

Flagstaff AZ
1:07:22
m4v

Tomas Lozano-Perez

An interview conducted by
Selma Šabanović
with
Matthew R. Francisco

August 28 2011

Q: So, we can start with you telling us your name and where you were born and when.

Tomas Lozano-Perez: I'm Tomas Lozano-Perez. I was born in Guantánamo Cuba, August 21, 1952.

Q: And where were you educated? And how did you come to the U.S.?

Tomas Lozano-Perez: We left Cuba when I was about ten, moved to Miami, lived there for a little bit, then moved to Puerto Rico, where I went to high school. After high school, I went to MIT. So, I stayed there for my bachelor's, master's and PhD.

Q: And you were always in the computer science department?

Tomas Lozano-Perez: Yes. Well, at MIT, it's electrical engineering, computer science. It had just changed right around the time I got there. I don't know exactly what date, but yeah, my degrees are mostly in computer engineering, computer science.

Q: And how did you get interested in computer science?

Tomas Lozano-Perez: Well, as a freshman, I thought that I was going to do electronics. I had no clue what that meant, really, but then I took a programming course and loved it and sort of stayed doing that. So, at MIT, they have a requirement to do an undergraduate thesis, and a guy I knew, one of my friends, was working at the AI lab at MIT. And so, my advisor, my undergraduate academic advisor was Patrick Winston, who was at the AI lab. So, I ended up doing a bachelor's thesis there, which ended up being in computer vision.

Q: And was it with Patrick Winston?

Tomas Lozano-Perez: Yeah. He was the supervisor of the bachelor's thesis. So, then I stayed and worked there for a year after my bachelor's, which is in '73, So, through '74. And we had a visitor for a year there, who was Hirochika Inoue, who spent a year at MIT. And he was working with a little robot that we had there, which was called Little Robot. It's a Cartesian arm that he used to doing some assemblies. And I was asked to help him with LISP. He didn't need much help, because he's a smart guy. So, but I got to talking to him. I got really interested in the robotic stuff. So, I went to graduate school, you know, after a year there and started working on – Winston asked me to work on a language for mechanical assembly. So, you know, Inoue had done his assembly using a very low level, you know, move here, move there, very simple kind of

language. And so, the question is: Could you do a high-level language for assembly? And at that time, there was this rash of problem-oriented languages or special-purpose, high-level languages for particular tasks.

In particular, there was something called PLANNER that had been developed at MIT. And so, I tried to start working on something like that and then realized that it kind of didn't make any sense, that that kind of a high-level control structures. The PLANNER had these backtracking control structures, and that didn't seem to be relevant to the robotics programming. So, basically I started doing basically geometric planning instead. Because I figured, like, even deciding where to put the fingers to pick up the objects was a hard thing. So, I started building little object models and trying to figure out how to do the grasping of the – and how to move them around. So, around that time in the early '70s, there had been a thesis at MIT by Terry Winograd, which had built this simulator robot that you could interact with in natural language, you could give it commands. So, it was called SHRDLU. And it had a little symbolic planner, you know, and it would pick up these blocks and so on. And in it there were two functions that depended on geometry. One was called FINDSPACE, and the other was called MAKESPACE. So, FINDSPACE is if you're going to put the block down on the table, where should you do it, where's a place that's free. And MAKESPACE was to move things out of the way to put something down. And they didn't know how to do it. They were doing it by picking random positions, which at that time was considered incredibly stupid. Now it's the state of the art, but at that time, random selection was considered very stupid. So, a number people had thought about algorithms for how to try to solve those problems. And so, I got interested in that. And I started – is this what you want?

Q: Uh-huh. I'm just taking notes.

Tomas Lozano-Perez: Okay. I started thinking about how to do the selection of grasping on these objects. I had done it by hand first. I did an assembly of aircraft, little model aircraft engine cylinder, which has a cylinder on a rod, and so on. And, I had programmed the robot to do it. And so, I was interested in how to figure out where to put the fingers when there were obstacles nearby, other objects and so, on. So, I started sort of trying to write programs to do this, and I hit upon the idea of using just the tip of the fingers as a reference point and then characterizing the set of points that were not accessible to the tip of the fingers due to the presence of other objects. So, I built these things that I called grasp sets, which were basically maps of the degrees of freedom and characterizing the places where you can put the fingers and so on. And then, I kind of ran out of time, and I did the design of how I would build a whole system, but the kind of hard-core piece was really on the mechanic grasping side.

Q: And what kind of robot were you using? Or was it mostly simulation?

Tomas Lozano-Perez: No. No. There was a real robot, but what I did is I did programming on a real robot. Then after that, I really did it in simulation. I did the computations geometrically. We had this robot, which was called Little Robot that was built by a guy called David Silver. And it was a Cartesian arm, which is an interesting set of degrees of freedom. There was a hand that could do Z and Theta, and then a base, which was the little table that did X and Y. So, it had four degrees of freedom but split. It was very cool. And it had four sensing fingers, which was not typical at the time. Inoue had used some early for sensing in Japan, but it really wasn't widespread. Well, robots weren't widespread, but this was. And they have a nice level interface to our computer, so, you could write LISP programs to run it, which was pretty cool back then. So, I had done a fairly elaborate assembly using that thing. And then I switched over to how do I plan for it? I never really closed the loop to the point of actually generating a plan for that robot. So, you know, at this point really it was a computational question: How do I pick the grasp? So, I got some way along that and wrote up a little paper, which appeared in IJCAI or something. And then I went off for another year in between. So, I had taken a year in between my bachelor's and my master's. So, I went off for a year between my master's and my PhD to IBM, T. J. Watson, there's a guy there, Peter Will, who was running the robotics group at T. J. Watson. And he'd accumulated an interesting bunch of people. And he used to visit MIT every once in a while. And I said, "I'm interested in kind of, you know, figuring out what IBM is like," because I thought I wanted to end up working in a research lab like that after words, and IBM was very interesting at the point. It had a monopoly, so it could afford a lab. You know, it's sort of like after the monopoly went, the lab went.

But anyway, so I went to T. J. Watson, and they had their project to build a language for assembly at that time they called AUTOPASS. So, I started working on that team. There's a guy called Mike Wesley, who unfortunately passed away, a British guy, very, very cool guy. He was working on a modeling system, a polyhedral modeling system. And I started working with him. And he said, "Oh, let's kind of try to build an example." And around that time, I ran across a thesis from Cal Tech by a guy called Dupa [sp?], who had done a motion planning system using what we would call kind of configuration space ideas at the time. And, you know, I said, okay, it turns out that this is exactly the idea that I had used for grasping, and he was using it for moving the arm around. I said, "Hmm, this kind of seems like a pretty general idea." So, we decided to kind of build a whole system that would try to do planning for little assemblies. So, it was the usual kind of, you know, without understanding it, Mike and I split the work. So, we said, "Okay, Mike, you write the A* algorithm, or the graph search algorithm, and I'll write something that computes the ground obstacles in the world and so on." Yeah, of course, that part was ten times bigger than the other one. But anyway, that's the way it works.

But we managed to do it by the end of the year. We had a prototype of something that, it had a bunch of typewriter parts that were placed in the environment and the set of obstacles. And we ended up using Styrofoam obstacles because the method that we ended up using was what we later called a visibility graph method, which was an idea that appeared originally in STRIPS, you know, that system. And but they hadn't known about ground obstacles in the way

that we were doing it. So, it ended up clipping the corners of the obstacles, so, you know, the Styrofoam got curvier over time. So, they had a Cartesian robot, too, which IBM marketed for a while. So, we put something together that would build these ground obstacles and then do visibility graph searches and do it over X, Y, Z, and a discreet set of rotations So, it could move those things around. And, you know, we demonstrated it. So, that was, you know, it was pretty cool, I thought. And then I got into a little fight with IBM because they decided that, oh, this is pretty cool. Why don't we make it confidential? So, we had a little disagreement around that point. I ended up talking to the Vice President of Research IBM about this disagreement. So, we reached a compromise that I could publish it as long as I didn't mention that it was ever implemented on a robot.

<laughs> So, okay. So, I wrote this paper, which later appeared in Communications of the ACM, jointly with Mike Wesley. So, it looked like a very abstract algorithm. It was good. The paper became very well known, people in computational geometry, which was kind of being born around that time, got interested in it. And we wrote an IBM tech report, and a guy called John Reif read it, and then he wrote this paper on the complexity of the piano movers' problem. So, it ended up having, I think, a fair bit of impact, people that didn't know it. And it was, you know, not mentioned that it was on a robot. It looked like it's an abstract algorithm, which probably was a good idea. It probably would have been a bad idea to kind of present it as a systems paper, as a robotics systems paper. It was really kind of an abstract algorithm, which was good. Later, Wesley managed to sneak in a picture of the IBM, of the path of the robot doing the simulation in one of the papers that he wrote on describing the geometric modeling system that they had built. So, he kind of snuck it past the lawyers. <laughs> Mike was good. He's a very kind of staid British guy with white hair, but he was funny. Anyway.

So, I went back to MIT to finish my PhD, and I wrote this paper and so on. And I showed it to Berthold Horn, one of the professors there, and he gave me a great hint that, in fact, that these volumes that I was building appeared in classical geometry. They're called Minkowski sums. This was actually Minkowski difference. And so, he pointed me at that. That was very useful. We were graduate students at the time then. He said he'd been reading physics, and he said, "Well, you know that this space that you're using is called configuration space in physics. So, what you're really doing is building obstacles in configuration spaces." I said, "Oh, cool, a brand, a name. That's perfect." So, then what I did in my PhD thesis was really kind of build on this idea, generalize the idea and implement some ways of computing them for bigger problems and so on. So, you know, published a couple of papers on that. And that became what people called, you know, became known as the configuration space approach, the idea of characterizing the constraints and the robot's degrees of freedom due to the presence of obstacles and then, you know, doing planning. And over the next, you know, few years, we showed that it could be used as the basis of all kinds of planning, for grasping and so on. And Matt and Russ Taylor wrote a paper on how to use it for thinking about uncertainty using an idea called pre-image backchaining.

And so, computing these pre-images into configuration space and thinking about the sensing that happened while the motion was going on, which was kind of my favorite piece of work that we ever did at the time. Because, you know, it was pretty clear and certainly it was pretty important, but it wasn't clear how to think about it. Now, it turns out that those ideas were also, present in the controls community, but we didn't really know about that. There's dynamic programming and belief space with some idea that was kind of being developed in the controls community. We weren't aware of it at the time, but I'd say, you know, it's something that I really liked. So, I didn't say, but my PhD thesis ended up being supervised by Berthold Horn because he was the more mathematically inclined faculty member at MIT. So, he ended up being the supervisor for my thesis, and Matt's thesis. And Berthold and I co-supervised Matt's thesis at the end, but at that time at MIT, there was this constraint that they weren't supposed – the government had told them not to support robotics because they go through, you know, eh, you know one day, they decided no, that's not useful work.

Q: When did that happen?

Tomas Lozano-Perez: This was, I think we had – it was ONR. And it was around the time that I was a graduate student, so sometime before 1980, in the late '70s or something. So, the AI lab had this money that came ultimately from DARPA the ONR. And somebody had decided that, you know, that was not a high priority thing. They were funded to doing AI. So, there were a couple of us that were already around, and we stayed, you know. They managed to kind of keep us going. So, I mean, as I remember, there were really only four graduate students in robotics at that time around the AI lab. The one who graduated first was Marc Raibert and then there was John Hollerbach and Matt and I. So, we all stayed in robotics and are still around. <laughs> But anyway. So, Berthold ended up supervising at least Raibert, Matt and I in one way or another. Are there any questions that come up out of that? Or I could keep going?

M1: I was curious about you already talked about the piston from the airplane and then also, the typewriter parts. Was there any way you kind of came up with those objects?

Tomas Lozano-Perez: The typewriter parts at IBM were natural. IBM had the Selectric typewriter, which was – I don't know if you ever looked at a Selectric typewriter. It was the biggest mechanical kludge in the universe. I mean, the fact that that worked was completely amazing. Because you'd see all these parts, the incredibly gnarly shapes. I used them later for when we were doing some computer vision stuff. I used those parts because they were just awesome. They were just, you know, mechanical engineering gone wild over the years adjusting to fit around corners and stuff. So, you know, when you thought about assembling stuff at IBM, it was, you know, the Selectric came to mind. Wesley was thinking about the covers of computers, of, you know, the big machines. And they had cabinets and stuff, and so he was

actually working on CAD for those things. But they were too big. So, we had a robot about yay big, and you know, the things you could pick, the Selectric typewriter parts were the right size.

Similarly, with the little robot, we had to look for something. And I think it was probably Winston and Horn who had come up with the idea of assembling a model aircraft engine. I think that's where it came from. So, the size of the piston, which is about that big, was about the right size for the fingers of the robot, perhaps a little small. I have that video some place. It's from 1976, and it was, you know, it was hard, because you know, they were kind of a little too small for the fingers. It would have been better to use something bigger. Even though I had built an assembly out of some ball bearings and a shaft, and those were a little bigger and probably would have been a better choice. I was, I mean, you know, just fresh. I didn't take graduate school; I didn't know better. But anyway, so, you know, there wasn't like a whole – robotics was very small then. You practically knew everybody. So there wasn't that much experience.

Q: And was that all funded by the ONR?

Tomas Lozano-Perez: We were. So, at the AI lab, we never even thought about funding, so, the whole lab was funded by one agency. It was ARP, DARPA, or I guess at that time it was ARPA, then there was ARPA then ARPA, then DARPA then DARPA, you know. And the guy who was in charge of ARPA had this funding model of basically funding labs. So, AI lab and MIT and Stanford and so on were basically, there was the one grant, and they funded everything. So, you know, they wrote a grant every three to five years or something, which was, you know, three hundred pages of stuff and saying all the things we were going to do and all the things that we've done. And we were there, so everything in the lab was paid basically from one place. And I mean, it was a radically different model for funding than we have now. It's much less constrained. So, if something came up that was interesting, you could follow it. It wasn't like you had to say that you were going to do this: And on the third day I shall have a breakthrough and – that wasn't something I had to say. So, you could just do it. So, you know, Winston was the head of the lab at that time, and he kind of liked the idea of robotics. There had been robotics going on there since Minsky started the lab. He was interested in robotics, and they had done some robotics stuff. So, Winston kind of kept that going. There were a few of us around doing it. It wasn't, you know, the major part of the activity as a laboratory. As I said, there were only, you know, three or four of us doing it. There were some more people doing computer vision and just, you know, basic AI stuff.

Q: How many people were around the lab at the time?

Tomas Lozano-Perez: It's hard to say. I don't know. We had about one floor at Tech Square, So, I don't know, fifty? It's hard to say. I mean, later it grew bigger. I don't remember. You know, there were maybe five or six faculty and graduate students and then some staff. So, I

mean, there were, you know, what we called hackers back then, as a compliment, not as a curse. And you know, they had built things like the first, one of the first chess programs that was rated, you know, it was taught by Rich Greenblatt, who was a guy working there. And, you know, it was a cool place. A lot of stuff was happening. And robotics was one of the things that we were doing. But there were only a few of us who were really working on that.

Q: Did you collaborate with people who were outside of the robotics group?

Tomas Lozano-Perez: Well, we talked to random people in terms of collaborations. I mean, the person I worked most closely with was Matt. We were graduate students together and, you know, we hung around. So, our supervisor, Horn, was like a vision guy. And he appreciated math, but he's A), a vision guy, and B), somewhat reticent, so, he was, you know, it's not like he, "Hey, come here guys." No. We barely ever saw him. So, we were kind of our own advisors. So, you know, Matt and I, a little bit more under sometimes with John, we'd talk about things and, you know, what's going on, and we'd try to read papers and talk about them. But it was pretty much, you know, kind of, we were going. And you'd talk to the other graduate students and, you know, other people working on planning and language and all kinds of stuff. So, you ended up absorbing a fair bit of the general culture from talking to random people. It wasn't as – because it was small, you know, you talked to everyone. It wasn't like when you have a big group, you only talk to the, you know, the people in your group, you know. And they believed in moving your office every year, so you'd talk to a different set of people kind of thing.

So, you know, I knew about vision. I knew about language. I spent a year doing linguistics as a graduate student just because it seemed like fun. And you know, Chomsky is around and seemed exciting, so, I did linguistics for a year. Then Chomsky woke up and said, "No. That's all wrong. Everything I said before was wrong. It's now this way." Oh, well, I'll just go back to robotics. <laughs> So, I think it was much less specialized, the education, and kind of environment. So, I think that's good. I think students today tend to be rather more specialized. But, you know, in terms of collaboration, I had worked with the people at IBM. Late in my time there, Russ Taylor arrived. He went to work there, and Russ was working on some of those same ideas when he was at Stanford. So, in terms of academic robotics at that time, Stanford was probably the biggest place under, you know, Bernie Roth and those guys there. The Internet started; it was working on there. And so, there was a fair bit of activity. Lou Paul and Bruce Shimano and all those guys. So, and Russ Taylor had been working there. So, I had kind of talked to him, and you know, I had talked to, worked with the people at IBM. But when I came back to MIT, aside from kind of random connections, mostly the guy I talked to was Matt.

Q: And you mentioned, I think, IJCAI, and Communications of the ACM. What other kinds of conferences or journals were up there for robotics at the time?

Tomas Lozano-Perez: There wasn't much. So, in terms of robotics itself, there was this conference, ISIR or something like that, International Symposium on Industrial Robots? Which was kind of a mish-mosh of stuff. A lot of it focused on industrial – it's called industrial robots. So, I mean, there had been kind of a giant disconnect between kind of industrial robotics and academic robotics from the AI kind of perspective. So, industrial robotics was about doing stuff fast and cheap, you know, <inaudible> and so on. So, it really wasn't a place where you really could present that kind of stuff. Any kind of more high-level AI-ish flavored stuff you'd go to an AI type conference. But it really was a problem.

So, Communications of the ACM was kind of one of the leading journals at that time. Now it's a kind of rag, you know, it's for the news and an occasional random article. But at the time, it was like the leading journal in computer science. So, we were very lucky to get our paper published there, and it was the cover article, and, you know, they had to prepare cover art and stuff like that. So, you know, it was very cool. And I ended up publishing that paper in my thesis, and the IEEE transactions in systems and cybernetics, which was kind of a hodgepodge of stuff. And the paper that, you know, that has configuration space on the title was in the IEEE transactions on computers, of all places. I mean, if you said, "What were the leading computer science journals?" ACM, you know, and IEEE computers were. So, I was trying to say, "Hey, you know, let's try to get wide circulation." So, I aimed for that. There weren't very many robotics venues in terms of publication, which was one of the other reasons that it later, you know, a few of us kind of talked about beginning a journal in robotics and, you know, IJRR eventually came out of that set of discussions. Lou Paul and Mike Brady took the lead on that and started that journal. But I mean, there wasn't – I think as far as I remember, ISIR wasn't really that kind of conference. I don't think I ever went to one of those. It really wasn't. So, I aimed at whatever random other journals were around.

Q: At some point, somebody told us that also, MIT, CMU and Stanford have these reports that they would circulate.

Tomas Lozano-Perez: Yeah. Yeah. So, those were not a rough read or anything like that. So, in AI and computer science in general at that time, each of the major institutions had a technical report list. So, IBM had one, and so, MIT in the AI lab, there was something called AI Memos. So, you would, you know, write AI memos and then try to publish them some place. That's usually how <inaudible>. The memos were in circulation earlier. So, for example, the delays in publication were enormous, so my paper on configuration space, whatever, took three years to appear. It had three reviews. Each one took a year. So, the AI memo version had been out, you know, long before. So, that was something that people looked at. So, Stanford had a series of technical reports. And they were physical, right. This was before the web, so there was a real use for a library. So, the library at MIT would subscribe to the, you know, CMU and IBM and Stanford and so on, their publications, these technical reports, and then we would use that as a, you know, means of trying to keep up.

Because there weren't really regular conferences that you could go to, at least, you know, the IJCAIs and later AAI's and so on for a while, were that until, you know, regular conferences in robotics started. Like, ICRA started in '83, something like that, no, '85? When was it? I don't know. It was in the early '80s. I was the program chair for the second one. Lou Paul had been the program chair for the first one. I did the second one. My experience with the IEEE wasn't the best, so I didn't do that again. But and then, out of the AI lab, we started the ISRR, and they came out as books. At that time, MIT Press was interested in publishing books like that, because there wasn't a lot in robotics, and they started publishing the IJR journal. So, yeah, so the technical report series played a very important role in all of computer science at the time. But it was especially true in AI, I think. So, you know, the standard thing you do is you – at the AI lab, we had the things called the AI memos, which were fairly polished, and then there's something called working papers. And previously they had a thing called vision flashes because they came out of the vision group there. So, you know, stuff that was less polished we'd throw in there. So, and then you'd publish a thesis that, you know, student theses and so on, so you'd get a nice little, you know, bound thing, and people got a hold of those. So, yeah.

Q: And how was the communication like in the community, then? Was it, you know, in these conferences that were not so much robotics oriented people but still kind of congregate in robotic circles, or not?

Tomas Lozano-Perez: So, we're talking here kind of before 1980. I graduated in 1980. I don't remember any place where you talked to people, except by maybe visiting their university. Now, I'm sure that faculty at that point and so on maybe had other places that they went and maybe talked to each other, that they probably had traveled more, but I don't remember really talking to people. I remember one time visiting Stanford and missing Russ Taylor but talking to Bruce Shimano kind of thing. And CMU I don't think I ever visited, although they were probably a little less active in robotics at that time. And in the Boston area, it was Draper Labs who was pretty active in robotics, this guy there called Dan Whitney, who had been very influential. So, we would see their reports and theses and so on. And so, we heard about that. And we'd look at the Stanford ones, and it was pretty informal. I really wasn't – I mean, we didn't have the Internet kind of thing. I mean, we had local e-mail, but not really – but there was at the time e-mail between the different, you know, the ARPANET was in place, So, you could send e-mail to Stanford and other places, but if you didn't know somebody already, you know, you didn't kind of cold call them. So, you know, you kind of knew who was doing what because of their, you know, the technical reports coming out and stuff like that. And there weren't that many people doing anything, so you could keep track of what was going on much more easily than you can now.

Q: When do you feel that changed? Because you mentioned ICRA and the—

Tomas Lozano-Perez: ICRA really – yeah, I mean, ICRA and ISRR were really the, I mean, from my point of view – maybe other people have other memories – but those are the times in which I remember actually talking to other people. So ISRR was by invitation at the – when it was created and the first few of us were great. I loved them and the – everybody was there and it was a lot of fun and you'd present stuff and there was a lot of discussion. I remember the first one I presented that worked that Matt Mason and Russ Taylor and I were doing on this backchaining stuff or uncertainty and I remember Jean-Claude Latombe getting upset, "Oh, this is ridiculous. It's a waste of time. We shouldn't be working on this" and Lou Paul came to my defense and it was an interesting discussion there. People went, "Oh, wow." <laughs> It was interesting. So there's a lot of sometimes heated discussion about things and it was – that was a lot of fun, but – so that was a change starting in the early '80s when – and when ICRA was put together, it was a fairly conference. The IEEE particularly wanted large conferences. They like large conferences because they make some money out of them so – which is – was the basis of some of our disagreement, but – and so then you started seeing lots of people and talking to lots of people being at conferences and being always – almost never listening to the papers but outside talking to lots of people. So that was a dramatically different experience from my – from earlier on and before, in the late '70s, when – being graduate students basically. You knew the people around you and that's it.

Q: Do you have a feeling for what factors led to this explosion of interest in robotics, the amount of people, the attention given to it?

Tomas Lozano-Perez: Well, I suspect that it was the idea that industrial robots were going to be important. I think the funders had this idea that somehow the industrial robots would be important and that it would be good to have academic research. That later led to a bust basically in which – because there really was relatively little connection between much of the work in robotics and what went on in factories, which was very stereotypical repeated kind of stuff. So there wasn't a great value placed on versatility and so a lot of us were working on that and that didn't really connect. So I think later on funding kind of dematerialized and I basically left the field around 1990 partly feeling that some things were stagnant and funding had gotten pretty hard at that time so I think there was the kind of optimism about industrial robots, and I think industrial robots in fact did pay off but the connection to academic research was not as strong as it could be. There are some parts of it that – some of the low-level control and kinematics and that kind of thing so that paid off in my opinion – but much of the work that was kind of computer science based there was really no connection. So I think that that led to a kind of drop. The thing that kicked it back up later was mobile robots but that was not for a while.

Q: So what kinds of things were you working on when you started at MIT as a professor?

Tomas Lozano-Perez: So I was interested in the – there were two kinds of things, the motion planning kinds of things that – I had a phenomenal group of students. My first group of students was out of this world so, Matt – we were graduate students together and I happened to be co-supervisor in his thesis at the end 'cause I graduated a couple years earlier. He had taken a couple years off but we were more – you could just as well say he supervised my thesis as I supervised his but then I – when I first started out of the blue I had this office of students, John Canny who went to Berkeley and is an amazing guy. There was Donald who is now at Duke, was at Cornell for a while, Mike Erdmann who is on the faculty at CMU, and this guy called Van-Duc Nguyen who didn't continue to a Ph.D., did a master's, but his master's thesis stuff is still referenced. If you look at any work on force-closure grasping and things like that, look at any paper on that now that are around and the odds are you'd see V.D. Nguyen as one of the references still today every time. He's amazing too so – and then after that Nancy Pollard who is on the faculty at CMU.

So I had a really phenomenal group of students and they were working on basically motion planning kinds of things, so Canny ended up doing his Ph.D. on the first provably singly exponential algorithm for motion planning. He's completely self-taught so he did it all himself and Bruce worked on – first on six degree of freedom motion planning for free-floating bodies and then he did a thesis on error correction and recovery, some interesting ideas, generalized configuration spaces and so on. Erdmann ended up working on basically – one was the modeling of friction in configuration space and then later he worked on randomized strategies, ended up doing some really cool stuff there. Nancy worked on grasping with three-fingered hand. Ken Salisbury was around and we had his hand and so she looked at algorithms for – planning for three-fingered hands so that kind of stuff. In parallel I was also working with a guy called Eric Grimson who was a vision guy. He did some work on object recognition using contact sensing and vision so we developed this thing called the Interpretation Tree, which is, was an approach – a constraint-based action approach to doing vision and object recognition so that became reasonably well known so we did that. I've always been kind of hands on so I always liked, actually did programming myself. I still do. So Eric and I actually did the coding for that project and we did that so – well, the students were doing a great job with the motion planning stuff. I did some myself also. I came up with an algorithm for doing planning for joint angles and so I – that's one that I did by myself, but I was also doing this stuff with object recognition so I did that through the '80s, papers in that area.

We did a system towards the end of that time called HANDEY, which was an attempt to put together a lot of the stuff that we had done in motion planning because a lot of that had been kind of independent little pieces so HANDEY could plan the whole thing, picking up an object. It would do localization using laser striping and then figure out where to grasp the object and then move it someplace and re-grasp it if it needed to and so on so it did kind of the whole thing so it was nice to bring it all together, but then at that point I started feeling like I wasn't – that we were kind of stuck. Partly the computers weren't very fast and we seemed to be reaching our limits and funding was getting hard so at that point somebody came to talk to me about modeling

molecules in configuration space and I said, "Whoa. That's a wild idea." So then I spent the next kind of, almost 15 years working on – doing discovery so I got, I worked at – with a startup and trying to use machine learning to do discovery. That was before the whole genome thing.

So proteins and peptides were the thing of interest so I got interested in that. Then it was good because it got me into machine learning, which I really didn't know at that time. This was also before the machine learning explosion so I got interested in that. And we came up with this problem called a multiple instance learning problem in machine learning, which arose out of my stupidity of – I had this representation for these little molecules and so – but the representation wasn't unique because – and so it was, "Okay. Somebody must know how to solve this problem," so contacted Tom Dietterich who is kind of reads every paper ever written on machine learning, still does I think. I said, "Tom, really how do we solve this?" [and he said,] "I don't know. That's a new one on me," so that was fun so we worked on that for a while, then worked on protein structure and stuff like that, then made the mistake of getting involved in academic administration for a while, hated it, <laughs> so then decided to take a sabbatical and said, "That robotics stuff looks like it might be fun again." So that was around – oh, I don't know – six years ago or something.

Q: What was the fun stuff that you were seeing then?

Tomas Lozano-Perez: Well – so one of the things that happened while I was away kind of thing was – had started to happen was the randomized motion planning kinds of things so that seemed to have done extremely well so, and of course the mobile robot stuff had really kind of done a great job so that was – it seemed like people were starting to become interested in using robots and for something and, besides industrial applications so might be interested once again in trying to do intelligence kinds of things, and some of the ideas – I had always thought that dealing with uncertainty was really kind of the critical thing because that's really where sensing and action come together and so there was a lot of work on using probability methods and so on. So I started talking to Leslie Kaelbling and she's one of the pioneers in POMDPs, which is an approach to doing – dealing with uncertainty and planning, and so we talked about a bunch of things. We also did some vision kinds of things but I said, "Why don't we use POMDPs for doing grasping kinds of stuff?" and she said, "No, POMDPs are too hard" <laughs> I remember, said, "I stopped working that. I'm doing machine learning now. They're too hard." And I think it might be interesting so we started working on that and it did prove out to be interesting so we're still working on that and we have a couple papers that are kind of outgrowths of that in this conference so that's been very productive.

So now we have a joint group and we supervise students together and stuff so that's worked out great so – yeah. So it felt like at least manipulation, which is the part of robotics that I've always been interested in, which is almost dead, so if you looked up robotics in the '70s and

'80s it was manipulation. When people talked of robot it was a robot arm that had been – Hans Moravec at Stanford had built his cart and stuff like that but you didn't do mobile robots then. Then for while there was only mobile robots. There were a few people who hung in there like Oussama Khatib and so on but doing manipulation was basically impossible and – but then it seemed like – back a few years ago it seemed like yeah, it might be time to start doing that again. I think I called that one right. All of a sudden there's been a real rebirth of interest in manipulation.

Q: What are some of the things that you think changed to make it reborn? Is it technological development? Is it conceptual –?

Tomas Lozano-Perez: So computers have, are just, oh, they're a lot faster. That's a transforming issue so problems that used to be- people used to worry about, learning the servo loop fast enough and doing geometric computations and so on. With computers today it's sort of "whoof" so really – that really transformed the things that were interesting. I think that was probably the biggest thing. I think conceptually the idea of using probabilistic representations and some of these ideas like – that came out of operations research and controls really do give you a kind of- a nice set of tools for thinking about uncertainty and so on so people had been using some of those for mobile navigation, the whole SLAM thing but they really hadn't moved into the manipulation end at all. The manipulation – the problems are somewhat different so there seemed to be kind of low potential for doing that so that's what we ended up doing.

Q: What are some of the challenges that you're working on now?

Tomas Lozano-Perez: Well, the thing that Leslie and I are working on is trying to find a kind of unifying way to think about planning, acting and perception so one of the weaknesses – conceptual weaknesses in much of robotics throughout the years has been that these things get compartmentalized partly because of education reasons. So there's the people who do controls and there's the people who do vision and there's the people who do SLAM, and usually what you end up with is solutions that don't play together very well because you make assumptions that are really incompatible with some of the other solutions. So what you really need is some way of really kind of thinking about the things together like AI planning kind of separated from robotics back from – when STRIPS was and they made a set of assumptions which happened to be incidental I think to some of the things that STRIPS – the way STRIPS was doing things and they've stuck with them. And now they consider it anathema to question those assumptions but they're just incompatible with the real world so that you can't have fully symbolic descriptions of the real world. That's not practical and – but they're stuck in that little local minimum and you look at the people who do POMDPs and they mostly – they say, "Oh, well, we have to find the optimal solution" and you can't find optimal solutions for any interesting problem so you have to kind of fake it so you want to kind of break away from that. And perception once again people

were – you make all kinds of assumptions about what it is that you want from perception that it's not clear how it integrates in a task and what you really need. So we've been trying to build, kind of conceptualize a system that – where perception and action and planning and execution are all integrated and trying to see if you can get some leverage on understanding what some of the assumptions were.

So for example, we ended up kind of using some of the ideas from high-level planning in robotics but turning it around a bit and we still get into arguments with any – when – we – ever we talk to anybody there because, "Oh, no, you can't do it that way." Why? And even here when we talked about it today some people says, "You can't do it that way. Everybody does it the other way. Right?" <laughs> So I think it – there are these conceptual silos that develop and you're starting to see some cross pollination but it's good to – so we're trying to find something that kind of cuts across so recently we were working with our colleague, Russ Tedrake, and a post doc, Rob Platt. We've been trying to understand how the control people think about the world and trying to – in fact, the, we had – Leslie and I have been doing a lot of discrete POMDP stuff and the controls is really continuous POMDPs and so how do you put them together and try to understand that so we've been meeting every week trying to kind of figure out how to cut across that boundary so that's been very productive so. Anyway, so I guess I'm old enough that I can – subject to finding a little bit of funding here and there I can kind of work on stuff that somebody starting their career might think a little risky.

So for example, the kind of thing that Leslie and I have been doing some of it has just – she and I have been doing it. We do the programming. There's no student involved because the field has gotten conservative so the review process tends to weed out any idea that's too – that we don't do it that way. So we are older. We can afford to have our papers rejected. Younger people get all upset about that but – so – anyway, so we're trying to do something which is cross cutting that really brings together the different aspects of the field and it'll take 20 years to get there but it's fun.

Q: You mentioned earlier that industrial application seems to be this thing that constrains robotics in certain ways and now people are talking even today about a much broader spectrum of applications whether it's in the world or with robots or NSF co-robots thing and things like that. Do those kinds of concerns play into your work at all or –?

Tomas Lozano-Perez: You need money. Right? <laughs> So yeah. The problem with industrial robots wasn't that we felt that was the only thing. It's just that that's what people thought they wanted to fund so you had to kind of at least move in that direction. If you had written a proposal in the '80s saying that you were going to do robots in the home that sounded a little strange so – There had been a little bit of work on medical robotics at Stanford at the hospital there but – so that wasn't mainstream and it wasn't funded. Pretty much the funding

came from DARPA and some industrial places and so it wasn't – that wasn't available. Now there's – it's a little more open. The only bright light for, not only – one bright light for funding was NSF in the early days. They were pretty open minded but little, tiny bits of money so you could fund one student but – so now with the increase, we are on the ascendant again, it seems interesting and there seems to be a broad range of applications and the military's very interested so that makes it easier to get funding to do a variety of things so that's good. If you can't get funding to do stuff, then the realities of the situation are you can't – so –

Q: And it seems like these kinds of questions feed into things that you're interested in like uncertainty. Being in the home is a much more –

Tomas Lozano-Perez: Right. I was brought up as an artificial intelligence researcher and that's always what interested me is getting robots to be intelligent and so the thing that's interesting is what do you perceive to be the critical blocking factors when you do robotics? So the early days of AI so – what people perceived as a blocking factor was things like planning and reasoning and they kind of viewed the robot as an output device, as kind of a printer on the 3-D world, but then when you start actually trying to do something with a robot then you realize that whoa, it's very hard to do anything and where everything breaks is because your sensors weren't good enough and your controls weren't good enough. And so you go here and the object's there and you missed or you go and you hit it and everything breaks. So once you start doing things in the real world then you immediately realize that uncertainty and – which brings a combination of sensing and acting is really a big problem. And one way to try to improve on that is by doing better sensing, right, and – or better control, and in fact those are kind of dominant lines in robotics so that my little take on it was that yeah, okay, you do better sensing, you do better control but ultimately you also have to reason about how you should act in the presence of uncertainty. So that was my little niche there for a while and roughly it still is. Thinking about that is kind of how do you – it's practical intelligence in some sense. It's not just abstractly reasoning about chess but it's producing actions that are effective even in the presence of uncertainty. That's the kind of area though where I got interested in.

Q: What do you see as the current obstacles in the direction of where the field is going?

Tomas Lozano-Perez: There are still hardware issues and that will always be but – so the hands in particular are problematic. There are some really good ones around but they're very expensive for example so we, a couple years ago, I guess a year ago we got this PR2, which is a Willow Garage thing. They gave out 11 of them or something and it was – it's an awesome device. It literally arrived in a big box and we pushed the box in, opened it up, pulled it out, and it works and it had sensing and control and everything but wow, <laughs> this thing is great but the hand is really stupid. It can only open eight centimeters and so it's a limiting factor of this really cool-otherwise really very cool device. Anyway, so I think hands are still a problem and cost. It still

costs hundreds of thousands of dollars to get one of those robots so I think as people get more interested in it somebody eventually will come down with something that's lower priced and more people will be able to work on it. That's what happened with mobile robots.

Right, initially, they were really very expensive and then I remember when Rod first started at MIT. I was his supervisor so he was a post doc for one year and so he was working with me on some manipulation stuff and- but he really always wanted to do mobile robots. And I remember one time this kid showed up. He was a 16-year-old who had built his mobile base, eventually started a company called Real World Interface, which actually sold those circular robots to everyone. This kid was amazing and so I remember he showed up at the AI lab and my office was close to the door and he said, "Hey, I built this thing. Would anybody be interested?" and I went "Ooh, yeah." <laughs> He had built it in his basement and it was cheap. You could produce lots of them. It really changed the face of the – anybody could work on mobile robots at that point. You had in your computer one of these little bases and you're off going so – whereas if you do manipulation it's a huge investment.

Q: What was the kid's name?

Tomas Lozano-Perez: I'm having trouble with that. Rod would know I think –

Q: No problem.

Tomas Lozano-Perez: Yeah. I can't remember his name. He turned out to be important. He really did this company and was eventually bought out by somebody at – it was very cool. So Rod used that base for some of the early work he did in mobile robots, yeah, because a lot of us were gifted mechanical designers. For a while Ken Salisbury worked with us at the AI lab and he would always produce these beautiful pieces of hardware and it was sort of like, "Oh, that's amazing," but once everybody else produced these kludges that didn't work very well, no. So I think we still have a little bit of a hardware problem. The PR2 has been kind of a great experience. It kind of works and the students can work on it and it's – it – so we need something more like that, maybe a little cheaper, and I think that that will help bring more people into the field.

Q: Four hundred thousand dollars isn't exactly affordable.

Tomas Lozano-Perez: No. No, I – but I think that's a temporary state of affairs because there's another – Mecha – have you – Mecha Robots. They produce a really cool platform now. They also are charging \$300,000 or something like that but that means we know how to do it so somebody's going to figure out some way of making them cheaper and so I think – just in terms

of community I think that the two issues I see are: One is the hardware at some level, and then the perception just got a lot easier because of the Kinect. Right. That's going to be – make a big difference. The other is I think education. I don't think we're going to make progress by being compartmentalized the way we are now. We have to figure out how to get out of that. I don't – it's hard just, you have to learn a fair bit but I think if we could decide and agree that that was an issue we could construct a curriculum and there is – there are stuff around that you could know and so that people are aware of at least, they don't have to be an expert in everything but they should be a little more aware of the broader picture so I think that – that seems like an odd thing to pick on but it – I kind of feel like there really is an issue there. There's a lot of stuff known but not by many people that know all of the pieces or even some subset of the pieces. I'm always learning. I've always loved – I have, every time, and it's good to kind of learn about what the other people are doing and I think it's kind of important just to formulate the problem in a way that makes sense so –

Q: You mentioned education. If you could give advice to young students, people who are interested in getting into robotics, what would it be?

Tomas Lozano-Perez: Go west, young man? Oh, no. That was – no. Go east. Well, I think spend some time learning about all the aspects. There is this thing that's happened in recent years with the advent of all these competitions in robotics, the FIRST and all those things like that, is that students come in thinking that robots is about hacking. You slap a bunch of things together and you can just go out there and do it so the idea that there are kind of principles and that's not something that really gets highlighted there very much. So it's really, there is a scientific – set of scientific principles that are good to understand and it's good to learn broadly before tackling it so that's probably the – at some point you want to actually do something but you also want to make sure that you're working on something that makes sense.

Q: You mentioned early on that Inoue was in the lab – did you have other connections with him or any other Japanese labs or other labs around the world or labs in the U.S.?

Tomas Lozano-Perez: I kept in touch with Inoue for some years on a personal level but I – not really much connections with the Japanese. I did have more connections later with LAAS in France so Georges Giralt used to visit. He used to take this pilgrimage throughout all the places that did robotics and so he kind of knew everyone. So I had met him and we got along. We used to talk in Spanish. He was Catalan but in the French side and so we had that in common and so he always invited me to go over and visit and so I did. I spent some time there at LAAS and met <inaudible> who was there and a number of other people and we kept in touch and I had a number of French students in the lab over time so I think that that was productive. They later became one of the prominent centers in motion planning. I think I might have had a little influence in that. I don't know. I used to keep in touch with them and talk so they really kind of

took to that stuff and did very well with it so – yeah. So I think from international places – I had some friends in Barcelona also but they were really working on other things but I spent there – some time there. In fact, I did my sabbatical six years ago, whatever, eight years ago at – there – Barcelona but not a lot. I used to travel around and visit labs but I didn't have a really close connection with them at the time. Anything else?

Q: No. I think those are most of the questions. Is there something you'd like to add?

Tomas Lozano-Perez: <laughs> It's 30 years' worth of stuff. Okay? <laughs>

Q: I know.

Tomas Lozano-Perez: But no. I'm fine. I think you got some idea.