, a fuller understanding of-of what neural networks were about, if-if they were pure Gibsonians, who had just looked at what information is being picked up.

Q: Was it your sense that the people interested in robotics were less concerned about the symbolic AI-versus-neural-network paradigm struggle, and were more willing to engage with neural networks than the more traditional CS people that might have been more interested in the debate between neural nets and AI?

Michael Arbib: I think one has to be careful, because my sample is completely biased. So, in other words, I interacted with people who were interested interacting with me. And if they wanted to interact with me, then they wanted to see what they could learn in the way of biologically inspired robotics. So I wouldn’t claim any typicality there. I think other people were – and if you look at the people doing mobile robots, there they-they probably put on some laser

range finders and didn't worry too much about the-the details of the visual processing. The-the controllers were very simple for steering the robots, and so a lot of their work would be more at the planning stage. An exception to that would be Ron Arkin's work, where he really was interposing a more potential field approach, based on the frog stuff, and other people did pick up on that. And then, of course, or – in the early 80's Rodney Brooks came up with, what I thought was a reduced form of my schema theory, and what, unfortunately, most people thought was something brand new – I'm sure you get a lot of this in these interviews – where he had a layered hierarchy, but-but, to my mind, the-the whole point was that you had to have representations to close the loop, in general, if you-if you were looking, for example, at navigation, where you had to know the layout of a-of a complex environment. So – but a – but the point is, that-that work of Rodney Brooks neatly complimented what Ron Arkin had taken from a more biological point of view, and sort of created sort of field of behavior-based robotics, and so on. But probably, if one writes the history of it, Rodney's strict hierarchy of representation-free schemas was more influential, at the beginning, at least. I think, in the end, everybody realized, as-as Ron himself pushed, that you had to see how to integrate the delegation of many functions to perceptual and motor schemas.

But still, that had to interact with long term knowledge about the world, and in the end, I-I think it's a moot debate. You just say, "Well, in some cases, the robot doesn't need much of a representation of its world, in other case, it does." And so, to be good at robotics, you really have to integrate the-the stuff which is more easy to relate to neural networks of the perceptual and the motor ends to networks that may tend to be more abstract interacting schemas, or even high-level programming languages that can bridge into-into representations. And we've done some modeling in the years since of rat navigation, and looking at models of the hippocampus, and those models actually go all the way back to the same year that Shun-Ichi Amari was a visitor. We had a-an Israeli visitor named Israel Lieblein, and we were looking at mice, finding behavior in rats, which had gone out of fashion. So we went back to the classic literature, and came up with something called the World Graph, the idea that the rat was, on the one hand, worrying about the affordances, it won't go through if there wasn't space to go through, but it had to know where it had got food, where it had got other things, and that, coincidentally, came out a year before a classic book called *The Hippocampus as a Cognitive Map*. So, in fact, we had a model of *The Hippocampus as a Cognitive Map* before <laughs> the book which declared this as a major theme.

But we were already influenced a little bit by the early finding on place cells in the hippocampus. I think that's another way in which – I wouldn't say that study has been very influential outside a relatively small circle, but for me, it-it gave a way of thinking about a hybrid system where we had – certain parts of the system were developed as neural networks for the – as it were, the immediate, if you will, the Brooksonian [ph] interactions with the environment to avoid collisions, find a place to go, things of that kind, and a representation this World Graph of knowing about your world, knowing where food is, and so on. And another thing interesting about it is that – I mean, just trivial personal history, but this, in some sense, was reinforcement

learning before reinforcement learning existed, because we had – if the animal was hungry, then the reinforcement of food was what mattered. If the animal was thirsty, the reinforcement of water was what mattered. In any case, avoiding a painful electric shock was important. And so, in fact, much later, after the theory of reinforcement learning had been well established on the basis of-of the Barto/Sutton work, then we were able to improve our original model by formally using some of the theory of-of reinforcement learning to fit things that had been done in an ad hoc fashion on the first-on the first pass.

Q: So that was mostly sort of perceptual side, but do you see this also in the motor side, in terms of central pattern generators being controlled by representations, or grasping and trajectories of the <inaudible>

Michael Arbib: Okay. So the central pattern generators were always interesting. We had some interaction, and I think the-the UCSD sabbatical, where I met – actually no, it goes-goes way back. Okay. So way, way, way back at Stanford, there were two people working on central pattern generators in insects. Don Wilson who-who died rather young in a boating accident on the Snake River, I think, and Don Kennedy, who went on to be chairman of the FDA and president of Stanford. And so, the study of central pattern generators was always an interesting theme for us, but apart from Tony Lewis' thesis and Auke Ijspeert's postdoc, that hasn't been something that we've really developed, as to seen from being part of just the general culture that I share with my students. And then with the-the hands, the whole point was that it wasn't a central pattern generator, in the sense of Ron going with them to be modulated, it really was a discreet movement. And, in fact, one of Stefan Schall's papers is where he claims to have differential brain imaging for discreet movements versus rhythmic movements. So the – you've got a target, and then you're coordinating two controllers, one to get the arm to the right place to carry the hand, the other to get the hand in the right shape to grasp the particular object. That has been a research issue that we've looked at, both in terms of neural networks and, in one PhD thesis with a guy named Bruce Hoff, in terms of adaptive control theory, where we, again, interacted with the Gennaro group, in terms of data on what happens when you perturb the position of an object, or perturb the size of an object. How can the controller adapt to that? And we came up with a-a very interesting model that explains some rather strange anomalies so that, if the hand is coming out like this, and you suddenly switch the size of the object, it may, in some case, not stay like this and go into the new size of the object, but it may actually sort of shrink down and open up again. And – sorry, that's if you displace the distance to the object. So we came up with an optimality criterion for the coupling, to say, basically, it cost you to hold the hand open. So if you-if you're suddenly told, "I've got further to go before I grasp," then it may be worth your while to close the hand for a while, in terms of that optimality criterion. So that allowed us to explain, what, at first, looked like some very strange behavior that Gennaro and his colleagues had observed in their work in France.

Q: You discussed a lot of different ways about how your work on different biological systems influenced things in robotics, and I was curious if your interaction about your work in robotics, what kind of effect it had on the rest of your work.

Michael Arbib: Well, as I say, there was a specific period at UMass, with the lab for perceptual robotics, and then, when I came here, interacting with Becky, where we were actually looking at robots, as well as continuing the brain modeling. And there, there was a real conversation, thinking about the optic flow, in terms of how the robot could use it, fed into our thought about the-the other system. And then thinking about hapsis, of-of the sense of touch, was influence by the robotics. But-but I would say that, in my personal history, it is more trying to understand the brain, and the robotics as a source of stimulation and occasional application, rather than many other people you'll talk to, for whom they're roboticists, and they learn a little bit from the biology occasionally, rather than centrally.

Q: And now there seems to be quite a push for cognitive robotics as a new area of research and expansion. Do you have any comments on that, or is that something that's interesting at all?

Michael Arbib: Well, that's what I've done for 40 years, and –

<laughter>

Michael Arbib: If they want to say it's new, good on them, but – oh, it's the same with embodiment, and so on. I mean, this is what we've been doing ever since we started work on the-on the frogs. And suddenly people in AI and cognitive science are talking about embodied AI, embodied robotics, but – I mean, not embodied robotics. Embodied cognition. But really, this has been a central theme ever since the late 60's in my work.

Q: So what year did you create the lab for perceptual robotics?

Michael Arbib: 1980-ish. Yeah, probably. Might have been a couple of years earlier, but 1980-ish would be right.

Q: And what year did you come to USC?

Michael Arbib: 1986.

Q: And the perceptual lab is still going?

Michael Arbib: Yeah. Well, it – when I was last at UMass, about three years ago, it was. And a guy name Rod Grupen was, at that time, leading it. And not only, as I said before, with a much wider range of robots than we had when I left, but nonetheless, building on the schema theory that I had initiated. So it was nice to see that those ideas that I had developed, that Ron Arkin developed, that Damian Lyons, who came up with our RS schema language, that that work that had been – from the early 80’s, basically, at-at UMass, that the person who had taken over the lab for perceptual robotics was, in some sense, building on that tradition, rather than just doing any old robotics, as it were.

Q: And did you actually collaborate with Ken Salisbury at all, or you just worked with his hand?

Michael Arbib: Well, in terms of getting the hand and learning how to use it, we interacted with him. But we didn’t, I think, feed back into the design of a new generation of Salisbury hands, or anything like that. It was more that we were customers and he was pleased to be informed of what we were doing with it. And I’m sure we showed him the touch sensors, but I cannot remember, at this stage, whether he said, “Oh, great. I’m going to add those touch sensors to my hand,” or whether he said, “That’s nice, thank you.” That-that is gone from my synapses.

Q: What do you think are some of the future challenges in the field of brain robotics? What do you think are some of the things that people are coming across now as big issues that they might want to deal with, or some of the things that are going to be coming up in the next five to ten years?

Michael Arbib: Just about everything. I mean, in terms of visual perception – when I worked with Hansen and Riseman – and this work peaked probably about 1980 on the visions system for seeing and perception – I thought that, in the next few years, we would have a general system for analyzing objects in the scene, and how they move. It’s not happened. The-the funding has all gone into, “Build me a very special purpose vision system that does a particular thing.” So I think general scene perception is still remote. We’ve seen a lot of progress lately in object perception, picking out a particular object in the scene. Mirror system stuff is, “How do you recognize what a hand is going to do to an object?” More generally, “How do you recognize actions?” I think we-we have very little guidance, and especially if we’re concerned with robotics for social assistance of-of patients and children, and others. Being able to flexibly recognize what’s going on around you, not just what objects are there, but what humans are doing, so that you can either get out of the way, or protect them from harm, or assist them in some way. I-I think that’s going to be – so I’d just say, on the-on the vision side, just that general thing of-of dynamic – what we might call dynamic scene perception, to be able to recognize agents, objects, and the actions that link them.

On the motor control side, I would suspect that the sort of work that I associated before, with Stefan Schall, on humanoid control, and so on, will continue. A lot of work in mechanical engineering of different systems. As I mentioned before, the question of how much of the robotics of the future will be so strictly humanoid, and how much will be able to, as it were, exploit different mechanical systems, I don't know. So that's a – the mechanical engineering end and the control theory end is tremendously important, but that's probably rather separate. Everything in between, you know? Planning is still pretty primitive. It's fine if you've got a very constrained environment, but going beyond that – and then language. I mean, that's not really robotics, but if we're going to have robots that are built, as some of them will be, for human interaction, then a decent language interface that can pick up on ambiguities, and use the current context to-to remove some of those ambiguities, are going to be very important.

So I think the general point is, I-I see robotics as blurring. It's not really a separate subject anymore. Just as computer science has blurred. When I-when I set up the department at UMass, when I persuaded some skeptics to allow Becky to hire me here, the view of computer science was pretty narrow. When we do operating systems, we do a little bit of theory of curing machines. And beyond that, it's not computer science. And now, we-we have – one of the things I initiated here was a seminar for our PhD students, just each week, to have one of our faculty come in and talk about their research. And I – for the first year or two of that, I would sit in on the lectures. And – I don't know. The barriers have gone. We-we've got Paul Debevec, who, as his computer science work, came up with a system for facial motion capture that powered the success of the *Avatar* movie, in part. We've got people doing senso-networks, so Gar-Gaurav Sukhatme is both looking at mobile robots, but he's also looking at what happens if you scatter a few thousand sensors in a bay to monitor the ecosystem. There's no line between those two-two projects. So the idea that computer science is no longer symbolic manipulation in a box, but-but embodies all these techniques of perception, all these diverse forms of effector control, and so on, means, I think, that-that, although there will still be a certain class of computerized devices that people will recognize as robots, probably the more humanoid ones, but the expertise that is in robotics will not be segregated from this very dynamic network of-of computer science research areas that's emerged in the last 20 years, oh, I'd say.

Q: And you yourself have had a very trans-disciplinary career in belonging to different parts of the institution. How did you kind of navigate that? Because – <laughs> if it makes sense. Makes it challenging.

Michael Arbib: Talk very fast.

<laughter>

Michael Arbib: No, I-I mean – I think there’s always a danger that, if you do too many things, you end up not being good at any of them. So I think the answer has been not to try everything all at once. I mean, in fact, I know when I came here 25 years ago, I dropped my work in-in algebraic theory of computation, for example, which was a sad thing to do. But I just couldn’t establish the program here. And then, as we said, later on, it became prudent to hire these great younger people like Schall and Mataridge [ph], and then Sukhatme and Iti [ph] to take over the robotics. So I think that – and now I’m obsessed with the evolution of language. I’ve always worked on linguistics, off and on, since I knew Chomsky as a graduate student – when I was a graduate student, not him. But I-I think that’s probably – that-that in any particular year, there have been real foci, grounded by having very PhD students, and – but at the same time, the base of knowledge is increasing, so that opportunities come up later. So, for example, I had a group of four PhD students at UMass working on schema theory approaches to language, that came to an end when I left there in ’86. But then when – about 11 years later, we did the brain imaging that said, “Oh, look. These mirror neurons in the human appear to be in Broca’s area, language area.” I was then ready to be able to-land on my feet running in returning to language, but now, tying it into an evolutionary scenario, as well as trying to think through whether the scenario gave us a new approach for thinking about the human brain, and neurolinguistics, which it has done.

<crew talk>

End of MichaelArbib01.mp4