Los Angeles, CA
1:04:05
m4v

Gurav Sukhatme

An interview conducted by
Peter Asaro

November 11 2010

Q: So, I just want to start by asking you a bit about where you were born, where you grew up, where you first began your studies.

Gaurav Sukhatme: So, should I just sort of free-associate and talk and – so, I was born and brought up in India, where I spent the first 21 years of my life. I came to USC after finishing my undergraduate degree at IIT in Bombay, in Mumbai. I came to USC as a graduate student in 1991 and did my master's and doctoral degrees here at USC. And subsequently I joined the faculty at USC, and now I'm a professor at USC now of computer science and electrical engineering, and I run a laboratory at USC called the Robotic Embedded Systems Laboratory. And so, in some ways I'm unusual in the sense that I've spent my graduate career, as well as my professional career, at one university, at USC.

Q: What did you study at IIT?

Gaurav Sukhatme: At IIT I majored in computer science and engineering, and so, roughly speaking, it corresponds to what at an American university would be a blend of computer science and computer engineering, but it's called one major, computer science and engineering.

Q: And when did you first become interested in robotics?

Gaurav Sukhatme: I had a casual interest in robotics growing up, but it was a mix of interests in computing and in robotics and in technology in general. So I can't say as a teenager I was necessarily more passionate about robots than I was about other facets of technology, but I was quite passionate about computer science. My interest in robotics really goes back to my early days at IIT Bombay, where I did some amount of reading on developments in artificial intelligence, particularly in physically embodied artificial intelligence, AKA robotics. And so it was roughly as a freshman or a sophomore, somewhere in that period, where through reading I became aware that there was this very, very interesting field, about which I knew very little, but I was an avid reader and I was also an avid reader of fiction, which had fired my imagination. But it was roughly at that time as a 19- or 20-year-old that I really began to sort of read what people had started to do in the field in terms of research. And so, when I came to USC as a graduate student in 1991, I had this vague, ill-formed idea in my mind that I would definitely specialize in artificial intelligence, in AI, and more precisely in physically embodied AI, so I knew I wanted to be the kind of computer scientist who worked on physical things, on robots. I didn't have a much deeper sense than that, but I had sort of had that, and then I slowly sort of built on that interest, but I would say my early interest was roughly when I was late teens or undergrad.

Q: Are there any specific books or researchers of robots that stand out that – early influences?

Gaurav Sukhatme: Yeah. Probably some of the earliest research robots I read about when I was 19 or 20 years old as a freshman or sophomore, there were several. I'll name a few. One that really captured my imagination – a series of robots that captured my imagination were the robots built by Grey Walter in England in the – I'm guessing the late '50s, though I can't be sure of that anymore. And these were the...

Q: Late '40s.

Gaurav Sukhatme: Late '40s they were. And these were the Grey Walter tortoise robots, right? And I saw an old BBC movie about them, which is very, very nice. I still have it on – and so if you'd like it, I would love to share that with you. It's an outstanding clip of what these machines did in the pre-digital computer age. But what fired my imagination about them was that they were marvels of computing, in a way. They were marvels to me of physically embodied computing. They were not necessarily marvelous mechanical creatures. They were ingenious, certainly, mechanically, but to me they sort of fired my imagination as a way of thinking about physically embodying computation, and that's what they meant to me. I was very, very happy to learn later that that's also what they meant to a few other people, but at the time I thought that was sort of clever about them. The other robot that really made an impression on me when I first became aware of these things in the late '80s was Shakey, which was the robot built at SRI in the 1960s. And that sort of captured my imagination, because it sort of showed me another side of robotics, as it were, and that was a very different kind of robot than the Grey Walter tortoise robots, right? It had a planner on it, and the robot executed plans, and it was able to change its mind, as it were, as a result of deliberation and not as a result of reacting to stimuli. And that really sort of was a different kind of robot, but it also was very, very impressive. And so these were really sort of two early ones that really captured my imagination, and they had really been around a long time when I was an undergrad, so I'm reading about them in the late 1980s, by which time one of them's already 40 years old and the other one is 20 years old, right? But they still represented to me sort of significant achievements in robotics. I was fascinated in a general sort of way also by robotic machines in automation, but probably it was the sort of mobile robots that captured my imagination more, because they held the promise of working, if you will, in unstructured environments, right, that they weren't necessarily working in pre-engineered setups, and Shakey was, to some extent, but certainly the Grey Walter robots were not. And so there was a promise there of robots functioning in environments that were not completely sort of adapted to their use, and that was something that I found very interesting. I didn't have enough of a background then to understand a lot of the underpinnings and the theoretical work that went into some of these things, but in terms of sort of firing my imagination, they were really quite remarkable.

Q: And did you know you were going to study with George Bekey when you came to USC, or did you find that out later?

Gaurav Sukhatme: I found that out soon after arriving at USC, so it was not something that was a done deal, so to speak, when I arrived at USC. I took a sequence of AI and robotics courses at USC when I arrived, and notable among them was an introduction to robotics class taught by George Bekey, and I believe he taught that in 1991, late '91 or early '92. I've forgotten. And it may've been one of the last times that class was offered by him. He was already very senior and very busy, but he would often teach this introductory course. And it was my first formal course in robotics, and it was soon after that that I began to work in his laboratory, or concurrently towards the end of that course is when I began to work in his laboratory. So, it was shortly after I arrived. I also took Michael Arbib's course on brain theory and artificial intelligence, which I enjoyed immensely, and Michael, in fact, was one of the people who advised me to talk to George. And so, it was part of my plan to talk to George anyway, and I already signed up for his course, so that was a sort of a happy coincidence. And I also took a very, very interesting AI course, which was not about robotics but was a fundamental AI course taught by Paul Rosenbloom. And it was a very strong intro to AI course, and so the confluence of these three courses in robotics and AI was sort of my beginning to graduate school. I took a bunch of other classes, but these three were up there. And that's sort of how I began.

Q: What was the first real robotics project you worked on?

Gaurav Sukhatme: I worked on two projects when I first joined George Bekey's lab in the early '90s. I worked on a project on robot hands, actually, on brassbane. I worked on a project to design techniques or algorithm for robot hands to grasp objects and a planning system that went along with that. And I also helped out in a small way on a larger project in his lab on autonomous robot helicopters. I was just one of many who helped on that other project, so the robot hands project was one in which I worked in a serious way. It was sort of my main thing, as it were, and I helped out on this other project. And as is the way, very often, in graduate school, I ended up writing my Ph.D. thesis on neither of these. I ended up writing my Ph.D. thesis on something completely different. So, after spending a year or two working on these, I maintained an active interest in both, but I moved away into yet a third project, which ended up becoming my Ph.D. thesis.

Q: What was that?

Gaurav Sukhatme: That project was a very interesting idea. It was a project on mobile robotics, and it came into being roughly around 1993 or 1994, when I was about midway during my graduate-student stay at USC, and it was a project funded by the Jet Propulsion Lab, by NASA, on asking the following question: If a small mobile robot was to perform various tasks of navigation and exploration in an environment, the kinds of robots that already were being planned by NASA to go to Mars, say, what would be a good way of asking how good is such a robot? How would you benchmark such a robot? Would it be possible to in some times create a science of asking how such robots are – how their performance is measured, if you will? And so the project sort of started off that way, and I devised, if you will, techniques to compare across various robot designs and various kinds of tasks in exploration tasks sort of inspired by Mars exploration. And I worked on a sort of family of techniques for that, and I ended up writing a thesis on that, actually, and it gave me – I was fundamentally fascinated by mobile robotics, and I found my way into it by this avenue, essentially. And so that was my most significant project as a graduate student.

Q: And what is your work focused on now?

Gaurav Sukhatme: So, let's see. So, my lab, called the Robotic Embedded Systems Laboratory here at USC, does several different things but primarily works on problems in estimation for robotics. And so, estimation is the science of making sense of sensory data, is to endow robots with the ability to, if you will, integrate noisy sources of information together into a meaningful sort of way, into something meaningful, informally. Formally it means something a lot more precise than that, and it's an old science. It's been around more than a couple of centuries, in some way or the other, and interestingly enough, almost all of us who have had a high school education have solved estimation problems quite formally, though we may not understand that, and, in fact, all of us solve them informally all of the time. Somebody tells you something and somebody tells you something else and you have to make a decision, and you sort of weigh information and trade off one versus the other, understanding that neither of your two sources are perfect, and you arrive at some sort of a conclusion based on disparate sources of information that are, in technical terms, noisy, right? We do this all the time. And if you took laboratory, if you took any kind of lab physics as a high school student, then of course you took a lot of data in the lab and then you fitted lines to it, and that's a process of estimation. Now, in robotics this is a very significant thing, because robots, sophisticated robots, tend to have a variety of sensors on them, and these sensors don't quite measure what – well, for starters, they're not all as accurate as you'd like them to be, and as another important consideration, they all measure somewhat different things, and no one sensor give you everything you want. So the science of combining these things together in interesting ways that are meaningful and actionable by the robot is really the estimation problem. Now, this is a whole big, general field. Within this field, my lab currently focuses on problems in one particular domain. We work very extensively on marine and aquatic robots that apply these techniques to exploring large-scale water bodies, rivers, oceans, lakes, and a lot of this work is very interdisciplinary. We do it in conjunction with colleagues in oceanography and colleagues in marine biology and people who model the oceans, and we work, if you will, in environmental robotics, right, but the technical contributions my lab makes are in estimation techniques particularly for robots that float or swim or maneuver in water in some way or the other like this robot boat that you see here. That's sort of one area. Another area that we work on is on estimation problems for fusing particularly visual information such as comes from cameras with information that comes from inertial sensors such as IMUs, gyroscopes and accelerometers. And these are devices that're typically found on robots like this or also increasingly on things like your cell phone, which have cameras and IMUs, and there're interesting information fusion and estimation problems on how you take visual data and inertial data and fuse them together in a meaningful way. And so for robots like this and even cell phones, which somewhat subversively we pretend are robots, we work on problems of fusing information together. And so the broad theme is information fusion and estimation, but we have a very strong focus on aquatic and marine robots and then some projects also on personal robots and doing some work on small embedded systems like cell phones.

Q: So, what do you think are the major problems confronting estimation and data fusion?

Gaurav Sukhatme: So, these estimation and data-fusion problems, we understand a lot more about them than we used to, but doing them in the large for robots that have to sort of do this over either long periods of time or over large spatial extent is still a relatively challenging thing. We've made a lot of progress, but there are definitely some outstanding issues when you try and scale them up. Taking them outside the lab is a significant challenge, and my group in particular is very interested and works very actively on the interplay between the algorithms needed for doing good estimation and the kinds of models that're available when one takes robots outdoors. So, I'll give you an example. What you don't see behind me is one of our underwater robots, which is currently patrolling the Pacific Ocean, and it's out on a two-week mission, where it's going to spend a very significant portion of the two weeks exploring the near coastal ocean near Los Angeles. And it's executing a variant of a particular planning algorithm, which in turn uses estimation techniques to gauge fairly accurately the robot's trajectory as it's moving through the water. And these algorithms in themselves would be fairly easy, but they operate under very severe constraints, and the constraints are that some of them are very sort of practical constraints. They operate on a small computer in an underwater robot with no other access to other computing. There are some power constraints and so on. But these algorithms, in some interesting way, have to be able to integrate information about what is known about the physics of the ocean, and it turns out that working with ocean modelers, we're able to infuse some of what they do into better estimation techniques for where our robots are. The science of doing this involves very large-scale estimation problems that need to be solved very quickly and in a timely way, so that the robot can more finely engage where it is. And so, if we were working in a perfectly still swimming pool, we wouldn't deal with this, right? The water would be, for all practical purposes, still, right? Even if it was a very big pool, it wouldn't really matter, but that luxury is not available in the ocean, right, so that has forced us over the past sort of three to four years to really try and understand how some of the algorithms we already knew about could better incorporate models that were designed by non-roboticists to do something entirely different. And how to interact with those models and how to use them in a meaningful way and in a sound way so that our planning techniques and our estimation techniques would work and would work better has been one major focus for us, and we continue to work on that quite a bit. But these robots are not just out there wandering around aimlessly, and precision is great and knowing where you are is great, but the point is to do something with the robot, and our robots are actually doing something very interesting right now. These are robots that are designed in the ocean to find and track plumes of algae species that are blooming in the water. So if you have an algae bloom in the water, we want to be able to find it and track it. So these robots are actually chasing microorganisms around in the water, if you will. The robots themselves are very big, but these aggregates of microorganisms form this huge bloom, and we want to be able to track it, right? Now, the interesting thing about this, of course, is that there's a set of models that predict to some degree of accuracy the physics of the ocean, which is what I talked about, and these need to be integrated into our robot estimation techniques.

Q: Currents and...

Gaurav Sukhatme: Currents and things like that. Then there's another set of models that says something of the ecology of the bloom itself. How are the population dynamics of this microbial species changing? And so, how fast is this bloom growing, for example? What is its spatial extent? And those are ecological models, so even though the water is perfectly still, these things are multiplying. The bloom's becoming bigger or smaller. And then there's a third set of models, which says something about the toxins that these things produce, and these toxins diffuse in the water, and those are chemical models about chemicals diffuse in the water. And so, in fact, it's an enormously complex undertaking to design estimation techniques that not only gauge for the robot its own position and do so in a meaningful way and an accurate way but also to gauge something about the state of the world that it's trying to observe, and in our case it's trying to observe this very particular biological process that's embedded in this very complicated environment, right, it's trying to follow and sample from in an efficient way. And so, this is what I meant by large spatial scales and large complexity problems when you try and do them for real, as it were. And, of course, we don't want to do this for one-hour test and claim success. We want to be able to do this over the lifetime of such events, and these events are typically a week or two weeks or three weeks at a time in the water, and so we want to be able to persistently be able to do this, right? And so, that's one major challenge, and where our laboratory and several other laboratories that are working on these issues, in monitoring the natural environment have really been addressing some of these problems. There are various facets to these problems, but that, I think, for us will represent – in the next decade or two will represent a tremendous advance once we've been able to sort of design such infrastructure, and our vision is very simple, and our vision is this. If you go into a biology lab today, there are people there using microscopes, but they don't have to know about optics. They don't have to know necessarily a lot about how these things work. They just work and they use them. I'd like my colleagues in marine microbiology to be able to use robots the same way. I'd like them to take a bunch of oceangoing robots, and they don't need to know necessarily the underlying planning algorithms, the underlying estimation algorithms. They need to be able to designate high-level goals for a group of these robots by pointing and clicking, as it were, and then let the system achieve those goals by reorganizing as needed. And so, as the environment changes, as this bloom moves, the robots have to move with it. As it splits into two, the robot group ought to split into two. As parts of the problem become insignificant, the robots need to refocus on the parts that are significant, right? And that's the kind of autonomy that we would like robots to have in exploring an uninstrumented environment, right, an environment that's out there, not an environment where we control everything. And it would be for me personally very gratifying if my colleagues were to use this networked infrastructure of robots the way they use a microscope, right, so without thinking how it works, but they know how to use it, and it helps them do their work better. And that allows us to make some progress, not just in thinking about algorithms in the abstract, but in the end sort of we advance by building a new scientific instrument. Potentially we advance what other people do. And by studying things like the environment using robots, I think we contribute more than just technology. We contribute something to society which is larger than just the technology, because, of course, if you can know something about how pollution moves in the water or how toxins spread in the water, then this has public-policy impact. And so, I've become more and more interested in thinking about larger robotics problems that have social impact. So the science application is a good driver for us, but we see a wider world beyond that, where environment monitoring, monitoring the natural environment, is a very, very important thing, because with better monitoring can come remediation and can come things like better public policy. Today when you go to the beach, if there's a beach closure probably there was something bad in the water three days ago, and that's why somebody's put up a beach-closure sign today, right, because the water sample got taken to some lab, and somebody said, "Oops, people really shouldn't be going in the water." If you had a fleet of robots that was monitoring the water, a red light on the pier would come on when the water quality was poor. Now, of course, I'm not just suggesting that this is important for recreational use, but it is tremendously important in parts of the world where water quality is a very, very serious public health issue, right, and we hope to be able to make impact at a large scale in that way by doing environmental robotics. So that's a big umbrella driver, and then there's, of course, technical problems that we work on shorter time scales.

Q: How did you first become interested in marine robotics?

Gaurav Sukhatme: I first became interested in marine robotics and aquatic robots more generally because I participated and still participate in a National Science Foundation center called the Center for Embedded Network Sensing, or CENS. CENS is an NSF science and technology center, an STC, which is a collaboration between several institutions. The lead institution is at UCLA, and the STC is headed by Deborah Estrin, who is an expert in embedded-network sensing. Participating institutions are USC, Cal-Tech, UC Merced, and the focus of that center from the get-go about eight years ago when it began – eight or nine years ago when it began, was to develop embedded network-sensing technologies to explore the natural world, to reveal things about the natural world that had previously been unobservable, as it were. And now, of course, not all this infrastructure needs to move around, but it was clear to us from the beginning that some of the infrastructure that the center would study would be mobile, because it simply wasn't possible to put sensors everywhere you needed to make measurements. And, of course, once you have moving sensors and they move autonomously, you have robots. And, in fact, when you have many of them, you have a multi-robot coordination problem and you have multi-robot estimation problems. So that was my role in the center. But what was very, very compelling about the center from an early stage was it brought together people doing these technologies for embedded-network sensing- some of the technologies were networking technologies and some of them were sensor development. Some were robotics and a variety of technologies, but it paired them with applications, as it were, with domains, with real things that needed measuring and monitoring. And in the way that this very often happens, through some early conversations, the particular domain that it became obvious to us might benefit through multi-robot monitoring, as it were, was the water. It's a bit hard to get sensors to move underground, say, if you want to make seismic experiments, right, or if you want to make other kinds of measurements. And so the water was very early on identified in the center right from the time we wrote the proposal to be one such area, and I had the very good fortune at the time of starting a collaboration with a colleague here at USC whose name is David Caron in the biological sciences department. And that collaboration through CENS was how I got started sort of on this field, and really about eight or nine years ago I had no experience doing robots in the water. I did a lot of multi-robot coordination. I did a lot of estimation work. Some of it was for robots on the ground, and a substantial amount, actually, was for flying robots in the 1990s, and I slowly moved away from those areas to an increased focus on dedicating my efforts mostly to robotic boats and also underwater vehicles. And that's primarily where my lab's been working now.

Q: So, what do you see as some of the central problems facing, I guess, these other areas, so the multi-robot coordination problems, like what are some of the central problems, and have there been big breakthroughs?

Gaurav Sukhatme: There have been big breakthroughs. Some of the breakthroughs have been foundational. So, people wrestled with really foundational issues in robotics all through the 1990s and the early part of the past decade on things like very fundamental problems, problems in allocating tasks across multiple robots. And there's been considerable progress on that. How do you do multi-robot task allocation? If you have a team of robots and some have to do certain tasks at certain given times and others have to do others, how do you do this? And there's been considerable progress in this area. There're variants of these problems that remain open, but we know a fair bit more about these areas than we used to. There's been foundational work on estimation in single-robot systems and in multi-robot systems, locating where robots are. The localization problem received a tremendous amount of attention in robotics, and really the state of the art in it now is dramatically different than what it used to be 15 or 20 years ago. Systems are a lot more accurate and can do amazing things. The problem of building environment representations, building maps of environments, is also a problem that's received tremendous interest in robotics over the past few years, and that, too, has matured, and there're highly capable robots that do this. In fact, sort of the autonomous robot that locates itself and builds a map of the environment and solves the concurrent problem, if you will, the simultaneous localization mapping problem, is another huge area that's received tremendous amount of attention and has made considerable progress in this area. But now, doing these things in more and more challenging environments sill proves to be tricky when you want to solve some of these problems, which are in some sense fairly basic problems, but when you want to solve them in environments which are very crowded or in environments that're highly unstructured or in environments that're very dynamic and are changing or where the robots are very severely constrained in their abilities like there maybe micro-robots, so they don't have a lot of computation, or where the robots are constrained in their abilities to communicate with each other. All of these variants still make a variety of these problems very difficult, and there's plenty of work to do in these areas and capabilities to build, right? And then there are areas beyond these foundational areas, but depending on how you want to think of them, they're sort of integral to the overall enterprise, which is most of us think of robots as things that ultimately get used to do something, do design algorithms for basic capabilities, but in some sense you also need ways to make the whole thing work for some purpose. And of course a huge field in that is the field of human robot interaction in which I am not an expert at all but it's an important field because it simultaneously asks the other question which is how are robots to be – how should they interact with people, right? And how do you design for that? How do you take a principled view of that? And there is a huge upsurge of activity in that area. And I don't know much about it but I read about it and I follow it and it's a tremendously important area. So there are several I think such areas. And now I think increasingly roboticists, many roboticists, see this very interesting confluence of communication computation sensing and actuation. And when I began as a roboticist about 20-odd years ago communication wasn't much in the picture. So in terms of sort of fundamentally technologically what has changed is it was rare to have robots that had really small working radios and it was rare to think of a multi-robot system where the communication was as effective as it would be today and that has changed. Sort of consumer grade small communication is everywhere so the cell phone explosion has happened in exactly the period of time that I'm talking about in the last 20 years. And so I very often even as a multi-robot researcher in the early 1990s thought about multiple robots each of which had sensing and each of which could command motions and each robot had computational capabilities. And yes, they communicated with each other and sometimes they didn't. But I didn't think very actively about the radio. I didn't think very actively about the communication systems. And much of my work in the late 1990s and the early part of the past decade was on network robots, on robots which explicitly used communication in interesting ways to coordinate their activity. And robots that had models of how communication of how communication was likely to work in certain areas and use that proactively in the planning that they did to do certain activities. So I used to work in this network robot area in multi-robot systems and I believe it still remains a challenge because there is increasingly sort of difficult environments in which one can think of several problems that are open in this area. And now of course one example of that of course is doing network robots in the water, the kind of thing we do but there are other areas in which networking and robotics is sort of there's been this very interesting area that's come into being in the last decade.

Q: *Robocop.*

Gaurav Sukhatme: *Robocop* is a great example and Robocop combines something else with it. It also combines a game theoretic sort of problem, right? So there's game strategy on top of everything else. So there is an adversary who is trying to do something different in the environment, trying to oppose you in fact. So that's another level of complexity. And so but that commoditization of small computation communication, sensing and actuation has also led to an explosion of more easily accessible robot platforms and *Robocop* is a good example of that. So it's just got a lot more people with their hands in robots being able to program them which is a very healthy thing, which is a very good thing in my opinion.

Q: What do you think are the biggest challenges facing robotics in general as a field?

Gaurav Sukhatme: Robotics is in some way an amorphous field. It has tremendous reach and it makes the question a bit hard to answer. To me now this is a biased view, obviously, but to me robotics is almost a there is virtually- there are very few exceptions but there is virtually no area of engineering that is not touched by robotics. If you enter the School of Engineering as a freshman and you don't quite know what kind of engineering you want to do, I would recommend you take a robotics course because you will see some electrical engineering and some computer engineering and some computer science and some mechanical engineering. And you may even see some chemical engineering occasionally and you will certainly see some biomedical engineering. It might be somewhat rare to see civil engineering I'll give you that, but there's a tremendous – there is plenty of breadth. So this problem is saying what are the biggest challenges in robotics is a bit hard to answer precisely because of this diversity. It means many things to many different people. I am a computational person. And so to me my universe are the computational problems in robotics. They are not mechanical design products in robotics of which there are many. And then there are problems that are not strictly robotics problems but they will impact robotics in a dramatic way. I'll give you one example. And that is the battery technology example. A lot of robotics will change dramatically if battery technologies change by a factor of ten, right? But that's not a statement about robotics. Many, many other technologies will change if we discovered how to do things where there was some other battery technology out there, right? And but robotics would benefit from it too but it's an issue. So if you look to problems that perhaps are not like the battery problem but are more sort of real robotics problems, then I think one area of challenge for robotics is really making the leap from algorithms that are well understood and whose mathematical properties are well understood. To sort of evolve and to be able to figure out which of those work in practice and to be able to design for more ambitious deployment if you will. There is a lot of work gets done in models and gets done in the lab and the subject is hard enough that it takes a long time to figure out which of the competing techniques really has legs when it comes to sort of going outside the laboratory. And it isn't a statement of fault but it is a statement about the hardness of the domain. And we labor under one issue here, one problem which is that it was a long time before people understood just how difficult it is to build physically embodied things that function autonomously in the unstructured world. We really grossly underestimated how hard it was and I say "we" generally. I mean just people generally underestimated. And it's not surprising why.

Biological systems seemed to function just fine in the natural world and we can't help but notice it and we can't help but feel like oh, that shouldn't be too hard but it is. It is. And so the transformation from understanding systems and from devising algorithms and doing good mechanical design in controlled circumstances to the full sort of uncontrolled, unknown environment in which you want the system to really operate proves to be a much bigger barrier than people really think. And we are getting increasingly better at crossing that. There have been dramatic improvements. But the more we do that, excuse me, I think the more successful we will be. And that still remains in some ways a big challenge to surmount that and many, many people try to surmount that and to push on that and I think that remains a problem. I'll tell you another thing that I think is not a problem but just a comment. In some ways robotics like some other technologies perhaps like A.I., has this property that when it's truly successful people cease to notice it. And I use the word "robotics" to really cover robotics and automation. I'm not really trying to distinguish too much from each other. People often say this to me they say "You work in robotics. What I really want is a robot dishwasher. How come there isn't one?" And I tell them you already have one. I mean you already have one in your house. It's a completely roboticized dishwasher. Once you press the button it does all the dishwashing. You don't wash dishes, right? But of course that's not what they really mean. What they really mean is they want a robot that can load the dishwasher, and can unload it, dry the dishes really well and put them away. That we don't know how to do. It's not the dishwashing that proves to be hard. It can be done and it's been done for decades but it is an automation technology. It is a robotic technology. It involves very sophisticated sensing and it involves very sophisticated use of limited resources to optimize a particular task. And the keyword there is "particular task". So I suspect what we hanker after are things that are a little more general and things that are a little more unstructured. And we don't quite see those. And so the push towards generality may not be fully general purpose robots but some generality, some ability to do something less than – something more than just the very narrow thing. And the ability to function in a sort of an unstructured environment is where we want to be. But along the way we have successes at doing things that are more limited, but then they cease to be recognized as robotic successes. There is I think a coming revolution in robots for monitoring the natural environment. It's going to happen. When systems like this are able to patrol or underwater systems of the kind we work on are able to patrol the oceans, and provide data and shape public policy about clean water, one of the things that might happen with such systems is that when they become successful, people will no longer call them robots. People will simply call them automatic boats or mobile sensors or something like that. There is another revolution coming in autonomous cars. I don't work in the area but there are several others who do. I imagine you are going to talk to them and it's dramatic how much is changing. But my suspicion is that 30 to 40 years from now we will all drive cars that will be significantly more autonomous than what we drive today but we will not call them robots. We will still call them cars. They will just be more autonomous. So robotic technologies like some technologies, do have the potential to be very successful but when they are very successful, very often they are not identified as being robotic technologies. Another area in which I think robotic technologies have had tremendous impact which is not widely recognized is the genomic revolution. So if you want to do large scale sequencing then a lot of the grunt work is done by robots. And these are robots that are in the backroom usually. You don't see them, they are bench top robots that do chemistry. They do the repetitive aspects of chemistry: handling micro vials and moving things around and mixing things and things like that. And there is a lot of sophisticated robotics that happens in order to streamline that smoothing it out undoubtedly in a structured environment. The field would not have advanced as rapidly were it not for that success. So there are some curious aspects to robotics technology. One is it's very broad. Two, it has tremendous impact but very often the impact tends to be siloed and then gets subsumed within some other bigger field and it's less clear to people that it's a robotic technology that's had the impact but rather something else. And three of course is the truly general purpose or at least somewhat general purpose robot still remains a difficult problem fundamentally so there is some confluence of each of these. But I want to say one thing. This is not a complaint. This is good. This is the excitement of robotics. It's the ability to make an impact in so many different ways is in fact what's interesting about it and I think it's part of the reasons why I originally became interested in robotics because it's the kind of field you can become interested in from almost any vantage point in engineering and in fact even beyond engineering to some extent but I'll restrict myself to engineering. It can be almost any kind of engineer would be interested in these technologies. It's a draw. There are so many open areas and so many ways to make contributions. So I view these things as positive. I don't view them as oh, it's too bad. I think this is it. This is one of the good things about robotics.

Q: How do you see the relation to ubiquitous computing?

Gaurav Sukhatme: Yeah so much of the work we did in the early part of the past decade on network robots had some strong connections with ubiquitous computing because we were and to some extent continue to work on robots that exploit the fact that there is other computation in their vicinity that they could use. So some of our work on network robots for example relies on computation that's happening outside the robot. Relies on a network to supply it information and say for navigation or say for finding things in an environment, things of that kind. And so pervasive or ubiquitous computing has been an interesting way to think about robotics. Particularly with the built environment being so instrumented today there is I think significant potential for thinking about robots that have just enough autonomy to do useful things in the built environment but they rely extensively on embedded infrastructure in the built environment. And such systems I think are making rapid advances. If you want a robot to function in an airport and you want it to ferry people from one gate to the other or you want it to do logistics of some kind, you want it to have autonomy so that it can navigate its way through crowds. But for it to figure out where it is needn't rely necessarily on itself but could rely on sensors in the environment because one can instrument the environment to some extent and there's no particular harm in doing so. And communication between these embedded infrastructures in the environment in the robot is now a reality. So I think there is an interplay in large built environments in particular where robots will need autonomy and will need computation on board, and will need to rely on their own sensing to solve certain problems but will be able to harness the power of ubiquitous computing to do other things. This is an active field people in network robots work on it, our labs work on it, and several other laboratories work on it as well.

Q: So apart from George Bekey who else do you look to as being your mentors or people who guided you in the field along your career?

Gaurav Sukhatme: So there have been a variety of people. Some people in the field have been inspirational in the sense that I didn't necessarily ever work with them directly but when I read their work or when I came to know about their work and when I interacted with them informally this just sort of inspirational, that was really nice. And so some people of that type are one person who comes to mind is Takeo Kanade at Carnegie Mellon. He's a computer vision person but of course he's made contributions to so many areas that labeling him as computer vision is not quite fair. And there's that sort of was when I was I met him first when I was a young Ph.D. student and heard him speak and really developed an appreciation for there's a tight interplay between computer vision and robotics, obviously so that's one person. Another person who was inspirational not because I worked with him or even do the kind of work he does but just was inspirational was Rod Brooks who did tremendous amounts of work in robotics and continues to do amazing things. And when I was a young student again just hearing him speak and sort of following what he did was inspirational in many ways. There was fairly early on when I was studying sort of basic robotics early on, I did a lot of reading and also had the good fortune of meeting Oussama Khatib who did foundational work in robotics in the late '70s and early '80s and is still incredibly active today. And so his is sort of an inspiration in the field for many of the things that he has done. Another person who these are all people I have never worked with actively but I've sort of known. I know many of them well now but I first came to know them at a distance as it were, people whose work I sort of followed and came to know and there are some connections between their work and mine but not necessarily that we've collaborated. Another inspirational person along this vein is Hugh Durant-Whyte who has been at Sydney for many years and really was an inspiration to many people in my generation for the kinds of things he did. There are just a tremendous number. I could name easily four or five other people possibly many more and I suspect once I go away I'll think of several others and say oh, I should have told you about so and so. But I'm not trying to be exhaustive in any way. I can literally name several others but at some point I would have to stop. So in terms of sort of actively working with people I actively of course work with George Becky who was my Ph.D. mentor and subsequently I haven't worked with anybody in that relationship but I've had the good fortune of collaborating with other people, some who have been somewhat more senior than I have been who have been inspirational. And some collaborators that come to mind who have been really fun to work with and an inspiration and also to some extent have mentored me are Deborah Estrin who is not a roboticist, who is a pioneer in computer networks whom I've worked with over many years. And Daniela L. Rus at M.I.T. with whom I continue to collaborate actively and have collaborated on and off for I'd say more than a decade now. And Vijay Kumar, Penn who I've collaborated on and off with for some time now. And then I'd like to mention another thing which is one of the reasons that I'm fortunate to be at U.S.C. and one of the reasons I really like being here is I collaborate actively with the other roboticists here at U.S.C. And so over 14 years I have collaborated with Maja Mataric on and off. We have a joint grant even now and that's been a wonderful collaboration. We really enjoyed it and I really enjoyed it and I collaborated actively with Stefan Schaal my other robotics colleague here. And to a smaller extent I've interacted with Wei-Min Shen who is the other robotics colleague here on campus. And so there are sort of- it's really nice to be at a place where you do your own work and your laboratory does its own work but you collaborate with the labs of colleagues who are nearby and they share some interests with you and there's enough overlap that you do interesting things together. That's been one sort of very, very nice thing. Again, in terms of collaborators I could go on and on. There are many and I tend to be fairly collaborative. I have a large number of collaborators but again, I don't think I'll try and be exhaustive. I'll just name a few people.

Q: Okay. In terms of people you've mentored how many Ph.D. students have you had? What kind of projects have they worked on?

Gaurav Sukhatme: I'm trying to count very quickly now but I think if I have the count right 14 people have graduated with Ph.D.'s from my lab. I've also mentored about seven or eight postdoctoral fellows. These are the people who have had the longest associations with me for whom I've been a mentor basically for 20 odd people. I've also mentored undergraduate and master's students who have been in the lab, but over the longer term it's really these fourteen Ph.D. students and six or seven postdocs, roughly twenty people who have been sort of the core people I've mentored over the last eleven years roughly, or twelve years roughly. They are spread across the spectrum. Some of them worked on network robots. Some worked on aerial robots. Some worked on estimation problems independent of which robots they were applied to so they've worked on a diversity of problems. They have gone on to various careers. Several of the ones have graduated just within the last year or two are postdocs at a variety of places. Some people who are slightly older than that are now faculty and running their own labs.

Q: Where at?

Gaurav Sukhatme: So one of my students Shrikanth Saripalli is a faculty member at Arizona State University where he's an assistant professor and doing very well. Ryan Smith who is just finishing his postdoc up with me will soon start as a faculty member in Australia in Brisbane. And Richard Vaughan who was a postdoc with me and also with Maja so we worked together the three of us is now an associate professor at Simon Fraser University in Canada so there's a bunch of people. Denis Wolf who was my Ph.D. student is now an assistant professor in Brazil. He is Brazilian and came to my lab and did his Ph.D. with me and then went back to Brazil and joined the faculty. And so there's a bunch of people like that. And then there's others who went on to careers in research at research labs or not necessarily academic research labs but they pursue research in research labs. And then there are still others who went on to pursue careers in other areas of computer science and not necessarily in robotics even. I have one student who graduated with me who now works for Microsoft and does completely different things. But that's I suspect most people have stayed in robotics and in fact almost all have stayed in research in some form or the other. Some do R&D and some do more academic research.

Q: What are some of the commercial R&D industrial uses around that?

Gaurav Sukhatme: Right so I'll name two just to give you examples. Sameera Poduri who did her Ph.D. with me and also postdoc with me briefly is now at Qualcom's Bay Area Research and Development Center in the San Francisco Bay area. Stefan Hrabar who did his Ph.D. with me several years ago on robot helicopters continues to pursue research on robot helicopters at C.S.I.R.O. in Australia which is the large Australian agency that coordinates Australian research in diverse areas of science and technology and other areas as well so he works there. There are in addition Maxim Batalin who worked with me on network robots in about seven years ago works now in an entirely different field. He works in wireless health systems and he's still here in L.A. locally across at U.C.L.A. in a new institute that they've set up for that purpose. And Boyoon Jung who worked with me on his Ph.D. thesis works for a company called NavCom who are experts in G.P.S. and many other areas of outdoor navigation solutions for robots and they are a wholly owned subsidiary of John Deere and so he works there and fulfills the R&D role. Another one of my students, David Naffin also works at the same company and so they both work in R&D in a company which is doing robotics but is part of a larger concern. So there are sort of several people who do many, many different things. I must be forgetting somebody but apologies to whoever I forgot. I am free-associating as I go.

Q: Sure. So looking back at the history of robotics, a lot has been done in East Asia, North America, Europe but yourself coming from South Asia having students that are now teaching robotics in South America.

Gaurav Sukhatme: Yeah.

Q: Where do you see emergence of robotics in these other countries that maybe didn't traditionally have robotics programs or when you were a student in India was there an opportunity to study robotics there or did you have to come to North America for your training?

Gaurav Sukhatme: I think there is now much more opportunity globally than there used to be. It's certainly the case that when I was an undergraduate in India in the late 1980s there was not a significant robotics focus. I took no robotics courses as an undergrad for example. None were offered and today that's not true anymore. There's a significant number that would be available. One thing that has changed and I alluded to this earlier when I spoke was the interesting thing is that robotics is a lot about doing and it is an engineering science that relies intimately on experimentation and sort of theory. And one of the positive developments has been that the commoditization of small robots has really sort of given access to many, many more people than formerly could afford them, could use them. And this is a very, very healthy and positive thing so just the ability for people in many, many different parts of the world to pursue robotics projects during their school life or their undergraduate career is just now much more possible than it was 20 years ago when there were very few robots and the development in some ways is not unlike the development of computing, right? So if you did computer science 40 years ago there was a computer at the university and everybody had some sort of access to it occasionally maybe. And then there were other universities where they didn't have a computer and then in the 1970s that changed and by the 1980s with the rise of the personal computer it changed completely. And now of course we no longer tell our students that you have access to computation because every student has their own access to computation. And yes, there's computational infrastructure universities provide and certainly for high end computation you still need that. But for a lot of personal needs to study computer science in so far as you need a computer you certainly don't need a university to provide it, do you? It's important in some cases but it isn't necessarily a show stopper. And that's slowly changing with robots as well where it would have been to me it would have been almost inconceivable as an undergrad to have had access to robots to program other than very simple toys. The sensing wasn't quite affordable enough or cheap enough or ubiquitous enough and the networking wasn't quite affordable enough or cheap enough or so on, small enough. And neither was the computing so everything was complicated and it was quite a big thing to have a robot and that is I think changing and that I think is leading to a lot of diversification. The other of course is that with the rise of the open source movement and with the rise of networking technology has come increased access to information which has proved wonderful. So if there's a simulation environment, nothing stops somebody in South America or India to be working on exactly the same problem as somebody in Europe or in the United States or in Japan, right? Whereas in the past the geography would have been a barrier. And today it's less of a barrier. I'm not trying to say that it's not a barrier. But it's been less so I think the spread of simulation tools and the spread of access to data and just the spread of commoditized robots themselves has enabled. Now of course these are just enablers. They don't do everything but I think these have been healthy things. They have been healthy things for robotics so I see a future where there will be increased participation from parts of the world where formerly there was not as much work in robotics and that's a very healthy thing for the discipline.

Q: Is there anything else you want to add?

Gaurav Sukhatme: Well, what can I say? What I'd like to add is that I am passionate about robotics and I continue to be passionate about robotics. It is not the kind of thing for which I have ever lost my fascination. So it is not a job. It is really for me something much more than that. I think so if I had to end I would end with an invitation. So if anybody watches this ever and they are thinking about what to do and they'd like to do some form of engineering, I'd say think about robotics. Robotics is a marvelous discipline. It is to me just a never ending source of interesting, challenging problems and has tremendous potential to change the work which is I think important as well. It's not an exercise in just sort of gratifying one's intellectual curiosity but it is a beautiful match of intellectually difficult questions with socially relevant things and I think that to an engineer, what more could you ask for?

Q: Thanks.