

Brussels Belgium
1:33:20
m4v

Luc Steels

An interview conducted by
Peter Asaro

July 8 2011

Peter Asaro: Start by introducing yourself, telling us where you were born and where you grew up.

Luc Steels: I'm Luc Steels and I was born in 1952, long time ago in a town near Antwerp in Belgium.

Peter Asaro: Where did you go to college?

Luc Steels: Well, I studied first in Antwerp. So I did undergraduate and graduate work there and my first degree was in linguistics in fact, but then I got completely enthusiastic about robots and about language, particularly through the work of Terry Winograd on <inaudible> and so then I got fellowship to, in fact, study computer science. So then I went to MIT and did, in fact, a new study in computer science. I worked in the AI lab where [Marvin] Minsky was my advisor and I retooled myself because I realized that you better know computer science in order to do something in this field.

Peter Asaro: What year was that that you went to MIT?

Luc Steels: This was in 1977 as far as I remember, yep.

Peter Asaro: And were there other people interested in robotics in Minsky's lab at that time or..

Luc Steels: Yeah. Well, so you see I come from the, you could say, the cognitive intelligence side rather than from electronics or motors, mechanical engineering and, of course, Minsky had always been interested in robotics, but there were lots of people there who were doing it and vision. You had the group of David Marr who was still there still active. There was also Horn was doing more work on robotics. There was Marc Raibert was doing his walking. So it was all intermixed, although I mostly worked on representation at that time and continuing with language. So when I was at MIT, I didn't do a lot of robotics per se.

Peter Asaro: What was your thesis on?

Luc Steels: It was about, it's in the spirit of The Society of Mind Theory that Minsky was working on at that time. So it was a model of distributed problem solving, you would say, or representation, which was applicable to planning or to language processing. It was general framework. Well, now we call it agent based, but this was long before the agent developments in AI, so, but there was basically the idea of agents cooperating to do cognitive tasks.

Peter Asaro: When did the idea of agents really come into the forefront of AI or become a concept of people using it?

Luc Steels: Well, I think there were several threats and so there's my personal threat, which was very much coming out of biology. For some people, it came more out of economics like economic agents, decision-making agents. So what happened was that after MIT I worked for a while in a company, Schlumberger who was at that time at the forefront of applying knowledge engineering to real-world problems particularly interpretation of oil measurements and geological measurements for oil exploration. And so then I came back to Belgium and I founded the AI lab here in 1983 and in the beginning we continued doing, well, you could say, symbolic AI, although there were already projects in vision and things like that, but the real switch for me came when I started cooperating with biologists, particularly David McFarland from Oxford University, who is an ethologist, and rather than looking at rational decision making, we started to discuss a lot about, well, animals and questions of autonomy, questions of behavior, real-time behavior with a physical body in an open unknown environment and so out of that – I'm now talking 1984 and 1985 the notion of an agent came about, an agent as being this autonomous entity that is not completely following instructions as programmed but has to survive or has to cope with circumstances and in particular doesn't know what is going to be confronted with because a lot of AI work before and still is ok, we have a large corpus here. We apply some machine learning technique and then we extract whatever it is, concepts or rules or behaviors or something like that. So the motivation for me came out of that and then I made contact with a new crowd of people, in fact, particularly Chris Langton, who then later started the artificial life movement. So this was more '88, he was here and also in particular Rodney Brooks. So with him, I organized a couple of workshops here in Europe to shape this way of thinking of autonomy, real-time behavior, minimizing representations and so we had, at that time, very strong connections because Brooks, who was at Stanford, had gone to MIT and with his team like Maja Matarić and so on, they all came here and we had intense relationships. A few Ph.D. students like Pattie Maes is one of them who then went to MIT to continue working on agents. So it was all kind of biologically inspired I would say and this was also the beginning of these simple LEGO vehicles. In fact in some sense, strongly going back to early cybernetics work, so in AI was always higher level planning and more and more complex representations and things like that and I think this is good. I'm not one that wants to throw out everything. I think that's a big mistake. But there was a balance too much in the direction of rationality and rational reasoning. So there was a correction going back to early cybernetics work. So one of the biggest projects we did here was inspired by this project of Walter, [William] Grey. He had these two robots, Elmer and Elsie, that had to recharge themselves in an ecosystem and so we set up an ecosystem here with small LEGO, these LEGO vehicles. At that time, LEGO Mindstorms did not exist yet and then we put pressures in this environment and introduced competition for energy and all these kinds of fairly biological experiments. So that was very exciting and the notion of agent came out of that and the goal of agents and also the idea of behavior based robotics, in fact, adaptivity, smoothness, more important than an accurate representation of the world. Raul Pfeiffer was also here in the lab as a visiting professor at that time and he

completely picked up on this embodied intelligence idea and pursued it with his students down in Zurich. So I think we were a real hop for this kind of stuff and the period between 1986 till the early 90s. Also, with the organization of, I think in my view, very important spring school in Italy that I organized with Rodney Brooks and so what then happened was that I made contact with the group Sony in Tokyo who then had this idea of the Aibo and so I started a long-term collaboration with that group, which is still going on, and so I went to Tokyo for a year and before the Aibo was actually a real product on the market.

Peter Asaro: What year was that?

Luc Steels: This was, I think, in the 90s already. Well, I think '94. And the Aibo actually has a lot of the ideas of agent and autonomy and adaptivity and development also. I still think that most people don't really know what's in the robot because it was never really published very well because it was in the object, the idea was to discover it in the object. So, but then I don't know when this happened, but let's say '94, '95, these ideas of behavior-based robotics were being picked up by more and more people and initially – well, we had this idea that this would be the root towards intelligence, an alternative route to intelligence, which is still my goal. I think it's fantastic to have these small robots doing all sorts of neat stuff, but I'm personally still interested in the question of language, for example, language interaction or intelligent behavior beyond what animals would do. And so then around '94 I made a big switch again to bring back language and so the idea is can we apply the same principles as developed in behavior-based robotics or some of the same principles, can we apply them to the problem of language interaction, and this already means for example that you're talking about adaptive language system because I think it is impossible to put it in the robot and then hope that it will be fluid interaction because language is an open system like the real world is open. So there was already this very important goal and also not only open but also evolution because language is not static. People invent new words all the time. There's pronunciations evolves, new grammatical constructions come in the language even if you are not aware of it but actually – so this biological perspective that we were exploring on autonomous robots or autonomous living systems I feel very strongly that it had to be applied also to language, but embodied language on real robots and so then I started a new research program, which is still going on which is to come up with the mechanisms for evolutionary adaptation and so on. And another thing is also that rather than implementing our language in the robots, I'm interested to see if robots are themselves capable to come up with a language and so that is adaptive to their needs, to their environments, to their sensory motor apparatus, etc. and so that's why it's about origins of language, evolution of language and evolution of language it's not just the syntax or the words, it's also the conceptualization of reality for language. And so how do we perceive the world like space, how do we come up with the concept front and back or left and right and because this is related to our body and our interaction with the world it's so fantastic to have robots that have a body and that can exploit their body for structuring the world around them. So at this point, I think we have many new technologies and many ideas, but they haven't become mainstream yet and I think partly because the world of roboticists is very different than the world of people

interested in language processing and/or concept formation. I think robotics people they realize, of course, that if you are ever going to have robots in real-world environments interacting with us, well, they better have some form of language. So everybody understands that, but I think problem of conceptualization, deciding what to say is as difficult, as complex as vision, perceptual vision. So vision is now more successful because a lot of people have worked on it and little aspects of the problem like line finding and so on and we have now enough computation power to run these algorithms, run them fast enough, so we can do fantastic things in vision today, but it's because of a lot of hard work. It's not one magical mechanism or something or a learning thing or whatever has been able to do the whole job. It's because lots of components now can be put together and the same thing for behavior, in fact, and so for language it's the same thing. We need a vast number of mechanisms that are put together like the way we conceptualize space is extraordinary and so just on that area that's lots of people have to work or enough people to deal with that and then I'm not talking about action or about time, for example, or just talking about objects in the world and their properties. So I think there's still – this is an area which is still very exploratory and very open.

Peter Asaro: What do you think are the biggest challenges for embodied language?

Luc Steels: Well, I think the biggest challenge – well, first of all, what we need are extremely flexible processing technologies. For example, for parsing, now you have had parsers already for 50 years, but most of them they are extremely brittle. They are designed to – to parse grammatical sentences – now, first of all, this grammar is elusive because we all speak differently and there is no de-grammar of a language, just doesn't exist, and it changes and most of the sentences we produce are not grammatical. So we have the problem of robustness, of how to cope with all this noise and incorrect input and etcetera, etcetera, which is a problem. In robotics before the people build a vision system that would work under very controlled conditions with light, like this, and in a nice clean environment with clear objects against white background or whatever, and so it's only by saying forget about that, we need to succeed. We need to confront the real world as it is and we need to find techniques for the real world. So the same is true for language. We cannot ignore the noise, the ungrammatical sentences. We need the flexibility to say, yeah, this is not how I'm saying it, but I still understand what you mean. So the parsing, production technologies, and this also requires a total change in attitudes from linguistics because the theoretical linguistics particularly like generative grammar, they focused on capturing grammatical language and judgments of grammaticality, which actually are totally irrelevant. We are not sitting here all the time thinking is this sentence grammatical. We want to understand each other. So there's that aspect of the problem and the other one is that meaning is actually the key to language. So language is we give hints to each other what we mean, but they are just hints. So it's not a programming language where I send you a program, you execute, and then you do what I would like you to do. It's rather I make suggestions to you or I give you hints and then you have to figure it out and the big possibility with robotics is that we have a body in the world. We have actions. We have goals. There is cooperation and meaning comes out of that and so the big problem with pure text-based language processing, even web processing, is

that there is no meaning. There is no relation to the real world that can be tapped to provide the meaning. So in that sense, I think it's absolutely crucial that robotics is used as a road towards meaning and then towards understanding and towards cooperative action, which relies on both partners actually being intelligent and figuring out what the other one wants even before he says it and then you just need a hint and then, yes, it's confirmation of what you already predict that the other one would do.

Peter Asaro: Which do you think has a bigger foundational role in language, the stability to ground the semantics or the stability to cooperative between agents or is it..

Luc Steels: Well, I think both of them are crucial. These are the two pieces and then the rest will follow. Then, we can have learning mechanisms. If you have very strong semantics, then you can build learning mechanisms already today, which will then do the job, I think. So what's missing is the grounding in the world of the conceptual apparatus and we don't need a few concepts. We're talking about hundreds of thousands of concepts. Now you see a paper and they're so happy when they've been able to do four concepts. Well, we need to scale up to vast amounts of concepts and it has to go quickly. So this is the problem, embodiment, the grounding, concept formation, and ground the concepts like colors, shapes, and the other one is the cooperative interaction. So the fact that you can guess the intentions of the other that you track the goals of the partnering in the dialogue, but we have lots of technologies and techniques. So I think a big problem is in fact the – because we're dealing with a huge system, how can we cope with the complexity of building it and of having many people cooperate to build it. I also think that the reason why we don't have yet any, in my view, no really impressive demos of language and robots is because in computational linguistics 90 percent of the people they work on statistical language processing. So this is all corporate, some statistical method and "boom" you get an ontology or you get some sort of relation and then you can do amazing things like Watson or search engines. I think it's all good. I think it's admirable all that work, but if you look at robots or people who are in a conversation, they're not understanding something in a particular way because in the past that was appropriate. No, but because it's here and now in this situation, grounded in this interaction that's the most appropriate thing. So we cannot – I think it would be dangerous to rely on this statistical methods for human robot interaction because it has to rely on real understanding otherwise the risk of misunderstanding is very high even though it would work 80 percent of the time, but then when it really matter there can be a failure, which cannot be allowed in human, robot interaction, particularly if you're serious about robots, I don't know, cooperating with people for healthcare or whatever in dangerous circumstances or something. We just cannot only rely on shallow language processing that doesn't really go into the deep structure of language and doesn't really try to understand and try to grasp the intentionality of the dialogue partner.

Peter Asaro: How important do you think it is to have some representation of the biological aspects of other agents in order to understand this cooperation of language and other intentions and things like that?

Luc Steels: Well, do you mean biological, how realistic from a biological point of view your system has to be?

Peter Asaro: Yeah or in terms of how do you model that line or its intentions or its goals?

Luc Steels: No, well, personally I believe that the goal of being very realistic with respect to biological processing is, how should I say, I don't want to say anything really bad, but I think it's not for us. It's not a good goal. You have some people, for example, this project that Henry Markram is doing in Switzerland called Blue Brain Project. So the idea of him but also of other people is that if we can reconstruct the brain structure at sufficient level of detail and simulate that, then we are done. Now, I think that is not a good approach, well partly because I firmly believe that the brain can do what it does because it's in a body and the real world is there and so on and others are there. So just rebuilding this brain and then hoping that now we have this brain and that's it our job is done sounds very wrong to me. Maybe we can learn a lot about brain science, but I don't think this is a root to intelligent behavior to artificial intelligence. So we're dealing with a hugely complex system, so I think we need the level of abstraction that we can handle that we can understand what's going on and then you have levels of implementation, which go down and down and down and down and then I think in one way you can go down to a neural biological system, so it's sort of compiler and in another way you could go down to, I don't know, silicon-based system or something. So I'm thinking here about the way we design the circuits. Nobody today is really doing circuits at the level of one block on a silicon and then here I put another block and then here I put another block. We design at a very high level of abstraction and then we have compilers who go down and translate this into then the, I don't know, how many millions or something lines and blocks and so on you have on your circuit. So I think we need to do the same thing for intelligence. We cannot go down to this level of the neurons and hope that we get any insight or any design possibility at that level. So this is another reason why I think hoping that just staying at this realistic neural modeling is not the right way. So, also, of course, I do think that in terms of just another part of your question perhaps, if you want to interact with humans then, well, you have to know something about other humans. So, of course, we use ourselves as a model of the other. Also in our speaking, I'm speaking to you hoping that your knowledge of English is roughly similar to mine and so you understand even though we never talked before. So we use ourselves as a model of the others. So in that sense it's very important – yeah, with a robot, it's always going to be different. The past seven years or so, I've been using these Sony QRIO humanoid robots. So I moved. We did a lot of things with the AIBO four-legged and we moved to a humanoid shape particularly because for embodiment it's closer to human embodiment. So if you start talking about an arm, while this robot has an arm, and raising the arm, well, he can raise the arm or it has a front and a back, etcetera. So in

that sense I think that being closer in embodiment is a way to have more relationship or to be able to, I wouldn't say experience similar things, but at least that more understanding is possible by the robot because he has a similar embodiment.

Peter Asaro: So do you think human robot interaction as a field has caught up with the idea that needs to really grasp this evolution of language?

Luc Steels: Well, not yet, I think. In my view, well, this may be partly my fault because I haven't been going around enough as a missionary. I'll give a talk at the RoboCup on Monday, so, to try and convince people of this point of view. No, I think this is still a wave that has to come and I think it's unavoidable. Of course, we need a lot of sophistication, a parser or grammar is complicated and there have been many people working on it, very good people, and we have to take their work into account. We cannot just say, forget about all of this complex stuff. I don't want to know about it. We have to confront the reality. So we have to convince both the robotics people to become part of that team that looks for the interaction between robots and language and for linguist or computation linguists who know a lot about the language processing technology. So they have to come together and that's very difficult because they're now two different communities.

Peter Asaro: And in terms of the foundations of computational linguists, do you see this work on the evolution of language as challenging the notion of universal grammar or perhaps discovering structure of the universe of grammar?

Luc Steels: Absolutely. In that sense, I should say that – well, I, at least, have been part of that debate. So I published quite a bit and being in conferences on evolution of language where my dialogue partners are linguists and anthropologists and psychologists and they are pretty fascinated by these robotic experiments, but they don't quite know what to make of it. For them, this is very strange and alien and they're not sure that to accept the argument because of the robot. Of course, if you are in robotics, like envision, you say, well, we built something for texture detection or for color. I think people working in computer vision have made major contributions to our understanding of human color vision, but in the case of language and they say, well, yeah, this is all for robots. They don't yet make this connection that this is giving them models for thinking about evolution and language and grounding. So this is still an uphill battle, I would say, with those people, but I find that fun to engage in, but at least I – well, one of the reason why it's an uphill battle is because we are telling those people that a lot of their beliefs are not valid or don't work you see. So usually we try. I think in AI and robotics, everybody is very pragmatic. If you come up with a solution, we try it. There's no you found this I don't want to try it. You see cooperation between teams and they found the better method for this. Okay, let's incorporate it. Let's try. And so here it's the same. We try something and then we find that it does not work and so then we say, well, this does not work for these reasons or this thing works.

And so, for example, this idea of universal grammar just does not work partly because nobody has been able to say what it is, this universal grammar, because if they would be, I think we would be most happy to use it, but it is not there. Language is an evolving system. There's huge differences between the way different languages conceptualize reality even for time and action and so on, for space. So you cannot just decide if universal comes mostly from people who know only one language and then they think that everybody speaks the same way or, of course, they use different words, but basically they view time in the same way, but it's not true. A language like Russian, for example, has a very intricate system of aspect, which is to talk about the internal structure of events and they categorize the world or time always in a way that involves paying attention to the structure of the event like is it repetitive or are we focusing on the beginning of the event, the ending of the event, all of these kinds of things are expressed in the grammar. Well, in English or Dutch or French, you hardly do that. So even languages you'd say Russian is maybe not that different from English but actually for this particular part it is profoundly different. So what then is this universal grammar is we forgotten about it or we're not using it or – I think it makes a lot more sense that every culture has developed their own way of conceptualization. There are many possibilities. They all work. Some cultures put more emphasis on this than on that depending on the tasks that they're confronted with, the history. So universal grammar doesn't seem to be a good idea and then also generative grammar is a bad idea. The grammar is about mapping from meaning to form, form to meaning. It's not about generating sentences that are grammatical. No, I give you meaning; you produce the form for it; or, I give you form; you reconstruct the meaning. We are not able to generate all the sentences of an English language. So there's lots of things with their perspective on language which would seem very odd when you come from the outside, from this more pragmatic view I would think and so in that sense there's a very big clash. Now, I don't know whether this is going to help linguistics get out of the little corner in which they painted themselves, but I certainly hope so.

Peter Asaro: Coming back to your work with the Sony AIBO, could you tell us more about that, what kinds of concepts that you were able to integrate and what it was like working with Sony?

Luc Steels: Well, the first thing I should say is that the team and Sony, which was directed by Fujita, Masahiro Fujita, and another person very important in AIBO development was Hideki Shimamoto and then the third one who is the father of the whole project and the champion is called Doi, Toshitada Doi. All these people and the team they have, they are really very brilliant people and very modest compared to what they have achieved and not known enough in my view. I think they deserve a clear spot in the history of the field and, of course, a lot of their technologies have since become integrated. When you have a camera and there's face recognition and then you take a picture and all of this, I mean, a lot of that is actually coming out of that project. So from the viewpoint of Sony, they did not continue with the entertainment robotics not because this was not successful, but I guess at some point they restructured their business and they only went for the big areas. They were doing lots of things and then had to make a selection rather and then robotics, entertainment robotics was there. I think most people

would say 150,000 robots is that not enough, and they were rather expensive, so by this it would be enough if that is the company's business, but if you do TVs, then you're talking about billions, anyway. So this is just about the team and about history and it's the same team that then built QRIO, the humanoid, which we have been working with this humanoid. It's just incredible. You take it off the shelf and you work with it for weeks without any breakdowns. It's just unbelievable in terms of the quality of engineering and the reliability and so on, which is very difficult if you have a university team building a robot or small company building a robot. There's so much know how involved in, not making a prototype – well, also it's already an enormous challenge to make a prototype, but then to go from a prototype to building 10,000 copies it's a skill in itself manufacturing, which I think is a big challenge for robotics in itself, but with the AIBO what – well, I did the experiments are also initially more from this perspective of the robot as a living entity that is autonomous and does not necessarily have a particular goal. It's just there like a dog you would say. If you visit somebody and he has a dog, you don't say, what is this dog for? Well, it's there. You didn't buy your dog because it would do the dishes or some other practical thing in the house. So this vision already, which is Doi's vision, I think is very new. So we kept on working on autonomy and on perception and concept acquisition with the robots, the AIBO, also spatial experiments and then interaction of language, but, of course, everybody said, this is funny. Dogs don't talk. Why do you do these experiments with language and then the AIBO. So this makes more sense with the QRIO to try that and then with the QRIO actually we have experiments with, for example, particularly for space. This particularly the work of Michael Spranger, one of my Ph.D. students who just finished his degree. I'll give you one of his thesis actually, which I think is absolutely brilliant work because he bridges the gap from – he has the two skills, the kind of new person that we need. So he knows about perception, knows about action selection and implementation of behavior on the robot completely, but on the other side, he knows about language processing and parsers and grammars and so he's able to bridge the gap completely between the two sides of the equation and so the spatial experiments that he has been doing is really about building up a group of robots, building up concepts like left and right, front and back, but also perspective because then spatial language if I say left could be from my perspective or yours or from the perspective of a third object. So this already requires that a robot is able to take the point of view of the other and then to express that point of view, not because we say you have to express it and here's the grammar but to discover that if you express the perspective from which you describe the spatial relation, then you have more reliable communication, more successful communication. So this is the driving force in our experiments. On evolution, it's not fitness or it's not biological selection, not genetically – so it's not this idea of language genes at all. It's language as a culturally evolving system and so the ideas that the robots adapt to each other's language all the time. After every interaction, they adapt their concepts, their preference for certain rules. They pick up words from the other. They pick up constructions that the other one is using. So there's this constant alignment going on and there's invention by the robots. If it doesn't have a word for something or no grammar or no concept yet for being able to say what he wants, there are mechanisms for, well, invention, you could say, of these parts, which are then learned by the other and then they do alignment. So he does this for spatial reasoning, spatial domain. We've also done experiments in action and here again the problem has always to do with the other one, which

robotics – some people look at collective robotics, but most people look at one robot interacting with the world. So we always look at-at least two robots interacting with the world and with each other. And so for action, if I ask you to do something and then you do something, I need to be able to perceive that what you are doing is what I asked you to do. It sounds obvious, right? But the question is how do you see that and how do you learn that in particular and also if you want me to do something, you have to be able to plan your action. Well, not your action, but you have to be able to choose actions so that the other one would then do what you want. So this means you have to be able to conceive actions of the other to achieve your goal and so we did lots of experiments actually with mirrors where the robot stands before the mirror and then learns about its own body and then once – as a way to then be able to interpret the actions of another robot's body and also to be able to learn that – by looking at himself, if this control, this sensory motor behavior is giving this visual image, so if I want to see this kind of behavior happening, then I have to talk about these commands that I want the body to do. So I think that's fantastic work also on this coordination. This all shows by the way that we cannot split the problem. We cannot say, we do robotics, then we add language. Language is actually helping the robots to learn about actions, to learn about actions of others, and what to ask for and how to say that and it's all very much interaction based. So this may be one of the big themes of the research that we do is what can you draw from interaction, not only about interaction itself but also about the world and about bodies and behavior of others.

Peter Asaro: So you mentioned Michael [Spranger], a recent Ph.D. student, who have been some of your other Ph.D. students and where have they gone to work?

Luc Steels: Well, so I had, I don't know, I think more than 30 Ph.D. students who graduated. Some of them MIT, Berkeley. Some of them in Japan. They've been everywhere. Well for robotics, another one is Tony Belpaeme, who is now in Plymouth. They a group, fairly big group, now on using the iCub for their projects. So he's more into visual perception and he did his thesis on color, so how a community of robots could evolve both the concepts for color and the words to go with color. I don't know whether it's useful to start picking out specific people. They're a bit everywhere.

Peter Asaro: You mentioned Raul Pfeiffer was here, were there other collaborations and what was your collaboration with him like?

Luc Steels: Well, the collaboration was mostly at conceptual level. We talked a lot and kept talking decades after that to – because what I usually – there's no big technology lab here as you see. So [audio cuts out] robots at the moment. We used to do more. Well, another one who was here actually is Andreas Birk who has now a robotics lab in Bremen, University of Bremen. There's another one maybe I should mention is Bart DeBouge. He's now in Amsterdam and he worked more on speech. So we had for a while quite intense activity in the domain of origins of

speech and the same idea actually, not putting into the robots the actual sounds but to see how they could come up with a speech system among themselves. Another student, now they are beginning to come. My memory clicked here. It was worthwhile to check out if you want is Pierre-Yves Oudeyer. I can write these names for you if you want later. He now has a fairly big group, very active group, in Bordeaux in France and they are also building small humanoids. When he was doing his thesis, he worked on phonetics also, so origins of speech, which is in a sense also robotic problem because you need to learn how to control the articulatory system, which is a very complex machine, in fact, with lips and tongue positions have to move and then you have to relate that to the acoustic signal. So it's a bit like with an action that you have to be able to recognize the action that you want to see done or you have to know that if I do this and he's doing the same this will roughly look like this. So with speech it's the same when I have to say an "ah." If you have to learn the speech, you have to be able to learn to move this to get the same sound. I think it's a very interesting domain also anyway. So Pierre Rieve, he particularly worked also on motivation for learning. In development robotics, you need a motivation to push the learner and his idea was particularly to focus on predictability so that the learning motivation was in terms of, not just to reach predictability, but if you increase in predictability that was the motivation for learning and so once you knew it, then you know it. I'm not saying you finish your learning, but you focus now on something else, which you cannot predict and so you try to learn about that and so this is how this robot autonomously explored the environment. He's still doing lots of experiments in that domain motivation. Another very good student I had was Frederic Kaplan. He's now at the EPFL in Lausanne, which does lots of robotics like with Dario Floreano, and all these people. He worked also on language with me and on the AIBO actually, a lot with the AIBO, and he did quite intriguing experiments with interaction between animals and robots, working with people in Hungary who are specialized in dogs, real dogs, and doing systematic experiments to see how dogs, when they are for the first time confronted with a robot, how they treat the robot and there's some very interesting behaviors that came out there. But, anyway, he has been looking now more at possible applications of robotics but not as a robot but in other forms. So, well, these are some of the students. There are others, but I just – these are the ones that just come to mind.

Peter Asaro: To go back a little bit, when you opened the AI lab here in '83 in Belgium, was that the first AI lab in Belgium and what was the state of AI across Europe and from your perspective at that time?

Luc Steels: Well, the thing is that Europe has a very strong tradition in cybernetics and also biological cybernetics, so connection with animals, but it does not have a strong tradition in AI. Now, it has a tradition in robotics, but the mechanical engineering side of it. So if there is any robotics, it is in mechanical engineering, electrical engineering departments, which consequently means you certainly don't get this idea of the robot as an autonomous living entity. It's more viewed as an engineering problem. You see that also in the shape of the robots that are built or were built and the interfaces with the robot and everything. So when I arrived here, it's a bit different in the U.K. because the U.K. always had much more AI. They are closer to United

States. So you have big difference between the continent, continental Europe. So when I arrived here, there was very little AI in Europe in general. Most countries like Spain or something, there was nothing. Progressively, things started to pop up, but in general, I would say AI in terms of as a research field in its own right, is not sufficiently powerful. There's some AI labs now in Europe. In Germany, you have DFKI, but they are very application oriented. So that's a bit of a problem. If there is any AI going on, it's either very strongly connected to pattern recognition, signal processing, so more the engineering side of it. If there's robotics, it's more the engineering aspect of it and if it's more AI, I would say it's very pushed towards practical application. You also see this in the funding. There are big programs. They call it cognitive systems, but actually it's a – a lot of it is psychology, cognitive psychology. So in my opinion, we still lack really good AI labs, where the creative work that – what I at least want to do and that we have been talking about – would be possible. Now, there are lot – of the last five to ten years, there's been a lot of activity to build new robots like the iCub project in Genoa. Yeah, so there's certainly more robotics. But in terms of what I would say the – what we need to work on the intelligence part – there's definitely not enough of it. Not enough creative work. Because people are – I don't blame anybody; it's because the money you can get for this research is immediately tied to applications. And then the projects with a strong psychology let's say, flavor, cognitive psychology, partly they want to do, use their own networks which is not going to lead you very far. It's good maybe for some aspects of pattern recognition, but if you're talking about language, or planning or something like that... So, like Pfeiffer has a big group in Zurich which is very successful, or Floriano also, they're building a new robotics institute. But they do not work on any of these topics. So either its robotics – you know, sensors and simple things and all that – very good, but no connection actually to intelligence. There's no language, no planning, no modeling of the other to any level of depth. So the gap is not being bridged in those places.

Peter Asaro: Some Japanese work on sort of conversational turn taking kind of stuff, but they're not really looking at the ground in your eyes.

Luc Steels: No, I think in Japan the situation is similar as in Europe; namely, AI was never very strong in Japan. The AI that we have been talking about here, right? So you have lots of robotics work; very sophisticated, fantastic actually, like Kuniyoshi at the University of Tokyo. What he's doing with robots as mechanical devices is incredible. But they certainly don't even consider the problem of interaction or cognition. So there it's similar, the gap is similar; and very little AI work. So the cognitive psychology people, they talk about it, but the technology is extraordinarily complicated so they lack the background of the computer scientist to manage the complexity, to really tackle these problems. So as a consequence, you get kind of small demos or it stays in the level of, the level ideas without any realization. So I think we really – I hope in the next decade we get a clustering of people who can – well, who can bridge that gap. And I don't know, at the moment, I haven't been in the US for a while. It's more difficult for me to judge what's going on.

Peter Asaro: What was it like working with Minsky as an advisor?

Luc Steels: Well, he's a very brilliant person. He's also a very – how would I say – he has his own mind. So if – that means he speaks his mind and he does what he thinks needs to be done and so on; which is a sign – well, which is behavior that many people who make breakthroughs have. But as a consequence, lots of people then feel that they have to attack this kind of person. Now, I have only good memories of it. So what I remember is a very kind and very intelligent and very supportive person who left a lot of freedom to his students. You could say you were on your own; that's bad, but I think that was good actually in my case. So he leaves complete freedom and takes the idea seriously of students and pushes them to be creative and to find new ways and so on. I mean this is – I think all very good as an educator. Now of course, I think that this is – I see this kind of attitude also with other AI leaders I should say. Now the side effect of it is that there's no – what we do not have in this – in AI is, for example, the history of ideas. If you look at biology, there are plenty of excellent books which don't go into the great detail of the molecular biology or the genetics; but they talk about the fundamental concepts of the fields. And sometimes I get a question like, "What can I read to learn about the basic ideas of AI?" Well, there are text books, but they are text books, right? There is not – whereas in biology I can immediately give you at least ten books that I think – written by leading biologists, that's the point. So that's a problem of the field, that we don't have this tradition, particularly from leaders in the field, to write books that are not popularization, but that capture the fundamental ideas of the field, and so this is a bit of a problem. And so somebody like Minsky – he wrote, of course, lots of stuff. But the others – I was thinking about him or McCarthy or something – none of them wrote actually such a kind of book. Maybe this is asking too much, but at least in biology, probably going back to Darwin, you see the *Origins of Species* and they kept this kind of tradition.

Peter Asaro: Howard Gardner has a good history of cognitive science, but AI is just a sort of chapter in there.

Luc Steels: Yes. Yeah but my feeling is that AI is not officially understood by other fields, and I interact a lot with biologists and physicists and all these people; and as a consequence they don't understand enough; or philosophers. They don't understand the key ideas of AI, and so they have all sorts of conceptions of it. Or they know about it through other work which is actually not a good reflection, like philosophers who have written about AI, usually against AI. But not knowing really what AI was about, or having misconceptions about – a good example is the notion of representation. In philosophy there's this one school of thought that said there are no representations. Or you don't need representations for intelligence, something like that. Okay, and then Brooks writes this paper, *Intelligence without representation*. And they all say, "You see? You see?" But actually when you read Brooks or when you look at his robots, what's in his robots, the robots, they do use representations. But they don't use the representations that let's say Shakey was using. So he was not trying to get to a logical, full, accurate representation

of the world. And so his point was more like we need to be – use representations if we can, but we don't need the accurate representations that some people we thought we needed. So this was his point, you see. But if I then talk to philosophers, then they get into all sorts of weird things, discussions about representations, by not understanding what we actually mean by representation. I think that there's a lot of work to be done by people in the field to clarify to the outside world what the key concepts are of AI. And maybe in an historical way like you are doing, I think that's very important because we have very little history. There's no book on the – as far as I know, real history, right? Which is difficult, and because we don't have it, every year it becomes more difficult, because there's more history so you have more work to understand the history.

Peter Asaro: Yeah. So, you did some of these early robots with the Lego brick, but that was the pre-Mindstorms sort of MIT brick, so were you engaged with Papert or with the other people that developed that? At what stage did you start becoming involved in that research?

Luc Steels: Yeah, we – there was a brick, but I think this was even before Lego was involved at all. I know – well I knew Papert because he was also at MIT when I was a student, and some of the collaborators of Papert. And like Maja Mataric worked in the Lego lab and other people. So some of the – we actually, this must have been quite early – built our first robots here with the MIT brick circuit, and we built our own. But this kind of programmable – because that's the point, it had to be – the key was in the programming, and particularly also there was a Lisp environment that Brooks had built to do this compilation from high level behavior language to actual behavior on this brick. This is an example of an innovation, if you want. Well, it's innovation – a technique that we use all the time, this compilation process. And I think this is why Brooks was able to build the robots that he could build. Of course, there's also ideas about representation, but there's also this idea of compilation as a way of mastering complexity which people in cybernetics did not have and still don't have. They still don't understand this idea. So they're still tinkering with parameters and building dynamical systems, and then another one and then see if you put them together it doesn't work, and so they never can push up the level of complexity of their robots sufficiently if you don't use these techniques from computation. So I think for me it was not just the brick, but also this methodology of programming in high level languages like the behavioral language that – with these augmented finites that – machines and so on – was very important.

Peter Asaro: Okay. So for young people who are interested in a career are pursuing robotics, what would your advice to them?

Luc Steels: Well I would say that they have to – I would definitely encourage them to do computer science. Not only electronics or if they come from that era, mechanical engineering. But to do computer science doesn't mean then that you learn how to program C. It means to

learn about abstraction, computational abstraction, and management of complexity. And I think that there are courses like that, okay, but it's that kind of computer science that you need. So this thing. And then a second thing is also to try to – you don't have to do this all at the same time, but at somewhere in your study, okay – suppose you start with signal processing, pattern recognition, some notions of electronics and motors, you know. Okay, you do all that for a few years; so you're not afraid, you have the grounding, the feeling. Because I think it's like working with clay. You have to do it. But then I think it's very important to also learn about, for example, to learn how to program in a real high level language like Lisp. Not because you want to, particularly of its use in industry; that's irrelevant. But to learn about abstraction and about tools for conceptualization, for language, for planning, problem solving, user modeling, all these kinds of things. So, and I don't think there's a study program where you can have this all neatly organized for you; so it's like going to the best people you can first to learn about the mechanics of real robots; maybe not to work on motors for the rest of your life, but to get a familiarity. It's like a musical instrument. You play it, or you learn how to swim – you jump in the water, right? You don't read books about it. So it's the same thing: you build. So definitely be in a place where you are allowed to build or you can build. And then there has to be this dual education I think. Some people like Michael Spranger did, I was talking about. There are some others. We have also a strong connection at the moment with Humboldt University, where there's a group directed by Manfred Hild. Well Manfred himself is very much of the first sort; a thinker, he builds the robots and does low level behavior control and all of this. But at least some people in his group also have this – have had the opportunity to look at planning let's say, and... so but I think that for young researchers I would say the field of robotics is still – it's not that it's wide open, but many big discoveries remain to be made, fundamental discoveries. But they are probably about, well like global architecture and also discoveries for small parts like object recognition or something, or there's little advances in all these areas. So there you can contribute. But there's also fundamental concepts I think which we are lacking, and concept – particularly about integration of let's say sensory motor stuff and concept formation; so integration things – architectural things. And now we can tackle these better because we have a lot more technologies already there. You see when we started here with robotics, there was nothing, right? You want a robot? Well you build your robot. You couldn't order a robot and the next week it would arrive through the mail. It was like building, and the students, early students who – they suffered a lot because they were doing these experiments like with the charging station and then “<noise>” another explosion, there was another robot going down the drain. So fortunately, thanks to enormous advances in mechanical engineering, now we have much more solid robots; it's easier to build them. We have smaller computers, less expensive. So the field is there, the basis is there to work on the real problems. So I think, the coming decade I think will be very exciting in robotics.

Peter Asaro: What do you think the biggest applications are going to be in the next decade?

Luc Steels: Well first of all I will already say, let's not worry about application, because we worry too much about application. We need more fundamental research. We need to give

freedom to people to think in an unconstrained way, the way it used to be in the I actually in the early 70s when the project at MIT with Minsky and Papert and then with Winston and Sussman and Horn and Marr and all these people – their concern was not about application. Their concern was to advance the field by saying, “What can we build now with what we have that would be beyond everybody’s expectation?” And that’s where the great things have come out of. But of course in terms of applications, I think there is actually, one is in health care. I’ve seen a robot project – it’s not really a robot, but a wheelchair – this was in Barcelona by a Ph.D. student of Ulises Cortes. The name of the student escapes me unfortunately. And this was about a wheelchair that was smart enough to help a person to navigate in his personal environment. Now there are many projects like this at the moment. But what impressed me is that they were doing – well, experiments you could say – with real people, and it was not in the lab. It was in the hospital, in the rooms where these people were rehabilitated for, after some – I don’t know, stroke or something. So there’s also in prosthetics, in helping people in all sorts of things, and our society is aging so I think it’s not to take away work from people because that’s very stupid. For robotics I think we should never try that even. We should try to think of applications where – which are difficult for people or where there’s not enough people or, which are dangerous; those kind of things. So I never think about robots as replacing people, but as – now there’s also the entertainment – the original vision that the Sony people had I think is still good, that entertainment could be – like computers – could be the biggest market for robotics in the future. But it has to be – it cannot be very expensive; reliable, all sorts of things. So it’s like toys for playing. Now of course, you can list all sorts of other things like space exploration, typical. But I think if we talk about robots in our environments, then to me, so in health care, in entertainment, robotics would be – or dangerous things like this nuclear accident in Fukushima – everybody says, “So where are these robots now?” Right? And then they used some robots, but you would have expected that there would have been more, and that they would have robots that would be able to operate in circumstances that – it’s possible to build such technologies; and so the world is getting more dangerous so we better think about it now.

Peter Asaro: Great. Is there anything else you’d like to add that we missed?

Luc Steels: Are you going to IJCAI?

Peter Asaro: To where?

Luc Steels: The AI conference IJCAI in Barcelona?

Peter Asaro: No. I’m heading actually to Switzerland next, so I’m talking to Rolf next week.

Luc Steels: Okay, good.

Peter Asaro: So the people in Zurich and Lausanne as well.

Luc Steels: Yeah. One – I'm also interested in art.

Peter Asaro: Ah, yeah. Would you like to talk about art robots...?

Luc Steels: Well, but the thing I was going to say is that I'm also interested in music, because I wrote an opera about a humanoid robot which will be performed in Barcelona on the 18th of July at the opening of IJCAI. So this is sort of the popular imagination of robots, so that's also interesting. Because robots as a vehicle for raising issues that concern people, and well, you can do that through art like a play, or we've had a few, or films. But also in philosophy you can raise a lot of issues I think, philosophical issues, through robotics. So you say, okay, what would it mean – well, this – like qualia or something, or consciousness or self or autonomy. It's very illuminating, I think, then to say, suppose we have a robot; now what would it take for this robot to have X, where X is one of these things we don't understand? And this is maybe a way in which first of all our research can become more relevant to humans, not in terms of application, but in terms of ideas, in terms of thinking about ourselves, and this is actually what interests me most in robotics, is the ability to explore issues of mind, like reflection. What is reflection? Or even, what is learning? And so we can translate that to operational models and then talk again with these philosophers and say well what about this? And I think it illuminates a lot of philosophical discussion that we do this. And so, and if you use robots, well there are robots being used in theater performances and things like that, but at least in this opera that I wrote together, my – the libretto is done by Oscar Villarroya, who is a Spanish, or I should say Catalan neuroscientist – you know, we kind of raise some of these issues of what happens when a robot is bought by somebody and this robot turns out to be an autonomous developmental robot, and of course, the person doesn't like that very much because he has to teach this robot everything and so on, and then – well things then – well I'm not going to disclose everything, but we had great fun in imagining. But I think that's – you see, there are plays about robots, but usually they are written by people who don't know anything about robots, which is another – then they come up with all this science fiction where robots are killing people, and – which I think is total nonsense, right? So, or maybe we can counteract a bit and have robots that have a more softer image. In this opera actually, the robot becomes more interested in music than in slavishly following what his owner wants him to do.

Peter Asaro: Well it was a play after all, it was the word robot.

Luc Steels: Yeah yeah, yeah yeah, exactly. And you have this whole Frankenstein thing, and you have films, Blade Runner and so on. So this is something I feel we from the field have to engage with, and have to – well, Minsky did that with Hal and so on, he was involved in Space Odyssey. But I think that's a fun part to get engaged in these things.

Peter Asaro: I taught a class last semester on robots as media...

Luc Steels: Yeah.

Peter Asaro: Half of it was on film and half was related to robotics issues...

Luc Steels: All right, yeah, interesting. Yeah, robots in themselves are indeed a medium that you can use. But that's not what I do in this opera; the robot is played by a human – well, we have an Aibo life, you know, he has a dog and he walks with the dog. But it's more about what – well a developmental robot, what does that mean? Autonomous... yeah. Well...

Peter Asaro: Great. Thank you.