First I would like to thank you for inviting me to give this address. It is an honor to address such an accomplished group of students. I have supervised three students who graduated from IIT Mumbai and Nitish, Aviral, and Rishabh were all amazing students. But i would especially like to congratulate the students who had difficulties with their studies but persevered and overcame them. I was one of those students. I went to Cambridge, which is probably the closest thing Britain has to IIT Mumbai, and I found university stressful and difficult and I did not do well. So my heartfelt congratulations go to the students who persevered and graduated despite having difficulties.

Perseverance is an admirable quality and in the long run I believe it is much more important than just raw talent. To be a good researcher you need a lot of perseverance because most new ideas turn out to be wrong and it generally takes a lot of commitment and hard work to discover why they are wrong.

Philosophers of Science, who are the closest thing that science has to priests, will tell you that you should always try to falsify your ideas, but that is not how people work. You try to prove your ideas are right and save the falsification for other people's ideas. So it is painful when you discover why your favorite new idea is wrong, but it is vital to recognize the shortcomings of your own ideas, because those shortcomings are the seeds of new and better ideas.

Perseverance is especially important if you have an intuitive idea that seems really good to you and the experts in your field tell you it is nonsense. My advice is to always trust your own intuition. Either you have good intuitions or you don't. If you have good intuitions you need to follow them and if you have bad intuitions it doesn't really matter what you do, so you might as well follow them anyway.

Doing science is a complicated process that combines intuition, reasoning, and practical experimentation. In the long run science is far more open minded that the other religions. It can establish that truly absurd ideas like quantum mechanics and black holes are correct. These ideas are much crazier than the idea I was taught in scripture lessons at school that there is an old white guy with a beard looking down on you from behind the clouds to make sure you behave yourself.

Senior researchers like me will tell you how to do science, but my advice is to ignore them and just do what seems right to you. My favorite philosopher of science was Paul Feyerabend who said that the correct method in science is whatever worked.

We still have very little understanding of how the mind actually works so I am always suspicious of people who claim to know where good ideas come from, but I think I have a few useful insights. Good ideas sometimes emerge from paradoxes: things that do not make any sense. For a good researcher, a paradox is unacceptable: there must be some way to resolve it. A paradox for a scientist is like a piece of grit for an oyster. The oyster is not comfortable until it has turned the grit into a pearl and the scientist is not comfortable until she has turned the

paradox into a pearl of wisdom. Good scientists treasure paradoxes and even if they cannot resolve them now they store them away as things to think about in the future.

One of my favorite examples of a paradox is an observation that Darwin made in his book The Voyage of the Beagle. He noticed that some tropical islands have a reef near the shore while other tropical islands have reefs that are a very long way from the shore and are in water that is hundreds of feet deep. The living coral in these reefs sits on top of big pile of dead coral and Darwin wondered how this dead coral got there because coral will not grow in deep water. The only solution Darwin could think of was that the coral reef started close to the shore and then the whole floor of the ocean slowly sank down making the island much smaller and the reef much further away. From a simple paradox about coral reefs Darwin inferred the very surprising fact that the ocean floor can sink.

It took another century before the idea of continental drift was accepted by geologists despite a mountain of evidence that supported it. Geologists thought it must be nonsense because they had a very strong and totally incorrect intuition that the earth was rigid and they found it hard to accept that this was untrue over periods of millions of years even though they knew about earthquakes.

David Attenborough was a student in the 1940's and he encountered the rigid beliefs of geologists of the time. He wrote "I once asked one of my lecturers why he was not talking to us about continental drift and I was told, sneeringly, that if I could prove there was a force that could move continents, then he might think about it. The idea was moonshine, I was informed."

The way geologists responded to the idea of continental drift was quite similar to the way symbolic AI, which was the dominant paradigm for AI in the 20th century, responded to the idea that a big neural network with random initial weights could learn to perform complicated tasks just by looking at data. The idea of learning millions of weights from data with very little innate knowledge was regarded as a silly romantic fantasy. It is hard for students now to understand just how strongly this idea was denigrated. Yoshua Bengio was told that there was no place for neural networks in the International Conference on Machine Learning and Yann LeCun had a very good paper rejected by the leading computer vision conference on the grounds that the whole system was learned so it taught us nothing about how to program a computer vision system.

I think the whole history of AI might have been quite different if the two major figures who believed in neural nets in the 1950's had not died before their time. One was Alan Turing and the other was John von Neumann. Neither of them could be accused of being ignorant of logic, but they were both convinced that it was not the way to understand how the brain computed.

The rejection of neural networks was even more vociferous in the school of linguistics led by Noam Chomsky. I remember watching a TV program called Nova about linguistics. One after another, the leading linguists of the day looked straight at the camera and said that there was a lot we do not know about language, but one thing we know for sure is that it is not learned.

Complete confidence in a manifestly absurd proposition is the sign of a scientific cult, but the good thing about science is that in the long run cults cannot survive rigorous reasoning and rigorous experimentation. Of course, every scientist with strong beliefs needs to bear in mind that they too might be part of a cult and in the early days of neural networks it was not clear whether we might be just another cult.

The remarkable success of neural networks for a wide variety of computational problems that could not be solved by symbolic AI gives me a lot of confidence that they are not a cult and they are here to stay. They are the best way to do speech recognition and to classify objects in images, but they are also the best way to do machine translation or to play Go or to figure out how proteins fold.

Neural networks with about a trillion parameters are so good at predicting the next word in a sentence that they can be used to generate quite complicated stories or to answer a wide variety of questions. This is far beyond anything ever achieved by building on conventional linguistics. These big networks are still about 100 times smaller than the human brain, but they already raise very interesting questions about the nature of human intelligence. They strongly suggest that we are not the rational beings that we thought we were.

People are obviously capable of logical reasoning, and this ability is crucial for civilization. We couldn't have mathematics or bank accounts or railways without logical reasoning. But it now appears that this is not our primary mode of reasoning. Inside a neural network, objects are represented by large vectors of neural activities and these vectors make it very easy for us to perform effortless, immediate intuitive reasoning by using massive analogies.

Let me give you an example of intuitive reasoning by analogy that clearly has very little to do with logic. Basic biology tells us that there are male and female cats and also male and female dogs. But suppose I say to you that you have to choose between two possibilities: Either all cats are male and all dogs are female, or all cats are female and all dogs are male. In the culture I come from it is just intuitively obvious that it is better to make the cats female and the dogs male. This is because the vector of neural activities that represents cats is more like the vector that represents women than the vector that represents men. This may be because dogs are big and loud and chase cats, but whatever the reason it is intuitively clear to me that cats are more female than dogs.

Logical reasoning is vital for correcting our intuitions so that they accord with reality, but intuitions based on analogies are our primary mode of reasoning. Politicians understand this very well. Trump's tweets made no attempt to present logical arguments or facts based on evidence and the liberal media completely misunderstood what was going on. Trump was not attempting to communicate facts. he was providing training examples for his followers. He was showing them how they should respond to various events or situations and by mimicking his responses his followers were learning to have similar vectors of neural activity and hence similar intuitions.

We can actually model this very well in simulated neural networks. It is a process that we call distillation. A student network attempts to mimic the outputs of a teacher network and this distills the beliefs of the teacher into the student. What is remarkable is that the student comes to have the same beliefs as the teacher even about things that were not part of the student's training.

Scientists may claim that their beliefs are based on applying logical reasoning to data, but I suspect that intuitive analogical reasoning is as important in doing science as it is everywhere else. Prolonged experience in a domain gives you neural activity vectors that capture the important similarities between objects in that domain and these vectors allow an expert to draw conclusions with no sequential reasoning. The conclusions are just obvious in the way it is obvious that cats are female. This means you need to put in a lot of hard work to get sufficient experience in a domain and then you need to trust your intuitions.

I now want to turn to the future. I believe that the rapid progress of AI is going to transform society in ways we do not fully understand and not all of the effects are going to be good. In areas like healthcare, AI will be enormously beneficial. Medical records contain huge amounts of information that is not being used and that could lead to far earlier and more effective treatments. If you have a fairly rare ailment and you go to see a family doctor you will be lucky if the doctor has any direct experience of a patient with your particular problem. Now imagine how much better things would be if that doctor had already seen a hundred million patients, including hundreds with your particular problem and thousands with closely related conditions.

AI is already comparable with experts at interpreting many types of medical image and it is likely to far surpass the experts as we get more data and as the technology of neural nets improves. An application that is of particular relevance to India is the early diagnosis of diabetic retinopathy. Neural nets developed at Google can now do this well enough to make a huge difference and save the sight of many people. I hope that many of you will choose to be part of the coming revolution in medicine that combines exciting new technology with helping your fellow human beings.

At the opposite end of the spectrum is the use of AI to create lethal autonomous weapons. I find this prospect much more immediate and much more terrifying than the prospect of robots taking over, which I think is a very long way off. Lethal autonomous weapons in the hands of someone like Putin are a very scary prospect.

It would be great if there was a Geneva convention prohibiting their use, like the one that has fairly successfully limited the use of chemical weapons. But this seems unlikely to me because America, China, Russia, Britain and Israel are all very unlikely to agree to anything that limits a technology that gives them a competitive edge. I hope that very few of you will spend your careers facilitating this nightmare.

Finally, I want to talk about the future of Science itself. The amazing technologies we have today are the result of curiosity driven research unfettered by religious or ideological

constraints. My generation just assumed that this is how things should be and that humanity would never go back to allowing religious authority or political expediency to dominate scientific discovery. But I now think that we cannot take this for granted just as we can no longer assume that fascism is a thing of the past.

I grew up in Britain in the 1950's and 60's and a major puzzle at the time was to understand how civilized and well-educated Germans could have allowed Hitler to gain power. Recent events in the US have made that seem a lot less puzzling.

I believe that highly selected and extremely well-educated students like yourselves have a moral duty to stand up for the truth in the face of political attempts to suppress it. I hope that this graduating class will use their talent and education for the good of humanity and will never put political expediency above the search for truth.