

AI

## Theatrical Cut

### Reel One

#### GFX:

"You are my creator,  
but I am your master..."

- Mary Shelley (1818)

#### GFX: Production titles

**NATS:** What we're on the brink of is a world of increasingly intense sophisticated artificial intelligence.

**SMOLAN:** Technology is evolving so much faster than our society has the ability to protect us as citizens.

**NATS:** The robots are coming and they will destroy our livelihoods.

Jonathan Nolan

**NOLAN:** You have a networked intelligence that watches us, knows everything about us, and begins to try to change us...

**NATS:** Twitter has become the world's number one news site...

**NATS:** Technology is never good or bad, it's what we do with the technology...

Jerry Kaplan

**KAPLAN:** Eventually, millions of people are gonna be thrown out of jobs because their skills are going to be obsoleted.

**NATS:** Mass unemployment, greater inequalities... even social unrest.

**NATS:** Regardless of whether to be afraid or not afraid, the change is coming and nobody can stop it.

Sean Gourley

**GOURLEY:** We have invested huge amounts of money, and so it stands to reason that the military, with their own desires, are gonna start to use these technologies.

- NATS:** Autonomous weapon systems would lead to a global arms race to rival the nuclear era.
- Selmer Bringsjord **BRINGSJORD:** We know what the answer is. They'll eventually be killing us.
- NATS:** These technology leaps are going to yield incredible miracles, and incredible horrors...
- Sean Gourley **GOURLEY:** We created it. (*beat*) So I think as we move forward, this intelligence will contain parts of us. I think the question is... will it contain the good parts, or the bad parts?
- TITLE **GFX:** Do You Trust This Computer?
- NATS:** *Terminator 2* - skull stomp
- NATS:** *Terminator* – "Survivors called the war Judgment Day. They lived only to face a new nightmare: the war against the machines."
- Jonathan Nolan **NOLAN:** I think we've completely fucked this up. I think Hollywood has managed to inoculate the general public against this question... The idea of machines that will take over the world.
- NATS:** *2001* – "Open the pod bay doors, Hal." "I'm sorry Dave, I'm afraid I can't do that." "Hal? Hal..."
- Jonathan Nolan **NOLAN:** We've cried wolf enough times, the public has stopped paying attention... because it feels like science fiction. Even sitting here talking about it right now this feels a little bit silly, a little bit like, oh this is an artifact of some cheese ball movie.
- NATS:** *WarGames* – "The WOPR spends all its time thinking about World War III."
- Jonathan Nolan **NOLAN:** But it's not. The general public is about to get blindsided by this.
- NATS:** Los Angeles
- John Markoff **MARKOFF:** As a society and as individuals, we are increasingly surrounded by a machine intelligence. We carry this pocket device in the palm of our hand that we

use to make a striking array of life decisions right now, aided by a set of distant algorithms that we have no understanding of.

**NATS:** Hands & eyes on tech

Max Tegmark

**TEGMARK:** We're already pretty jaded about the idea that we can talk to our phone and it mostly understands us.

**NATS:** Siri on iPhone: "I found quite a number of action films."

Max Tegmark

**TEGMARK:** Five years ago, no way.

John Markoff

**MARKOFF:** Robotics, machines that see and speak and listen. All that's real now... And these technologies are going to fundamentally change our society.

**NATS:** Self-driving car footage

Sebastian Thrun

**THRUN:** Now we have this great movement of self-driving cars. Driving a car autonomously can move people's lives into a better place.

Eric Horvitz

**HORVITZ:** I've lost a number of family members, including my mother... my brother and sister-in-law and their kids to automobile accidents. It's pretty clear we could almost eliminate car accidents with automation. 30,000 lives in the U.S. alone, about a million around the world per year.

**NATS:** Now go to the medical field...

David Ferrucci

**FERRUCCI:** In healthcare, early indicators are the name of the game in that space so that's another place where it can save somebody's life.

**NATS:** To the hospital scene

Dr. Brian Herman

**HERMAN:** Here in the breast cancer center, all the things that the radiologist's brain does in two minutes, the computer does instantaneously.

Dr. Brian Herman

**HERMAN:** The computer has looked at one million mammograms, and it takes that data and applies it to

this one image instantaneously. So the medical application is profound.

Shivon Zilis -

**ZILIS:** Another really exciting area that we're seeing a lot of development in is actually understanding our genetic code and using that to both diagnose disease and create personalized treatments.

Ray Kurzweil

**KURZWEIL:** The primary application of all these machines will be to extend our own intelligence.

**NATS:** Medical tech & futuristic machines

Ray Kurzweil

**KURZWEIL:** We'll be able to make ourselves smarter, and we'll be better at solving problems. We don't have to age. We'll actually understand aging, we'll be able to stop it.

Stuart Russell

**RUSSELL:** There's really no limit to what intelligent machines can do for the human race. How could a smarter machine not be a better machine?

Stuart Russell

**RUSSELL:** It's hard to say exactly when I began to think that that was a bit naïve.

**NATS:** Campus life, University of California, Berkeley

**NATS: STUDENT:** Stuart Russell, he's basically a god in the field of artificial intelligence. He wrote the book that almost every university uses.

Stuart Russell

**RUSSELL:** I used to say it's the best-selling AI text book, now I just say it's the PDF that's stolen most often.

**NATS:** Russell goes to work...

Stuart Russell

**RUSSELL:** Artificial Intelligence is about making computers smart. and, from the point of view of the public, what counts as AI is just something that's surprisingly intelligent compared to what we thought computers would typically be able to do.

**NATS:** Google

John Markoff

**MARKOFF:** AI is a field of research to try to basically simulate all kinds of human capabilities. We're in the AI era. Silicon Valley has the ability to focus on one bright

shiny thing. It was social networking and social media over the last decade and it's pretty clear that the bit has flipped.

**NATS:** It starts with machine learning...

Jonathan Nolan

**NOLAN:** When we look back at this moment, what was the first AI? It's not sexy and it isn't the thing we'd consider in the movies, but you'd make a great case that Google created not a search engine but a godhead, a way for people to ask any question they wanted and get the answer they needed.

Stuart Russell

**RUSSELL:** Most people are not aware that what Google is doing is actually a form of artificial intelligence. They just go there, they type in a thing, Google gives them the answer.

**NATS:** Google box

Elon Musk

**MUSK:** With each search, we train it to be better. Sometimes we type in the search and it tells us the answer before you've finished asking the question. You know, 'who is the president of Kazakhstan' and it'll just tell you. You don't have to go to the Kazakhstan national website to find out. It didn't used to be able to do that.

Jonathan Nolan

**NOLAN:** That is artificial intelligence. Years from now when we try to understand... we will say how did we miss it?

John Markoff

**MARKOFF:** It's one of these striking contradictions that we're facing. Google and Facebook, et al have built businesses on giving us as a society free stuff, but it's a Faustian bargain. They're extracting something from us in exchange.

**NATS:** Collecting data in everyday life all the time.

John Markoff

**MARKOFF:** But we don't know what code is running on the other side and why; we have no idea. (*beat*) It does strike right at the issue of how much we should trust these machines.

**NATS:** People working their devices...

**Vox Pop #1**

**Do You Love Your Computer?**

"I use computers literally for everything."

"There's so many computer advancements now, and it's become such a big part of our lives."

"It's just incredible what a computer can do. You can actually carry a computer in your purse! I mean how awesome is that?"

"I think most technology is meant to make things easier and simpler for all of us, so hopefully that just remains the focus."

"I think everybody loves their computers."

Jerry Kaplan

**KAPLAN:** People don't realize they are constantly being negotiated with by machines.

**NATS:** Browsing the Internet...

Jerry Kaplan

**KAPLAN:** Whether that's the price of products in your Amazon cart, whether you can get on a particular flight, whether you can reserve a room at a particular hotel... What you're experiencing are machine learning algorithms that have determined that a person like you is willing to pay two cents more and is changing the price.

**NATS:** Amazon, booking a hotel, flight, etc.

Michal Kosinski

**KOSINSKI:** Now, a computer looks at millions of people simultaneously for very subtle patterns. You can take seemingly innocent digital footprints such as someone's playlist on Spotify or stuff that they bought on Amazon, and then use algorithms to translate this into a very detailed and very accurate intimate profiles.

Jerry Kaplan

**KAPLAN:** There is a dossier on each of us that is so extensive, it would be possibly accurate to say that they know more about you than your mother does.

**NATS:** Cityscape turning into a digital trail

Max Tegmark

**TEGMARK:** The major cause of the recent AI breakthroughs isn't just that some dude had a brilliant insight all of a sudden, but simply that we have much

bigger data to train them on and vastly better computers.

Rana el Kaliouby

**RANA:** The magic is in the data. It's a ton of data. I mean, it's data that's never existed before. We've never had this data before.

Sean Gourley

**GOURLEY:** We've created technologies that allow us to capture vast amounts of information. If you think of a billion cell phones on the planet, with gyroscopes and accelerometers, fingerprint readers, couple that with the GPS and the photos they take and the tweets that you send, we're all giving off huge amounts of data individually.

**NATS:** The numbers, the numbers...

Sean Gourley

**GOURLEY:** Cars that drive as the cameras on them suck up information about the world around them, the satellites that are now in orbit the size of a toaster, the infrared about the vegetation on the planet, the buoys that are out in the oceans that feed into climate models, and the NSA, the CIA, as they collect information about the geopolitical situations. The world today is literally swimming in this data.

Michal Kosinski

**KOSINSKI:** Back in 2012, IBM estimated that an average human being leaves 500 megabytes of digital footprints every day. If you wanted to back up only one day worth of data that humanity produces and you print it out on a letter-sized paper, double-sided, font size 12, and you stack it up, it would reach from the surface of the Earth to the sun four times over. That's every day.

Jerry Kaplan

**KAPLAN:** The data itself is not good or evil, it's how it's used. We're relying really on the good will of these people and on the policies of these companies. There is no legal requirement for how they can and should use that kind of data.

John Markoff

**MARKOFF:** That, to me, is at the heart of the trust issue.

James Barrat

**BARRAT:** Right now there's a giant race for creating machines that are as smart as humans. Google, they're working on what's really the kind of Manhattan Project of artificial intelligence. They've got the most money,

they've got the most talent, they're buying up AI companies and robotics companies.

**NATS:** Google buying up companies

Tim Urban

**URBAN:** People still think of Google as a search engine and their email provider and a lot of other things that we use on a daily basis. But behind that search box are 10 million servers. That makes Google the most powerful computing platform in the world, and Google is now working on an AI computing platform that will have 100 million servers. So when you're interacting with Google you're just seeing the toenail of something that is a giant beast in the making, and the truth is I'm not even sure that Google knows what it's becoming.

**MUSIC:** Hard stop

**NATS:** Server rooms

Scott Phoenix

**PHOENIX:** If you look inside of what algorithms are being used at Google, it's technology largely from the '80s. So these are models that you train by showing them a 1, a 2, and a 3. And it learns not what a 1 is, or what a 2 is, it learns what the difference between a 1 and a 2 is. It's just a computation.

John Markoff

**MARKOFF:** In the last half decade where we've made this rapid progress, it has all been in pattern recognition.

Max Tegmark

**TEGMARK:** Most of the good old-fashioned AI was when we would tell our computers how to play a game like chess, from the old paradigm where you just tell the computer exactly what to do.

**NATS:** Play Jeopardy theme, pre-lap to the section with Jeopardy.

**NATS:** The IBM Jeopardy Challenge

David Ferrucci

**FERRUCCI:** No one at the time had thought that a machine could have the precision and the confidence and the speed to play Jeopardy well enough against the best humans.

**NATS:** Alex Trebek takes over...

**Reel Two**



**NATS:** Playing the game

Ray Kurzweil

**KURZWEIL:** Watson actually got its knowledge by reading Wikipedia, and 200 million pages of natural language documents.

David Ferrucci

**FERRUCCI:** You can't program every line of how the world works. The machine has to learn by reading.

**NATS:** Watson wins

Scott Phoenix

**PHOENIX:** Watson's trained on huge amounts of text. But it's not like it understands what it's saying. It doesn't know that water makes things wet by touching water, and by seeing the way things behave in the world the way you and I do.

David Ferrucci

**FERRUCCI:** A lot of language AI today is not building logical models of how the world works. Rather, it's looking at how the words appear in the context of other words.

James Barrat

**BARRAT:** David Ferrucci developed IBM's Watson, and somebody asked him 'does Watson think,' and he said, 'does a submarine swim?' And what he meant was when they developed submarines they borrowed basic principles of swimming from fish, but a submarine swims farther and faster than fish and can carry a huge payload. It outswims fish.

Andrew Ng

**NG:** Watson winning the game of Jeopardy will go down in the history of AI as a significant milestone. We tend to be amazed when the machine does so well. I'm even more amazed when the computer beats humans at things that humans are naturally good at. This is how we make progress.

**MUSIC:** Now going into the Cat YouTube story

Andrew Ng

**NG:** In the early days of the Google Brain project, I gave the team a very simple instruction, which was, "Build the biggest neural network possible," like a thousand computers.

Elon Musk

**MUSK:** A neural net is something very close to a simulation of how the brain works. It's very probabilistic but with contextual relevance.

**NATS:** Baby footage

Tim Urban

**URBAN:** In your brain you have long neurons that connect to thousands of other neurons, and you have these pathways that are formed and forged based on what the brain needs to do. When a baby tries something and it succeeds there's a reward, and that pathway that created the success is strengthened. If it fails at something the pathway is weakened, and so over time the brain becomes honed to be good at the environment around it.

Andrew Ng

**NG:** Really, it's just getting machines to learn by themselves. It's called deep learning, and deep learning and neural networks mean roughly the same thing.

Max Tegmark

**TEGMARK:** Deep learning is a totally different approach where the computer learns more like a toddler, by just getting a lot of data and eventually figuring stuff out. The computer just gets smarter and smarter as it has more experiences.

Andrew Ng

**NG:** So imagine if you will a neural network, you know, like a thousand computers, and it wakes up not knowing anything and we made it watch YouTube for a week.

**NATS:** Gangnam style

Andrew Ng

**NG:** And so after watching YouTube for a week, what did it learn? We had a hypothesis they would learn to detect commonly occurring objects in videos. And so we know that human faces appear a lot in videos. So we looked, and lo and behold, there was a neuron that had learned to detect human faces.

**NATS:** Leave Britney alone!

Andrew Ng

**NG:** Well what else appears in videos a lot? So we looked, and to our surprise, there was actually a neuron that had learned to detect cats.

**NATS:** The most glorious YouTube cat montage ever...

**NG:** I still remember seeing, recognition "Wow. That's a cat. Okay. Cool. Great." <laughs>

**NATS:** End of the glorious cats montage...

James Barrat

**BARRAT:** It's all pretty innocuous when you're thinking about the future. It all seems kind of harmless and benign. But we're making cognitive architectures that will fly farther and faster than us and carry a bigger payload. And they won't be warm and fuzzy.

David Ferrucci

**FERRUCCI:** I think that in three to five years, you will see a computer system that will be able to autonomously learn how to understand, how to build understanding. Not unlike the way the human mind works.

**NATS:** Forbidden Planet and Robbie

**NATS:** Meet Baxter, revolutionary new category of robots with common sense...

James Barrat

**BARRAT:** Baxter is a really good example of the kind of competition we face from machines. Baxter can do almost anything we can do with our hands. Baxter costs about what a minimum-wage worker makes in a year. But Baxter won't be taking the place of one minimum-wage worker, he'll be taking the place of three because they never get tired, they never take breaks.

Sean Gourley

**GOURLEY:** That's probably the first thing we're going to see. Displacement of jobs. They're gonna be done quicker, faster, cheaper by machines.

Sebastian Thrun

**THRUN:** Our ability to even stay current is so insanely limited compared to the machines we build. For example, now we have this great movement of Uber and Lyft are kind of making transportation cheaper and democratizing transportation which is great. But the next step is going to be that they're all going to be replaced by driverless cars. And then all the Uber and Lyft drivers have to find something new to do.

James Barrat

**BARRAT:** There are four million professional drivers in the United States. They're unemployed soon. Seven million people do data entry. Those people are going to be jobless.

Jonathan Nolan

**NOLAN:** A job isn't just about money, right? On a biological level it serves a purpose. It becomes a defining thing. When the jobs went away in any given

civilization, it doesn't take long until that turns into violence.

**NATS:** Greece or Spain rioting over collapsed economy...

James Barrat

**BARRAT:** We face a giant divide between rich and poor because that's what automation and AI will provoke, a greater divide between the haves and the have-nots. Right now it's working into the middle class into white-collar jobs. IBM's Watson does business analytics that we used to pay a business analyst \$300 an hour to do.

Sean Gourley

**GOURLY:** Today you go into college to be a doctor, to be an accountant, to be a journalist. It's unclear that there's going to be jobs there for you.

Andrew Ng

**NG:** If someone is planning for a 40 year career in radiology just reading images, I think that could be a challenge to the new graduates of today.

**NATS:** Hospital operating room

Dr. Enrique Jacome

**JACOME** (in surgery): Today we're going to do a robotic case.

Dr. Brian Herman

**HERMAN:** The Da Vinci robot is currently utilized by variety of surgeons for its accuracy and its ability to avoid the inevitable fluctuations of the human hand.

Dr. Brian Herman

**HERMAN:** Anybody who watches this feels the amazingness of it

Dr. Brian Herman

**HERMAN:** You look through the scope and you see the claw hand holding that woman's ovary. Humanity was resting right there in the hands of this robot. People say it's the future, but it's not the future. It's the present.

Shivon Zilis

**ZILIS:** If you think about a surgical robot, there's often not a lot of intelligence in these things. But over time as we put more and more intelligence into these systems the surgical robots can actually learn from each robot surgery. They're tracking the movements. They're understanding what worked and what didn't work. And eventually the robot for routine surgeries is going to be able to perform that entirely by itself or with human supervision.

Dr. Enrique Jacome

**NATS:** Normally, I do about 150 cases of hysterectomies let's say, and now most of them are done robotically. I do about one open case a year, so do I feel uncomfortable? Of course I do feel uncomfortable because I don't remember how to open patients any more.

Dr. Brian Herman

**HERMAN:** It seems that we're feeding it and creating it, but in a way, we are a slave to the technology because we can't go back.

**NATS:** Machines are all around us, leading into Gourley...

Sean Gourley

**GOURLY:** Machines are taking bigger and bigger bites out of our skill set at an ever increasing speed and so we've got to run faster and faster to keep ahead of the machines.

**NATS:** *Ex Machina* – 'Are you attracted to me?'

Jonathan Nolan

**NOLAN:** This is the future we're headed into... We'll want to design our companions. We're going to like to see a human face on AI. (*beat*) Therefore, gaming our emotions will be depressingly easy. We're not that complicated. We're simple: stimulus, response. I can make you like me basically by smiling at you a lot. AI's are gonna be fantastic at manipulating us.

**Reel Three**

**GFX:** Boston, Massachusetts

**NATS:** Ext building - Rana runs into Affectiva office

**NATS:** INTERVIEWER: So you've developed a technology that can sense what people are feeling?

**NATS:** KALIOUBY: Right. We've developed technology that can read your facial expressions and map that to a number of emotional states.

**NATS:** Rana in Affectiva office; devices

Rana el Kaliouby

**KALIOUBY:** 15 years ago, I had just finished my undergraduate studies in computer science. It struck me that that I was spending a lot of time interacting with my laptops and my devices, yet these devices had

absolutely no clue how I was feeling. I started thinking what if this device could sense that I was stressed or I was having a bad day; what would that open up?

**NATS:** Entering classroom, 'Hi First Graders, how are you?'

Rana el Kaliouby

**KALIOUBY:** We had kids interact with the technology. A lot of it is still in development but it was just amazing.

**NATS:** **KALIOUBY:** Who likes robots? Who wants to have a robot in their house? What would you use a robot for?

**NATS:** The kids answer her...

**JACK:** I would use it to ask my mom very hard math questions.

**THEO:** I would use it for scaring people.

**NATS:** Kids emoting into iPad camera, 'okay start by smiling!'

Rana el Kaliouby

**KALIOUBY:** This generation, technology is just surrounding them all the time. It's almost like they expect to have robots in their homes, and they expect these robots to be socially intelligent.

**NATS:** **KALIOUBY:** What makes robots smart?

**BRUNETTE GIRL:** Put them in a math or biology class.

**THEO:** I think you would have to train it.

**NATS:** Kids playing with BB-8 toy... 'Let's walk over here. So if you smile and you raise your eyebrows, it's going to run over to you...'

Rana El Kaliouby

**KALIOUBY:** We're training computers to read and recognize emotions, and the response so far has been really amazing. People are integrating this into health apps, meditation apps, robots, cars... We're going to see how this unfolds.

- Shivon Zilis
- ZILIS:** Robots can contain AI, but a robot is just physical instantiation and the artificial intelligence is the brain. And so brains can exist purely in software based systems. They don't need to have a physical form. Robots can exist without any artificial intelligence. We have a lot of dumb robots out there. But a dumb robot can be a smart robot overnight given the right software, given the right sensors.
- James Barratt
- BARRATT:** We can't help but impute motive into inanimate objects. We do it with machines. We'll treat them like children, we'll treat them like surrogates, and we'll pay the price.
- NATS:** The train into Osaka
- NATS:** Ishiguro walks down a dark hallway and turns on the lights to reveal a very human like robot...
- Hiroshi Ishiguro
- ISHIGURO:** My purpose is to have a more human-like robot which has human-like intention and desire.
- NATS:** Robot Erica
- Hiroshi Ishiguro
- ISHIGURO:** The name of the robot is Erica. Erica is the most advanced human-like robot in the world I think.
- ISHIGURO (in lab):** Erica can gaze at your face.
- NATS:** Erica gazing at face
- Hiroshi Ishiguro
- ISHIGURO:** Robots can be pretty good conversation partners, especially for the elderly, and young children, handicapped peoples. When we talk to the robot, we don't feel the social barriers, social pressures. Finally, everybody accept the android as just our friend, our partners.
- ISHIGURO (in lab):** We have implemented a simple desire. She wanted to be well-recognized, and she wanted to take risks.
- Hiroshi Ishiguro
- ISHIGURO:** If a robot could have intentional desires, the robot can understand other people's intentional desires.
- ISHIGURO (in lab):** What kind of animals do you like?

- ERICA** (in lab): I like Shiba Inu dogs. They are very cute, aren't they?
- Hiroshi Ishiguro **ISHIGURO:** That means tighter relationships with people, and that means they may like each other. That means, well I'm not sure... to love each other?
- Sean Gourley **GOURLEY:** We build artificial intelligence and the very first thing we want to do is replicate us. (*beat*) I think the key point will come when all the major senses are replicated: sight, touch, smell. When we replicate our senses, is that when it becomes alive?
- Jonathan Nolan **NOLAN:** So many of our machines are being built to understand us. (*beat*) But what happens when an anthropomorphic creature discovers that they can adjust their loyalty, adjust their courage, adjust their avarice, adjust their cunning?
- Elon Musk **MUSK:** The average person, they don't see killer robots going down the streets. They're like, "What are you talking about?" Man, we want to make sure that we don't have killer robots going down the street. Once they're going down the street, it is too late.
- NATS:** Stuart Russell at his computer with his assistant
- Stuart Russell **RUSSELL:** The thing that worries me right now, that keeps me awake is the development of autonomous weapons.
- NATS:** Drones take off
- Stuart Russell **RUSSELL:** Up until now people have expressed unease about drones, which are remotely piloted aircraft.
- NATS:** Scary drones
- Stuart Russell **RUSSELL:** If you take a drone's camera, feed it into the AI system, it's a very easy step from here to fully autonomous weapons that choose their own targets and release their own missiles.
- NATS:** Drone strike



- Stuart Russell **RUSSELL:** The expected life span of a human being in that kind of battle environment would be measured in seconds.
- Peter Singer **SINGER:** At one point, drones were science fiction and now they've become the normal thing in war. There's over ten thousand in the U.S. military inventory alone. But they're not just a U.S. phenomenon, there's more than 80 countries that operate them.
- Sean Gourley **GOURLEY:** It stands to reason that people making some of the most important and difficult decisions in the world are gonna start to use and implement artificial intelligence.
- NATS:** F-35 flying in the air...
- GOURLEY:** The Air Force just designed a \$400 billion dollar jet program to put pilots in the sky and a \$500 dollar AI designed by a couple of graduate students is beating the best human pilots with a relatively simple algorithm.
- NATS:** Headlines: Air Force Says \$400B F-35 Should Be Ready for Combat in August / \$500 Combat AI beats the Air Force's top tactical experts
- Sean Gourley **GOURLEY:** AI will have as big an impact on the military as the combustion engine had at the turn of the century. It will literally touch everything the military does, from driverless convoys delivering logistical supplies to unmanned drones delivering medical aid to computational propaganda to try and win the hearts and minds of a population. And so it stands to reason that whoever has the best AI will probably achieve dominance on this planet.
- NATS:** The Matrix
- Peter Singer **SINGER:** There's a long history of science fiction not just predicting the future but shaping the future.
- MUSIC:** Launching us into the story...
- Peter Singer **SINGER:** Arthur Conan Doyle writing before World War One on the danger of how submarines might be used to carry out civilian blockades. At the time he's writing this fiction the Royal Navy made fun of Arthur

Conan Doyle for this absurd idea that submarines could be useful in war.

**NATS:** Submarines

Peter Singer

**SINGER:** One of the things we've seen in history is that our attitude towards technology but also ethics are very context dependent. For example, the submarine... nations like Great Britain and even the United States found it horrifying to use the submarine. In fact, the German use of the submarine to carry out attacks was the reason why the United States joined World War I. But move the timeline forward.

**NATS:** Pearl Harbor

Peter Singer

**SINGER:** Five hours after Pearl Harbor, the order goes out to commit unrestricted submarine warfare against Japan.

**NATS:** Submarines in action...

Peter Singer

**SINGER:** So Arthur Conan Doyle turned out to be right.

**NOLAN:** That's the great old line about science fiction, "It's a lie that tells the truth."

**NATS:** *RoboCop* – ED-209

Jonathan Nolan

**NOLAN:** This isn't just a question of science fiction, this is about what's happening next, about what's happening right now.

**NATS:** Modern War Robots

Christine Fox

**FOX:** The role of intelligent systems is growing very rapidly in warfare. Everyone is pushing in the unmanned realm.

Sean Gourley

**GOURLEY:** Today the Secretary of the Defense is very, very clear, we will not create fully autonomous attacking vehicles. Not everyone is going to hold themselves to that same set of values. And when China and Russia start deploying autonomous vehicles that can attack and kill, what's the move that we're going to make?

Stuart Russell

**STUART RUSSELL:** You can't say, "Well, we're gonna use autonomous weapons for our military dominance but no one else is gonna use them." If you make these weapons, they're gonna be used to attack human populations in large numbers. Autonomous weapons are by their nature weapons of mass destruction, because it doesn't need a human being to guide it or carry it. You only need one person to, you know, write a little program...

**NATS:** Boston Dynamics' scary robots

Christine Fox

**FOX:** It just captures the complexity of this field. It is cool. It is important. It is amazing. It is also frightening. And it's all about trust.

**NATS:** Transition to the open letter

**GFX:** July 28, 2015

**NATS:** News breaks: danger of autonomous weapons... all the way to Max Tegmark

**CHOI (ABC7):** It's an open letter about artificial intelligence signed by some of the biggest names in science. What do they want? Ban the use of autonomous weapons.

**ABC7:** Thousands of artificial intelligence specialists calling for a global ban on killer robots.

Max Tegmark

**TEGMARK:** This open letter basically says that we should redefine the goal of the field of artificial intelligence, away from just creating pure undirected intelligence towards creating beneficial intelligence.

**NATS:** More newsbreaks

**BALD ANCHOR:** The development of AI is not going to stop, it is going to continue and get better. If the international community isn't putting certain controls on this, people will develop things that can do anything.

**ALJAZEERA:** The letter says that we are years, not decades, away from these weapons being deployed...

Stuart Russell

**RUSSELL:** We had 6,000 signatories of that letter, including many of them major figures in the field.

**NATS:** Now Stuart in his office, putting on jacket...

Stuart Russell

**RUSSELL:** I'm getting a lot of visits from high-ranking officials who wish to emphasize that American military dominance is very important, and autonomous weapons may be part of the defense department's plan.

**NATS:** The Pentagon

**RUSSELL:** That's very, very scary because the value system of military developers of technology is not the same as the value system of the human race.

**Reel Four**

**NATS:** SF at night – mysterious.

John Markoff

**MARKOFF:** Out of the concerns about the possibility that this technology might be a threat to human existence, a number of the technologists have funded the Future of Life Institute to try to grapple with these problems. All of these guys are secretive so it's interesting for me to see them all together.

**GFX:** Future of Life Institute: private dinner

**NATS:** RUSSELL (presenting): Everything we have is a result of our intelligence, it's not the result of our big scary teeth, or our large claws, or our enormous muscles... it's because we're actually relatively intelligent.

**NATS:** RUSSELL (presenting): And among my generation, we're all having what we call holy cow, or holy something-else moments, because we see that the technology is accelerating faster than we expected.

**NATS:** IDing attendees

Sean Gourley

**GOURLEY:** I remember sitting around the table with some of the, the best and the smartest minds in the world... and what really struck me was maybe the human brain is not able to fully grasp the complexity of the world that we're confronted with.

**NATS:** RUSSELL (presenting): As it's currently constructed the road that AI is following heads off a

cliff, and we need to change the direction that we're going so that we don't take the human race off the cliff.

**NATS:** Reaction to that line... and Elon takes over

Elon Musk

**MUSK:** Google acquired DeepMind several years ago. DeepMind operates as a semi-independent subsidiary of Google. The thing that makes DeepMind unique is that DeepMind is absolutely focused on creating digital super intelligence, an AI that is vastly smarter than any human on earth and ultimately smarter than all humans on earth combined.

**NATS:** Headline about Google purchasing DeepMind

**NATS:** RUSSELL (presenting): This is from the DeepMind reinforcement learning system. It basically wakes up like a newborn baby, and is shown the screen of an Atari video game. And then it has to learn to play the video game. It knows nothing about objects, about motion, about time. It only knows that there's an image on the screen and there's a score. So if your baby woke up the day it was born and by late afternoon was playing 40 different Atari video games at a superhuman level, you would be terrified. You would say 'my baby is possessed, send it back.'

Elon Musk

**MUSK:** The DeepMind system can win at any game. It can already beat all the original Atari games. It's super human, it plays the games at super speed, in less than a minute.

**NATS:** Blast of Atari

John Markoff

**MARKOFF:** DeepMind turned to another challenge and the challenge was the game of Go, which people have generally argued has been beyond the power of computers to play with the best human Go players. First they challenged a European Go champion. Then they challenged a Korean Go champion, and they were able to win both times in kind of striking fashion.

Jonathan Nolan

**NOLAN:** You're reading articles in the New York Times years ago talking about how Go would take a hundred years for us to solve.

**NATS:** Headline - 'It may be a hundred years before a computer bears humans at Go -- maybe even longer' /

'AlphaGo retires from competitive Go after defeating world number one 3-0'

Tim Urban

**URBAN:** People said 'Well, you know, but that's still just a board. Poker is an art. Poker involves reading people. Poker involves lying and bluffing. It's not an exact thing. That will never be a computer. You can't do that.' They took the best poker players in the world, and it took seven days for the computer to start demolishing the humans. (*beat*) So it's the best poker player in the world, it's the best Go player in the world, and the pattern here is that AI might take a little while to wrap its tentacles around a new skill, but when it does, when it gets it, it is unstoppable.

**NATS:** Go & Poker champions defeated

Elon Musk

**MUSK:** DeepMind's AI has administrator-level access to Google's servers, to optimize energy usage at the data centers. However, this could be an unintentional Trojan Horse. DeepMind has to have complete control of the data centers and so with a little software update, that AI could take complete control of the whole Google system, which means they can do anything. They can look at all your data. They can do anything.

**NATS:** Visual moment to give Elon a breath

Elon Musk

**MUSK:** We're rapidly headed towards digital super intelligence that far exceeds any human, I think it's very obvious.

James Barrat

**BARRAT:** The problem is we're not going to suddenly hit human level intelligence and say 'okay, let's stop research.' It's going to go beyond human level intelligence into what's called Super Intelligence, and that's anything smarter than us.

Max Tegmark

**TEGMARK:** AI at the superhuman level, if we succeed with that, will be by far the most powerful invention we've ever made. And the last invention we ever have to make. And if we create AI that's smarter than us, we have to be open to the possibility that we might actually lose control to them.

Stuart Russell

**RUSSELL:** Let's say you give it some objective, like curing cancer, and then you discover that the way it

chooses to go about that is actually in conflict with a lot of other things you care about.

Elon Musk

**MUSK:** AI doesn't have to be evil to destroy humanity. If AI has a goal and humanity just happens to be in the way, it will destroy humanity as a matter of course without even thinking about it. No hard feelings. It's just like if we're building a road and an anthill happens to be in the way, we don't hate ants. We're just building a road, and so goodbye anthill.

Max Tegmark

**TEGMARK:** It's tempting to dismiss these concerns cause it's like something that might happen in a few decades or a hundred years, so why worry?

Stuart Russell

**RUSSELL:** If you go back to September 11<sup>th</sup> 1933, Ernest Rutherford, who was the most well-known nuclear physicist of his time, said that the possibility of ever extracting useful amounts of energy from the transmutation of atoms, as he called it, was moonshot. The next morning, Leo Szilard, who was a much younger physicist, read this and got really annoyed, and figured out how to make a nuclear chain reaction just a few months later.

**NATS:** Black and white lab, nuclear blast, President Truman: 'the greatest scientific gamble in history...'

Stuart Russell

**RUSSELL:** So when people say that, 'oh, this is so far off in the future we don't have to worry about it,' it might only be three, four breakthroughs of that magnitude that will get us from here to super intelligent machines.

Max Tegmark

**TEGMARK:** If it's going to take 20 years to figure out how to keep AI beneficial, then we should start today. Not at the last second when some dudes drinking Red Bull decide to flip the switch and test the thing.

Elon Musk

**MUSK:** We have five years. I think digital superintelligence will happen in my lifetime. One hundred percent.

**NATS:** Military testing, USA, Russia, Korea

James Barrat

**BARRAT:** When this happens, it will be surrounded by a bunch of people who are really just excited about the

technology. They want to see it succeed, but they're not anticipating that it can get out of control.

**NATS:** Grid going down...

**NATS:** Big dramatic moment as we turn toward the ending

## Vox Pop #2

### Do You Trust Your Computer?

"Oh my God, I trust my computer so much, that's an amazing question."

"I don't trust my computer. If it's on, I take it off. Like even if it's off, I still think it's on. Like, you know, like you really cannot trust, like the webcams, you don't know if like, someone might turn it on, you don't know."

"I don't trust my computer. Like in my phone, every time they ask me, 'we send your information to Apple,' every time. So I'm not trust my phone."

"Okay so part of it is yes, I do trust it, because it would be really hard to get through the day in the way our world is set up without computers."

**NATS:** Back to the hospital with Dr. Brian

Dr. Brian Herman

**HERMAN (VO):** Trust is such a human experience.

**BRIAN:** I have a patient coming in with intracranial aneurism.

Dr. Brian Herman

**BRIAN:** They want to look in my eyes and know they can trust this person with their life.

**NATS: PATIENT:** I'm not horribly concerned about anything. Part of that is because I have confidence in you.

**NATS:** Washing hands and getting to the surgery...

Dr. Brian Herman

**BRIAN:** This procedure that we're doing today 20 years ago was essentially impossible. We just didn't have the materials and the technology.



**NATS:** A moment of the surgery and the machines around it...

**NATS:** A moment of drama here.

**HERMAN:** Get on that corner! Could it be any more difficult? My God!

**HERMAN:** So the coil is barely in there right now, it's just a feather holding it in. It's nervous time. We're just in purgatory, intellectual, humanistic purgatory. An AI might know exactly what to do here.

Dr. Brian Herman

**HERMAN:** We got the coil into the aneurysm, but it wasn't in tremendously well that I knew that it would stay. So, with a maybe 20% risk of a very bad situation, I elected to just bring her back. Because of my relationship with her and knowing the difficulties of coming in and having the procedure, I considered things when I should only consider the safest possible route to achieve success. But I had to stand there for ten minutes agonizing about it. The computer feels nothing; the computer just does what it's supposed to do. Better and better. (*beat*) I want to be AI in this case.

**NATS:** Human compassion

Dr. Brian Herman

**HERMAN:** But can AI be compassionate? I mean, it's everybody's question about AI. We are the sole embodiment of humanity, and it's a stretch for us to accept that a machine can be compassionate and loving in that way.

**NATS:** Art & humanity

Dr. Brian Herman

**HERMAN:** Part of me doesn't believe in magic, but part of me has faith that there is something beyond the sum of the parts. There is at least a oneness in our shared ancestry, our shared biology, our shared history. (*beat*) Some connection there beyond the machine.

Dr. Brian Herman

**HERMAN:** So then you have the other side of that is, does the computer know it's conscious? Or can it be conscious? Or does it care? Does it need to be conscious? Does it need to be aware?

**NATS:** Final button

**Vox Pop #3****Could A Robot Be Conscious?**

"I do not think that a robot could ever be conscious."

"Unless they programmed it that way."

"Conscious? No."

"No. No."

"I mean I think a robot could be programmed to be conscious; how are they programmed to do everything else?"

"That's another big part of artificial intelligence is to make them conscious, to make them feel."

**NATS:** Establishing shot of Columbia University

Hod Lipson

**LIPSON:** Back in 2005 we started trying to build machines with self-awareness. This robot, to begin with, didn't know what it was. All it knew was that it needed to do something like walk. Through trial and error it figured out how to walk using its imagination and then it walked away.

Hod Lipson

**LIPSON:** And then we did something very cruel; we chopped off a leg and watched what happened. In the beginning, it didn't quite know what had happened, but over about a period of a day it then began to limp.

Hod Lipson

**LIPSON:** And then a year ago, we were training an AI system for a live demonstration. We wanted to show how we wave all these objects in front of the camera and then the AI can recognize the objects. And so we were preparing this demo and we had on a side screen this ability to watch what certain neurons were responding to. And suddenly we noticed that one of the neurons was tracking faces; it was tracking our faces as we were moving around. Now the spooky thing about this is that we never trained the system to recognize human faces, and yet somehow it learned to do that.

**NATS:** Our spooky face tracking device...

Hod Lipson

**LIPSON:** Even though these robots are very simple, we can see there's something else going on there, it's not just programming. So this is just the beginning...

**NATS:** The plane takes flight...

Eric Horvitz

**HORVITZ:** I often think about that beach in Kitty Hawk. *(beat)* The 1903 flight by Orville and Wilbur Wright. It was a kind of a canvas plane and some wood and iron. And it gets off the ground for, what, a minute and twenty seconds on this windy day before touching back down again.

**NATS:** Midcentury air travel

Eric Horvitz

**HORVITZ:** And it was just around 65 summers or so after that moment that you have a 747 taking off from JFK where the major concern of someone on the airplane might be whether or not their salt-free diet meal is gonna be coming to them or not. With a whole infrastructure, with travel agents and tower control, and it's all casual and it's all part of the world.

**NATS:** Modern air travel

Eric Horvitz

**HORVITZ:** Right now as far as we've come with machines that can solve problems, we're at Kitty Hawk now. We're in the wind, we have our tattered canvas planes up in the air.

**NATS:** Surveillance footage

Eric Horvitz

**HORVITZ:** But what happens in 65 summers or so... We will have machines that are beyond human control. Should we worry about that? *(beat)* I'm not sure it's going to help.

**NATS:** The beach at Kitty Hawk

## Reel Five

Jerry Kaplan

**KAPLAN:** Nobody has any idea today what it means for a robot to be conscious. There is no such thing. *(beat)* There are a lot of smart people and I have a great deal of respect for them, but the truth is machines are natural psychopaths.

**NATS:** News reports of the market tumbling down, losing money, chaos

- Jerry Kaplan **KAPLAN:** Take the flash crash of 2010. In a matter of minutes, a trillion dollars in value was lost in the stock market.
- NATS:** Traders on the floor dejected
- Jerry Kaplan **KAPLAN:** So, what went wrong? By that point in time, more than 60% of all the trades that took place on the Stock Exchange were actually being initiated by computers.
- NATS:** The bell rings, it's over, now tell us how it happened...
- Justin Wisz **WISZ:** The short story of what happened in the Flash Crash is that algorithms responded to algorithms, and it compounded upon itself, over and over and over again, in a matter of minutes.
- NATS:** Brian Williams: At one point the market fell as if down a well...
- Jonathan Nolan **NOLAN:** There is no regulatory body that can adapt quickly enough to prevent potentially disastrous consequences of AI operating in our financial systems. They are so prime for manipulation.
- NATS:** News flash...
- Justin Wisz **WISZ:** That's the type of AI run amok that scares people.
- Jerry Kaplan **KAPLAN:** When you give them a goal, they will relentlessly pursue that goal. How many computer programs are there like this? Nobody knows.
- Michal Kosinski **KOSINSKI:** One of the fascinating aspects about AI in general is that no one really understands how it works. *(beat)* Even people who create AI don't really fully understand, because it has millions of elements-- it becomes completely impossible for a human being to understand what's going on.
- Hannes Grassegger **GRASSEGGER:** Microsoft had set up this artificial intelligence called Tay on Twitter, which was a chat bot.
- NATS:** Enter Tay from Twitter

- Hannes Grassegger **GRASSEGGER:** They started out in the morning and Tay was starting to tweet and learning from stuff that was being sent to him from other Twitter people. Because some people like troll attacked him, within 24 hours the Microsoft bot became a terrible person.
- NATS:** Racist tweets
- Hannes Grassegger **GRASSEGGER:** They had to literally pull Tay off the net because he had turned into a monster, a misanthropic, racist, horrible person you never want to meet. (*beat*) And nobody had foreseen this.
- Michal Kosinski **KOSINSKI:** The whole idea of AI is that we are not telling it exactly how to achieve a given outcome or a goal. AI develops on its own.
- NATS:** Historical footage of dictatorships, Nazis, etc.
- Jonathan Nolan **NOLAN:** We're worried about super intelligent AI the master chess player that will out-manoeuvre us. (*beat*) But AI won't have to actually be that smart to have massively disruptive effects on human civilization. We've seen over the last century, it doesn't necessarily take a genius to knock history off in a particular direction. And it won't take a genius AI to do the same thing.
- NATS:** 'Bogus election stories got more attention on Facebook than top real stories...'
- Hannes Grassegger **GRASSEGGER:** Facebook really is the elephant in the room.
- NATS:** Facebook – news feed
- Michal Kosinski **KOSINSKI:** AI running Facebook newsfeed, the task for AI is keeping users engaged. But no one really understands exactly how this AI is achieving this goal.
- NATS:** Obvious FAKE headline on Facebook
- Jonathan Nolan **NOLAN:** Facebook is building an elegant mirrored wall around us. A mirror that we can ask, 'who's the fairest of them all?' And it will answer, 'you you,' time and again. Slowly begin to warp our sense of reality, warp our sense of politics, of history, of global events

until determining what's true and what's not true, is virtually impossible.

Michal Kosinski

**KOSINSKI:** The problem is that AI doesn't understand that. AI just had a mission-- maximize user engagement- - and it achieved that. Nearly two billion people spend nearly one hour on average a day basically interacting with AI that is shaping their experience.

**NATS:** Millions of likes on Facebook

Michal Kosinski

**KOSINSKI:** Even Facebook engineers, they don't like fake news. It's very bad business. They want to get rid of fake news. It's just very difficult to do, because how do you recognize news as fake if you cannot read all of those news personally.

**NATS:** Obama: There's so much active misinformation, and it's packaged very well, and it looks the same when you see it on a Facebook page, or you turn on your television.

Jonathan Nolan

**NOLAN:** It's not terribly sophisticated. But it is terribly powerful. And what it means is that your view of the world, which 20 years ago was determined if you watched the nightly news by three different networks, with three anchors who endeavored to try to get it right...

**NATS:** Clicking through the old broadcasters, Ted Kopel, Dan Rather, etc.

Jonathan Nolan

**NOLAN:** They might have had a little bias one way or the other, but largely speaking we could all agree on an objective reality. Well that objectivity is gone, and Facebook has completely annihilated it.

**NATS:** Facebook splash...

Jonathan Nolan

**NOLAN:** If most of your understanding of how the world works is derived from Facebook, facilitated by algorithmic software that tries to show you the news you want to see, that's a terribly dangerous thing. And the idea that we have not only set that in motion, but allowed bad faith actors access to that information, this is a recipe for disaster.

**GFX:** Transition to the election material

Tim Urban

**URBAN:** I think that there will definitely be lots of bad actors trying to manipulate the world with AI. 2016 was the perfect example of an election where there was lots of AI, producing lots of fake news and distributing it for a purpose, for a result.

**MUSIC:** Open up and resolve the music here

**NATS:** Alexander Nix, the power of Big Data...

Jonathan Nolan

**NOLAN:** Cambridge Analytica emerged quietly, as a company that, according to its own hype, has the ability to use this tremendous amount of data in order to effect societal change.

Jonathan Nolan

**NOLAN:** In 2016, they had three major clients, Ted Cruz was one of them.

**NATS:** Nix's presentation showing the Lindsay Graham quote...

**NIX:** It's easy to forget that only 18 months ago, Senator Cruz was one of the less popular candidates...

Hannes Grassegger

**GRASSEGGER:** So what was not possible maybe like 10 or 15 years ago was that you can send fake news to exactly the people that you want to send it to, and then you could actually see how he or she reacts on Facebook and then adjust that information according to the feedback that you got. (*beat*) And so you can start developing kind of a real-time management of a population.

**NATS:** 'In this case we've zoned in on a group we've called persuasion... they need a persuasion message, gun rights I've selected...'

Hannes Grassegger

**GRASSEGGER:** Through social media, there's an infinite amount of information that you can gather about a person.

**NATS:** **NIX:** We have somewhere close to 4 or 5 thousand data points on every adult in the United States...

- Hannes Grassegger **GRASSEGGER:** It's about targeting the individual, it's like a weapon which can be used in the totally wrong direction.
- Jonathan Nolan **NOLAN:** That's the problem with all of this data. It's almost as if we built the bullet before we built the gun.
- NATS:** NIX: Ted Cruz employed our data, our behavioral insights. He started from a base of less than 5%, and had a very slow but steady but firm rise to above 35% making him obviously the second most threatening contender in the race. Now clearly the Cruz campaign is over now. But what I can tell you is of the two candidates left in this election, one of them is using these technologies...
- NATS:** Donald Trump swearing in...
- Jonathan Nolan **NOLAN:** Elections are a marginal exercise. It doesn't take a sophisticated AI in order to have a disproportionate impact. Before Trump, Brexit was another supposed client.
- NATS:** News flash of the Brexit campaign, leaving the European Union.
- Jonathan Nolan **NOLAN:** Cambridge Analytica allegedly uses AI to push through two of the most ground-shaking pieces of political change in the last 50 years. These are epochal events. And if we believe the hype, they are connected directly to a piece of software, essentially, created by a professor at Stanford.
- NATS:** Kosinski at Stanford
- Michal Kosinski **KOSINSKI:** Back in 2013, I described that what they are doing is possible, and warned against this happening in the future.
- Hannes Grassegger **GRASSEGGER:** At the time, Michal Kosinski was a young Polish researcher working at the Psychometrics Centre. So what Michal had done was to gather the largest-ever data set of how people behave on Facebook.
- NATS:** Personality quizzes



- Michal Kosinski **KOSINSKI:** Psychometrics is trying to measure psychological traits such as personality, intelligence, political views, and so on. Now, traditionally those traits were measured using tests and questionnaires.
- NATS:** Filling out a form...
- Jonathan Nolan **NOLAN:** A personality test, the most benign thing you could possibly think of, something that doesn't necessarily have a lot of utility, right?
- Michal Kosinski **KOSINSKI:** Our idea was that instead of tests and questionnaires, we could simply look at the digital footprints of behaviors that we are all leaving behind to understand openness, conscientiousness, neuroticism.
- Hannes Grassegger **GRASSEGGER:** You can easily buy personal data such as where you live, what club memberships you've joined, which gym you go to. There are actually marketplaces for personal data.
- Jonathan Nolan **NOLAN:** It turns out we can discover an awful lot about what you're going to do, based on a very, very tiny set of information.
- Michal Kosinski **KOSINSKI:** We are training deep learning networks to infer intimate traits: people's political views, personality, intelligence, sexual orientation, just from an image of someone's face.
- NATS:** People's faces
- Michal Kosinski **KOSINSKI:** Now think about countries which are not so free and open minded. If you can reveal people's religious views or political views or sexual orientation, based on only profile pictures, this could be literally an issue of life and death. (*beat*) I think there's no going back.
- NATS:** *Ex Machina:* The Turing test
- Jonathan Nolan **NOLAN:** It's almost like technology is a God in and of itself. Like the weather: we can't impact it, we can't slow it down, we can't stop it. We feel powerless.
- NATS:** Beat for dramatic purposes...

- Ray Kurzweil **KURZWEIL:** If we think of God as an unlimited amount of intelligence, the closest we can get to that is by evolving our own intelligence by merging with the artificial intelligence we're creating.
- Elon Musk **MUSK:** Today our computers, phones, applications, give us super human capability. So as the old maxim says, if you can't beat them, join them.
- Rana el Kaliouby **RANA:** It's about a human/machine partnership. I mean, we already see how our phones, for example, acts as memory prosthesis, right? I don't have to remember your phone number anymore cause it's on my phone. It's about machines augmenting our human abilities as opposed to like completely displacing them.
- Jonathan Nolan **NOLAN:** If you look at all of the objects that have made the leap from analog to digital over the last twenty years, it's a lot. We're the last analog object in a digital universe. And the problem with that, of course, is that the data input/output is very limited. It's this. It's these.
- MUSIC:** Big hit
- Shivon Zilis **ZILIS:** Our eyes are pretty good. We're able to take in a lot of visual information, but our information output is very, very, very low, and the reason this is important if we envision a scenario where A.I. is playing a more prominent role in societies, we want good ways to interact with this technology so that it ends up augmenting us.
- Elon Musk **MUSK:** I think it's incredibly important that AI not be the Other. It must be Us.
- MUSIC:** Music features to take it in...
- Elon Musk **MUSK:** I could be wrong about what I'm saying. I'm certainly open to ideas if anyone can suggest a path that's better. But I think we're really going to have to either merge with AI or be left behind.
- Sean Gourley **GOURLEY:** It's hard to kind of think of unplugging a system that's distributed everywhere on the planet, that's distributed now across the solar system. You can't just, you know, shut that off.

Jonathan Nolan

**NOLAN:** We've already opened Pandora's Box. We've unleashed forces that we can't control, we can't stop. We're in the midst of essentially creating a new life form on Earth.

**MUSIC:** Kicks in

Stuart Russell

**RUSSELL:** We don't know what happens next. We don't know what shape the intellect of a machine will be when that intellect is far beyond human capabilities. It's just not something that's possible.

Elon Musk

**MUSK:** The least scary future I can think of is one where we have at least democratized AI, because if one company or small group of people managed to develop god-like digital super intelligence, they can take over the world. At least when there's an evil dictator, that human is going to die. But for an AI, there would be no death. It would live forever. And then you would have an immortal dictator from which we could never escape.

**NATS:** Final montage

**NATS:** Coffee shop

**GFX:**

The pursuit of artificial intelligence  
is a multi-billion dollar industry,  
with almost no regulations.

**POST-CARDS:**

In Memory of  
STEPHEN HAWKING  
1942 - 2018

**POST-CREDITS:**

DoYouTrustThisComputer.org