

Los Angeles, CA  
0:41:20  
m4v

Stefan Schaal

An interview conducted by  
Peter Asaro

12 November 2010

**Peter Asaro:** Where were you born, where did you grow up and where did you start your education?

**Stefan Schaal:** Okay. First start talking?

**Peter Asaro:** Yeah, introduce yourself.

**Stefan Schaal:** Okay, so I'm Stefan Schaal, where shall I look at, your face normally?

**Peter Asaro:** Like this to the camera, and try to be pretty close to the camera.

**Stefan Schaal:** Okay. So I'm Stefan Schaal and I was born in Germany, in Europe, and I got my primary education at the Technical University of Munich, where I did my diploma and my Ph.D. in mechanical engineering and artificial intelligence. And I simultaneously doing my Ph.D. also it was at MIT, where I was in cognitive sciences and artificial intelligence.

**Peter Asaro:** When did you first become interested in robotics?

**Stefan Schaal:** I think during mechanical engineering I liked the dynamics of robotics already very much, so that was kind of, I liked moving systems and controlling moving systems. And then when I got my Ph.D. I got increasingly more interested in how to build these systems and also how to add intelligence to them, which kind of added the artificial intelligence component. And then with MIT, when I was there I was essentially, I was in brain and cognitive sciences in the artificial intelligence lab at MIT. When I went there it was actually the late 1980s, this is kind of where the new wave of neural networks and computational motor control and, kind of intelligence robotics was created, to some extent at least, there was other universities as well of course. And with this kind of excitement I got very interested in robotics as well. Actually I got interested in both biological motor control and robotics at the same time and tried to understand the biological components from a technological viewpoint and the other side, tried to understand from biology a little bit what we should be doing in technology.

**Peter Asaro:** What was the first robotic project you worked on?

**Stefan Schaal:** The very first robotic project. So when I arrived at MIT, so I started to work with Chris Atkinson who was a professor at MIT at this point, and the first thing we did is essentially work on little juggling machines so the very first thing I did was actually a five ball juggling machine, which was called the bounce jugglers where the balls basically bounce on the

floor once before they come up, and it was originally invented by Claude Shannon actually. Claude Shannon had a lot of toys in his basement or wherever he kept them and he liked juggling machines so we basically, he created a three ball bounce juggler and we basically created a five ball bounce juggler, it was a very primitive robot, but I tried to make a point that you can create some very complicated behavior with an extremely simple machine, actually it had no computer at all, it was just a machine. And then later on we started to do more intelligent machines which were bouncing balls and tennis racquets and we do devil sticking which is, you know devil sticking is a more complicated manipulation task. And these were the first projects we worked on. And then it continued essentially in 1992 with a collaboration where we started with a Japanese research lab, the ATR labs, and the person in charge over there is Mitsuo Kawato who's been a prominent person in computational motor control both with biological and robotics interest, and there we started actually to work with anthropomorphic robot arms which were very expensive at this time, and he was capable of affording some of these robots and I started to actually spend a lot of time in Japan and work with these robotic systems. And this was then starting to be about manipulation, doing complicated tasks with robots, fast tasks, catching balls, dancing movements, manipulating the world.

**Peter Asaro:** Who did you study with in Germany, in Munich?

**Stefan Schaal:** In Germany this was the Department for Engineering Design, that was Professor Klaus Ehrlenspiel who's more known in the German environment, he's truly a mechanical engineering design professor focusing about how to be creative in the design process and working with computer-aided design systems and things like that.

**Peter Asaro:** Who did you work with at MIT?

**Stefan Schaal:** The prime person was Chris Atkeson who was there as a professor. Other people in my environment included Marc Raibert who's well-known for his hopping robot, and Emilio Bizzi who is a very established neuroscientist in motor control of monkey neurophysiology.

**Peter Asaro:** Did you take neuro-models and put those into the robots?

**Stefan Schaal:** What we tried to do it much more understanding the principles, how the brain generates motion and that is. Okay, so going back a little bit, in the late 70s there was a very prominent vision researcher you might know of David Marr, who basically started to try to put more structure in the way how you address problems in terms of talking about theory, algorithms and implementation. And I think what we've been doing in motor control at MIT was primarily we were interested in algorithms and theory of motor control, and then basically looking at

implementations in neuroscience as something which we would like to understand, and implementations in robots which we would like to synthesize. And that was kind of the back and forth, but there's a branch of neuroscience, which is systems neuroscience, where people really try to understand how does the brain deal with the physics of the motor system, and how can it learn, and how can it basically get better, how can we do imitation learning and how can we really bootstrap autonomous motor systems out of some basic computational principles.

**Peter Asaro:** What were some of the central problems that you encountered?

**Stefan Schaal:** I think what we were very interested in at this point, which we still are, is actually how to bring learning and autonomy into robotic systems. So, there was not too much learning going on in the 90s and still, if you look, the field of robotics has two branches, the one is mobile robotics, the one is what is called manipulative robotics. I think in mobile robotics people have been doing much more machine learning techniques towards data filtering sensing and kind of state estimation in general what people would be calling that. Learning in manipulator robotics was less, it's also more annoying because you work with real time systems in much more stringent ways, so a mobile robot can stop and think, a walking robot cannot just stop and turn off its motors it just falls over. So getting machine learning techniques into these systems was really something which we cared about and lots of the work we did focused essentially about skill learning, how can you learn a complete motor skill from like throwing a basketball into a hoop or doing devil sticking, or grasping a cup on the table, how can you learn that, how can you teach a robot how to do these things totally by itself, and also improve by itself in the test.

**Peter Asaro:** What do you think the big breakthroughs have been over the last couple of decades?

**Stefan Schaal:** Big breakthroughs. Faster computers. Besides that I think machine learning has matured tremendously, in comparison to the late 80s, early 90s to what we have now it's not a field of statistical learning, machine learning which used to be neuro-networks in its original idea, and what we can do in terms of data processing in the kind of empirical inference on data has become profoundly solid in good work. That's one component. I think getting into new ideas, how to approach robots, I think. Okay, inappropriately telephone. I think the idea of imitation learning and basically learning how to bias robots was a very important component in robotics that you, instead of just try to preprogram everything from human intelligence you show a robot something and it can repeat that, it's an idea which is old, it has been around since at least from the beginning of the 80s, but it found a very new wave I think at the end of the 90s and has now become a very prominent field in robotics again. Another component which I think is important is to go away from, okay robotics is plagued by the idea that people always reinvent how they approach a particular motor skill or manipulation task, they come up with new

representations, they come up with new maybe adaptive or machine learning control algorithms, but from one task to the other there is not a lot of similarity so you start from scratch. And I think people, at the end of the 90s, early 2000s, started to work a lot on how to create uniform representations, and one key word here is kind of the idea of motor primitives, how can you find some generic building block of motor control which you can reuse in as many scenarios as possible, which basically boils down to a representation which has some parameters which can be reused, but which can also be adapted to the new task in an intelligent way. I think that was very useful. Model based robotics, which has been around also for a long time, is a very important topic, the idea of getting away from high gain control as originally in robotics to what is called low gain control and using internal models to get the control accuracy, this is I think currently one of the most important topics to get robots out in the real world. If you want to collide with a robot when it's out there in the world it should give in, it should not be stiff like an industrial robot and hurt you, and in order to create this controllability you need very good models of the system which you often can only get with machine learning techniques if the system is sufficiently complicated.

**Peter Asaro:** What are some of the biggest problems that you're facing right now?

**Stefan Schaal:** What is? Okay if we really want to go and bootstrap a robotic system out of machine learning we still, I think having not all the tools we need to learn reliably, so I personally care a lot about kind of this called black box learning, that you can have a learning system where you feed in data and you're guaranteed to get a good and reasonable outcome. That does not work in general in machine learning so far, in particular not in real time learning for robot systems. The robotic systems, when they learn have this big problem that they have to explore the environment in order to gather data and then make, from this data better, get better ideas how to improve, and that is still a very complicated process, to guide this exploration, how to explore, robotic systems live in, the systems I care about live in very high dimensional spaces, that means just random guessing in those spaces is totally impossible, and you need to really find out ways how to explore the environment, how to extract useful information and then insert this into your control mechanisms to get better. So I think this is remaining a big problem. In general I think we're still having a problem with actuation and power supplies in robots, we are far away from humans and it's a little bit handicapped by this technology, which is something which hopefully will be resolved in the future. And getting reliable and fully autonomous perception into the system is also something which has not really been accomplished in a sufficient way, and then getting everything together, like a complete perception action loop where perception tells the robot what to do and actions then tell the perception system that it needs to do something else otherwise we don't perceive well enough or we don't, we haven't acted appropriately, and getting a complete autonomous system going is still quite a bit far down the road.

**Peter Asaro:** What was your Ph.D. project at MIT?

**Stefan Schaal:** What I worked at MIT was primarily about using machine learning techniques for control, it was not my Ph.D. project, my Ph.D. project was still in Germany which was about computer-aided design and artificial intelligence, MIT was really about machine learning for robot control and understanding biological motor control with similar tools.

**Peter Asaro:** Where did you go after your Ph.D.?

**Stefan Schaal:** After my Ph.D. it basically became primarily machine learning and robotics and computational neuroscience, basically going between these three areas, doing behavioral experiments, partial imaging experiments, humans, to study how humans actually accomplish certain motor behaviors using data analysis techniques for machine learning to look at biological data and using all those insights to put them into new technology, into robotic systems. So a lot of things which we've developed in the last 15 years I would say are actually motivated by biology like, there's a lot of imitation learning which we've been doing, representations which are useful for more motor primitives and machine learning techniques which mimic some component of what the brain might be doing which have become not purely technological approaches but they got their ideas often from biological research.

**Peter Asaro:** And you went to Georgia Tech also?

**Stefan Schaal:** I went to Georgia Tech for a short time, after my postdoc at MIT I was a research professor at Georgia Tech, but simultaneously I spent a huge amount of time in Japan so I didn't really spend a lot of time at Georgia Tech at this time. I was simultaneously also a research professor at Penn State University in psychology for behavioral experiments, I did that with a collaborator. But during I think the times 1994 till, actually until 2002, I spent most likely four months a year at least in Japan.

**Peter Asaro:** What was the research institute in Japan?

**Stefan Schaal:** ATR is a telecommunications research lab but it's divided into several branches, one branch was basically the department of Mitsuo Kawato who is working on computational neuroscience and basically computational approaches, neuro-network approaches, learning approaches to motor control, and who's always had a very strong interest in trying ideas from computational neuroscience on real robotic systems. So they were, I think in 1996 we got a very big funding program, Kawato essentially got this program which is called a ERATO program, it's like a science and technology incentive of the National Science Foundation, which gave us five year of generous funding to basically start working actually with humanoid robots to scale up a little bit, and I think in 1998 we had one of the first humanoid robots in the world, I think

Waseda University has all the robots and Honda came shortly afterwards out with their robot, but that was kind of the way for humanoid robotics actually started to gain attention.

**Peter Asaro:** And they received a lot of funding from the Japanese government for that kind of research, or was it mostly university or corporate funding?

**Stefan Schaal:** No, the ATR is funded by some Japanese basic research agency similar to the NSF, so we were funded by this Japanese Science and Technology Corporation, and that's basic research, not industrial money. That was the key funding we got so when I got, started my faculty job here at USC it was essentially from the Japanese point of view that, I was still part of the Japanese funding system in this ERATO project, and for them it was just like taking a group inside of Japan and putting it abroad, which is actually permissible within this project, and for USC it was just an external grant. But it was only this one source of funding which I've been getting, it was always run through ATR essentially that I could get this money.

**Peter Asaro:** When did you come to USC?

**Stefan Schaal:** It was 1997.

**Peter Asaro:** Can you talk a bit about what the robotics lab was like when you got here and started your research?

**Stefan Schaal:** Let me see.

**Peter Asaro:** Who was here.

**Stefan Schaal:** At this time, I was simultaneously hired with Maja Mataric, who you interviewed just a few days ago, and there was primarily people doing robotics George Bekey, who basically worked on many different branches of robotics, he can most likely talk about this much better, but he was interested in mobile robotics, he was interested at some point in grasping and manipulation, he was also interested in biomedical engineering. I think he founded the biomedical engineering department here at USC at some point. Another person, professor working here was Ari Requicha who at that point started to move towards nanorobotics which was a very new wave, very aggressively new idea in robotics. And Gaurav Sukhatme was also shortly later starting his career here and moving from mobile robotics more into sensor networks. So it was a nice but a small environment than we have here, I think in comparison to about 12 years ago I would assume we at least doubled or more the robotics program here.

**Peter Asaro:** Do you do a lot of collaborative projects with the other faculty here?

**Stefan Schaal:** We have some joint programs so with Gaurav Sukhatme we actually have currently a joint program on grasping and manipulation and particular in novel and unknown environment, it's about basically manipulating objects and grasping objects which you've never seen before, and the idea is that you learn to do this in a test scenario, or in a training scenario I should say and then later on in a test scenario you get confronted with new objects that you hopefully can generalize to those. And the idea is to go from simple grasping objects, putting them somewhere, to extreme things like take, putting your hand to the gym bag, fumbling around, figuring out that's the triangle.

**Peter Asaro:** Do you collaborate with Michael Raviv at all?

**Stefan Schaal:** Michael and I we collaborated initially a little bit on topics of computational neuroscience which were closer, so we talked about behavioral experiments, we wrote some grants together but it was never, actually some of his students worked with me much more so we worked more through students together rather than other things, which was about doing simulation studies together and basically exploring certain aspects of computational motor control. Particularly imitation learning is something which Michael, through his interest in neurons, was very interested as well.

**Peter Asaro:** Do you have other strong collaborations in the US?

**Stefan Schaal:** I have a strong collaboration with Chris Atkinson who is now at CMU, so that's just my former mentor and we've had basically never-ending collaborations in the end. Let's see, other, some collaborators at MIT occasionally, I have actually a fair amount of collaborations abroad still, which is people who used to work with me at some point like the EPFL in Switzerland, various people, now the Max-Planck Institute in Germany. There's other people who are in Scotland actually to be precise.

**Peter Asaro:** Which Max-Planck Institute?

**Stefan Schaal:** That's Max-Planck for Biological Cybernetics in Tuebingen.

**Peter Asaro:** How many Ph.D. students have you graduated?

**Stefan Schaal:** I think I've graduated about six Ph.D. students and I usually have, in my lab, on the order of, between four and six Ph.D. students and then two postdoc somewhere, so it's a moderate size lab.

**Peter Asaro:** Are most of your former Ph.D. students still working in robotics and academia?

**Stefan Schaal:** They're not all robotics people, they are machine learning people, they're people who are interested in computational motor control, they are, and where are they? So, one, two, three, four of them have faculty positions in different universities, others are postdoc for instance at CMU or in Canada, others went to industry, actually the minority went to industrial jobs.

**Peter Asaro:** One of the things we're looking at is the mentorship networks and operation networks. What do you see as the central professional organizations for the kind of work that you're doing in robotics?

**Stefan Schaal:** Central professional organizations, so.

**Peter Asaro:** Or conferences or journals.

**Stefan Schaal:** Well, from the world of robotics it's a relatively easy game to play so it's, IEEE International Robotics, Conference on Robotics and Automation ICRA, then IRIS as the other big conference in this area. The new conference, which I helped cofounding is Robotic Science and Systems which has been around now for five years, which is a smaller conference, little bit pushing reviewing standards in robotics and trying to highlight kind of the coolest robotics research in every year. I spent also a lot of time at machine learning conferences like neural information process systems, NIPS, and ICML. And then lots of other little conferences, including actually Society of Neuroscience meetings and neural control of movement, those are the kind of primary meetings for computational motor control in biology and neuroscience.

**Peter Asaro:** What are the robots that you're working on right now?

**Stefan Schaal:** Okay, lots of robots. So here in the background we still have some robots, this is, these are old arms from Sarcos which are hydraulic robot arms, one is a master robot, the other one behind me is a slave robot, we are interested in ocular motor control and active vision so we have heads. We are getting actually a complete walking humanoid in the next few months from Sarcos, which is a hydraulic robot with, just similar to human appearance, legs, arms, heads, fingers, for grasping and manipulation. We, just a few months ago got a PR2 from Willow Garage as a personal robot to work with. We have some small scale educational robots,

the NAO robots which you might have seen from Aldebaran in France. We have some Barratt arms for studying manipulation and grasping, particularly interested in tactile manipulation and tactile perception, in the context of grasping and manipulation. And then in about one and half, two years, well one and a half years, we will get a Korean robot which is the HUBO robot, which is a humanoid robot about, I think one meter and 60 high, which is part of a National Science Foundation major research instrumentation grant. I think that's pretty much all of them, it's a lot of robots actually these days.

**Peter Asaro:** And these dog robots?

**Stefan Schaal:** Oh the dog robot yeah, that's a DARPA project which ended just about a year ago, which was about a machine learning project, how to use machine learning techniques to teach a robot how to walk over rough terrain.

**Peter Asaro:** What are some of the crucial problems facing grasping and manipulation and tactile?

**Stefan Schaal:** Okay there's a few things, one is dexterity, that means having squishy fingers which can give in and using actually a lot of force control, how to manipulate objects. Tactile sensors are still relatively few in hands, and there's no really super good robotics hands out there so far. And learning how to basically now use massive perception from tactile sensors, force sensors, and a hand together with appropriate receptive sensors in the arm, and use those signals to guide manipulation, to know when, manipulation is not just grasping something and putting it somewhere, it's really, if it's not good can you re-orientate it? If you notice it slips in your hand can you re-grasp, can you put it down? If you need additional information can you just rotate it in front of your eyes in order to get more information? I think working with a huge amount of sensors, and using them intelligently to guide grasping and manipulation is a very interesting topic, and I think this is partially what we have been working, starting to work on in this one new project which we're having.

**Peter Asaro:** Have you been working on grasping and manipulation prior to this project?

**Stefan Schaal:** We've worked a little bit on that, it's part of our other robots, they all have hands. We've never had a hand with tactile sensors so it's, for the first time we got one of the new Barratt arms which has a hand and they have tactile sensors in all the fingers and the thumb, and that will give us some very new amount of information. We also have a collaboration with a professor of biomedical engineering, Gerald Loeb, who developed some new tactile sensors for the fingertips which are very similar to human skin sensors, so it can, they can give you heat perception, they can give you pressure perception, they can give you sheer perception, they have

actually quite amazing properties, and we are just starting to put them on a robot arm and try to explore those. Don't worry.

**Peter Asaro:** Have you also worked in walking robotics, you said you were getting a full humanoid?

**Stefan Schaal:** We've started to work with walking robots, well with quadruped walking robots on our little dog robot, and with biped robots we've worked primarily in simulation and with a humanoid robot in Japan at ATR, we have a very similar robot. We haven't done too much walking, it's much more balance control, stepping control, doing basically full body manipulation while standing on two feet and balancing. Walking is, but the DARPA we're going to get, the Sarcos robot is very special, it's a torque controlled robot, which is relatively rare in humanoid robotics, and it will give us an ability to do compliant control in a walking robot which is not so much explored so far. So we are getting essentially into it but we also are approaching it maybe from a slightly different angle than other people have been doing before with more traditional ZMP controllers and other walking mechanisms.

**Peter Asaro:** Have you done work in human-robot interaction as well?

**Stefan Schaal:** We are interested in imitation learning essentially which is a form of human-robot interaction. We are not doing a lot of work in terms of studying how a robot should interact with humans in order to be appealing so, I think what we care about is to create humanlike motion which is our interest in biology kind of projected onto humanoid robots, and that's a component of human robot interaction but it's not a major focus on what we are doing, so Maja Mataric does that for instance and I think she has a real focus on that.

**Peter Asaro:** What have the best strategies been so far for imitational learning?

**Stefan Schaal:** I think there's many philosophies. We did some about eight years ago with collaborators who are now at the EPFL in Switzerland, Elke Eisberg primarily. We developed an approach to imitation learning with motor primitives, which is kind of a cute approach which uses differential equations or dynamical systems to kind of capture imitation, and then this approach is able to reuse it and generalize it to new motion, or to new scenarios. And I think this has left a little bit of an impact in the robotics community as many people try to do either our approach or they kind of used techniques which were inspired by this approach. So motor primitives and imitation learning has become I think something very prominent. Now there's a huge number of additional projects which I would not be able to enumerate without forgetting at least three or four so I wouldn't try to do that. But I think the idea of motor primitives has

become pretty central to imitation learning as one component, and increasingly more people do machine learning in this area which is also part of our approach.

**Peter Asaro:** What have been the most fruitful biological and neurological inspirations?

**Stefan Schaal:** I think from our point of view it's actually very related to what I just said, it's this, we created these motor primitives out of the idea the dynamical systems approach is very interesting to generate behavior where you can guarantee that it will end up in an attractive state, which is in the end the goal of the movement, and that you can also use them to couple perceptual information into the system, so for instance if you're walking and you see some other person walking you could synchronize with this person so that you both walk the same fast. Or if you're catching a ball you can synchronize with a perception to catch the ball at the right spot. These are, in the language of psychologists they're called perception action coupling phenomena, and they're very beautifully modeled with differential equations. And so all what we've been doing in this regard actually was inspired by the behavioral literature on this topic which is at least 20, 30 years old.

**Peter Asaro:** Do you use central pattern generators or?

**Stefan Schaal:** That is essentially the idea behind that essential pattern generator which is an idea which now goes at least a hundred years back. That was originally explored in invertebrates or lower vertebrates, and I think it was taken into more human-like behavior much later by the psychologists essentially who tried to model behavior in general like in humans or animals with pattern generator like approaches.

**Peter Asaro:** How does that play out in your research?

**Stefan Schaal:** It works very well.

**Peter Asaro:** Can you give an example of how you utilize those?

**Stefan Schaal:** Well we just basically, our motor primitives are basically a form how you represent periodic motion and how you represent kind of point to point reaching or discreet motions with the same representation, and these are all dynamical systems, they have either point attractor, they have limit cycle attractors, and with one representation you can literally capture pretty much all behaviors which you would like to generate so you, we don't have to think about representations at all anymore, it always works what we're doing.

**Peter Asaro:** And you chain them together?

**Stefan Schaal:** You can chain them together, you can just create one which is very big depending what you want so you can basically have, for instance a tennis forehand as one bit motor primitive, if you wanted to you could kind of partition it into first going backward and going forward. In our world this is not really necessary, we can do it all with one and we can actually now create really, relatively complicated behaviors out of just one behavioral representation.

**Peter Asaro:** What's the hardest task that you've successful accomplished with a robot?

**Stefan Schaal:** I think the little dog walking over, the little dog robot walking over rough terrain is most likely among the most impressive complete behaviors we've generated since it's really, it's rock climbing, it's going up staircases, it's going up, or jumping over, little hurdles and things like that, it's quite a complicated scenario, and it's basically all now running off the same kind of software and the same representations.

**Peter Asaro:** It's more difficult than devil sticking.

**Stefan Schaal:** It's much more difficult than devil sticking, actually devil sticking is surprisingly simple.

**Peter Asaro:** Which is why we do it for recreation.

**Stefan Schaal:** What's easy for robots is not always easy for humans and vice versa unfortunately, so devil sticking is, from a technological point of view it's what is called a regulator task, you just have to keep an object at a certain repetitive behavior, that means you just have to hit the one side of the hand in a certain position, then on the other side of the hand in a certain position. There's very simple kind of control algorithms which can do this really well.

**Peter Asaro:** And normal juggling is similar or is that more complicated?

**Stefan Schaal:** Normal juggling is actually similar except that. People have studied juggling for a while, so three ball juggling is pretty straightforward, the interesting part is the catching with your hands and the coordination with your hand and having actually compliant hands so, if you have stiff robot hands and the ball drops on it it just bounces off before you catch it. So having the right kind of hands helps tremendously, one can replace it actually, you can take, instead of

the hand take a funnel, and that will work just beautiful, take beanbags which have a soft impact, and you can make a robot juggle relatively easily. Three ball juggling is not so hard, it has been done, actually Chris Atkinson did this with one of the humanoid robots in Japan. Five ball juggling is harder, because of timing constraints the accuracy becomes harder, and if you want to go seven or nine you just need a very fast and accurate robot, and then you would be able to do it as well.

**Peter Asaro:** Do you juggle yourself?

**Stefan Schaal:** Three balls. Not very hard.

**Peter Asaro:** What would be the technology that would be the biggest breakthrough for the kind of work you do, that are available?

**Stefan Schaal:** Okay. Besides hardware where I think power and actuators are one big missing component and a huge amount of sensing on the robots is another missing component so, biological systems has sensing in abundance, robotic systems are extremely sparse. So there's hydro-components. I think machine learning, using machine learning, a massive amount of machine learning in robotics, and learning how to autonomously bootstrap robotics systems from almost birth so to speak, to fully functional behaviors would be really, really exciting, and it's still not very well done. Lots of manual programming still in robotics.

**Peter Asaro:** Were there any other roboticists or books or specific robots that really inspired you or influenced your work over the years?

**Stefan Schaal:** Well there's a ton of textbooks in robotics, again I will not dare enumerating them because I would always forget some. What is the most inspirational? I think in neuroscience the idea of mirror neurons most likely was kind of a cool finding in the mid-90s which got people again thinking about imitation, a very loose connection actually between imitation and the mirror neuron system. But that was kind of a, it's not a book it's just research papers. I think they were kind of inspiring. In general I think the progress of machine learning is really among the most fascinating to make progress in autonomous systems in general and try to understand how to put these things into robotic systems. Big difference between learning in robotic systems and what's normally done in machine learning is that most machine algorithms get a bag of data and then just analyze it, robotic systems generate the data themselves and have, live in an online, real time scenario, usually we can generate data in abundance, it's not an issue how much data you can generate it's much more how can you generate intelligent data where you learning something from. So there's a lot of work to be done to learn in perception action loops or just learning control in general, reinforcement learning is one special instantiation of that.

Yeah otherwise I don't know any particular component which is the most inspiring, it's lots of bits and pieces which come together.

**Peter Asaro:** Okay. That's most of it, let's see. Where have you gotten most of your funding?

**Stefan Schaal:** I think most of my fund, well let's go through history, I think initially most of my funding came from the Japanese Science and Technology Corporation together with ATR in Japan, and then afterwards primarily National Science Foundation. And then it's a little bit of NASA and recently a fair amount of DARPA funding which is, there are some projects which relate to machine learning in robotics like the little dog walking project or now the new project we have about grasping and manipulation. But I think the bulk of funding usually was National Science Foundation.

**Peter Asaro:** Do you remember how you first got interested in machine learning?

**Stefan Schaal:** I think I was actually excited about neural networks when they basically gained popularity in the late 80s. I was interested in self-organization and I was interested in learning, machine learning and especially neural networks were kind of the cool thing to do at this point. Self-organization, complex systems theory was I think another very cool thing to do at the late 80s, early 90s, and trying to understand how to bring these two things together that you can at some point think about machine learning algorithms as being completely self-organizing, which has actually been realized to some extent, not entirely in the current work of what people have been doing, it was kind of fascinating.

**Peter Asaro:** What was the first work in self-organizing systems that you looked at?

**Stefan Schaal:** Well I went to the Santa Fe Complex Systems summer school where you just get exposed to the staff example of complex systems like diffusion equations, self-organization in fireflies, I think Steven Strogatz at MIT was very popular at this time with his work in showing how just animals can self-organize into coordinated behavior without really directly communicating. Just analyzing dynamical systems, attractor dynamics, trying to use attractor systems to generate coordinated behavior like pattern generators do was one part of that.

**Peter Asaro:** When did you go to Santa Fe?

**Stefan Schaal:** 1992. A long time ago.

**Peter Asaro:** Okay. Was anybody else that was there when you were there that was talking about robotics?

**Stefan Schaal:** Who was giving lectures at that point? There was a lot of excitement at this point I think about the upcoming wave of synchrony in the brain which is Rolf Singer and people who worked on this topic, these are German people at Max-Planck Institute. That was I think one component. Low dimensional dynamical systems and analyzing them was very popular, just a whole bunch of lectures on that, nobody in particular, and beginning of machine learning and neural networks obviously becoming more and more popular. But I don't recall anybody who should be mentioned over the others, it was just a great series of lectures essentially and I don't recall everybody anymore at this point.

**Peter Asaro:** What have been some of the breakthroughs in vision systems or some of the big challenges you've faced in vision?

**Stefan Schaal:** So the thing is we don't do too much vision ourselves, this is just another discipline and we usually just try to use vision, we're very interested in active vision, ocular motor control, that means how to move eyes around in order to gather information, in order to get better perception and then basically do better movement afterwards. Yeah I'm not an expert in this area to really make particular claims. I like currently the development of basically increasingly more depth sensors from the very beginning, so range finders or cameras which actually give you a complete depth image of the environment that helps a lot to do, unlike what most vision people are doing, starting to do kind of point cloud reasoning, which is a newer topic, I think the mobile robotics people have been doing this for a while since they work with range finders for a longer time, and now I think in big robots it's becoming more popular as well.

**Peter Asaro:** Is there anything else you'd like to add?

**Stefan Schaal:** No, I'm happy. I assume you will edit this anyway.

**Peter Asaro:** Yeah. Okay.

**Stefan Schaal:** All right, all done? Well it was fun.

**Peter Asaro:** Yeah.