Brain Science Podcast

Episode 25

Interview with Rolf Pfeifer, PhD, on Embodied Intelligence

Transcribed by Lori Wolfson
All errors or omissions responsibility of the transcriber

(Music)

INTRODUCTION

*This is the **Brain Science Podcast** - the podcast for everyone who has a brain - and I'm your host, Dr. Ginger Campbell. On the **Brain Science Podcast** we explore how recent discoveries in neuroscience are unraveling the mysteries of how our brains make us who we are. For more information including Show Notes, links to previous episodes, and information about how to subscribe, please go to the website brainsciencepodcast.com. We also have a Discussion Forum at brainscienceforum.com and you can send me email at docartemis@gmail.com.*

(Music)

Welcome back to the ***Brain Science Podcast***. This is Episode 25. Today I have an interview with Rolf Pfeifer, the author along with Josh Bongard of the book, *How the Body Shapes the Way We Think: A New View of Intelligence*, which was published in 2007 by MIT Press. Dr. Pfeifer is the Director of the Artificial Intelligence Laboratory at the University of Zurich in Switzerland. He also wrote *Understanding Intelligence*, which was published in 1999.

The main topic of our interview is the idea of embodied intelligence, and in particular the new field of embodied artificial intelligence. We are going to discuss why artificial intelligence is relevant to those of us who want to understand how our brains work, and we will review some of the key ways in which computers and other man-made models differ from our brains.

If you are new to the ***Brain Science Podcast*** I want to welcome you. I would also like to invite you to visit the website brainsciencepodcast.com. In the Show Notes you will find links to previous episodes that relate to today's interview. Of course, many new listeners just go back and listen to all the episodes in order, but that list is getting kind of long, so this will give you some shortcuts.

After this interview I'm going to have a few announcements, but let's get on in to the interview.

(Music)

INTERVIEW

GC: Rolf, I want to thank you for coming on the ***Brain Science Podcast*** today, and I was wondering if you would start out by just telling my listeners a little bit about yourself.

RP: OK. I'm directing the Artificial Intelligence Laboratory here at the University of Zurich. And I think it's important to know that artificial intelligence has significantly changed. There have been important paradigm changes since it started about half a century ago. Originally the idea was that intelligence can be viewed as computation; so the idea of intelligence as an algorithm, as abstract single manipulation.

Now, this approach has led to many useful applications. For example, if you use one of your favorite search engines, a data mining system, or consumer electronics you will find algorithms in there that have their origin in this classical perspective. But it has basically become indistinguishable from computer science in general.

This approach has also led to a number of failures. For example, with ideas on perception, manipulation, locomotion, and movement in general, this approach—which always has a centralized perspective; you have a central controller and you're trying to control the robot or your system—and these applications have not been successful at all. We still don't have very good walking robots, we still don't have very good systems that can, for example, recognize a face in a crowd, or things like that.

Now, in the 1980's in the field of artificial intelligence the concept of embodiment has been introduced by Rodney Brooks of MIT. The concept had been around for quite some time before, but I think in artificial intelligence it was Rodney Brooks in the 1980's. And there the idea was that it doesn't make sense to consider or to view intelligence as abstract single manipulation, but that you always need an embodied system that can physically interact with the real world.

GC: That really challenges the viewpoint that the artificial intelligence people have had way back to the time of Alan Turing. Right?

RP: Absolutely. Yes.

GC: Like I mentioned to you in our correspondence, I don't know how much my listeners know about artificial intelligence but I thought you might want to tell them a little bit about why Alan Turing is important, since they're likely to see that name.

RP: Right. He is basically the one who had this vision of computation. You could say he is in some sense the inventor of computation and of the modern digital computer. And he had a research program, in a sense, for developing intelligent computers. He predicted that within 50 years we would have intelligent computers—of course, a predication like many others in artificial intelligence that has not materialized. I think artificial intelligence is a bit a field with many false predictions.

Interestingly enough, Turing also was working on dynamical systems. That idea of dynamical systems is actually reemerging now in artificial intelligence concepts of self-organization, of morphology, of emergence, of synergetic interactions, and ideas of that sort. So, he basically has two faces, so to speak: a computation face and a dynamical systems face.

GC: Now, dynamical systems are going to be non-linear?

RP: That's right. When we talk about dynamical systems we usually mean non-linear dynamical systems. Or, we also use the term complex system, or chaotic system.

GC: And that means that we can't just break the thing down into its parts and add them back together and get the same thing.

RP: That's exactly the point. I mean as Strogatz was saying, you take your two favorite songs, and if you play them at the same time you don't get double pleasure.

GC: So, especially if we consider the limits of artificial intelligence, why is it that you think that artificial intelligence is relevant to trying to figure out intelligence?

RP: I think we have to come back to the distinction that I made at the beginning between this kind of classical computational view and the more modern embodied view. I think the classical view does not contribute significantly to our understanding of natural forms of intelligence, whereas the embodied approach does.

And I think there it's really important to see and to understand that we never have a computation or some processing going on somewhere in isolation, but if you look at the brain, the brain is always embedded into a physical organism which then also has sensors, actuators, that mediate the connection of the brain to the real world. I think this is important to understand, and also in the embodied view of artificial intelligence this is one of the essentials. I mean it's obvious—it's absolutely obvious and almost trivial when we say it like that—but then it's still surprising how many people don't take it into account.

GC: Absolutely. So, would you say that would be one of the most important ways in which a brain is different from, say, a computer or a neural network?

RP: Right. I mean that's of course an important question. I have to say that I'm not a neuroscientist. My background is basically in physics and mathematics, and then I work most of the time as an engineer. But we do cooperate closely with neuroscientists.

So, I think one essential difference is, as you just mentioned, that the brain always is part of a complete organism. You can never understand it if you look at it in isolation. I think that's a fundamental difference to a computer, because a computer doesn't have very interesting interactions with the environment. It has very impoverished interfaces.

But there are other really important differences between computers and brains. I think one is that it's really massively parallel. Also its enormous plasticity, I think, is very different, at least from today's computers. And then it's also a complex non-linear dynamical system which has largely analog characteristics. Or, it's largely an analog continuous dynamical system with, of course, digital characteristics—you know the individual spikes can be understood as digital, or abstracted as digital signals.

And there is also a lot of chemistry. You know? It's not just about signals, but it's a lot about chemistry. I mean basically the brain is in some sense a pharmaceutical industry. And this is quite in contrast to traditional computers.

GC: You've probably—and I apologize for this pun—you've probably embedded your answer into a previous answer to this question. But I want to really bring this out, so I'm going to ask it explicitly. Why is it that embodiment is so important to understanding intelligence?

RP: I think there are a number of points. One is I think if we really want to understand intelligence we have to understand it at least at three time scales. One is the here and now time scale, where we talk about the actual mechanisms. For example, you see a red traffic light and you push the brakes. So, that's basically what's happening now. Then you can ask how this ability came about in the individual. That's the ontogenetic time scale.

And then we have the phylogenetic time scale—how human beings have come about, and how this ability to learn about these things has come about. And an evolutionary perspective is always an embodied perspective. I mean these organisms' brains have evolved over evolutionary time, always as part of complete organisms that had to interact with the environment. I think this is one point.

Now, to be a bit more specific, I think embodiment is about the interaction between physical processes and information processing in the brain. For example, if I grasp this bottle of water which is standing right next to me on the desk here, then through this physical interaction—grasping is a physical process—through this physical process I'm actively generating sensory stimulation in my hand and on my fingertips.

And also when I grasp it, typically I bring the bottle into the range of the visual fields, I have it in front of me, I can look at it. I can also drink from it. So, what I'm doing through this interaction is I am generating sensory stimulation in various sensory channels. Oh, I forgot the proprioceptive one; I can feel about how heavy the bottle is.

And I think that's the core, or one of the really essential ideas of embodiment. These patterns of sensory stimulation that I'm inducing contain what we call information structure; that is, they contain correlations. And these correlations are absolutely essential for categorization and for learning. I think this is actually at the foundation of learning.

And it turns out that if this interaction is in addition sensorimotor coordinated—that is, there is a sensorimotor coupling, meaning that the sensory stimulation is influenced by the motor action and the motor action in turn induces sensory stimulation—if the interaction with the environment is sensorimotor coordinated, then the information structure that is induced is particularly high. So, sensorimotor coordination not only enables us to manipulate the environment, but also, from an information theoretic point of view, it is there to provide the brain with, let's say, structured sensory stimulation, or good raw material, so to speak. Right?

GC: Yes. I like your example of picking up something in your hand, because whenever they try to make a robot pick something up it's so challenging. It's so easy for robots to break things because not only do they have no sensory feedback about the weight, but also just the fact that our hands are soft so that when we pick something up we kind of deform to the object.

RP: Exactly. So, now you're touching a really essential point, also—and I think that the importance of this can hardly be overestimated—you said our hands are soft. So, it's not only morphological properties, it's morphological and material properties that are really essential.

Let's look at the morphology of the hand for a bit. Now, I'm going to stand up, because I have to feel the embodied interaction with the environment. So, if you stand and you have a loosely swinging arm—you know you just let your arm hang down and it's kind of loosely swinging back and forth—now the neural control for this is actually relatively simple, whereas if you look at the trajectory of the hand it's a really complex trajectory in 3D space.

Now, you could say, well it's nice that this is easy to control, but it is also useful. Yes, it's very useful because of the anatomical properties—or in our jargon its morphological properties—that is, because of the particular anatomy of the arm, shoulder, hand system, and because of the distribution of the sensors on the hand. We have a very high density of touch and temperature sensors on the surface of the hand and on the fingertips. We have much fewer sensors on the back of the hand. But that's exactly what we want.

Now, imagine that the arm is coming into the range of the visual field, touches some objects, then it's very easy and natural to grasp. It's much harder to bend the fingers backward. But hands have evolved over evolutionary time for grasping. And then you are grasping a bottle, and because you have the high-density sensors on the hand and the fingertips you are generating rich material—sensory stimulation—for the brain.

So, basically this embodied interaction, if you will, is a very good exploration strategy. Just bending the arm behind you upwards is not a very interesting exploration strategy. So, it's kind of like this interplay of morphological properties and material properties that give you a particular dynamics, and provide biomechanical constraints for a particular movement, and these are very good exploratory movements.

So, I think during the ontogenetic development of the child this particular exploration strategy leads to extremely useful sensory stimulation in multiple sensory channels that can then be exploited for categorization and for learning. Now, you can imagine that because you have the correlations in the different sensory channels, then it's very easy to learn these multi-modal associations. You can use just variants of simple Hebbian learning, because the correlations are there and Hebbian learning can pick up the correlations.

GC: Would you tell my listeners what Hebbian learning is?

RP: What I learned from you neuroscientists is there is this slogan, 'Fire together, wire together.' And I think basically Hebbian learning is about that. So, if you have within a particular time window two neurons active, then the strength —I mean this is very, very roughly speaking—then the strength of the connection between the two is reinforced.

GC: Thank you. I have talked about that in the past, but I try to make sure that my episodes can have somewhat of a stand-alone quality to them. I think that example really brings home why the embodied approach is, like you say, a simple idea once it's pointed out to you.

(Music)

GC: Now, your work is involved in robotics, and more specifically biorobotics. Can you tell my listeners about what biorobotics is?

RP: On the one hand we have industrial robotics. And industrial robots are programmable; you can program them to do various sorts of things. But once they're programmed to do a particular task, like at an assembly line or something, then you want these robots to do exactly what they're programmed for, and you don't want them to come up all of a sudden with an interesting idea.

In biorobotics we target completely different applications. So, we look at robots that have to function in the real world. We look, for example, at tasks that humans do. In humanoid robotics people would like to be able to do these tasks, like grasping a bottle, or arranging dishes, or whatever humans can do. And then it makes, of course, a lot of sense—because humans can do these things very well—to try and understand how humans actually do these tasks.

And then from this understanding we try to make certain abstractions which then should hold not only for the biological system, but also for artificial systems. And then, of course, we try to employ these principles. So, always of course we do have to make abstractions—we have different materials; you know we don't have biological tissue—but I think at some level of abstraction we have some kind of biological realism, if you like.

And we do that not only for humans. We work a lot also with insect biologists. I think we can learn a lot from insects. We also, for example, build robot fish, we build robot dogs—quadrupeds—and of course we build humanoids—bipeds.

GC: In your book you mention that one of the fundamental discoveries from embodied artificial intelligence is that the close coupling between sensory and motor systems is essential for intelligent behavior. And you've kind of alluded to this idea. Do you want to say anything else about that?

RP: The coupling of the sensory and motor systems: Well, maybe I should give another example, because in my opinion the importance of this idea can hardly be overestimated. If you think about walking – let's say you are walking – and I think there's a very nice example that you can actually exploit. You can of course exploit material properties. Maybe we can come to that also—that we can exploit material properties. You can also exploit the interaction with the environment. I don't know whether we should talk about this now, or –

GC: Well, did you have another example you wanted to give?

RP: Yes. I have a couple of examples. One is I think it makes sense when you look at intelligent behavior to not only look at what the brain does—what the role of the brain is in the generation of behavior—but it makes sense to look at how the brain cooperates with the rest of the body. I think it's better to look at the kind of cooperation, and in some sense that the brain orchestrates the interaction with the environment, but it's not such a tight control.

For example, if you are running, then you come down on your feet on impact, then the knee makes a very rapid damped oscillation: very rapid, very small damped oscillation. Now, this oscillation is on the time scale—I'm not exactly sure—but I think it's on the order of microseconds. Now, this is way beyond the time scale of the neural system, which is more like the millisecond range. But how then is this movement controlled?

Well, we cooperate a lot with people from biomechanics, and what they tell us is that the brain dynamically changes the material properties of the legs; that is, the muscles. So, it changes the stiffness of the muscles. And then on impart the stiffness has to be high, but then the brain is in some sense outsourcing this functionality, this control, to the material properties of the body. I think that's a really essential message, that the brain can outsource functionality to the morphological and material properties.

Or, just to give you another example of what I mean by this, normally when you have your arm on your desk—let's say your right arm—then the hand is facing left. Now, you can also put the right arm on the table such that that the hand is facing right, but this requires much more effort. Also, we don't grasp things like that normally—though we can. So, this requires a lot more effort because the muscle-tendon system is like a spring system. So, it's like winding a spring.

Now, if you let go, then the arm will turn back into its normal resting position with very little control from the brain, by exploiting the material properties of the muscle-tendon system as a dynamic spring system. So, again the brain is outsourcing a large part of the control to the morphological and material properties of the body.

GC: I don't know if this is the right time for this example, but I can't get it out of my mind. What about the example of the insect walking?

RP: Yes. I think it's a good time. So, you can exploit morphological and material properties, and there is an interaction with the environment that you can also exploit. Now, the insect walking, it has been known for quite some time that in insect walking there is no center in the brain for controlling globally the coordination between the legs in walking.

And this is a really complicated problem. You know insects have six legs, they have three joints per leg, so that makes 18 joints. If you want to solve the inverse kinematic problem, which is the standard way of doing these things in classical robotics, then I think you have to do a lot, a lot, a lot of computation. And we don't assume that ants are actually solving the inverse kinematics problem. Although we don't know—they might—but we assume that they don't.

Now, Holk Cruse, a German biologist, has been studying insect walking for about 30 years, and I think he knows how this actually works and how this global communication between the legs functions. So, imagine that the insect is standing on the ground with all its legs, and then it's pushing back with one leg. Now, what happens with the other legs—because the insect is an embodied system and there is gravity, it's connecting to all the other legs—so, if one leg pushes back, all the other joints in all the legs move in the forward direction. And instantaneously they move in a forward direction.

Now, because these angles change, all you need is to have angle sensors in the joints and there you have global communication between the legs. But not through the neural system, but through the interaction with the environment. So, there are basically two advantages to this. It's really fast, and you don't need neural resources, which are much more expensive. So, basically you're just exploiting something which is there anyhow, which is a principle that we call the principle of cheap design.

GC: So, that's the reason why people in biorobotics are trying to do things like make robotic insects. Because isn't walking one of the big challenges in robotics?

RP: Absolutely. Walking is very fundamental. And it's a challenge in robotics, but it's also a challenge for artificial intelligence.

By the way, there is a beautiful quote by British biologist Louis Wolpert, who said, 'Why do plants not have brains?' The answer is actually quite simple.

GC: Because they don't have to move.

RP: Right. Exactly. And so, the idea is that if we really want to understand the emergence of intelligence over evolutionary time, or over also ontogenetic time, then we have to understand movement, we have to understand locomotion. And now if you think about this, whenever you move—assume that you're walking—you're also inducing sensory stimulation; again, the sensorimotor coupling.

So, for example, you walk, you have pressure sensors on the feet and you have force sensors in the tendons and the muscles, so you're generating regular patterns of sensory stimulation in your body. But also you are inducing optic flow—optic flow meaning the speed at which the environment moves across your visual field—and this optic flow turns out to be extremely useful for many, many different tasks.

(Music)

GC: I'm going to move onto a different topic now, although obviously we could talk about locomotion; that would be like a whole podcast. In your book this line struck me very strongly. You said, "Behavior is always emergent." And that's kind of a buzz word, so I was wondering if you could describe emergence.

RP: Yes, I know. Emergence has a bad reputation. I mean people used to say—and people in artificial intelligence are maybe a bit guilty themselves—they say if people in artificial intelligence don't understand a phenomenon they call it emergence. I think that's actually complete nonsense because emergent phenomena can be understood in a perfectly rational way. And what do we mean by that? I have an underlying mechanism—let's say we can look at the brain, or in a robot we have an algorithm on a microprocessor—and then this has to be translated into a physical behavior.

Now, we normally think—and that is really the classical thinking; I think it has probably very strong Cartesian roots—that there is a direct correspondence between the control signal and the actual behavior. Now, you can imagine that for example you want to move your arm, and there is an obstacle in the way of the arm. And even though the control signal is pushing and would move the arm if the obstacle weren't there, it doesn't move the way the control signal wants the arm to move. So, there is always this interaction with the environment, and the behavior is control signal plus interaction with the environment.

I need to give you another example, because also material properties play an essential role in the emergence of behavior. For example, we have a very primitive quadruped robot. It can run very fast, but it's very simple. It doesn't have any sensors, and it just has four servomotors that make an oscillation movement back and forth. And then each leg has two passive joints, and the links are connected by simple springs, and then the fore and the hind legs just move. You put the robot on the ground, and then it starts smoothly picking up its dynamics and it starts running in a very nice and relatively natural way. I mean as natural as these robots can be.

Now, the movement is just an oscillation, but the movement of the robot dog is a nice running gait. Of course once you understand the material properties, once you understand the environment, then you can make a connection between the control and the robot. That's what I mean by understanding emergence. But you cannot derive from the neural signal the behavior of the robot.

And that holds for biological systems as well. If you want to understand the purpose or the function of a certain brain area, or some brain circuits, you need to understand how this brain is embedded into the organism and how it's connected to the real world. Otherwise just looking at the brain doesn't give you that. I mean that would be making a category error—philosophers told me that this is called a category error. So, basically behavior is something different from the underlying mechanism.

GC: So, would you say it's fair to say that emergence is a result of the fact that the system is dynamical, or complex, or non-linear?

RP: Yes. I mean there are different types of emergence. So, you have this kind of emergence in individual robots. Maybe it's good to give an additional example, because this is really important. We have one robot which is really a very strange construction. It's made of some very elastic material and it's kind of really bent, and then it's got two like skis—not feet, but more like skis—on the ground. And then at the shoulder you have kind of like an elastic blade, and at the end of these elastic blades we have two rotating electrical motors.

Now, when these motors start rotating, the whole creature, or the whole construction, starts to oscillate. It begins to oscillate, and by the way it's constructed it turns into a hopping forward movement. Now, if you look at the control signal it's just the rotating motors, but this through the morphology and through the interaction with the environment is translated into a forward movement. If you just look at the control signals for the motors, which just turn these motors—they don't turn wheels, they just provide the oscillation—you have absolutely no chance of predicting this forward movement.

GC: I find myself thinking about a baseball pitcher throwing a curve ball, and how dependent it is on the fact that the ball has seams.

RP: OK. I've never thought about that.

GC: Well, I have a lot of listeners who are not Americans, so they probably won't get much out of that either.

RP: Right.

GC: I think you may have a picture of this thing you just described in your book. You've got a lot of good pictures of different robots in your book, *How the Body Shapes the Way We Think*.

I really can't do justice to everything that's in your book. One of your themes that I was particularly interested in, though, was that you had the goal of trying to create a theory of intelligence.

RP: Right. A slightly ambitious task.

GC: Yes, but you know this is an ambitious podcast. So, what would a theory of intelligence like this accomplish?

RP: I think a theory of intelligence should on the one hand enable us to understand biological systems and how the behavior of biological systems comes about, and on the other hand it should enable us to build intelligent systems. So, the approach that we like to use is a synthetic approach, and so I think the theory should not only provide an understanding of existing biological systems—I mean psychologists, biologists, neuroscientists, they will be more than happy if they have an understanding of the biological system—but we also would like to be able to use the theory to build systems. In a sense we would like actually to marry science and engineering.

GC: By synthetic you mean you build stuff.

RP: Right. So, we build stuff. We call it the synthetic methodology. It actually goes back a long time. I think Grey Walter, a neuroscientist that was also building robots in the 1950's, called it a synthetic approach. The idea is the following: You have a phenomenon in nature that you're interested in—it can be anything, like people recognize a face in a crowd, like an ant finds its way back to the nest after it's found some food, or take your favorite example—and then we try to build a system that mimics at least certain aspects of this behavior.

And it's amazing. It's an extremely powerful method, because if you try to build something, you immediately see it doesn't work. Or, it does work. It's as Richard Feynman said—I think he once said, 'If I can't build it I don't understand it.'

GC: Yes, I think he was the one that said that.

You talked in the book about how you thought that Karl Popper's criterion of falsifiability is too restricted. But that doesn't mean you don't think a scientific theory should be testable?

RP: No, no, it has to be testable.

GC: What does testable mean in the context of your synthetic approach?

RP: Well, testability is one thing. And what does that mean? Testability, you set up an experiment, you make certain predictions, and you see whether your predictions actually are fulfilled or not. Of course, once you've built a robot then you can start predicting things. And it turns out that even if you have a robot that you built yourself, it's not always easy to predict. But I think predictability, or falsifiability, is one criterion.

But I think there are actually more important criteria. I think it's more like a process. And one extremely important point that other people made – and Gerry Edelman made that point very strongly. He said, 'Well, I'm not interested so much in whether a particular model is correct or not, or is true or not.' I mean that would be then how well the model can predict something in reality. But he said, 'It's kind of like a sequence.' So, he builds one model, and then that gives him ideas on what to do next.

So, I think the generation of ideas is at least as important for the scientific progress as whether a particular model can be tested or not. And if you actually build things, it's just unbelievable. I mean almost without thinking the ideas are generated. It's an extremely powerful method, and you have to experience that. And that's why more and more neuroscientists are also getting into building robots.

GC: So, it's another version of that whole theme that it's the questions that matter, not the answers. That's more of a philosophical point of view.

RP: Well, I think being able to ask interesting questions is at least as important as finding the right answers.

GC: I mean if you have a theory that doesn't generate any interesting questions, it's probably not going to generate any new ideas either.

RP: That's right. In our case the theory should in addition—and that's very important for us—provide principles that we can apply to the engineering of systems.

GC: Yes. And you went a great deal into the idea of design principles in your book. Since my podcast is not an engineering podcast I didn't really go into those specifically. But I do have some listeners who are engineers, so I'll point that out to them.

(Music)

RP: Now, maybe I should give you still one more example. I mean these design principles also apply to biological systems.

GC: Right.

RP: So, maybe it's of interest to actually give one or two examples of these principles.

GC: I wanted to ask you a question that I think will allude to one of the principles. In neuroscience there's a lot of controversy about the role of modularity. It seems to me that you had something you called the complete agent principle. Doesn't that kind of address that question?

RP: Yes. It addresses the question, at least in a kind of general way; you know, the complete agent as a dynamical system. And so, basically, of course, in some sense there is modularity. But in another sense there is also the complete agent, which is a complete dynamical system. And whenever I do something with one part—I do something with the hand—it's always the entire body that's involved, even though it does make sense to look at the hand separately. But the hand is always part of the whole body dynamics. You can't do anything with the hand without all the other parts in your body being affected by it.

GC: And when you start designing a robot you can't cheat this fact, right?

RP: Yes. I mean it's a tricky business because of course we have to use components. Ultimately then we use cameras, we use springs, we use joints, we use electrical motors, we use links. So, we do use components, so in some sense there is kind of a modularity.

But I think when you're designing a robot you have to always keep the entire robot in mind, because you cannot just fabricate the individual modules, put them together, and then sort of think that they interact linearly so that there's something like linear superposition. But you have to take the entire agent into account. And then you're always in for a surprise.

GC: Did you want to say something else about the design principles?

RP: Well, just to make it maybe slightly more specific. There is, for example, one principle—which is also I think in neuroscience very important—the redundancy principle. There is a trivial meaning of redundancy, which means duplication of components. Now, I don't think that's very interesting. Because assume that instead of two eyes you have 100 eyes. If it all of a sudden gets dark, then the additional 98 eyes are not going to do you very much good.

However, because we also have our touch senses, for example, or we have our acoustic senses, they also give you spatial information and enable you to move around. Your touch sensors that you have on the hand, for example, also give you the possibility to acquire geometric information about the environment. And these touch sensors, because they're based on mechanical touch rather than electromagnetic waves, continue to function when it's completely dark. They don't for locomotion or orientation in space; they don't function so well because you can reach only as far as your arms. But you still can continue to move.

So, there is a partial overlap in the information that you can acquire from the different sensory systems. You can get very good geometric information from the visual system, you can also get very good geometric information from the haptic system, but the information that you can extract is not identical. Though over time then, through cross-modal learning, one sensory modality can become a partial predictor for the sensory stimulation in another modality.

So, over time babies can learn that when we just look at an object—say a bottle of water, or a can of beer—we can already just by looking at it make a partial prediction of what it will feel like when we actually touch it. And I think this kind of prediction is very important, and for this prediction to work we need this partial overlap in the information so that we can extract the mutual information.

GC: Yes, I think it's been shown that if you hear a sound then you can pick an object out of a scene quicker, because you are expecting to see that particular thing.

RP: Exactly. It's that kind of thing that really is at the heart of this redundancy principle. And I think it's very important, to understand biological systems, that we really apply this principle

GC: You've already talked about your quadruped robot. That's the same as what you refer to as 'Puppy' in the book?

RP: That's right.

GC: There's a link in your book to a video, but the link is not good, and I was just wondering have you put that video up? And if you have, maybe you could let me have the link so that I can put it in the Show Notes.

RP: I think they can all be seen on our web page. On our web page there is a link called 'Robots' and I think they can be seen there.

GC: OK, I'll put a link to your web page up on there. And also if they want to know more about the engineering aspects, I guess your interview on *Talking Robots* was a little bit more along that design line.

RP: From EPFL. Yes.

GC: Yes. Because that's actually how I discovered your book; I heard you on *Talking Robots*, and that's why I bought your book. So, you need to know that being on a podcast can sell books; at least one.

RP: Oh, wow. OK.

GC: And I'll probably have a few listeners who will buy your book, too.

RP: OK. Excellent.

GC: We're just about out of time, but what's the most surprising discovery that you've made so far, or has come from your work?

RP: Well, I think maybe the most important—I don't know whether it's the most surprising—I think it's the interaction of the physical processes and the information processing of the brain, or of the control, if you like. I think this coming together of all these things, and then the morphological component. You can take the example of the grasping. When you see all the ideas that are coming together in this process then you start realizing the importance of this concept of embodiment.

You know it's the morphological and material properties, the biomechanical constraints that constrain your movement so that the useful exploratory movements are the easiest ones to control. And then how it comes together with the distribution of the sensors on the organism, and that they are stimulated in the right sorts of ways because this is already constrained. And then when you grasp and you generate sensory stimulation, and this sensory stimulation contains correlations, which facilitates learning. So, there is an entire story. And I think it's the beauty of this story which I like best about the concept of embodiment.

GC: Now, before reading your book I was kind of like most people, 'Why can't they make computers that have better vision, or walk better?' But now it seems obvious what the problem is. First of all we've come to appreciate that not only is it not computational, but it's complicated in a different way.

RP: Right. That's exactly right. I mean people try to apply computation. That's the thing we understand. You know we don't understand how else it could be. And I think we have to give up the idea that the brain is really controlling everything. What the brain does, it sort of outsources control. It sets the parameters of a dynamical system.

For example, when I have a loosely swinging arm the brain sets the parameters of the dynamical system in the sense that there is a certain muscle tone, there is a certain posture of the body. If I don't have the muscle tone—as in paralyzed people, for example—the arm swings in a completely different way. But the brain doesn't have to control the details of all these movements. And I think that's very fundamental to understand. And it gets us to think about ourselves and the world around us in a completely different way.

GC: And then, I know that you had that example about how ants find their way back to food, and how something that appears to be very complex doesn't necessarily require a complex control system.

RP: Exactly.

GC: That's the other side of the coin.

RP: That's the other side of the coin. Actually that example goes back to Herb Simon, who was one of the initiators of artificial intelligence, and that's why we call it Simon's ant on the beach. So, you see an ant crawling on the beach there. If you look at the path it makes—the trajectory—it looks very complicated.

But maybe the way this comes about—and that's again this idea of emergence—is this behavior emerges from an interaction of the ant with the environment. And the underlying mechanism is maybe just: if there is an obstacle on the left, turn right; if there is an obstacle on the right, turn left; and otherwise go straight. And that's it. And then you have a complicated environment. And then the ant has a complicated path. So, a very simple generating mechanism, but seemingly complex behavior.

GC: Well, we've only scratched the surface of all the ideas in your book. Is there anything else important that you want to share with my listeners before we close?

RP: Well, of course, there are many other ideas that are now beginning to pick up momentum, one of them being modular robotics. I mean we talked about modularity. Now, of course, if you look at something like the human body, it's built from modules. I mean if we look at biological cells, in a sense you can look at them as modules.

Now, because we feel that a lot of the intelligence can be outsourced to the morphology of a system, being able to build arbitrary morphologies is a really important thing. Not only building morphologies, but once they're destroyed we want to regrow them. I mean the human body has an enormous capacity for self-repair and things like that.

So, we started studying—and many people in the field started studying—modular robots, which are built from modules that can assemble themselves in a process of self-organization. So, we don't want to control everything in a top-down manner, but just as biological systems have some developmental, let's say, conditions—you know, like genetic regulatory networks, or things like that—if you want to build really complex systems I think you have to grow systems.

And so, one direction there is to study self-assembling robots—robots that can assemble themselves—and then the global structure also emerges from local interactions between the robots, which are not globally controlled. And then—and we don't understand this very well—often this structure that spontaneously forms through a process of self-organization has a certain functionality.

For example, we have made experiments where we have tiles swimming on water. You know they just have vibrators on them, non-specific energy. They assemble into a pizza-like structure, or a pie-like structure, and then the structure starts rotating, although the individual parts by themselves don't rotate. So, I think this is a really exciting field that tries to make a very close connection also to the growth in biological systems.

GC: And you also mention in your book, I think, about artificial evolution.

RP: Right.

GC: And that's definitely more than we can get into today.

RP: Maybe just very briefly. I mean the most brilliant designer that we know is evolution, even though—and that's the fascinating thing about it—evolution is completely blind. Evolution doesn't want anything, but things emerge. Also here we can apply the term 'emergence.' Humans have emerged over evolutionary time. And so, it seems natural. It's an idea that actually goes back to Ingo

Rechenberg of the Technical University in Berlin, and roughly at the same time John Holland in the U.S., who tried to exploit evolution for engineering problems.

So, trying to make a very abstract model of evolution—now most people call this genetic algorithms; I think that's a term that most people have heard—using genetic algorithms to solve hard engineering problems. And that actually, I have to say, works very, very well.

There is each year an Evolutionary Computation Conference, and at this conference there is a human competitive session where evolutionary programs compete with human designers, in a sense. And I think it was like a year or two ago where NASA was trying to design antennas for their satellites; I think it was communication satellites, but I'm not sure. And the designs of the artificial evolutionary systems were superior to those that human engineers had come up with.

GC: That's pretty amazing. But I think it's an example of how your synthetic approach can bring us knowledge that we wouldn't have any other way.

RP: I think so. Yes. I think it's a very nice example. And we can learn a lot from artificial evolution. I mean we have to be aware of the fact that biological evolution is much, much, much more complicated than what we have in artificial evolution. Even more surprising is the fact that we can do very, very interesting things. We can design systems with a level of complexity that would be very hard to design otherwise.

For example, we use a lot of artificial neural networks—which are inspired of course by biological brains—to control our robots. And especially if these are like dynamic neural networks with recurrent connections so that the networks have their own intrinsic dynamics, it's very hard to know how to design these networks, let's say, manually or by just thinking about it. And for that we always use artificial evolution.

GC: Well, Rolf, I really appreciate you taking the time to talk with me today, and I wish you the best with all your work in the future.

RP: Well, thanks for asking me. It was a great pleasure.

GC: And I hope maybe some of my listeners will drop you an email.

RP: Yes. That would be fantastic. I would like that. OK. Well, thanks very much.

GC: Thank you, Rolf. Bye.

(Music)

I hope this interview will stimulate some conversation at the Discussion Forum, which is at brainscienceforum.com. You can also leave comments at brainsciencepodcast.com or send me email at docartemis@gmail.com.

Recently I have been encouraging many of you who have been sending me emails to participate in the forum. Even though I don't respond to every post on the forum, I do read them all. And if you put your writing there you're more likely to get some feedback from other listeners.

Before I sign off today I need to talk a little bit about money. We are approaching the first anniversary of the ***Brain Science Podcast***, and I really appreciate all of you who have written to me to let me know that you're listening. If you listen to other science podcasts you know that the ***Brain Science Podcast*** is fairly unique in that it is not supported by any organization or magazine. And it's a solo effort. As you might imagine, each episode takes quite a bit of time to prepare.

My long-term goal is to be able to reduce the time that I work in the emergency room so that I can devote full time to podcasting. That means that I have to figure out a way to generate income. But obviously I want the podcast to remain free. So far I've made a few hundred dollars from advertising, but advertisers seem to be more interested in shows that have much larger audiences.

Here in the United States many of us make regular donations to public radio. This has allowed NPR to survive several attempts by congress to kill it. We are willing to do this because we know that NPR gives us content we won't get anywhere else. So, I'm hoping that you will feel the same way about the ***Brain Science Podcast***.

Recently I added a new page on my website that is labeled, 'Donations and Subscriptions.' It is now possible for you to make donations to support the ***Brain Science Podcast***. You can make single donations of any amount, or you can sign up for a monthly subscription, which you can of course cancel at any time. For example, if you gave $2.00 a month, that would be $1.00 per podcast.

Of course, cynics would say that this approach can never work, but independent musicians are also showing that this can work—that people are willing to pay for things that they think are valuable. All I ask is that you consider what you get from the ***Brain Science Podcast*** and what you think that it is worth to you.

Meanwhile, looking forward to the future, I have several exciting interviews planned for upcoming months, and we're going to continue to have book discussions. I'm going to start a thread in the forum about upcoming interviews. If you're interested in submitting questions to upcoming guests you'll definitely want to check that out at brainscienceforum.com.

Again, I want to thank you all for listening. I look forward to talking to you again soon.

(Music)