


Can I Trust my Robot?

AI: Ethics and Morals



Agenda

- Part 1
 - Introduction
 - Issues: Jobs
 - Issues: Safety
 - Issues: Society
- Part 2
 - How it works (v very briefly)
 - Conscious Machines

A Big Question

- ✦ AI is...
 - ✦ ...not artificial consciousness.
- ✦ Is an artificial mind possible?
 - ✦ More in part 2.

Meanwhile...

- AI is already affecting our lives

Neural Network Finds 8-Planet System

By: [Monica Young](#) | December 15, 2017



Scientists have applied a brand of artificial intelligence to data from the exoplanet-hunting Kepler satellite, resulting in the discovery of the first eight-planet system outside our own.

Neural networks, an application of artificial intelligence, can do a little bit of just about anything, from [diagnosing breast cancer](#) to [driving cars](#). Now, they can help detect planets, too. A unique collaboration between an astronomer and a Google engineer has resulted in a neural network that combed through a preliminary sample of Kepler-detected exoplanets, making two Earth-size discoveries along the way.



Artist's depiction of the Kepler-90 system. Kepler-90i, one of the two exoplanet discoveries NASA announced, is the third planet from the left. It's Earth-size and likely rocky.
NASA / Ames Research Center / Wendy Stenzel

There will be big benefits

- ✦ Dangerous work
- ✦ Tedious and unpleasant work
- ✦ Constant help and care for the vulnerable





Click to enlarge

Alpha 1S Humanoid Robot

Product Code : RB-Ubt-01 by UBTECH Robotics

★★★★★ (1) [Add my review](#)

✓ **In stock**
Free shipping on this item


- Compatible with iOS and Android
- Finger-friendly structural design
- APP programming by positioning the servos manually, easy to learn
- APP control program features
- 16 high precision servos

Risks & Issues

- ✦ Take away human work
- ✦ Profound effects on our societies
- ✦ Safety
- ✦ Responsibility when things go wrong
- ✦ The treatment and ownership of robots themselves

Robots in Fiction and Film

- ✦ Asimov's Robots ("I, Robot")




A robot may not injure a human being or, through inaction, allow a human being to come to harm.

WHAT WILL YOU DO WITH YOURS?

Three Laws of Robotics

A robot must obey orders given to it by human beings, except where such orders would conflict with the First Law.

A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.





Robots in Fiction and Film

- ✦ Asimov's Robots
 - ✦ "3 Laws of Robotics"
 - ✦ Nice idea - not possible.

Robots in Fiction and Film

- ✦ AI, Bicentennial Man
 - ✦ Exploring the meaning of Personhood



DARLENE

AT LAST - A LOVE OF YOUR OWN

AT LAST - A LOVE OF YOUR OWN

DAVID

AT LAST - A LOVE OF YOUR OWN

AT LAST - A LOVE OF YOUR OWN


DAVID

- Height 50"
- Weight: 75lbs
- Multi-irregular
- Easy to follow instructions
- 5 yr. full service warranty



Robots in Fiction and Film

- ✦ Blade Runner, Ex Machina
 - ✦ Safety and Slavery





Part 1: Some Issues



Issue 1: “They're taking our jobs!”



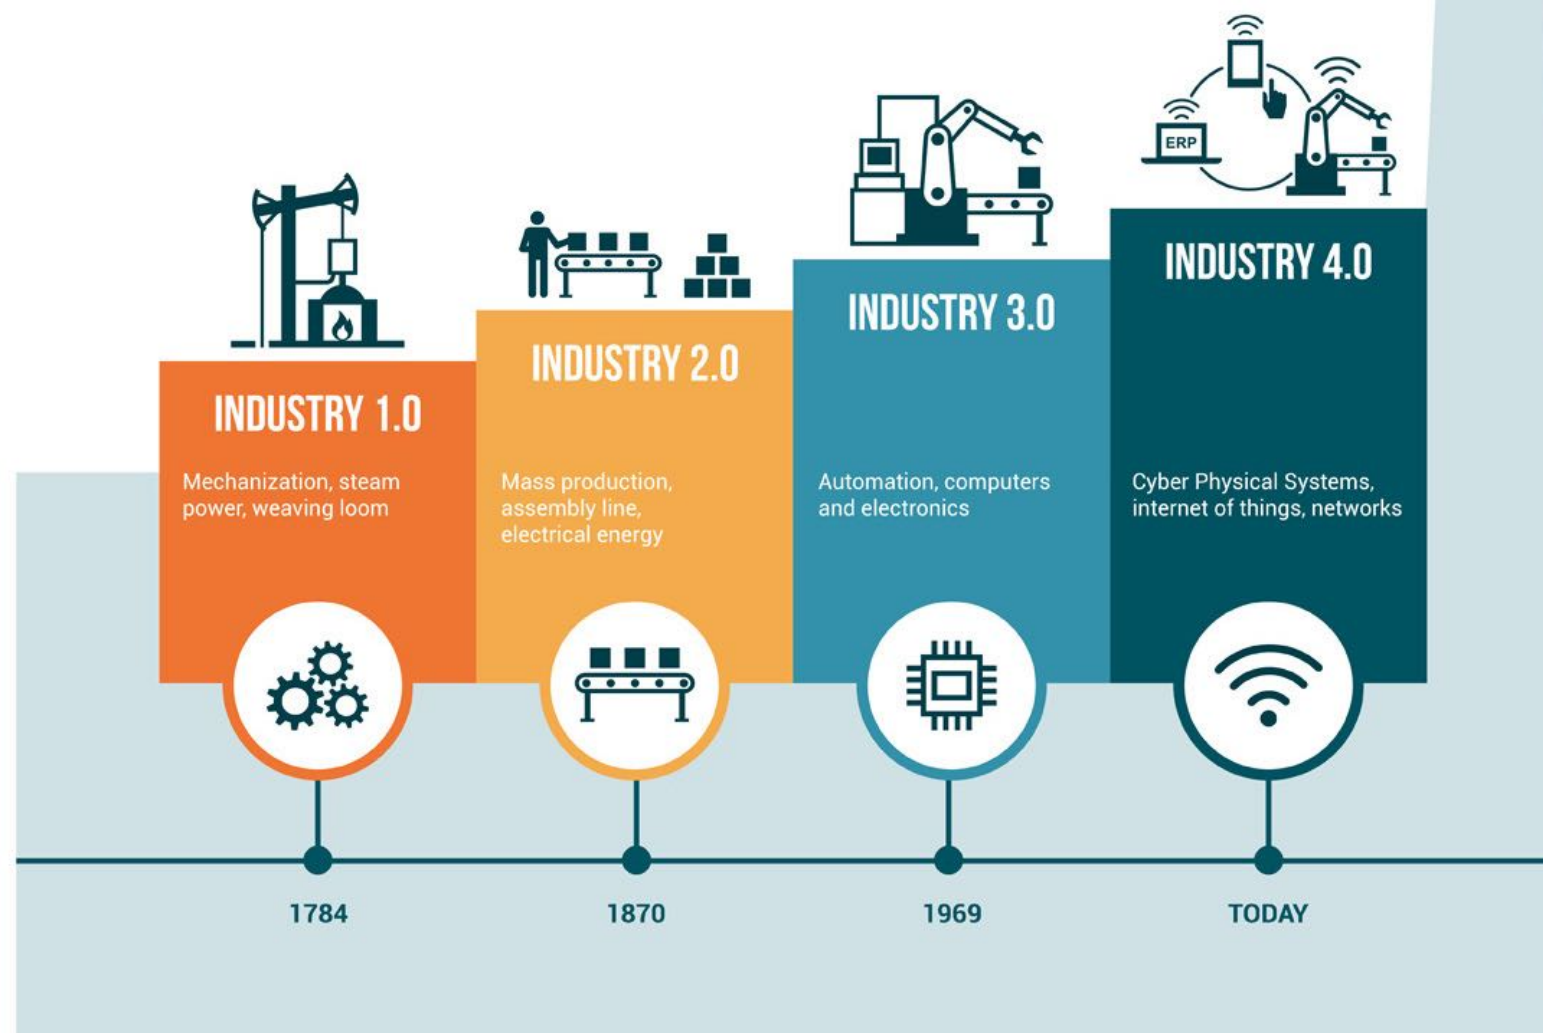
Impact of automation

- ✦ Similar things have happened before:
 - ✦ Industrial Mechanisation
 - ✦ Mass Production
 - ✦ Computer Automation

4th Industrial Revolution

INDUSTRIAL REVOLUTION

TRANSFORMING INDUSTRIES AND INNOVATION



4th Industrial Revolution

- ✦ “New jobs will arise”

The world will always need human brilliance, human ingenuity and human skills.

”

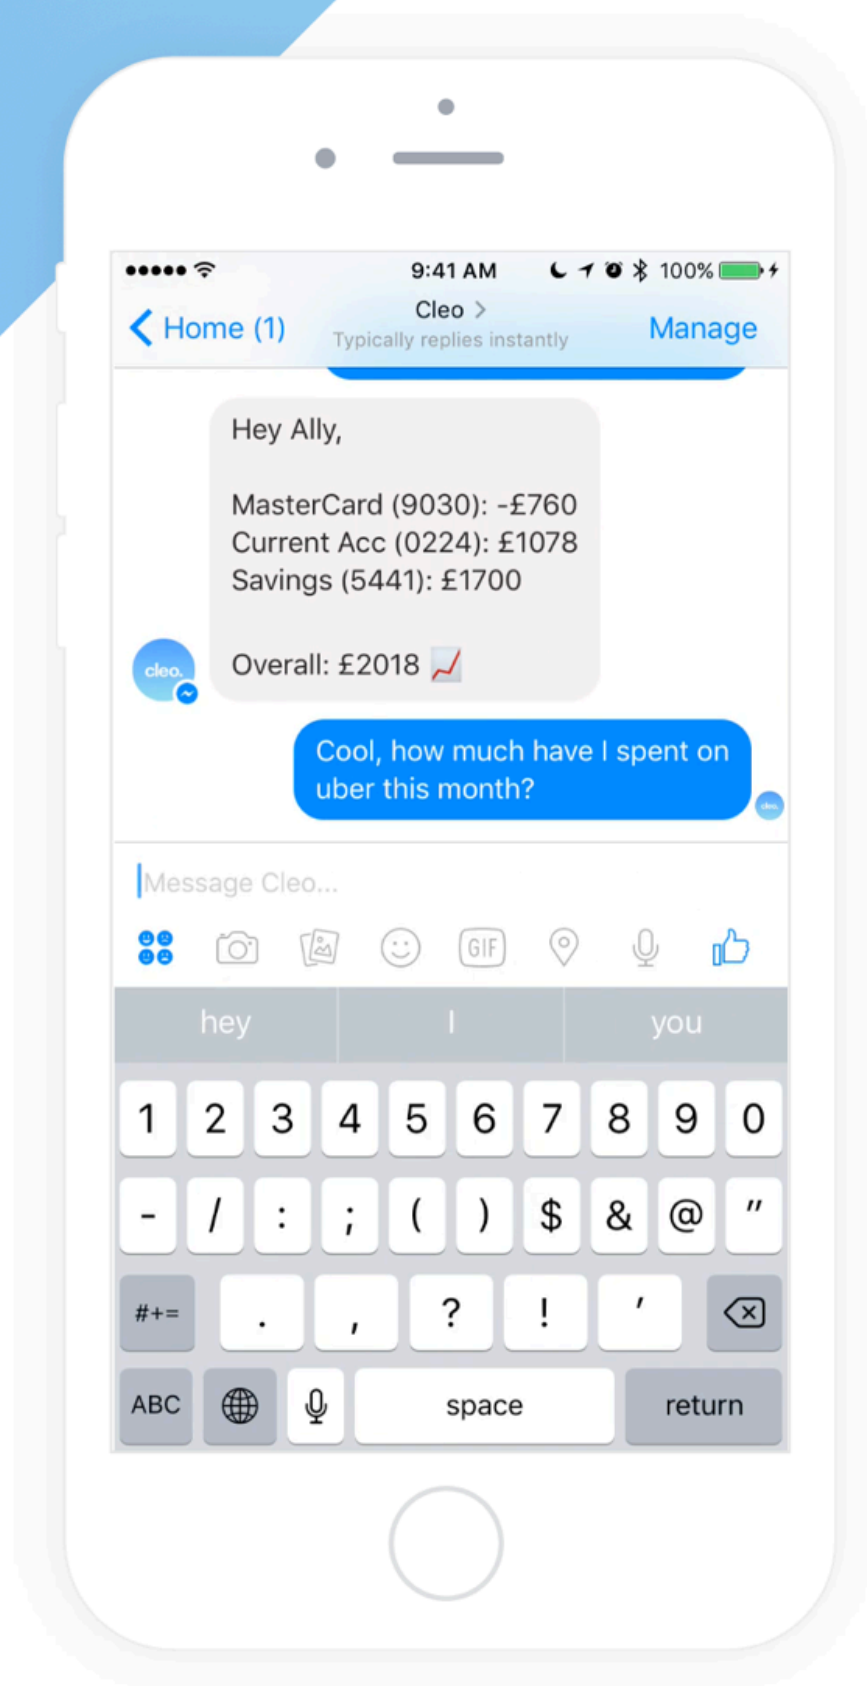
Brad Keywell
Co-founder and CEO, Uptake

WORLD
ECONOMIC
FORUM

- ✦ What if *all* human work can be done more cheaply and equally well by a robot?

What jobs could they do?

- Taxi driver?
- Nurse?
- Domestic servant? (everyone loves ironing!)
- Doctor?
- Finance?



Meet Cleo.

Your AI friend that looks after your money.
Sign up for free, it takes less than 2 minutes.

[Sign up →](#)



[Watch our video](#)

What jobs could they do?

- Taxi driver?
- Nurse?
- Domestic servant? (everyone loves ironing!)
- Doctor?
- Finance?
- Engineer/inventor?
- Teacher?



Animal Math Games for Kids in Pre-K, Kindergarten and 1st Grade Learning Numbers, Counting, Addition and Subtraction



Premium 4+

Eggroll Games LLC

★★★★☆ 47 Ratings

\$4.99

Screenshots [iPhone](#) [iPad](#)



What jobs could they do?

- Taxi driver?
- Nurse?
- Domestic servant? (everyone loves ironing!)
- Doctor?
- Finance?
- Engineer/inventor?
- Teacher?
- Artist?
- Church of England Vicar?

Is meaningful work destined to be the preserve of a lucky few people?

- ✦ There is a clear Biblical command for people to work.
- ✦ *“The Lord God took the man and put him in the garden of Eden to work it and keep it.”
Gen 2:15*
- ✦ *“Six days you shall labour, and do all your work,” Ex20:9*
- ✦ Idleness is not how God intended us to live.
 - ✦ Even if all our physical needs can be provided.

Implications

- ✦ Safety
- ✦ Liability
- ✦ Who is responsible in an accident?
- ✦ What if an AI reacts in an unpredictable way? Perhaps by crashing the stock market?
 - ✦ Who pays? Who gains?

On the other hand...

- ✦ We could work less and better.
 - ✦ Less drudgery, more fun?
- ✦ Less isolation.

Question

- Are there jobs that robots should never be allowed to do?

Issue 2: Is it safe?

- Robots running amok?
Might we accidentally
make this a reality?



Case Study: Autonomous Vehicles

- ✦ Potentially much safer than humans:
 - ✦ More and faster senses
 - ✦ 360° cameras, Radar and LiDAR
 - ✦ React faster

Complex problem!

- View the scene
- Measure ranges, directions & speeds of everything in it
- Map the environment surrounding the vehicle
- Predict the movement of everything in the map
- Understand the behaviour of people, vehicles, animals...
- Decide what to do.
- REPEAT THE ABOVE ABOUT 25 TIMES EVERY SECOND!

If your autonomous vehicle makes a mistake...

- ✦ Whose fault is it?
 - ✦ Yours, as the car owner?
 - ✦ The car manufacturer?
 - ✦ The people who made the camera or computer?
- ✦ Who do we sue?

Mistakes

- ✦ What are you going to think when one of these AIs makes a mistake?
 - ✦ Your AI doctor?
 - ✦ Your child's AI teacher?
 - ✦ Your domestic appliance, help or valet?
 - ✦ Your AI investment manager?

Ethics for AI?

- ✦ Research: to mimic human decision making
 - ✦ Are human ethics appropriate to imitate?
- ✦ "Zero risk" approach does not work
 - ✦ AI "Ethical Paralysis"
- ✦ Possibility of malicious human exploitation of AI conditioning?

Issue 3:

Human society will change

New technology changes us

- ✦ Consider...
- ✦ The Printing Press




New technology changes us

- ✦ Consider...
 - ✦ The Printing Press
 - ✦ Industrial mechanisation




New technology changes us

- ✦ Consider...
 - ✦ The Printing Press
 - ✦ Industrial mechanisation
 - ✦ Fast and affordable transport




New technology changes us

- ✦ Consider...
 - ✦ The Printing
 - ✦ Industrial me
 - ✦ Fast and affordable transport
 - ✦ Computers, Internet and Mobile Technology



“I’m just going down the pub with the robot.”

- ✦ AIs will be very interactive, intelligent, responsive devices



- ✦ Humans will build them into our social interactions.
- ✦ We are compulsive communicators.
 - ✦ How will we behave when we have a device we communicate *with* (not just through)?

Inevitably

- ✦ New technologies and media too often start with pornography.
- ✦ AI will be no different.
- ✦ We need to be ready for the damage this will do.

The Telegraph HOME | NEWS | S

Technology

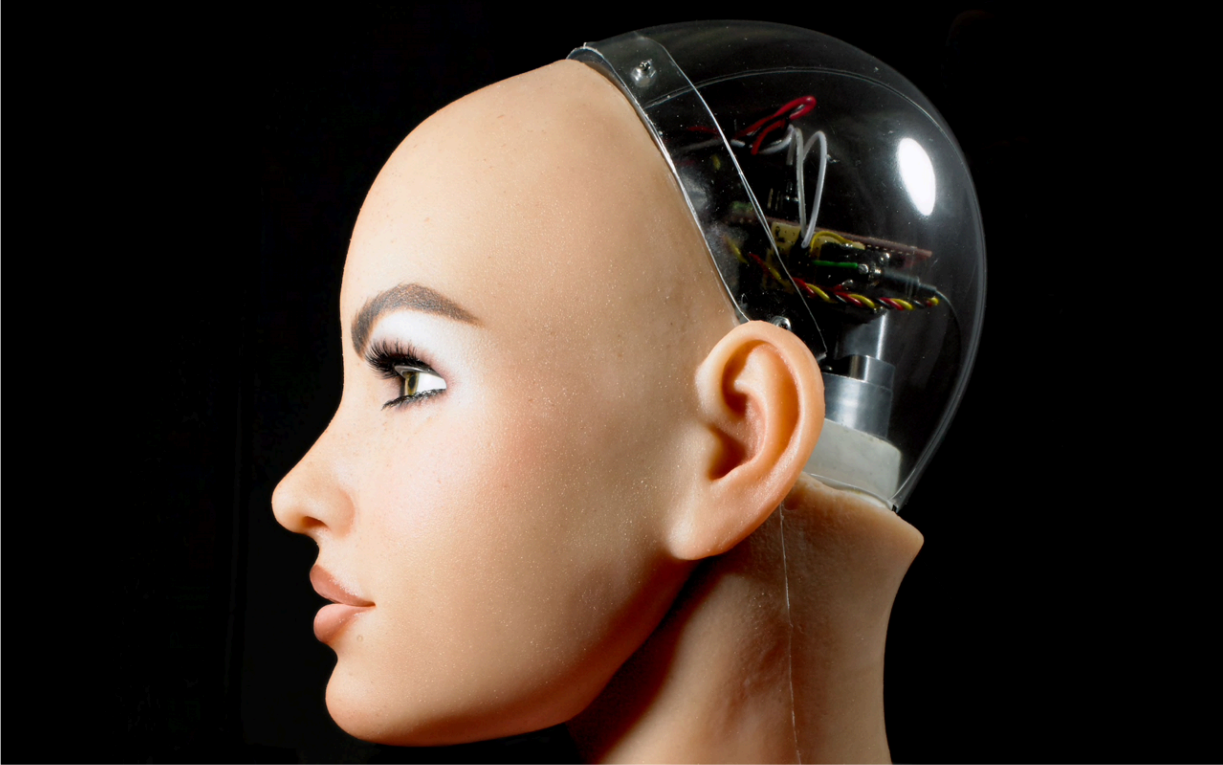
News | Reviews | Opinion | Internet security | Social media | Apple | Google | New

PREMIUM

🏠 > Technology

Meet Harmony: Could a 'sexbot' reboot your unhappy relationship?

[f share](#) [🐦](#) [📌](#) [✉](#) 71



The head of Harmony, an AI-enhanced sex robot. She will go on sale in January for £15,000 CREDIT: NAOMI HARRIS


Part 2:

THIS IS YOUR MACHINE LEARNING SYSTEM?

YUP! YOU POUR THE DATA INTO THIS BIG PILE OF LINEAR ALGEBRA, THEN COLLECT THE ANSWERS ON THE OTHER SIDE.

WHAT IF THE ANSWERS ARE WRONG?


JUST STIR THE PILE UNTIL THEY START LOOKING RIGHT.



How does machine learning work?

- Not really “programmed”

First: a “normal” computer



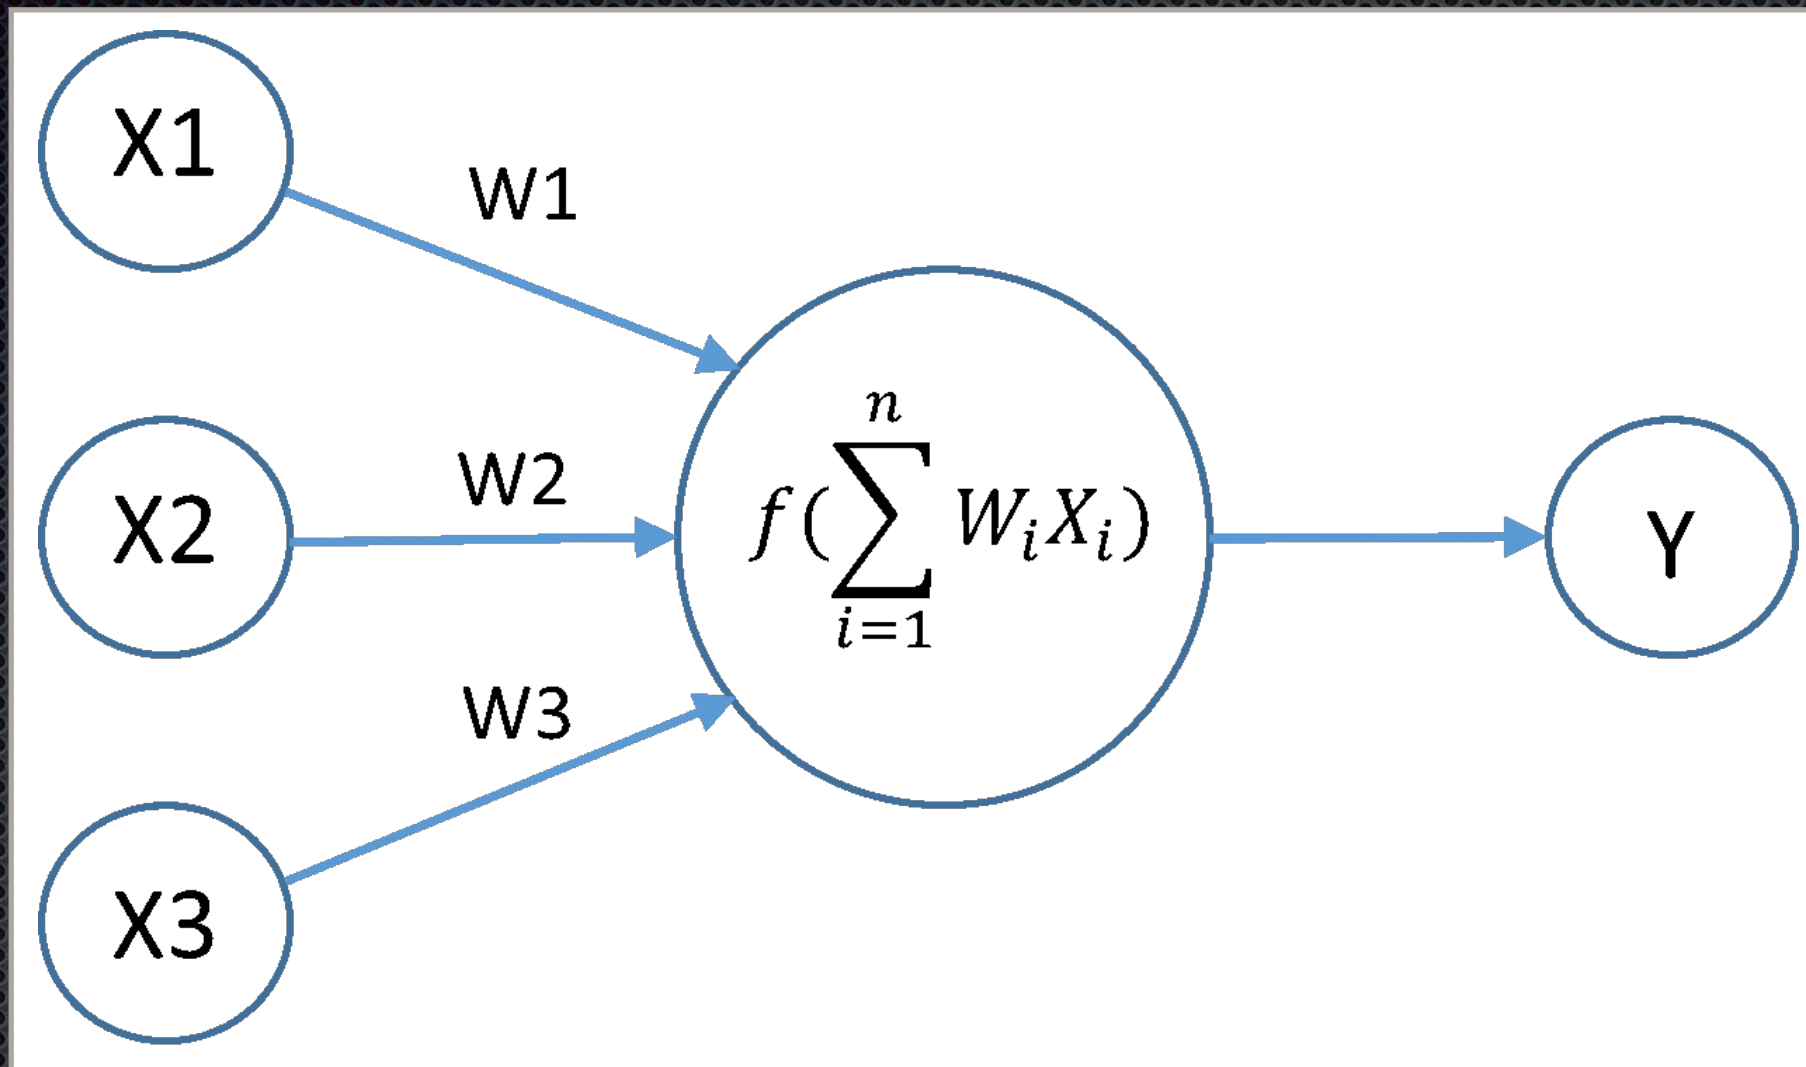
```
C = A + B;
```

```
for (i=0;i<100;i++)  
    C[i] = A[i] + B[i];
```


Human Neuron (Brain Cell)




Artificial Neuron



- Imitates the Biological Neuron with a Mathematical Function

Connect them together



How does machine learning work?

- ✦ Not programmed: trained or conditioned.
- ✦ Each NN is trained to do only one thing.
- ✦ Deterministic.

That Big Question...



The Question: Is a truly conscious machine possible?

- ✦ Maybe. Maybe not.
- ✦ Penrose's argument:
 - ✦ There is something special about the physics of our brains that enables consciousness

A Thought Experiment

- ✦ Just in case...
- ✦ Let's assume that we can and will build conscious machines.

A Thought Experiment

- ✦ How would we know?
 - ✦ Turing Test
 - ✦ "Wit, will and worry"
 - ✦ Conscience and compassion:
 - ✦ "Made in the likeness of God"?

IMHO:

- ✦ If we were to succeed:
 - ✦ We won't know how it works (in detail).
 - ✦ There will have to be some non-determinism build in, deep down.
 - ✦ We will not be able...
 - ✦ ...to program it
 - ✦ ...to determine its behaviour
 - ✦ ...to be confident that it shares our sense of right and wrong or has a conscience.
 - ✦ This is a serious problem.

Questions

- ✦ The questions that follow are provocative.
- ✦ I do not offer answers.

Questions (No answers)

“You shall not make for yourself a carved image, or any likeness of anything that is in heaven above, or that is in the earth beneath, or that is in the water under the earth.” Ex20:4

- Could we be deceived into believing that a machine is a person and treating it as such?
- Could we be setting ourselves up as gods?

Questions (No Answers)

“But you, O Lord, are a God merciful and gracious, slow to anger and abounding in steadfast love and faithfulness.” Ps86:15

- Is there a way that we could instil compassion and a conscience into a CM?
- Are emotions necessary for compassion or a conscience?
- What if CMs turn out to be psychopathic?

Questions (No Answers)

Is it possible that a CM is made in God's image?

- Only God can instil His image in a being.
- But it is not the case that the mind of such a robot would be the creation of the people that built it. They will only have built the machine that it resides in.
- Therefore, if it has a consciousness, who put it there?
- Could a robot sin?

Questions (No Answers)

- Can one be cruel to a robot?
- Would a conscious machine have rights?
 - The right not to be kept as slaves?
 - The right to payment for their labour?
 - What do you pay a robot with?
 - The right not to be switched off?
 - The vote?

Finally: Should we do it?

- ✦ A slave race of powerful but biddable creatures?
- ✦ We may be closer to making Caliban or Nexus-6 a reality.
- ✦ Are we ready for what it entails?

Finally: Should we do it?

- Ray Kurzweil: “Our sole responsibility is to produce something smarter than we are; any problems beyond that are not ours to solve.”

Leading technologists disagree

Cautious:

Elon Musk
Stephen Hawking
Nick Bostrom
Bill Gates

Optimistic:

Mark Zuckerberg
Larry Page
Sam Altman



Finally: Should we do it?

Cody Volkers

“In Genesis 3, there are a multitude of sins wrapped into Adam and Eve’s disobedience. But behind their disobedience was a basic sinful desire: to achieve equality with God. The Bible teaches that we are to be God-like in a moral sense, but not in an axiological sense (that is, in terms of absolute value). Throughout the Old Testament, God’s people are taught to imitate God in character — to “be holy, as I am holy” (Leviticus 20:26). This carries over into the New Testament where Christians are exhorted to “imitate Christ” (Philippians 2). But this moral obligation rests on a basic ontological fact: We ourselves are creature, not God. We are to imitate God’s moral nature precisely because he is God, and we are not.”

<http://www.ecclesiam.org/journal/a-christian-view-of-artificial-intelligence>