Oral History of Raj Reddy

Interviewed by:
Hansen Hsu

Recorded via Zoom April 14, 2021
Pittsburgh, PA

CHM Reference number: X9458.2021

Hsu: It is April 14th, 2021. I am Hansen Hsu, here with Raj Reddy, and so first, we have a question called "I am," so sort or repeat after me but sort of insert your name in place of mine. So given your worthy accomplishments in many areas, how would you describe yourself? For example, I am Hansen Hsu. I am a historian, museum curator, et cetera. Can you please state your name and choose a few titles to describe yourself? Say, "I am," state your name, "I am a..."

Reddy: I am Raj Reddy. I'm a professor of computer science and robotics at Carnegie Mellon, and I've been at Carnegie Mellon for slightly over 50 years. Before that I was at Stanford.

Hsu: Okay. Do you-- is there anything else you want to add that, that description?

Reddy: Me? You want me to add anything?

Hsu: Yeah. Yeah.

Reddy: Yeah.

Hsu: Is there any-- yeah. Oh, you don't have to.

Reddy: Yeah.

Hsu: Oh. Actually, can you say-- could you say it again just a little bit? I think that was a little slow. I think Jon wants to get another take of that.

Reddy: I'm Raj Reddy at Carnegie Mellon University. I'm a professor of computer science and robotics, and I've been here for over 40, 50 years, and I primarily work on A.I. research and human computer interaction. Most of my work has been on speech, vision and robotics on-- in A.I., the perceptual side of it.

Hsu: Okay. That's great. So next we're going to start with sort of a brief-- well, sort of a-- we're going to go for an overview over your life, and so we're going to start with where and when were you born?

Reddy: So I was born in India, in a small village of 500 people in the south [of] India near Chennai. It used to be called Madras at that time, and in Andhra Pradesh, and the name of the village is Katur, and I was born in June 13th, 1937.

Hsu: Okay. Could you talk about growing up in the village?

Reddy: So I-- and I grew up in this little village for the first nine years of my life and went to primary school there, and the interesting thing about primary school in the 1940s around the world is the Second World War, was there were shortages everywhere and in the school, we didn't even have paper and pencil, pen, pencils, so they taught us how to write the letters, the script, by using sand. We would kind of clear the sand and write on the sand, and that worked okay at the time, and it's not uncommon if you look at anyone in India at that time who went to school during the Second World War, and that's the time they would've had the same experience.

Hsu: Did your village have electricity or running water?

Reddy: At that time it did not have electricity or running water or any of the normal facilities we take for granted today. Now they have electricity, they have internet connection, but they still don't have running water into every home, but they're-- have water taps and water supply solutions in the village. But I think really soon they're predicting they'll have running water.

Hsu: Okay. And you were like the only member of your family to go to elementary school; is that correct?

Reddy: Yeah. No, basically, at the age of about nine, because we didn't have elementary school in the village, I was sent away from home to the nearby city to live with friends of the family and go to school there. That town is called Srikalahasti, and I went to elementary school and high school there and graduated at the age of about 15, 16 from the high school in Srikalahasti, and then I went to Chennai or Madras at that time. Madras, to Loyola College, where I did my junior college. It would be what you would think of as 11th and 12th grades, and then joined polytechnic or technical school called College of Engineering in Chennai at that time.

Hsu: Hm, I see. Could you talk, tell us a little bit about your family, your parents and also your siblings? How many siblings do you have?

Reddy: Yeah. I had three brothers and three sisters, a total of seven of us, and it was not a very large family but not small either. These days almost everyone just has one or two kids. In those days it was normal to have 5 to 10 children in a home.

Hsu: Mm-hm. And what was your father's occupation?

Reddy: My father was-- had land. You know, we were farmers, and my-- he never went-- finished high school, and my mother also didn't go beyond primary school, I think, and that's how I grew up, but normally in the house we had books and it's not uncommon for us to have lots of books to read, but they're not English books. But they were books in local language, Telugu, and I had an adequate access to library books at that time.

Hsu: So you were actually fairly well versed in Telugu literature; is that correct?

Reddy: Yeah. At that time, you know, all the way to-- until the graduating from the school, high school, most of the medium of instruction was in Telugu and therefore-- and also the literature and everything. At that point I would not say I was well versed in the Telugu language. I was comfortable with it, you know.

Hsu: But you were fully literate and as-- your siblings were also literate as well. Okay. Yeah. Could you talk a little bit about how your early life in the village influenced the-- sort of the trajectory of your career later on in terms of your interests?

Reddy: Yeah. My early life mainly meant I was kind of living away from home after I was nine years old, in-- both in elementary school and high school and college, and so most of my life I, after 9 or 10, I've been away from home, you know. That's not usually the case for most people.

Hsu: And is there anything else from your college days at Loyola and in the engineering school that you'd like to discuss with us?

Reddy: Yeah. <laughs> The most interesting thing that happened when I went to Loyola College is all the medium of instruction was in English, right. So I went to the Telugu medium school, and many of the instructors, we had Jesuit priests from Scotland and Italy and Ireland, and so not only did I have to deal with the conventional Indian English accent, but also Scottish accent and Irish accent and so on, so for the first six months I didn't understand a word of what they were <laughs> saying, but fortunately I could read books in, you know, books in physics and mathematics and so on in English because my English background in-- from school was okay. Just that we didn't-- that was not the medium of instruction. We, you know, the class was all taught in Telugu, except the language, English language classes. But even there you never spoke in English other than to repeat some sentences. If you remember the "Slumdog Millionaire" scenario where they talk about the Three Musketeers, that would be the kind of English instruction I would have had.

Hsu: Mm-hm. Okay. Yeah. So let's move on to your time in Australia then. How did you end up going to graduate school in-- deciding to go to graduate school in Australia?

Reddy: Right. So I graduated with a what they called BE, like a BS here in BTech.

Hsu: And which year was that? 1958?

Reddy: 1958. And it turned out India was part of the British Commonwealth, which it still is, I think, and Australia had a number of internship programs for students graduating from India, and so three of us applied and got in and were selected, so all of us were classmates together, so the three of us ended up in Australia at that time. This was, would've been, 19-- in December of 1958, and so I was in Australia for five-- four years, I would say, until 1963.

Hsu: Okay. This was-- which university was this?

Reddy: No. The university, I was, you know, kind of-- I was actually sent to the Ministry of Roads and Buildings to be an intern, and after one year, then I joined University of New South Wales to do a master's degree in Civil Engineering, and at that time, my advisor, the head of the department, Stan Hall, had just returned from Britain, and there he was exposed to the use of computers in large structural matrix calculations, and he, when he came back it so happened University of New South Wales also just got a new computer, it's called English Electric DEUCE Mark II, and it was a mercury delay line memory computer, and it was a vacuum tube [based computer]. Everything was vacuum tube. Was a huge, you know, it was a-- would've taken a whole room like this, and I was kind of exposed to this computer. Now, before that I'd not even seen one, so I didn't know any better at that time that-- what this computer was like, and so this computer had one language, at least the one we used, called GIP, other than machine language programming. GIP [was a high level programming language,] it's a matrix language. Every [data item was represented as a matrix] it was a single data structure. Everything was a matrix, and so if you wanted to add two numbers, you had to make them into two one-by-one matrices and then add them and that's how you did all the calculation, and it turned out to be very powerful. It's a three-address machine language and so GIP was a very powerful, high-level language for 1959-'60, when I was using it.

Hsu: Mm-hm. So this was G-I-P? Is that GIP G-I-P?

Reddy: Yeah.

Hsu: Oh, okay. What does that stand for?

Reddy: Something Interpretive Language [General Interpretive Program]. Inter-- no, I don't quite remember what it was. I believe, you know, you'll find English Electric DEUCE Mark II in Wikipedia entries. <inaudible> Wikipedia entry. I found lots of entries for that, those kinds of computers. It was basically a enhanced version of Turing's original ACE computer. ACE design was A Computer Engineering-- Engine-- Computing Engine. That's how it got the name ACE. ACE. And so when commercial version of that was manufactured by English Electric, they enhanced it a little bit and called it DEUCE, and so that was the computer we were using at that time. It had a display, micro, you know, CRT display, and it had a speaker, which would kind of have all kinds of sounds, and all the computations the computer was doing at any given time had a unique sound and so we could actually tell what was going on, and somebody actually then figured out how to make it sing songs and so on. There were computations that you could do which would then say, "Daisy, Daisy, <laughs> give me an answer too," or something. You know, these sort of songs from after the Second World War, I guess, so...

Hsu: Yes. Thank you. And then after you finished your master's at New South Wales, you went to work for IBM, is that correct?

Reddy: Yeah. Yeah, I-- at that time there were not many people that even knew what a computer is and how you'd use it and so on, so when I applied to IBM they were very happy to immediately hire me and then at that time the IBM computers were IBM 650, which is a magnetic drum memory machine, and IBM 7090, which is a vacuum tube computer, 36-bit machine. You might have some of them at the Computer History Museum, but they were the early versions of IBM computers, and after I joined them the first semiconductor machines were manufactured. The first ones they ever brought out were IBM 1401 and 1620. The 1620 was a scientific machine; 1401... 1401 was used for commercial processing, and they had-- at that time there were no keyboards and monitors. You would do-- use punch cards to program whatever and data and feed it into a punch card reader, and the output was a line printer that the-- laser printers did not come for another 20 years at that-- for a long time, all computer output was on large sprock-- line printers. Again, you probably have them at the History Museum, but you don't see them anywhere else at this point. <laughs>

Hsu: What did IBM hire you to do?

Reddy: I was-- because of my technical and scientific background, I was hired as a applied science representative to kind of help them, but mostly I ended up doing systems programming for-- because that was something that also needed to be done, and so that's what I was doing. I was a technical representative. Not in sales or marketing or any of that, although they trained us for many-- in all such things. One of the interesting parts of that, that part of my life, is I probably got as good an education at IBM as I got anywhere else prior to that. Almost half of my first year I spent in some training program or the other and it was all off, you know, off site at some remote location and it was very unique and bonding experience with all the other people.

Hsu: Great. And what-- well, while you're at IBM, what led you to eventually decide you wanted to get-- to go back to graduate school for another degree?

Reddy: Yeah. Basically-- <phone rings> excuse me one second. I'll turn this off.

<00:21:19 off-topic conversation>

Hsu: So while you're at IBM, what made you decide that you wanted to return to graduate school and get a higher degree?

Reddy: Right. So at IBM I would come across various journals and articles, and one of the articles I read there at that time was Minsky and McCarthy's and Newell and Simon's papers on A.I., predicting how A.I. will solve many of the problems of-- within 10 years, like Simon had famously predicted at that time, that a computer will beat a world champion in chess in 10 years, and so it didn't happen in 10 years. It took 40 years but it did happen, but what he was kind of saying was whatever we consider to be intelligent behavior in human beings can also be done by machines, and it turns out that turned out to be a much harder problem than we-- any of us thought at that time, mainly because we did not understand some of the complexities of sensor processing, speech and vision and robotics, because we do them but we don't have algorithms and-- to tell us how we understand speech or how we see and perceive the world, and so you had to build whole models and theories of speech perception and visual perception and very quickly we discovered we're very far from equaling human performance. So the short answer is I came across lots of papers on A.I., and I thought perhaps if I was going to do a graduate school that I should not continue in engineering but in computer science, and at that time there was no computer science school, so I applied to programs at Stanford and other places, and so I got into Stanford and ended up there by 1963.

Hsu: Mm-hm. So you had already decided you want to get a higher degree though, but you could've also stayed at IBM; is that correct?

Reddy: Yeah, yeah. In fact, they offered me jobs in other parts of the world, IBM world. Mainly as they said, "You can go back to India and start, you know, work for IBM India." But I-- by that time I'd already decided to continue with advanced graduate education.

Hsu: Hm. What was the primary reason why you wanted to go, I mean, even before you discovered all these papers? Had you already decided to go back to graduate school?

Reddy: Yes. You know, basically when you're 20 or 21 or 22, some-- you know, you don't know what you want to do in life, and usually what happens when you don't know what to do, you ended up going to—in a graduate school, do something else to see if you like that, you know, and that's what happened with me. I knew that whatever I was doing was interesting but I definitely did not want to be a civil engineer and building bridges or something, but I also didn't think I wanted to continue at IBM. I would've been quite happy being at IBM. There's nothing wrong with. It's just that if I wanted to understand the field better I should probably go to graduate school.

Hsu: Mm-hm. I think you've described sort of staying at IBM, going to work for IBM in India as sort of the road not taken. If you had gone that direction, what-- how do you think your career would have progressed, and what do you think you would be doing now?

Reddy: I have no idea. <laughs> I, if I went to IBM India, I would've probably been in commercial data processing, and it would've been a very different life, yeah. It was-- another way of saying it is I was not passionate about whatever I was doing. I was okay do-- you know, it's a job. You have to do it, and I did a reasonably good job of it, but that was not the passion.

Hsu: Hm. Well, that's a good-- that's sort of a good segue to, you know, what was it about A.I. that really grabbed you, that really made you go, "This is really what I wanted to do with my life"?

Reddy: Right. It seemed very esoteric, right. "I wonder if you can make computers intelligent?" And here are some people saying, "Yes, they can be intelligent," and Turing and Newell and Simon and McCarthy and Minsky, and I didn't know much about A.I. and-- at that time, other than having a vague idea that computers can be intelligent or play chess or something. But the issue was if I was going to do something different and new, this seemed like the most likely path which I might explore. Oh. One other thing I should say is while I was at IBM I was also continuing my studies with my advisor who introduced me to computers, Stan Hall, in University of New South Wales, and I was doing research on various, you know, engineering problems, structural engineering problems, complex, and so I was already-- even though I was working at IBM I was also doing research at that time, and so that kind of made me feel so if I was going to continue and do a PhD or some advanced degree I should probably do it in computer science or computing at that time, rather than engineering or civil engineering.

Hsu: Hm. So I guess could you share a story about a specific experience or maybe a particular person that influenced you to make the decision to go to graduate school to pursue A.I.?

Reddy: So one of the things a-- I'm not sure this made me go into computer science, but I was good at kind of coming up with interesting things to do by myself without anyone asking me to do it, right. So here is a computer. People would come and look at it and say, "What is-- what can I do? What can you do?" and so on, so I had to invent a little game and the simplest game you can think of is how fast are your reaction times, right, and so I made up a little program where it'll kind of say something and-- on the screen and then you have to hit the spacebar or touch any key and that told you how long it took you to react, and then you say, "Oh, your reaction time is 210 milliseconds," and then somebody else would come along and so you can have a little game among multiple people to see who's faster and agile. So those kinds of things, and the other thing at that time I did was because I was an applied science representative I was the technical support person for the astrophysics lab, which was looking at galaxies in the southern hemisphere at the Australian National University, so I would come across many people there who were kind of looking at the skies to see if they can find any galaxies. You must remember, at that time, astrophysics, they only knew there were three or four galaxies. Now we know there are billions of galaxies, and this was the beginning in the early '60s.

Hsu: Hm. And when you were applying to graduate schools, besides Stanford, how many schools did you apply to?

Reddy: I only applied to Stanford and Carnegie Mellon.

Hsu: Oh, wow.

Reddy: And... <laughs> Because that's where the people I knew were, right. Newell and Simon were at Carnegie Mellon, and I didn't apply to MIT for some reason, but I applied to Stanford and that probably because one of my cousins was already there finishing his PhD in semiconductor design, device design. So I knew of the place and so I applied there, and you must remember, in 1963 there were no computer science departments in-- they got set up, started around 1965, so... But, you know, there was-- John McCarthy was there and he was one of the founding fathers of A.I. and actually created the name. "Artificial intelligence" is attributed to him.

Hsu: Yes. So how did you ended up deciding to work with McCarthy at Stanford instead of Newell and Simon at Carnegie Mellon?

Reddy: No. They put me on waitlist or-- at Carnegie Mellon, <laughs> saying, "We can't give you a fellowship," or something, and Carnegie-- you know, Stanford also didn't give me a fellowship, but they said, "You're admitted, and you can join," so I went there, and that's all it was.

Hsu: Okay. <laughs> When you first arrived at Stanford, did you ever-- what was your experience as an Indian immigrant to Silicon Valley in the 1960s?

Reddy: Yeah. The Silicon Valley was, at that time, you know, apricot orchards, actually. There was no Silicon and there was no Valley. You know, there was a valley, but <laughs> it was mostly orchards. But Stanford was kind of a amazing place even then. You know, the Palm Drive and the dramatic quad and the buildings, and we ended up-- I know. The computer science department was in one corner. It was called Polya Hall, and-- but I had no clue of what to expect, right? I just ended up there and so as a graduate school, it was-- at that time we did not have a major computer, we-- they had a computer center and that was using not IBM computers that, you know, that I knew, like 7090, but they had a Burroughs 5500 or something, and that was fine and all that you did was wrote programs in FORTRAN or ALGOL or one of the languages, and-- but fortunately as soon as I-- no sooner than I landed there, John McCarthy got a DARPA—(ARPA [at that time]) grant to do research in A.I., and one of the things the grant funded was to buy our own dedicated computer for the A.I. labs. It was a PDP-1. They were just coming out at that time, 1963, and so there're number of people like Steve Russell, who did the Spacewar! and a whole bunch of other people who had come with McCarthy to-- from MIT to Stanford, and so all of them would work from morning, you know, nine to five or, you know, some hours, and after that the computer would be completely empty. Nobody else would be using it. You know, I was used to being-- because, you know, when you're working at IBM, computer time is very expensive, a thousand dollars an hour or something, even at that time, and so you used the computers whenever you could get your hands on it. You know, for example the story goes that Art Samuels, who did the first checkers playing program, who used to work at IBM, essentially would use his program to debug to see if the computers are working on the pipeline of the computer manufacturing. <laughs> So it turns out, you know, access to computers was at a premium and so here I was at 1963, and I had access to the whole PDP-1 with no one to bother me from about 5:00, 6:00 P.M. in the evening until 7:00, 8:00 A.M. in the morning. So I would work every night, seven days a week, all night. You know, unlike others who probably had other social obligations and social engagements and so on, I had nothing. I was just-- and so I worked, you know, 7 hours a day, maybe 14 hours every day, and so, you know, and there, you would not do it unless you were passionate about something, and that's what it is the thing. It becomes a challenge when you're programming the computer and it doesn't work and, you know, why it didn't work. You need to debug it and you keep working on it.

Hsu: Yes.

<00:38:49 off-topic conversation>

Hsu: Great. Did you ever experience any discrimination coming to the U.S.?

Reddy: No. No, that's very interesting, you know, basically unlike newspaper stories you read today and at that time, there was zero discrimination, either in Australia or at Stanford or in USA, probably because I was mostly around highly educated people in intellectual environments, and at that time, by that time, you know, once they know you and accept you, then the idea of color or race goes away. You know, it seems to, so I have never had any discrimination of any type even today, up to today, and I consider myself lucky that way, and on the other hand, I don't go out looking for trouble either. You know, I almost never go to strange places and where you might be mugged or something, and so mostly I think it's a question of what it is that you're inclined to do. If you're kind of curious and restless and need to go out and experience the world, then probably you might get into some trouble. In my case, I was all kind of-- all my restlessness was occupied by the computers. <laughs> That's, you know, so you just never had that problem. For example, we had a student whose father called me from India saying, "We're desperate."

You know, this guy ended up in a hospital, and apparently he went to some party in the strip district and then instead of walking back with somebody he kind of decided to walk back at two o'clock in the morning by himself and somewhere along the way he was mugged and his phones and other things were stolen. So these things happen, and I probably never do those, did those things, and so that's how I think you might say I was in a bubble. I was just in a sheltered environment. In that environment, there was no such problem. <pauses> So the-- my whole life experience is between Stanford, Palo Alto and Carnegie Mellon and Pittsburgh area, Oakland. That's it, <laughs> other than when you go on vacation. But even on vacation you only go to places like Niagara Falls or Grand Canyon and things like that, which are touristy places hopefully no different.

Hsu: Yes. Earlier you mentioned the PDP-1. So that leads to my next question. Can you talk about how you ended up going into speech recognition and understanding as a subfield?

Reddy: Yeah. Right. So it turned out, you know, in 1963 in the-- maybe it was September-October time frame, the computer was up and working and I was one of, you know, in the-- in John McCarthy's class and he said, "Look. Unlike other computers this one has an A-to-D converter, so that means you can actually speak and then when you speak it'll kind of-- you can record and play back," so-- and he said, "Why don't you see what you can do with PDP-1 and A-to-D converters?" and so that was my first introduction to digitization and using A-to-D converters and D-to-A converters, and so I think largely my getting into speech recognition and understanding was because at that time PDP-1 came with a built-in A-to-D and D-to-A converters. Most computers at that time, up to that point, did not have the speech input. So if you'd remember, when the PCs came, which was 20 years later, one of the first things that somebody did is they built a whole microphone interface, A-to-D and D-to-A converters, and that became a big business for them, and lots of people were buying it, and I used to say, "Why would anybody buy just an A-to-D converter?" Even the fact that you can just make it say something was interesting for the PCs that was 20 years later. But in 1963, that was the first time I think anyone had computers with the ability to have sensory input. For example, I remember as late as 2003 I was saying we have all these personal computers, and what they are missing, although we can have it, is integrated multimedia input, and now that we have them, we can do things like Zoom calls and so on. Even as late as 2003, none of the laptops and computers, PCs, would come with built-in sensory-data input, microphones and video cameras, and I remember I was saying create-- one of the things we created at that time was a thing called a PCTVT, which is not only a PC and a TV and a video camera and so on, and John Markoff picked up on it and created-- wrote an article for New York Times. At that time I was saying these are in our future. We're just absolutely going to have integrated multimedia input and output devices on all computers, and what I did not anticipate is we would have something like an iPhone, which not only would have all of those but would have a lot of other sensors like GPS and a whole bunch of other things, and now it's amazing. Most people take them for granted.

Hsu: Yes. Besides just the technical component of the A-to-D converter, what was it about your own experiences that may have led you to be interested in speech?

Reddy: Yeah. So basically two parts. One is the PDP-1 not only had-- it's your first interactive computer, personal computer, and you could-- there was a display and a keyboard, and you could actually program it and throw the programs and things on a little micro-tape, DECtape, and so the PDP-1 came with things that were unique, that were never part of conventional computer architecture, and that-- they came, because the origins of PDP-1 are from Lincoln Labs, where TX-2 and LINC and other computers were built. They were already experimenting with these kinds of other I/O devices than just conventional data in and data out, and the second thing is my understanding [a] little bit about the structure of language, it turns out if you learn any Indian language or Indo-European language, they are based on Sanskrit. Sanskrit is the first language, where everything is based on the sound structure, your alphabet. You learn thing where-- all the different things you can say, where your articulators are more or less in the same position like ka-kha-ga-gha-na or pa-pha-- so depending on which articulator-- so the way the alphabet is organized is very structured and is based on deep understanding of acoustical [events]. I didn't understand it, but Panini, the Sanskrit grammarian, understood it, and the reason they needed to do it was at that time, all the literature and the scriptures were all communicated orally, and that means there was nothing written, and so whenever they communicated orally, it was very important they articulate it exactly the way it is intended. So that meant they had to teach them different sounds and way they [sound] and how to articulate those sounds, and so that's-- so I thought I could actually build a alphabet recognizer, and the recognizer I tried to build was not English alphabet, but simply a vowel recognizer. But except in Sanskrit, every vowel has a short form and a long form, and that's how I was trying to see if a computer could actually recognize all the vowels of Sanskrit language. I wasn't thinking about Sanskrit; I was thinking about Telugu, which is essentially a derivative of Sanskrit.

Hsu: So then you could say that your multilingual background was a direct influence.

Reddy: Yes, multilingual background and understanding of the structure and the [pronunciation]. so basically I had to create a task and demonstrate the computer can actually recognize something, right, when you speak, and the task I chose was vowel recognizer. If I could only recognize vowels reliably, and-- which I did, right, and so it turned out the vowel recognizer was working very well, and I had an interesting experience, and-- which continues to haunt people that build recognition systems even today. That is when you say the right thing, it will recognize you, but supposing you don't say the things. It still has to classify it as 1 of 10 things or 15 things, and supposing it's not any one of those, it still will pick one of them, and so it does not-- at that point in time, we did not know how to build a classifier which says none of the above; it is not one of them. We probably could have built it, but we didn't. So, Ivan Sutherland, who did the early work on TX-2 on Sketchpad and so on, moved on to ARPA and was one of the program managers. So he was coming to visit the Stanford AI labs in '63, so John McCarthy brought him to the lab, and I was working there and said, "Oh, Raj is working on this speech-recognition thing." So I said, "Yeah, say any vowel, and it will kind of recognize it," and it did, and so Ivan said "ah," and it recognized you said "ah." Ivan said "ee," and then he whistled into it, and he said and PDP-1 said, "You said "ee"," and we all had a good laugh, and I was kind of chagrined, and so you learn immediately the limitations of trying to build many of these AI systems. They do not have the common sense and the capabilities human beings develop over a lifetime, 20 years or 50 years or something, and they do not have the computational capabilities of humans who happen to have 10 to the 10 computational elements in neurons and 10 to the 14 connections. So that's the difference. Even in all of '60s and '70s, we could not even get 1 image into the computer, because scanning and digitizing at the rate of 30 hertz-- an image every 30th of a second-- turned out to be very hard. There's not enough bandwidth to get the data in. So you kind of do random sampling and do multiple sampling to get the data, and that was okay. You could do that, but the image you're scanning had to be stationary. It could not move, and so there was a whole bunch of technical issues, which stopped us from doing much of the speech and vision and robotics work until around the turn of the century, actually. Computers are not fast enough.

Hsu: Yes, and so you graduated at Stanford what year?

Reddy: 1966. I was there as a graduate student from '63 to '66, and the Stanford became a computer science department, the program I was in, computer science program, in '65, and a year and a half later I was done with my teacher's research and finished Ph.D. in 1966, and I was immediately hired at Stanford as an assistant professor, mainly because at that time there were no computer science Ph.Ds., and if you wanted a computer science department with faculty who understood the field, had at least an understanding of the field, then you needed to hire those people. So two of us graduated that first year, Bill McKeeman and I, and so that's where we are.

Hsu: Yes, and so you were faculty at Stanford from '66 to '69, is that correct?

Reddy: Yeah. I was an assistant professor there, and they wanted me to stay on, and, in fact, unanimously recommended to the dean that I should be promoted. But it turns out Stanford has a rule. They don't want their own Ph.Ds. to stay on and never experience anything else. So the dean said, "Look, I already made an exception for Raj when you wanted to hire him. Now he should just go somewhere for a year, and then he can come back," and that's how I ended up leaving Stanford and going to Carnegie Mellon in 1969, and I could've gone back several times, in fact, to Stanford, but it turned out I was having a lot of fun at CMU, and there was no need to move.

Hsu: So when you had to go elsewhere, you applied to CMU and also to Berkeley. Is that correct?

Reddy: Yeah. I applied to CMU, Berkeley, and I think Irvine, maybe, UC Irvine, and I didn't follow up on Irvine, because I was not interested, and both Berkeley and CMU made me an offer, and it so happened, since Newell and Simon were at Carnegie Mellon, it was natural I would kind of go to Carnegie Mellon.

Hsu: Yes. Could you sort of briefly outline the major milestones of your career at Carnegie Mellon from '69 to now?

Reddy: Yeah. There are the technical milestones and administrative milestones. Let me first talk about technical milestones. So it turns out in '69 when I went there, ARPA, Advanced Research-- it was not called DARPA at that time-- was trying to decide whether they should have a major program in speech recognition, and a number of us made representations, me and Art Samuels, a number of people from MIT and Stanford and here, and so what ARPA said was-- or at that time it was Larry Roberts, I think, kind of said, "Tell us what it is that you think we can do in five years if we had significant funding," and in order to do that, they asked Allen Newell, who was one of the leaders in the field, obviously, at that time, to chair a committee to produce a report on setting up breakthrough goals, five-year research goals, that might be attempted by funding from DARPA. So ARPA program, it was called Speech Understanding program, SUR program, Speech Understanding Research program, got started in ’71 and went on until ’76, and there were five different groups that were working on the problem. There was MIT, CMU, Lincoln Labs, BBN and SRI, and what was unique about these programs compared to other kind of research which are kind of a curiosity-driven research-- here it was goal-directed. The idea was you have to demonstrate a system in five years which would understand connected, spoken speech, recognize at least 1,000 words and perform some task of understanding the request or query with better than 90 percent accuracy or 95 percent accuracy. I don’t remember exactly what, but the report is online. You can find it, but the Speech Understanding Report became the bible in 1971, and ARPA agreed to that and funded these four, five centers. The main difference is they were all working to demonstrate this one particular target, whereas in most other places, maybe one group is working on something, and somebody else is doing a different thing. They’re not all working on the same thing. That was a unique change in the research methodology or research paradigm, saying we have this goal-driven research and then say, let us see if anybody can deliver on that goal, and so there was a report that was produced at the end of ’76, which kind of showed that almost all the groups achieved the goal of understanding speech at a certain accuracy, but the main difference was the speed at which it took them. Given the computers of the day, some systems took 1,000 times real time, or maybe you speak a sentence. You come back an hour later, and then finally it says you said, “How are you?” or something, and so then you suddenly realize human beings are amazing. No sooner than you finish saying the sentence, they not only recognize it, but they’re able to then immediately respond to it, understand and respond. So the issue is, how do we get computers to that stage of performance, and we did not quite make it, except there were two systems at Carnegie Mellon that were the fastest of all of them and equally accurate, and the Harpy system actually was working in real time, and it would recognize and produce an answer immediately, immediately means maybe within a few seconds, still not instantaneous like 200 milliseconds, which human beings can do, and so what happened then was there’s a video that was created on the Harpy system and the Hearsay system that are on YouTube. You’ll be able to find them, and these systems actually were able to demonstrate the thing, and at the beginning of the Harpy movie, Allen Newell comes on the screen and says in 1971, ARPA said, “Do this, this and this,” and five years later there’s at least more than one system that was able to demonstrate that performance, and that became a benchmark-driven performance, became a paradigm for ARPA since then. For example, all the autonomous vehicle trials are good examples, where they said, “We want anyone that can demonstrate a car can drive itself can join, but it must be able to do the following, XYZ.” Both the off-road navigation and then urban navigation, those were the two systems that were a grand challenge and urban challenge that were done I think in 2006 and 2008 or ’11 or ’09, and they all came out of ARPA research. They came out of early work that we did at robotics in Carnegie Mellon in 1980s. We had a working system in 10 years. ’95, we demonstrated No Hands Across America, where a car would drive from here to San Diego. But there had to be a driver in the driver’s seat, and he may not have a hand on the steering wheel, but it wouldn’t know how to handle many of the exceptions, like, “What do you do when you come to an intersection of the road?” kind of thing. So going back, then, we fortunately have a clear record of the Speech Understanding project accomplishments in 1976, and we continued on that, and then we also-- ARPA at that time started image understanding research, and we did a number of things there, except, it turned out, image understanding was 100,000 times-- or maybe 1,000 times more harder than speech, because it needed that much more computing power, which we did not have at that time, and so everything we tried would barely work, and then it took until 2010, when they could demonstrate high-quality recognition using ImageNet database, and so they all needed-- almost all AI systems needed a billion times more computing power than we had when we started working on it. You could certainly play chess, and you could certainly play-- prove theorems, which were the two original goals of AI, because everyone thought that's what intelligence was. But intelligence was more than that, right? Human beings do lots of things. We learn from experience. We use a lot of knowledge. We respond in real time. We use language, and we use symbols and abstractions. We still don't have any system that does all those things.

So that brings me to 1980s, and that's when we started the Robotics Institute at Carnegie Mellon, and-- in 1979. So the main things we were trying to do at that time were systems, autonomous systems, not just cars that drive themselves, autonomous land vehicles, autonomous air vehicles and autonomous sea vehicles, and, for example, the autonomous helicopter that was built by Takeo Kanade in the Robotics Institute was actually used for reconnaissance of the 9/11 crash in Johnstown to see-- because people didn't know what happened and what the situation was, and so it so happened we could actually take the autonomous helicopter to do a reconnaissance and survey and transmit the images back. So in the '80s we mainly did robotics. In the '90s-- we started in mid-'80s a language-- Center for Machine Translation, and we got converted to Language Technology Institute in '93, and Jaime Carbonell, one of my colleagues, was in charge of it. By that time, I was the Dean of the School of Computer Science, but the main [task] in the '90s we were working on whether we can understand language, and, again, if you're looking-- look back today, look forward to today, there is a language-modeling system called GPT-3, which apparently has 175 billion connections or something. The numbers are not that meaningful other than to understand the magnitude of the task, and it is able to predict and also write articles in a natural language. But more interesting thing that happened was we used to think language translation is impossible, and so when we were working on the Center for Machine Translation in '86, when we set it up, we knew it's a very hard problem. We didn't know how much progress we could make, and-- but when Google, using completely different approach-- up to that point we were mainly using knowledge-based systems, tried to encode the knowledge about the language, and they said, "We don't know how to encode the language. We don't even have a model of the language. We're going to just use statistical models," and that's what characterizes most of AI today. Most of the knowledge that's used in AI systems is data-driven, statistically derived knowledge, which seems to work pretty well, or at least it's working 30 percent better than the earlier systems we could build by handcrafting the knowledge, and so these systems of today, statistical, data-driven systems, have their beginnings in the mid-'80s and the understanding of-- it turned out language technology is not just translation. It's also information retrieval and search, and it's also summarization. It's also finding entities in a text, entity detection. So there are lots of different aspects of language that we take for granted, and so there was research going on in all of that, and now all of that is happening purely statistically, using neural-network-based systems, and they are much better in the sense-- they are at least 30 percent better and make fewer errors than systems that we were able to manually handcraft. The main difference is by that time, we already have had a lot of experience on what the issues are and what needs to be done, and so in the '90s we set up a Center for Automated Learning and Discovery, CALD, which was transformed to Department of Machine Learning, headed by Tom Mitchell, and so those are the main events that happened in my life, where I kind of went administratively from being the head of robotics in '79 to the Dean of School of Computing in '91 and where I stepped down from that in '99, and since then I'm just a university professor.

Hsu: Okay. So those are the main administrative milestones you just mentioned? Yes, okay, and so what's your main area of work now?

Reddy: For me, my main area of work is-- I've kind of been looking at saying it's all well and good to make all these advances in AI, but who cares? How can we help humanity in the problems that they face on a day-to-day basis? So much of the work in the last 15 years or 20 years has been technology in service of society. So the early work in there, we started with work on digital libraries. Can you actually get all the books online? The idea was all authored works online, accessible anytime, anywhere, for anyone, and free. It's not just text and books; it's also music. It's also videos. It's also paintings and everything else, and we-- with funding from NSF on digital libraries, we made substantial progress. But you ran into the copyright problem. It turns out a lot of the books that are interesting are still in copyright, and, however, many of them are what we call orphan works, and so there's a huge discussion. Brewster Kahle, one of the people that wrote the early search engines, has been able to continue to make progress at this time, and so even Google, which started scanning the books in five or six years after we did, had to stop because of the lawsuits and saying, "You can use it for fair use, but you cannot make it available randomly to public." So at some point, Google has kind of stopped or downgraded their scanning activity, mainly because they may be violating copyright law. The copyright law as we now have it is very difficult to implement. For example, most books have a copyright of 50 years after the death of the author. First of all, there is no way to find out whether the author is alive or dead and nowhere to ask for permission. So most of those books, even though they may be in copyright, are not making any money for the author, and they will kind of die and will be lost to society forever unless it happens to be a very popular book like Shakespeare or Austen or whatever, and those will still be around. But at this point, about a million books are published every year. Only about 50,000 of that see the light of day in that they're sold more than the minimum number of copies. Out of that, maybe only less than 1,000 are what you would call bestsellers, where they sell more than a million copies, and so it has to-- but if you're a bestseller and you want to have access to it, the copyright gets in the way, and then at some point there's an ambiguity between the publishers and the authors on the copyright. So there's a huge set of problems that still need to be resolved and society have to work out. So one of the first things we did was digital libraries.

Then the other two things I became interested in, mainly because of my work in language and AI, is what you would call a literacy divide and language divide. There are about 30 or 40 percent of the population who are either illiterate or semiliterate. By that I mean they are semiliterate in that they may know how to read the letters of the language. They might even recognize some words, but they can't read a whole paragraph in English or any language that they know and understand what the paragraph is saying, and these are people that are semiliterate, and about two billion, maybe, two to three billion people in the world, are either illiterate or semiliterate. So what we are trying to see is, how can technology help an illiterate person? Right now these three billion people are left out from benefiting from the use of our technology, IT information technology and smartphones. So if you own a smartphone, if you can't read English, you are not able to do most of the things. But you can learn enough to recognize icon-- pictures to make-- phone your grandson or something, a few things you can do, but you are kind of substantially limited; however, if you can speak to the smartphone and if the smartphone can speak back to you and carry on a dialogue, you can do a lot of things. You can do online shopping. It turns out if you live in a village, the number of products you have available are probably less than 1,000. On a Amazon website, online shopping, you have access to over a million products, and if you go to in-between-- supermarkets have 20,000, 50,000 products, and huge department stores are-- Wal-Mart and so on might have 100,000 or 200,000 different items. So the flexibility and the cost available to people that are literate is significantly higher than if you are an illiterate person. You can break that barrier, the literacy barrier, and then the second one is the language barrier, language divide, it's called. The language divide comes from the fact most of the things on the Internet are in English, and if you are born and educated and if you are literate, even if you're born in India, if you knew English, you can probably benefit from the Internet. But only five percent of the people in India can use it, and in countries like China it used to be even less, but now they're better, and-- but if you are in Europe or USA, you're able to benefit from and access the facilities on the Internet. So the most important one, from my point of view, that everyone needs access to is the encyclopedia, the Wikipedia, and basically it gives access to the knowledge of the world. Anything you want to know, if you can go to Wikipedia and ask-- search and find the entry, you can learn about any field, anything you want, and invariably, I end up going to Wikipedia at least once a day, mainly because I want to-- I hear some phrase that I didn't know or person or thing that I didn't know, and that's the best way to find. Wikipedia has most reasonable explanations of everything, and so-- but only people with-- who are English-literate are able to benefit from the Wikipedia with seven million or eight million articles. The rest of the people also-- there are other-language Wikipedias. For example, none of the Indian-language Wikipedias have more than one to two percent of the size of the English language. Whereas in English is like 8 million articles, the Hindi is only 120,000, and Telugu is only 60 or 70,000, and so that means that they are being deprived of access to the knowledge, and so by breaking the language barrier, I think we can actually have-- benefit 3, 4, 5 billion people who cannot read and write English fluently, and those are the people now I think we are on the verge of being able to empower. They will be able to read any book or watch any movie or read any website, because you can get immediate, instantaneous translation. So, for example, every movie in English can be watched by anybody in the world, because they can get immediate-- the translation, dubbed translation of that into the local language? How? Because we have already demonstrated technology that Google or Microsoft-- everybody has, where you can go English to-- speech to the text, and the text can be then translated. Then the translated text can be then synthesized, and that can be played out immediately, and so this demonstrated where-- I think in 2012 Rick Rashid gave a talk in English in Beijing, and it was immediately recognized, translated to Chinese and synthesized and played out, and the Chinese-- that whole demonstration worked reasonably well. But today, in order to go from any language to any language, you need a large amount of data in both languages. You need about 100,000 hours of speech and about 100 million words of text to be able to kind of go speech-to-speech translation of any movie, and then you also need computers that have about a billion times more computing power. All of that, we are on the verge. Whether it happens in 5 years or 10 years or 20 years depends on the resources made available, because the technology is there, and we understand how to do it. So that's where we are currently working on. So the language barrier and the literacy barrier are two things I think we are almost on the verge of that would transform the world. Instead of two billion people who can only read and write English, maybe everybody can benefit from-- if we break the language barrier, and if you break the literacy barrier-- so that all that you're doing is only speaking to the computers. The great thing about human beings is they talk to each other all the time, and it's spontaneous unrehearsed speech, and everybody on the planet speaks and hears and can think and follow instructions that are in spoken language. But if you can't read and understand, then you can't do many jobs today, and that can be overcome. Anybody can be taught to do any job, and-- anytime.

<01:27:49 off-topic conversation>

Reddy: So that-- those are the two things in-- literacy divide and language divide. The other main thing I've been working on is kind of motivated by my own background coming from a little village. It turns out if you're born in a rural environment to parents who never went to college and never went even to finish school, then they're at a significant disadvantage. Not only do they not have anyone to guide them or direct them, but their schools are not that good, their teachers that they have access to are not that great, and furthermore, most of the testing is done based on national tests like SAT and TOEFL, and that means people that have better education will get into places like Harvard and Stanford, and even though you may be very capable, right, that is a friend of mine used to say, "We reward people that are better educated, not the best people, but best educated people." So, you know, that means if you don't have a level playing field and you don't have access to education, then you're, no matter how smart you might be, you will never reach your potential. So the question is how do we arrange for people? And both California and Texas have done something very interesting, and if you may remember in California in 2001 or 2, the Supreme Court of California said affirmative action is unconstitutional. You have to pick the best people, not based on-- not the people based on some color or ethnicity. So the problem is if you don't get-- have a level playing field and you don't have the equal opportunity, then a group of people will never get to become the best people and never get the opportunity to go to college. So what-- to their credit, what the Regents of UC system did around I think 2003 or 4, said, "We're not going to use SAT only. If you happen to be the best student in your school, the top four percent of the students in your school, you still take the SAT test, but you're not compared against somebody going to Brooklyn Polytechnic or, you know, Crystal Springs School in Silicon Valley, but, you know, if you're the top, among the top four percent of your school, we'll admit you," and so it turned out that experiment was extremely successful. Now, I believe I was told, that they've increased the number of people that are automatically admitted to nine percent, you know, and there's a institute at UCLA which studies performance in the tertiary, in a college student performance, and they've found these students that are admitted from these schools have a lot of problems. They still don't have enough money, they still have to work and they have all kinds of other limitations. The most important problem is their education was not as good as the ones, people, other students, in the same class. So for about six months to a year, they actually struggle, and then they catch up and then they seem to be well, and so you have a four-year education performance and the overall, the performance of the students that get in through the special system, is as good as, the average distribution is almost the same, as the rest of the cohort. So translate that to India and where I come from, right. So it turns out instead of 15 percent of the people being below poverty level in the U.S., there is more like 60 or 70 percent of the people are below poverty level and their kids cannot go to college. They can't afford, and not only that, they, even if they can, they cannot compete in the nationwide competitive exams. So talking to the head of the state from where I come from, I said, "Look. I come from a village, you come from a village, somehow we escaped, but a lot of very good people that are not getting the opportunity." So to his credit, he immediately seized on the opportunity and we set up schools, residential schools, for the top one percent of the students in the state, which are-- everything is free. The tuition is free, the accommodation is free, and the food is free, the clothing is free, and the shoes and everything else is free, and it turns out women, girl students, don't even go to college because their parents can't afford to send the children because of the safety and so on, so they stop after their high school or elementary school, whenever the school stops in the village, and therefore, you know, only about 20 percent of the women go to college, or maybe less. In this particular experiment we did, because we were admitting the students that were the best students after the 10th class, we found more than half of the students that were admitted were girls, because at that point they were still part of the educational system, did not drop out, and so this is a unique experiment, result, you know, unexpected, unanticipated consequence of this admissions program. So 53 percent of the girls in this college where we're admitting less than one percent of the top students in the whole state, and everything is free for them, gifted rural youth, and that schools are going and they're functioning. They get free laptops and free access to tools and other things, so the main problem is there's still a lot of structural, functional problems they're working through, and it's now 10 years. It was start-- they were started about 2008, and it's going well, and so my hope is we can provide a level playing field for a much larger part of the population and create models where they can all be educated and given equal opportunity to succeed. So those are the kinds of things I've been trying to do. Not always successful, <laughs> because there are a lot of social and structural problems that take a long time to overcome.

Hsu: Yes. Great. I guess I'm-- yeah, that's a really, really fascinating goal. I'm wondering how do we increase access to knowledge and information for the entire world while also making sure that people aren't misled by misinformation?

Reddy: Right. So the-- let's take the first one. How do we provide access to information? First thing is, to change the current structure of the education. Right now, until COVID, the students in the primary schools did not have access to information technology. They did not use it. So what we need is instead of three R's, we need four R's, reading, writing, 'rithmetic and real-time access to information, <laughs> and the last one, now, finding, you know, the-- what many of us do, I'm sure you do and I do, is every time-- any time I hear something that I didn't understand, immediately I type in or speak in and I get a-- very quickly understand what that phrase is and go on and go forward, and so I think we need to change the curriculum to make digital literacy an understanding of how to beneficially use all the resources on the internet to every student, every primary school student, and this will also help a lot in countries like India and Africa, because many of them don't go past fifth grade. If at least they knew they can find information they want, and they had the basic stills, because most of them will have access to the-- a smartphone or can find one in the neighborhood or in the home, and so anyone can find out what they want, but they must have been taught to know they can do it, they-- and so they-- I think we need to work on that problem of digital literacy at the second, third, fourth-grade levels, and give them-- the way I would do it is we spend a lot of time in arithmetic and long division and long multiplication and that may not be as important in the future because we can get those long multiplications using a calculator or smartphone, and we should reallocate the time so that at least one quarter, or if not more of the time in the first five grades, or maybe if you add kindergarten and pre-kindergarten, first six or seven grades, where you're lear-- having access to digital literacy is a most important thing.

Now the question is how do you deal with fake news and misinformation and how do you kind of-- two things. One, I think, I'm very hopeful in the next 5 to 10 to 20 years, we will have A.I. assistants, which given a statement, will come back and give you a probability of how accurate, whether it is true or not, and in order to do that, they need to be able to access all the information globally of everything that everybody knows and then filter it in such a way so that you can actually determine whether a particular statement is true or not. So we can see this currently in the trial of Derek Chauvin going on, right? Basically the prosecution presented their case and now we have the defense, saying, “That’s not always true. Here are all the other possibilities,” and now you need to come up with a probabilistic statement saying, “Given all the parameters we know, what is the likelihood any given statement is either correct or not?” I believe computers will be able to do that, and if we can do that--and why do I say that? If you remember what happened with Watson and Jeopardy competition, in order to answer random questions in Jeopardy questions, and in less than, what, 200 milliseconds, by pressing the button, the computer has to-- had to kind of combine knowledge from every source it could find and then come up with a probabilistic decision saying, “This is the most likely answer,” right, and so I think that’s what’ll happen. You know, maybe what IBM should do is instead of spending a lot of time on medical Watson, they should do a-- kind of build a Watson for fake news and misinformation.

So coming back, I believe in the meantime, there’s also the problem of critical reasoning and critical thinking capabilities human beings need to be given, saying, “When somebody comes and says, ‘Whatever is being said about me is not right, and here’s the truth,’ you need to somehow be able to absorb all of that and then make a decision.” So it turns out a large part of the information we get, it is irrelevant whether it’s right or wrong. You know, other than one or two people being canceled, as they say, or ostracized, because now they’re seen as a pervert or something else, the rest of the world doesn’t care, you know, and so most of the information of the kind that you see in the news, ultimately, even after three, four, five years, won’t make any difference, and then the only question is, “What is the truth? How do we know the truth?” and unfortunately there is-- that problem may be undecidable. You know about the Gödel’s and undecidability and unsolvability of Turing. There are certain problems, “I always lie,” or something like that, a statement where whether it’s true or not, we don’t know. We cannot-- it cannot be uniquely determined. The same is true with some fake news, and the-- ultimately what happens is some individual or persons or-- actually are impacted by it. Maybe they lose their life, and that’s the world we are living in, and unfortunately... If we can kind of create individuals who have critical thinking capabilities and provide them with an A.I. assistant, which can also give them some confirmation of various statements, we’ll be most of the time right. We may be wrong, because a lot of circumstantial evidence can cause people to make a decision which ultimately may be wrong. So I think the fake news problem will be with us, and I think our technology can actually provide some beneficial solutions to it.

Hsu: Great. Thank you. Let’s move on to talking about mentorship. So who-- which mentors have been most influential in setting the foundation for your work, influencing your work or directly shepherding your research in your career?

Reddy: Yeah. So in universities, my advisor at University of New South Wales, Stan Hall, without his initiative I would not be in the computing field at all, and so he has been very supportive, and at Stanford John McCarthy and Ed Feigenbaum. Both of them, they’re very, you know, helpful and I would not say they were mentoring me. They may be mentoring me broadly by providing me support in various things, but then in the area that I was working on was unique that they didn’t know any <laughs> more than I did.

Probably they knew less than I did. In that sense, they could not specifically guide me in my area, and that's when I was, you know, sometimes thinking I would've made lot more progress if I had a mentor who had kind of said the right sentence, the right word, saying, "Do X, Y, Z." You know, for example, in speech research, we extract parameters of the speech, right, and for a long time we extracted formats and various other kinds of features and later on there was something called LPC and mel cepstral coefficients, and so that evolution took like 20 years for us to discover. There are certain representations that are more efficient and more accurate than the others, and if I were-- instead of <inaudible 01:47:39> at Stanford, let's say I was at Bell Labs, where there were many more people who were knowledgeable in speech, they-- I might have actually made the right decisions. The same is true in statistics. For example, I got A+ in Emanuel Parzen's course in Statistics, but I never understood how to apply the knowledge in statistics I had learned-- even though it was an A+-- in my real problems in speech. I said, "Oh, that has nothing to do with statistics. Speech is--" you know, as though it is completely different. So those kinds of things sometimes, but the main mentors for me have been McCarthy, Feigenbaum at Stanford, Newell and Simon here at Carnegie Mellon, and many other colleagues and students and mentors. It's almost an [axiom]-- you learn as much from students as you learn from senior people sometimes, but you must be open and willing and to assume that they know as much or better. The best example of that is Jim Baker, who was one of my PhD students, went on to create Dragon Systems, which was bought by Lernout & Hauspie, which was then bought by Nuance. In the last week, Nuance was bought by Microsoft for $16 billion. But Jim Baker didn't benefit from it, but the technology he invented, the Hidden Markov Model of technology for learning of features of speech, is still used. It is the main standby for almost every speech system that anybody has used, and when he first came and said, <laughs> "Raj, I want do to X--", I said, "I don't think it's going to work, Jim, but if you want to do it, go ahead and try it," <laughs> you know. So in that sense, you have to respect each other and especially your students, and the same thing happened with Kai-Fu Lee and what he wanted to do, so you find you must be open to ideas and diff-- with people that want to do something different than what you want [them] to do, and that has worked out well for me, basically, because no one person probably knows all the things. That's why people work in teams, and teams work well, teamwork, especially in places like Bell Labs when you read the discovery of the semiconductor, by Shockley, Bardeen and Brattain. It's very interesting story of how people with different skills had to come together to make that invention.

Hsu: Yeah. That's, you know, you bring up your students. You know, a number of your students have mentioned your specific style of mentoring and support, how you're willing to let them do their project and even give them support in funding, even when you disagree with their fundamental approach. How, I mean, that's-- is that unique? How important is that?

Reddy: I think it's very important. You know, anyone that thinks they know all the answers is obviously wrong, and they're missing-- But there're-- I think in the leadership skills, if you look at people like Lincoln and Franklin Roosevelt and, you know, they seem to know what they want to do but they also know that there are other points of view and you cannot simply force your rule.

Hsu: Yeah. Let's talk a little bit more about your students. So you already mentioned James Baker. Could you talk a bit more about James and Janet Baker and your connection to them?

Reddy: Yeah. Jim and Janet Baker were two of my PhD students, they're early PhD students from speech. They were already at Rockefeller University doing their PhD program, so they came one day and said, "There's nobody there. We're working on speech and there's nobody at Rockefeller that can do it, but we're afraid to come here and start all over again," and so I was-- I said, you know, "No problem. We'll give you credit for the situation [ph?]. Maybe you can start doing your PhD research." So those are the two people. You know, I went to our department head, Joe Traub at that time. I said, "Joe, I don't want these two guys to go through all the qualifying, taking the courses again and qualification exams and so on," and he said, "No problem," you know, so we worked out a scheme whereby they could complete the PhD at CMU, and so they transferred here, and that worked very well, and that's one of the great things about the CMU environment. Newell and Simon and Perlis kind of set the stage, you know, where it's-- they called it the reasonable person principle. They would empower people to do anything you want, you know, except you can also make mistakes and screw up, and they say, "If a reasonable person could've done that, that's okay," <laughs> and so I've been very fortunate to have had, well, absorbed those lessons from them by osmosis, so to speak, and I believe that has stood us in good stead at CMU. It still goes on. You know, basically we're really flexible. We don't kind of stand by rules and so on, and if something is good to do, we'll do it, you know.

Hsu: Yes. Could you discuss maybe Lee Erman, Rick Hayes-Roth and Victor Lesser?

Reddy: I'm sorry?

Hsu: Could you talk about mentoring Lee Erman, Rick Hayes-Roth and Victor Lesser?

Reddy: Yeah. So basically, Lee Erman was my student at CMU-- oh, no, Stanford. So when I transferred from there, he came along, and Victor Lesser was also a student at CMU but he was working with me, but by that time he was finished his PhD with Bill Miller, and so he also came with me when I came to CMU, and--

Hsu: Oh, so he was at Stanford also originally?

Reddy: Yeah. Lesser and Erman were there. Rick Hayes-Roth was at University of Michigan, and then he came one day and said, "I'm not making much progress at Michigan. Can I come and join your group?" and so it worked out, and the three of them were-- worked on the Hearsay system, and there's a video on Hearsay on the YouTube. If you can't find it, I'll send you a link. And that-- they worked on a different architecture, the so-called blackboard model, which we believe is the model of human brain, of how things work, where whenever we see a lot of stimulus, we try to interpret them, and they're not always easy to interpret, and the model we use is, you know, three or four experts speaking different languages, like French and German and so on, trying to work on a common problem, and so in speech you have people that know acoustics, the people that know linguistics, people that know grammar and people that know tasks, specific semantics and so on, so the issue was, "How do you build the system which can actually work with all these different experts representing different kinds of knowledge?" and the model we created for Hearsay system was what's called blackboard model. The blackboard model is essentially a metaphor, blackboard metaphor. You write something on the board. Everybody can see it, and somebody else comes along and writes something on it, and it's what's called a hypothesis and test model, and every time you see something on the blackboard and then everybody else, all the other knowledge sources that know something about that particular item or event or something, say, "Oh, if that is the case, I believe this is the case." So hypothesis and test on the blackboard where knowledge is widely available, usable, turns out to be. I believe there are some new models of cognitive machines or-- that Manuel Blum, another Turing Award winner, has been working on, which use the theater metaphor. The theater metaphor is another variation of blackboard model, where lots of people are sitting in the theater in the auditorium, and on the theater, some action is going on, and everybody's looking at it and internalizing it and then they make certain decisions based on that and then the whole... So this is a model of psychology and short-term memory and long-term memory and how all of them interact with external sensory data and of speech and vision and other things, and so there are some new advances coming along. I don't know what the implications might be, but the blackboard model that we used in Hearsay was in 1974-- '73, '74, was one of the first examples of an attempt to do that kind of thing. It comes from Newell and Simon were trying to understand how people solve problems, and they were using something called post-production, production systems, where condition-action rules fire and every time something fires the context changes and now you know something more, and that's one specific way of doing the blackboard model, so-called production systems. But in general, I think there are many different variations of that idea of how people see, hear, reason and act in reality, and it all seemed to be centered around the model of the short-term memory. Now, what is-- what are you paying attention to right now is the thing that seems to be crucial.

Hsu: Okay. Thank you. Could you talk about mentoring Bruce <pronounces> Loweer or is it <pronounces> Lowery or--

Reddy: Yeah, Bruce Lowerre <pronounces Lowery>, yeah, was the guy that did the HARPY system. Basically, he was like the lucky person, basically. He was in the right place at the right time. By that time, we had lot of work going on. We had a lot of speech data and other things, and Jim Baker came and did a thesis on Hidden Markov Models and built a graph system and demonstrated that if you do an exhaustive search of this graph you will come up with the right answer, when-- or the most probabilistic, highly, most likely answer, and there was something quite not right and reduces dynamic programming and so on. So whereas you can prove it's complete and it's actually correct and optimal, human beings are not optimal. They never can absorb all the information and make all the right decisions to find the optimal solution, and [Herbert] Simon calls it satisficing. You know, you settle for a good enough solution, not the best solution, and in fact, Herb Simon got a Nobel prize for that statement. His result was in economics, and where he studied human beings, how people make decisions, and he found that humans simply cannot handle more than a small number of facts. Maybe few tens or maybe hundreds. If you give them thousands of facts, they will simply ignore many of them and find a good enough solution, and you do that all the time when you do, you know, travel planning, you're driving from A to B downtown, you make multiple choices depending on the conditions, and you try to find that good enough solution to get there if the-- so this problem keeps occurring over and over again in computer science, the traveling salesman problem and so on, and so in the-- Bruce Lowerre's case, you know, he was my student after Jim Baker, and I said, "I think you can do this, get almost the same result, but only spending 10 percent of the work," you know, doing a tenth work, to find a good enough solution rather than the optimal solution."

So he came up with this model for what is called beam search where at each stage of the decision you say, “I have a hundred possible choices. I can follow all of them. I cannot follow all of them, and so what I’m going to do is only follow these 10 out of the hundred, the most probabilistic choices, and then follow them from there,” and it turns out that works very well. It’s called beam search. That’s what HARPY system used. Because of it’s only doing 10-- 1 percent to 10 percent of the work, it was able to almost work in real time. Almost all the other systems took much longer, and so there’s a background to that that happened at CMU in the Computer Science Department where in the next door we had Hans Berliner and Jim Gallogly kind of building what are called chess technology, Tech systems, Tech-- chess playing programs, called Tech, which followed Richard Greenblatt’s program at MIT, to play chess, and they were able to get much better performance by brute force search, except they had to figure out what paths not to follow, and so normal branching factor is like about 35 in chess, so you can’t look at all the paths. So you say, “Okay. These are already not promising and we’re going to follow only these,” and so we already understood the concept of, you know, best first search or something and-- but they’re not exactly the same. In chess, you know, you follow the best and then kind of try to figure out which path to follow. In speech, because you have to make the decision in real-time, you just go, only follow everything, but only the best paths, not only one, but maybe 10 or 20, and then when you come to the end you pick the best one, and it turns out most of the time it was correct, and that was Bruce Lowerre’s contribution for the HARPY system.

Hsu: Yeah, and I find it fascinating that in all of this, all around the same time you have Hearsay and HARPY and Dragon, all these three separate speech projects all going on at the same time. You know, what was unique about CMU that allowed three separate projects rather than having everybody work on the same thing?

Reddy: Yeah. Yeah, basically it is my probably personal research strategy saying, “I know enough to know that I don’t know all the answers,” right, and therefore when two or three different things were appropriate, you know, we had a similar system. When Kai-Fu Lee was doing Sphinx, we had another system that was Angel system, that was doing knowledge-based speech recognition. But at that-- at the time, in the ‘70s, it turned out Hearsay system was the main system, you know, which multiple people were working on. Maybe there were about 10 people, whereas Dragon System was just two, with Jim Baker and Janet Baker, and later on it was just one person, you know, Bruce Lowerre, working on the HARPY system, and all of them could proceed because, you know, and I didn’t know it at that time but we now have-- we have a A.I. law, or a principle that is based on, it’s what we call “when in doubt, sprout.” What happens is, in A.I., search is very ubiquitous. Everything is search, and so when you don’t know what to do you do three or four things and hope, you know, one of them turn out to be right, and this is a very commonly occurring theme in most A.I. systems, and the sound bite for that is “when in doubt, sprout,” right, <laughs> and so that’s what was happening here in real life and real research lab paradigm, trying to see, “If this person thinks he has a better idea, well, let’s go try it,” and so I think you’ll find most of the research in biology, same thing happens. There’s no one answer and people are kind of trying different things and when something right happens, then everybody picks up and goes there and start doing more work on that particular path.

Hsu: Thank you. Can you talk about mentoring Kai-Fu Lee? Can you talk about Kai-Fu Lee?

Raj Reddy: Yeah. So Kai-Fu Lee came I think in ’85 or ’87 or something, and the was here as a student and then we hired him as an assistant professor and a year later he went off to Apple where he did some interesting things in the early ‘90s, pre-Steve Jobs, and then from there he went to Silicon Graphics, I think, and then from there he came to Microsoft where he spent like 10 years. He was the head of Microsoft Labs there, and then ultimately ended up at Google and from there he’s now a venture capitalist. So that’s his trajectory or path, but when he was in-- he was always full of ideas and when he came as a student, he was already very creative and very smart and knew exactly what kinds of things he wanted to do, so I just said, “Okay. Go do whatever you want,” and so he kind of talked to various people and came back and said, “Everybody else is doing X--”. “I want to do Y--”, and what he wanted to do was kind of build a system that’s purely statistically based, and at CMU, all of us were kind of knowledge-based systems, A.I. types, so... But I said, “Okay. That seems like a reasonable thing to want to do. It may not work, but that’s okay. You know, you can come back and try something else.” So that-- and what he was trying to do, until that point, the rule of-- the common belief was the only way to get a speech recognition system to work is to make it speaker dependent, his master’s voice, to kind of have them speak only to you and not to anybody else, and-- no, not, you know, have the computer system recognize only your voice characteristics. Turns out, that’s much easier to do, get higher accuracy, than if you had to do speaker independent system. But at CMU, I’ve always felt that that speaker dependent characteristic thing is an anomaly. Because human beings, there are all kinds of people, and there are all kinds of accents, and still recognize them. And so the only question is, how do we build models that are speaker-independent models? So that required us to collect data from multiple speakers, rather than just one speaker, and kind of build a statistically-- a training system which kind of works based on all of them. So the Sphinx system was built on the Hidden Markov Model, except it was using statistical models that are multivariate Gaussian, as we call them, which had captured variability of speakers, up to some point. So the system that Kai-Fu Lee built, the Sphinx system, was speaker independent, the first speaker-independent system that we were able to demonstrate. And he was unique in that, and then from then onwards he went off and do a number of other things. But his early work on speech, speaker-independent systems, was the one I think most people remember him by.

Hsu: Yes, thank you. From your experience, what makes a good mentor, or leader, or role model?

Reddy: There’s no one good mentor or role model, I think they’re different. My best role model, I think, is Allen Newell. He was very good in the sense-- he was very empowering. If you had an idea and you want to go try it, he would say “Okay, I don’t know that we have the resources, but I’ll find it for you. Go try it.” And I remember a similar example from Bell Labs. One of my colleagues in speech is Bishnu Atal, and he was working on cepstral coding and needed supercomputer-like power, and apparently he used up all the Bell Labs computer budget for his acoustics department in one night <laughs>, or one day. And so there was a big human [ph?] cry. And so they went to the Director of Bell Labs, I think it was Ed David, or someone, and they said “Okay, we can’t let him have any more computer time because that computer’s needed for everybody else. But I have another budget here where it turns out I have lab equipment, and I’m going to buy him his own computer so he can use it all the <laughs> time. And we’ll call it ‘lab equipment,’” okay? So that is to me a role of a leader, of a mentor, who will find solutions to solve problems so that you can do the things you want to do. It’s easy to say “No.” A lot of people say “No.” And what you want in a mentor is someone who empowers you, doesn’t kill your enthusiasm on day-- as soon as you tell him by saying “No, it can’t be done.” And that kind of kills the whole thing. But if you can find a way of empowering your colleagues, then I think you may have right kind of a role model. Mentally, not everybody is like that. Some people will go beyond their horizons, and I think the phrase is “Think out of the box,” a willingness to think out of the box. And I think that probably would characterize a role model as a mentor; people that are willing to consider options that you are-- that are beyond your own-- your interest or experience.

Hsu: And what do you think makes a good partner or collaborator?

Reddy: Usually partners that accomplish great things, each one of them is an expert in one thing, and together they do things that no one of them could’ve done. And that was the experiment in invention of the semiconductor, William Shockley, and Bardeen and Brattain, and many other examples of that kind, where you’re kind of-- come together as a team where each of you bring an expertise and capabilities that the others don’t have. And so you divide up the thing. Sometimes, it turns out, we all know what to do, and we have to divide up the thing and do it, because there’s not enough time for one person to do it all. And we may all have the same talent, but then the team comes together and say “Okay, I’ll do this, you do this,” and it works out. So ultimately, it’s the project leader who kind of decides who should do what, and sometimes it’s not always right, sometimes it doesn’t always work. For example, when we were working on the autonomous car problem, there were multiple people working on the autonomous car problem, and the two main ones were Red Whittaker working on his own solution, and then Chuck Thorpe, and a number of others who are working on automated learning of the image features, and both of them were working on [it]. And so it looked like they were competing with each other, and they were, and so I had to kind of get involved and saying “Hey, do whatever you want to do, and don’t try to stop the other guy.” And nobody’s stopping you, and if you have the best ideas, you do it, and if the other person wants to try something else, that’s okay, too. He should simply accept it. And so these complications happen in labs, competing approaches sometimes think-- and you have written examples, again, in research histories of people with different approaches. If you read the making of the atomic bomb, there were lots of different approaches to fission and things we discovered in the labs, and they were all slightly different, and not all of them would work. But ultimately, once something works and it’s replicateable [sic], everyone else jumps on the bandwagon.

Hsu: Great. What makes a good student or a follower?

Reddy: What makes a good what?

Hsu: A good student or follower.

Reddy: I’m not sure what you want are followers. What you want are people, students, who are super bright, and who can take an idea and then do their own thing. I had a student from China, unfortunately he was only here for a year, he was the best example. He would come in the afternoon and saying “Okay, what do you want me to do?” I would tell him what I think he should do, and go away home, and then come back next morning, and he would be there waiting out in front of my door, saying “I’ve done it! Here’s the work. [ph?]” <laughs> So I didn’t know how exactly what he was going to do and when he was going to do it, but the fact he was able to get it done. And so that's what happens, usually you know-- you don't know exactly how it's to be done as a team leader. You think you may be able to do it, but you haven't done it, so you don't know, and you give the problem and point them in the direction, hopefully, and they have to figure out all the rest of it. And most people aren't like that. In research, that's what research is about; namely, solving problems without a teacher, or learning without a teacher.

Hsu: If you think about passing the baton to the next generation of people working in your field, what advice would you give them?

Reddy: That is usually, in my case, it's self-selection. Basically, people at some point start doing their own thing, and then you say "Okay, he's no longer working in my lab, and he's working on his own and doing his own thing," and that's fine. And Alex Waibel, one of our students, one of my students, is a good example of it. He's gone on to do a lot of interesting things, and after about five or ten years, he started doing things in Germany, because he was from Germany, and which was different than anything we were doing at CMU. And I said "That's fine. If that's what you want to do, we will--" the main difference is if people want to do something different than what you want to do, then they have to have their own funding. And so passing the baton is different in the sense of you're saying "I'm retiring, and here's the lab, and here's the funding, here are the people, now you can start doing the same thing." And that's true too, sometimes that happens. And usually, the passing the baton in my case happened when I had to step down as the Director of Robotics Institute and then appoint somebody else. And then I appointed Takeo Kanade, who was a brilliant scientist and an expert in vision, and he then appointed Chuck Thorpe, and on. So it turns out every Director usually kind of finds their successor, and that's what happens. So that's reasonable, and the other passing the baton happens when you step down, you need a new department head, then somebody even higher than you comes up and says "I'm going to set up a committee and the committee will review all the candidates, and then make a selection." And that works also, but not necessarily as better, because many times the person who can take over from whatever you're doing is a person who is already there and knows all the problems.

Hsu: Thank you. So next I'm going to-- a question about decoding technology. So CHM, our mission, we help to decode technology for everyone. So can you explain, or help decode your most important breakthrough, in a way that a high school student can understand?

Reddy: Yeah. So basically, the most important advance that happened in speech and image understanding technologies, speech especially, is the fact that we have computers that are a billion times more powerful than we had forty years ago. A billion. With that, we're able to use new approaches to understanding of signals, continuous signals. When you speak, all what you're seeing is changes in air pressure, and that is being digitized, sampled and digitized. And so a computer is seeing a whole sequence of numbers, and in order to interpret that, it has to have a model of how to interpret those things. And what we discovered is we can actually build models, extract features and do recognition, but it's almost never as good as human performance. So what we discovered since then, now, is if we had a computer that's a billion times more powerful, we can let it learn by itself. And that's where the neural network-based learning, that is more or less simulating what is happening in the human brain, where there's an auditory cortex which literally has ten-to-the-nine, ten-to-the-ten cells, and these are selforganizing themselves into synaptic clusters that are able to recognize individual phrases. And that's what is unique. Basically, when you hear the names of your parents or something, you immediately recognize it because that's something you've heard many times, and that's what the human brain does, it's able to recognize things instantaneously. To build systems that can do instantaneous recognition we need to have computers that are maybe more than a billion times more powerful in the future, which will discover the principles of organization without human beings explicitly having to code it. And so that is the main lesson from the AI research; namely, there are things in science and technology which can be completely understood, crisply understood, and programmed as algorithms. And there are other things that we don't fully understand, we may never fully understand, but which can nevertheless be solved. So, for example, multiplication and division. Originally, the number systems that were invented in Egypt and Babylonia four thousand years ago, maybe five thousand years ago. Until the first turn of the century, in the first millennium, around 100 AD or something, people could not-- did not have an algorithm. They could add and multiply because they were pretty smart, that was because they knew a lot of facts. But what they did not have is an algorithm, and they did not understand the concept of zero. And it took them five hundred years to invent the concept of zero, and then someone to come up with a step-by-step procedure that could be taught to primary school students on how to do addition, and multiplication, and division. And so the same thing happens; right now, many things that we're doing, like how you speak, and hear, and see, are all in the continuous domain. They're all continuous signals, not discrete numbers that we are adding and multiplying. And in order to interpret them as discrete symbols, you need a way of doing what we call signal-to-symbol transformation. And that transformation only happens if you have a model for classification of the signals. Right now, it is done by deep neural networks which are simulated network-based solutions. But in the future, there may be more compact networks that can perform as well, and we have to wait and see, and that's where the research goes. That's what we call explainable deep neural networks. Not only can they do what they're doing, but they can explain how they made that decision. Right now, most of the learning that's happening, they get the right answer, but they don't know why they get the right answer, how they get the right answer.

Hsu: Right, thank you. As you consider the challenges facing people and our planet, what new technological innovations are you most optimistic about making a positive difference, and why? What new innovations are you most optimistic about making a positive impact? Yeah, society and the world.

Reddy: Right. To me, those are all in this-- if you can change the lives of half the population of the world, right? Basically, three, or four, five billion people, then you've done something substantial. And they're all in this language divide, and literacy divide, and providing access to knowledge to anyone, anywhere, in any time, in any place, for free. Knowledge should be free, and that's where the copyright is-- problem is currently holding it back. And I have solutions for that. Basically, copyright problem, if you create an intellectual property, somebody should pay for it, and there should be a way the government comes up with a model saying the intellectual property becomes public. We're initially going to give you something, but we'll continuously monitor seeing how much impact it had, and over a period of five, ten, twenty years, we'll give you additional resources. And now it becomes part of the society. It's like building roads and bridges; you build roads and bridges because they're needed for society. And so you can't say "Sorry, you can't go on my road." You can, they are called toll roads, but most of the infrastructure is free. And so the way we do it is by kind of creating a tax system where we collect money from society as a whole to pay [for] the intellectual property. Rather than current system of people that create the intellectual property, which is not only books, but art, and paintings, and music, and movies, and everything. And so somehow, we can now keep track of every time-- if I make a movie, which is an Academy Award-winning movie, if I make it freely available to everyone in the world, lots of people will watch it. Every time somebody watches it, I can keep track of it. Not I, the systems can keep track of it. And then based on that, I can be rewarded by having created such a wonderful movie that everyone wants to see. Right now, that's not the way it works. You kind of have to go and market it yourself, and then intermediate people like Amazon get in, and Netflix get in, and they collect money, and then they pay you some. It may not always be fair because that's what movie studios do. They rob Paul to pay Peter because they're trying to kind of amortize across many pictures. So the person that creates the most blockbuster movie may not get all the benefit that you might get. And so I think there are ways, new models that can be created that would work, but what I'm saying is those are societal models, and we in computer technology carried on [ph?] saying it, we can't do much about it. But I believe what we are doing, and what will cause to happen, will fundamentally change the lives of people much more than it has. Right now, the technology we have created is primarily helping people who speak English language, and can use devices, have digital literacy and understanding. And if you don't have digital literacy, and don't have access to a device, and don't have access to internet connectivity, and then you don't have access to--

<pause in thought>

Reddy: The net effect is you cannot benefit from things. Only about two billion people can right now. I believe as English language is understood by other people, or translation systems are provided from any language to any language, maybe four, or four-and-a-half billion people will be empowered, instead of just two. The remaining two-and-a-half billion, three billion people, are illiterate or semi-illiterate, they can't even read or write. So you can only speak to them, just like the Amazon, Alexa agent, right, if you have-- where you're only speaking. There's no keyboard, no nothing. You say something, you carry on a dialogue, and it'll get done. If it didn't understand, it'll ask a clarification question. So those kinds of systems ultimately will empower the remaining three billion people, and it's just that they-- and the systems don't yet speak the hundreds of languages, and thousands of dialects that we have to handle. But that will happen within the next 10 years, 20 years, or so.

Hsu: Great, thank you. Next on-- the next question is about the perils of technology. What new technologies and their applications, and implications, are you most concerned about or worried about?

Reddy: It is the same issue; namely, every knowledge that we create can be used for good and can also be used for evil. And what the hackers are doing and-- is an example where they want to benefit from kind of a lack of knowledge or gullibility of people. So people say "AI is going to take over the world and we'll all be slaves for the robots." I don't think any of that will happen, but what is more likely to happen is human use of human technology, human-- or some humans becoming superhuman capabilities, and where the rest of the people don't have the technology, and they may in fact be disenfranchised, and we need to protect against that. And the way these technologies will come to be, it is not just AI, it's all technology. And if you have access to those technologies, biological technologies, and the biology and gene splicing may have as much impact on the future as information technology and computer technology. But all of them are currently understood and exploited by a small number of people, and the rest of the people are at their mercy, I think. However, we have governments of laws, and rules, and regulations, and if something like that happens, then it hopefully would come under antitrust of some sort, monopoly, unless the person gets so-- person or entity becomes so all-powerful that they can take over all the systems of the world. And that's not yet the case, but it could, and the best example I can think of is what happened in Iran, right? Supposing we create a technology which will simultaneously take down all the internet of the world, and simultaneously hack into all the energy power systems of the world, then you can essentially bring the whole world to standstill. There are books about it, and it has nothing to do with AI per se, but AI is part of that, all those advanced technologies which will be used to kind of cause harm, take over the world. And we need to have systems and people-- I believe most of the countries have different teams that do that, anti-hacking surveillance. So hopefully we'll be better than the people that somehow are able to take over this world. So we just need to assume that society as a whole-- and clearly there's no doubt, and, namely, even today we can do things that we could not do even 10 years ago, 20 years ago, let alone 100 years ago. And you can't roll the clock back, you can't take them back and say "We're going to live like Pennsylvania Dutch and not use any of the technological innovations." And at the same time, the problem, of course, is that societal systems change, and what is right 20 years ago may not be right today. And even Bill Clinton signed a bill that said same-sex marriage is illegal in 1997, or '8, and then 20 years later, he said "That's wrong." So basically, society as they change, some things which are wrong at one point may turn to be right.

Hsu: Thank you. So then how do we ensure that new technologies are inclusive, accessible, and ethical for the benefit of everyone, everywhere?

Reddy: Yeah, the most important thing is inclusive, and, namely, the good thing is our technology is getting cheaper and cheaper every day, and my solution was to provide all this technology for free to everybody on the planet. And the way I justified it before, I was saying this now for 10, 15 years, but in the last year, I think I'm being much more effective, and my answer is economic. Basically, the whole global GDP, as a result of COVID, has gone down by 15 to 20%. And if you take any one country, 10% of the GDP of the United States, which is 20 trillion dollars, right, is 2 trillion dollars. For roughly 1% of the 2 trillion dollars, which is 20 billion dollars, you can give every person a laptop, right? So let's say it's not 1%, let's say it's 5%, whatever it is, so for a small part of the cost of the loss to the economy, we can give everyone a free laptop, and free phone, and free internet connection. And it's only going to cost 5% of what we lost last year, and it's obvious to me that the government should do that. Right now, the main justification for infrastructure, spending a lot of money, is everybody when they drive their car through the potholes and so on, costs $600 because of-- per person, because of bad infrastructure. Why? The car repairs, and tires, and everything else costs that much per year. And so let's say if we spend that money, then hopefully we'll reduce that cost of-- 10% of the cost. And so "How do we provide inclusivity?" which is the main part of your question, is essentially for the government to decide this is an essential service. For better or worse, the society decided things like water, and electricity, and roads, are essential infrastructure, and we've had them for, what, a hundred years or more? And we need to decide that information technology, broadly, and access to smartphones, and internet, and power, are essential technology. We just have to kind of give everyone-– we decided that about education, every child should be educated, and they should be literate. All we're saying is they should be digitally literate, and then have access to these things. It'll happen, I think, as people discover, or realize, how essential all these things have become, I think we'll have an inclusive solution.

Hsu: Great, thank you. Next, I'm going to ask you about the award, the Fellows Award itself. What does it mean to you to join the other pioneers as a CHM Fellow?

Reddy: I'm sorry, what?

Hsu: So what does it mean for you to receive this award, and to join the other pioneers as a CHM Fellow?

Reddy: Right. I don't know, you should talk to the people-- <laughs> decided on what aspect of my contribution. Mainly, in some sense, I've been fortunate to spend almost 60 years of my life in computer technology, and over the period I've been active, I've been at the forefront of many of the developments. Starting from things like advances in memory, from mercury delay line memories to magnetic-core memories, to semiconductor memories, and lots of advances. So I was right there, in the right place at the right time, and I think as a result of that, I've been able to make contributions to many things. Perhaps most important thing that has impacted society today is whenever you speak to Alexa, or any other technology, the underlying technologies were created by my students, or their students, and the foundations began long ago, in '60s and '70s. And I think that represents the main unique contribution that they were looking at. They may have been looking at the fact that we are about to enter into cars that drive themselves. That's another technology that we created starting in 1984, '83, in the Robotics Institute at Carnegie Mellon. And other things I may have been instrumental, or a catalyst, like the language technology, and the machine learning technology, and so on, where other people did most of the work, I enabled what they tried to do.

Hsu: Thank you. What would you like your legacy to be?

Reddy: <laughs> I don't have any pretentions about legacies. I don't think anyone will remember all the things, except the people that go to visit Computer History Museum. There they will kind of see all these people did-- So you have a lot of Fellows, almost 50 Fellows, and each one of them had a major contribution to the computer history, and so that becomes their legacy. But in the overall scheme of things, maybe one or two people were remembered. Probably Steve Jobs, and mainly because he has created an instrument building on all the technologies that were created before him. But he was able to synthesize an instrument which has become ubiquitous, right? It did not exist 13 years ago, 2007, it did not exist, and now everyone uses it. And we're not saying not [just] everyone that's educated and can afford [it], but we should make it inclusive and give it to every person on the planet, and if you do that, the benefits will be significantly higher. So I think the legacy of any one of us individuals can say "Yeah, I was instrumental in getting Steve Jobs to do the things he did which helped the world." But most people won't remember that, right?

Hsu: Interesting. What role do you see for the Computer History Museum to help inspire and inform technological citizens today, to shape a better future?

Reddy: Absolutely. I think that the Computer History Museum is very unique and central for people to understand progress of technology over the last 60, 70, 80 years, and before then. The Computer History Museum used to have Babbage's machine, which is now taken back by Nathan Myhrvold. Hopefully he'll contribute it back when he passes away, or something, which will be good. Ultimately, it should be in the museum. But the impact of the Computer History Museum, to me, is future generations, when they come, they can go to this one place and see all the things that we did not know how to do, now we suddenly are able to do. And I think most of the new generations will take things for granted, and they just have to put themselves in the shoes of the pioneers, going back to Babbage, or Turing, or Eckert and Mauchly, or whoever, and they will find-- and more recently, people like Gordon Bell, who kind of has helped to set up the Computer History Museum early on. They will all have left a very important legacy that is kind of captured in the Computer History Museum, and I think the future generations will be grateful and thankful for everything you've done.

Hsu: So what one word of advice would you give to a young entrepreneur or innovator?

Reddy: I have two one-word <laughs> advices [sic]. One soundbite would be--

Hsu: No, that's the plan for now.

Reddy: -one soundbite will be "think out of the box," right?

Plutte: Can we start that again? Sorry, there's-- start by pulling it down, and then hold it, and just saying "My one word of advice is," and...

Reddy: Yeah, my one word of advice would be to "think out of the box" and be willing to consider all kinds of options, rather than the obvious ones. So that is my-- think out of the box. My second word of advice is what Google used to say, but maybe not [be] practicing fully, is "Don't be evil," or "Do no harm." I think that's something we all subscribe to, and it's a good thing to think about. Any time we're doing something, if that's going to cause any harm to society, or people, others, we shouldn't try to do that.

Hsu: Great. Can you tell us a story that illustrates why you chose that word?

Reddy: This one?

Hsu: Yes.

Reddy: Yeah. Yeah, I--

Hsu: Or both, actually.

Reddy: Yeah. This first one, I chose that because we are not always right. We don't know all the right answers. So you must have an open mind to be willing to listen to others, other ideas, that are not your own. And if that idea, it happens to be different, you shouldn't reject it out of hand just because it's not an idea that you thought of. And that's what it means to think out of the box, you know? And the second is we're all working on technologies that can do immense amount of good for society but can also be used by some people for harm. And usually, such people are people in power, and want to kind of create perhaps a different kind of world than what you're comfortable living in, and it's very important for us to be always alert to see how such things could happen. And we should, number one, refuse to work on such things if it's not-- and many times, it's not obvious that whatever you're creating is going to create some harm. But when it does happen, we have to come out and say "That is not acceptable." And even work towards making such uses of technology inoperative and unusable.

Hsu: Thank you. So people are also interested in how-- to learn how innovators think. So can you describe maybe a time when you had an "Aha!" moment, a eureka moment?

Reddy: I think I have a lot of "Aha!" moments. Basically, the best example-- these things don't happen completely out of the blue. But most recently, the "Aha!" moment I had is we're entering a world where we need to know a lot of things about a lot of things, and the current educational system is simply not set up for educating us on thousands of different topics. So if I go to college for 10 years, or even 20 years starting from kindergarten, I have 20 years times 2 semesters, times maybe 5 topics. So at most, I can learn 20 times 10, that's 200 topics, I may need to know 10,000 topics. And so the idea I came up with, which suddenly came to me, is perhaps we should change the whole structure of education so we'd-- every day we learn a new topic. However, in order to learn a new topic, we need to be literate, and we need to have some basic skills. And so the question is how do you organize? So if every day you learn a new topic, and you learn for 250 days in a year, and so you study for the whole life, you're learning all the time. And maybe the first 20 years you're learning maybe 5000 topics, maybe 1000 of them are foundational, and the rest are kind of elective, you can learn anything you want to learn. And if we had a new kind of educational systems where you learn a new topic every day, and if you need to learn more than one day's worth of something, let's say you need to learn how to cook. You may need one week or two weeks, fine, you spend ten days of learning how to cook. It doesn't have to be all done at the same time. It could be done, but the issue is to structure learning so that the chunk of time you're looking at is not one semester, but one day, and whatever you can learn about it in one day's good enough. Like if I'm learning diamond appreciation, I don't need to learn more than one day's worth of diamond appreciation. And that may be true even for things like art appreciation. I may never want to be an art connoisseur, or may not have the money to buy any art. However, I'd like to have at least one day's worth of stimulating education that tells me basics of art, and what to look for, and how to appreciate it, and that's all I need to maybe know. But I never get the opportunity today because the educational system says I must need thermodynamics if I'm going to be an engineer. And the fact I may never use thermodynamics in the rest of my life, or calculus, is not important. So the idea of somehow restructuring the entire education is something I came up with. But that's part of an earlier solution, what I've been proposing, what I call one-day courses, and one-week courses, and one-month courses, instead of semester-length courses, which are one month. So these are all interesting ideas. We all come up with them, not just me, all of us. When we see a problem, we say "Yeah, if only." But those are little ones, right? The big ones, like what Steve Jobs came up with when he created the iPhone, required for him to be living in that. He had to kind of build the desk-- Macs, and then build iPads, and iPods, and finally, built iPhone. And each one of them are building blocks for ultimately coming up with iPhone, and it also needed access to lots of other technologies, and knowledge of lots of technologies. So what happens is if you're in a system, in a lab, or in a situation where you are being exposed to all the future technologies, then you will in fact come up with your own versions of out-of-the-box thinking, and new ideas. And if you're not exposed, then you'll never think of it. And so if you are living in a forest as a tribal [society] and you have never seen any of these things, and you may innovate how to do hunting better, like ten thousand years ago, or twenty thousand years ago, when they figured out how to hunt mammoths and bisons, and so on, you'll do things that you need to do, creatively, in that environment. So where we'll be in a hundred years from now is anybody's <laughs> guess, but people will then come up with interesting solutions for that time.

<03:03:10 audio break>

Hsu: Okay, so next I'm going to ask about demystifying myths. So this is kind of a chance to set the record straight. What do you think is the biggest misunderstanding about the history of your field, of AI?

Reddy: The biggest misunderstanding is people listen to the phrase "artificial intelligence" and everyone makes up their own version of what it means. And unfortunately, most of them are wrong, and they are-- the common misinterpretation is somehow we're creating an intelligent human being, or a person, a robot, artificially, which has all the capabilities of the human beings. And the reality is, when we started, when the original phrase "artificial intelligence" was created, all that they wanted to do was to see if we could do things that human beings could do when we see them being intelligent, that is, play chess and prove theorems. And that was expanded over a 10-year period, from 1956 to '66, where we said we really want to be able to do anything that a human being can do, not just some highly intellectual tasks. That means, can they walk? Can they see? Can they hear? Can they talk? All of the things that human beings do, with no effort at all. And then we would've expanded it further to kind of see whether you can be creative, whether you can be emotional, and whether you can understand a joke. All kinds of things that nobody's really working on at this point. Not yet, anyway. So when somebody says "Human beings discovered new things, can a computer discover new things?" They do, and it's usually something that-- from big data that is not obvious as a pattern for human beings, but can be done by computers. So in that sense, we are--

<pause in thought>

Reddy: I've lost the thought here. You were asking the question that-- can you repeat the question again?

Hsu: What's the biggest misunderstanding of the field, of the history of AI--

Reddy: Oh, <overlapping conversation> misunderstanding, yeah. And so, basically, the common misunderstanding in the history of AI is people misunderstand the phrase and make up their own versions of them. And then they go on to say "That cannot be done," or "Something is wrong," and so on. But people working in AI are struggling to make even simple things that human beings can do, and nowhere are we working on things that will replace human beings in any way, shape, or form. And so a lot of the advances that we made are simply things-- to kind of do things that if done by human beings-- it's no different than any other technology. When you invented an automotive, a car, you're now able to go 50, 60, 100 miles an hour, whereas before, you could only go 4 hours-- 4 miles an hour at most, or less. So the same thing is happening now with thinking. We're asking the question "Whatever I do with my brain, can I use a technology, create a technology that will make it easier for me to do the same thing?" And so ultimately, if a computer can read, and understand, and see, and do everything that human beings can do, then there comes a time, you can imagine, where you don't have to do anything. You don't even have to learn how to read or write. But we are not anywhere close to <laughs> that because they-- and if the computer can read and write and understand, and tell me what I have to do, all I have to do is sit around and do nothing, that would be one world you can think of. It may be possible, but I don't think that's where we'll go. But more likely scenario is there are boring things that we don't want to do that a computer will take care of it for us. Like, for example, paying bills, and reading newspapers, and paying-- answering e-mail and, so on, the things we do which currently takes up a lot of our time, which could be done with the assistance of computer, faster, better, cheaper. That's what we're looking for, and that's all that's being done. So there's no reason at all to think any time in the future computers will take over the world. But they will take over the world in that sense of people using technology can cause a lot of grief if they had access to all the tools. And the tools we're talking about, essentially disabling the internet, disabling the electricity grid, and disabling the usual stuff; food security, water security, and transportation, and communication. So in so far as that can be done, all that's happening now is whereas in the past it may have taken someone much longer, now it may be done-- doable faster. And so we need to kind of also create solutions which are equally fast. When something like that happens, "How quickly can we get back up to speed?" is the question.

Hsu: I think that brings up a-- something that I-- that you touched on earlier. How do you define "intelligence"?

Reddy: Yeah. So we used to have a standard definition of intelligence for a machine—AI research purposes, and which had four or five attributes. A machine that is intelligent, are intelligent, first of all, can learn from experience. And number two, it accumulates and uses large amounts of knowledge over a period of time. Number three, it is able to operate in real time and communicate in language, and use abstract symbols and abstractions for what you might call chunks in psychology. And those I'd say are five of the attributes that we use for defining "intelligent". At this point, we have not had any system-- we have had systems that use one or two of them, but no system that has all those five attributes, or six attributes. Whereas human beings everyday do use all the attributes. That's what makes them unique, in some sense.

Hsu: So would you say that machine intelligence and human intelligence are very different? Or are they-- are there similarities?

Reddy: Machine intelligence, we need to define what we mean, if it has all these attributes, right? Human beings also have the same attributes and they're able to do it, whereas machines don't have all those attributes simultaneously in one program or one app, and that's what makes them different. And machine intelligence, in the narrow sense of the word that's used today, is for a machine to do something which is uniquely human, but not everything that human beings can do. For example, the recent experiment in GPT-3, where it was able to write a whole article about some subject pretending it knows the subject. It was mainly using statistical knowledge accumulated from large numbers of articles on the subject. So the fact it is doing that, writing an article, doesn't make it human intelligence. It simply says there are tasks that human beings do that machines may also be able to do, and thereby providing respite from boring things you have to do, human beings do. Every day most of us are doing things, many of them are boring, and the only reason we do it is we need to do it for a living, and that's where I think machines can be helpful. Then the question is supposing we can get to a stage where human beings-- machines can do most of the things I want done, and just what do I do? Some people have laughingly said "Go to the beach and have a martini," or something, or a cocktail. So those things may make sense, but I don't think that's the way human mind works, and how the human beings organize themselves. My suspicion is if machines take over most of the routine things we do, we will have more time to create, do creative things. And in particular, another way of looking at it is purely in an economic sense. We work because we want to create wealth, and when we work on some useful task, somebody pays us for the task. If we get to a stage where most of these tasks already can be done by machines, then we will hopefully-- one concern is, at that point, a lot of work-- we'll be unemployed, we won't have anything we can do. I don't believe that is the case. I mean, just like before the Industrial Revolution, 95% of the people used to live on the farms, right? And now only 4% of the people produce all the food that we consume, and therefore the rest of the people are doing other things, and that's what I think will happen. Namely, if we get to a stage that many of the boring things we're doing today will be done by machines, then we will invent other things to do, and what we'll do is, in the process, create more wealth. And in particular, my expectation is we'll create a society where there's plenty of everything, whatever you need, and therefore you will be in a different kind of society where the market economics of today, and all the political problems we're facing today, might go away.

Hsu: Sounds very "Star Trek". <laughs>

Reddy: Yeah. It's realistic, and the only question is when will it happen? Namely, what kind of jobs are going to become unnecessary in the future? We're already discovering that after COVID, many of the people that were having jobs before no longer are needed for those jobs. However, there are other kinds of jobs that have opened up that there are serious shortages of people to do the jobs, those jobs. And so there's kind of a reorientation of workforce that's going on right now where there's some jobs that don't have enough people that can do them, other jobs are no longer needed. And so my suspicion is it's more than "Star Trek", it's essentially technology-driven alternatives that we're creating. <laughs> But I guess that's what "Star Trek" is, a technology-driven alternative. So you could imagine ways of actually getting there.

Hsu: Yes, thank you. We've talked a lot about the intersection of your work and society. When you were first working on speech back in the '60s, you were explicitly thinking about the social impact of that?

Reddy: Not at all. I did not seriously start thinking about social implications until the '90s. And even then, I was simply saying "I'm working on talking to computers," and-- or something like that. But as I got older, I kind of realized that perhaps publishing more papers is not the only thing I could be doing, perhaps I could do something that might help society. And so in the last 20 years, more of my time has been going to technology in service of society, rather than simply creating technologies.

Hsu: So would you do anything differently today, if you were developing your technology today?

Reddy: No, I don't think they're different. It's just that when I see an announcement of a technology, immediately I think about how it can help society. Whereas in the old days, it was solution for the sake of the solution.

Hsu: And what technology has made the biggest impact on your life?

Reddy: Computers, <laughs> as a whole. Essentially, in the last 60 years I've been living with computers, and they have transformed the way I live, learn, and work. Everything I do, from morning to evening, is around computers.

Hsu: Can you comment about maybe the promise and perils of India's digital identity system?

Reddy: Yeah, so they have something called Aadhaar Card. Is that what you're talking about?

Hsu: Yeah.

Reddy: I don't fully know all the promise and perils of it. They needed a unique identification for every person so that they can provide social benefits. Free rice or free whatever. So there's whole societal delivery systems, like pensions, and healthcare, and so on, all of which needed a way, and it turns out there-- a lot of common names in some-- like "Reddy" is a very common name. There are probably millions of "Reddys" in India, and then other millions of "Reddys" in Ireland. But the unique identity problem is something everybody has, and India only woke up to it just recently in the last 20 years or so, and introduced this ID, unique ID. Because of the advances in technology, they were also able to add biometric identification to the conventional ID. In the old days, like when you got a social security card, all you had is a social security number, that's all. If I took your card and used your number, the systems wouldn't know any different. Now that you're able to kind of precisely tell, and therein lies the promise and peril. Namely, if you're using a unique ID for every transaction, everything I know, or everything I do during the day may be knowable, right? For example, even now because of the phone, as I move around, because of the GPS system exactly-- my location is known all the time within a few meters resolution. And that is a concern for many people, especially if you want to preserve your privacy. Everybody wants privacy but not-- different people have different definitions of "privacy". The younger generation does not seem to have the same metrics of privacy. For example, they will disclose all kinds of things on Facebook that I would never put in, mainly because I'm an older generation and I don't feel the need to tell everything to everybody that I'm doing. And so my privacy is not the same as privacy of a 10-- 20-year-old, and the issue is can both of us maintain our privacy? I think so, but we are already in a society where a large part of what we're doing is already known. Our phone calls are monitored all the time. You may not think so, but I believe that is the case, and everything we do, every place we go. And even if you walk on the street, the cameras, surveillance cameras on the street are capturing your picture. So that most of the actions you do, every-- things you buy and sell, and whatever you do, are already known, but they're only known today to a small group of people with a lot of resources. In the future, it is possible that anyone can find out by buying an app what so-and-so is doing at this point. You may want that to be known, or may not, and especially if there're some stalking, or personal hazard involved, and having a unique identity makes it difficult. So I know, for example, there are many famous people, especially actors and so on, need to have a different identity when they're traveling because they don't want people to know when they're traveling on a plane, or something. And I believe that is being worked out somehow, it's just that most-- the rest of us can't have multiple identities. I can only have one identity, one passport, one driver's license, and I have to use them wherever I go, or buy. So there are some disparities in the system, but that's okay. Basically, I think the society will evolve, and new rules and new regulations will come up, and what was okay, what was not okay 20 years ago is now okay, and what is not okay today will become okay 20 years from now. So we'll have to wait and see.

Hsu: Thank you. One of the things a lot of your students and collaborators have really talked about is your unique ability to obtain funding. Can you talk a little bit about that?

Reddy: I'm not sure I'm unique. Basically, I was lucky, in the right place at the right time. ARPA has a history, it's kind of grown from small funding in '63 to hundreds of millions of dollars. And so they have to fund activities, and so if you're working on things that sound promising or interesting, then they will fund it. Second thing they probably look for is track record; namely, have you delivered on what you promised, right? And so fortunately people at Carnegie Mellon have a record of delivering on their research targets, and so much so, we uniquely get our share, or maybe-- of our funding. But it is partly the location, if you are from MIT, or CMU, or Stanford, you're immediately assumed to be capable and competent, and so on, so you get funding. The same people in some other less known place, when they apply, they won't have the same benefit of doubt. Rick Hayes-Roth is a good example; namely, he would tell me when he was at Michigan before he came here, he would have a hell of a trouble getting funding, or papers published, and all kinds of things. The minute he appeared at Carnegie Mellon and started writing papers with Carnegie Mellon address, his papers are accepted. So that kind of thing happens, too. So you sometimes benefit from association with good institutions.

Hsu: How are you able to help Carnegie Mellon weather the AI funding winters?

Reddy: Basically, Carnegie Mellon has always been good, as I said, and I happen to be at some points, both in the Robotics Institute and as the Dean of the School, in the right place at the right time, and was able to kind of facilitate that. And I think I was more of a catalyst, rather than someone I [who] magically got things...

Hsu: Okay. Can you talk about how important your connections to ARPA, or DARPA, program managers were in this process?

Reddy: I'm not sure I had any unique contact into program managers. But over a period of time-- see, one of the good things about being an early bird in computing is you know most of the people in the field. In the '60s, everybody and everywhere I knew, right? And so when you needed to get some idea explored or funded, you just went to them and said "I have this idea, what do you think?" and "Can you fund it?" And many times, in the '60s, and '70s, and even early '80s, that's all it took. They said "Okay, send us a proposal." Now it's much harder, they have to go through some announcement process, and you have to submit proposals for some particular activity, and then you have to compete with the whole world. You get a little bit of benefit given that you're-- have the track record, but not always. Basically, it's an interesting challenge.

Hsu: And you've known a number of DARPA Directors and other people; Larry Roberts, Ivan Sutherland you mentioned, Bob Kahn, Cordell Green. Could you discuss relationships with them? Just sort of talk about your connections to these important DARPA people--

Reddy: <overlapping conversation> Yeah, it was more of a friendship. You know them before they were in DARPA, and you met them before, and then they go to DARPA, and they come out of DARPA, and they're back in some university. And so I think over a period of time, you professionally develop links and friendships. And, for example, Bob Kahn, he's a good example. I've known him from '71 or something, maybe earlier, and I still interact with him now these days. And so it's a long-term interaction where he was only at DARPA for maybe 10, 15 years of the 40 year life cycle of our relationship.

Hsu: Yeah. So then that speaks to the fact that the DARPA community is really sort of this wider computer science community that somehow DARPA kind of brought together. I think Alan Kay's mentioned that. How important is that, the fact that there's this community that everybody knew each other? And partly because-- through DARPA connections?

Reddy: Yeah. I think that's extremely important. Basically, now we've become so big, and also, we have trillion-dollar companies. There were no trillion-dollar-- <laughs> not even billion-dollar companies, other than IBM, in the '60s and '70s. Maybe DEC, and maybe Data General. And so comparatively speaking, the number of people that are active and doing great things now is much larger. However, from the point of view of community, in the '60s and '70s, we're a small community. I would say less than a thousand people, maybe, and everybody knew everybody else. And so, for example, even though Gordon Bell was never at DARPA, I knew him from Carnegie Mellon. He was a professor here, and then when he went to DEC, and then subsequent enterprises, I've stayed in touch with him, I still know him. Same for Jim Gray, who's no longer here with us. And same for Ed Feigenbaum, who was on my thesis committee at Stanford, and colleague, and then friend for the last 50, 60 years. So many of these contacts you have are contacts that have been developed over a period of time. And many of them, you may not be working on exactly the things that they're working on, but they know what you're doing, and you know what they're doing, and there's some mutual respect, and so on. And that helps a lot, and that in turn comes from track record. And that's also because the number of people at that time was small. Now, there are probably thousands of people that are doing great things I don't know anything about. When I suddenly see a name, I say "Who the hell is he? Who is that person?" And the only way I go find out more about them is by going to Wikipedia and see they're there, and many people are already there, and you can know what they're doing.

Hsu: Thank you. Have you ever had any issues with the ethics of getting funding from DARPA, or being funded essentially by the US military? How would you navigate such issues?

Reddy: Delicately <laughs>. It turns out we had two very good people, I have high respect for them, who we lost because they did not like being funded by DARPA. One of them is Geoff Hinton, and Geoff was at Carnegie Mellon in the early ‘80s, and he left around ’87 or so to Toronto. He invented the backpropagation algorithm when he was at Carnegie Mellon. And he was great, a great guy, and good friend, and-- but he was not comfortable knowing that many of the computers he’s using, and all of the resources he needed, are all coming out of a DARPA grant, even though it was an umbrella grant that I was the PI of, or Allen Newell was the PI of. And so they didn’t have to answer to anybody or do anything, but they just felt uncomfortable. And my answer to them was “Look, these funds are going to be spent somewhere.” It will be better to spend at CMU doing good fundamental research than going into some weapons area, or something. And they can hear that, and they can maybe accept it, but it’s difficult. Another person that had similar problems was Ravi Kannan who’s another serious Professor in Theory at Carnegie Mellon, who ultimately left and kind of-- still around, does various things. I think he’s working for Microsoft at this point. And the main issue there is either you’re working for CMU getting DARPA money, or working for Microsoft, which indirectly is getting, as of last news item, twenty billion dollars providing some cloud resources. So the question is what do you do?

Hsu: So you never-- you saw it as the money is going somewhere anyway?

Reddy: No, but, <overlapping conversation> basically, I had no problem, personally, because I thought doing what we were doing is the right thing to do. But there are people, like Geoff Hinton, that were uncomfortable, and other people, they may have even started the campaign that we should not take any money. And as a result, one of the things we did early on is we-- the university made a policy saying “We will not do any classified research.” And so everything we do has to be publishable. And this sometimes got us into trouble with industrial funding, where we would say, "Everything we do has to be publishable." And they'd say, "No, this is a trade secret. We need to keep it secret." And so, finally, we came up with a compromise saying, "We'll keep it secret for a year or two, some known time, and then it'll be published." And the only thing we agreed on is we will not publish their proprietary data, that's not our data to publish. <inaudible> <phone rings 03:40:00>.

Hsu: But say, you never had any issues with technologies that could have like a dual-use technologies, that could have a military application, like say, autonomous vehicles, for instance.

Reddy: Yeah, no, we worked on many dual use things, and we were happy to do that. It turns out in robotics and manufacturing there are a lot of factories that Defense Department has and maintains for making different equipment would just sit around idle doing nothing. So, we worked on ways of dynamically repurposing all that machinery using software so that you can be making, you know, tanks one day and Toyotas the next day. Or you know, or Hummers the next day, whatever. And I think all of those is the right way of actually thinking about it. You know, we live in a society where defense is essential. And given that defense is essential, we have to support whatever is being needed, and if we can figure out a way of not wasting the resources when they're not defending the country, that's true of people, that's true of equipment, that's true of resources in general, that would be wonderful. And that's what many leaders did. For example, if you go back and look at Franklin Roosevelt, he essentially kind of asked the army to essentially set up camps for all the people working on social service projects in the forests, environment and conservation, and so on. They created half-a-million jobs or something, and they needed a place to stay, and the army was tasked to go set up the camps. They had tents and they'd do that already. And so, they did that. So, I think more-- we should do more of that, if there are people sitting around doing nothing more than retraining and training again, perhaps we need to have dual-use function for many of the soldiers, not just equipment.

Hsu: Yeah. So, we spoke with Kai-Fu Lee, and one of the things he told us was that you were able to actually create, to influence the National Bureau of Standards, NBS at the time, which is NIST now, to create a national database for speech recordings? Could you talk a little bit more about that and how important was that for the speech research?

Reddy: Yeah, I think it's called LDIC, Linguistic Data Consortium, LDC. Which was started at National-- NIST, and then <coughs>, and then transferred to University of Pennsylvania. And NIST continues to contribute in a lot of other ways. And it is not me doing it. It's a community effort. We all recognized we are going to be needing to use common data to produce compatible results. And that meant everybody had to have a system that uses the same data and produces results based on my system, and so you can actually compare side-by-side which system is doing better and which is really not. And so, you know, what kind of errors are they making? Are there errors that are kind of-- if you combine the two systems would you get a better result, for example? All kinds of things. See, it was a DARPA effort, community effort, to produce the Linguistic Data Consortium and other things. And that's the sad part now. For example, if you think-- look at Apple and Microsoft and Google and so on, they all collect speech data. They have millions of hours of speech data, but they don't share it. And if they shared, if they kind of created a national archive, national database, of all that data shared, they could produce systems, even much better than currently they are. But that did not happen.

Hsu: Great, thank you. Could you talk a little bit about your connections to Xerox PARC?

Reddy: Yeah. So, I was an early consultant for Xerox PARC from 1970 to '75 or seventy-- and so there was, you know, people there working on--

Hsu: How did you make that connection?

Reddy: Again, in a small community, right? Basically, Xerox PARC had Bob Sproull Sr. I didn't know him because he was a physicist and kind of got recruited. He [with the help of Bob Taylor] then recruited all the computer scientist people from DARPA and BBN and Berkeley, you know, like and so on. Almost everybody that was there came from DARPA community and I knew all of them. So, but I didn't want to join Xerox PARC, so I became a consultant commuting from Pittsburgh, and it worked out pretty well. And one of the good things that happened as a result of that, is the creation of a pre-laser printer, you know, using selenium drum, and that was called XGP. Xerox had those devices, but mainly for using in Xeroxing technology, copies, and but we realized if you can put an image, or paint an image, magnetize, sensitize the image onto the drum from a computer, then you can actually produce a page. And that's how the pre-laser printer technology worked. Namely, you know, you first code it, and kind of scan the image, and produce the copy on the drum and then instead of doing it from an analog page, which is what Xerox pages were copying-- were here, we took a bit pattern, a bit from a computer, and then painted it onto the thing, and that was possible because they're all electronic signals anyway. And so, I knew they had the device, and so we did a-- because it's a conflict of the following kind, so we said, "If I give you the data from CMU for your speech research or whatever, you need to do something, we can barter, so in return if you give me--," so they gave us an XGP, a Xerox printer, and which we then interfaced with 1401 and built the real-time printing system, you know? That was the first demonstration going away from the old-style IBM line printers, right? And this was done in '71. If you think of it, the laser printers did not come for ten more years later. So, in '71, '72, '73 the team at CMU built the interface and we were printing things. And so, one of the people that was very in a good-- or beneficial for us was Chuck Geschke, who went onto Xerox PARC, and from there went to found Adobe. Chuck Geschke printed his thesis on that machine, so by the time he went to PARC, he already knew the technology and what it could do and how you could produce page-ready, printer-ready pages using laser printer. The main mistake we made that Adobe didn't is in the design we said, "The painting had to be done in real-time. The image creation has to be done in real-time." So, if I had a page I wanted to print, and it should be painted at the same speed or less, or faster speed as the printing speed of the page. And what we-- what happened was that kind of limited, if you had a complex image, you could not paint it, and that meant the page came out garbled, or you know, streaked. And so, ultimately, that particular design we invented, which is real-time scan conversion, was not the right solution. The right solution was the one that Adobe adopted, which is you scan convert first in offline, and then print it. <laughs> Which is what is obvious solution. And so, you live and learn. There’re lots of things like that. Where something, you know, you're kind of trying to accomplish something much harder, instead if you're willing to compromise and say, "okay, I'm not going to do the scan conversion in real-time, but let it take as long as it want, and the only thing I need is a bitmap memory," right? So, if I have an image, which is 10,000 by 10,000 bits, I need a hundred million bits of memory. And at that time in '71, that was unthinkable, right? Ten by ten megabytes or something. And today, it's nothing. You know, all of us have it in a smartphone. So, those are the kinds of things technology kind of makes obvious that the solution that you came up with is no longer relevant.

Hsu: Yeah. Can you talk a little bit about your friendship with Alan Kay?

Reddy: Yeah, Alan and I were friends, you know, because of the community, he was at Stanford when I was there, and joined us as a fresh PhD-- Dynabook and so on. And then he was at PARC, Xerox PARC for many years. And for ten years or so. And then around 1980, you know, in-between, he also was at CMU as a faculty member for a year. And then went back. And came and said, "Raj, you know, we're-- a bunch of us are going to France for this Centre Mondial, which is trying to use computers for disadvantaged in Africa and other places, to teach them how to read and write and so on." And so, my problem was, I was just starting the Robotics Institute and all kinds of other things. I couldn't take off even for one month or so I agreed to commute. I became the Chief Scientist for the Center, but my role was to kind of go there once a month for ten days or something, for a week, and then come back. And so, Alan kind of quickly disappeared, kind of disenchanted because the bureaucracy of the French, and so did many others. And I think I was probably the last one of the starting group that stayed on for four or five years, and that's mainly because I didn't have to live there and deal with them on a daily basis. So, Alan was instrumental in kind of recruiting me to go join the Centre Mondial, which is how I came in contact with Jean-Jacques Servan-Schreiber and President Mitterrand, you know. We went to visit him in his weekend resort, or weekend retreat for a few day-- for a day or two and came back with him. So, and that was also the place where I met Shimon Peres, who was also involved with Centre Mondial and I didn't know Shimon before. Not Shimon-- but the person that introduced me to Shimon Peres was Sam Pisar. Sam is a holocaust survivor and so one day, you know, I was at the Centre Mondial and he brought Sam Pisar, and said, "Hey, you should meet So-and-So," I knew the name and it was great to meet him. And Sam was-- Shimon Peres was probably one of the smartest and brilliant men that I have come across. He's from the out-of-the-box thinking. He was one of the best examples of that. He was all the time looking for ways to bring peace and prosperity, too, for Israel. And he understood there were some unsolvable problems and so he was trying to kind of make that, I don't know, bridge-- build a bridge one step at a time. So, it's too bad, he did not become the Prime Minister. I think he was the Prime Minister, but I forget now, and he was also the President of Israel, but somehow he never got the same recognition as Rabin or subsequent Prime Ministers. So, that gets you back to Alan Kay. All these things happened, because of Alan Kay. <laughs> I would not have gone to Paris, or I would not have met with all the people there, and then contacts with Israel and so on.

Hsu: Yeah, was it the work with Le Centre Mondial, was that work what you received the French Legion of Honor for?

Reddy: Right. So, it turned out in 1984, President Mitterrand was making a grand tour of the United States and so he was talking to people. And so, when Schreiber said, "One of the people that kind of stuck around and tried to help us out was Raj Reddy, and we should do something." And he said, "What about if we gave him a Legion of Honor medal?" <laughs> And I think that's what happened. So, and my own feeling is if a medal should be awarded, it should be awarded for the whole team that went from here. That is, Alan Kay and Sam [Pisar], Nicholas Negroponte and several other people, and Terry Winograd. There were four or five people that were here from there. But for some reason, various people dropped quickly. I think Negroponte was the Secretary General. He was kind of in charge of all of that. I'm the only one that didn't drop out, but other people either did not work out what they could do, and did not get a satisfactory assignment, so they just left. So, that is the reason, I know in '94-- 1984, when Mitterrand was visiting CMU, not CMU, he was visiting USA, and he made a one-day stop, or a two-day stop at Pittsburgh and came and visited CMU. And that's how I got the Legion of Honor award.

Hsu: Hmm, and could you describe a little bit more like the actual work that you did for Le Centre Mondial? I think were you like involved in a pilot program for India?

Reddy: Yeah, but it's not-- basically, the pilot programs, at that time we were mainly thinking about use of computers for education, and the use of computers for healthcare. And the idea is if you had a computer in a village and had a medical, say a medical technician, and the medical technician had 40 most common-- you know, 30/40 most common medicines that would be needed in the village, and then he would type in the symptoms into the computer and the computer would come up with a diagnosis, which, you know, the symptoms may involve his taking heartrate and blood pressure and sugar levels or whatever-- sugar can't be done then-- but you could do blood pressure and heartrate. And so, all of that <clears throat> was a prototype at the Centre, and then we transferred some of them to African countries. Senegal was the one that was most widely used, because Senegal is a French-speaking African territory, and so they had close connections with France, and so that's what happened. And there was the Minister for Technology and Education of Senegal called Jacques Diouf, who was instrumental in getting Senegal involved. And so, I went and visited them a little later, and saw they're using computers in classrooms and schools. But you know, the whole thing, if you think of 1980, you're looking at '82/'83, that was about the time PCs were just coming in. And they were-- and there were other 8-bit computers from-- but all of them were just-- so, getting them into Senegal and connecting them with monitors and training the students, training the teachers, and making sure that the apps are working, all of them took time. It also taught us what we didn't know how to do, and all the issues involved. So, you could see, "Okay," and it just showed it's okay to have great ideas, but when it comes to implementation, there's a lot more detail that goes into it that we had not planned for. Now, we did plan for it ultimately. But it takes all of that to make it happen, not just the idea.

Hsu: Thank you. I want to get back to talking about something that I think we mentioned earlier. So, it was partly your connections and CMU's connections with PARC that-- did that lead to sort of the push at CMU to develop the 3M Machine, the 3M workstation?

Reddy: Yeah, exactly. So, you know, basically there were two people that were closely involved with PARC. I was involved and Allen Newell was involved. You know, he was working on kind of predicting the best way of inputting data. And there is a book called "The Keystroke Model," by Newell and Moran, and one other person, anyway, it'll come back to me. Anyway, this book kind of predicted that the fastest way of predicting how long it would take to do a job is to count the number of keystrokes. So, if you kind of type a capital A or something, it's two keys, cap key and-- whereas, you know, so any attempt to kind of ask how are the fastest way of doing something, you could predict by looking at the number of keystrokes as a first order of approximation. And so, when we came to use of speech, so you say, "Don't type. And so to minimize the keystrokes, just speak it. Speak the message." And that's what you do now these days, if you're speaking into iPhone using Siri or something, so you don't have to use the same keys or you don't have to type. And so, anyway, both of us were there. And both of us knew about the Alto machine and what they have done and they're one machine per person, or personal computer model. And Steve Jobs went to visit them, I think, in '78 or '89-- '79, and saw this personal computer model. And then, you know, he and Gates were working on personal computers of different kinds, so he decided he wanted to build a bitmap display computer. And that's what Mac was. And we were not worried about building a workstation out of PCs, or PC chips, so there were a whole bunch of others like Sun Microsystems, there were two or three other companies on the East Coast, including one in Pittsburgh called PERQ, and all of them were building personal workstations. And so, CMU, you know, basically, Allen Newell and I, I think, convinced the President of CMU that we should adopt one personal computer per student model. And he bought into it, and Newell called it "The Greening of CMU with Computing." <laughs> There was a book called "The Greening of America," and I don't know if you came across it, and so this was what we were trying to do. And so, we came up with one computer per person. And so, we went back to DEC and IBM and IBM said, "no, we have this 360 Model 50," or something, "But it can be used by ten people." So, we said, "No, it can't be used by ten people. One computer per one person." You know, a megabit and a MIP, and a megapixel, you know, whatever 3M Machine, megabit bandwidth. And that's what the 3M Machines were. And we were already predicting we will have 3G machines, gigabit memory and gigabit computing power, again, gigamips, a million of instructions. And gigapixel display. And that happened, actually the giga instructions per second happened, I think, around 1995 or so. So, that's where that whole thing came from--

<04:06:50 audio break>

Hsu: Thank you. Did that also-- was that also part of why CMU invested in NeXT?

Reddy: NeXT, yeah, NeXT was not a-- yeah, yes and no. We not only invested in NeXT, but we were early users of it, and I was one of the people, and for my lab I had like ten NeXT machines or something. And <clears throat> but we tried to keep our buying the equipment away from our investment, because of the conflict problem or whatever. Because we were using DARPA money, ARPA money to buy the things. And so, you know, when I bought them there was no pressure from anybody to buy those machines. I bought it because they were the best you could get at that time. They were more expensive than I wanted them. But there was no way to make all of those technologies available for a lower price. It took time of exponential growth. So, that's how we became both a customer and a user of NeXT. And it had a lot of advances. You know, it had a laser printer, it had the PDF document display, you know, what you see is what you get kind of display. At that time on PCs, you would see one thing on the display, and you would see a different thing on the printed paper, which was kind of annoying and there was no-- you had to spend a lot of time to get them to look the same that you're getting exactly what you want. And so, that's what we did.

<04:08:50 audio break>

Hsu: Yeah, so next, I'd like to talk about the sort of starting the Robotics Institute. Could you talk about that?

Reddy: So, it turns out Robotics Institute was started in '79 with funding from two sources. Westinghouse Corporation, that there was the President of Westinghouse Public Systems Division called Tom Murrin and was also responsible [for] manufacturing automation across the company; and Admiral Baciocco was Head of ONR, Office of Naval Research, both of them together provided the initial funding. But in order to get to that point, we had a series of meetings. The original initiative for the Robotics Institute came from our President, Dick Cyert. He said, "Hey, I hear about all these things. Is there anything we're doing at CMU?" And Newell and I and one other person, we went to-- excuse me one minute.

<04:10:14–04:10:40 off-topic conversation>

Reddy: And so, the Robotics Institute was started with the initiative from Dick Cyert, who wanted to know what can be done and I was the closest thing at that point. And I was also working on a different project with Carl Sagan and NASA. There was a NASA report on A.I. in space. I don't know if you came across that, there was a group of people, all the key people in A.I. were on that committee, except for Carl Sagan most of the people who were in computing or A.I. And Minsky was there, and Pat Winston and so on. Accidentally, I think, I became the Vice Chairman of that committee and we produced this report which was kind of widely used at that time. It is now available as a PDF. On my website you will find it; you'll probably find it in other places. And coming back to Robotics Institute, so, you know, at that time, I was already thinking about robotics and automation for space and manufacturing and other things. So, in the discussions, we said we can-- we're already doing some things, but the scale at which we are doing is very small. You know? We only have a few people and not enough money. And so, the-- what Dick Cyert did was essentially put me in touch with Westinghouse, in particular, Tom Murrin, and he was on the board or something, so there was an approval from the Westinghouse to jointly fund this activity at CMU, so we got like five million dollars from Westinghouse, which is probably what, like 20 million now, in 1979, and we got a million dollars a year funding from ONR to do autonomous systems. And so, we started Robotics Institute in '79, and so we had a number of people like Paul Wright and Fritz Prinz-- Paul Wright is at Berkeley now, and Fritz Prinz is at Stanford <clears throat>-- from Mechanical Engineering, and a whole bunch of people from Computer Science, especially Takeo Kanade from Japan, who was a full professor by then. And so, all of-- and then we had some younger people like Red Whittaker and Hans Moravec who joined. So, initially, there were about ten people to start the Institute and very quickly we were doubling, you know, every few years, and not at semiconductor pace, but every six or eight years. And so, so much so we are currently over 100 million dollars a year, you know, basically, the Robotics Institute. So, I was the Director for about 12 years, and then I handed off to Takeo Kanade, who was there for another ten years, and then others like Chuck Thorpe and Matt Mason came along. So, we-- you know, it's been a very good initiative.

Hsu: Hm, yeah! And the Institute really seems like there was a lot of-- can you talk more about the Strategic Computing Initiative, and how that influenced a lot of funding around autonomous vehicles and other things?

Reddy: Yeah, basically, the strategic computing at DARPA had funding for autonomous vehicles, or they had funding and we went and made a case that you should fund autonomous vehicles at CMU, and they agreed and that they felt it could be justified under Strategic Computing Initiative. <coughs> So, that's how we ended up getting the funding from DARPA for that. And the big problem we had was there was no space. So, we were building this NavLab vehicle outside and it was snowing, and I remember Red Whittaker was in the snow kind of wandering around in 1984 and finally I brought the President and said, "Look, you have to get us some space." So, space turned out to be the biggest political barrier at CMU at that time, because everybody wanted space and there's not enough space to go around.

Hsu: Hm, yeah, thank you. You know, when you became-- I'm going to go a little bit forward a bit. So, after you became Dean of Computer Science, you helped set up all the other various centers that were separate from the Computer Science Department. And, I think, you know, we spoke to Pradeep Khosla, and he actually told us that a lot of the reason that you had to set up all these separate centers was because a lot of these things couldn't be done in Computer Science itself?

Reddy: That's not-- you know, that's one way of saying it. <laughter> You could have done them in other ways, but it takes a long time. You know, think about MIT. MIT Computer Science could never become a separate department until very recently. And they said, "No, no, Electrical Engineering and Computer Science are unbreakable." And they called themselves EE and CS, same is true with respect to Berkeley. At some point, they set up a separate Computer Science Department at Berkeley. But EECS at MIT was there until very recently. And you know, so what happened, of course, is things like A.I. Labs and CSAIL, subsequently, were autonomous units inside the Department of EECS. Only, you know, about two years ago when they set up the College of Computing, then they merged all of these things onto College of Computing. And now it's not EE and CS, and EE is separate and all of them are separate. So, the main thing-- the great thing about CMU, not me, is the flexibility. If they needed to set up a Robotics Institute, and the President said, "Okay, set it up!" you know, and that was it and there was not a lot of political haggling with a lot of stakeholders. And the same thing happened when we set up a Center for Machine Translation. And so, by the time I became the Dean, we knew what things-- there are lots of things that needed to be done. The most important one at that time was Human Computer Interaction, HCI Institute, besides already having this Language Technology and Machine Learning. So, what happened was we had a couple of professors and psychologists, John Anderson and his two associates were threatening to leave CMU to go to University of Colorado, because they were not getting <coughs> the resources and funds-- you know, that they needed and the promotions and positions in Psychology Department. So, they resigned and it went to the Provost, Paul Cristiano, who said, called me and said <laughs>, "Can you help us out? We're going to lose these very important people, what can we do?" So, I didn't talk to anybody, we just, you know, made them-- we created-- we were just creating this Human Computer Interaction Institute. And we made them-- these younger people, faculty members there. Tom Andersson was already adjunct professor, or had a courtesy appointment in Computer Science. So, the way these things get created is the ability to decide something is an important area and move decisively and not seek a lot of consensus and compromise, and you could do that because there was this empowering principle at CMU, it was called--

<04:21:03 audio break>

Hsu: All right, are we ready to go? Okay. I think you were talking about the HCI Institute and the other centers.

Reddy: Yeah, so we were able to set up the HCI Institute because there were some needs where we were going to lose faculty that had this interdisciplinary character of using computers in their field. And these were John Anderson and Sarah Kiesler in Social Sciences and others. Sarah Kiesler and Elise Prow [ph?], for example, showed in their study that people that use computers all the time become loners, you know, antisocial kind of thing. And that was an interesting obse-- result from their study. And the problem was many of these things were not true psychology or not a-- everything that was being done, it was computer-- at the intersection of computers and psychology. So, there was a need to create something between design and psychology and computing, and it turned out, as the Dean, I had the flexibility and freedom just to do it and set it up and do it. I don't believe I ever asked for permission of anybody, whether the Provost or the President. And I also don't believe I got any permission from faculty. Normally, these things you cannot do without going through a whole bunch of hoops. And we had a system at CMU of empowerment, where we could actually do things, and if you screw up royally, then somebody might say, you can't do that, but otherwise, you just did it, you know? So, and that's what happened. So, basically, we created HCI Institute, we changed the name of Center for Machine Translation to Language Technology Institute to include things like search and information retrieval, and summarization and entity detection and so on, in language. So, not just translation. And then we created--

changed the name of CLAD, Center for Learning and Automated Discovery, to Machine Learning Department. That was created in the '90s and then changed to a Department later in the-- and then we also, but so we had Robotics Institute, we had all these four other departments, Robotics and Computer Science, LTI, HCI, and then there were two new institutes we created, Software Research Institute and Computational Biology. So, by now, we have seven different departments in School of Computer Science, of which maybe five were created when I was the Dean, and two more after that.

Hsu: Wow, yeah. That's a lot. That's pretty significant. You also-- did you also have an important role in creating Carnegie Mellon in Doha?

Reddy: Yeah, catalyst. You know, basically, there is an example where it could not be done by me. <laughs> And so, it had to go all the way to the Board of Trustees and everything else. So, what I did and what-- it turned out about that time, I was the Co-Chairman of the Clinton's President, you know, PTAC, Presidential Information Technology Advisory Committee, Irving Wladawsky-Berger from IBM and I were the Co-Chairmen of that PTAC. And as PTAC Chairman, or Co-Chairman, I was at Davos, you know, World Economic Forum. Not I was-- I was invited, and then I was also invited to go to Doha to help them with their educational strategy. And so, they were already setting up a camp-- college on-- it was Virginia Commonwealth University and another college with Texas A&M. So, you know, there was set up a college with Cornell Medical School. So, they asked me if we could set up a school of computer science or College of Computer Science in Doha. I said, "I think it's certainly doable, but it would require the Provost and the President to get involved." And so, in 2001, I literally-- not lit-- the Provost at that time was Mark Hamlet, and I went to Doha, and in Doha we met with Chairman of the President of Qatar Foundation, Sheikh Amoza, and we agreed in principle we'll do it subject to approval by the President and the Board of Trustees. And that took a long time, because there were a lot of trustees worried about, "What happens if there's a terrorist attack or what happens if we have to close down? Or safety of the lecturer or the faculty and thing?" So, all of that had to be worked out, and so it was mostly done by the President and the Provost to kind of make it happen, and there was also huge legal effort. You know, they had to agree on a contract and terms. And again, those are not my cup of tea. <laughs> I don't have time to go through a word-by-word legal contract. So, we-- fortunately we had our Legal Counsel, Mary Jo Dively. So, it was the three of them that made it happen. But whenever there was a hitch or some problem, I kind of got involved and resolved the problem for them. And so, ultimately, we had a CMU School of Computing and Business. We have two degrees. We have a business degree from the Tepper School, and a Computer Science degree from School of Computer Science.

Hsu: Yeah, but you were-- I mean, is it safe to say you-- were you important, instrumental in convincing the Provost and the President to do this?

Reddy: Yeah, I was kind of the catalyst is the right word. <laughs> But for me, it would not have happened, but you know, like everything else, there were a lot of people, and a lot of things that have to be in place to make it happen.

Hsu: Yeah, and so what did you say to them that convinced them this was a good idea? That CMU needed to do this?

Reddy: I didn't have to say anything. You know, I was there, and they said, "Can we do this?" Mainly because they were trying to figure out who to partner with with Computer Science. And they already knew CMU was the best, or one of the best. And so, if we had said no, they might have gone to MIT or Stanford. But since I was there and they were asking, I said, "You know, it's doable, let me go see if I can talk to the Provost and bring him here. And then you can work it out."

Hsu: Great. So, next, I want to talk about Rajiv Gandhi University of Knowledge Technologies. Talk about how you set that up. Where did the idea come from? And then the process of creating it.

Reddy: Yeah, so the basic idea, I think we talked about a little bit earlier came from the fact that rural children are at a disadvantage or in this-- if you think of it from USA, people living in low-income areas or below poverty level are at a disadvantage when it comes to educational opportunities. And why? Because the current system of merit-based educa-- admission kind of says, "Well, you take this exam, and if you do well, we'll admit you, right?" SAT or something. The problem with SATs is it assumes you know all the answers to the SAT. But that means it assumes you have been educated properly and had good teachers and they actually made sure that you learned what you were supposed to. And unfortunately, in low-income neighborhoods, the schools are not that good, and they're substandard, the facilities are not that good, and the teachers are not that good. And when you go home, the parents are not educated. They're not able to nurture you or anything else. And if you need some help, and you can't get private tutoring, you're not rich enough, you may actually have to go work somewhere to make ends meet. So, there are all kinds of reasons. I was starting with the assumption-- we used to have this big debate with Shockley and others about well, somehow, some people are more intelligent than others. And for the first order of approximation, all things being equal, everybody is equally capable and equally intelligent. And but some people have better white matter in the brain, and therefore it takes-- they can learn faster, but others need slower [instruction]. There used to be a person called Benjamin Bloom, he's a senior well-known Professor in Education at Chicago, <coughs>, and he famously said <clears throat>, "I can teach anyone anything-- almost anyone almost anything, given enough time." And the basic idea is, and the second thing they demonstrated is, if you can have a one-on-one education, the students that are D students will become A students. A students will become A++. The whole curve shifts. It's called the two-sigma problem. And there's a famous paper by them on that. And so, the-- I knew about all of this. Mainly because of at Stanford, when I was a graduate student, there was another Professor called Pat Suppes, who was using a neighborhood lab to set up computers to teach children how to do math. You know, Computer Aided Education. And-- <clears throat> <drinks water>-- so, the use of computers in education has been kind of in the back of my mind throughout. And when I saw this result, when as Chairman of the PTAC, Co-Chairman of the PTAC, I was familiar with the results of Benjamin Bloom and also what happened with the California system and University of Texas-- or not-- State of Texas system on educating people. And so, all of that was in the back of my mind, right? So, basically, when I met with the Chief Minister of Andhra Pradesh, <clears throat>, in '70-- 2007, I kind of said to him, "Look, you are from a rural area, and I'm from a rural area. There are a lot of bright people that are still there in the villages because they never had the opportunity to go to school and college and whatever." And so, he said, "What should we do?" So, we set up this RGUKT, you can go look at it, RGUKT.in, for India, is the website, if you haven't already seen it. And so, the bottom line there is this whole system was set up at scale, you know. On Day 1, we had 6,000 incoming students, and it was trying to fix a lot of problems.

One of the problems was giving all of them a computer on Day 1, so that you can-- even if you don't have competent teachers, you can show them the MOOC lectures, type lectures-- this was ten years before MOOC, or five years before MOOC. And so, the system, you know, the-- so they're all the-- we were trying to eliminate all the problems that rural kids have from why they can't go to college. First, they can't take the conventional exams and pass them. So, what we said was, "If you're the top student in your school, we'll admit you," kind of thing. And it doesn't matter how many marks you had. And so, you know, so we admitted 6,000 students based on their relative performance in their school. And the net result was we had people coming from throughout the whole state who had problems with-- they had-- did not only have proper education, they did not know a lot of stuff, but they also had no shoes, and no proper clothing, and actually never had three proper meals. And so, they were anemic and had all kinds of health problems when they came in. And so, but we- to their credit, the Chief Minister provided all the funding needed to build up the facilities and everything. So, we were, you know, doing this experiment and I said-- Rajiv Gandhi University, and it was kind of-- by now we have had five or six graduating classes. It was started in 2008 and it's a six-year program, Grades 11 and 12, and four years of Engineering. And so, the whole program took off. They're a lot-- there were a lot of problems I won't go through. Most of them are bureaucracy, you know, getting the funding and construction and construction not happening properly and so on. But the basic ideas were implemented and they were not good quality teachers, and teaching. So, there was a whole bunch of implementation problems that are still being worked out after ten years, but we've had four graduating classes. And many of them given this initial push, because basically you can think of them as getting educated-- education they didn't get in the school. And finally, when they graduate with a degree, they were the first child in their family and extended family, to ever go to college and graduate from school. And so, many of them-- because they were the top students in their school, were naturally bright, or capable. Left to themselves, they would have become smart taxi drivers, you know, able to kind of produce-- go to the thing. Now they have a engineering degree and college education and so there, you know, many of them are actually being successful. So, we'll see, I think, over a period of time how things worked out.

Hsu: Hm, that's great. I think you-- in a previous interview you mentioned that you had to actually have the university teach things like social etiquette as well.

Reddy: Yes. Basically, the problem with these kinds of people, these simple things, like you know, young women, well, they were 15/16 years old, didn't know hygiene and sanitation and all the other things that they needed, because they grew up in these villages. And so, that means we needed to have women teachers and women wardens who would kind of show them, you know, the whole hygiene and other things. And so, it was-- and the social etiquette basically they're all, you know, socially okay in the following sense, namely, they were part of a school, and they had 50 to 100 students in the school, in the class, so they were okay, but when it came to English medium instruction, learning how to use computers, navigating the environment, and finding food and so all of that takes-- and they were also for the first time they're moving out of their home. Up to that age of 15, they were all at home going to a local school. For the first time, they were kind of displaced, and there was problems, psychological problems. Some of them had felt homesick and wanted to quit and go home. And we said, "Okay," you know, we sent them home, but they came back after a month anyway, because the parents said, "You're missing out on a great opportunity, you know? Where else would you get free education for four years, including free food and accommodation?"

Hsu: Hm, yeah. So, was this-- was this part of like, you know, a program to try to address some of the class and caste issues?

Reddy: Yeah, no, it was mainly poor versus rich, rather than caste. If you happen to be in a city, even if you're, you know, of a lower caste, you know, then you can go to City College, you know, and stay at home. So, the cost of education is almost half, and the chances are in the city you have proper clothes and your parents are doing something, you know, and are able to kind of support you. So, I think it's more economic backwardness rather than caste society, and so I did-- we did a calculation on the cost of educating a rural kid is roughly twice as much as cost of educating somebody who's already living in a city. Because not only-- both of them had to pay tuition, but they don't have to pay the hostel fees, and accommodation fees, and all the other requirements that you have, when you're moving from a village into a city. And more importantly, most of them don't have any place to stay. And it's not safe. And all of that leads to-- so anyway, you were asking the question, where do good ideas come from? As a result of doing this, I have now come up with another idea, which I'm trying to sell to people, and it may happen in a few years, is rather than taking the kids to a college, you should take the college to the kids wherever they are. And so, the idea is with the-- our technology, you can now have a college in every village. And if we only have a small number of students, ten or so, in every class, or ten or twenty, and that's okay, because you know, much of the education, many of the expenses are not there. The faculty are all online somewhere else. And there may be TAs in place. And there many of the normal costs of going to college go away, because you're staying at home and going to coll-- school. And there's also safety and security, and social connectivity that you miss when you go the-- so we call this KG to PG, in every village. KG to PG College in every village. So, that the upper classman can be TAs for the lower classmen, so people-- it's like the old idea of one room schoolhouse in Montana where people would go to this school in 50-mile radius, and there'll be one teacher. And they have to answer all the questions. And if she's busy with some student, some other upper classman student will kind of go answer the questions. So, it's kind of a collective education with only one teacher. And I think that's the same idea, namely, you can actually be receiving high quality education all the way to Masters and PhD degrees without leaving your village. All it takes is connectivity and a laptop and a Zoom link. <laughs> And so, we have to wait and see if, you know, the government will understand it and take off. So, that's where we are.

Hsu: Yeah. So, what do you think RGUKT, you know, where has it succeeded and where does it need to be improved?

Reddy: It is, you know, it’s the cup is half-empty or half-full, right? It's succeeded in that it has actually graduated a number of people and many of the people who might not have otherwise got an opportunity of getting in. It's half-empty in that the faculty were not ready. They didn't understand computer-based education. You know, you had to hire lots of people in a hurry, and now we're kind of stuck with them. And there are other people trying to be hired and then said-- so the faculty hiring and staff hiring is a big problem. Second is the content and the curriculum. The curriculum was set up based on computer-based education and at the same time it has to comply with conventional college curriculums. <clears throat>

Sometimes they're not always consistent and the students learning have to be digitally literate, and when they come in, they're not. And we try to get them literate, and sometimes it works, sometimes it doesn't. And finally, we're taking them from a local language medium, Telugu medium, to English medium, and half of the students don't understand what's being said! So, they just sit there completely blank. And so, ideally, a one-on-one type of education would be the right thing to do, where you set up a one-room schoolhouse type thing and one teacher has to kind of educate all of them. And so, there are lots of things that could have been done differently or better, but it's an experiment. But the benefits are happening already, better than if we did not have it. But I think there are lots of things we can do differently.

Hsu: Hm, great. What lessons can you impart for others on the challenges of building and maintaining institutions?

Reddy: So, basically, my current lessons, you know, basically I was thinking about, it was the President of-- one of the-- I think it was Bernie Sanders or somebody was saying in the-- that they'd provide free education for everyone. Was that Sanders or Buttigieg, or somebody.

Hsu: Yeah, I think that was Sanders. Yep.

Reddy: And everyone said, "No, a pipe dream, it can never be done." I think with our technology it can be done. What it requires is every person to have access to connectivity and computing. And we already know that's going to be infinitely-- it costs less than ten percent of the cost of COVID. And so, assuming that is there, and assuming that we can create learning environments, there's a very interesting result that happened. I don't know if you're familiar with it, called the XPRIZE on Learning. You know what XPRIZEs are, right? Where kind of impossible target is set and the first group that meets it. So, it was kind of launched by Larry Brill-- I think, Brilliant, you know, is the name. And if you Google XPRIZES you'll see all the ones that you have been offered. One of the more recent ones about 12 months ago is the success of a thing called Learning XPRIZE. The Learning XPRIZE said, if you can teach a child to read and write and do arithmetic at the first-grade level, after the end of the first grade, with no teacher, just only, you know, not leaving their village, there's no teacher. And the computer is the teacher. It's your teacher. And so, two groups actually won the prize of ten million dollars, and the key thing is, you know, supposing-- the first thing you need to get a computer in the hands of a kid. And it has to work. So, you know, so what they found is many of these villages didn't have power so they had to actually put a solar panel in the village or in their place, so each kid had a working computer. Then, you know, there's no teacher. And the parents don't know how to teach it, and nobody else can teach it, so the kid had to kind of learn by themselves. So, they don't know what to do, right? So, there's a laptop. They open it, and they see all these keys. They don't know what they are, and how to do. And as soon as you open, the computer comes on, you know, and that's part of the operating system, I guess. And then a little animated stick figure comes and wiggles and wiggles on one corner. And so, you wait for five minutes to see anything happens, it keeps wiggling. So, finally, you touch it or something, suddenly it speaks to you, you know, in your local language what to do-- you know, and then it take-- walks you through step-by-step, what you have to do, and then essentially it gets you to do the lessons and so on. So, it's a fairly impressive accomplishment and I don't know where they are now. The group from USA was originally from Korea.

They were called KitKit Labs or something. But you'll find all of that under XPRIZES in Learning. <inaudible 04:53:54> And so, my idea is we can provide-- if everyone, every child, you know, and even every adult has access to a computer and connectivity and the appropriate software, they can all learn without a teacher. But that's too much of a stretch. It's easy enough to provide teachers in a single room classroom in every community, every block, you can have-- and you go there and you sit with other teachers and you're kind of providing, you know, someone who just graduated from school, as our volunteers from-- mothers who are coming and acting as teachers for two hours a day. Then you can do it with no cost. There are no other costs other than the cost of the computer and connectivity. And which I'm saying we're already going to be given, because of COVID-like problems. And then you-- and everybody can learn what they need to learn! And you know, you can do free education for everyone. And as an idea, it can be demonstrated, it can be done, but it's not being done yet. But it will happen. Mainly because the COVID has kind of woken up everybody to remote learning and learning with our-- on the computer. So, slowly people will see all the interesting possibilities without my-- us having to explain it. Before there was all kind of mythical. They say, "How can you do that?" you know? Now they can discover for themselves.

Hsu: Yeah. Yeah.

END OF THE INTERVIEW

