

Audio file

[119037-TheHorizonGuidetoAI.mp3](#)

Transcript

Artificial intelligence. What is it? How does it work? What can it do? By looking back through almost 60 years of Horizon Archive, we can trace its growth. And we can also track how our relationship with artificially intelligent technology seems to be following a pattern. First comes bright-eyed optimism and idealism. The only thing we can be sure of about the future is that it will be absolutely fantastic. Then fear and caution. The artificial intelligences of the future might keep us as pets. Then it's back to hope and belief. This pattern seems to have emerged because of the practical applications and limitations of the evolving technology around us. Hello, Ken. Say, I was wondering... With each new technological leap, we're able to invent machines that can do new and different things. He called it. And we therefore develop different fears to deal with them. Human life has been wiped out by the nanobot. You've got to admit that computers can get out of control. It wasn't so long ago we thought artificially intelligent robots would take over the world. Now it's more likely to be a mainframe with access to our banking details. Are we about to enter a new decade optimistic about all that AI can deliver? Or fearful of its ability to control vast areas of our lives? The roots of the modern computer were largely developed during the Second World War. But it was really during the 1950s and 60s that scientists, engineers and governments recognised just how important this new technology was going to be. These we call computers. They can be made as big and complicated as that installation over there. or as small but still complicated as that. And that little computer is capable of logic. This is a typical, small, modern commercial computer. It costs about 35,000 pounds, or you could hire it for something like 20 pounds an hour. It runs off an ordinary mains plug, and it's capable of doing something like a thousand calculations every second. These were much more than just fancy adding machines, and were becoming capable of an ever-increasing array of tasks. Scientists use them more and more to simulate real-life situations, from the acoustics of a concert hall to electronic circuits. To study the nature of the human voice, they've programmed a computer to sing a song and accompany itself. This is the master science fiction writer, Arthur C. Clarke, on the set of the film 2001, A Space Odyssey. For him, the best scientists were those inspired by classic sci-fi. One mark of a first-rate scientist is an interest in science fiction, and conversely, one mark of a second-rate scientist is a lack of interest in science fiction. Science fiction was important in shaping our visions of the future. In the top of his house in Newark, New

Jersey, Sam Moskowitz has filed away copies of every science fiction magazine and book ever published in the English language, and he's read most of them. Science fiction was an expression of 20th century's man's hopes, dreams, and aspirations. A heavy percentage of science fiction is merely rubbish, bug-eyed monsters, space battles, or something of that nature. However, science fiction has one thing in its favor, that even a story, which may be from a literary standpoint absolute trash, may prove very prophetic. The George Washington of science fiction, the father of science fiction, is generally acknowledged to be Hugo Gernsback. Hugo Gernsback published, in April 1926, the first science fiction magazine in the history of the world. It was called *Amazing Stories*. Anything you can think of that is being done today in the Spade Age, we have done before in some way or another. in my former magazines. If, by some miracle, a prophet could describe the future exactly as it was going to take place, his predictions would sound so absurd, so far-fetched, that everybody would laugh him to scorn. The only thing we can be sure of about the future is that it will be absolutely fantastic. This is a city of the near future, planned by scientists and designers for the General Motors exhibit at the World's Fair in New York. They see a future where man will be making fuller use of the world's at present untapped resources. The future was a shiny place in which anything was possible, and all thanks to science. Improved technology will make it possible to penetrate jungles and build roads with tools so efficient that from tree cutting with a laser beam to laying rd foundations will be a matter of only a few short hours with equipment like this. These tools, it was imagined, would be powered by some kind of machine intelligence, which was surely just around the corner. Science fiction was already becoming science fact. The development men are always on the lookout. Recently, they produced a new version of an old idea, the picture phone. And last year, they put it into experimental service between three cities and the World's Fair in New York. Hi, John. Hello, Ken. Say, I was wondering if you had yet seen Winteringham's demonstration of frame storage. Society hoped technology could improve everything. Even this man's driving. We're driving along a section of Route 128, and I'm trying to get an estimate of the demand made upon a driver by this particular section of road. It was a time of scientific optimism. A technological utopia was emerging in which our every material need would be met. In the garden of his house, Professor Thring tests out the other half of the robot, its legs. An engineer himself, he believes that engineers are the only people who can really change the future, for good or for evil. The evil he wants to eliminate is the boredom of drudgery. And the robot, he hopes, will do just that. Oh, Mr Ford, would you set the robot to clear the table? Yes. Some Americans may believe there is no future for a domestic robot. But here in England, in this house on the outskirts of Epping Forest, there is one actually working. This was a future promising unbridled leisure time, a society freed from the shackles of work. For example, you can envision in the future a housewife dialing up a grocery store, then pressing some more buttons to indicate what groceries she wants delivered to her house. She can authorize the bank, again by dialing a code, to automatically

remit from her computerized checking account to the checking account of the grocery store and so on. And the primary agent of all this change would be... the robots. Every day in our homes and offices, as well as in our motor cars, hundreds of these little robots are doing more things for us than we realise, taking care of the routine tasks and leaving us free to live and work and play in greater ease and comfort and safety. Of course, in the 60s, it was only women who did the housework. Is the housewife also safe from losing her time-honoured jobs, scrubbing, dusting, washing? How can she look forward to having a robot? A robot about the house? If I put the robot into the house and I finally teach the robot to program and dust the entire house, let's say I've taught it to walk around and perform all the dusting features, how long do you think it'll be before the wife wants to change the furniture around? Even the drudgery of applying lipstick will be confined to the history books. The most intelligent inhabitants of that future world won't be men, they'll be machines. the remote descendants of today's computers. I suspect that organic or biological evolution has about come to its end, and we are now at the beginning of inorganic or mechanical evolution, which will be thousands of times swifter. Then there's a biological frontier, the development of synthetic life, the control of living organisms, genetic control, tremendous possibilities and tremendous pearls here. and the cybernetics, the coming of the robots. In fact, to some extent, the robots are already here. They certainly were. And like a toddler standing up for the very first time, they were learning to keep their balance. Now this machine works just the same, in the same way that you gain your balance when you're standing on the ground. When you fall forward, you push down on your toes to regain your balance. Then if you were falling backwards, you would push on your heels so you could come forward. I need not look to have spatial correspondence. All I have to do is close my eyes, because you and I naturally can balance ourselves with our eyes closed. And I can feel myself going back and having the same pressure on my heels as I would when I was standing on the ground. Some visionaries, like writer and professor of biochemistry at Boston University, Isaac Asimov, were way ahead of their time, even envisaging a distant future where robots and humans would merge and become one. We will have our robot becoming less metal, more organic. We may have a society in which robots will drift away from total metal toward the organic, and human beings will drift away from the total organic toward the metal. Asimov's point is perfectly illustrated in this BBC television play broadcast in 1966. So you see, nobody must know about Tony. I mean that he's a robot. Tony? His official designation is TN3, but he answers to Tony. Him. If you think Tony is human, then try and feel his pulse. No. Go on, try. Oh, no. Tony, will you give Mrs. Belmont your wrist? I wonder if we will make robots so much like men and men so much like robots that eventually we'll lose the distinction altogether and have a combined culture. This may be the best after all. Maybe humanity itself will die out as humanity and sort of melt into. this machine culture, which won't look like a machine culture, but to an untutored Martian, may look like a human culture all along. As technology changed the world around them, scientists of the time knew that robots

were becoming a reality. Asimov even invented 3 laws of robotics to restrict what they'd be capable of. The first law is as follows. A robot may not harm a human being or through inaction allow a human being to come to harm. Number two, a robot must obey orders given it by qualified personnel unless those orders violate rule number one. In other words, a robot can't be ordered to kill a human being. Rule number three, a robot must protect its own existence, after all, it's an expensive piece of equipment, unless that violates rules one or two. A robot must cheerfully go into self-destruction if it is in order to follow an order or to save a human life. Although the technology was relatively primitive, it was developing at breakneck speed. So the first tangible fears about what it could do began to set in. Would robots take our jobs and our livelihoods? It was certainly looking possible. Industrialists of the day wondered how their employees would react to automation in an interview clearly of its time. I mentioned before the three laws of robotics. We're designing the robots so that they'll obey those laws. They won't do away with people. And what they'll do is they'll enable people to do more human tasks. The human being at the moron level. And I bring that up only because so many people say that, well, gee whiz, automation and things like that, what are we going to do with our lower class citizens, so to speak, lower class citizens? Well, a moron is actually a highly developed creature. And only a person who's designed a robot can develop a real respect for a human being. People will somehow feel that robots will turn upon them. Now, this I don't think is very likely in the old sense, where a robot turns upon its maker and rends him limb from limb. However, in a very abstract sense, this is precisely what may happen. Robots as a whole will turn on humanity in the sense that a robot culture will be developed that will make human culture meaningless, let us say. and they will come to play a part in more and more specifically human environments. But you must recognize that behind every act of that robot, there's human thought that prevailed, and there is nothing in the robot's current intelligence that would enable it to come even close to independent thought. And we don't even contemplate giving the robot independent thought, there's no need for that. What you can't have in a robot is any motivation. It cannot have its own aim, which it sets for itself. Because for this, an emotional brain is necessary. And we could never put emotional brains into a robot, in my opinion. And certainly, if we did know how to, we wouldn't be so foolish as to do so. By the 1970s, things were changing. Computers were getting smaller and were running increasingly complex programs, able to outperform humans on many tasks. And with true artificial intelligence a big step closer to reality, our concerns deepened. Because it seemed computers were catching us up. Benjamin Landy has been playing tournament chess with his fellow humans for 26 years with great success. Now he's pitting his 10 million nerve cells against the 40,000 microcircuits of a computer. Okay, I'll type in the command to play black for the computer, and then I'll type in his move, which is Barna King 4. Last year, the computer, a fellow member of the American Chess Association, beat Mr. Landy. There was an outcry. American chess journals describe Landy's defeat as a disgrace to the human

race. What sort of mechanical brain can challenge us in this way? Some of us think of them as inhumanly dull. To others, it seems that computers are baffling their creators with prodigious intellectual feats. We're warned that machines must be restrained lest they take over. Computers draining all the fun out of chess sparked fevered debate in the nation's pubs. Where would it all end? To these programmers, the computer presents a constantly changing image. Is it a good servant or a bad master? Is it more than it seems? It's something to fear. You've got to admit that a situation can arise with a computer which gets out of control and computers can get out of control. But we must have more strict control over the way in which they're programmed, which is an entirely different kettle of fish. Oh, absolutely. Well, and so we should keep half an eye on what they're doing so that when they make silly mistakes we can catch them. Of course they should be afraid. These things can take over. All you should do is build up complete and utter skepticism about their abilities. Oh, no. And let them know you can't. Let them move it from theirs. The thing is, nobody's really come up with an intelligent machine yet. Yes, but how do you define intelligence? Which always, back to this one question, always invariably back to the question, how do you define intelligence? An intelligent machine is a machine which can do something which, until the machine did it, was regarded as the perquisite of intelligent man. What were once thought of as purely human abilities were now actually being programmed into computers. One particular feature of a computer program makes it a really powerful basis for creating artificial intelligence, its ability to make a choice. If you have a program which you feed to a computer, there will be instructions all over the place saying, if such and such is now the case, then do this. But if not, then do that. And to do this or do that may mean move to some other point in the program and start from there. It's rather like a recipe which runs if the milk boils over, then, but if not, carry on. The hand-eye machine in the computer science department of Stanford University is entirely computer-controlled. There are no human hands at the other end. By handling objects like this, it's building up its own body of knowledge about the effect of its actions. But the implications of a machine which is building its own model of the environment is that it could operate independently of man. Such a device would not be built in man's image, but it might be better to think of it almost as an executive rather than just a robot. If I hire an executive, I don't expect to have to tell him exactly what to do about everything that comes up. I don't, in short, have to program him. This is the point of the executive. He has a degree of autonomy, and you let him rip. Only when he fails, do you fire him. Now, if you have a very intelligent machine, I think one would like to do the same thing. And mathematically, this appears to be a possibility that programming, as we know it today, whereby every movement of the machine has to be put in by a human being, will get put in several orders removed by the machine itself, and this is perhaps part of what we mean by artificial intelligence. However, if we suppose that this machine is going to be many times cleverer than we are, we really can't afford to let it go away and carry on, because we shan't understand its behaviour.

It was beginning to dawn on the scientists of the 70s exactly what computers could be capable of. They could make simple choices. Some had a degree of autonomy. But one day, it was hoped they might achieve much more, perhaps even abstract thought. Suppose there were to be machines which did have something corresponding to what one regards to facial expression and another human being, which did, as it were, cry out in pain, which did talk a great deal about their inner life, which counted the dreams they'd had and so on. If machines performed like this, then one might, I think, be as inclined to attribute a consciousness in that sense to them as we are, as one is to one's fellows. Let me say a ill word about today's ill word. It's much easier to imagine a machine creating works of art than appreciating them, in fact, because whereas you can, to a certain extent, create works of art by formulae, there already is a poetry machine in Paris, which produces not bad symbolist poetry, not very good, but not bad. I once had a, I mean, a critic mistake on his performances to one, which isn't too bad. It's very difficult to have machines being aesthetic criticism, because aesthetic criticism is something you can't easily formalise. And machines operate according to rules, and there are really up to now no rules about aesthetic criticism. But if you had aesthetic criticism formalised, which I suppose is conceivable, it might lead to some judges nobody would agree with. If you could do this, then I don't see why a machine shouldn't appreciate works of art in that sense. The more the technology advanced, the weirder the experiments became. Well, it was the 1970s. As art embraces random association, the random creation of a computer becomes more relevant. For this dance, long bits of movements, directions and tempo came spilling from the computer like so many beads, with nothing to hold them together. It's up to the dancers to provide the string. Rules for creating by computer are well known, but we don't know how to make the computer evaluate what it produces. For appreciation, there are no rules. The fact that they are programmed I don't think does make an essential difference because we are, after all, a program too. Our programming is much more complicated than that. It's a question of learning a language and of being indoctrinated in all sorts of ways by the people who do teach us. And we have a certain innate endowment, but so, of course, does the machine. The more they're like people, then the more they're like people. And just as I don't very much mind where another person was born or how he was born, or wouldn't very much mind if I found out he was synthetic. I mean, supposing it suddenly turned out that you were the first synthetic person, I wouldn't feel thought threatened by you. I'd be fascinated. The more they're like people, then the more one uses them like people, and why should one worry? But people were worried. They could see how automation and increasingly intelligent machines could ruin livelihoods, and not just those of the working class, but white-collar workers too. At one time, 20 men monitored the nation's power stations in the national grid from this control room. But the job became too complex for them. One man and a computer took over. A 100 years ago, Samuel Butler speculated that the time might come when man shall become to machines what the horse and dog are to us. What is becoming

possible will profoundly affect our future evolution. We're on the road to producing something which is, of course, dangerous. Those concerns deepened as the rate of technological change continued to accelerate. The advance in the behaviour of machines in the last 100 years was like a billion years of biology. Its speed has been millions of times faster because we can combine separate improvements directly when nature depends upon chance recombinations. What's more, there may not be a point at which we can say, Let's stop the technology of machine intelligence here. What sort of relationship will man have with intelligent machines? Are they to have independence? What if an independent machine should make a mistake? Do you blame the machine? Now, how does one look at this? One might say that the machine is bound to make a mistake every so often, therefore it would be unfair to blame it. One might find that a fault had developed in the machine, some of its components had gone wrong, and then you send the machine back to be repaired. But in neither case, surely, would one want to say, Let's punish the machine. But now, if we're saying that machines of an intelligent kind, in the future, if you like, or essentially like people, Are we being equally irrational to say, let's punish this person in retribution if the person has done something wrong? Dream your force. Answer your phone. Protect against. By the late 1970s, entrepreneurs were also getting in on the action with their own vision of what AI could do. Clay 2 is a product of Quasar Industries, Rutherford, New Jersey. No relation to the TV sets. Who programmed you, Clay 2? I will point. Daddy. Daddy. Daddy is Tony Reichel. He says Quasar will soon sell robots \$4,000 each. In December of 1979, we're going to production of approximately 125 robots a day that'll be used for domestic purposes for the home. Domestic purposes meaning it'll do what? Answer the door when guests arrive, take the wraps and store them, vacuum the rugs, polish the floors. But Clay 2 is streets ahead of any other robot ever built. The experts are convinced he's a fraud. This, on the other hand, is genuine. He's called the Wabbit, and in some ways he's the most ambitious robot ever built. The Wabbit and Friends, a group of students who built him at Waseda University, Tokyo. The Wabbit has eyes, stereo cameras buried in its body, which look for an object and direct the hands where to find it. There are touch sensors on the back of the hand that tell the rabbit it has located the object. Picking it up is another matter. The rabbit can also speak. In Japanese, of course. There was also another giant leap forward. By this time, computers could take in and process information at incredible speeds. So adept were these problem-solving machines that we began to call them super-intelligent and feel very small and puny by comparison. We will have almost no communication with super-intelligent computers because nature of what they're thinking about will be completely foreign to us. I've thought of as an example two machines, maybe each the size of a desk, next to each other, communicating with each other. I've thought of one maybe as George and the other as Sam. And they're talking to each other and you walk up and you say, Hey, George, what are you talking about? Well, from the time you started knocking until you finished asking your question, George has already said to Sam, more words than all the utterances of

all the humans that ever lived. So the question is, well, what would it say to you? This is how the computer sees the human brain, the intelligent machine that enabled our ancient ancestors to survive in their brutally dangerous world. The brain's rate of evolution is extremely slow, although its intelligence has stretched to its limits to produce some incredible achievements. Off the ground, on the floor. There we go. By contrast, the dangers of the world have begun to evolve explosively fast. The human brain may not be intelligent enough to cope for much longer, unaided. Never before has a computer attempted anything so profound. We seem to be witnessing the spawning of a new kind of intelligence, artificial intelligence. Already, many computers are behaving like infant prodigies. What will they do when they grow up? Faced with such a rapid rate of evolution, it will be reckless for us to make any prediction. As these machines evolve, and as some intelligent machines design others, and they get to be smarter and smarter, it gets to be fairly difficult to imagine how you can have a machine that's millions of times smarter than the smartest person, and yet is really our slave doing what we want. We've been training chimpanzees to talk in sign language. And of course, a lot of progress has been made there. But if we ever get a chimpanzee that can really communicate, and we tried to talk to him, what we'd find is that he's interested in talking about things like where can he find a banana, and will you tickle me, or playing games that chimpanzees like to play. But if you want to talk to them about nuclear disarmament and who's going to be elected president, they simply won't be interested. And on the other hand, we would be very little interested in discussing for long what a chimpanzee has on its mind. Likewise, I think that the artificial intelligences of the future will be worried about weighty problems that we simply can't understand. And they may condescend to talk to us. They may amuse us on occasion or play games that we like to play. And in some sense, they might keep us as pets. I think that what that means really is they might solve some problems for us, like curing certain diseases. They might find it necessary to take some of our toys away, some of our hydrogen bombs and things, but there's no reason that they would want to go after the same things we want, because they won't be interested in them. I once owned a Porsche, a very high-powered sports car, but I wouldn't have let, say, any old 14-year-old boy drive the car. A Porsche in those hands is a dangerous instrument. I think the state of moral wisdom, say, of our society is such that it is at best a 14-year-old boy, perhaps an 11-year-old boy. And under those circumstances, the very, very powerful tools that we're making, and I'm thinking particularly of computers, I think have to be looked at as at least potentially very dangerous instruments. The thing we call the computer does not grow by a tablespoonful of gray matter every 100,000 years, which is the case in the rapid growth of our brain, but grows a factor of 10 in power every seven years, the computer generation. There's no question but that it will match us in narrow reasoning power by 1990 and go beyond us to become the great new intelligent race of the future. Once the shock passes, of the shock of knowing machines can do things better than we can. That beneficial effect will be that we will have a burden removed off of our back,

which is the burden of being the sort of supreme intellectual creature on the planet. We have to worry about everything today that isn't worried about by God. And when computers come along, they'll be able to worry about a lot of big questions that we're basically incompetent to worry about. It seems that as we journey through the horizon archive, we see the same old pattern and emotional cycles again and again. Hello, Igor. Hello, Igor. Hello, Igor. We start by feeling really optimistic about the prospect of great change. Then, when we actually achieve it, we're filled with fear. However, the fear subsides once we get used to the applications of all that new tech. We start feeling positive again. That's exactly what happened in the 1980s. The fear began to subside and optimism reigned again. We'd have friendly robots and keep them as pets. Genghis is an intelligent micro-robot that can sense what it touches and sees. This robot has quite a lot of sensors. It's got belly sensors, so it can detect if its belly is contacting something. It has force sensing in each of the legs, so it can tell if the motors are stalled. It has pyro sensors up front so it can detect if there's a change in the heat field. And it has whiskers or antennas on the front so it can detect if the front of the robot comes into contact with an obstacle and take an appropriate action. So he's moving back now. Why? Because he said the sensor was an obstacle in front of him that was too hard for him to climb over, so he's decided to go around it. A minute ago, you referred to it as he. Well, Genghis was a male mammal wardian conqueror, so that's why we refer to him as a he. We don't want them to have a problem with a small ego and feel very slight, so we give them powerful names. Scientists and engineers began to re-evaluate intelligence to understand what thinking really was. and to program truly intelligent machines. The theory of AI says that if you analyze the world in symbols and put the right rules in the machine, then it will have a mind and understand the world as we do. It will be a thinking machine. The machine will have to know all about the world, the totality of facts. It has to know that helium balloons fly upwards, that children need looking after. The policemen wear blue helmets, the professors wear special clothes. Easier said than done, but for a brief moment in the 80s, it seemed possible. This was largely due to mass market production of the core of all computers and AI, the microchip. During the 1980s and early 90s, the microchip became smaller, cheaper and ubiquitous. Computers really entered our own homes for the first time. We became accustomed to having calculators and video games, and as home computers became part of everyday life, we worried less that artificial intelligence would take over the world. I think the fear of intelligence in robotics has subsided because By and large, the reality has not fulfilled the claims, so that people have begun to be more relaxed about it. But I think those fears will re-emerge as the reality catches up with the advertising. Scientists pushed ahead in their pursuit of truly human-like behaviour in machines, trying new and different strategies all the time, like simulating biological organic brains rather than mechanical ones. Yes, I want to make a brand intercultural dish. Yes. Aizawa is convinced that conventional computer technology will not be able to keep abreast of the increasingly complex demands of the electronic society. In spite of the

huge advances being made in today's information technology, Aizawa believes that the needs of the next century will only be met by soft, brain-like processes. If we succeed in artificial design brain at the first stage, and then my ultimate goal should be to create the life, but it's a long way to achieve that. It certainly was a long way off. As Asawa was trying to create an artificial brain in the lab, many were still trying to get their robots to achieve even the most basic of tasks. So this is COG. COG is the first serious attempt at building an intelligent humanoid robot. What we're trying to do here is build a robot which is able to interact with people in the way that people interact with people. The quest that all us science fiction fans grew up with of wanting to build a robot which was just like a human. The enormity of the challenge was becoming apparent to everyone that tried. We got it. As we entered the 21st century, the pace of change and the power of computing was set to reach unprecedented levels. While computer power was increasing, the size of its hardware was becoming much, much smaller. Being able to make computers smaller and more portable meant whole new technologies using AI could be dreamt up. We're making a motorcycle with the dress itself. Like, can you get any cooler than that? That's why everybody on the team puts in so many hours, because when you see it work, you say, wow, I made that. But much like the decades that came before, things weren't quite perfect. While some of the vehicles have mastered the art of not crashing, others clearly haven't. But artificial vision is notoriously complicated. After decades of research, it's still unproven, and at higher speeds, dangerously unreliable. What's going on? Sorry. But despite the setbacks, progress continued. As aspirations got bigger, technology got smaller. One day, scientists could manipulate molecules to create tiny computers that will fuse with our bodies. Today, it's very hard to imagine things like direct brain links, you know, You just get dismissed as a nutter if you talk about these things. But in 40 or 50 years' time, it'll actually be quite possible to do that. We can imagine molecular computers, molecular transistors and so on, being small enough, we can get these things into contact with every synapse in your brain. I mean, supposing 20, 30 years from now, I'm wearing a really smart shirt and I have an accident. That shirt knows I've had an accident because it can measure the G-forces. It might even measure that I'm bleeding. The shirt could tell the ambulance in great detail while they're on the way exactly what's wrong with me, so they've got the equipment ready for when they arrive. It might save my life. As a result of molecular computing and those sorts of technologies, over the next few decades, we're going to see far more change than we've seen over the last few hundred. So life will just be beyond recognition. Ultimately, we're not going to go to doctors and have visits in the way that we do today. We're going to have systems in our body that are continually monitoring our body, detecting problems, and fixing them immediately. But very quickly, our hopes for tiny, artificially intelligent robots were replaced by a paranoia that they may start to replicate on their own. Human life has been wiped out by the nanobot. We're talking about instruments that are designed literally atom by atom and molecule by molecule. So they're below what you can see. Miniature, artificially intelligent robots with the

power to create or destroy whatever we ask them to. The intelligence of nanotechnology will not be in one nanorobot or nanobot. It will be a collective intelligence of millions of, actually trillions of nanobots working together and pooling their thinking resources. And if that gets out of control, we would have essentially a non-biological cancer that could just eat up the natural world. That's the so-called gray goo problem. So great is this fear of the gray goo that eminent figures around the world, such as Prince Charles, have raised concerns about it. Luckily for us, and Prince Charles, the fears of grey goo were never realised, and as the dust settled, we continued to focus on the applications of an intelligent computer's power rather than its size. The American engineer and co-founder of Intel, Gordon Moore, had made a startling prediction. Every year, the power of computers was doubling. And Moore was convinced that this unprecedented growth was set to continue, without end. Known as Moore's Law, the prediction had become true. Computers are about a billion times more powerful than they were 1/4 century ago, and they will become a billion times more powerful than they are today in 1/4 century. We'll have both the hardware and the software to recreate human intelligence in a machine. By creating non-biological intelligence machines, that are ultimately billions of times more capable than human beings today. And we will integrate with this technology and it will enhance human potential. Since Horizon has been following scientists trying to create artificially intelligent robots through the decades, it's clear that there has been a number of big breakthroughs. It's also clear that progress is much harder than first imagined. Stair, please fetch the stapler from the lab. Stair, please fetch the stapler from the lab. I will go get the stapler for you. So you didn't tell me you could talk. Yeah, turns out getting a robot to talk is not the most difficult thing. It's blindingly obvious to you and me where the stapler is, but to a robot to figure this out is actually surprisingly difficult. Oh, he's missed it. Stare knows the basic layout of the building and has other sensors that mean it can avoid crashing into unexpected obstacles. Come on, left. Left. It knows what a stapler is, and it knows which room the stapler is in. But other than that, it's on its own. look, he's going down, he's going for it. And when we pick up things, there are lots of ways to do it. When you pick up a coffee cup or when you pick up a bottle of water, the motions you would make with your hands are very different from the motions you would make with your hands when you're picking a stapler. And it has to choose for itself how to do that. It's done it? Yes, and it's picked up a stapler. Stair's done it. It's picked up the stapler. It can find and pick up a stapler. Incredibly difficult task, but not quite the home help of the 1950s techno-utopia. These robots have been set up to develop, much like a young child does. They're beginning to understand how their bodies work by looking at a reflection of themselves. You've got a mirror here, so what's the... Well, the experiment is about that the robot would learn something about its own body. Because in order to move in the world, in order to control it, in order to also recognize the movements of another, you need to have some sort of model of your own body. Right. And the way that the model is going to be built up is that the robot is doing actions and watching itself

performing these actions. So to get a relation between the visual image and movement of the motor. So here you see this looking at its hands. And you also see very much how it's trying to keep balance. It's pretty impressive, actually. How all these motor commands are sort of in an early phase, right? It's extraordinary, isn't it? really does look like it's sort of encountering itself for the first time. But what's even more remarkable about Luke's robots? is that once they're able to recognize themselves, they start to evolve their own language and communicate with each other. One of them is going to speak, and then he's going to ask the other one to do an action. He is going to invent a word, because he doesn't have yet the word to name that action. Right. Okay, so then he says the word, and this one isn't sure whether... What it could mean, right? This is a brand new word. So then he will make a guess. And if the guess is OK, totally by luck. Right. Well, they both know this word, and they know for the future they can use this word to communicate with each other. OK, so he's speaking first. He's doing the action. That's fantastic. OK, you notice how he looked. Okay, so now he's recording what? He's getting the real action. All right. Okay. So now there's another interaction going to happen. Again, I don't know which one is going to speak. Okay. Oh, is that the word he just learned? Yes, yes, of course he knows already. So he's doing, he's doing it. And he will say yes, presumably, will he? Yeah. Over the decades, we've dreamt of a future with artificially intelligent machines. As the Horizon Archive has shown, we've weaved our way through cycles of optimism and pessimism, fear and wonder, and back again, as our future alongside AI has been brought into focus. But while engineers and scientists have worked to achieve the reality of a truly humanoid robot, an unexpected form of AI has been creeping up on us. And this surprising shift has changed everything. Our future would be built not on robots, but on data. Now, dreams of robots are being replaced by the realities of cloud computing. of networks that stretch to all corners of the globe, connecting people and information in ways we'd never dreamed possible. Colossal amounts of data, uploaded and downloaded across the world every second, allow us to communicate and navigate, to monitor our health and diagnose disease, and to track payments and shipments, even to find love. The reason why I believe it's good to be a technology optimist is throughout the entire history of the human race. Technology has empowered us. From the very early days, the Bronze Age, the Stone Age, to today with smartphones and modern medicine, it has freed us, it has levelled the playing field for everybody, and has empowered us as a human race. Why stop that? Artificial intelligence has integrated itself into the very fabric of our lives, almost without us noticing, and in ways we could have never imagined. Smart homes are now a reality where lighting, heating and other electronic devices can all make seemingly intelligent decisions. Wake up in the morning, the house knows that I'm waking up, it can wake up with me. When we walk in the kitchen, it will play the local news and sort of greet us into the day, tell us the weather forecast so we know how to dress. Artificial intelligence is appearing everywhere and powering everything. In warehouses, AI is enabling robots to organize, move, and prepare millions of online orders. This is basically a very large

warehouse where orders come in and they need to be fulfilled. So the system figures out which robots need to go to which pods, pick it up, and bring it to the perimeter of the warehouse, where then people take things off of the pods and put them into the orders, which eventually go out. There's a lot of robots in action here. How come they don't collide? The robots have to generate trajectories and plans, and those plans are then shared to a coordinator, which then figures out how they should go and execute their plan so that they don't hit each other. Artificial intelligence has changed the world of medicine, and some old fears have crept back in. Should I be worried as a GP? The thought that a robot or artificial intelligence could take my job just seems crazy. I'm going to pose as a patient and give myself an imaginary condition. And then we can see just how accurate the machine really is. May I ask, please, what's troubling you today? I'm feeling tired all the time. So as well as feeling tired, I've been feeling kind of weak. Let's tell the computer that. Do you get breathless on exertion? Yes, I do. Thanks, I've noted this. So I've given the computer all of my symptoms now, and it's come up with a diagnosis. You can see that I have put down fibroids, and the computer has said uterine leiomyoma, which is actually the same thing. What those circles represent are diseases, symptoms and risk factors. And what those lines represent are the relationships between those. So based on that, the computer has taught itself actually how strongly related those diseases, symptoms and risk factors are. Okay, so that's how it determines the probability is from looking at past real life cases. Absolutely, and that's why this is machine learning. But machines can only learn from information that they're given. That information comes from us. Facebook today could not exist without AI. It's as simple as that. Over a billion people use Facebook every day, and they load their newsfeed a few dozen times every single day. And, you know, if you imagine how many people it would take, if you were to sort of line out all of the pieces of content that are available to you every day, if I had to sort of sort out how relevant is this story going to be to this person, you know, I'd multiply that by a billion people, we would sort of see, for a human, this would be a task that is absolutely impossible to do. Every day, AI technology is learning from us. The more it learns, the better it becomes. But as this rapid change is occurring, they're once again becoming fearful of its potential impact on our lives. They thought that this was going to be funny, and they were teenagers, and so they didn't think about the implications of deleting everything someone owns and how much kind of precious data you may have in your life. You know, I mean, data is quite precious to people now. It's valuable. Artificial intelligence. Machines that seemingly think for themselves are already changing our lives. Looking back through the Horizon Archive, it's clear our relationship with AI has followed a pattern. Optimism followed by a fear of the unknown. If this cycle continues, then today's anxieties about AI using our data in a bad way soon may evolve into something more positive. Scientists are again dreaming of the cities of the future and the wonderful things we'll be able to achieve. However, creating a robot that can truly think for itself may be decades or even centuries away, it may never happen. But the merging of man and machine, the organic

and the metal, human and AI, is happening right now, just as Isaac Asimov predicted. It seems to me that as robots become continually more advanced, that People will not try to keep it entirely a matter of metal and electrons. We will have human beings who will make more and more use of artificial organs of metal and plastic. Meet Eric Sauter. Deep inside his brain are two arrays of electrodes. I suffered A gunshot, which left me paralyzed from the shoulders down. Eric's spinal cord was severed, stopping the signals from his brain that control movement reaching his limbs. I'm A C3, C4 complete, quadriplegic complete. To try to restore movement he has lost, Eric is part of a trial to merge his brain with a robot arm. In short, we may have a society in which robots will drift away from total metal toward the organic, and human beings will drift away from the total organic toward the metal and plastic. And that somewhere in the middle, they may eventually meet. Now, when we have a kind of metal-organic hybrid. Will it matter that he was originally metal and became metal-organic, or that he was originally organic and became metal-organic? Or will it not matter? Will we then have formed a kind of mixed culture, which perhaps might be higher, more efficient, better? In the beginning, I was very conscious of them. Now, I completely forget they're there until somebody reminds me, like, what's that on your head? And I'm like, oh, yeah, that's right. I have two pedals, those are sticking out of my head. Now, Eric is ready to try to pick up a bottle of beer, using just his thoughts. It's a big moment. Let's do it. You ready? All right. Okay, here we go. He thinks only of the goal of the movement, bring hand to mouth, and the robot arm works out the rest. There you go. First step done. When you go reach for something, you don't walk it step by step, you just do it. Once the arm has grasped the bottle, Eric thinks, bring hand to mouth again. Is there anything essentially horrible about thinking that man has the right to create a pseudo-living system just as nature did? All right. Hey, you finished that thing off? That's good. His progress in this extraordinary trial has extended what it is to be human. The beginning was my brain, my arm, and the robotic arm. Now when I go in there, it's my brain and the arm. We are one. And it feels like my arm. I think the brain is, it's a part of us that it's ready to use any tool available to keep us moving forward and helping us live a better life.