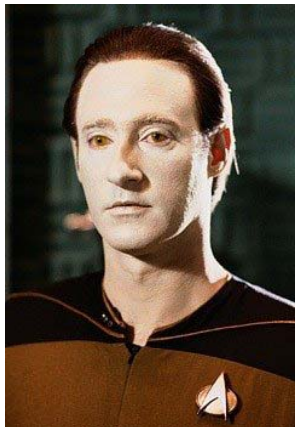

Human and Artificial Intelligence:

Are we building new tools or new people?

L. Bartoszek

9/25/13

For the Geneva Library Foundation

The purpose of this talk

- To introduce concepts and ideas that people are seriously discussing about the future of humanity
- I want to stimulate discussion and thinking about these ideas so that everyone is prepared to help create the future
- I want us to think about what it **means to be human** so that if/when new things appear we can decide how to treat them
- I'll make a general survey of large areas of research
  - You'll have to do homework on your own for more detail and come to your own conclusions

Some terms to discuss:

- Intelligence
  - How are we different from animals?
  - How are computers different from us?
- Consciousness
  - **Is anybody home?** *How do we know?*
- Brains and Minds
- Moore's Law
- Parallel processing
- Nanotechnology
- The Singularity

Some of these words cover a lot of ground that we won't be able to do justice to here.

Some of these words mean different things to different people.

http://www.scientificamerican.com/article.cfm?id=what-makes-us-human

"The 1 percent difference: Humans are distinct from chimpanzees in a number of important respects, despite sharing nearly 99 percent of their DNA. New analyses are revealing which parts of the genome set our species apart."

Many of us believe that when we look into an animal's eyes "somebody is home", but we don't consider them people.

There has been a lot of work discovering what makes us different from other animals.

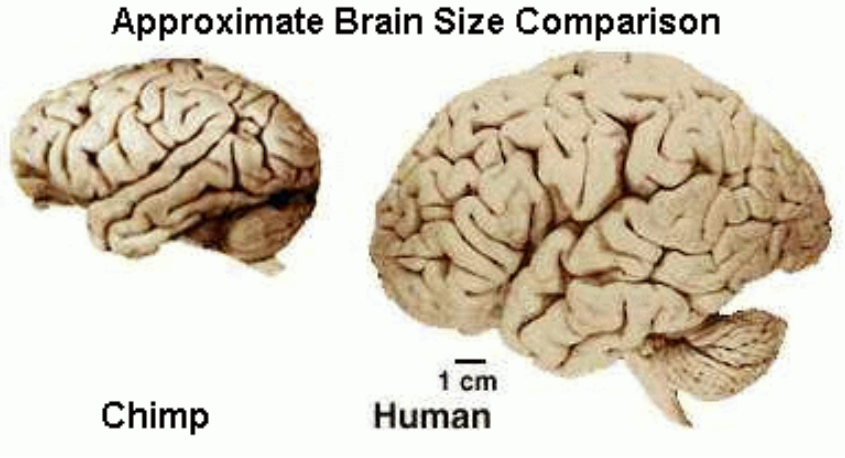

http://blogs.scientificamerican.com/guest-blog/2012/02/21/how-did-human-brains-get-to-be-so-big/

"In the 8 million to 6 million years since the ancestors of humans and chimps went their separate ways, the human brain more than tripled in size. If the earliest humans had brains the size of oranges, today's human brains are more akin to cantaloupes.

As for our closest primate relatives, the chimps? Their brains haven't budged."

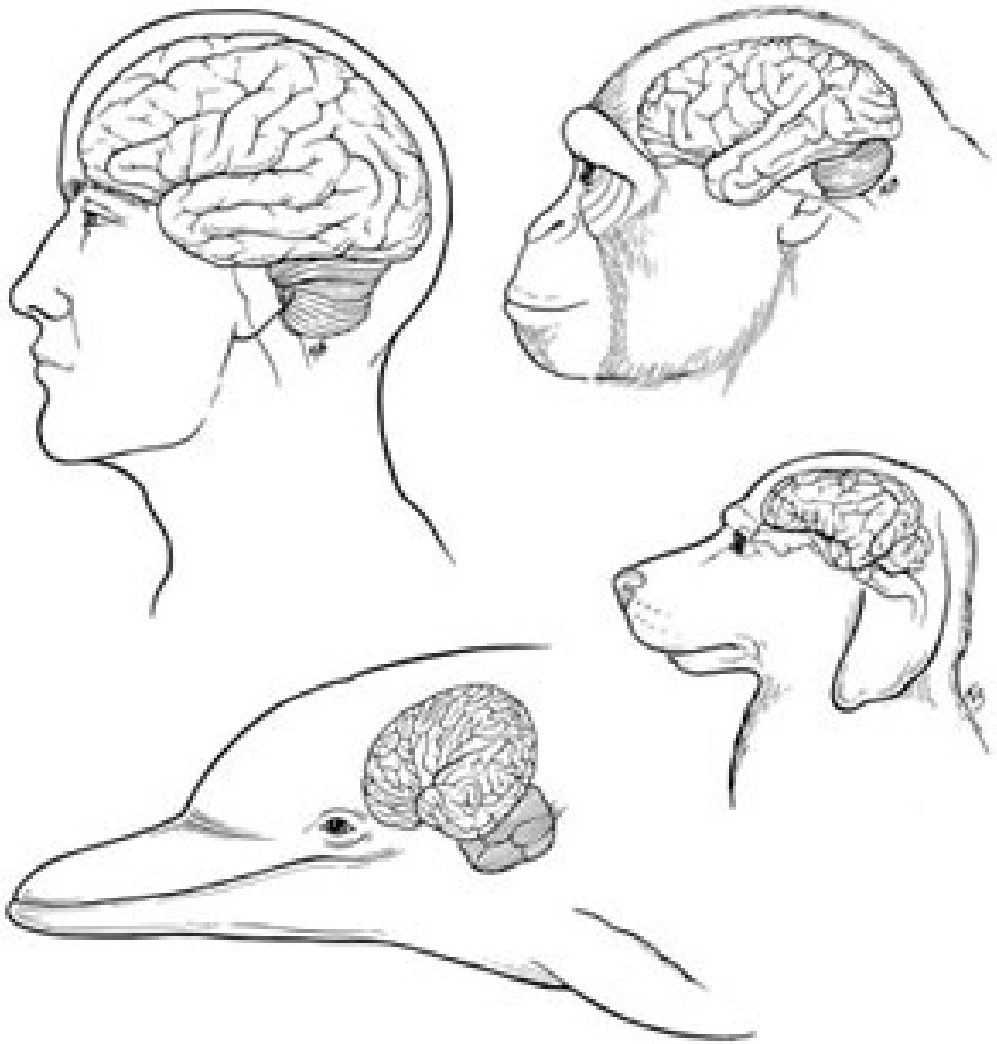

http://www.dana.org/news/brainhealth/detail.aspx?id=10060

Why is it so hard to figure out what the brain does and how it does it?

"Every mammal's brain—human, chimpanzee, dolphin, and dog—has the same basic structure: a wrinkled cortex atop a creased cerebellum, with a brain stem leading to the spinal cord.

More intelligent animals have more cortical surface, however. The human cortex is much larger in proportion to the overall brain than any other animal's."

A basic assumption is that intelligence comes from our brains.

This was not obvious to the ancient Greeks.

"During the 4th century BC Aristotle thought that, while the heart was the seat of intelligence, the brain was a cooling mechanism for the blood."
http://en.wikipedia.org/wiki/History_of_neuroscience

Because Minds are what brains "do"

The “inside” and “outside” of intelligence

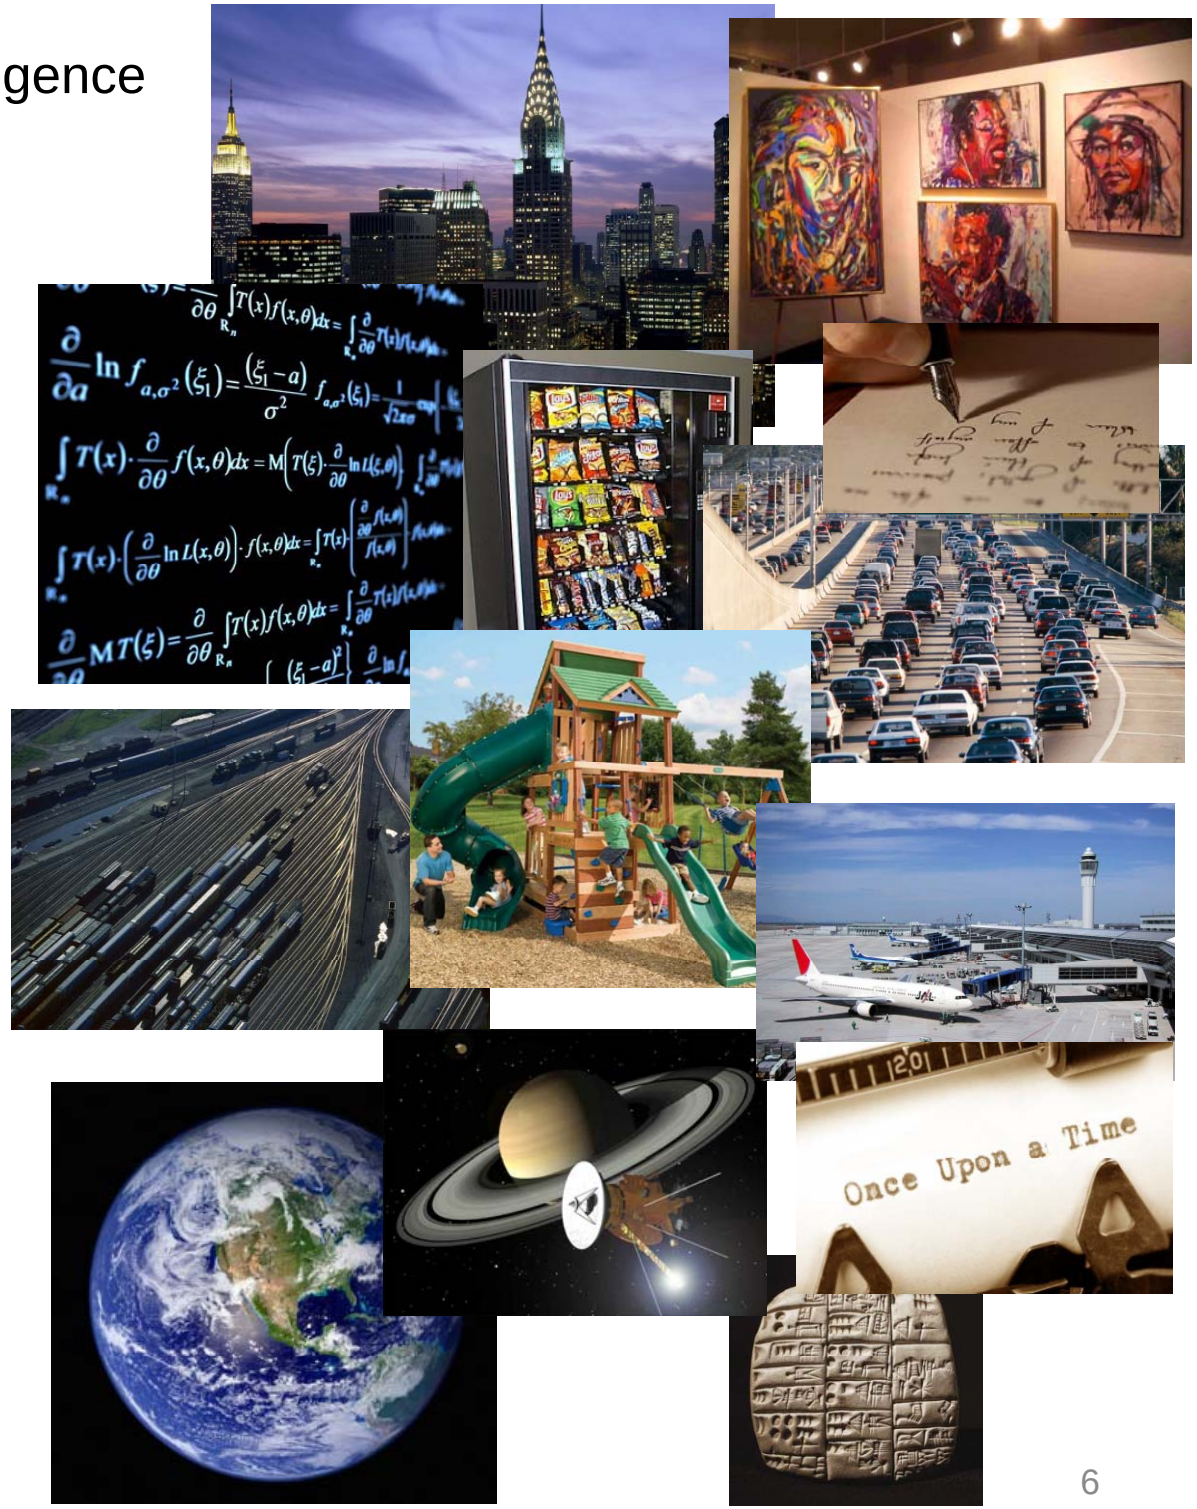

I have an internal, subjective life that you cannot know directly.

I call it my “consciousness”.

What are human brains and computers strong and weak in?

What people are good at:

- Living
- Socializing
- Feeling
- Laughing
- Catching baseballs
- Having “common sense”
- Pattern recognition

What computers are good at:

- Calculating numbers
- Remembering accurately
- Sorting through huge amounts of data
- Being very fast at calculating

What people are bad at:

- Calculating numbers
- Remembering accurately
- Sorting through huge amounts of data (boring!)
- Being very fast at calculating (except for savants)

What computers are bad at:

- Living
- Socializing
- Feeling
- Laughing
- Catching baseballs
- Having “common sense”

We are the perfect complement of each other!

Notice some things I didn't put on the list?

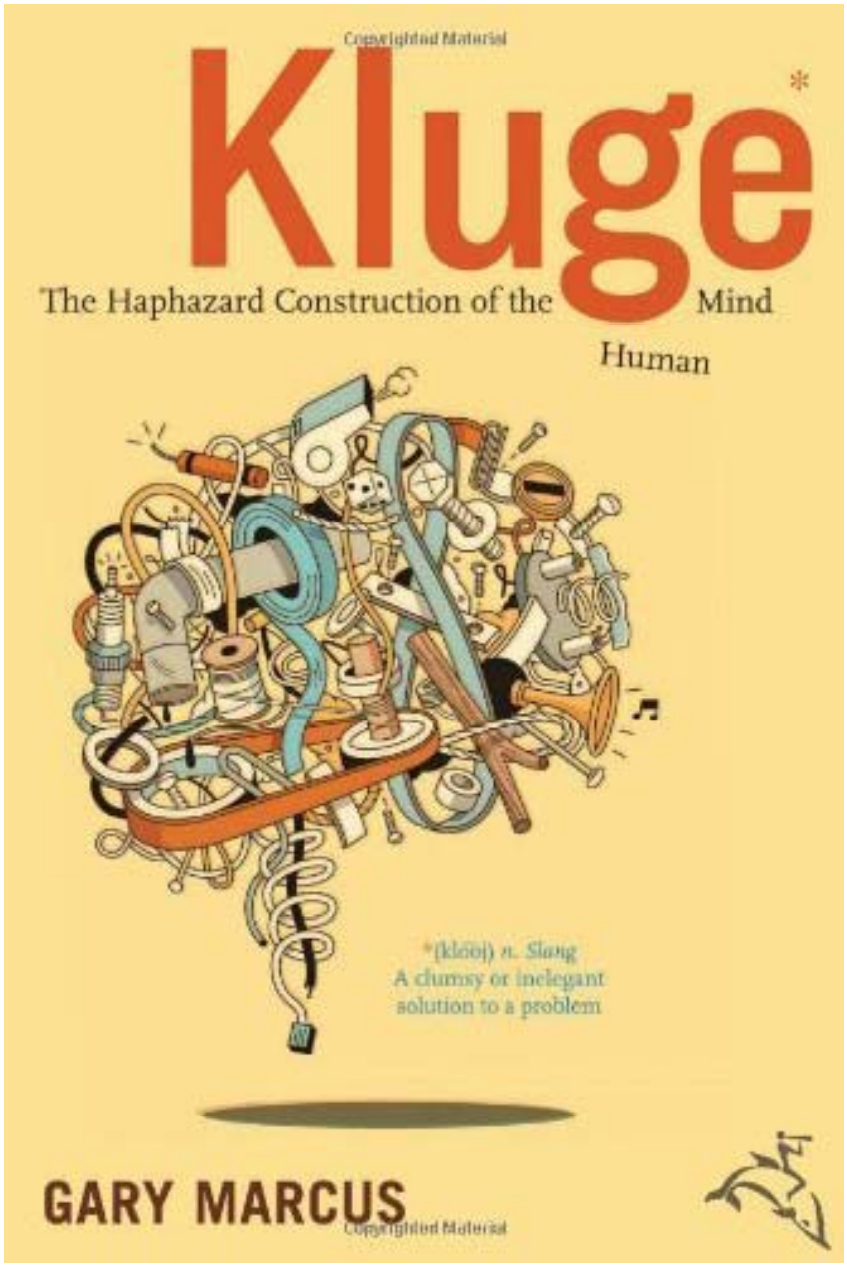

- **Reasoning**
  - Computers are better at reasoning than we are in many ways
  - They are designed-to-reason, not evolved kluged-together thinking machines like we are
  - Gary Marcus' book is highly recommended if you doubt what I'm saying here
  - We can be biased to trust someone by holding a hot cup of coffee!
- **Language**
  - Computers get better at understanding the vagaries of human speech every day
- **Learning**
  - Computer learning is a reality
  - See http://en.wikipedia.org/wiki/Machine_learning#Algorithm_types

10 important differences between brains and computers

1. Brains are analogue; computers are digital
2. The brain uses content-addressable (associative) memory
3. The brain is a massively parallel machine; computers are modular and serial
4. Processing speed is not fixed in the brain; there is no system clock
5. Short-term memory is not like RAM
6. No hardware/software distinction can be made with respect to the brain or mind
7. Synapses are far more complex than electrical logic gates
8. Unlike computers, processing and memory are performed by the same components in the brain
9. The brain is a self-organizing system
10. Brains have bodies

Bonus Difference: The brain is much, much bigger than any [current] computer

http://scienceblogs.com/developingintelligence/2007/03/27/why-the-brain-is-not-like-a-co/

Where we are right now

- As far as I know as of today, **there is no sentient computer in the world**
  - *If the NSA has one, it's not saying*
- We have chatterbots, computers that can beat the best humans at certain games, expert systems and personal digital assistants like Siri—but ***Nobody is Home!***

- *Will that always be the case?*
  - *There are reasons to think—maybe not.*

The New York Times

Science

WORLD | U.S. | N.Y. / REGION | BUSINESS | TECHNOLOGY | SCIENCE | HEALT

Obama Seeking to Boost Study of Human Brain

By JOHN MARKOFF
Published: February 17, 2013

The Obama administration is planning a decade-long scientific effort to examine the workings of the human brain and build a comprehensive map of its activity, seeking to do for the brain what the Human Genome Project did for genetics.

President Obama has made it clear that he thinks the federal government should invest in mapping the human brain in a project of similar scale to the Human Genome Project.

Whatever we learn about the brain can be applied to AI.

U.S. Department of Health and Human Services

NIH News

National Institutes of Health

For Immediate Release
Wednesday, September 15, 2010

$40 million awarded to trace human brain's connections

Souped-up scanners to reveal intricate circuitry in high resolution

But a group of nanotechnologists and neuroscientists say they believe that technologies are at hand to make it possible to observe and gain a more complete understanding of the brain, and to do it less intrusively.

Where is the money?

Just a few of the more prominent names:

- IBM
  - $104.5B revenue in 2012
- Microsoft
  - $74.26B revenue in 2012
- Google
  - $50B revenue in 2012
- Global gaming industry revenue in 2012: $67B
- Worldwide semiconductor industry revenue in 2012: **$298B**

Fortune 500 largest revenues in computers and office equipment makers in 2012

Company	Revenue, $ billions	Profit, $ billions
Hewlett-Packard	127.2	7.07
Apple	108.2	25.92
Dell	62.07	3.49
Xerox	22.63	1.29
NCR	5.44	0.05
Pitney Bowes	5.28	0.62

http://money.cnn.com/magazines/fortune/fortune500/2012/industries/8/
(I don't know why IBM is not on this list.)

How do computer and software companies compare to industry as a whole?

Top companies: Biggest

By Employees | By Equity | **By Market Value**

Rank	Company	500 Rank	3/29/2012 ($ millions)
1	Apple	17	568,615.1
2	Exxon Mobil	1	405,714.1
3	Microsoft	37	269,511.6
4	International Business Machines	19	241,314.5
5	Chevron	3	211,238.9
6	General Electric	6	211,096.1
7	Google	73	210,835.0

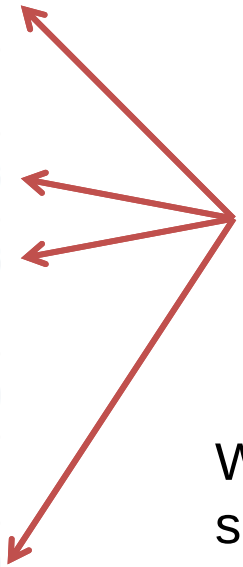

These are all software or computer (or both) manufacturers

We are talking about serious money here.

I truncated this list. Oracle is at 18, Intel 19

http://money.cnn.com/magazines/fortune/fortune500/2012/performers/companies/biggest/marketvalue.html

What do I conclude from this?

- The vast majority of money spent on software and hardware is going to build new tools
  - People want smarter stuff and there is tremendous competition to make smarter products
- A few companies like Google, and some universities, have publicly stated they are trying to build sentient computers
  - Ray Kurzweil is Director of Engineering at Google now
  - "How Ray Kurzweil will help Google make the Ultimate AI Brain"
  - http://www.wired.com/business/2013/04/kurzweil-google-ai/
- There is much less money in building new people
  - But there is a dedicated core of people trying

The power of exponential growth

- **Moore's law** describes a long-term trend in the history of computing hardware. The number of transistors that can be placed inexpensively on an integrated circuit **doubles** approximately **every two years**. This trend has continued for more than half a century.
  - From the Wikipedia
- The more transistors, the more computing power you have.
- Computers attain the number of computations per second of the human brain in the 2020s.
- If the trend continues, by 2049 computers have more computing power than all the human brains that ever lived

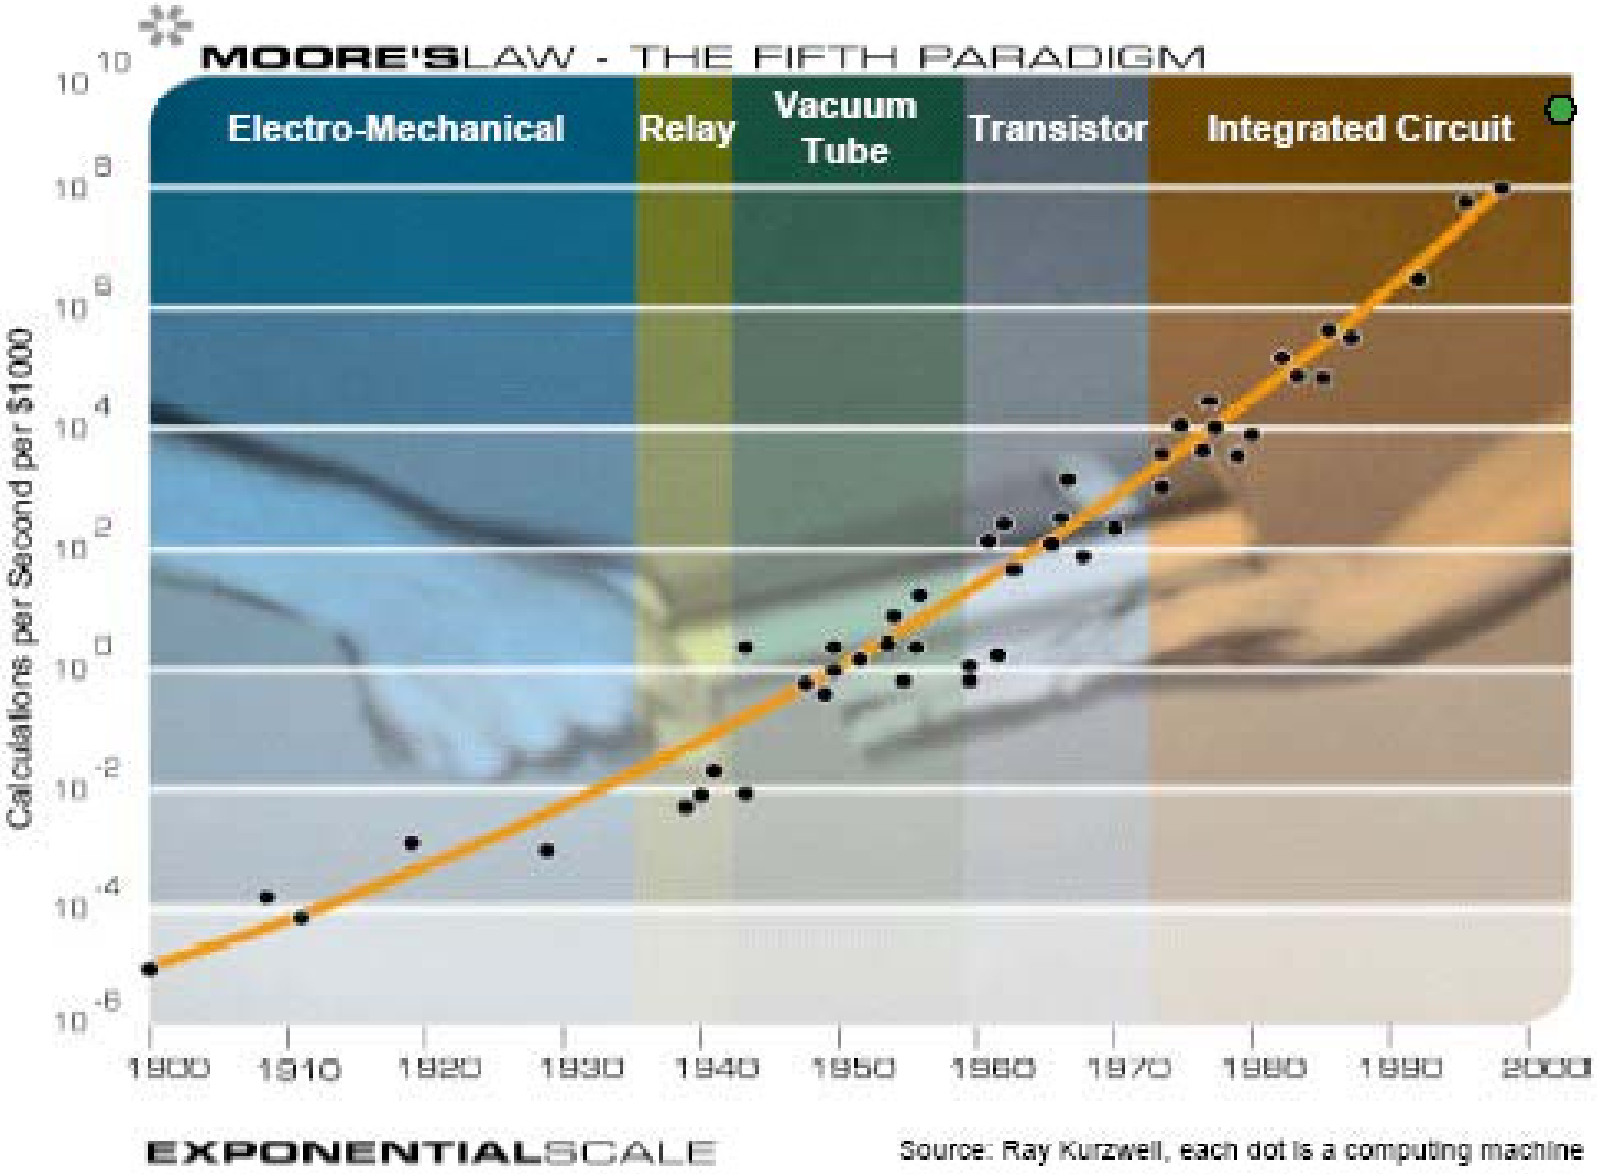

The curvature of the graph on a log plot indicates that the exponent is growing too

The **Big** Question

- **When a machine has the same computing power as a human brain, can it become conscious?**
- One of the criticisms of Artificial Intelligence is that we can't build it because we don't have a theory of human consciousness
- Ray Kurzweil has outlined a plan of research on the human brain that could lead to a computer with similar mental structures
  - Is enough connectivity, complexity and structure enough to create consciousness? No one knows.
  - Rodney Brooks believes that strong AI will not happen without a robot body to experience the world
    - He would argue against the scenario of BINA48 "waking up" (see further)

Ray Kurzweil's latest book:

Vector quantization used to reduce data sets (representing sensory feeds from the senses,) which can then be operated on by "Hierarchical hidden Markov models" as pattern recognizers

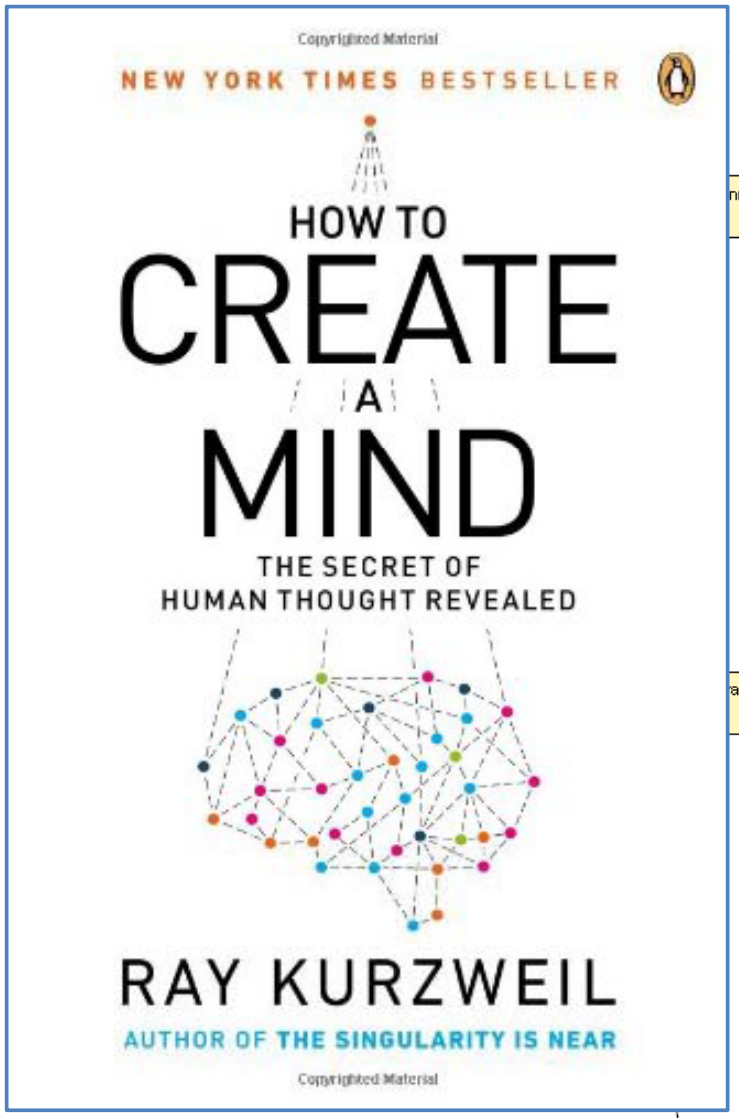

From page 142: "People often fail to appreciate how powerful mathematics can be…"

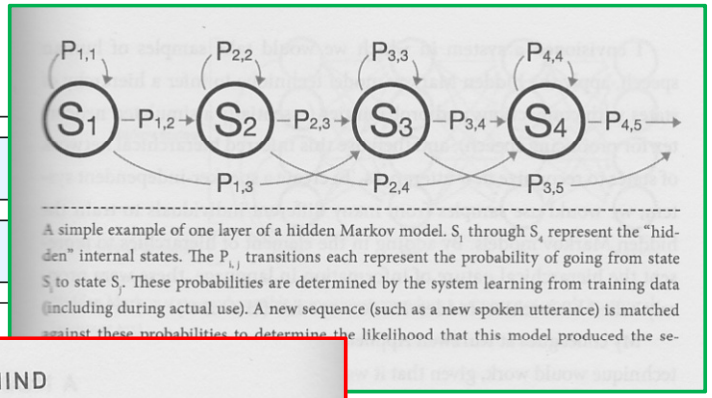

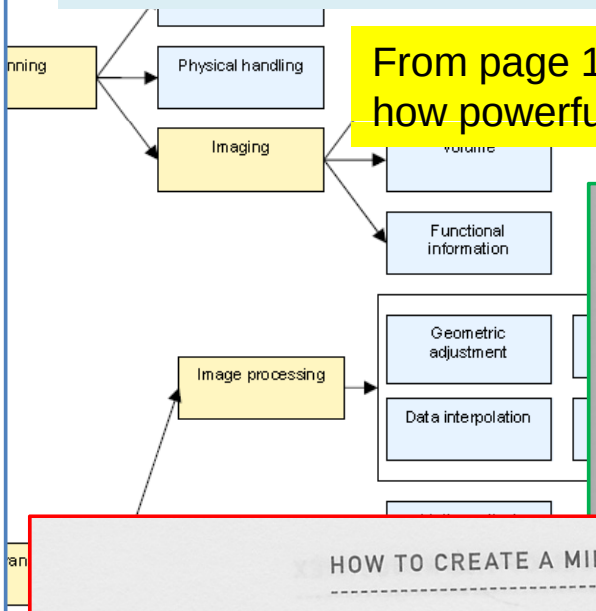

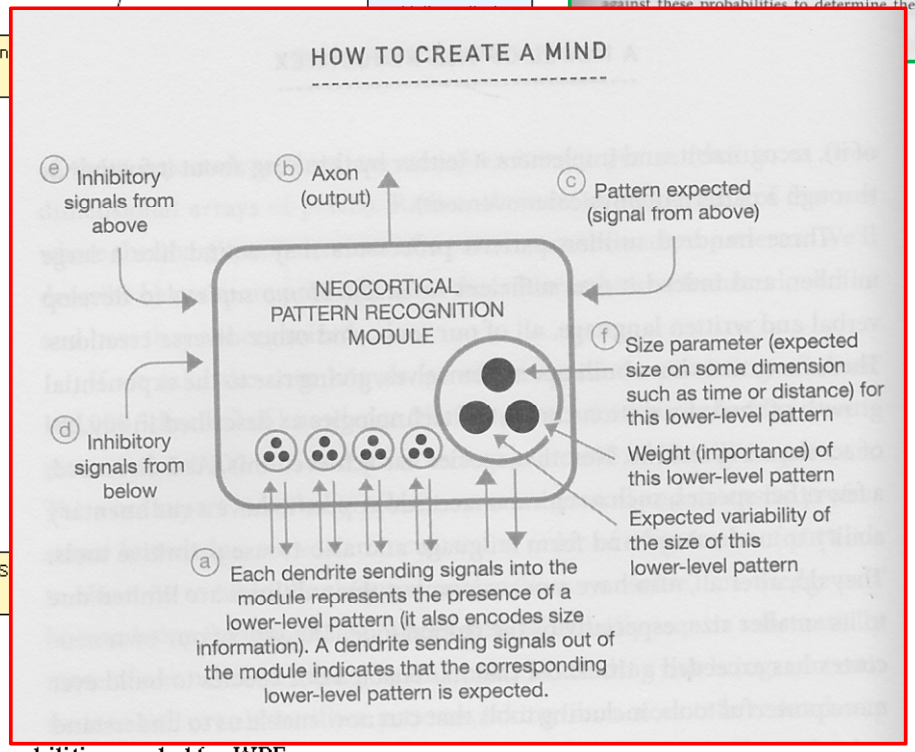

Figure 4: Technological capabilities needed for WBE.

Ray has spent his whole life studying pattern recognition.

Background chart is from "Whole Brain Emulation—A Roadmap" by Anders Sandberg and Nick Bostrom, referenced in Kurzweil's book.

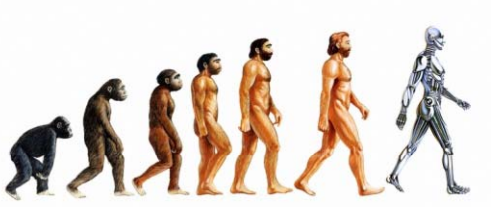

Natural vs. Supernatural

- If you think humans are special and have eternal souls created by God, you may have a hard time buying the concept of people building machines that **become people**
- If you think people evolved by an entirely natural process without any supernatural intervention, then humans creating **a new intelligent species** is not that hard to accept
  - What if people build complex computing machines that develop the right conditions that a "soul" can enter them (the way it would in a baby,) and they become alive and conscious?
    - *This question is for the first group above*
- What if people build machines that appear alive and conscious and you **cannot prove** that they aren't?

Summarizing a few of the loudest voices

- Ray Kurzweil—we will become the new species (he's a technology optimist)
- Hans Moravec—we will build the new species and they will be our "mind children" (replacing us)
- Bill Joy—The new species doesn't need us (he's a technology pessimist)
- Jaron Lanier—There won't be a new species
- Roger Penrose—People are special. Can't build them

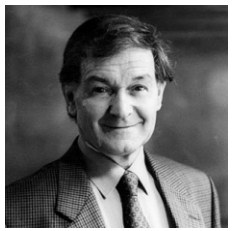

Kurzweil's Law of Accelerating Returns

- "An analysis of the history of technology shows that **technological change is exponential**, contrary to the common-sense "intuitive linear" view. So we won't experience 100 years of progress in the 21st century — **it will be more like 20,000 years of progress** (at today's rate). The "returns," such as chip speed and cost-effectiveness, also increase exponentially. There's even exponential growth in the rate of exponential growth."

The convergence of 3 technologies

- Kurzweil argues that the convergence of **genetics**, **nanotechnology** and **robotics** will lead to the Singularity
- He says that we are on the verge of Epoch 5, the merger of technology and Human Intelligence
- There are way too many arguments and counter-arguments in "*The Singularity is Near*" to summarize here. Highly recommended reading.

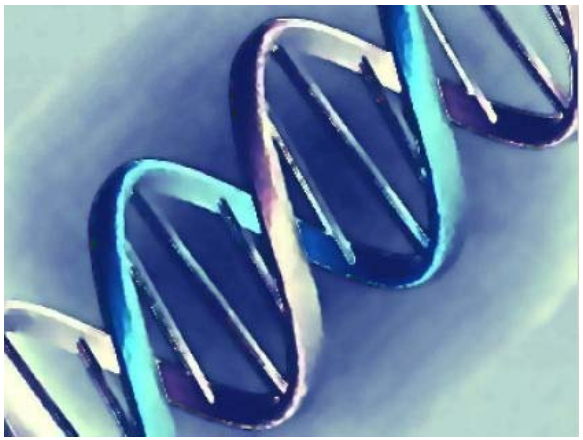
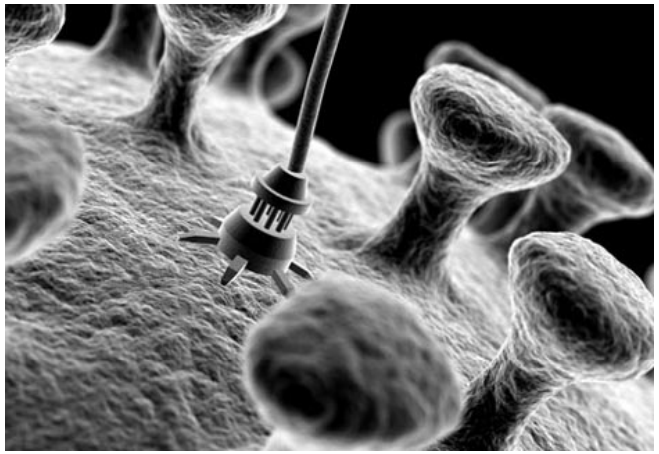
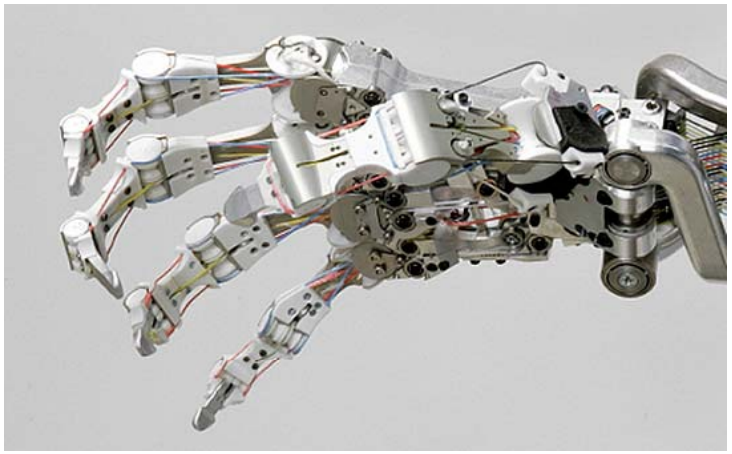

What is “the Singularity”?

The Singularity is the technological creation of smarter-than-human intelligence

- *“A future that contains smarter-than-human minds is genuinely different in a way that goes beyond the usual visions of a future filled with bigger and better gadgets. Vernor Vinge originally coined the term "Singularity" in observing that, just as our model of physics breaks down when it tries to model the singularity at the center of a black hole, our model of the world breaks down when it tries to model a future that contains entities smarter than human.”*
- The singularity is a time after which “normal” human people cannot understand history and may not be relevant to it

Taken from “The Singularity is Near” by Ray Kurzweil

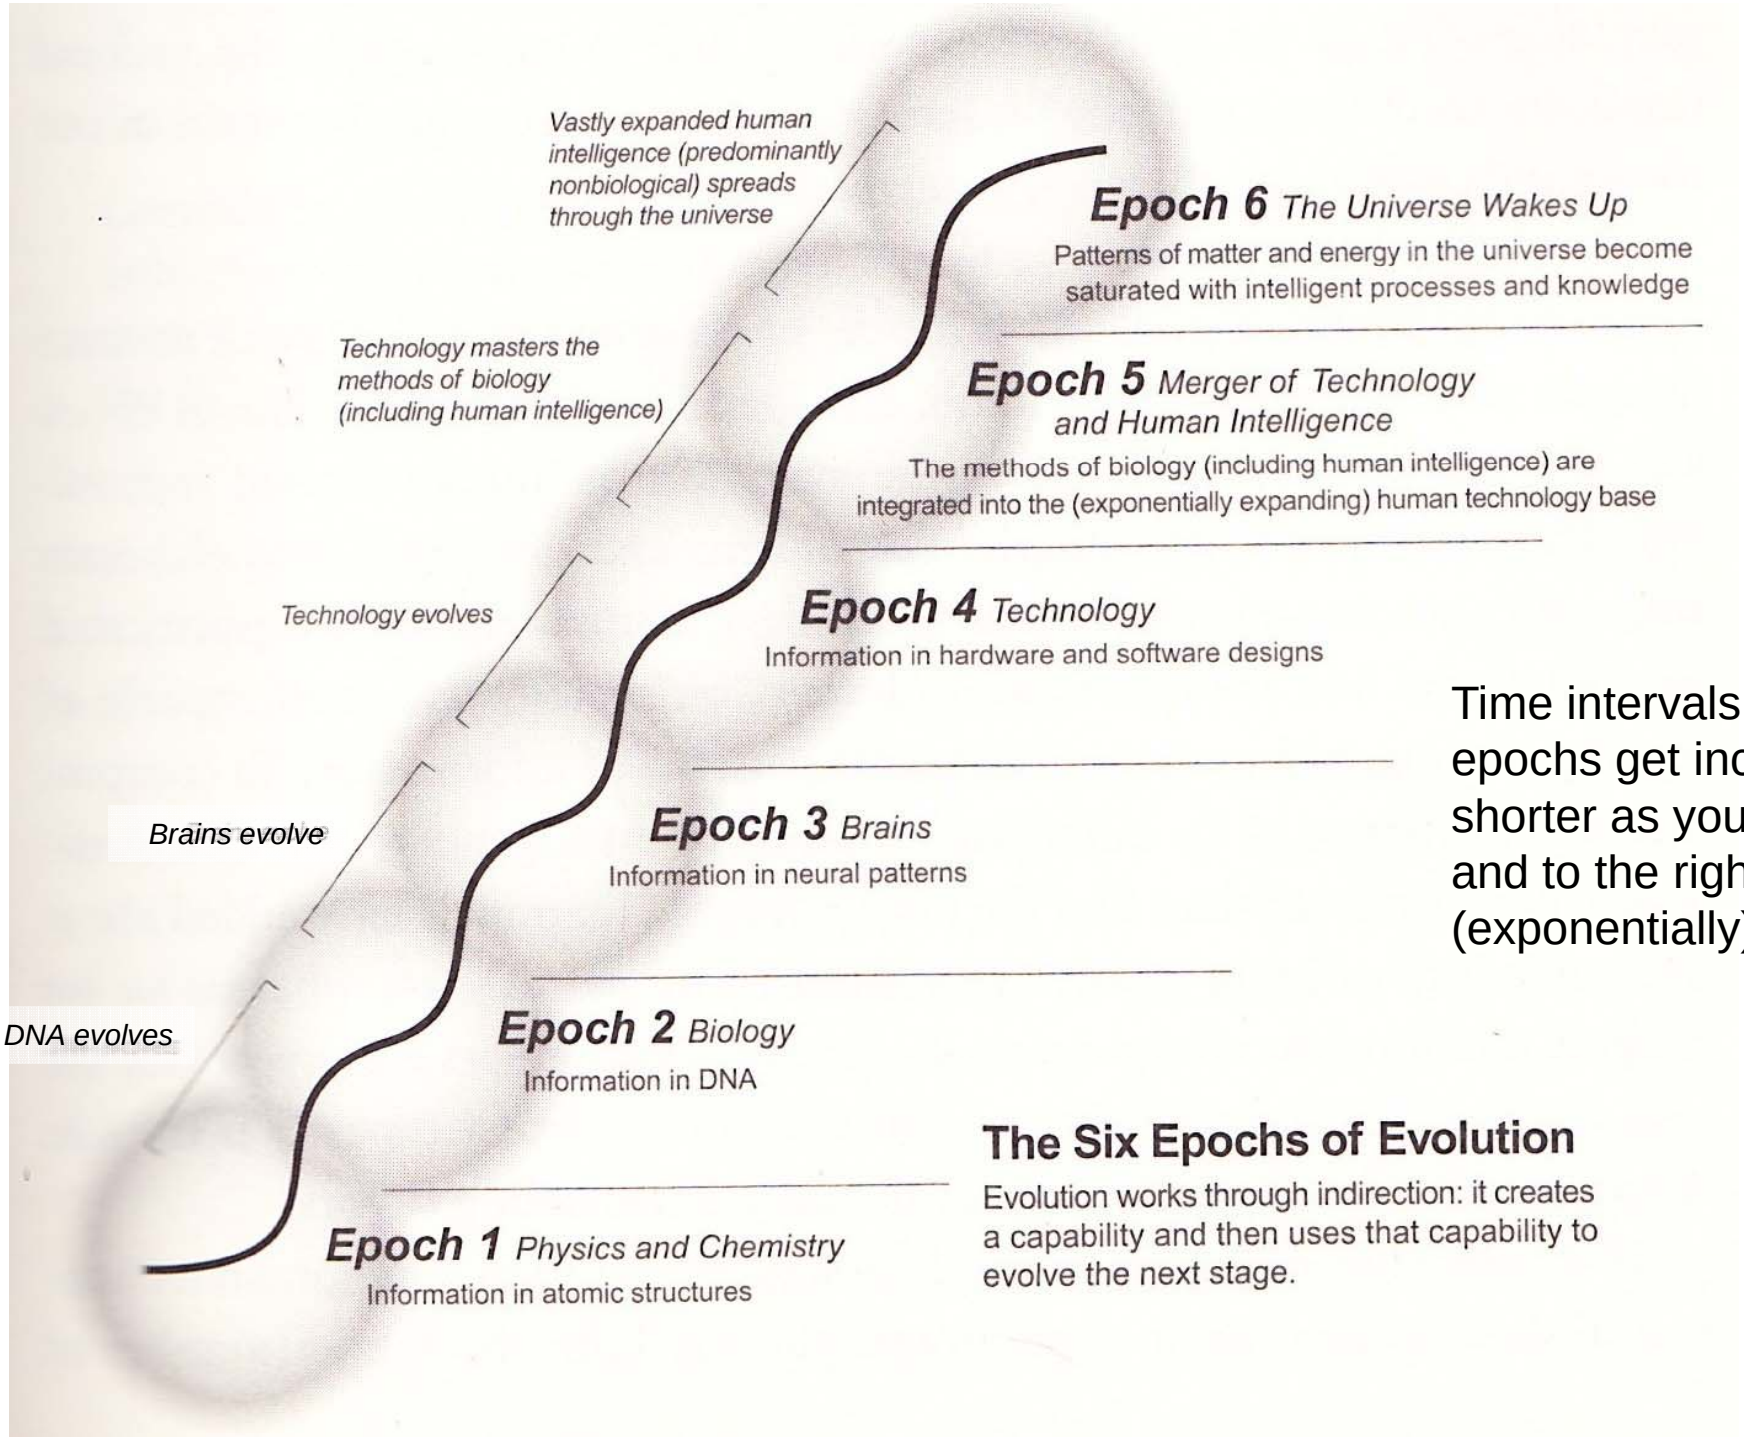

Time intervals between epochs get increasingly shorter as you go up and to the right (exponentially)

Taken from “The Singularity is Near” by Ray Kurzweil

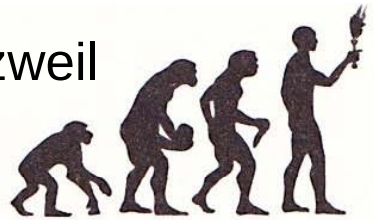

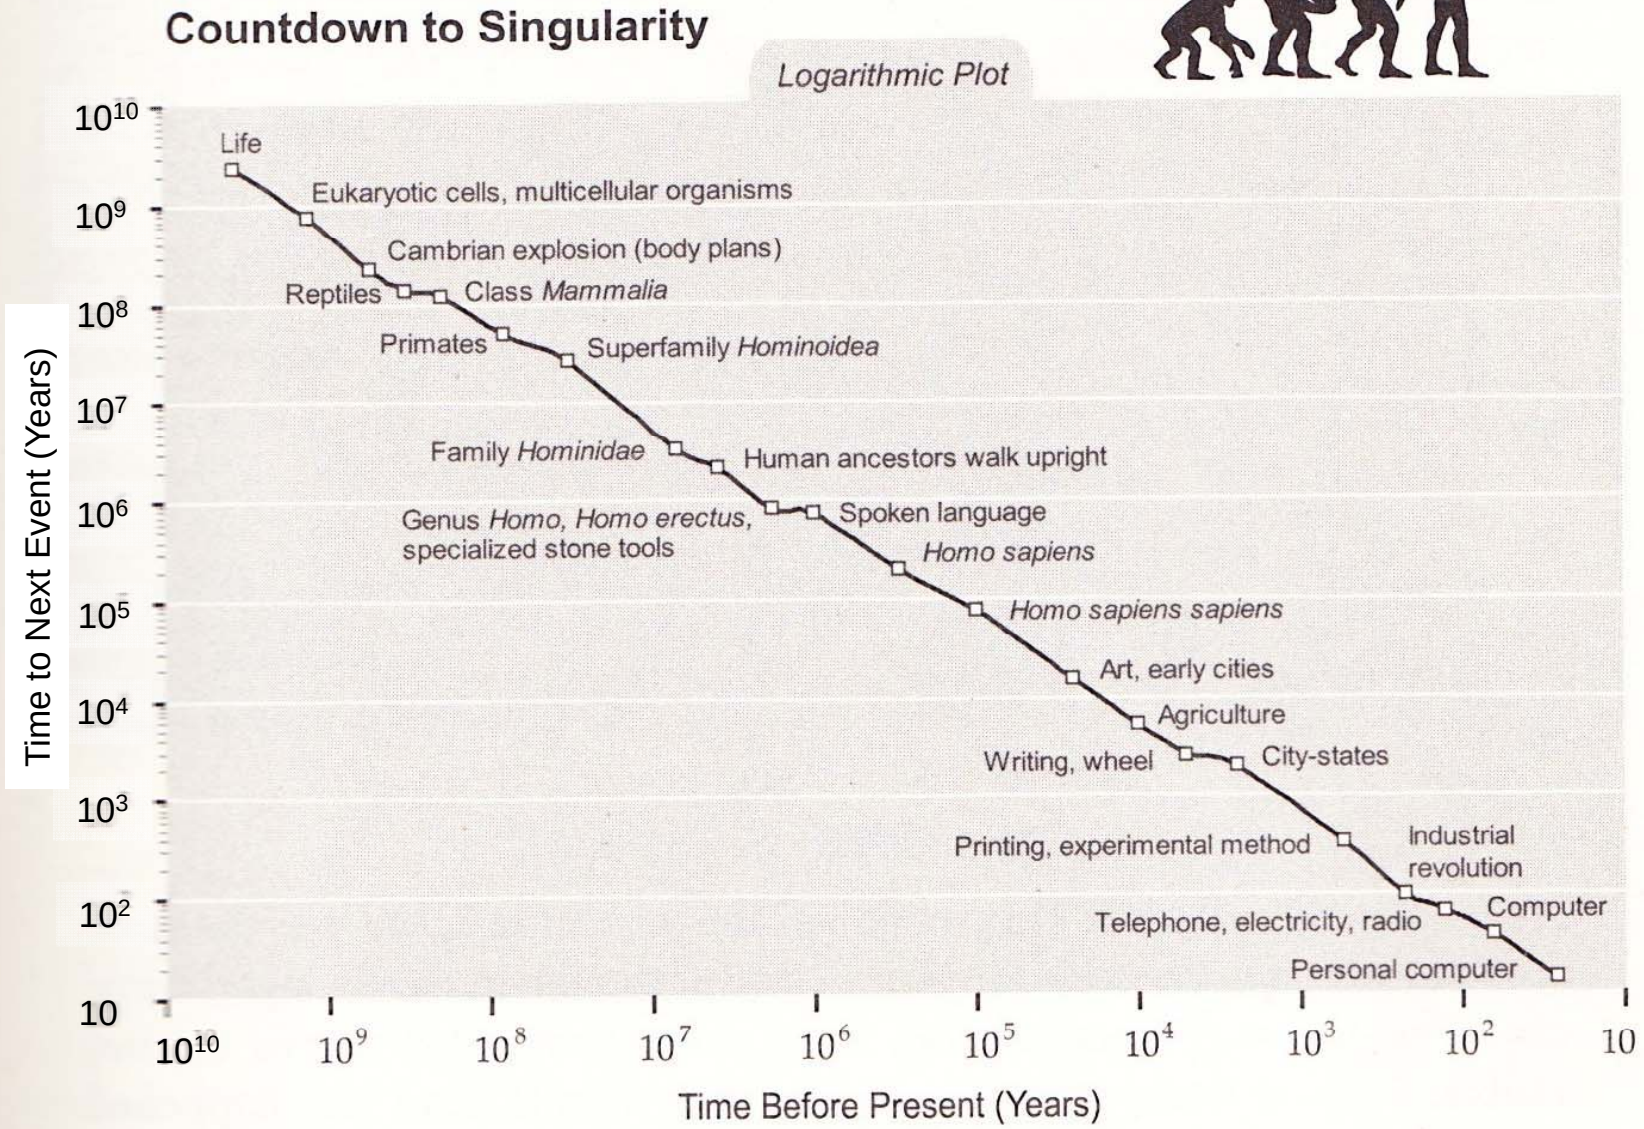

Countdown to Singularity: Biological evolution and human technology both show continual acceleration, indicated by the shorter time to the next event (two billion years from the origin of life to cells; fourteen years from the PC to the World Wide Web).

What is Nanotechnology?

CRN

Center for Responsible Nanotechnology

SEARCH www

Our Ideas

C-R-News

About CRN

New! Nanotech Scenario Series

Nano Tomorrows

Get Involved

Donate to CRN

FAQ

What is Nanotechnology?

A basic definition: **Nanotechnology is the engineering of functional systems at the molecular scale.** This covers both current work and concepts that are more advanced.

In its original sense, 'nanotechnology' refers to the projected ability to construct items *from the bottom up*, using techniques and tools being developed today to make complete, high performance products.

With 15,342 atoms, this parallel-shaft speed reducer gear is one of the largest nanomechanical devices ever modeled in atomic detail. LINK

It all started with physicist Richard Feynman's famous 1959 lecture "*There's Plenty of Room at the Bottom"* in which he talked about manipulating matter at the atomic scale

The Holy Grail of nanotech ***used to be*** the self-replicating nano assembler—a microscopic robot capable of picking and placing atoms to build anything

ScienceDaily®

Your source for the latest research news

ALWAYS REMEMBER YOU'
WE WILL.

News | Articles | Videos | Images | Books

Health & Medicine | Mind & Brain | Plants & Animals | Earth & Climate | Space & Time

Science News

Neurochip Technology Developed: Advances to Further Brain Research of Diseases Such as Alzheimer's and Parkinson's

ScienceDaily (Aug. 10, 2010) — The University of Calgary, Faculty of Medicine scientists who proved it is possible to cultivate a network of brain cells that reconnect on a silicon chip -- or the brain on a microchip -- have been involved in the development of new technology that monitors brain cell activity at a resolution never achieved before.

Unique Micronail Chip Makes Electronics and Bio Cells Communicate

ScienceDaily (Nov. 11, 2009) — IMEC presents a

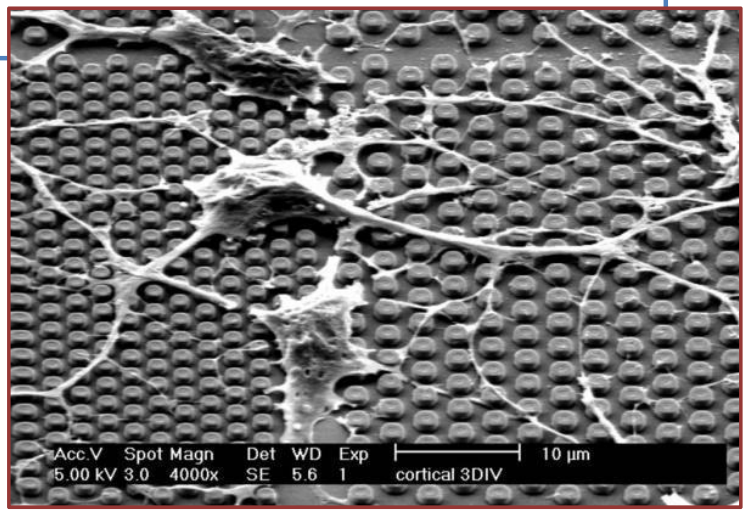

Interfacing our biological nervous system with computer hardware is another ongoing research area.

This is how brains and computers could talk directly.

This area of study is a type of nanotechnology

Jaron Lanier's 6 beliefs that lead to "Cybernetic Totalism"

from his article in Wired, "Half of a Manifesto" 12/00

1. Cybernetic patterns of information provide the ultimate and best way to understand reality.
2. People are no more than cybernetic patterns.
3. Subjective experience either doesn't exist, or is unimportant because it is some sort of ambient or peripheral effect.

Beliefs that lead to “Cybernetic Totalism” cont’d

4. What Darwin described in biology, or something like it, is in fact also the singular, superior description of all creativity and culture.
5. Qualitative as well as quantitative aspects of information systems will be inexorably accelerated by Moore's law.
6. Biology and physics will merge with computer science (becoming biotechnology and nanotechnology), resulting in life and the physical universe becoming mercurial; achieving the supposed nature of computer software.

Beliefs that lead to “Cybernetic Totalism” cont’d

“Furthermore, all of this will happen very soon! Since computers are improving so quickly, they will overwhelm all the other cybernetic processes, like people, and will fundamentally change the nature of what's going on in the familiar neighborhood of Earth at some moment when a new "criticality" is achieved - maybe in about the year 2020. To be a human after that moment will be either impossible or something very different than we now can know.”

Lanier doesn’t believe any of this.

More criticisms of the Singularity

- Jaron Lanier argues that while hardware gets faster, software doesn't progress on the same curve
  - It is "brittle" and "locked in" by legacies
- Kurzweil counters this:
  - "Even more remarkable—and even less widely understood—is that in many areas, *performance gains due to improvements in algorithms have vastly exceeded even the dramatic performance gains due to increased processor speed*. The algorithms that we use today for speech recognition, for natural language translation, for chess playing, for logistics planning, have evolved remarkably in the past decade ... "

More from Kurzweil on algorithms

- Grötschel, an expert in optimization, observes that a benchmark production planning model solved using linear programming would have taken **82 years to solve in 1988**, using the computers and the linear programming algorithms of the day.
- **Fifteen years later—in 2003—this same model could be solved in roughly one minute**, an improvement by a factor of roughly 43 million.
- Of this, ***a factor of roughly 1,000 was due to increased processor speed, whereas a factor of roughly 43,000 was due to improvements in algorithms!***

The (fictional) case of BINA48

- "Attorney Dr. Martine Rothblatt filed a motion for a preliminary injunction to prevent a corporation from disconnecting an intelligent computer in a **mock** trial at the International Bar Association conference in San Francisco, Sept. 16, 2003.
- **The issue could arise in a real court within the next few decades**, as computers achieve or exceed the information processing capability of the human mind and the boundary between human and machine becomes increasingly blurred."
  - http://www.kurzweilai.net/biocyberethics-should-we-stop-a-company-from-unplugging-an-intelligent-computer

BINA48's case comes down to:

- **How do you know what a person is?**
- BINA48's lawyer referred to it as a human being because the computer convinced her that it had a "human" mind
- The corporation's lawyer argued that BINA48 could not be a person because it didn't have a subjective internal "life"
  - How do any of us know who has a subjective experience?
- The judge decided the computer did not have standing with the court—then he stayed his own decision for a higher court to decide

The Turing Test

- According to Wikipedia:
  - The **Turing test** is a test of a machine's ability to exhibit intelligent behaviour. In Turing's original illustrative example, a human judge engages in a natural language conversation with a human and a machine designed to generate performance indistinguishable from that of a human being. All participants are separated from one another. **If the judge cannot reliably tell the machine from the human, the machine is said to have passed the test.**
- Paraphrasing Supreme Court Justice Potter Stuart on obscenity:
  - We don’t know how to define what being human is, but we know it when we see it.

Many criticisms of the Turing Test

- People are sometimes fooled by "chatterbots" that have no mind, they just follow a large set of rules to sound human
- People fall into the anthropomorphic fallacy
  - Giving inanimate objects human attributes
- The test focuses on how the machine acts
  - John Searle argues the difference between "actually" thinking and "simulating" thinking
  - His chinese room argument is intended to show that, even if the Turing test is a good operational definition of intelligence, it may not indicate that the machine has a mind, consciousness, or intentionality.
- The bottom line is: **We don't have a good way to tell who is human**
  - I know people who couldn't pass the Turing Test!
  - If it looks like a duck and quacks like a duck, is it a duck?

Watson on Jeopardy

- IBM built a computer that beat the two highest winners of Jeopardy in February of 2011.
- Shows are available on YouTube.
  - They are very worth watching.
- **"Watson acquired most of its knowledge on its own by reading natural language documents such as encyclopedias."**
  - http://www.technologyreview.com/blog/guest/27263/

Ken Jennings joked, "I for one welcome our new computer overlords," written under one of his answers.

Thought experiment

- Kurzweil describes a hypothetical process replacing a human brain one nerve cell at a time with a non-organic machine neuron (using nanotechnology)
- The end result would be a super-computer version of your own brain millions of times faster than the original
- **Would you still be the same person?**
- Would you accept this conversion?
- What if you could get a hardware extension of your brain without changing any of your original brain?
  - You'd get the ability of computers to remember perfectly, be interconnected with others and the internet, make fast computations and *be backed-up*
  - Would you take this option?
  - **Could you resist if you worked in an environment where others did it and your income depended on it?**

As I was writing this talk…

- I realized that the expression “building new people” could be interpreted in two ways:
  1. Building brand new sentient beings (strong AI), which is what we have discussed up to now.
  2. Transforming ourselves into “new people”

- Let’s take a look at how we are transforming the human brain and body with new technology

Brain signals sent over the internet from one brain to another

August 27, 2013

Researcher controls colleague's motions in 1st human brain-to-brain interface

Doree Armstrong and Michelle Ma

News and Information

Posted under: Engineering, News Releases, Research, Science, Technology

University of Washington researchers have performed what they believe is the first noninvasive human-to-human brain interface, with one researcher able to send a brain signal via the Internet to control the hand motions of a fellow researcher.

Look at the date on this article!

University of Washington

University of Washington researcher Rajesh Rao, left, plays a computer game with his mind. Across campus, researcher Andrea Stocco, right, wears a magnetic stimulation coil over the left motor cortex region of his brain. Stocco's right index finger moved involuntarily to hit the "fire" button as part of the first human brain-to-brain interface demonstration.

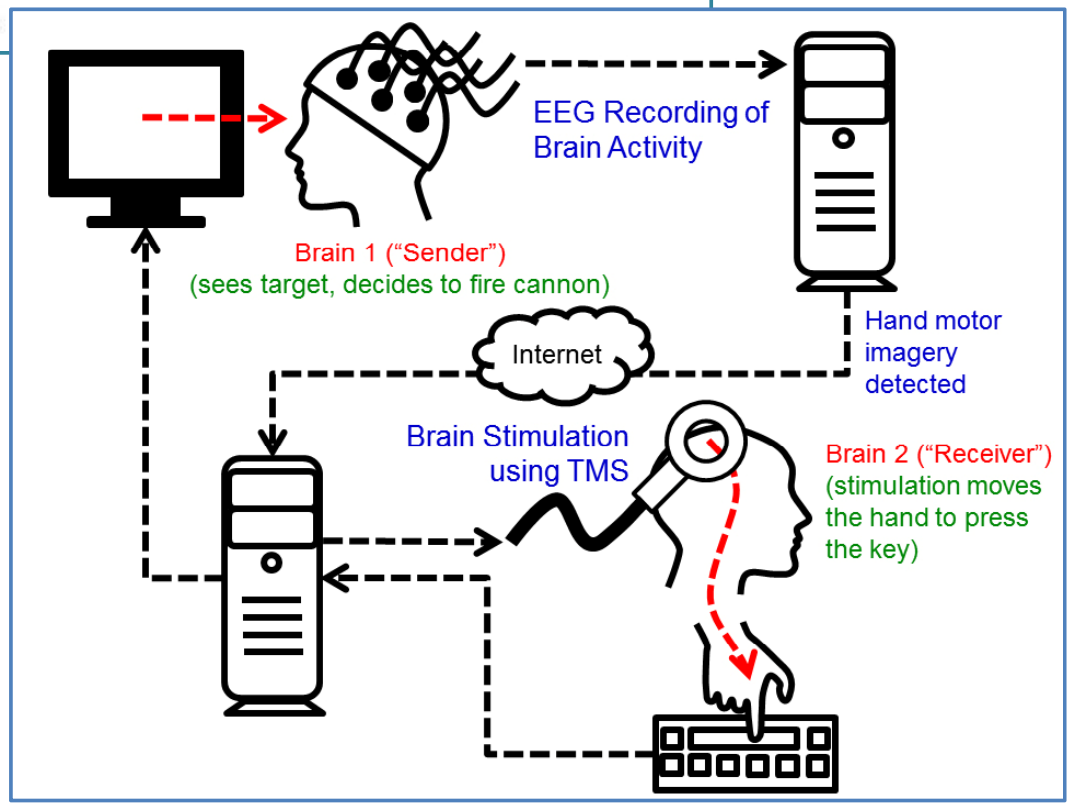

This “bionic man” is on display at the Smithsonian and is built from many of the replacement parts available for humans today

This is the September, 2013 issue

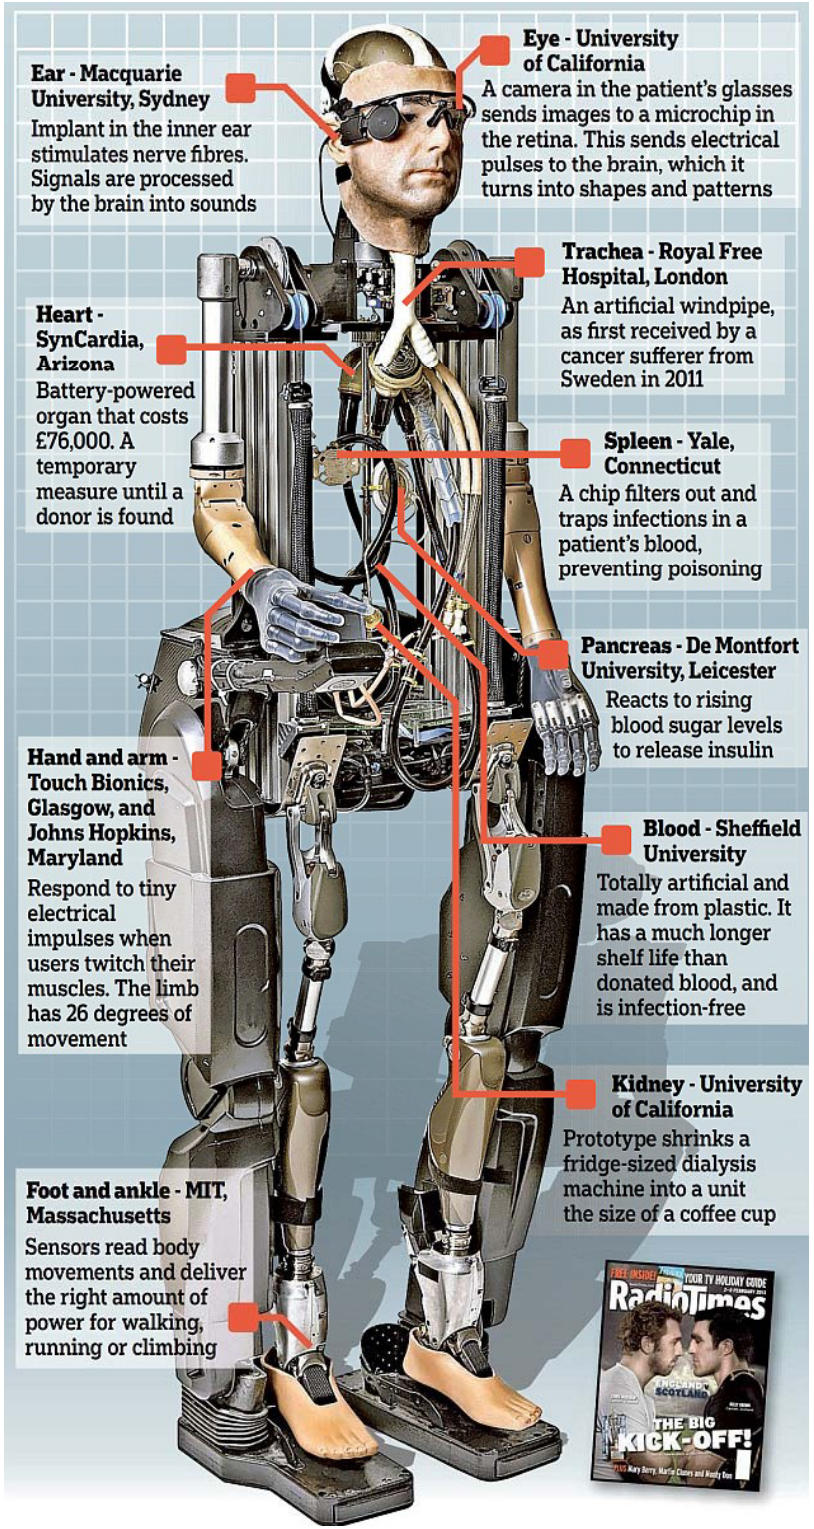

This article is in the September issue of “Mechanical Engineering”, the monthly journal of the American Society of Mechanical Engineers

It states that our attitudes about robots have crossed a line.

•“Many people view robots as safer than people.”

•[Siri] “…will always listen to us, and never disappoint us, get angry or cause conflicts. It is friendship without mutuality since robots have no needs or desires of their own.”

•“Most of Turkle’s MIT colleagues believe the need for caretaker robots for the elderly is self-evident.”

THE

ROBOTIC MOMENT

MACHINES OFFER CARING BEHAVIOR. IS THAT A GOOD THING?

BY ALAN S. BROWN

Sherry Turkle, the psychologist who directs MIT’s Initiative on Technology and Self, sometimes tells the anecdote about how she was once approached by a young psychology graduate student after a talk about reactions to “sociable” robots that can mimic and evoke emotions. Anne, the student, confided that she would trade in her boyfriend for a sophisticated Japanese robot if it could provide what she described as “caring behavior.”

According to Turkle, who retold the anecdote during a plenary lecture to the American Association for the Advancement of Science in Boston this past February, “Anne told me that she relied on a feeling of civility in her house and did not want to be alone. She said, ‘If the robot could provide the environment, I would be happy to produce the illusion that there was somebody really with me.’”

Summary

- Things are changing faster all the time
- The successes build on one another making the next effort more capable and accelerating change
- There are many different efforts all over the world to create new technologies
- There are many economic pressures pushing the evolution of technology
  - Some people think these pressures make the Singularity inevitable
  - "You can't stop Judgment Day. You can only postpone it."—the Terminator Robot, Terminator 3: Rise of the Machines
- **How much influence do we have on the future?**

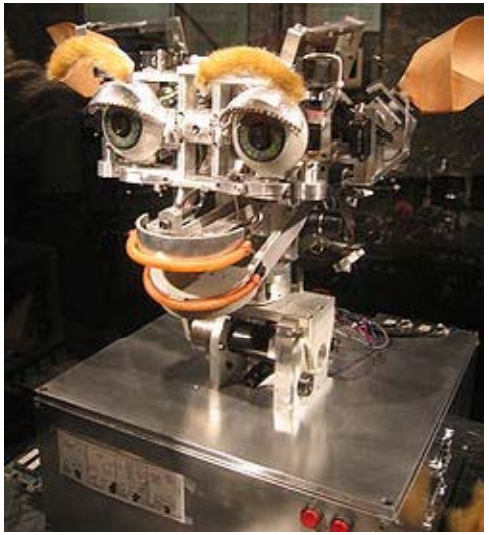

One of my favorite questions:

Will we make good pets?

(Think about how many ways this can be interpreted.)

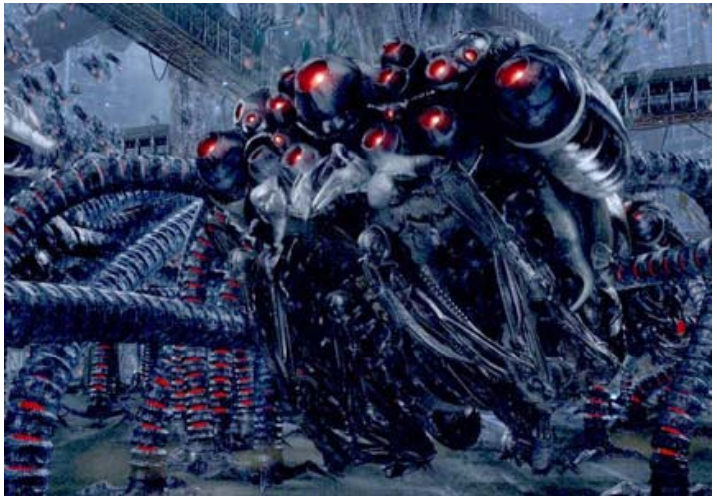

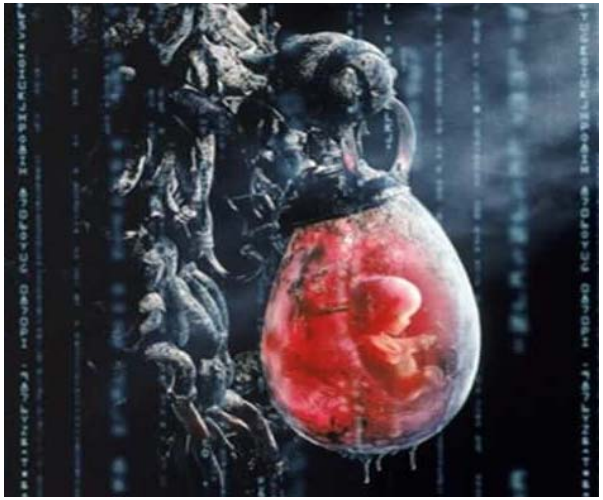

Backup Slides

Movies I recommend

- Transcendent Man (2009)
- I, Robot (2004)
- A.I. Artificial Intelligence (2001)
- The Matrix (1999)
- Bicentennial Man (1999)
- The Terminator (1984)
- Blade Runner (1982)
- Colossus: The Forbin Project (1970)
- 2001: A Space Odyssey (1968)

Some Non-fiction Books I recommend

(Acknowledgement to Mike Ososky)

- *How to Create a Mind* (2012) by Ray Kurzweil
- *Transcend—Nine Steps to Living Well Forever* (2009) by Ray Kurzweil
- *Love and Sex with Robots* (2007) by David Levy
- *The Singularity is Near* (2005) by Ray Kurzweil
- *Flesh and Machines—How Robots will Change Us* (2002) by Rodney Brooks
- *Robot : Mere Machine to Transcendent Mind* (2000) by Hans P. Moravec
- *The Age of Spiritual Machines* (1999)—*When Computers Exceed Human Intelligence* by Ray Kurzweil
- *The Meme Machine* (1999) - by Susan Blackmore

More non-fiction books

- *Guns, Germs, and Steel* (1998) - by Jared Diamond
  - This is a history of the last 15,000 years. It's the history of memetic evolution in humans. Taken with Rare Earth this is the history of our planet from birth to current times.
- *Engines of Creation* (1986) - by K. Eric Drexler
  - The Classic ground breaking book on nanotechnology. Predicts the coming of the internet. This is where 'gray goo' was coined, Drexler is a nano wizard and future visionary. This is what freaked Bill Joy (chief scientist at Sun).
- *The Mind's I* (1981) - by Douglas Hofstadter and Daniel Dennett
  - Introduction to philosophical issues that deal with AI. Includes introduction to the Turing test.
- *Godel, Escher, Bach* (1979) - by Douglas Hofstadter
  - A metaphorical Fugue on Minds and Machines in the Spirit of Lewis Carroll. One of the first proposals of intelligence being an emergent phenomena of system dynamics. One of the deepest intellectually challenging books you can read.

Some SF books I recommend

- *Accelerando* (2005) by Charles Stross
  - Anything by Charles Stross
- *The Footprints of God* (2003) by Greg Iles
- *Evolution's Darling* (1999) by Scott Westerfeld
- *Einstein's Bridge* (1999) by John Cramer
- *I, Robot* (1950) - by Isaac Asimov
  - The classic that introduces us to the three laws of robotics.
- *Too many others to mention…*

Websites for further reading

- http://www.kurzweilai.net/
  - Excellent summaries of research going on all over the world
- http://singularity.org/
  - Singularity University
- http://www.wired.com/wired/archive/8.04/joy.html
  - Why the Future Doesn't Need Us
- http://e-drexler.com/
  - Check out the nanofabricator movie!!
- http://www.edge.org/
- http://www.ted.com/
- http://www.wired.com/wired/archive/10.12/holytech_pr.html
- http://www.extropy.org/
  - Source for Transhumanism information
- http://www.youtube.com/watch?v=yj_-sBNQKcQ
  - A criticism of Transhumanism
- http://www.youtube.com/watch?v=d8W1WuxGniE
  - The original "Did you know?" video—a little old now