

RAS_2024_Interviews_02_Kris-Hauser_v1.mov
2:00-2:36

Kris Hauser

IEEE RAS Interviews

Q1: Tell us about your work experience.

Kris Hauser: Right after my postdoc with Ken Goldberg, I also worked with Honda Research Institute as a postdoctoral scholar. I had done an internship with them when I was a grad student. And I continued on some of my work from that internship and expanded upon that. I then started a faculty position at Indiana University. And I spent five years there in the Computing and Informatics Department, I believe. And after that, I moved to Duke University and I was part of the Electrical and Computer Engineering Department and the Mechanical Engineering Department there. And then five years ago in 2019, I started at University of Illinois in Computer Science.

Q1: Do you have any advice for students looking for advisors?

Kris Hauser: As I advise grad students currently, it's very important to pick an advisor that matches you not just in terms of research interests, but also advising style. And the presumed advisor that I had at Stanford didn't quite match the style that I had hoped for. And so I ended up looking for new advisors and then I found my actual advisor, Jean-Claude L'Etonne, who was much more of the style that I preferred. And the skills that I learned in the graphics discipline, whereas working on motion generation and simulation, very, very quickly transferred over into the field of motion planning.

Q1: What was your first experience with robotics?

Kris Hauser: Well, so I wouldn't say that there was any, I'm not one of these people who had a childhood dream to make robots. And in fact, I sort of saw these movie robots and C-3PO and R2-D2 and so on, and Terminator and things like that. And I thought it was a little bit silly until I started actually getting into the research of robotics. And so falling into robotics is maybe not the usual path of a roboticist, but I found that it helped me solve problems. And that's what really enthused me about the field of robotics and that's just an enormous source of interesting problems that kept me going. And kept me enthusiastic about trying to solve tough problems.

Q1: What was the first robotics project you worked on?

Kris Hauser: I actually have a colleague who I met him in grad school. He's my academic older brother. His name's Tim Brettel. He's now also a professor at University of Illinois. He had a project, he was an avid rock climber, and he ended up working with JPL on a rock climbing robot. So he built this robot called the Lemur. It was a climbing robot that had four limbs and it could climb up very steep surfaces. And my first project was to try to make the planning for that

much faster. So the previous techniques that were used were somewhat brute force. They had this computational complexity that exploded because of the number of possible options about where you might place your hands and feet. And my goal was to actually use machine learning to try to make that accelerated. So right now the topic of integrating machine learning and motion planning is a very hot topic. But back 20 years ago I had started working on that.

Q1: How would you describe the relationship between probability and cybernetics?

Kris Hauser: I actually had to learn a lot of probability in grad school that I had not known before. So the main difference between computer graphics and robotics is that in robotics you have to deal with the real world, which has a lot of uncertainties. Whereas in computer graphics you get to control that world. And so I had kind of, I'd learned a lot of skills in terms of modeling, spatial relationships and motions and simulation from my computer graphics background. But then switching to robotics, I had to start thinking more deeply about probability and uncertainty that actually exists in the real world. But in terms of predicting the future, we had back then taken more of a deterministic kind of approach where we modeled the world very, very carefully. And ended up using that to predict the future steps of the robot. And the planner would be able to reason deterministically over those steps. We hoped that then that would be kind of all the uncertainties and unexpected events that happened in the real world would then be happened in control. We just kind of pushed, kicked the can down the road a little bit. But they are extremely important. And you can see that nowadays with these integrated robot systems becoming more and more important that dealing with that from the head on is becoming a more important topic.

Q1: Describe the evolution of robotics projects that you've worked on.

Kris Hauser: The first projects that I worked on were deterministic planning problems that had really no basis in reality. Let's just be honest. Back then we were very lucky to even get our hands on a robot. You're considered a superstar if you got anything working on a robot in those days. Nowadays you pretty much see robots in every single lab and I try to make sure that all of my students are getting experience on real robots by the time they graduate. But back then it was a little bit different. So near the end of my PhD I started actually integrating with perception. I worked with some folks who were experts in perceptions. So Radu Rusu who invented the point cloud library. He was the main developer. He was a collaborator on a project that we called Leading Flatland where we tried to get the Rex robot, which is a six legged robot, to navigate up and down stairs in multi-level environments. So we did some mapping on board and then also did the planning. That was my component. But I didn't have much of an exposure to feedback control. Although I took many classes with Osama Khattib back at Stanford, I never really had the chance to do really serious feedback control work until I went to Ken Goldberg's lab.

And so there I learned about model predictive control and we worked on some needle steering work. And there we had to deal with the realities of inserting needles through soft

tissue, which are that you can't really predict how they're going to behave. They had a really interesting kind of technique. He and other collaborators, they had these flexible steerable needles that could go through tissue and they curved as you went through tissue. And we worked on some of the first planners and controllers that were able to steer in 3D. Some of the foundational work had been done in 2D by Ron Alterovitz, who's now at University of North Carolina. But then me and some other postdoc in Ken's lab started working out how to control these in 3D.

So that was very interesting having to work with feedback and work with real time planning in that setting. And so that exposure got me interested more in systems and how robot systems were built. And fortunately, around the same time when I first started my faculty job, we started seeing things like the Microsoft Connect come out, which was a huge, huge change in the way that we viewed robotics. The fact that we could now work with three dimensional information was just a game changer. And so that got me more interested in how we could use perception and real time feedback to control these robots. So I started to become less of a motion planning researcher and more of a robot systems researcher as I started to kind of mature through my early days as an assistant professor.

Q1: When exactly was it difficult to get your hands on robotics?

Kris Hauser: Probably before, let's say about 2005 or so, we could buy these big industrial robots. There's also some folks working on surgical robots as well, but they're very expensive and they're also scary. So when I first started at Indiana University, I ended up having a couple of large industrial robot arms and we had to build these cages for these robots so that nobody could step in without triggering the safety stops. And it was just not a really fun time to work with these robots. Basically, my students did not like working with them at all. And it wasn't until the 3D sensors came out and then the cobot revolution came about where we could start to put these robots embedded in the lab on a tabletop and feel safe around them. So once the UR5 came out, that was a big change of things and it made people more and more comfortable with working with robots. And they were also much, much cheaper too. So approximately a third of the cost of a standard industrial robot.

Q1: Which research questions were you interested in, and how have they evolved?

Kris Hauser: This was not an advisable thing to do as an assistant professor, but I basically changed the tack of my research quite a bit from my grad school days up until around tenure time. So I started out being a very algorithmically focused motion planning researcher. I had proven various things with regards to complexity of various algorithms, which is what I was most known for during those days. But I started becoming very interested in actual robot systems and how they could be integrated into reality. And this idea that we could actually make robots a reality started becoming more and more prevalent through these advances of cheaper sensors and cheaper actuators, drones starting becoming reliable and so on. And one of the

problems that I took on back then was the issue of human robot interaction in terms of the planning of robot motion. And so my career award was based off of this idea of cooperative motion planning, which is to have humans designate goals for robots. And then the robots would try to achieve those goals through the automation of the motion control of the robot, motion itself.

And so that was very interesting because I started working with human subject studies and I started working with real time systems. And safety, of course, is a really big issue with human robot interaction. But then also understanding the behavior of the human was a big problem and that required trying to advance this kind of idea of AI for human robot interaction. And one of the first things that, well, we did back in, I believe it was 2012 or 2013, is we had the first symposium at the AAAI conference for AI for human robot interaction. So I was one of the co-organizers of that first symposium and it was very interesting because we saw that people had to develop artificial intelligence to understand people. And especially for human robot interaction researchers, they need to be able to have access to this sort of thing. And it was a bit of a change to what had been currently this sort of state of human-robot interaction research, which was mostly human subject studies or retrospective surveys and so on of humans. Whereas we saw this big kind of, the technical side of research wanted to get into human-robot interaction but didn't know what kind of tools they needed. And so this kind of idea of using AI for learning about humans to then put into robots was a very interesting idea. And I thought that that was a really cool kind of offshoot of what I had done there.

Another direction that I have continued to work in is integration of robot systems. So in the kind of classical robotics sense plan act framework, there's this kind of sensing component and perception. There's the acting component, which is the planning, which is planning and control. And then there's acting, which is the actual motion of the robot. And what was also very interesting there is learning. So learning is a part that's not usually talked about with sense plan act, but it's really important for all of these different pieces to work well together. And so looking at how we could integrate all of these different pieces together got me to think about the borders between these different components. So planners need to have input from perception. Perception needs to produce outputs that planners can use. Planners need to have suggesters or things that if you want to pick up an object, where do you pick it up? These are called affordances. So where does affordance prediction come into this? And then there's control. So how do you then create, take these plans and then control them accurately? So then I saw that there are all these missing pieces in the boundaries between the kind of classical sense piece, planning piece and acting piece. And I've continued to start to, continue to work on these problems until today. And the problems are always changing, but they're always there.

Q1: What technical challenges have you faced in your career?

Kris Hauser: Yeah, and in terms of technical challenges, I'd say that one of the biggest hurdles, which I briefly touched upon, was trying to connect pieces together. We have a lot of great theory that has come about since my grad school days and before, which they're really taught, they're still taught today. They're kind of the foundations of robotics in many ways. But when

we try to put together some kind of system to use those components, we very quickly find that they just don't work as advertised or that they take enormous amounts of time to get to work as we hope for them to work. So customization or kind of basically encouraging that kind of slot A fitting into slot B nicely has been a major challenge throughout my entire career. And that is, it's not because we can't do it once. It's because that if we keep on sticking with the same paradigms, then people are gonna still, every time that we had to do this, it was like a student would have to spend several, several months just trying to get pieces to work together. And if we don't teach in our classes and in our fundamental material how these parts should be fit together, then everyone's going to face the same problems every time they try to put something together. So a big goal of mine in my current research is to try to make the boundaries between these different components broken down or more compatible with one another.

As an example, the kind of prevailing paradigm in motion planning from 2000 onwards has been kinematic motion planning, which is to create a path. These paths, they create a path for the robot to avoid obstacles, but they don't. So first of all, motion planning researchers didn't look at anything but minimizing path length, if anything, as a cost. And they also didn't take into account the dynamics of the robot while trying to generate these paths. And so you end up getting these very jerky, nasty looking paths. And it was people in motion planning were kind of wondering, well, how come people aren't using motion planners?

The real reason was because the output of these planners was not really compatible with what robots should do, which is to generate smooth, nice looking motion. Robots have inertia, so they can't just change direction on a dime. And so it took a very, very long time for people to start to realize that you needed to have different types of models for generating these paths. And so I worked on some models for using acceleration bounds in motion planning. And so that became more widely adopted back around the 2010s or so. And it ended up becoming like a default setting in some other package called Open Rave. And people would come up to me and be like, oh, you're the person who came up with a parabolic path smoother package! We use it in Open Rave. So that was like an interesting kind of challenge. Breaking down the barriers lead to interesting kind of research questions and also opportunities.

I guess I'll also say something about perception. So back in the old days, we didn't have perception embedded into our problems that we were solving in the motion planning community. And until the connect came about, and then we started being able to build these three-dimensional models. And back in the past, we were working with CAD models. We had just used some kind of CAD designer would create some sort of model of the environment that an industrial robot would work in. And we'd have a model of the object that we were going to try to manipulate. And so we would just plan using those as our input. When we start to get a motion planning person together with a perception person and get them to try to work together, the motion planning person would be like, "Okay, we'll give me a CAD model." And the perception person would say, "Well, I don't have a CAD model. Here's a point cloud." It's all noisy, it's partial, it looks ugly. Now what are you going to do with it? "Well, I can't grab it because I don't have a model of it," as the motion planner person would say. So that whole field of working with imperfect knowledge, imperfect and partial perception data and putting those into planners also bore a lot of fruit.

So a lot of the more recent work on GRAS planning has used machine learning for these types of tasks. So you can see that learning came in as a solution to a big problem in real robots, which is that you'd have to somehow guess how you might have to pick up this object, even though you didn't have perfect information. And so that kind of classical paradigm of assuming full knowledge of everything just kind of broke down. And now we routinely use learning when doing GRAS planning, because that's the only thing that can really deal with these types of situations.

Q1: Tell us about the tools that have opened up new possibilities in your research.

Kris Hauser: Yeah, so the Connect was a game changer because we all of a sudden did not have to rely on just pure camera vision to infer things about the three dimensional world. And instead we just got the data right there. And computer vision back then also was a bit spotty, I'd say. So we couldn't trust vision inputs very well. There was stereo vision, but it's still got problems. And the Connect just kind of gave us the data right at our fingertips. And that was just amazing. We could just calibrate how far things were and just get that information right off the bat and put those into our models for planning, GRAS planning, etc. And it really made things a lot, lot faster.

The other thing that happened around then, which I don't think it really affected my work so much is it changed the way that we thought about robots, but the robot operating system - ROS. It is essentially, it's a standard now in academic labs and also many industry labs. It was developed during that time by Willow Garage. And I had some folks that I knew at Willow Garage and collaborated quite closely with some of them. But it brought the knowledge that integrating systems was important and having your work be used by others was important to the robotics field. And I think that was a really big change. It is like a paradigm shift in how we viewed robotics. And you could see that very quickly after a few years of Willow Garage being around. They were around for about five years. Companies shudder because of various reasons, but funding being the most important one. But those folks and the know-how and the kind of ethos of building robots that worked, that used all these different components, really changed things in that all these startups started coming together with different types of robots, drones, ground robots, etc. that could make use of all the tools that others had built. And so this really brought out this idea of like an open source ecosystem.

And that really did change how I thought about what robotics was. So it wasn't just an academic discipline in that we needed to share our ideas through papers with other academics, but that we also had to put products out so that others could use them. And it made me realize that, as a person working on algorithms, I was really doing a service to the world. I wasn't to others, actually not directly impacting the world, but providing a service to other engineers that would then be able to use my tools and then take them out and put them into practice on robots. So I started becoming more and more interested in open source software contributions. And so I make it a routine practice in my lab to take the work that we've done and convert them into usable software tools. So I've done that quite a bit. I've been working on one software package since 2012 or so. It's called KLAMPT, K-L-A-M-P-T. I continue to work on it. And the other

part about that is that I looked at ROS, and I saw that it was making--I actually have jumped around between many different departments. I was in computing, then I went to engineering. Now I'm in computing and engineering. What I do find is that no, offense to engineers out there, but the average engineer is not a good software engineer. And software development is something that even good computer science students have trouble with. So learning about new pieces of code, learning about how a new architecture works, how a new robot works, took so much time.

And so I also found that tools like Python, the Python language is very easy to pick up. They've basically switched over from when I was in undergrad learning Scheme to--or C, over to basically teaching Python for introductory courses. And it's cognitively a lot simpler than more kind of professional grade programming languages. But it leads people to be able to pick things up much, much more quickly. And so that was also a really big change to my kind of thinking about how to work with the world, which is to provide tools that are easily accessible to novices. Because, let's be frank, all of us when we're picking up a new software package, we're all novices. And so I tried to make things basically Python-oriented completely in my lab, try to make sure that code is very well documented. And so that kind of just ethos shifts also kind of hit me as well.

You also mentioned learning. And so this is more in the last five or ten years or so. But learning has--I will give the field of machine learning credit where it's due. It is very, very easy now to pick up learning. So you can probably spend six months to a year in an undergraduate curriculum learning how to use machine-learning tools. And you can do a very credible job of using machine-learning in computer vision tasks or any kind of other task out there. And so it's made that kind of accessibility that this field has pushed, especially for some of these tools in Python that you can basically download tutorials and just tweak them a little bit and get them to generally work okay, is a massive change from the way that we did things before, where we had to have an expert in each discipline or each topic, each component come in and try to work together and fit things together. Now we can get an undergrad and get a reasonable job of doing object detection. And that's just amazing. So it's really changed the kind of focus that we put our effort into. And it makes perception not quite easy, but it makes getting off the ground a lot easier than it used to be.

Q1: Who has funded your research, and how have you developed those relationships?

Kris Hauser: Over the years, I've been funded by a huge variety of sources. Most of my funding comes from National Science Foundation. And that's kind of the scientific granddaddy of university funding in the US. It is a really nice source of funding in that most of the judgment of who gets funded is basically your peers. It's great panels that spend a lot of focused effort and very solid review of each proposal. I tend to say that not every good proposal gets funded, but the proposals that do get funded are good. So I've enjoyed that kind of relationship. And there it's basically relationships with your peers that gets you funding there. I've also gotten funding from National Institutes of Health. I've been dabbling a bit with various forms of robot-assisted medicine. We've been working over the past many years on eye surgery and eye exams because

robots have a really nice precision compared to human's hands. So that's been a very fruitful collaboration that I've had with some folks at Duke University.

Early on in my career, I was involved in the DARPA Robotics Challenge. So I got some DARPA funding from that. Those competitions, those very large competitions at DARPA seem to be somewhat over at this point in time. The latest one was the subterranean challenge that I did not participate in, but I guess we'll see if they decide to revive some of these large multi-tens of millions of dollar funding programs at DARPA again. And I've also gotten some industry funding from Amazon and also Mitsubishi Electric Research Laboratory.

Q1: How has funding influenced your work over the years?

Kris Hauser: It's been very interesting how funding has opened my eyes. So usually as a graduate student, you are focused on some research topic and they just have this kind of laser view of that topic and you don't see the forest for the trees. You can get locked into thinking that your little corner of the world is the most important thing and everyone else should care about it. But funding really makes you actually think a little bit more honestly with a bit of a better perspective about how your work might influence society. I view it as a good way of aligning social need to researchers. A lot of researchers don't look at social need as the primary driver of what they do. So I saw that I was getting funding for these system projects, and that really opened my eyes to this being an opportunity for robotics to make a difference. The fact that people were interested in larger integrative projects meant that was where the action was. That's where the fun stuff happens and that's also where the impactful stuff happens.

And so that really changed my view on how research was done, but I also wanted to keep a foot in basic research and so I kind of swung very far to systems research funding as an assistant professor and then, once I got tenure I was like, wait a minute, I can do a bit of both. I can do, I can have one foot back in the basic research game, one foot in the kind of applied, integrative systems game and that's where I've kind of been ever since. But it does get you to think about different problems. It does get you to think about how do you get stuff to work together, teams of people working together. I've found that for at least for the last ten years of my life, one of the biggest sources of strain and excitement has been working with other people and working with large teams of people. Everyone has different views of how a project should run. Everyone has a different view of what's exciting to them, how to spend their time and finding good collaborators, finding people whose style you can work with, whose vision aligns with yours, has been a very interesting and eye-opening bit of the kind of professor game for me.

Q1: Is there anything you'd like to pursue that you haven't yet?

Kris Hauser: I guess one thing that I always see is that learning research is, has just exploded and if I had focused on machine learning back in grad school, I probably would've had a lot more, let's say citations, maybe some kind of impact in the academic world. So some part of me

is like, “wow, I see these assistant professors who are getting like 5,000 citations or something by the time they go up for tenure,” and I’m like, “man, that’s just like a whole level more than us folks working on “classical” robotics.” To some extent, I also want to say that, at some point, it’s a bit more...I’d say the learning approach to robotics is where the hype is at. It may not be where all the impact will turn out. There’s a lot of impact in terms of getting people excited about this idea, of getting the robots to actually work. It’s where I think my soul lies in getting the robot to work. Not just generating papers or getting citations.

And so, I’m much more of this engineer at heart. I really feel like my problems that I like, the problems that I pick are driven by the actual challenges, and trying to put together some kind of system to do something. To large extent, it’s not about, let’s say, getting the learning to work. That’s sometimes the easy part. It’s oftentimes other questions that motivate me. So I can say that, although I look with a bit of jealousy at those people, I also think that I probably wouldn’t have been able to look at myself in the morning the same way. I view a bit of pride in the types of things that I’ve done over my career.

Q1: What is your most important contribution to robotics?

Kris Hauser: In terms of important, I’d say there’s a bit of a difference in terms of what I view, again, getting into citations and so on. My favorite paper is not my most cited paper. I think this is really important for people, for students to realize. The paper that you put your soul into, it’s your baby, but may not be the one that makes the biggest impact in the world in terms of the citations and so on. I’m very proud of the work that I’ve done with, it’s called multimodal motion planning. This is something that came out of my PhD work. Basically, it’s the study of how motion plans that are engaged in multiple steps of some kind of plans, such as manipulating something, grabbing it, then moving it, then grabbing something again, or like a leg robot to be able to climb some kind of terrain, taking multiple steps. I proved some deep theoretical results on those types of planning problems, and how you could solve them using sampling-based planning techniques. And I thought that work was very deep, and in terms of pride in the work, I view that as being some of the work that I am most proud of.

I’m also proud of some of the work that I did with what’s called the minimum constraint removal problem, also a theoretical motion problem, which I don’t think has ever been used in any kind of practical way, but it’s really influenced some of these algorithmic researchers in the theoretical complexity theory. It’s a theoretical computer science, and they’re picking up on this problem and studying it more deeply than I had. So, in terms of things that I’m proud of, I am proud of those. In terms of real world-impact, though, there’s a lot of systems work that doesn’t get quite as much maybe publicity or impact in terms of other researchers picking up on these topics.

The work that we’ve done with, let’s say, eye surgery robots, so guiding needles to perform corneal surgery. I feel that is having a lot of potential impact in terms of human health. Some of the work that we did with picking and packing as part of this Amazon picking challenge, we started working on packing and that has a lot of potential impact on the

efficiencies that could be put into place in terms of like Amazon shipping boxes that don't just contain a bunch of air and a little thing inside of them. Those types of problems that have impact have potential impact on applications, I also find to be very exciting and they do motivate me quite a bit, whereas they seem as just engineering work to some people, and maybe don't sound as appealing from an academic impact point of view, but I think in terms of real world impact, I am very proud of those as well.

Q1: How do you anticipate the future of robotics looking in the next decade or two?

Kris Hauser: That's a very good question. In the last five or ten years, the idea of big data has been, of course, a very influential idea throughout machine learning as well as robotics, computer vision. These large language models are really great at what they do. The reason why they work is because they are able to have a ton of data generated by humans that lies on the internet. These types of data sources just don't exist in robotics. We end up having very different sources of embodiment of the robot. The robot body is different from a human body. The sensors available on the robot are different. Rather than thinking about big data as being the future of robotics, I think we solve a lot of problems with big data, but the sticky problems, which is why big data has not really had as much of an impact in robotics as it had in other fields, is that we just won't - we can't and won't have cheap data available on these different sources of modalities that robots will need to use. Tactile data doesn't exist on the internet. Body positioning data doesn't exist on the internet. That's proprioception. We're going to need to have some developments in not just big data learning but small data learning.

This is a topic that a subset of the machine-learning community is very focused on this, using zero shot or few shot learning to try to solve problems with much, much smaller data sets than before. Trying to transfer data sets from, knowledge from one data set to another operating regime. Also, it's going to be very very important. The difficult thing is that the key theoretical problems are not quite solved yet from the machine learning perspective. I think this is going to be a really important area for many years to come. My group's research has been starting to look into this field for the last few years or so and finding that the tools are getting better, but they're still not there. You can't just hand wave away the problem by saying, "Well we'll gather more data." It's just not possible for a robot to gather thousands of hours of data in your own home without actually being in your own home. It's really just a very important problem for all of roboticists. I think it's under explored and under appreciated in the field right now.