


Hal in 2001: A Space Odyssey (1968)




<https://www.youtube.com/watch?v=HwBmPiOmEGQ>

Her (2013)



<https://www.youtube.com/watch?v=WzV6mXIOVI4>

Ex Machina (2015)



The Singularity

Lyle Ungar

Computer and Information Science



Within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended.
- Vernor Vinge 1993



The History

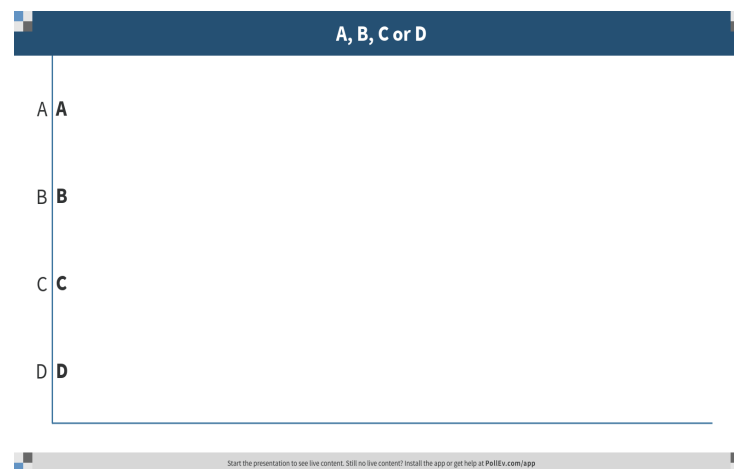
The ever accelerating progress of technology ... gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue.

- John Von Neumann 1950s



Do you think *human level* computer intelligence will most likely occur

- A) within 10 years
- B) 10-30 years
- C) 30-100 years
- D) never



Do you think the singularity (superhuman intelligence) will most likely occur

- A) within 10 years
- B) 10-30 years
- C) 30-100 years
- D) never

A, B, C or D

A


B

C

D

Start the presentation to see live content. Still no live content? Install the app or get help at PollEv.com/app

Predictions for human level AI



The end of humans?

...I fear the development of full artificial intelligence could spell the end of the human race.

-- Stephen Hawking


I think we should be very careful about artificial intelligence. If I had to guess at what our biggest existential threat is, it's probably that.

-- Elon Musk

I am in the camp that is concerned about super intelligence. ... I agree with Elon Musk and some others on this and don't understand why some people are not concerned."

- Bill Gates

The Premise



Modified from
Hans Moravec

Outline

◆ Thesis

- Computers are getting faster (Moore's law)
- Faster computers are smarter
- Therefore computers are getting smarter

◆ Antithesis

- Crying 'wolf'
- It's easy to look smart but hard to be smart

◆ Synthesis

- Different kinds of intelligence
- We are all cyborgs

Thesis: Brains vs. Computers

◆ Brain: memory and speed (high end estimates)

- 100 billion neurons each with 10,000 synapses → 1000 TB
- Each neuron firing at 100 Hz → 10^{13} bits/sec
 - per synapse would be 10^{17} bits/sec

86 billion neurons
in the human brain;
11 billion in cortex

◆ Computer: memory and speed

- 100 GB RAM
- CPU (2 GHz) → 10^{11} bits/sec

◆ Computer power doubles roughly every 18 months

- giving a thousand-fold increase every 15 years

Prediction: in 15 years, a \$1000 computer will have the compute power of a human brain

Thesis: Brains vs. Computers

◆ Brain memory

- 100 billion neurons each with 10,000 synapses → 1000 TB
- 300 words/min * 2 bytes/word * 1/60 min/sec = 10 bytes/sec
- 30 years = 1 billion seconds
- → 10 GB of words in a lifetime
- 0.1-100TB data (images) in a lifetime
- → 10 TB is a better guess

◆ Computer memory

- 10 TB = \$300 on amazon
- But we want *fast* (on chip) memory
 - 5 GB/GPU → 2,000 GPUs for 10 TB (\$500,000—including GPU)

Thesis: Brains vs. Computers

◆ Brain speed

- 100 billion neurons * 10^4 synapses firing at 100 Hz → 10^{17} bits/sec
- Perhaps 10^{15} FLOPS

◆ Computer speed

- CPU → 10^{11} bits/sec
- GPU: 1,000-5,000 gigaflops → 10^{12} FLOPS


1,000 GPUs = 1 brain?

Moore's law

- ◆ **Initially:** The number of transistors in an integrated circuit doubles roughly every two years
- ◆ **More generally:** Compute power doubles every 18-24 months.

Cray XMP - 1983

1 Mhz clock
64 Meg RAM/32 Gig disk
1.5 tons
250 Kilowatts
\$35,000,000




iPhone - 2015

1.3 Ghz clock
64 Gig RAM
4 oz.
0.125 Watts
\$650



Not to
scale



Gordon Moore

Moore's law = exponential growth

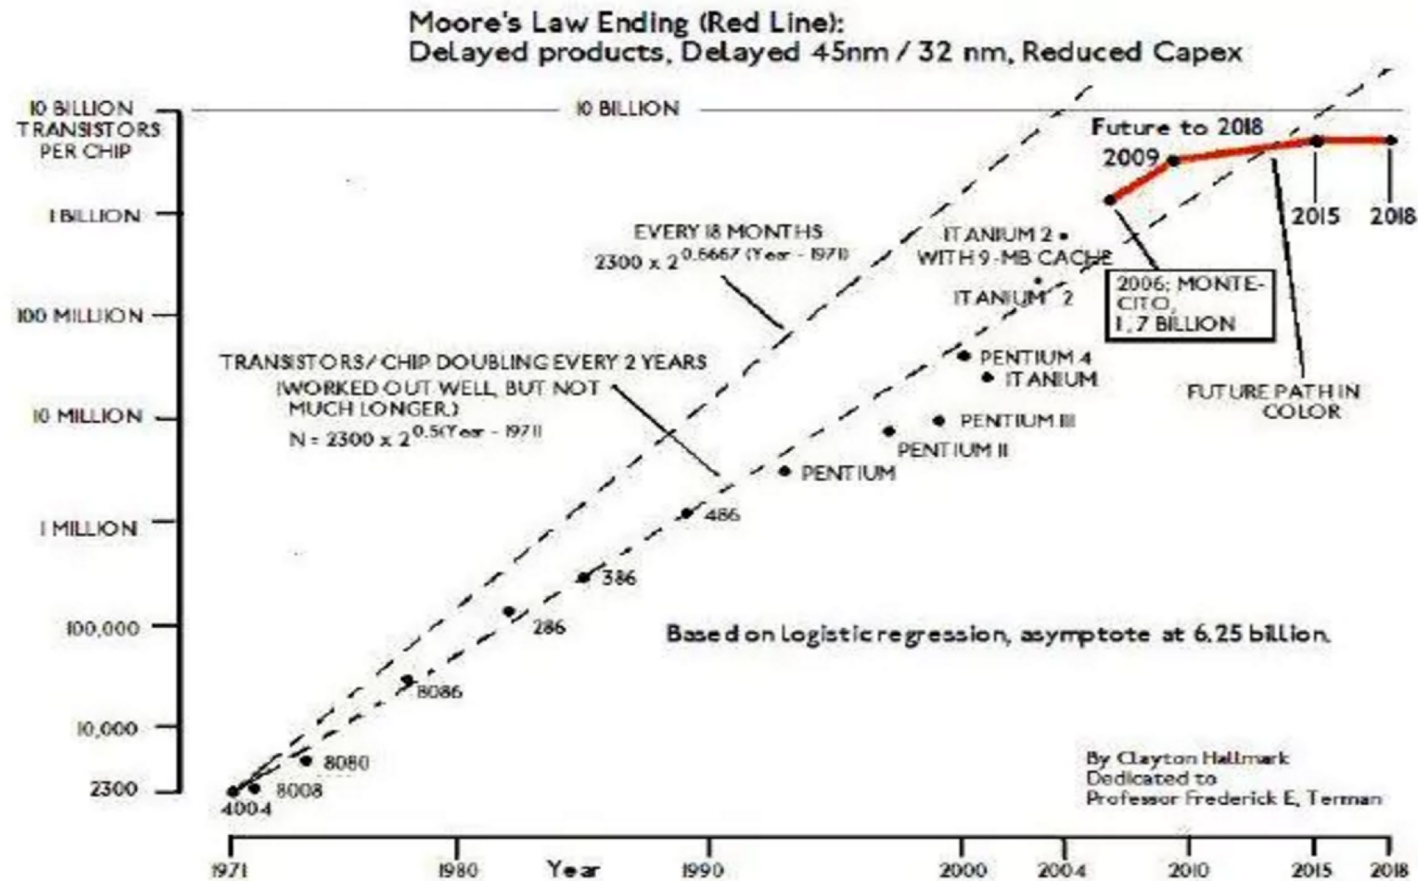
- ◆ Currently: circa 1,000 GPUs to match a human brain
- ◆ 10,000-fold speed-up over the 20 years
 - i.e. doubling every 18 months

1992	2012	Times improvement
128 MB	144 GB	17,000
1-core 25 Mhz	12-core 3.5 Ghz	11,000

- ◆ At that rate a 1,000-fold speed-up will take 15 years
 - $2^{(15/1.5)} = 1,000$

Prediction: in 15 years, one GPU equivalent will have the compute power of a human brain


The end of Moore's Law?



Exponential Moore's Law curve CREDIT: CLAYTON HALLMARK

<http://www.telegraph.co.uk/technology/2016/02/25/end-of-moores-law-whats-next-could-be-more-exciting/>

But remember the longer story



GPUs!

GPUs work in parallel

<https://www.youtube.com/watch?v=-P28LKWTzrl>


- ◆ For some tasks (e.g. deep learning) 10-50 fold speedup at only 2x the cost
 - I.e. runs at 10x the performance per dollar
 - for computation (same cost for memory)

<http://www.redgamingtech.com/why-modern-gpus-perform-faster-than-cpus-good-at-parallel-computing-part-1/>


But how does compute power translate into intelligence?



10 billion neurons in cortex

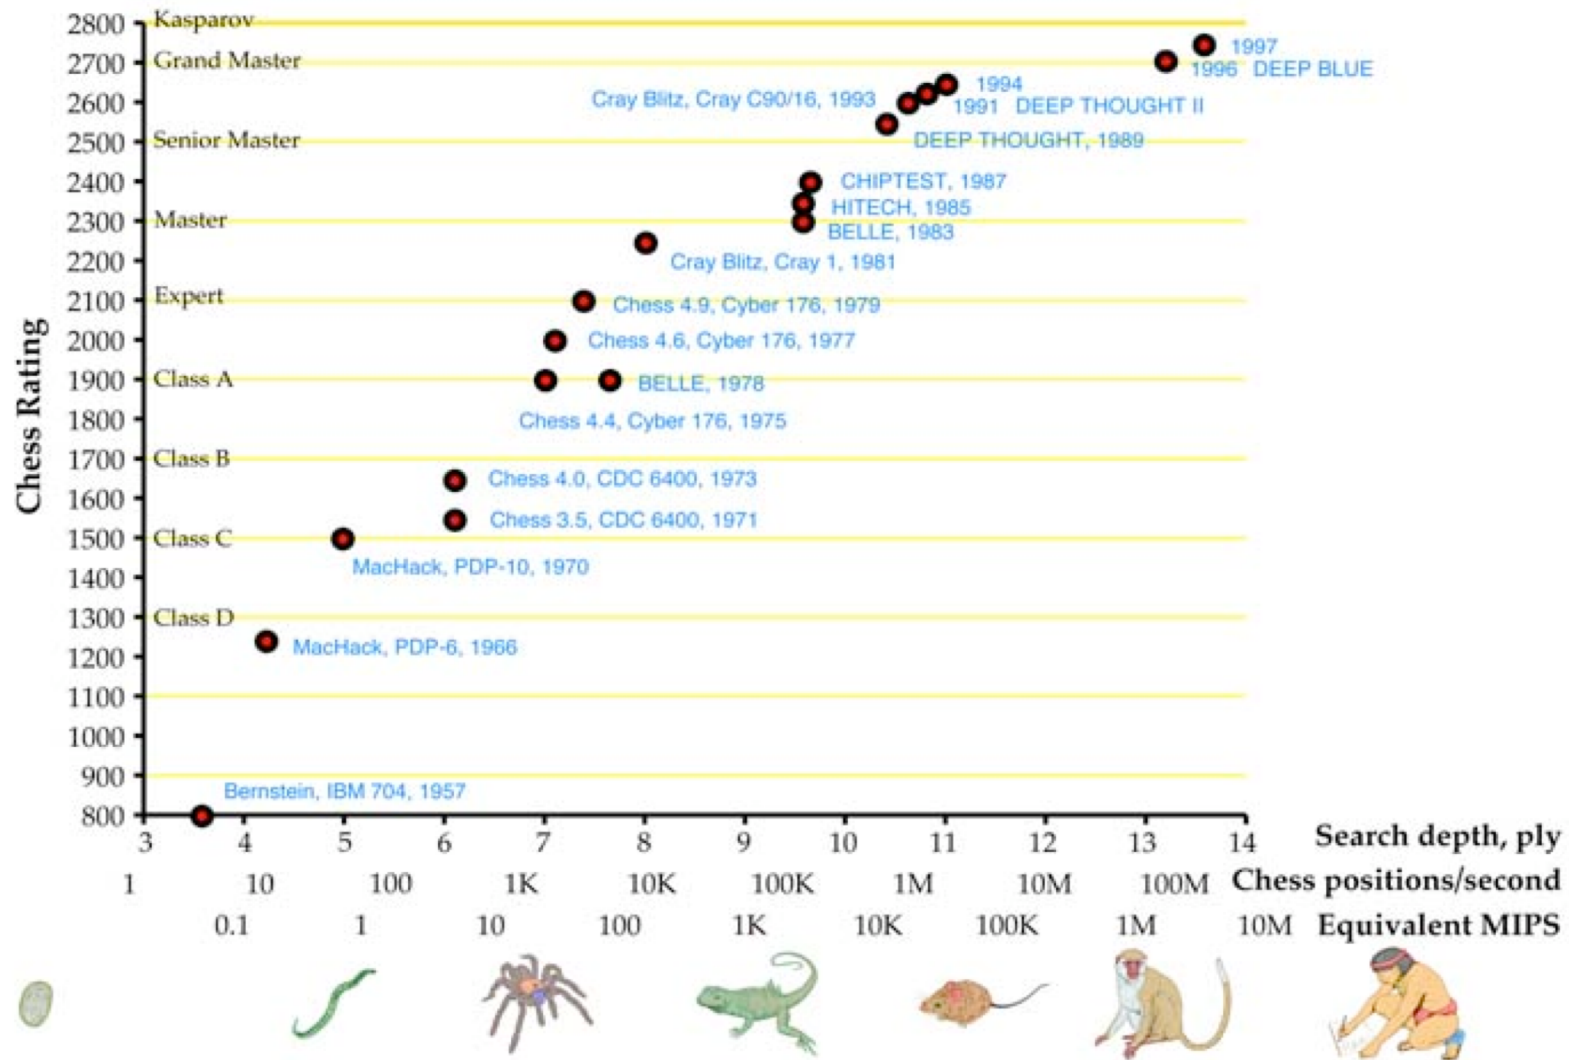


100 million neurons



20 million neurons in cortex

Chess Machine Performance versus Processing Power



Speech Recognition

	1985	1995	2000	2010	2020
Price	\$5,000	\$500	\$50	\$0	\$0
Vocabulary size	1,000	10,000	100,000	10^5	10^6
Continuous speech	No	No	Yes	Yes	Yes
Minutes Training	180	60	5	0	0
Accuracy	Poor	Fair	Decent	Better	Good?

From The Singularity Is Near, p. 103 (updated)

Artificial General Intelligence (AGI)

◆ Possible paths to AGI

- Specialized AI and Machine Learning
- Simulate Human Cognition (high level)
- Simulate Brains (neural level)
- Simulate Evolution

◆ What could weak Artificial General Intelligence do?

◆ What could superhuman intelligence do?

The Singularity: when computers are better at **everything** than people

What could weak AGI do?

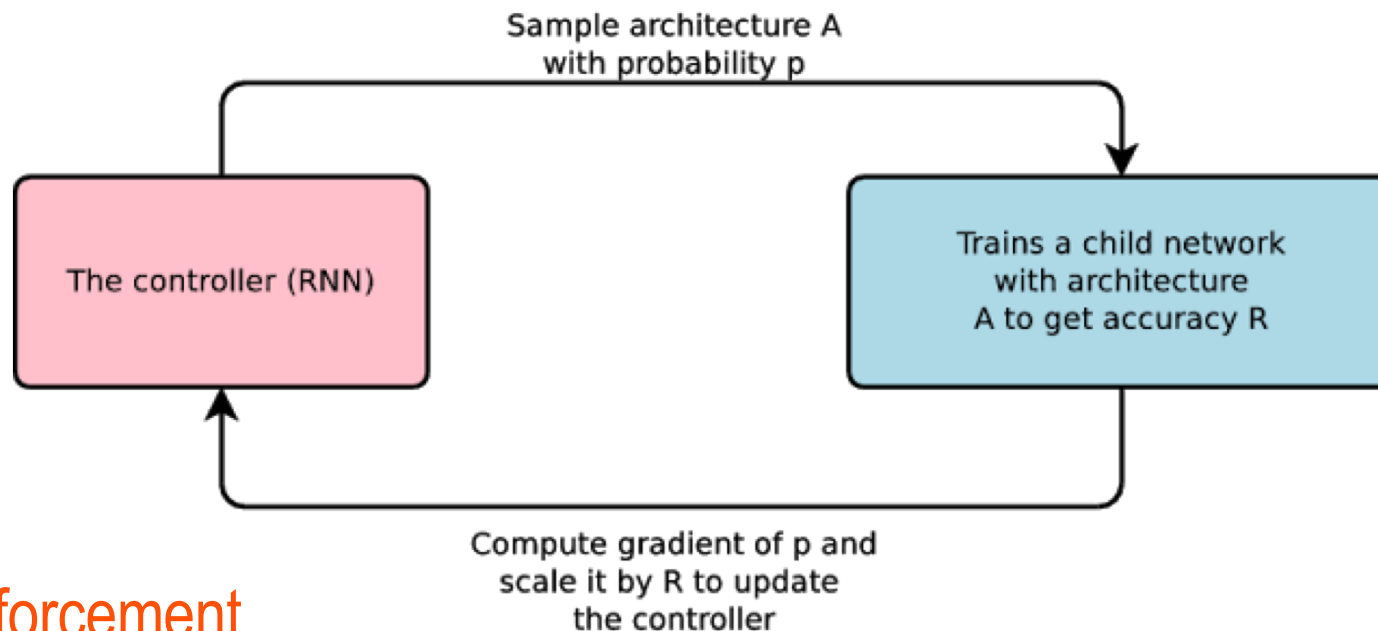
- ◆ **Watch everything you do**
 - Where did I leave my keys?
 - Monitor the elderly
 - Eliminate crime
- ◆ **Interactive entertainment**
- ◆ **Turn rough dictation into good text**

What could a *smart* computer do?

- ◆ Anything anyone can do – but better
 - Not good for job security
- ◆ Design a *really* smart computer
- ◆ Solve aging
- ◆ Upload a copy of your brain
 - What would life be like if you could be copied?
- ◆ Destroy the world by mistake

Design a really smart computer


Using Machine Learning to Explore Neural Network Architecture




Use Reinforcement Learning to search for the 'best' neural net architecture

<https://research.googleblog.com/2017/05/using-machine-learning-to-explore.html>

Our Distorted View of Intelligence





The good, the bad and the bogus

◆ A future so bright you'll need sunglasses

- Marvin Minsky (and Ray Kurzweil)

◆ The future doesn't need us

- Bill Joy

◆ Robots are our children

- Hans Morovic

I'm as fond of my body as anyone, but if I can be 200 with a body of silicon, I'll take it.

- Danny Hillis

Antithesis

◆ The Singularity can't happen because . . .

- Moore's law will end soon
- Software isn't improving as fast as hardware
- AI has overpromised for 50 years
 - It is just a fantasy of old men who want to "upload" on to computers and live forever

Douglas Hofstadter on Kurzweil's and Moravec's books: "It's as if you took a lot of very good food and some dog excrement and blended it all up so that you can't possibly figure out what's good or bad. It's an intimate mixture of rubbish and good ideas, and it's very hard to disentangle the two, because these are smart people; they're not stupid."

Software improves slowly?

◆ Software improves slowly

- I still run emacs because Word is still slow
- Programmers still write 10-20 lines of code a day

◆ Software writing is massively faster

- One line of matlab saves a month of programming in C
- How long ago did it take 20 years ago to write a video game like Penn freshman do now right?
- How long would it have taken 10 years ago to write a decent Google mash-up?

Overpromising

Machines will be capable, within 20 years, of doing any work a man can do.

- Herbert Simon 1956

Within a generation . . . the problem of creating 'artificial intelligence' will substantially be solved.

- Marvin Minsky 1967

Extrapolation is dangerous

◆ In 1966 Time Magazine predicted

- By 2000, technology would have advanced enough that no one in America would work for a living.
- Each American would receive \$30-\$40,000 (1966) dollars every month simply for being American.

Intelligence is harder than it looks



Deep learning is a better ladder but a better ladder doesn't necessarily get you to the moon.

- Gary Marcus

It's easy to *look* smart but hard to *be* smart

◆ Size isn't everything

- Wikipedia knows more than you, but is it smart?
- What have computers ever invented?

◆ Understanding knowledge is hard

1. Today was Jack's birthday.
2. Penny and Janet went to the store.
3. **They** were going to get **him** a present.
4. Janet decided to get **a top**.
5. "Don't do that" said Penny.
6. "Jack has **a top**."
7. **He** will make you take **it** back."

Charniak '72

Need inference, not just facts

The **city council** refused to give the **students** a permit for the demonstration because ...

1. **they** feared violence.
2. **they** advocated revolution.

Winograd '70

It's the representation, stupid!

◆ Representations make reasoning easy or hard

- IV + IX
- IIII + IIIIIIII
- 4 + 9

◆ We only have representations for some problem types, and don't know how to learn new ones

- HMMs or deep learning for speech recognition
- But what is the right representation for a scene?
for a partnership? for love?

Some deep net people claim they can learn any representation

Most thought is simulation?

- ◆ I know how you feel by my feeling the same way
- ◆ I understand speech partly by “knowing” how a mouth and tongue would move to make the sounds
- ◆ I know how a frisbee will go by visualizing throwing it

System 1

Reactive

Parallel

Unconscious

Embodied

Emotion & Simulation

System 2

Reflective

Sequential

Conscious

Abstract


Logic

Consequence: smart computers may need to simulate humans?

Synthesis: *Intelligence Amplification*

- ◆ There are different kinds of intelligence
 - Which is smarter:
 - A cat or a dog?
 - An octopus or an ant?
 - A computer or a person?
- ◆ Computers are in many ways smarter than us
 - How good are you at multiplication, playing chess or go, chip layout, gene sequencing, or airline scheduling?
 - But they are mostly still *much* dumber than us
 - They will gradually take over more functions
 - e.g., translation, surgery, driving cars, teaching, ...

We are as in 19th C for artificial flight



Airplanes don't fly like birds do.
Computers won't think like humans do.

Universal components of intelligence?

◆ Flight


- Power to weight ratio is key
- Wing shape key for lower power to weight ratios

◆ Intelligence

- Memory
- Processing speed
- What else??



We are already cyborgs




You can't do much thinking with your bare brain

Will computers be our cortex?



External computation



Cortex

Subcortex

Brainstem,
Cerebellum

Image credit: <http://siimland.com/what-to-feed-your-neocortex-to-boost-brain-power/>


Neither people nor AI are very smart?

- ◆ **Current AI's are very specialized**
- ◆ **Individual people don't know very much and can't do very much**
 - Imagine being dropped by yourself on a nice desert island
- ◆ **Societies are smart**
 - And increasingly have humans working with AIs

Or maybe we are in a simulation?

- ◆ If Moore's law continues and the singularity occurs, then (eventually) the entire 20th century of human life could be simulated on a home computer.
- ◆ So post-singularity, most people will be simulated.
- ◆ So we are probably ...

Would you like to run a simulation of the 20th Century with minor modifications to see what happens?




Qualia and Reality

Would you still experience the color red and the feeling of anger if ...

- 1) Your retina were replaced with a high quality photo-receptor.
- 2) V1 was replaced with a high quality computer that simulates the neurons there.
- 3) V2 was further replaced with a simulation.
- 4) Your amygdala was replaced.
- 5) $\frac{1}{4}$ of the neurons in your brain were replaced.

- A) Yes
B) No



Consequences

◆ If the singularity occurs

- It is unclear what niche humans will occupy

Even if it takes twice as long as Kurzweil says (40 vs. 20 years) and is half as dramatic, the singularity, if it occurs, will be the biggest event of the 21st century (unless humans kill ourselves off first).

Consequences

◆ Even if no singularity occurs

- Computers will get faster and smarter
- Many jobs might go away
 - Secretaries, travel agents, store clerks, truck drivers
 - Lawyers, radiologists ...
 - *Computers will do to white collar jobs what the industrial revolution did to physical labor jobs?*
 - ◆ *But faster*

<https://www.youtube.com/watch?v=NrmMk1Myrxc>

“about 47 percent of total US employment is at risk”

The Future of Employment

Carl Benedikt Frey & Michael Osborne



OECD:
14% lost

Forrester:
10% lost (net)

Job impact

◆ Many jobs are at risk

- US: over 4 million people drive for a living
- US: 4.5 million retail sales workers
- India: customer support, IT, ...


◆ But mostly *tasks* are automated



Reality

- ◆ Job loss rate is (so far) unchanged
- ◆ Mostly *tasks*, not jobs are automated

ATMs did not replace bank tellers



James Bessen

**Will AGI change
our idea of what it
is to be human?**

Yuval Noah Harari

New York Times Bestselling

Author of *Sapiens*



Homo Deus

A Brief History
of Tomorrow

What you should know

- ◆ **A brain ~ 1,000-5,000 GPUs (in 2018)**
 - 1 brain ~ 10 TB ~ 2,000 GPUs (at 5 GB/GPU)
 - 1 brain ~ 10^{15} FLOPS ~ 1,000 GPUs
- ◆ **Exponential growth is non-intuitive**
 - Slow, then really fast
 - Moore's law = 1,000-fold increase every 15 years
- ◆ **Consequences of the singularity**
 - Unbelievable scientific and economic progress
 - Humans superseded??
- ◆ **Controversy over if/when it will happen**
 - And what the consequences might be

Thank you!