

Hofolding, Germany
1:35:14
m4v

Ernst Dickmanns

An interview conducted by
Peter Asaro

June 21, 2010

Ernst Dickmanns: Yep. I'm Ernst Dickmanns born in Newcastle, near Cologne in 1936 January 4th, and I went to school in Lutztoph, a neighbor – a neighboring village, where my father was a – a teacher. And then I went to the Gymnasium starting in 1947 in Porz which is not part of Cologne, the city of Cologne, at that time it was a city of its own. And of course, like in Germany, we went to Gymnasium for nine years. And then after that I did a practical work in industry, which you had to do at that time if you wanted to become an engineer. I wanted to do aerospace engineering. And this I started 1956 in Aachen.

Q: What kind of industry were you working?

Ernst Dickmanns: I did my, my practical work here in near Munich with Dornier Aircraft Engineering, and another part which was – what is it called in English? Mining, the ores ...

Q: Oh, mining, yeah.

Ernst Dickmanns: Well, no, if you process the ores...

Q: Metallurgy.

Ernst Dickmanns: Metallurgy, yes. This I did in Kulstoph, which is close to Cologne again.

Q: Okay, and then?

Ernst Dickmanns: And then in 1961 I made my diploma from engineer, and I went to Agentur, at that time Deutsche Forschungs-und Versuchsanstalt, which was something like the German NASA at that time, in Mülheim. And 1 year later in 1963 we moved on to Oberpfaffenhofen where we're still located right now. And from 1962 to 1971 I was essentially associated with the aerospace research, and I did a study on NASA fellowship at Princeton University, New Jersey in 1964-65, essentially control engineering and aerospace applications. And when I returned I started working on an optimization project numerical methods for trajectory optimization, and with this I did make my PhD in 1969 at the Technical University in Aachen with reentry vehicles and turning the plane of reentry vehicles by tipping into the atmosphere and flying out again at that time. This was one of the future maneuvers which I thought would be used possibly in the space shuttle, but it turned out that the shuttle was so delicate with respect to heating that they never used maneuvering when reentering. So this was more an academic work.

Q: Oh.

Ernst Dickmanns: But, being in the U.S., I met one of my co-students from Aachen. He was with NASA Huntsville and on a drive down to a sitting lodge in 1971 – no, in 1964 when I was in Princeton, I met him, and we kept this connection for the next 15 years or so. And he invited me after I had my PhD to come for a postdoctoral research associate-ship to Huntsville, Alabama, which I did in 1971-72, and there I worked on the shuttle orbiter reentry. This was the field I had been working in with my dissertation. So we did some investigations and several

publications in this area, and then I came – went back to Germany, but at that time I think all the activities in this direction of future launching systems in Europe, they were very delicate, so there was not enough money so they abandoned it pretty soon. And that's when I started working with satellite control and my group did – what do you call it – launch and positioning of the first European communication satellite symphony in 1974-75. And after that at the 1st of January I think it was in 1975, I became acting head of the research center Oberpfaffenhofen DLR. And I was in charge of about 700 people and, I don't know, half a dozen or two dozen institutes and smaller installations at that center.

In 1972, so one year earlier, the Federal Army University – we visited the Bundeswehr where the Bundeswehr was founded in Neubiberg which is close to here, and I got an invitation to give a presentation and possibly become a professor in the aerospace department of this university. And after some hesitation, I decided in October of 1975 to go there. During my activities at the Oberpfaffenhofen Research Center of DLR, I had become acquainted with satellite imaging essentially for remote control, no, not remote control, remote sensing, and being a control engineer. I had the impression with progress which you could observe increasing computer power per volume, per performance unit, and also the prices, so that it looked like you could afford sufficient computing power for real-time image processing somewhere ten to fifteen years down the road. So I decided to do all the installations at the university in the direction of developing vision systems for mobile systems, and I had the application areas of aircraft, space vehicles, and above-ground vehicles. And with this topic, we started in 1977. The first dissertation started in 1977. So we had the pole-balancing first, then we did a satellite control, I can show you in a movie if you like. This was an air cushion vehicle floating on a table two by three meters and doing, by reaction jet control, doing docking to another satellite on the table. And after this had been successful, we decided to go for grants and that's when the development of our vision system as a relatively big line of development started here in Germany.

So we were approached by Daimler-Benz for the – we could do this together, and we decided to cooperate with Daimler in developing the sense of vision for vehicles. At that time, possibly you should know that 1986 was the 100-year anniversary of the first car being built by Daimler and Benz. So they approached us and said – there was a big European framework, the Eureka Framework, where industries should cooperate in order to compete with the U.S. and the Japanese developments in these areas. So Daimler proposed to do a large-scale research, half a billion marks at that time, for the development of technologies for the second century of car developments. And vision became one of these issues. So we were successful in eliminating buried cables for driving autonomously on the Autobahn, and we said we could do it by vision. And we got this contract, and – well, there were several other steps in between. I think I skipped these. You can see them in the video. And we made a final demonstration of this project, which ran for seven years, from '87-94, at the Paris demonstration in October 1994. For this demonstration, Daimler had equipped two S500s, cars, with vision systems, and, um, we drove in normal traffic around the airport Charles-de-Gaulle, three lane traffic, with speeds up to 130

kilometers, doing autonomous lane changes, with guests onboard. So this was the first demonstration, I think, in public to a large audience. In total our two vehicles drove several thousand kilometers on the Autoroad 1 in France.

And the next step was a continuation of this from European funding, so we switched to the next generation. The first one was done with transputers, and then we switched to the next generation of transputers, we hoped, but this didn't, didn't realize, the nine-series never came to existence. So we switched to conventional processing with the methods we have developed, which have derived from control engineering recursive estimation methods quite contrary to what has been done in the AI and computer science community where they looked at image processing essentially as a cause – a static image evaluation, and by evaluating sequences of images, they wanted to arrive at an interpretation of the motion along the timeline. Our approach was quite different. We said we want to do a high-evaluation rate right from the beginning. I made or set the limit at .1 seconds, so 10 frames a second. And then we decreased the complexity of the task, and we wanted to increase the complexity when more computing power became available.

And of course, what you could observe was at every PhD generation, I say 4-5 years, the computing power increased by a factor of 10. So within 15 years, we started 1977, from '80-95, this is a factor of 1000 in computing power. And then, at that time, there was sufficient speed and the systems were smart enough so we could switch to off-the-shelf systems. Up to then we had custom-designed systems. But not, as in the U.S., developing systems completely new with hundreds and thousands and ten thousands of processors, maybe single-bit processors, but we decided right from the beginning to go for conventional microprocessors because we saw that there's a big difference between the biological systems, which was one of the predecessors where the development in your country, in the U.S., they looked at and said the technical should look similar to the, um, the biological eye. And quite a bit of the development was going in the direction of coming from pixels to processing structures behind that, and we said there is no need for combining image-taking with processing because in silicon, you have 10^5 , 10^6 times higher bandwidth, so you can send not just 2, 3, 4, but dozens of images in on frame-time. So there's no need to do this, and we could stick to the conventional microprocessors.

But of course, the evaluation has to adjust to these properties, and this of course, could be done with the recursive estimation technologies which have been around since 1960, since the Kalman filter, but not as has been done in the computer science community by looking at static images, but by looking at real image sequences with models for motion of 3D objects. So these were 4D, 3D motion of 3D objects in time, over time. This would be called our approach, the 4D approach. And I could show you afterwards the basic idea of if you graph, just systematically build up, you can see. And this allowed us to achieve maybe 10 times the performance of autonomous vehicles with .1 or .01 computing power onboard the vehicle. So the community was quite surprised when we demonstrated our first running on the Autobahn in 1987 at speeds, and the maximum speed of our vehicle VaMoRs, at 96 kilometers per hour. When in the DARPA

project, the ALV, and the NAVlab, they crawled around 2, 3, 5 kilometers per hour, and they evaluated, say for one image, between 1 to 10 seconds. And this was only possible by developing these methods, which we called 4D approach. And if you like, we can look into this.

Q: Yeah, so let's talk a little bit about, more about that. So how would that relate to James Gibson's theories of perception or David Marr's, sort of – he has this 2D, 2½D, 3D. So how do you construct, reconstruct 3 dimensions out of 2 dimensions is what he –

Ernst Dickmanns: Yes.

Q: –sort of thinks of as the vision problem. Do you buy, do you endorse that perspective?

Ernst Dickmanns: No. No. It's not. We, we reject it because we feel you have to have available all 4 dimensions right from the beginning, and your models should have the time component and the 3D space components. And then we talk about motion of 3D objects with features over time in 3D space, and we're looking for features. Essentially we were looking at corners and edges, that we could extract relatively simply, and from this we have to reconstruct the objects. In addition to the edges it is important to, in order to make distinctions between objects, to have the, what is it, the grey value at the side of the edges, so this we talk in addition. And then we were able to look at motion from one frame to the next. We made predictions. And what we did right from the beginning was prediction-error-feedback, so according to the model, we assumed a 3D object with a 3D motion, and then we made a prediction how should this develop over time. But we did not just take the model as the nominal model, but we also used a set of linear models around the nominal one. And then, in the next frame, we compared the prediction with what really happened in the real world. And that's why we call it prediction-error-feedback. And this was a breakthrough in real-time processing.

So most of the computer science people didn't believe the results when we first showed them. And of course it was very simple. It was white lines and dark-to-bright transitions at the side of the road. We didn't have to have white, white lane markings on the road. But we could do it with a normal transition from macadam to just grass or whatever. And um, then of course, we have to be very careful to check whether the initial assumptions about the world and the objects are correct or not. And it turned out that this seems to be very similar to what the biological systems are doing, and I think right now the developments you can observe in the theory of mind is according exactly to these models because the Spiegelneuronen – what is it, the mirror neurons – what they do is essentially store developments of appearances over time. And from this you can see directly not just a snapshot, but if you see a snapshot for you in perception, it's a snapshot of a motion, so you perceive the motion. And that's a big difference to reconstructing 2D, 2½D, 3D. I don't believe in that. 2D plus time? We have to work to retrieve 3D plus time. And of course you do have to continue the conditions over time which is very essential. And you do have to, in the interpretation, you right from the beginning.

Q: Are you familiar with James Gibson, J.J. Gibson's –

Ernst Dickmanns: A bit.

Q: – theories of affordances, 'cause that's a lot about this, that action is what gives you that third dimension as you move through space, but that wasn't influential on your work?

Ernst Dickmanns: No, no. We directly came from control engineering, so this is the Kalman filter approach. And we said what has been done wrong in the computer science community is to look at images and say the edges move linearly or maybe non-linearly so you add some noise in the image. We say no, what you see in the image is a perspective distorted mapping of the 3D world, and the interpretation is directly done in the 3D world, not in the image. Not in the image plane. It's directly done in the 3D world.

Q: So in terms of recognizing obstacles that are off the ground versus just a colored patch on the ground, so is your method then a lot more efficient than –

Ernst Dickmanns: Oh, yes, of course. And the good example, a nice example you can do is if you stand on a road with a camera in the vehicle, and you look at the road in front, you have a certain appearance of, like a pencil. And if you take an artist, and he puts a large sheet of paper on this, on one of the locations, say, 3 meters in front of the vehicle, and you paint the outside road exactly as it is, as it appears to you at this position, then there is no way to make a distinction between the real world and the painting. But if you move, your vehicle moves along the road, and the sidelines on the picture, they move to the side while looking at the real world, they shift along the road. So over time you see the discrepancy between a static interpretation which is wrong and a 3D interpretation which is right, which is correct. So this is very essential.

Q: Were there other sorts of visual tasks that were difficult for this methodology that you encountered, that were –?

Ernst Dickmanns: Well of course obstacle detection is one of the difficult ones, and the most difficult one we met later on, maybe ten years later. These were negative obstacles, which has been part of research in the U.S. in the military field, not so much in the university. Also not in the DARPA activities of Urban Challenge and Grand Challenge, but the U.S. Army research had, they had projects going on, and we've been cooperating with them since 1987-'88.

So there you do have the big difficulty that, in order to understand depth, of course you have to have stereo vision or some basal range finders. But if you drive at higher speeds, then you don't detect the obstacle early enough to do close reaction. So we did an interpretation with stereo interpretation, and with laser range finders, and with image interpretation, and we came to the conclusion that you should have a combined system. Both laser range finder and stereo interpretation and the visual interpretation because if there is a hole in the ground, you do have some certain properties of the horizontal surface, and you do have a, usually, a completely different visual property of the vertical part of the hole. So what we can do by normal vision detect the small vertical part. If you look at 10 meter distance on a road that is maybe 1 meter

wide and 1 meter deep, you just see a small, small upper fraction. And if there is grass in front of this, you're not able to detect it by, at least not at that time, you're not able to detect it by stereo. And also laser range finders, you do have problems detecting it. So we made the combination that the first detection is done by normal vision, and then if you close in, then of course you're lost with vision and then you have to look at the stereo interpretation or at the fine structure which you get from laser range finding.

Q: What –?

Ernst Dickmanns: So that's the most difficult problem we did.

Q: What about, like, water-filled potholes or –?

Ernst Dickmanns: Well of course –

Q: Are those even more challenging?

Ernst Dickmanns: – it's mirrored. It's the same as like for humans, I believe. It's just shallow water or it's deep water. You cannot detect what it is. Quite a bit of discussion has been going on in this field. My feeling is if you are on a normal road and up to then the road has been usual, well, almost flat, then there will be puddles, and usually you can assume you can drive through it. So if there are these appearances, we just assume it's not critical. If we go cross-country, which we did in the late '90s and early 2000s, the feel is completely different. Then you have to stop and you have to look what it is.

Q: Yeah. Yeah, it's a good assumption in German roads, but we have a lot more holes in American roads I'm afraid. So what about moving obstacles? So if a bicycle or a pedestrian comes out in front of the vehicle, how do you try to deal with those sorts of –

Ernst Dickmanns: Well you see, how the features shift from one to the next, you see the set of features which moves in conjunction, and then we come to the conclusion that this is an obstacle, a moving obstacle. The first investigation we did was the pole-balancing. Maybe you should have a look at the film so you'll really understand much better what is going on. I do have a 50 minute film which gives a survey on the approach with all the applications we've done. I don't know whether you have time, or if you would like to do it. It's also on the CD.

Q: Yeah, we can look at that –

Ernst Dickmanns: Shall we do this first?

Q: Well, I think I have – I think we can go on for now, but I'm –

Ernst Dickmanns: Okay.

Q: But as far as, so then from the control perspective, was it ever a problem to figure out when you should be braking, when you should be turning? How do you maintain the stability of the vehicle if you're trying to avoid a collision at high speed?

Ernst Dickmanns: Well of course these are difficult problems, but usually if you do the 4D representation then you also make predictions of not only how you are going to move, but also how the other objects are going to move, and you do have a representation of all objects in parallel, and then you see; that's a straightforward analysis will there be a possibility of a collision. Then of course you have to react. And car control, that's simple for an engineer. That's been done for a hundred years. That's not a new field.

Q: No, but you didn't have any issues with, like, hunting or over steering or anything like that then?

Ernst Dickmanns: Well this you see, essentially, if those people deal with car control who start from static images and who don't have the notion of maneuver. I didn't emphasize this, but this is very essential. If you do control application, I've been working in aerospace maneuvers extending over not seconds and minutes, but hours, so you always should have in mind what is the overall maneuver you are doing. What are the maneuver elements you are applying? Then you have a list of maneuver elements which you have to work upon one after the other, and then you have to find the transitions from one maneuver element to the next. But within the maneuver elements, there are control engineer and feedback loops, or if you want to make a lane, a turn-off, there is a feed-forward control. You do this for a lane change, so this means you do a direction change by cars, it's an Ackerman steering like all cars. And then if you do the other way you have an offset to the side, but going parallel if it's symmetric.

So this is, again, it's relatively simple, and if you do have some errors, most people, I've seen many people that try to correct this immediately. No, we say I'm in a maneuver, and how do I change the maneuver in order to achieve my goal? So the control corrections, the magnitude of the corrections is much smaller. I've seen very nervous control applications and these applications to ground vehicles, and I've sometimes had the impression that there's an order of magnitude difference between what you see in many approaches and what you see in our approach.

Q: Interesting. In terms of organizing a high-level maneuver, are you using a hierarchical control system –

Ernst Dickmanns: Yes

Q: or are these – at what level are these being represented, I guess?

Ernst Dickmanns: Maybe I should, I prepared some information here for you, so let's see. This is the overall system, and it's hard to explain just by words. I think one should look at the

different realizations. So these are the maneuvers, the maneuver capabilities, action and behavioral capabilities. This is a lane change. This is a position of non-linear feed-forward control and linear feedback control. And then here you do have the network of capabilities which indicates which other subsystems you need for realizing a certain behavior. And usually like in say, vision control, you do have one motor turning in yaw and the other in spin. And then you can combine this to do all the different maneuvers and saccades and whatever you like.

And you asked for the hierarchy of the system; this is the vision system. So first we do feature extraction just on the overall image without any meaning behind it – this is edges and corners. And then we immediately jump to the assumption of objects and motion of objects. And then we can set on top of this looking for special features. If it is this type of object I should see, and then we look for that, and either you see it or either you don't. And then we write the spatiotemporal coordinates in this, what do you call it, dynamic object base. And here is all the information on the objects that are relevant to the task you have to solve. And of course the data you have to handle is 2 to 3 orders of magnitude less than the image plane down here.

And on top of this there is the situation understanding level. It looks at the combination of objects or subjects. We call an object which is able to sense the world and do action according to what it sensed, this is called a subject. And of course we assume that we do know reasonable behaviors of other subjects. They want to survive and they behave essentially like how I control my car. So this is being done here, and the overall control system, that's how it's called in the loop, the overall control system looks like this.

So we do have vision control, which is with multiple cameras usually. Here is the capabilities of the vision system, capabilities of the vehicle. This is the behavior of the vehicle, and based on the collection of sense data, we come the conclusion of a situation. We do have a mission element over here. And there is the decision for gaze control and decision for vehicle control. And this is, there is of course an overall goal which you have to achieve.

This one shows the number of feedback loops. It's more, well, just an abstract representation of which we close. It's about half a dozen on the different levels, both image processing and feature extraction and feature interpretation in object control and stabilizing the interpretation of the situation. So this is very essential to this. Well I think that's about it for that.

Q: Yeah, that was very helpful. So in terms of other kinds of applications, how robust is that methodology, would you think for say, for a robot who wanted to walk through a crowded plaza instead of drive down a road? 'Cause on a road, you have rules that, you know, you stay on the right, there's traffic signals –

Ernst Dickmanns: No, I think it's completely applicable, and as far as I can see, it's the best method available to do any kind of control task of public systems. We did applications in satellite docking, applications in aircraft landing. I could show you we've been recognizing the relative position to a runway within a real flight in 1991 with experiments in Brunswick. And in

the loop simulations, we did helicopter control around the airport of Brunswick with different aspects and finally landing on top of the helicopter H, this mark on one of the taxiways. So it recognized it, maneuvered to it, hovered above it, and then came to a stop.

And if you do have, of course if you do have many objects moving not so nicely like usually cars do, but more erratic like pedestrians, you need computing power. We started looking at pedestrians in the early '90s, but at that time we were not able to do it in real-time. So I do have a video film showing this. But since 1990 to now, it's 2 decades, it's 4 times 5, so 10^4 this increase in computing power available now. So I'm pretty sure we could do it now, in real-time, with a dozen other subjects or objects.

But of course, the difficult task, and there we do have, did have, quite a bit of discussion with the industry, what are the proper sensors in order to do this? Do you have 20 or 40 cameras around the vehicle, or should you develop a system like the – like biology did? Moving the head and moving the eyes? And we came the conclusion that maybe the best compromise you can find is having fixed focus systems, about 3 to 4 cameras in conjunction, and then do gaze control into angular directions. And by this we reach say 120° , 130° viewing range should be sufficient in horizontal direction, and maybe 20° or 30° or 40° in vertical direction. And then you can do the rest just by looking where the interesting things are. And the same – one big discussion was well, we could afford, nowadays, with the cost of cameras, put several cameras on the vehicle and we'll have high resolution everywhere. And we'll scale down to lower resolution. But, again, my impression was doing the navigation through all these images, it's much more complex than developing an eye and have a wide viewing range with 2 cameras, with an overlapping central area where you can do stereo interpretation, and then direct the gaze control to the area where you want to see something more closely. And we proposed to have a color camera with medium range, say 100 meters range for the wide angle cameras have 20-30 meters good resolution. And then on top of that, in order to be able to have a resolution similar to the central part of the human eye, we say we should have another camera with a higher focal length, say at 300 meters, we should have a resolution of about 5 centimeters per pixel.

So these combinations of focal lengths, that's what we call the EMV vision system, MarVEye multiple ray? I could tell you later. My memory's – vehicle eye, vehicle eye, they call it the vehicle system expectation-based multifocal saccadic vision system. That's the one we developed, and I think this is a good compromise between the different aspects. Realizing the basic functioning of biological systems in silicon, on a silicon basis, not doing all which is necessary because in carbon, you are not able to get high-speed data transportation, but you're able to make thousands and ten thousands of cross-links, that's what we have in our head, which is very difficult to do in silicon, at least with the technology we have available. So we decided to make these simple systems, but then use the high data transfer rate in order to organize an overall system which does have the same functionality, very similar if not the same functional basis like biological systems, but realized completely different.

Q: That was the feature of the first design with the transputers, when you didn't have very much computation power even.

Ernst Dickmanns: Yeah, well the first system was even 1-2 orders of magnitude less than the transputers. So the first drive we did in 1987. It's hard to believe, but this was 8086 processors with 2 Megabyte cycle rate – basic frequency.

Q: Not 2 Megahertz?

Ernst Dickmanns: 2 Megahertz, yeah. Of course the world today is completely different, now we are to 2 Gigahertz, and you do have even higher, much higher. Data communication was a problem at that time too. You were not able – since you were not able, maybe I should mention this, since you were not able to look at the entire video image, a colleague of mine, Volker Graefe, he developed a system that I had proposed, he realized it, where the video signal is transferred line after line. And he developed a system that was able to grasp those pixels which were in a certain window marked by the upper-left and lower-right corner. And we were able to define about a dozen of those windows. So we could disintegrate the image into 12 subwindows, and then just extract it from those windows. But we were able to switch the window from one frame to the next. So this was very essential.

Q: Oh, and they were analog video.

Ernst Dickmanns: They were analog video, of course, at that time the cameras were this size.

Q: Yeah, so you had to choose the rate, the Hertz of the scans, and interlacing and de-interlacing, and drawing these boxes –

Ernst Dickmanns: Yes, yes, yes. We had to work with American and Japanese and German cameras, and we had to switch between 30 hertz and 25 hertz, all that stuff. NTSC and –

Q: – PAL.

Ernst Dickmanns: Yeah.

Q: So, how did – who – what was the inspiration for the work in vision. Like, did you come up – you said you saw that work was accelerating very quickly in that field when you were doing control systems, and you saw that convergence coming. Were there particular researchers that you collaborated with, or people whose work inspired you?

Ernst Dickmanns: No, at that time I was just picking up the Lehrstuhl, professorship for control engineering at the University of der Bundeswehr, and I met my colleague who's been working in measurement technology, measurement science, photograph, and he was also interested in that. And we decided since we couldn't get a system on the market which, which satisfied our requirements, so he said, “Well, I do, I think I do have an idea.” And he came up with a dozen microprocessors, Intel industry microprocessors, and he developed these systems

grabbing these windows and integrated into a system. And the first relatively high-dynamic stabilization we did was balancing a pole. And you know, depending on the length of the pole, the frequency goes down. So we started with a 1 meter pole, and we were able to come down to 50 centimeters; I think 30 centimeters was not possible at that time.

And we did this pole balancing on an electrocart, 3 kilogram, and acceleration of 0.8 G, 8 meters per square second with about 1-0.1 second frame time, with 4 Intel 8085 8-bit microprocessors. And when we came in 19 – what was it – 1981, we went to the first international conference and showed this, people didn't believe us. They said that's impossible. Because if you look at the edges you get from the pole you are balancing, they of course, you have to integrate over time in order to get the image, so we had to, to look at the front part of the edge because there was quite a bit of blurring going on.

Q: Yeah.

Ernst Dickmanns: And Meissner did this dissertation, and, together with Hans who was the other PhD student with Volker Graefe, they really developed a fine system. And this was the basic start for our method, what we learned from this simple application. The next application was the satellite docking, as in the air cushion vehicle, which of course was a very slow system by jet propulsion from compressed air. But these were relatively complicated hexagonal 3D satellite bodies, real bodies in 3D, and the system had to move around, and you can see it in the video. And it decided which corner to look in order to get the best interpretation with a limited amount of features available, 4 or 5. So it decided to look here, and when it was self-occlusion, something called “catastrophic event” at that time in the literature, it decided, well, yes, this is going to disappear, so I have to abandon this and go to another one.

And looking at the matrix of the linear approximation – what are they my memory faults on the names – Jacobian matrix. If you look at the Jacobian matrix – from the Jacobian matrix, you can set up a feature which tells you which combination of features to select in order to get the best interpretation. And this is what Wünsche developed, and he's now my follow-on at the University of der Bundeswehr, so he is, he is now working with the modern autonomous vehicles, which are now VW Touregs.

Q: Okay. So did Mercedes-Benz make any practical applications from your work –

Ernst Dickmanns: Yes.

Q: – that they use in their vehicles?

Ernst Dickmanns: The start, how to proceed, it was, well, it was quite some discussion because the question is vision is so complex and you have to invest quite a bit in order to get started. Before you get the first results. Is it really worthwhile doing it for just 1 function, or should we wait until after we have developed a system which is able to do 4 or 5 functions, and then you

can justify the cost of the system. But there was a divergence with respect to application to real-world application in production line around the mid-'90s when industry decided we should go very simple functions like lane departure warning. And we should install very simple systems in the car, and then do it on a minimal cost basis. And that's what was started, and these systems have been around since the early '90s. And you can buy some on the market right now. You can even buy them in middle-class cars right now.

First was lane departure warning. The second one was distance-keeping, and industry decided because the number of computations that you had to make and the computer systems you had to have available onboard were much lower for radar systems. So they decided to go for radar in distance-keeping. What it's called, adaptive cruise control. But then of course, in radar, you usually, you don't see the road. And it took some time until one realized that maybe, and when computer availability of low-cost computers became better, then the idea came that why don't we combine radar with vision. And these are the systems that have been developed in the late '90s-early 2000, and they are coming on the market right now. I don't know whether there is one on the market; I've been retired since 9 years.

Q: Well the Mercedes has this tension-assist, so it tries to monitor – I don't know exactly how the control system works, but it's supposed to indicate if you're getting sleepy? But it's looking at the driver input as well as the surrounding area?

Ernst Dickmanns: This is different. This is a different system, yes.

Q: Yeah.

Ernst Dickmanns: There are many, many different systems which have been single-eye in different systems that have been developed. What has been demonstrated in 1994 also in Paris was traffic sign recognition, but in order to do a real-time demonstration, they had a separate van with separate van with separate computers onboard, and they were able to analyze just 2 types of systems, passing not allowed, and I think 1 or 2 speed limits. So it was not possible to integrate this into a system at that time, which was sufficiently small to put it into a passenger car. In the meantime it's completely different. You can buy these systems reading traffic signs and showing them in the display on the dashboard. So these are available.

And what else? I think in preparation are warning systems for if you access a crossing to see whether there are vehicles coming from the side. So first single investigations have been done, started during the Prometheus project, but these were cameras mounted looking to one side, and camera mounted to this side, and one camera looking ahead. And of course if you look at all of this in a real-world system that you wanted to sell somebody, costs so high you couldn't do it.

And that's why we came to the conclusion we should develop a vehicle eye with the properties I mentioned before, and then develop the entire software around it to be very flexible.

And since industry was not willing to do this, or they didn't expect getting the money, so we decided to leave this type of civil development and join the more military developments which were going on both in Germany and the U.S. And since that time, since 1995, we did have cooperation with several American institutions like some of research, and National Institute of Standards and Technology, and we developed these systems. The EMS system we developed in connection with these, and the German Ministry of Defense.

Q: So what were some of the military applications? Or have they made applications out of it?

Ernst Dickmanns: Well of course, vehicle guidance on road-running, this is one. And detecting negative obstacles is another one. What is really available in application, I guess it's very little. It's still too early. It's too complex to be introduced and to be competitive. So it's still in the research area.

Q: What was my other question? So as far as the autonomously driving vehicles, were you ever riding in the vehicle while it was driving itself?

Ernst Dickmanns: Sure! It's a strange feeling the first time I got into the – it was a 5-ton van. And when this accelerates like a human driver, really, with all its power it has, for the first 10 minutes, you are very anxious what's going to happen, and then you see, well, he behaves like a human behaves, and you get accustomed to it. And we always had a safety driver in the seat. And he was able to intervene when something happened. But we've been driving not only in the car since 1986, but we've been driving in public traffic since 1992 in on the Autobahn, on Bundesstraßen, in on small roads, so. It was, well we had to have a special license from the German military organization to do this, but after they have observed us for I think about 10 years, yeah, something like 10 years, and they said, well, yeah, you can do it. But there have to be 2 or sometimes 3 people in the car. So we were not allowed to have a fully autonomous vehicle, nobody onboard, running in public traffic because I think they assumed that those people in the car, they wanted to survive, and they would be very careful.

Q: Good insurance policy. But you took it onto the Autobahn at very high speeds, correct?

Ernst Dickmanns: Well, in the European project Cleopatra, this was in 1995-'96, there were other groups. A Danish group, they did Schweißen, in ship-yarding, ship-building. What do you call this Schweißen?

Q: Sails?

Ernst Dickmanns: Fuse plates of iron?

Q: Oh, welding.

Ernst Dickmanns: Yes, welding. And there was a meeting in Odense, in Denmark, which is 1600 kilometers north of Munich, so they decided, 2 people decided in, after, in April or March,

the CMU group which was our competitor in the U.S., they had done the Hands All Through America, so the longitude control was done by a human driver, but the lateral control was done autonomously. And then we had a meeting in, I think it was September-October, and we decided to do a fully-autonomous drive to this meeting up there. And we wanted at the same time to collect material in order to see what are the most essential parts that have to be improved in the third-generation vision systems. This was the final of the transputer systems. We knew we had to switch to a different computer system, so we wanted to know what are the parts which have to be improved. And during this ride in the northern plains in <inaudible> the top speed driven autonomously was 180 kilometers an hour.

Peter Asaro: That's pretty fast.

Ernst Dickmanns: It's pretty fast. That's 110 miles an hour.

Peter Asaro: Were there ever moments where you were unsure of the robot or overrode it?

Ernst Dickmanns: Well, of course you sometimes see a misinterpretation of the situation. It needn't be dangerous. One thing I recall is there was one student looking at other cars as obstacles and then obstacle avoidance, and the system was looking on the Autobahn, two lanes and one car in each lane, and they suddenly came up with a hypothesis. It didn't see the left car, it didn't see the right car, but it interpreted one object in between of the two, so in looking for some time of course they moved to slightly different speeds, and then after a second or so or maybe one and a half seconds it made a decision, "No, that cannot be a correct hypothesis," so it started from new and it got two objects and then it converged towards the two objects. Of course it always takes some time, maybe a second or so, in order to recognize until it's sure what it interprets, and this was one of the conclusions we took from this test drive to the north. We should have system capability on all levels, both the feature extraction level, the object interpretation level, the situation interpretation level so that the system itself comes to a conclusion, "How sure am I about the data? Are my predictions correct? Are they getting worse?" And then there should be warnings. So this type of recognition and then calling for help, it's very essential. And one of the interesting points was that if you look at statistics of human traffic accidents there is a surprising lack of people that got a heart attack while driving. And people said "Oh, something must be wrong," and then somebody else came to the conclusion "Well, those people, maybe they feel that something is coming, and so they slow down and maybe stop at the side." And then they looked at the statistics and came to the conclusion "Well, if you include these cases, the statistic is okay." And this is the same conclusion we came to. The system itself should observe its properties, its performance on all the different levels and then come to a conclusion, "Am I sure that I can go ahead, or should I stop or should I slow down?" And of course on the Autobahn with a car you can always slow down. If you are in the air it's a little bit more difficult.

Peter Asaro: Fewer obstacles, though.

Ernst Dickmanns: Fewer obstacles. That's right. <laughs>

Peter Asaro: Do you foresee that eventually all cars are going to have autonomous driving capabilities or the majority of cars? How far off do you think that is?

Ernst Dickmanns: I think if we think in decades or centuries, yes, this is going to come. I tended to make prediction in 1985 or so, the end of the '80s, maybe in 2010 or so, 2015. I thought if we did have enough computer power that this would be the case, but seeing the difficulties it is a long way, and I'm pretty sure because of one meeting we had with the organization which is responsible for developing the traffic guide rules in Germany. There's a special institute, and we had a meeting with the legislative side, the automotive clubs, some legislators and the press, and then of course the same question was asked, and of course industry said "We are willing to do this, but there's one precondition. Everything what's happening in the car should be written down or should be protocolled so that later on you can see who did which control input." And then the people from the car <inaudible>, what is it, car clubs representing the drivers, they said "This is never going to happen as long as we are in charge and have something to say, so there will be no full documentation of what's going on in the car, not on a legalized basis." And I left this meeting with the impression that maybe the technical problems will be solved before the legislative problems will be solved, and this of course has to do something with the processes if there is an accident, and the situation of course is much worse if you can nail down what really happened. And this is the same situation right now. People are very hesitant to do full documentation of what's going on in the-- but if this is being done I think it's maybe a few decades in order that you can do this.

Peter Asaro: Do you think the robots will be safer drivers than people?

Ernst Dickmanns: Yes. In the average, yes. They don't drink alcohol. They don't get tired. So I think these monitoring systems-- and that's the first application-- this will be very useful, and they will work for the first two or three decades as monitoring systems. And then when maybe on the military side there's quite a bit of experience with real autonomous driving it may also come into private cars, but it's a long way to go.

Peter Asaro: Are some of your applications being used for actual satellites in space?

Ernst Dickmanns: Yes. We did some satellite docking. We looked at it, and the only application we've done was with <inaudible> grasping a free-floating object in space. This was in 1993 onboard of space shuttle Columbia, but this was inside the D2 box, which is inside the

shuttle. So there we did the remote control, because computing power was not sufficient onboard. We had to transfer the images to the ground, and this went through several satellites, through I don't know how many kilometers, 1,000 kilometers of cables, then there was a computation done in Oberpfaffenhofen, and of course there was a three-second time delay when the computer started working, and of course the signal going back to the satellite was another three seconds, so overall we did have six to seven seconds time delay, and we were able to compensate for this by predictions, because in outer space of course there are nice Newtonian conditions. It's a second-order system, except for the corrections which are being done on the shuttle, but the system itself, it's just Newtonian motion. And we've been able to show that-- and there's one film also on the CD where you can see that onboard March 2 or 3, 1993, there was the first grasping. In the meantime I think some activities have been done, both by the Japanese and the US and the Russians, but this was the first one. Real application I think wouldn't be too difficult, so the docking with the slow speeds, it will happen. It's just a question of "Would you do it?" because you lose one more justification for having humans onboard, because some agencies are very happy that they have justification for humans onboard.

Peter Asaro: Who are some of your intellectual influences as far as what inspired you to go into control systems or what inspired you to do some of these vision control systems for self-driving cars?

Ernst Dickmanns: Well, I was intrigued when I first learned about differential equations when I was 16 or 17. I recollect that learning differential equations and what you can do in describing the world. This was really a big event to me. And then I decided to become a control engineer and to do control engineering. Of course the most difficult things to control were aircraft, so I decided to do aircraft engineering, and after I've done it in Germany I went this one year to Princeton, where I looked in-depth into control engineering, and there it was Dunston Graham and <inaudible>. He wrote a book on dynamics and flat mechanics. And then when I came back, as I mentioned, we did this symphony positioning on the equatorial orbit, geostationary, and then the rest was just my own idea. I've seen what you could do with this recursive estimation part, which is nothing but Gauss's idea of least squares approximation, but Gauss did the least square approximation to batch numbers measured, and what Kalman did was he redefined it for a sequence of measurements coming in one after the other for space applications. This was very essential. So he did it in a recursive way, so now you should have dynamical models. Again, there are other differential equations, and this was the breakthrough. So I tried this, and, I don't know, it was the early '80s when we had a meeting, somewhere in <inaudible> I think it was. People from computer science, they didn't believe in it. They had heard Kalman filtering, but they didn't realize that what you have to do is formulate your models according to the real world in 3D space and time, everything in 3D, both the object shape and the motion, and correct in time, and then take the perspective model and use linear approximations to these strongly non-linear processes. So this was the breakthrough in the mid-'80s, so essential at that time. Maybe it's more the discussion with the people who didn't believe in it. I recall when we first made a presentation in 1987 at Santa Cruz in California there was a NATO Advanced Study Institute. At

that time they had these institutions where they brought together people from physics, from biology, from physiology, engineering in order to exchange what might be necessary for developing intelligence for vehicles. And when I first said "We've used Kalman filtering for interpreting motion image sequences" they said "Well, forget it. We've tried this for years, and you're not going to get far with that." And then the day after I showed my films already with the car running, and, well, that was a big surprise. So several names-- I don't want to quote them here, but they said "Well, that's surprising." So then we continued developing this, and one of those in America, which we had quite a bit of exchange, who was the most open one from all the American colleagues was Takeo Kanade, so I really appreciate the way he handled all of that. It was different from other experiences we had. That video equipment was gone when I had to give a presentation and things like this, but-- so he invited me to give a talk at CMU, and we had discussions, and I could discuss with the students, and they were open, and some students started in the same direction, and, well, they essentially said "Yeah, well, it looks like this is a good idea." And when I was at Caltech as a visiting professor in 1998 I recall that one Italian guy came back from a conference in Hawaii where he said "Well, it looks like all the developments of the special architectures for vision systems based on these multiple 110,000 million-fold single-bit processors-- it's going to be abandoned. It's no more seen as the way to go." And they switched back to the-- what is it? What do you call this? There's a short term. General purpose processor, GPP, and the development, which then just went from 60 to 32-bit or from 32-bit, the next step. Anyway, the communication also was high enough at that time that they came to the conclusion "We can do all of it with conventional microprocessors." And this is the way how things develop. I don't know whether you're aware of the first developments in the DARPA project on autonomous vehicles. They had about a dozen different architectures, specially designed architectures for vision. One of the most well known and widespread was a thinking machine, and I obviously had the impression that handling numbers and getting to an intellectual level is different from just working bottom-up. You have to have a high-end component. And very influential to me-- I think you asked this question-- were the thoughts of Schopenhauer and of course way back Kant. Kant has been misinterpreted here by the German idealists, and Schopenhauer put them from the head on the feet again, and he made the distinction between the real world and the world within our head, *Die Welt Aus Wille und Vorstellung*. That was the title of his main work. And I think this gave me some <speaks German>, confidence, into the approach, and I suddenly thought "Maybe this approach is not just good enough for technical interpretation, but you can really recognize or explain quite a bit of biological systems of the human mind and this world by going this way." So you have to start from a high-level representation and how this was generated. What was the hardware and the time involved, millions of years and I don't know how many hundreds of generations of development? What it boils down to is that there has to be some preconceived notion of a spatial-temporal world, and then you can combine sets of features with object classes, individuals of objects and of situations, how to behave in which situation. And this is what we did then. I had up to 20 PhD students in parallel, and so this was very interesting, this period. And we made quite a bit of progress. At that time people came to visit us from all over the world, from Japan, China, US of course also, all of Europe. We had the common project. PROMETHEUS was the European project. There were somewhat like, I don't know, 20 or 30 universities involved in this project,

and we switched around from one place to the other and had discussions, so this was very, very interesting.

Peter Asaro: Were you much influenced by the works of the cyberneticians, of Norbert Wiener or Warren McCulloch?

Ernst Dickmanns: Well, of course, and Norbert Wiener, he was essentially a control engineer and said that the idea of feedback, that you use information in order to derive your proper control in order to achieve a goal, that's the basic idea which has been followed right from the beginning. Yeah, I think that's very essential, but then of course when all the predictions have been done, what you can achieve with this very general approach people noted, especially in the computer science community, that it's very hard to represent background knowledge on this level, and that's when the separate computer science, artificial intelligence direction developed. And they started with more or less static frameworks, but this was-- and I've seen that right from the beginning-- this was not going to solve the problem. What we need is both parts. We need the one side and we need the other side. And what I could see was that what is necessary in order to get powerful systems is to combine <inaudible> static representations but not as static facts but as models for dynamic processes. And if you do have articulated bodies you do have certain laws how they can move, and all of this you have to represent. And once this is available you can expect what the other guy is going to do whom you observe if you've seen a situation like that. So this was very essential, but in addition to Wiener's work you need the static representation but not of course of static states but of what is the collection of models which you need in order to derive at an overall system, which can be goal-directed and based on momentaneous data you measure.

Peter Asaro: Do you think those models are learnable?

Ernst Dickmanns: Yes. They have...

Peter Asaro: They have to be.

Ernst Dickmanns: We learn in this way. I don't believe that there is some insertion from another world or so, so the biological systems discovered that they have some data about the world by all the senses and vision, and they tried to get a good interpretation. And from my point of view maybe <speaks German> to say it a little bit over-- oh. You get an understanding of why the notion of space and time is necessary and available in all animals right from the beginning if you look at the way how the inner ear works and how the eye works. And on one of these NATO Advanced Study Institutes I noted and I learned that in the eye you do have a lag, until the interpretation of the world is there, of several hundred milliseconds. In the inner ear you do get

the derived measurements of accelerations and turn rates with, say, a magnitude of millisecond time delay. And in order to get these two combined you have to have a notion of time. How is this going to be developed? And one of the nice things is if you go onto <inaudible> we say, like the Oktoberfest, there are certain carousels where you have separate motions around separate axes, and what happens there is that the visual impression doesn't fit the inner-ear impression, and then you get...

Peter Asaro: Dizzy.

Ernst Dickmanns: Dizzy, yes. And this is one of the actual inputs where I came to the conclusion the internal representation of time is due to these different time delays between the inertia part of the inner ear and the eyes. And if you get these two in conjunction and you get them consistent then you feel well. If it's not the case you get dizzy. And of course that's the very nice thing that the inertia sensors-- from the inertia sensors you get the impact of perturbations right on the acceleration basis, right on the lowest time level. If you want to interpret an acceleration by some perturbation from vision you first derive-- vision is on the second integration level, so you look at states. You have to have the first derivative to get speeds, and the second derivative is accelerations. So, again, the combination of both is very essential, and that's why we propose you need a 4D model in order to combine this. And I consider this as our main contribution to dynamic vision. That's why we call it dynamic vision.

Peter Asaro: One of the machines Heinz von Foerster built was a dynamic signal analyzer for sound, but he simultaneously did this multiple-derivative analysis of sound processing.

Ernst Dickmanns: Yes, if you have one sensor you are not in the same situation, and if you do have only sensor data and have to work with it you're also not in the situation, so the biological systems had to come up with some representation of the outside world, which allows you to combine the large time delay from the visual system with a very short time delay at least on your system. And what helps you also is that in vision relative motion between two bodies-- you can make a distinction. Is he moving random? And, again, here the inertial system helps you, because your own motion is being sensed here too, so you get another input, and you can make a better discrimination. And from my point of view this is the starting point for developing higher integrated systems with the notion of 3D space and time integrated.

Peter Asaro: That's great. In your systems do you ever try to extract that three-dimensional representation and save it as memory? Obviously over time spans you have this...

Ernst Dickmanns: You know, there's one dissertation, and that's maybe the most complex representation, where one of my students did a generic class of car models, and by combining

certain parameters-- he had about 20 parameters I think-- by selecting value ranges for these parameters he could make distinctions between a passenger car, a sedan and a combi and a van, a bus and, as I mentioned, 12 different types of cars. And we were able to show that this works with real systems. So models and the adjustment of model-- that's the essential part, generic models. So we do have models which are quite adaptable, and then you have to have the procedures available to adapt it to the corresponding situation. And we mostly worked with edges and corners, and we very soon came to the conclusion we have to have some area-based information like texture and color, but computer power wasn't there. That's why we couldn't start with it. But nowadays it's available, and one should do it. And once this is available I think we pretty soon achieve the performance level of humans. My guess it's one or two decades away, not more.

Peter Asaro: What do you think the biggest transformations in robotics have been during your career, apart from the increasing computing power?

Ernst Dickmanns: Well, of course I hope that a transition from quasi-static representations to spatial-temporal dynamic representations is one of the essential parts. On the other side-- no, I think all the basic laws have been known. Motion laws have been known, perspective projection has been known. Realizing that systems in silicon have to be because of the properties of the substrate-- in carbon you do have a reaction time in the rate of milliseconds. In silicon you do have ranges of nanoseconds. So this has been recognized from my feeling where money went for developing systems only by the late '80s, early '90s, so this was a very essential part. And now of course once this is going to be developed you can see that with progress in electronics also now you can make subsystems, which you combine with conventional computer systems. Maybe I should mention at this point when we did this I did have this cooperation in negative obstacle detection with the US Army. There was the task of detecting these ditches. In '92 or '91 I think we had a system available which was 30 liters like this, and that was the first system capable of doing full video frame, full video rate stereo interpretation like this. So we did some developments over here and we showed that it worked. It also worked in connection with the EMS vision system, and then they said "Yeah, well, maybe we should proceed," and they got a funding that the system was developed further. And one or two years later we had a Europe Card-sized 16 by 10 centimeters electronic card, something like this one here, which you put into the normal PC system, and we could show that we could do full 3D interpretation and detect these ditches with something just put into one of the four PCs we had onboard. We had four PCs with two dual processors, eight processors onboard, and in addition this unit which they had was somewhat like 80 billion operations per second, so it's tremendously powerful. And this development of course is going on, and you may have half a dozen or a dozen of this or similar systems in future simple, conventional computers. I recall when we started in the late '80s, I think it wasn't '90s, the first international discussion. The goal was one coffeepot, two-liter size for the entire computer system onboard of an autonomous vehicle, and we are pretty close to that.

Peter Asaro: What influence do you think the DARPA challenges had on autonomous driving systems and research and development?

Ernst Dickmanns: Well, I'm sorry to say, but I don't believe in the first one, because this was just vehicles pulled through the desert by GPS waypoints, and they didn't have to look for negative obstacles. All they had to do is drive along these GPS paths and avoid positive obstacles sticking out of the world. Yeah, this they did well. I think the robustness of the system was a very good demonstration, but with respect to recognition of the world it wasn't too much. And the second one, the urban challenge, if you look at it from a perception point of view it's also a little bit more complex but not the real task, because the entire information on the road lane markings, on the traffic signs available, on crossings, where to go, how to move from one direction into another on crossroads, everything was stored on a CD, which was put into the computer. And one of the focus areas, the talk-- what is it? Where the people exchange. One guy asked "We do have a vision system onboard, and if I notice that the lane marking is different from the one given on the CD, which one is valid?" The answer unfortunately was "The one on the disc." So, again, I would say with respect to perceiving traffic situations there's not so much progress. Good progress has been made in detecting obstacles sticking out of the ground and adjusting to obstacles and also to a relatively complex situation at crossings. But also with vehicles only-- and, again, the paths they drove were given by GPS, so this is a tour through the city by rather strict pre-described routes. It was impressive how they could do it, but there were some cars without any vision, and I think even one did finish without vision, so not too much of like visual perception what has been done there.

Peter Asaro: Do you think as a general strategy having these big competitions is a good way to try to advance research?

Ernst Dickmanns: Well, no, you have to look at the justification, and the justification was if we have to do-- how do you call this?-- mission support in an area which has been not covered where the Army has rule over the terrain, where they know the roads are good and you want to move support-- what is it?-- maintenance stuff from one location to the next. And this is the type of task which could be shown that can be done, because then you know that there are no negative obstacles and you need a relatively small perceptual capability. For this it was good, but if you look at the general task of perceiving a 3D world there's quite a way left to go. There's some progress, but it's not the general task which we would like to solve.

Peter Asaro: What were some of the other major breakthroughs in the last 20 or 30 years of robotics that you think are really influential?

Ernst Dickmanns: Well, with respect to what you can do with human-like articulated bodies I very much admire the work of Herzinger here in Germany and also in-- what is it?-- near Salt

Lake City there's another US institution. They also did some very nice developments. Yeah, this was really good. What else? The most impressing to me was the small vehicles on Mars. This is really something where I have the feeling that things are able to perform much better than engineers dare saying. So they predicted a lifetime of, what, six weeks or something like this, and it worked for four years or five years. It traveled I don't know how many kilometers, and it's really amazing what has been achieved there. So that's really very impressive. And I'm also aware of what's going on in, say, outer space with the sondes [probes] which are sent to other planets. There may be similar things, but this one here is really-- yeah, I think to me this is the most impressive one.

Peter Asaro: What else would you like to add? Any other thoughts or reflections on robotics and where it's going, where it's been?

Ernst Dickmanns: Well, I think the main part is that one should try to get a better understanding of robotics, that it's not going to be competitive with life development, that it's going to abolish life after some time. As you can see in some of the science fictions, one should really try to promote the idea that the cooperation between technical systems and biological systems is what we are looking for and that life also for humans and other biological systems may be improved quite a bit if we do the proper application of these technologies. So there's quite a bit that can be done in household care, in elder people care, and developments are going on in this direction. Maybe I should add that also. What impressed me quite a bit was the work which has been done in your country with respect to support in hospitals by robots. But then you see there are always some individuals that try to fool the system. I don't know whether you have heard about these marks for electronic positioning have been removed, and one of the robots tumbled down the staircase and things like this, so it's amazing, but apparently this is part of the nature of some humans. They would like to show that they can fool systems, biological ones and technical ones. And as long as you have to deal with these-- system vandalism I think it's called-- and you always have to be aware that those people tending to vandalism will be around, and this is one of the main difficulties for technical systems, because it will be almost impossible to predict which type of vandalism will be possible with these systems. And if the manufacturer is going to be made liable for a product even under these conditions progress will be rather slow.

Peter Asaro: Yeah, I think that's a very good point. What do you think the importance of modeling robotic systems after biological systems is? How important is it to really pay attention to how biological systems work?

Ernst Dickmanns: Well, I think the functional idea, look at the functions and how it is being solved, like instead of having hundreds and several hundreds of eyes or even tens of eyes like some animals even do have today, nature developed the head and neck and the eyes, so the function is important, and we can learn from that, but we shouldn't make the mistake to rebuild

similar systems on a completely different technological background. As I mentioned, the carbon systems and silicon-based systems are completely different, and you are not able to transfer the solution in carbon-based systems like our bodies onto silicon. This is different, and this has to be recognized. There have been big projects in your country and in my country also where they wanted to develop an electronic eye, because it's known that there are somewhat like 120 million light-sensitive sensors in the eye, but the number of communication lines going to where the vision process is being done in the back of the head-- it's only 1.2 million, so there's a factor of 100 compression. So that was the basic idea why people started looking at combining sensors with processors in the first stage in order to reduce data communication after that. But on silicon that's silly. There is no problem. Take all the information you want, you need to have and just send them, because bandwidth is so high. Don't fool your system around with requirements which are not due to technical necessities but due to some example based on a completely different hardware or wetware or whatever you call it.

Peter Asaro: In terms of robotics research in Europe and in Germany specifically, where do you see the major research centers being and...

Ernst Dickmanns: In robotics.

Peter Asaro: In robotics within Germany and Europe.

Ernst Dickmanns: Well, I think one of the best ones in Europe is Herzinger DLR in Oberpfaffenhofen, and there are other ones in France and good ones also in Sweden. I'm not so aware of what's going on in England, because we didn't have so many connections. They pretty soon dropped out of the cooperation in PROMETHEUS, so I'm not aware of what's going on there. Maybe there are some activities going on in Oxford, Cambridge maybe, but I'm not sure. We did have some cooperation early in autonomous vehicles, but they have been abandoned as far as I know in England. With respect to autonomous vehicles, Finland has been doing interesting work every now and again.

Peter Asaro: Which labs?

Ernst Dickmanns: Oh, it's close to Helsinki, Epso, Eso? Epso I think it's called. <inaudible> maybe also, but I'm not sure.

Peter Asaro: What were the labs in Sweden and France?

Ernst Dickmanns: In Sweden it's the-- what is it?-- KTH Konigliche Technische Hochschule in Stockholm, and in Lund -- no, what is it? Is it Lund? No. It's Linkoping. There's a university in Linkoping it is I think, yeah. And in France maybe in <inaudible>, and of course Toulouse, the-- oh, what's the name?

Peter Asaro: Well, there's the...

Ernst Dickmanns: No, that's industry. I thought of the research institution, and there's a large research institution which has been in Paris first and was then moved down to Toulouse, and we had quite a bit of cooperation or exchange with those. I'm sorry, but my memory with respect to names, it's not too good anymore, so I forgot the names. But you'll find if you look at Toulouse...

Peter Asaro: Yeah, okay.

Ernst Dickmanns: And some activity have been in the southern part, also in <inaudible> Institute, <inaudible>. But this is 15, 20 years ago, so I don't know what's going on right now.

Peter Asaro: I think that covers most everything I could think of, unless you have more to add.

Ernst Dickmanns: No. Maybe you should have a look at...

Peter Asaro: We can look at some of the videos.

Ernst Dickmanns: ...some of the videos, yeah.