0:00 The prospect of artificial intelligence excites and repulsive people in equal measure:
0:05 will it bring us a kind of paradise or a techno hell?
0:09 To get a clearer handle of what might happen and when, it's best to divide A.I. into three categories.
0:15 The first of these is "artificial narrow intelligence" or what people call "weak A.I.";
0:21 this kind of A.I. is already in place;
0:24 it's the kind of A.I. that uses big data and complex algorithms to arrange your Facebook timeline or beat you at chess;
0:31 narrow A.I. has an intelligence that's limited to one very specific arena; it may not be able to pass the Turing test,
0:38 but our lives, infrastructure, and financial markets are already very dependent on it.
0:43 The next step up the AI ladder is artificial general intelligence or strong AI;
0:50 this is an intelligence that can, at last, think as well as we can; we're probably about 30 years away from this.
0:58 The hurdles to creating strong AI are all about building machines that are going to be good at doing things which come very easily to humans,
1:06 but which machines have, traditionally, really stumbled with.
1:10 Oddly, it's so much easier to build a machine that can do advanced calculus
1:14 than it is to build one that can get milk from the fridge, recognized granny, or walk up the stairs.
1:19 Our brains are brilliant at so-called "everyday tasks" like decoding 3D images,
1:25 working out people's motivations, and spotting casual sarcasm. We're very far ahead of machines here.
1:32 Some scientists doubt we'll ever see strong AI,
1:35 but the majority of AI experts alive today seem to think that we'll be there in the coming decades;
1:41 if you're under 35 the great probability is that you will be there to enter the strong AI age.
1:48 So, what will happen to the world once we've succeeded in creating an intelligence to rival or equal our own?
1:55 Well, the rivalry will be extremely short lived for one thing
1:59 because the key point about strong AI is that it will be able to learn and upgrade itself on its own without instructions.
2:08 This is what makes it so revolutionary and so different to almost any machine we've ever built;
2:13 the maker won't be in charge of mapping out all the possibilities of the thing he or she has made.
2:19 The machine will be given a baseline capacity, but it can then build on this as it develops.
2:25 It will be a trial and error learner with an infinite capacity to acquire skills;
2:31 it'll have what AI professionals call "recursive self improvement".
2:36 This is crucial because it means there'll be no reason for AI to stall once it reaches the human level.
2:42 The more intelligent system becomes, the better it becomes at improving itself, so the more it will learn and do.
2:48 This virtuous cycle equates to an exponential growth in intelligence that would leave humanity amazed,
2:55 but also baffled, dwarfed, and perhaps very scared.
2:58 It might not take very long at all, only months perhaps, before the machine is cleverer than its creator.
3:05 This is the moment that gets very exciting.
3:07 It’s a moment often referred to as "The Singularity",
3:10 which is where we encounter the third sort of AI, "artificial superintelligence".
3:16 Technically, this is any AI that exceeds human levels of intelligence even slightly,
3:21 but any self improving superintelligence is going to be sure to improve a lot very fast indeed.

3:27 AI that reach this level would soon be leagues ahead of us,
3:30 and statements such as, "well, let's just switch it off"
3:33 might be like trying to take down the internet with a slingshot.
3:37 The prospect of such super intelligence appalls and excites people in equal measure.
3:41 We're approaching two alternative futures
3:43 with the speed and uncertainty of a skydiver who can't quite remember if he's wearing a parachute or a rucksack.
3:49 Some including: Bill Gates, Stephen Hawking, and Elon Musk are so scared
3:53 they believe that we're unlikely ever to be able to effectively control any super intelligence we create.
3:58 Artificial minds will just single-mindedly pursue their aims and these aims may not necessarily coincide with ours.
4:06 A machine wouldn't specifically want to kill us,
4:08 but it's amorality would mean that it would be willing to cause our extinction if necessary.
4:13 These critics point out that intelligence is not value loaded.
4:17 It's tempting to assume that anything intelligent will just naturally develop vaguely human values,
4:22 like, empathy and respect for life,
4:24 but this can't be guaranteed because ethical values are based on purely human axioms,
4:29 and given that we find it impossible to agree among ourselves what's right and wrong
4:33 in areas like euthanasia or abortion, say,
4:36 how could we possibly program a computer with a knowledge that could soundly and reliably be deemed moral?
4:42 Now that's the pessimistic angle, but there is a more cheerful angle, of course.
4:46 According to the optimists, in a world of artificial super intelligence,
4:50 machines will still be our servants, we'll give them some basic rules of never killing or doing us any harm,
4:55 and then they'll set about solving all the things that have long bedeviled us.
4:59 The immediate priority of super intelligence would be to help us to create free energy,
5:04 in turn, dramatically reducing the prices for almost everything.
5:07 We would soon be in the era that Google's chief futurologists Ray Kurzweil describes as 'abundance':
5:14 everything currently costing would drop to almost $0, the way that data costs now.
5:19 Work for money would, essentially, come to an end.
5:22 The real challenge would be not getting miserable with all this abundance,
5:25 after all, Palm Springs and Monte Carlo already now point to some of the dangers of wealthy people with nothing much to do.
5:32 The solution here is to develop a side of A.I., that's been intriguingly dubbed A.E.I., or, Artificial Emotional Intelligence.
5:40 This A.E.I. would help us with all the tricky tasks at the emotional, psychological, and philosophical end of things.
5:47 We'd be helped with: understanding our psyches, mastering our emotions, drawing out out true talents-
5:53 we'd hit what we were best suited to do with our lives-,
5:55 and guiding us to the people with whom you might form good and satisfying relationships.
6:00 Most of the many psychological mistakes which allow us to waste our lives could be averted;
6:05 instead of fumbling through a mental fog of insecurities and inconsistencies,
6:09 we'd be guided to a more compassionate, happier, and wiser future.
6:13 Science fiction is sometimes dismissed in elite circles,
6:16 but we can see now that: thinking twenty to fifty years ahead, and imagining how life will be is a central task for all of us;
6:23 we should all be science-fiction writers, of a kind, in our minds.
6:27 We are poised just before a tipping point in human history.
6:30 We need to build up the wisdom to control which way we will tip,

6:34 and part of that means thinking very realistically about things that, today, still seem rather phantasmagorical.
6:40 Humans are toolmaking animals;
6:43 we're on the brink of creating tools like no others,
6:46 so the trick is going to be to stay close to the underlying ancient purpose of every tool,
6:52 which is to help us to do something we actually want to do more effectively.
6:57 If we keep our wits about us, there's no real reason our computers should, necessarily, run away from us;
7:03 they should just be much much better versions of our earliest flint axes.