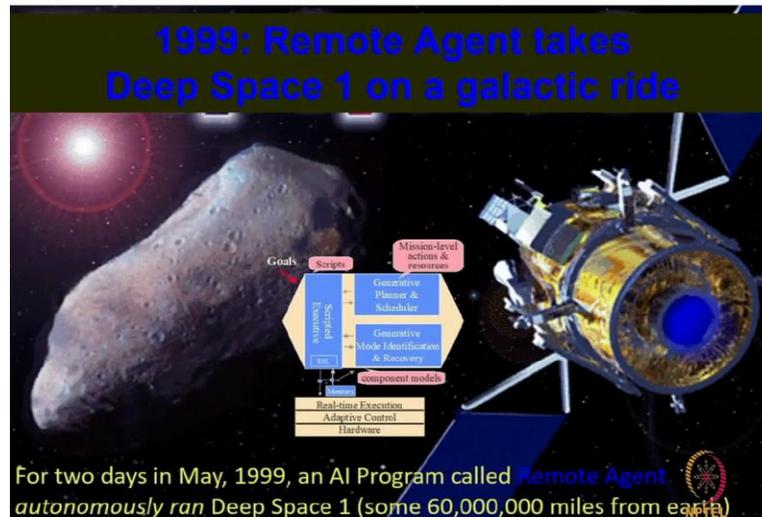


Artificial Intelligence
Prof. Mausam
Department of Computer Science and Engineering
Indian Institute of Technology Delhi

Lecture-4
Introduction: History of AI in NASA & DARPA (2000s), Part-4

(Refer Slide Time: 00:29)



So then more things happened. And a place where AI was developed very rapidly and was being deployed successfully was NASA. This is a, I have some history because I interned at NASA as a graduate student. And also, I was part of a field called AI planning. Which was very central to NASA at the time? Maybe it is still today. So, what are these applications of for example in 1998, there was a deep space mission, which was launched by NASA whose job was to fly by some asteroid.

Now, the remote agent system was a system which was supposed to manage this spacecraft. If a fault happens, this remote agent system will understand that there is a fault and we will it will correct it. So it was a diagnosis and repair system. And this was based on AI. And AI researchers worked very hard to put this the remote agent on to this spacecraft. Now, for most of the time of its space mission, it was operated and controlled by earth.

But for 3 days in order to test the AI system, they said let us give the control to the remote agent. So let is have the remote agent controls the whole spacecraft and let us see if it succeeds or fails.

So the goal of remote agent is to identify any bugs and to correct them. Now they only gave control for 3 days. What happened there were no bugs, the spacecraft did not make any mistakes, and there was no problem? So, nothing was supposed to be diagnosed.

And if nothing was supposed to be diagnosed, it was very difficult to say that remote agent succeeded or failed. So, in order to do that, they simulated some faults in the spacecraft. And so, one of the faults the simulated was a failed electronics unit, which remote agent fixed by reactivating the unit, a failed sensor providing false information, which remote agent recognized as unreliable and correctly ignored and last but not least a particular thruster was stuck in an off position that is what was shown to remote agent and remote agent detected it.

And compensated for it by switching to a mode that did not rely on that tester. So in other words, remote agent did its job but in a simulation on spacecraft because in reality nothing but really happened. But we are happy to take that as a success of here. Because the fact that an AI system went as far as an asteroid is quite creditable and quite a success in those days, you have to go back to the days where AI had just gone out of the winter. And things are starting to just lo up very slowly.

(Refer Slide Time: 03:16)



What I have done in my internship was working on an algorithm for control of the rover on Mars. So everybody knows about the Mars rovers, spirit opportunity, curiosity, there are many versions of the rover. And in the very first generations of the rover, there was almost no planning on

board. I have been told that when you send something to, let us say, Mars, you have you cannot send your latest GPU from Nvidia and latest Intel CPUs.

You have to test the device, the processes, for radiation for this for that for shock, bla, bla bla and it can only send a very robust and very heavily tested chips. So they will basically 10 to 15 years behind on what chips they could send to Mars. But nevertheless earlier what would happen is that all the scientists will come together on a round table, they will say what should my lower do today, Mars can contact Earth twice every 25 hours for, you know 15 minutes apiece.

So basically, they will say together, they will save I think my rover should do this action, then this action, they will sequence all the actions, they will give it to software to make sure there are no issues with it, there are no constraints that are getting violated, and then they will upload it to the rover. The rover will start executing but the rover does not know everything about Mars it only observed something from the camera and senses.

So it starts to execute something and it suddenly identifies a small piece of pebble or something which you did not know existed earlier. So at that time, this sequence did not incorporate for it. So if the rover says, I am asked to go straight, but there is a small which I should not go through the ditch. So what should I do? I do not know. I am an expensive piece of equipment, let me stop and the next time I can contact earth I will tell them there is a (FL: 05:18) what should I do? And so, then all the scientists will come together who there is a (FL: 05:23).

Let us have the rover go around it. So they will again write everything you know manually they will sequence all the instructions, make sure everything is verified and uploaded next time, then they will start executing it will see another the point of the matter is most of the time rover was sitting idle. Now what is a simple solution for it? Give the robot some ability to land on the device on board.

It should not be asking every little thing for scientists from Earth and most of us Scientists from Earth should also not be you know sequencing everything manually they should have a software which is giving them the plans, which a planner is producing. And so this planner was devised

by an AI researcher working at NASA at the time. And, other team off of course, there is always a team. So they created a software.

And that software was supposed to given a goal which the scientists would decide, you should reach that rock or you should do that experiment with the rock. The planner automatically will generate the plan, and then it will give to the human scientists to say, now you figure out is this or not? Now, when they created this software, the first reaction of the operations people was how do I trust this AI system? My rover is a multi, I do not know, billion dollars multimillion dollar at least piece of equipment.

I cannot get over make mistakes. And if AI system makes a terrible mistake, then we will be hosed. And so it went on and on because their operations people did not understand what is happening inside the AI system. They did not trust it. And by the way, this notion of trust is going to stay with us in the course and towards the very end of the course, it will talk about ethics of failure, trust and transparency are extremely important and is a major issue today in the modern AI system.

So the best is NASA scientists could say is that a there is this button. As soon as you press the button, the planner will be turned off. It will no longer sequence anything for you, you can still do everything manually. We have given you full autonomy. Let us get the software to you. Let us put it on the deployment set up. And then if you want to turn it off, you can always turn it off and moreover, it will do all the constraint management for you.

So at least the constraint parts you can trust that it will do because it is just doing some basic check of weather all the safety constraints are satisfied. So operations people finally bought it. They said we will turn it off if you are not satisfied. And at least for the first few years, they do not turn it off. And so that was I do not know, today what is the status, I have sort of lost touch with them.

But I do know that you know, for the first few years, they were now regularly using the AI planning system for doing the control of robot. And later they started giving rovers some motion

planning on board so that it does not keep waiting all the time. So all of those things have started to happen. Hopefully, the latest versions are even further along. So this was another success of a because after coming out from a time where AI was considered a bad word.

We suddenly had AI systems working in NASA and getting deployed in real missions of importance. However, we still did not reach you and me NASA chess these are very sort of specific and niche areas, but it did not touch the lives of people at large still. And so in 2005,

(Refer Slide Time: 09:16)

2005: Cars Drive Themselves

- Stanley and three other cars drive themselves over a 132 mile mountain road



One other major event happened, and this you must already know a little bit about, and this was the DARPA Grand Challenge, you know about DARPA Grand Challenge. This was the first working demonstration of a long trip taken by a car completely by itself. Now we talk about driverless cars. Everybody knows about driverless cars, but you have to realize that this driverless car and this Tesla car and so on this auto drive feature and auto.

All of these features that we have, they have a long history of research. And even before 2005 there was a long history of research, but in 2005, we had the first working demonstration with DARPA thought that we are ready and created this challenge where 220 miles stretch, I think approximately, of the Navajo desert in Nevada was to be traveled by these cars completely automatically. With no manual control, there was not a single human sitting in the car itself.

You can follow it so that you know everything is safe. You may have a control at the end saying that, stop, you know, you are doing something crazy. But the cars are supposed to drive it completely, automatically autonomously. Not 220 132. And there is their story after stories of the stories about this because this was a huge event. And when this got successful, I mean, everybody was really amazed and excited as to the possibilities of the future.

One of the very interesting stories is that that team that was supposed to win this competition and the team that actually won the competition, but 2 different teams. So does somebody know who won the competition? Which university? Know you guys are too young for this. So Stanford won the competition. And it was won by Sebastian thrun group Professor Sebastian thrun at the time was a professor at Stanford later.

He started Udacity and if you and he was the first person to teach an online AI course, along with Peter Novick is you are the coauthor of the bible that you have of the AI book. So if you think about it, AI course was the first course. Which was done in an online fashion and 100,000 people registered for it? And no, some 5000 I think finished it. 10,000 finished it again. I have to check the numbers.

I do not remember So, the fact that suddenly, you are AI course which you were teaching to Stanford students, was taken by people from all across the world, people who may be working in industry, people who may be single moms, people who may be working moms, people who may be, you know, old people who just want to learn about, everybody to this course. And then there were people who completed.

The course and went on to you know, increase the quality of their life by getting the next job that opened the eyes of all academicians to the possibilities of online courses. And now you do you guys do not even attend my classes because you know that, you know you can go online and check out the course content that you are interested in. So to make my classes funnier than funnier so that you are dead.

So, this all started because of AI course. And the person who started this was Sebastian thrun and Sebastian thrun went on to form Udacity, which was an online course with platform, and its competitor later came to be Coursera or I do not know which one came first. And then they competed and then, you know, we sort of know which is more dominant today and Udacity sort of to a different role.

And so, but there is more interesting story about how Sebastian did it and why CMU lost and you know spend 2 minutes quickly telling with some of those very nice nuggets of you know history trivia. For example, Sebastian did not have a long history of doing driving cars, self driving cars. He got a team of undergrads together with researchers and postdocs; a lot of the work was Believe it or not done by undergrads.

Not unlike you, maybe in their fourth year I do not know. And they all work on this thing day in and day in, day in and out and of course, they got a car and they started instrumenting the car and they started putting in the sensors, bla, bla, bla. But they also did not have a long history and so they could use some of the more latest algorithms and they could develop some more interesting probabilistic reasoning algorithms to identify.

Where exactly they are in the big picture because they were in a you had to localize yourself, you had to understand what the terrain is in front of you etc. There are a lot of questions to be solved and very interesting. A lot came out of this particular exercise. Their team was called Stanley. Now the team that was expected to win however was CMU's team led by Professor there to Whitaker, who was very, very famous for exactly self driving cars.

He had a long history, years of history of building driverless cars and working on the technology for it. He did not have one entry in the competition; he had 2 entries in the competition. One was called Highlander one was called sandstorm, and the 2 cars were very different. Sandstorm job was to just follow the leader and come second. Now, I cannot exactly understand the rationale. But it is what it is they probably wanted to make sure that one of the participants finish the competition, because eventually only 4 teams finished the competition.

So, Sandstorm played a perfect race. It kept following the leader, and came second that the car that was expected to win was Highlander. And believe it or not, Highlander started having engine troubles. To the extent that at some point, it is said that the Highlander will go forward, come back, go forward, come back, go forward, come back, and it was stuck. And finally, Stanley carefully over to Highlander ended up winning the competition and Sebastian got a lot press it was an underdog wins story.

But nobody knew what happened to Highlander, including red Whitaker himself. Dead Whitaker said, Let bygones be bygones. Chucked Highlanders into the garage made a new car and won the next competition which is called the DARPA urban challenge, where the cars had to drive streets and roads in the city where they could be people because on the desert, there were no people. So then for 12 years and diversity.

I believe in 2017 or whenever, they were having some fun, or maybe 2019. They got the car back from garage, cleaned it up, and in the process of cleaning figured out what went wrong. They realized that they were filters, which was connecting the engine control module and the fuel injectors. And if you touch the filter, it stopped working. And if you removed your any touch from the filter, it will start working that is it.

It was this one, not even I believe in electronic component or maybe it was an electronic component. I do not know there was just one component that went wrong.

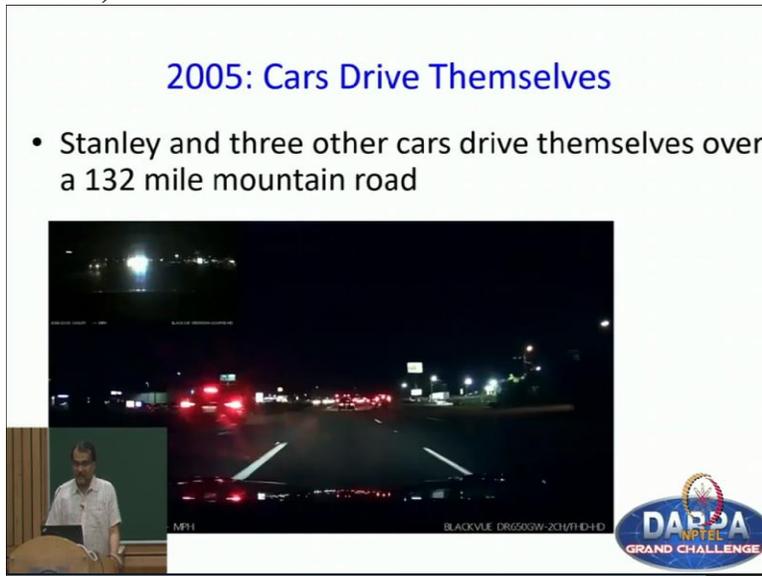
And they lost the competition. Nobody really talks about Highlander and the fact that they won the DARPA Grand Challenge, because the first demonstration is the first demonstration. If you win, if you are forced to do something, you are the leader. If you are the second to do something good.

Think about this so that is where in research, you can do research where you are the first to do X or you can do research where you are the second to do X and you better do you do better than the first. And well if you do better than the first, and eventually if you can make a product, you can make millions. But if you want fame, you have to be the first to do something like the first to do demonetization. That is a bad joke. But the point stands so good. So let us move forward.

(Refer Slide Time: 18:33)

2005: Cars Drive Themselves

- Stanley and three other cars drive themselves over a 132 mile mountain road



And that led to Sebastian thrun going to Google and Tech transferring that technology to Google. So Google worked on self driving cars for a long time before finally, Tesla also benefited from the that kind of work. And we have a Tesla car which have automatic control feature and you must have seen some of these videos where the guard makes a decision to stop because there is an animal coming in and If you see something bad happening like this car.

On the left is coming to the and you go on the shoulder and so on so forth and this is all automatic driving capabilities and in the night it is also pretty awesome particularly because humans can see but it can sense through laser find sense. So I will move forward. So now we have come to 2010 is and IBM Watson has happened.

(Refer Slide Time: 19:46)

2011: IBM's Watson



And Ken Jennings pledges obeisance to the new Computer Overlords



IBM Watson, so again, notice how important IBM has been in this whole picture because IBM has a long history of building big AI systems and the next big a system they built. What was the Watson system where the Jeopardy players IBM Watson defeated Ken Jennings and another contestant, who were the you know, well known world champions in jeopardy on answering Jeopardy questions.

I have a video but I do not think my audio is hoed up so I just skipped but you can check this video link on your computer basically Watson quickly answered questions any question you throw at it Watson would answer faster than Ken Jennings in the end, even in the real in the real thing. Watson ended up getting 77,000 dollar and Ken Jennings said, what is going on I have defeated every human who has come to Jeopardy, but this is too much for me. And finally he says, I welcome our new computer overlords reference that comes from Simpsons