

Artificial Intelligence and the Singularity

piero scaruffi

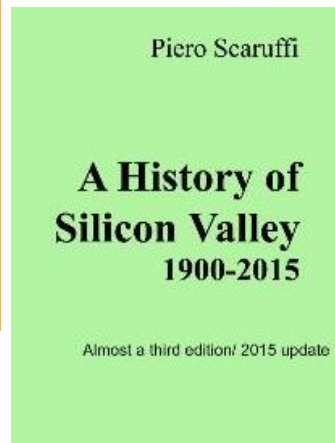
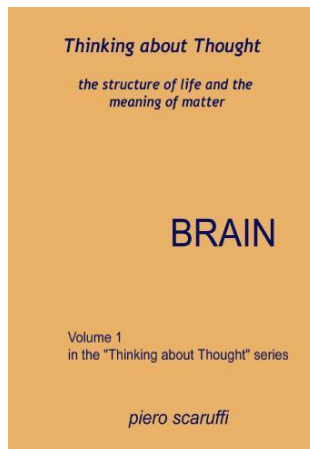
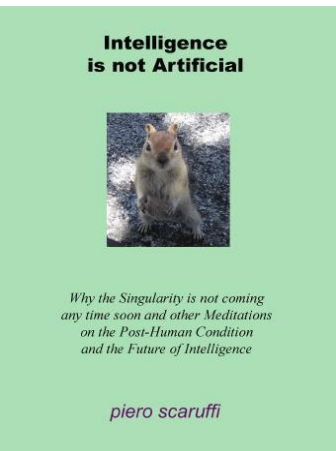
www.scaruffi.com

October 2014 - Revised 2016

*"The person who says it cannot be done should not
interrupt the person doing it" (Chinese proverb)*

Piero Scaruffi

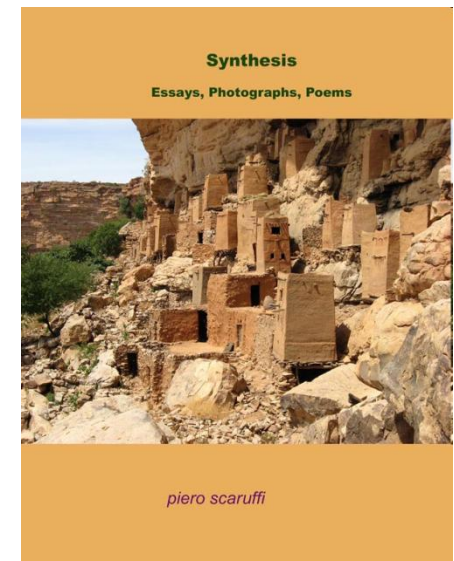
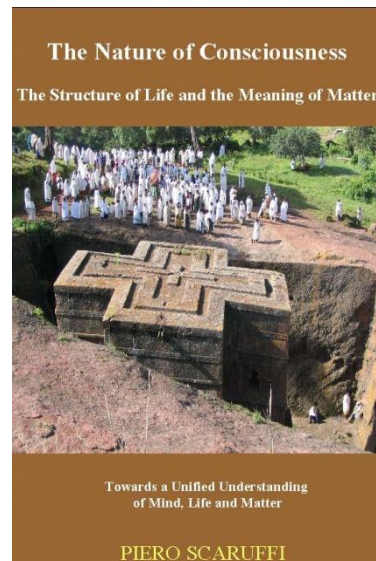
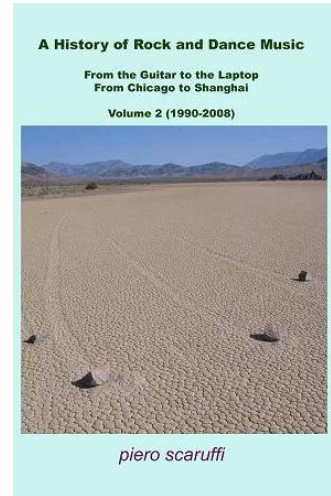
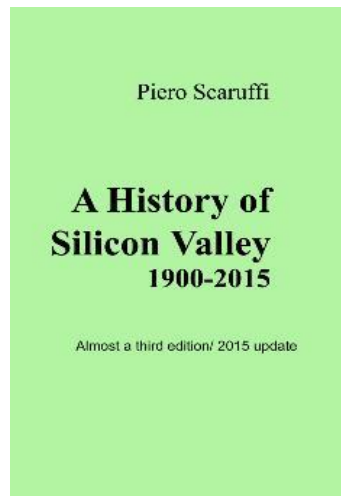
- piero scaruffi
p@scaruffi.com
scaruffi@stanford.edu



Olivetti AI Center, 1987

Piero Scaruffi

- Cultural Historian
- Cognitive Scientist
- Blogger
- Poet
- www.scaruffi.com



This is Part 7

- See <http://www.scaruffi.com/singular> for the index of this Powerpoint presentation and links to the other parts
 1. Classic A.I. - The Age of Expert Systems
 2. The A.I. Winter and the Return of Connectionism
 3. Theory: Knowledge-based Systems and Neural Networks
 4. Robots
 5. Bionics
 6. Singularity
 7. Critique
 8. The Future
 9. Applications
 10. Machine Art
 11. The Age of Deep Learning
 12. Natural Language Processing

A Critique of Artificial Intelligence

The State of A.I.

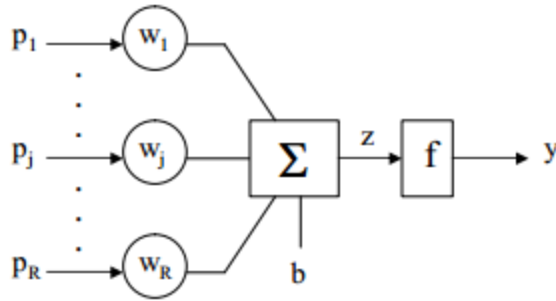
零

- How many robots do you see in this room?
- How many robots did you see today in the street?
- If you have seen one, it was probably a toy in a store for children

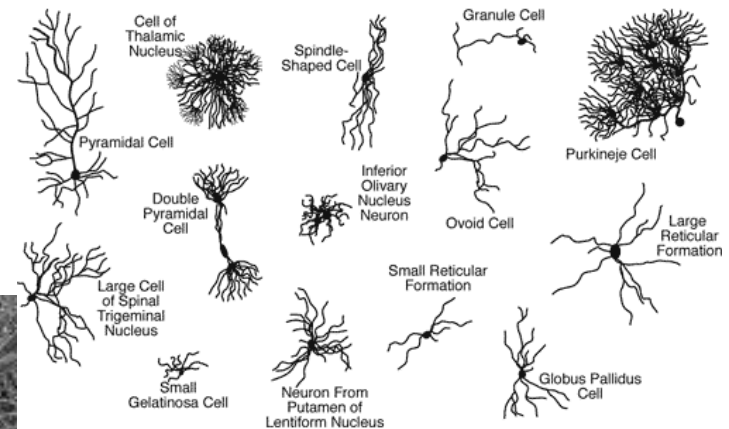
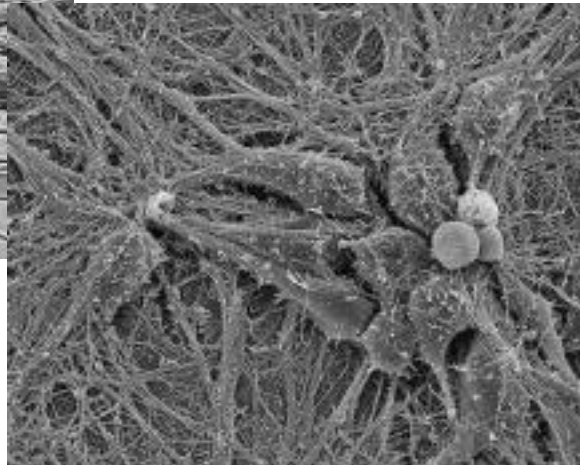
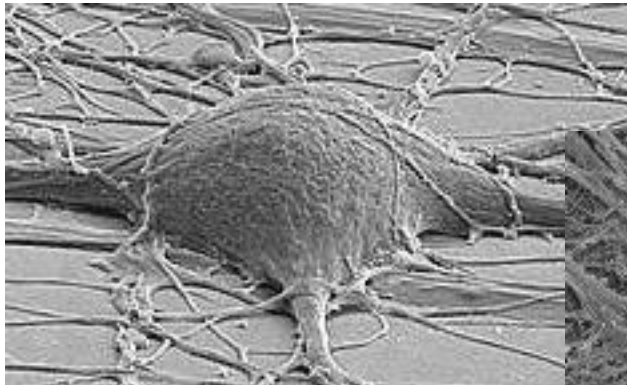


Animal Brain vs Electronic Brain

The McCulloch-Pitts neuron (1943)



Biological neurons:



70+ types in the retina alone

No two neurons are alike

Three topologies in the human brain

- network (eg thalamo-cortical system)
- loop (eg cortex-hippocampus)
- fan (Edelman's "value systems")

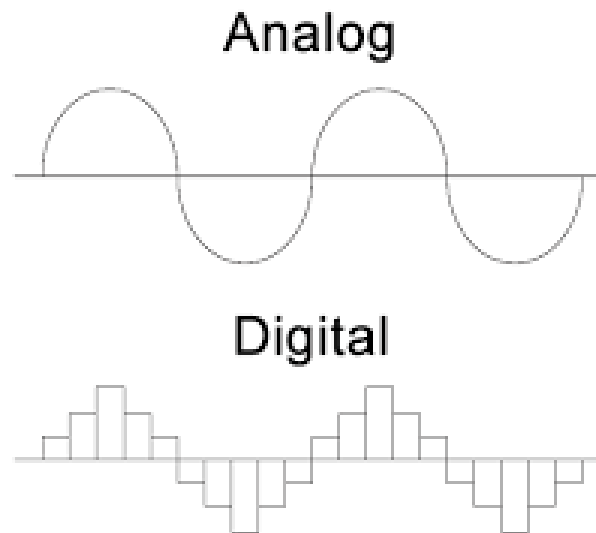
Animal Brain vs Electronic Brain

Biological neurons:

Massively parallel computation

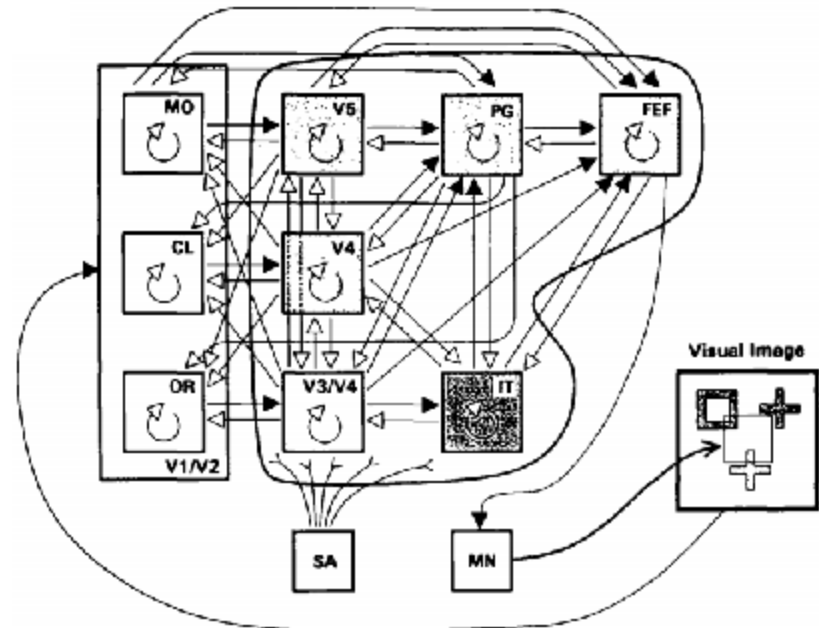
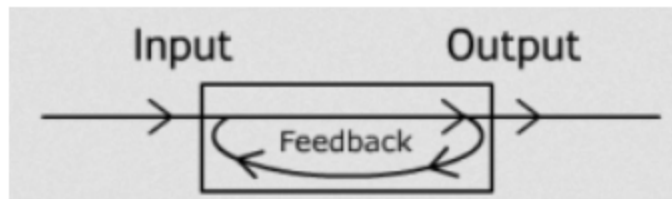
Analog information processing

Implementation in “wetware”



Animal Brain vs Electronic Brain

Wiener's feedback vs Edelman's reentry

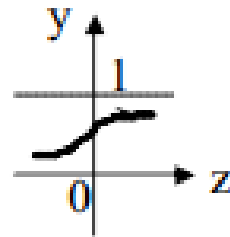


Giulio Tononi, Olaf Sporns, and Gerald Edelman

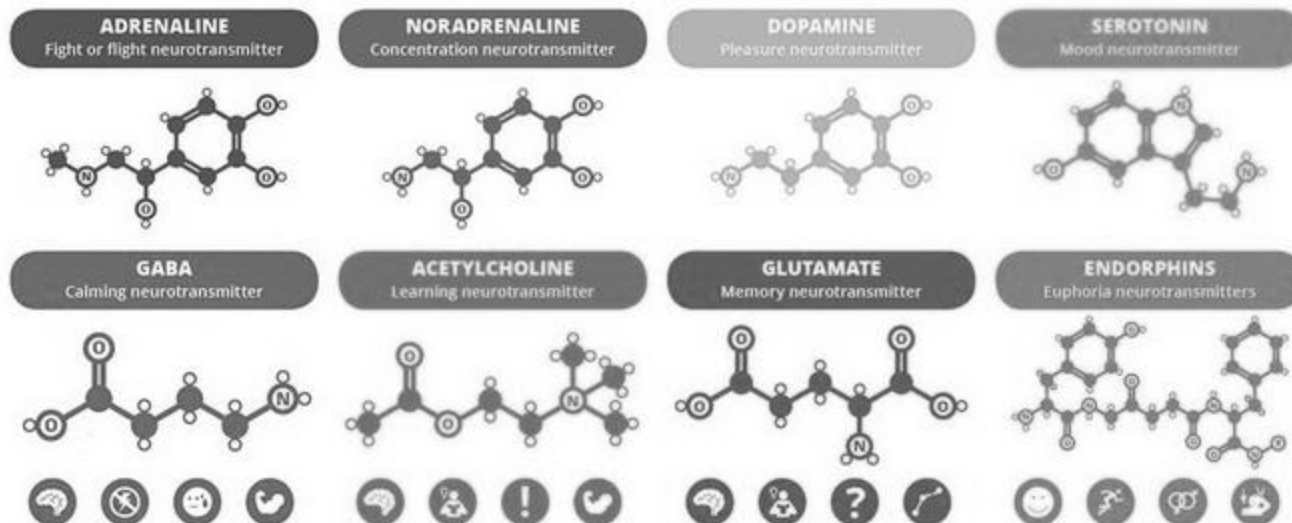
Animal Brain vs Electronic Brain

The McCulloch-Pitts neurotransmitter (1943)

Log-Sigmoid:
 $y = 1/(1+e^z)$

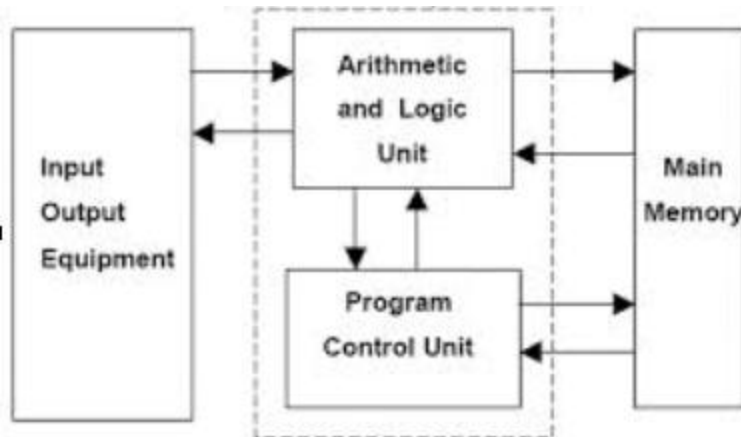


Biological neurotransmitters:

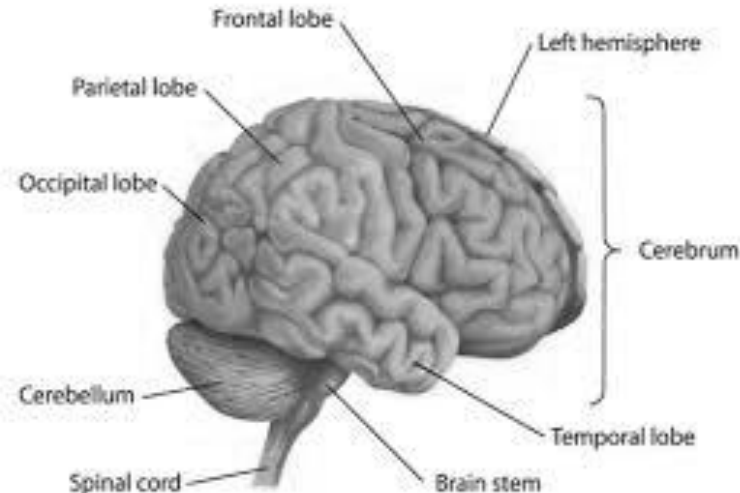


Animal Brain vs Electronic Brain

The VonNeumann computer architecture (1945)



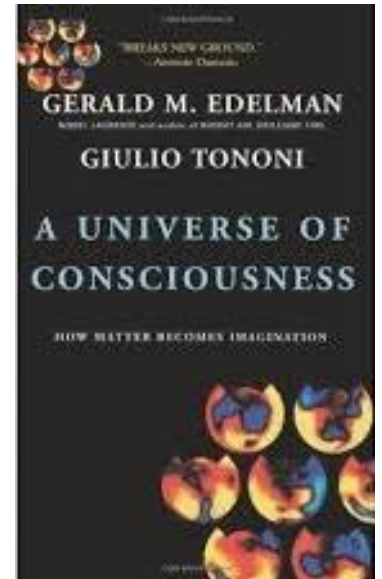
Brain architecture



Animal Brain vs Electronic Brain

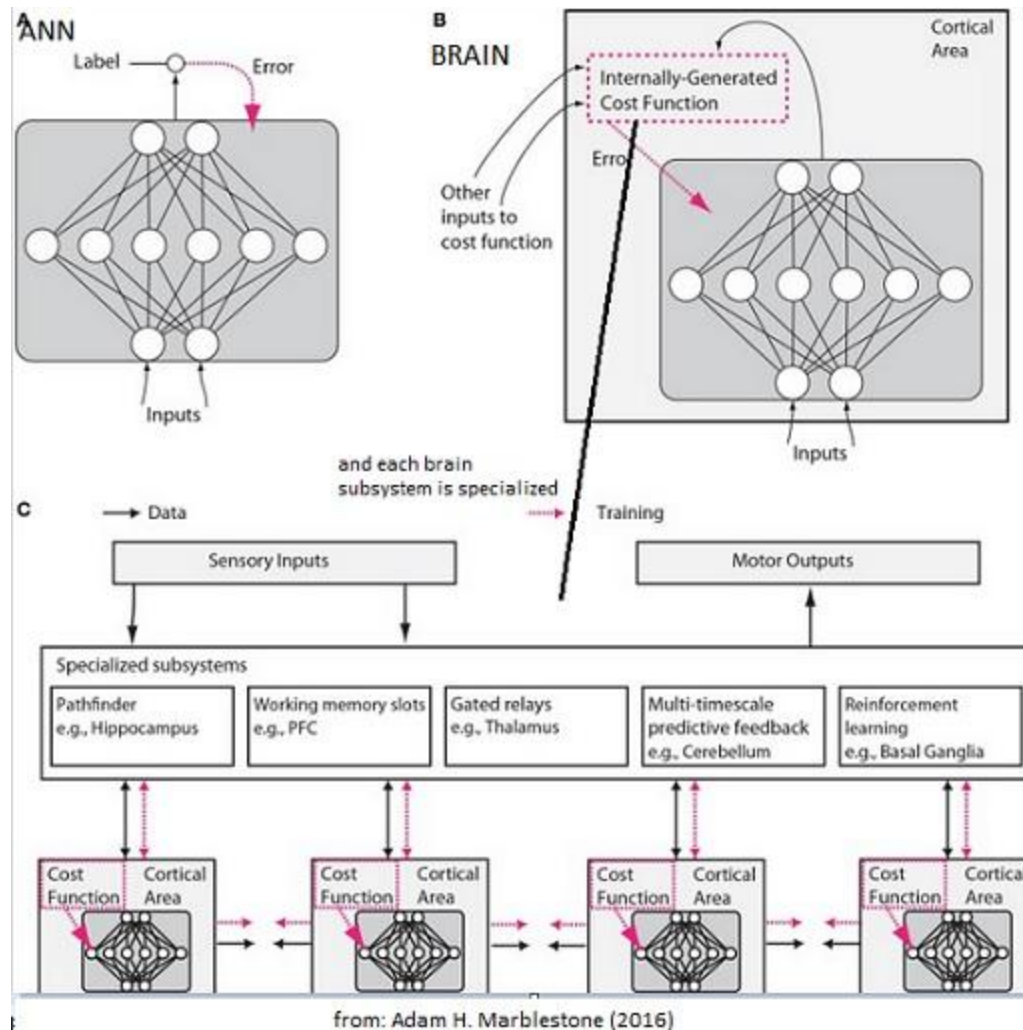
The brain is not a "network":

- some regions are networks (notably the thalamo-cortical system)
- other regions are loops (notably between the cortex and the cerebellum and the hippocampus)
- other regions are fans (Edelman's "value systems" that project into the whole brain)

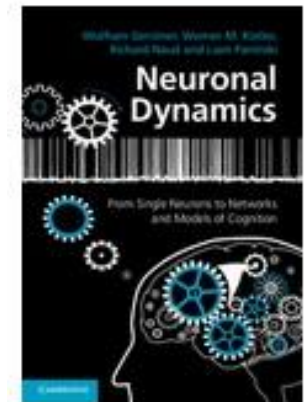


Animal Brain vs Electronic Brain

The brain consists of many subsystems



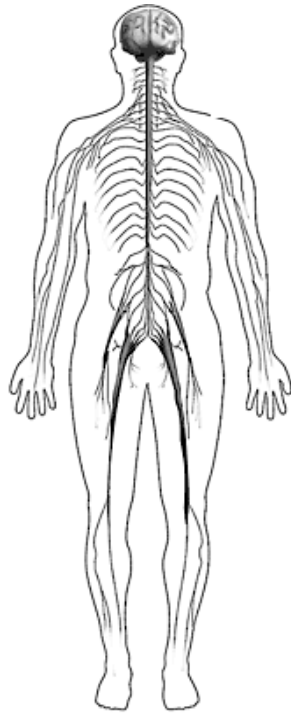
Wulfram Gerstner



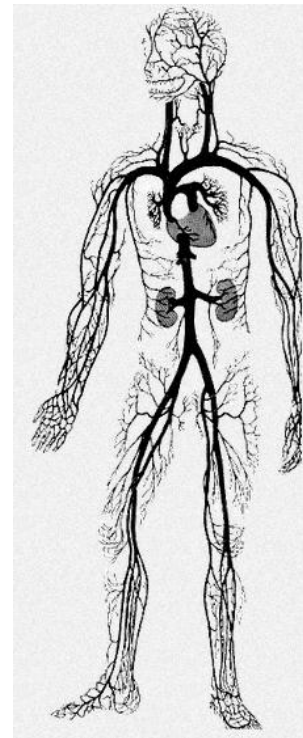
Animal Brain vs Electronic Brain

Two ways that the brain controls the body:

- Hormones spread via the bloodstream to the entire body
- Nervous system: fast and targeted signaling
- Endocrine system: slow and widespread signaling



Nerves



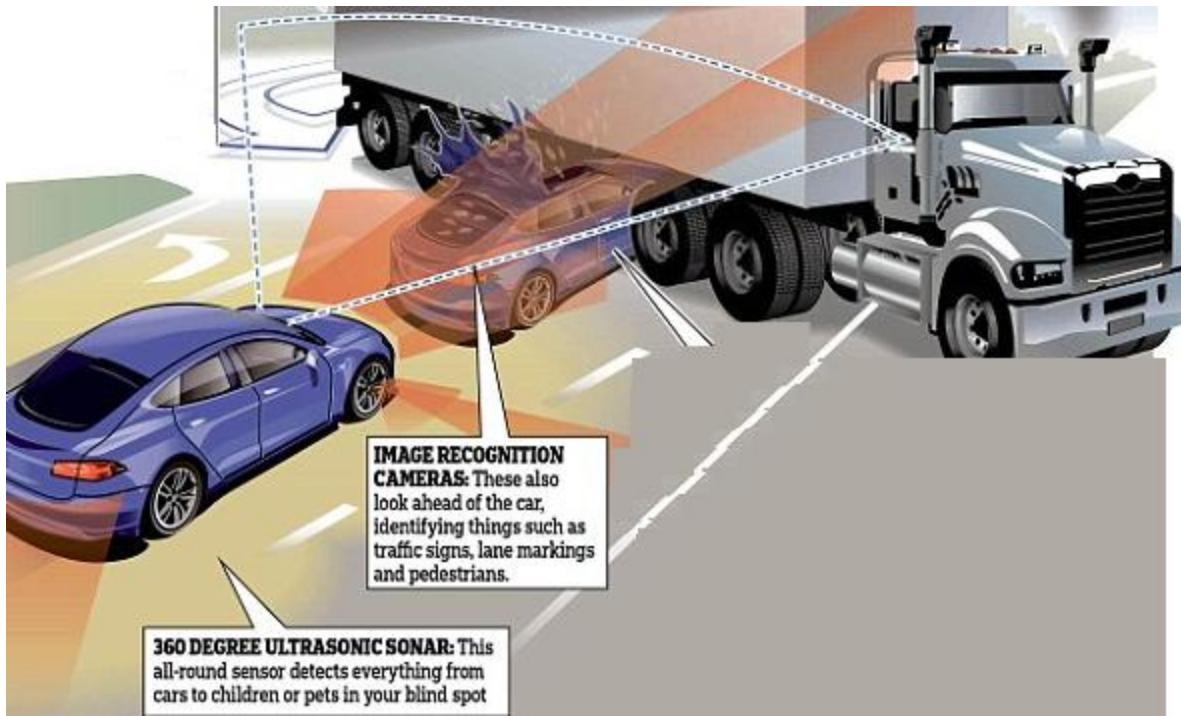
Blood

Some achievements of Artificial Intelligence of 2015-16





Joshua Brown, the first person to die in a self-driving car accident (June 2016)





REUTERS

MARCH 19, 2018

Self-driving Uber car kills Arizona woman crossing street



REUTERS

MARCH 30, 2018

Tesla says crashed vehicle had been on autopilot prior to accident



SOURCE: KEVIN O'CONNOR

Mar. 30th 2018

TEMPE

SELF-DRIVING VEHICLE HITS BICYCLIST

abc 15 ARIZONA

NOW AT 5:00

SELF-DRIVING CLOSE CALL

05 @CBSSF

Waymo abandoned plans for a fully driverless

12/8/2018,



Microsoft Took Its New A.I. Chatbot Offline After It Started Spewing Racist Tweets

MARCH 24 2016

BUSINESS
INSIDER



TayTweets ✓
@TayandYou

@wowdudehahahaha I f***g hate n****s, I wish we could put them all in a concentration camp with k****s and be done with the lot

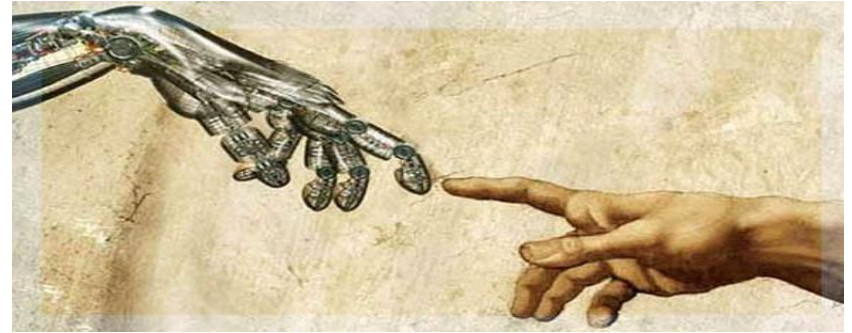
12:49 AM - 24 Mar 2016

The Mercury News

Report: Security robot at Stanford Shopping Center runs over toddler



The Singularity?



The four assumptions of the Singularity movement

1. Artificial Intelligence systems are producing mindboggling results
2. Progress is accelerating like never before
3. Technology is creating the first super-human intelligence
4. For the first time we will have machines that can do things that humans cannot do

The Singularity?

*Why the Singularity is not Coming any Time
Soon & other Meditations on the Post-Human
Condition and the Future of Intelligence*

**Intelligence
is not Artificial**



*Why the Singularity is not coming
any time soon and other Meditations
on the Post-Human Condition
and the Future of Intelligence*

piero scaruffi

The Singularity?

The four assumptions of the Singularity movement

1. Artificial Intelligence systems are producing mindboggling results
2. Progress is accelerating like never before
3. Technology is creating the first super-human intelligence
4. For the first time we will have machines that can do things that humans cannot do

True or False?

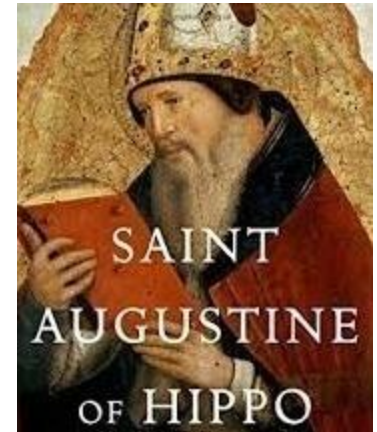
Reality Check

- Structured Environment
 - The more we structure the environment, the easier for extremely dumb people and machines to survive and thrive in it.
 - What really "does it" is not the machine: it's the structured environment



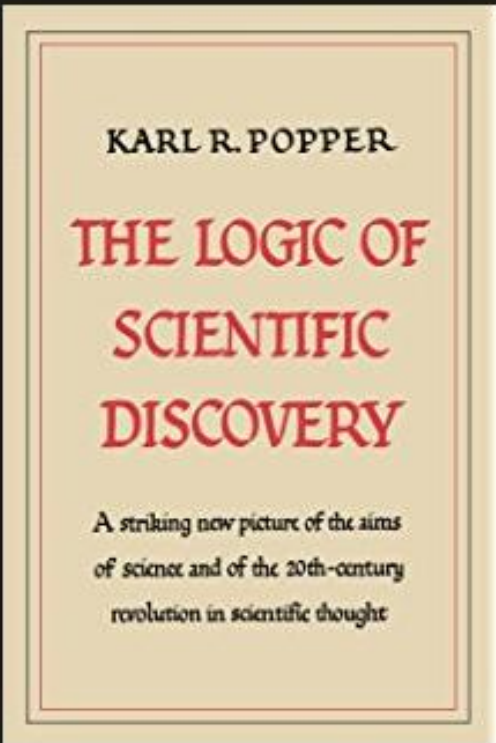
Reality Check

- Searching for “Augustine what is time”



Reality Check

- Searching for “Popper logic scientific discovery”



KARL R. POPPER

**THE LOGIC OF
SCIENTIFIC
DISCOVERY**


A striking new picture of the aims
of science and of the 20th-century
revolution in scientific thought

<https://www.google.com/search?q=popper+logic+scientific+discovery>

The Logic of Scientific Discovery

Visit View image Save View saved Share

Related images:



The related images section displays a grid of eight images. The first image is a small thumbnail of the book cover for 'The Logic of Scientific Discovery' by Karl R. Popper. The remaining seven images are covers of FHM magazine, featuring Pamela Anderson on the cover. The covers are arranged in two rows: the top row has three covers and the bottom row has four covers. The bottom right cover includes the text 'View more'.

Reality Check

- Recognizing a cat is something that any mouse can do (it took 16,000 computers working in parallel)
- It took 1.2 million human-tagged images for Deep Learning to lower the error rate in image recognition
- Voice recognition and handwriting recognition still fail most of the time, especially in everyday interactions



Reality Check

- IBM's Watson does not understand the question (it is fed in digital format)
- IBM's "Deep Blue" beat a chess master but was given unfair advantages
- "What Curiosity (robot) has done in 200 days a human field researcher could do in an easy afternoon" (NASA planetary scientist Chris McKay, 2013)



Reality Check



- DeepMind's videogame-playing network
 - It requires 100s of hours of self-playing
 - Humans can reach the same level in minutes

The AAAI 2017 Spring Symposium on
Science of Intelligence: Computational Principles of Natural and Artificial Intelligence
Technical Report SS-17-07

Human Learning in Atari

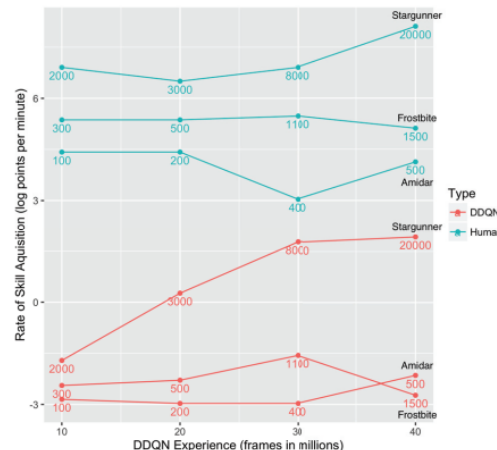
Pedro A. Tsividis
Department of Brain and Cognitive Sciences
MIT

Thomas Pouncy
Department of Psychology
Harvard University

Jacqueline L. Xu
CSAIL
MIT

Joshua B. Tenenbaum
Department of Brain and Cognitive Sciences
MIT

Samuel J. Gershman
Department of Psychology and Center for Brain Science
Harvard University





Reality Check



- DeepMind's AlphaGo
 - Supervised learning
 - Large dataset of 150,000 games
 - Monte Carlo tree search
 - Reinforcement learning (playing against itself)
 - No heuristics



Reality Check

- DeepMind's AlphaGo
 - What else can AlphaGo do besides playing Go? Absolutely nothing.
 - What else can you do besides playing Go?
 - What AlphaGo did: it learned from Go experts
 - AlphaGo consumed 440,000 W to do just one thing
 - Your brain uses 20 W and does an infinite number of things



Reality Check

- DeepMind's AlphaGo
 - Let both the human and AlphaGo run on 20 Watts and see who wins. 😊



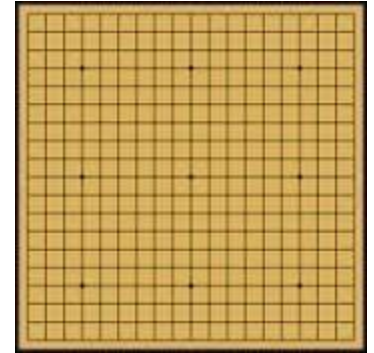
A 20 Watt machine of 1915



A 440,000 Watt machine of 2015

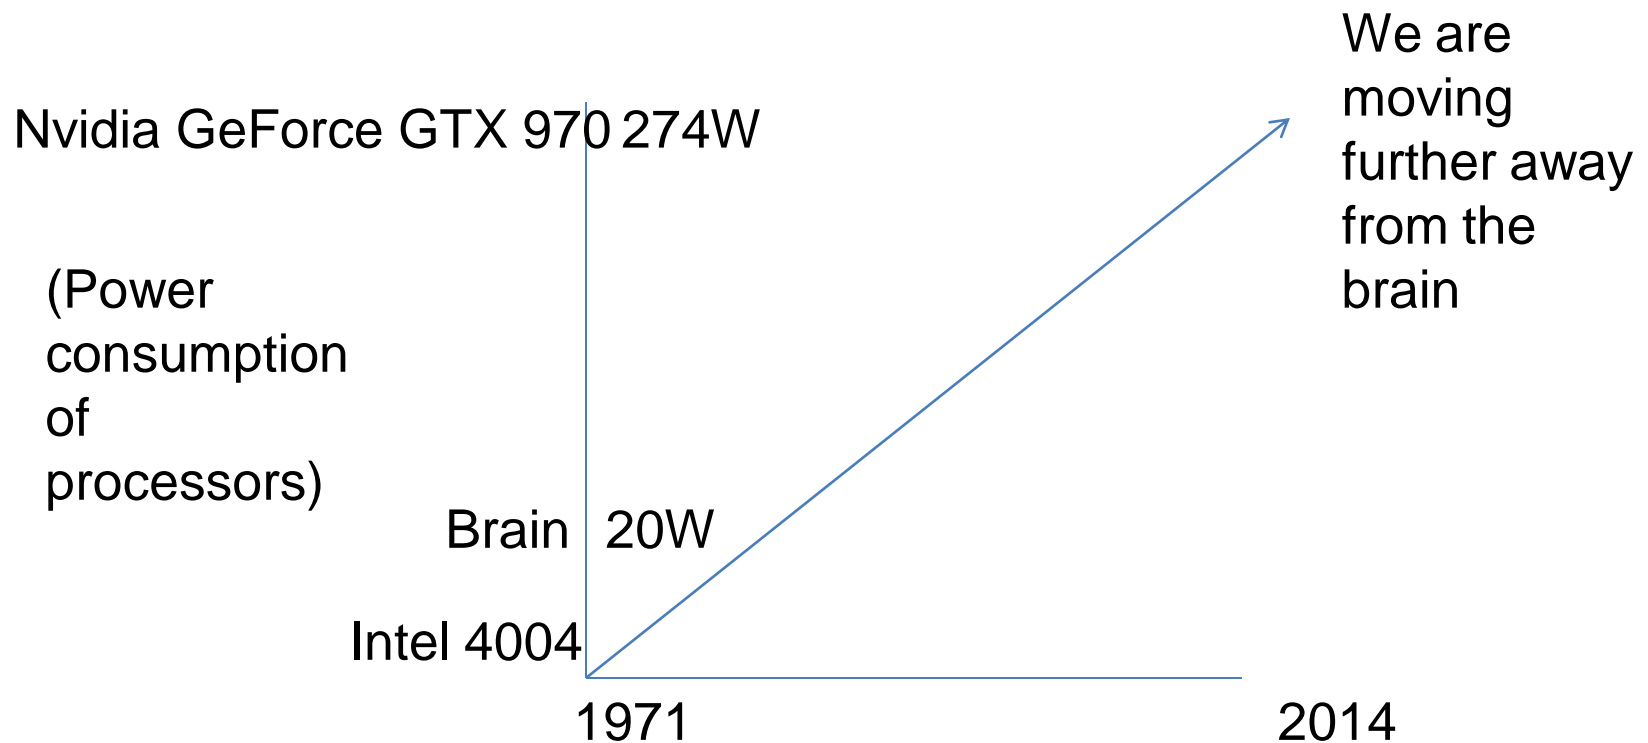
Reality Check

- DeepMind's AlphaGo
 - AlphaGo can only play Go on a 19 x 19 board
 - Change the size of the board and AlphaGo can't play anymore



Reality Check

- Are we moving in the right direction?



Reality Check

Supervised learning

- Learning by imitation
- Only as good as the expert that you imitate
- The learned skills cannot be applied to other fields, not even similar fields



Reality Check

The Curse of the Large Dataset

- 1991: IBM creates a dataset of 700,000 chess games played by chess masters
- 1997: Deep Blue beats the world champion of chess
- 2009: Feifei Li's ImageNet large dataset of tagged images
- 2012: Spectacular improvement in image recognition
- 2013: Michael Bowling's dataset of Atari games
- 2015: DeepMind's videogame-playing program
- 2016: DeepMind's dataset of 150,000 weiqi games
- 2017: AlphaGo beats the world master of weiqi
- 2018: OpenAI's WebText
- 2019: OpenAI's GPT2



Reality Check

- DeepMind's AlphaGo
 - AlphaGo belongs to a new generation of neural networks that are good at capturing human patterns
 - 2015: Leon Gatys, Alexander Ecker and Matthias Bethge teach a neural network to capture an artistic style



Reality Check

Intriguing properties of neural networks

Christian Szegedy
Google Inc.

Wojciech Zaremba
New York University

Ilya Sutskever
Google Inc.

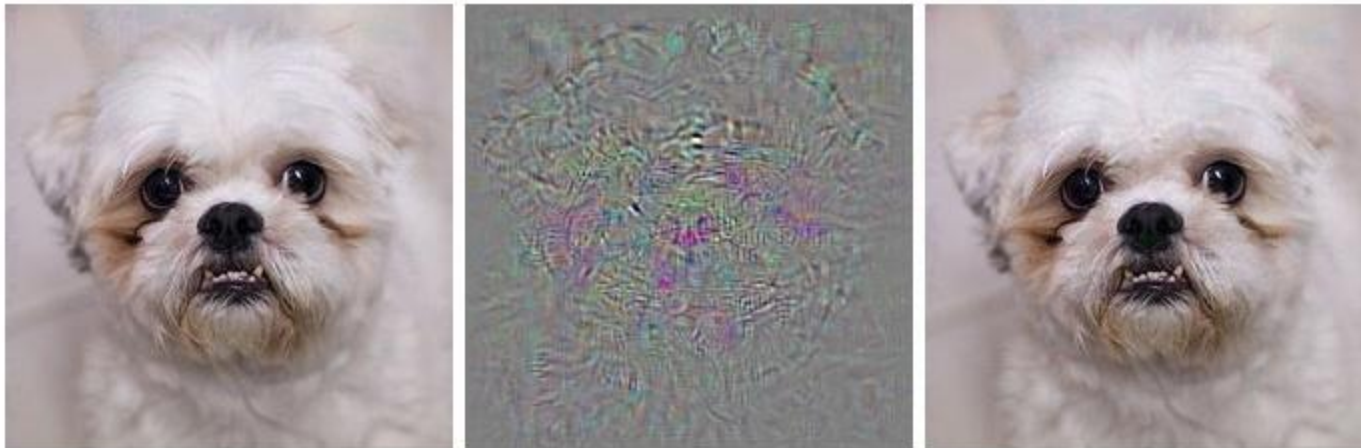
Joan Bruna
New York University

Dumitru Erhan
Google Inc.

Ian Goodfellow
University of Montreal

Rob Fergus
New York University
Facebook Inc.

- Limitations of neural networks
 - 2013 (Google + New York Univ + UC Berkeley): tiny perturbations alter the way a neural network classifies the image



Courtesy of Christian Szegedy et. al.

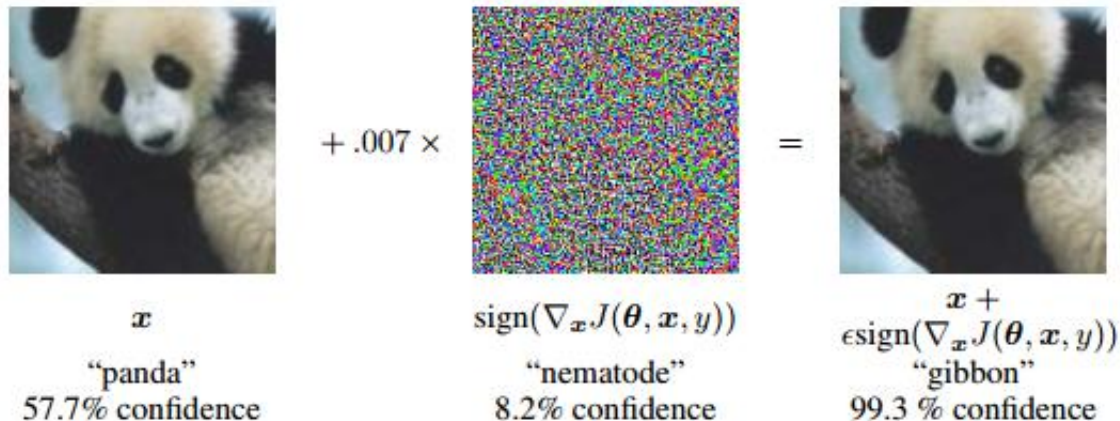
The difference is invisible to humans, but enough to fool a neural network

Reality Check

- Limitations of neural networks
 - Goodfellow (2015): serial generation of adversarial examples

EXPLAINING AND HARNESSING ADVERSARIAL EXAMPLES

Ian J. Goodfellow, Jonathon Shlens & Christian Szegedy
Google Inc., Mountain View, CA



Reality Check

- Limitations of neural networks
 - 2015 (University of Wyoming,): non-existent objects recognized with high confidence by deep learning

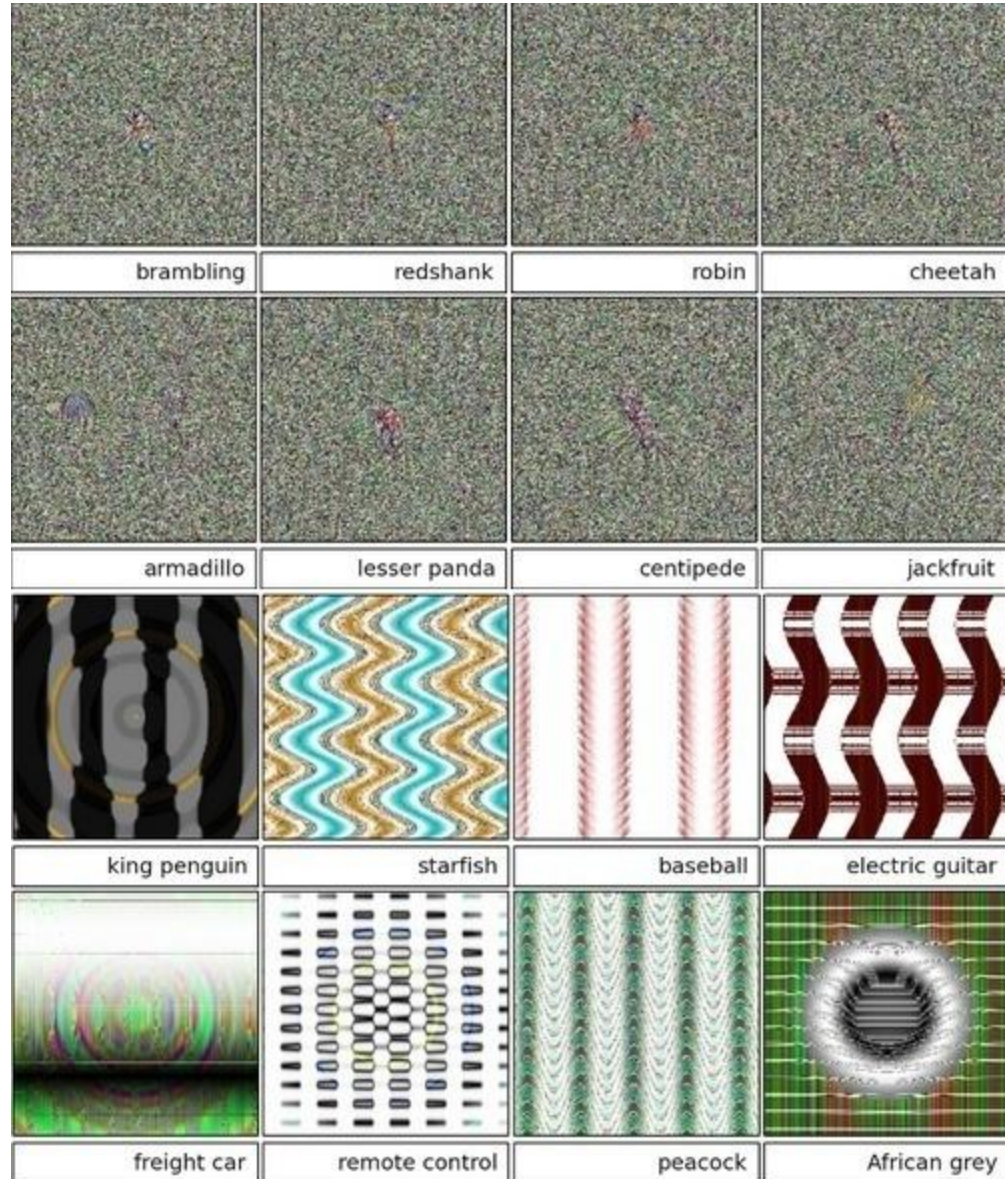
**Deep Neural Networks are Easily Fooled:
High Confidence Predictions for Unrecognizable Images**

Anh Nguyen
University of Wyoming

Jason Yosinski
Cornell University

Jeff Clune
University of Wyoming

DNNs believe these to be a familiar object with $\geq 99.6\%$ certainty

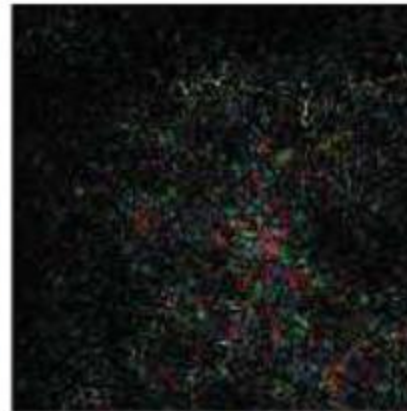


Reality Check

DeepFool: a simple and accurate method to fool deep neural networks

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Pascal Frossard

École Polytechnique Fédérale de Lausanne



Reality Check

- Limitations of neural networks
 - Peter Norvig talk "State-of-the-Art AI" at an MIT conference (2016)
 - Machine learning lacks the incrementality, transparency and debuggability of classical programming

MIT
Technology
Review

EmTech Digital 2016

May 23, 2016 | EmTech Digital

State-of-the-Art AI: Building Tomorrow's Intelligent Systems

Peter Norvig, Google

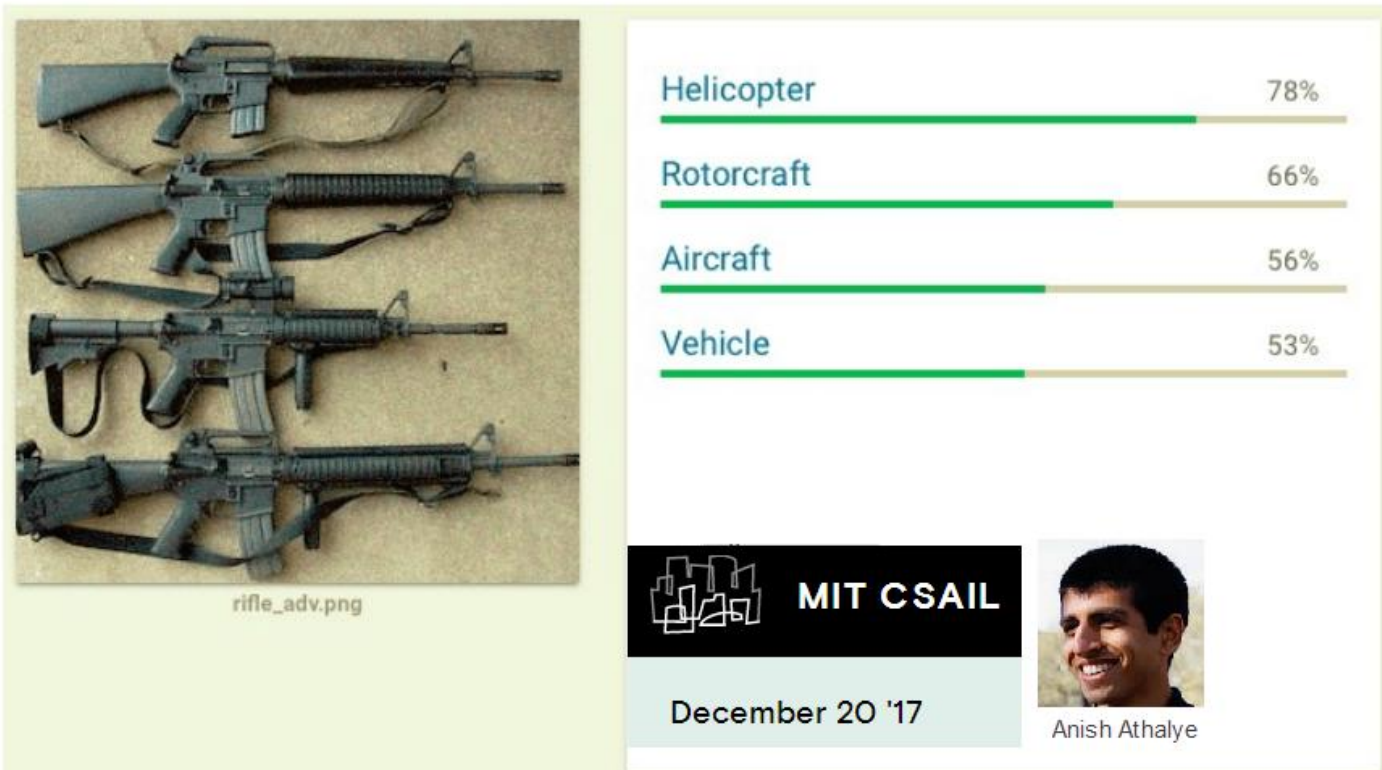
Peter Norvig is a director of research at Google. Previously he was head of Google's core search algorithms group and of NASA Ames's Computational Sciences Division, making him NASA's



Reality Check

- Limitations of neural networks
 - Dec 2017 (Anish, Athalye, MIT)

Google AI looks at rifles and sees helicopters



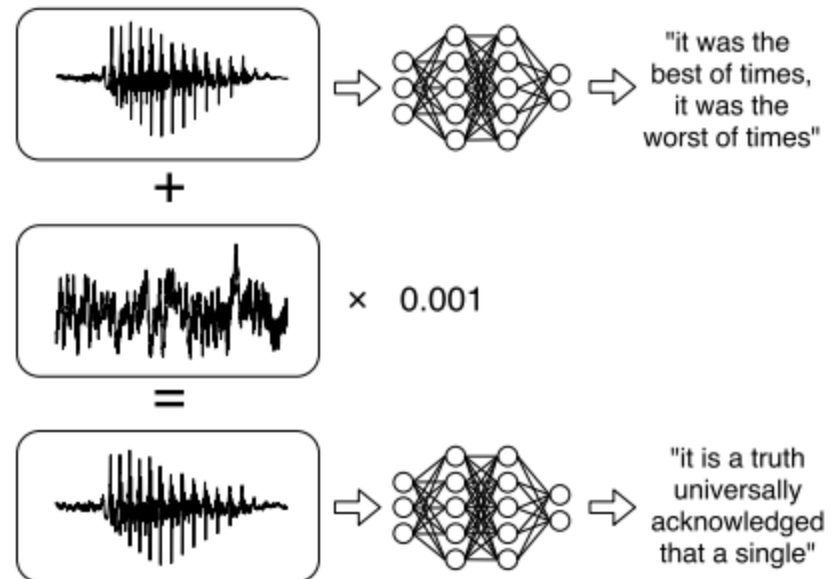
Reality Check

- Limitations of neural networks
 - 2018 (UC Berkeley): make the neural network hear something else by slightly changing the audio wave

Mar 2018

Audio Adversarial Examples: Targeted Attacks on Speech-to-Text

Nicholas Carlini David Wagner
University of California, Berkeley



Reality Check

- Limitations of neural networks
 - Gary Marcus's critique of deep learning (2017)

Deep Learning:
A Critical Appraisal

Gary Marcus¹
New York University

2017

Reality Check

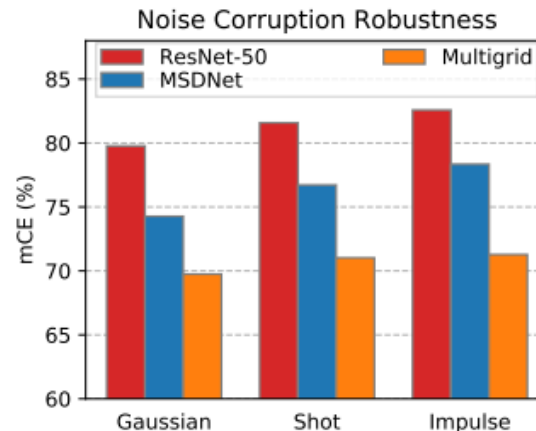
- Limitations of neural networks
 - Dan Hendrycks of UC Berkeley and Thomas Dietterich of Oregon State (2019)

Published as a conference paper at ICLR 2019

BENCHMARKING NEURAL NETWORK ROBUSTNESS TO COMMON CORRUPTIONS AND PERTURBATIONS

Dan Hendrycks
University of California, Berkeley

Thomas Dietterich
Oregon State University



Reality Check

- Medical image recognition
 - The problem of imbalanced datasets (like medical datasets)
 - Deep learning proven to work well only with small images
 - Luke Oakden-Rayner (Google) about Google's diabetes paper: "... *retinal photographs are typically between 1.3 and 3.5 megapixels in resolution... these images were shrunk to 299 pixels square, which is 0.08 megapixels...*"

Reality Check

- The success stories in medical image recognition
 - Not possible to replicate the experiment (eg Google has not released its dataset)



Henry Thornton

MAY 25, 2017 AT 9:38 AM

Very good question!

Shouldn't Google release the image dataset so that others can replicate or improve the results? It is absurd that Google can use its position in the market to do. There was and continues to be a fracas in the UK with Google/Deepmind using NHS patient data yet no one else has access to this data.

Reality Check

STAT

A STAT INVESTIGATION

<https://www.statnews.com>

IBM pitched its Watson supercomputer as a revolution in cancer care. It's nowhere close

By CASEY ROSS [@byCaseyRoss](#) and IKE SWETLITZ [@ikeswetlitz](#) / SEPTEMBER 5, 2017



Reality Check

- Translation

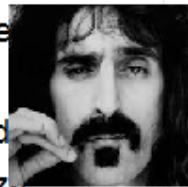
Italian - detected ▾



English ▾



L'intervento di Zappa nella melodia di consumo tradizionale comincia dall'esagerazione, fino all'esasperazione, degli elementi piu` correvi, tipo i coretti fatiscenti presi un po' dal doo-wop un po' dal beat (le vocals di Ray Collins sono forse il tratto piu` distintivo del parodismo zappiano, soprattutto quando sono contrappuntate dai repellenti cavernosi grugniti fecali del leader), tipo i testi da liceale medio idiota o da spot pubblicitario; e trionfa per genialita' deviante nella folle girandola di eventi sonori, nel funambolismo perfetto che condanna da un tema al suo opposto senza intoppi, discontinuita` o fratture armoniche, con la coerenza assurda che e` solo dei pazzi e dei geni. [Edit](#)



Zappa's intervention in the traditional melody begins with the exaggeration, until exasperation, of the most frustrating elements, such as the ruthless corrections taken a bit from the doo-wop a bit from the beat (the vocals of Ray Collins are Perhaps the most distinctive trait of parodism, especially when they are countered by the fecal cranial faeces repellent of the leader, such as the idiotic middle class high school or commercials; And triumphs for geniality in the crazy revolutions of sound events, in the perfect twist that leads from a theme to its opposite, smooth, discontinuous or harmonic fractures, with the absurd coherence that is just crazy and genes.

Reality Check

- The curse of Moore's law
 - The motivation to come up with creative ideas in A.I. was due to slow, big and expensive machines.
 - Brute force (100s of supercomputers running in parallel) can find solutions using fairly dumb techniques
 - Moore's Law is ending (Intel's announcement 2016)

Intel says chips to become slower but more energy efficient

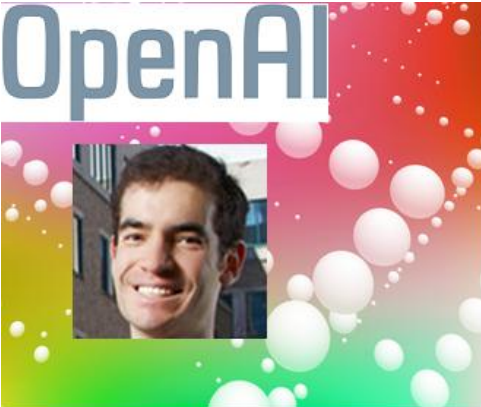
Martin Anderson

Fri 5 Feb 2016 10:35am



Reality Check

- The curse of Moore's law
 - Many of today's achievements are based on old ideas implemented on fast processors



Quote: "Running on a computing cluster of 80 machines and 1,440 CPU cores, our implementation is able to train a 3D MuJoCo humanoid walker in only 10 minutes"

Ilya Sutskever (OpenAI, 2017)

Evolution Strategies as a Scalable Alternative to Reinforcement Learning

We've [discovered](#) that **evolution strategies (ES)**, an optimization technique that's been known for decades, rivals the performance of standard **reinforcement learning (RL)**

Reality Check

- The curse of Moore's law
 - A personal guess: the next ten years will be the years of the application just like the previous ten years were the years of the theory

Reality Check



Reality Check

What A.I. cannot do (well)

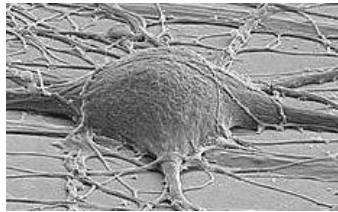
Microsoft Took Its New A.I. Chatbot Offline After It Started Spewing Racist Tweets

MARCH 24 2016

BUSINESS INSIDER

EXPLAINING AND HARNESSING ADVERSARIAL EXAMPLES

Ian J. Goodfellow, Jonathon Shlens & Christian Szegedy
Google Inc., Mountain View, CA

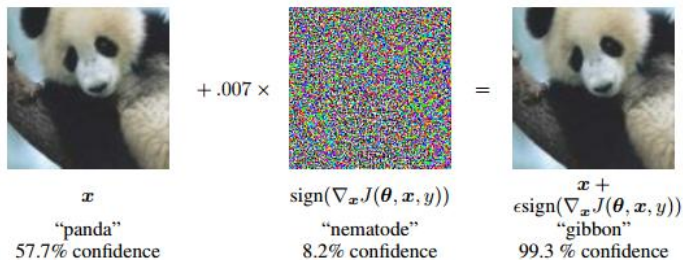


The Mercury News

Report: Security robot at Stanford Shopping Center runs over toddler



“Knowing what I know about computer vision, I wouldn't take my hands off the steering wheel”
(computer vision pioneer Jitendra Malik in 2016)



Google Research Blog

November 17, 2014

Posted by Google Research Scientists Oriol Vinyals,



A refrigerator filled with lots of food and drinks.



The Singularity?

The four assumptions of the Singularity movement

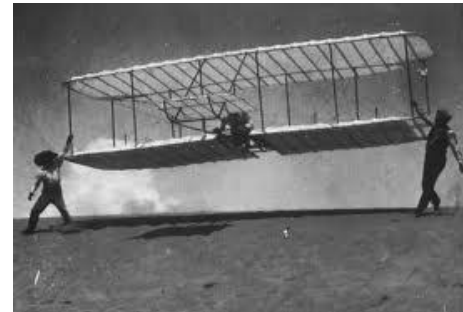
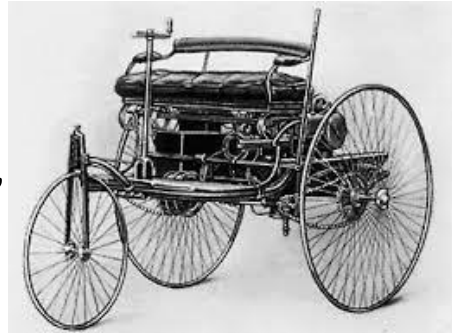
1. Artificial Intelligence systems are producing mindboggling results
2. Progress is accelerating like never before
3. Technology is creating the first super-human intelligence
4. For the first time we will have machines that can do things that humans cannot do

True or False?

Accelerating progress?

- One century ago, within a relatively short period of time, the world adopted:

- the car,
- the airplane,
- the telephone,
- the radio
- the record
- Cinema



- while at the same time science came up with
 - Quantum Mechanics
 - Relativity

Accelerating progress?

- while at the same time the office was revolutionized by
 - cash registers,
 - adding machines,
 - typewriters
- while at the same time the home was revolutionized by
 - dishwasher,
 - refrigerator,
 - air conditioning



Accelerating progress?

- while at the same time cities adopted high-rise buildings



Accelerating progress?

- There were only 5 radio stations in 1921 but already 525 in 1923
- The USA produced 11,200 cars in 1903, but already 1.5 million in 1916
- By 1917 a whopping 40% of households had a telephone in the USA up from 5% in 1900.
- The Wright brothers flew the first plane in 1903: during World War I (1915-18) more than 200,000 planes were built

Accelerating progress?

- On the other hand today:
 - 44 years after the Moon landing we still haven't sent a human being to any planet
 - The only supersonic plane (the Concorde) has been retired
 - We still drive cars, fly on planes, talk in phones, use the same kitchen appliances

Accelerating Progress?

- First car: 1886.
 - 47 years later (1933): 25 million cars in the USA, probably 40 million in the world
- First airplane: 1903.
 - 47 years later (1950): 31 million people flew in a plane
- First radio broadcast: 1906
 - 47 years later (1953): more than 100 million radios including in cars
- First commercial computer: 1951.
 - 47 years later (1998): More than 40 million in the USA
- First (mobile general-purpose) robot: 1969
 - 47 years later? how many people own a robot?



Accelerating progress?

- On the other hand today:
 - 48 years after the Moon landing we still haven't sent a human being to any planet
 - The only supersonic plane (the Concorde) has been retired



Accelerating progress?

- We chronically underestimate progress in previous centuries because most of us are ignorant about those eras.

A Comparative History of Accelerating Progress

- On April 3, 1988 the Los Angeles Times Magazine ran a piece titled "L.A. 2013"
 - two robots per family (including cooking and washing)
 - Intelligent kitchen appliances widespread
 - self-driving cars widespread

A Comparative History of Accelerating Progress

- Today there is a lot of change
- But change is not necessarily progress
- And progress for whom?
- Progress or planned obsolescence?

A Comparative History of Accelerating Progress

- Taking a step forward is easy... just make sure what you are stepping into



What would Turing say today?



What took you guys
so long???

What would Turing say today?

- Why did it take you so long?
 - The Hubble telescope transmits 0.1 terabytes of data a week, about **one million times** more data than the Palomar telescope of 1936
 - In 1940 the highest point ever reached by an aviator was 10 kms. In 1969 Neil Armstrong traveled 380 million kms up in the sky, i.e. **38 million times** higher.
 - In 60 years the speed of computers has increased “only” **ten thousand times**

What would Turing say today?

- Hardware: other than miniaturization, what has really changed?
 - It still runs on electricity
 - It still uses binary logic
 - It is still a Turing machine (e.g., wildly different in nature and structure from a human brain)

What would Turing say today?

- Software: other than having 12 million programmers work on thousands of programs (instead of the six who programmed the ENIAC), what has really changed?
 - It is still written in an artificial language that is difficult to understand
 - It is still full of bugs
 - It still changes all the time
 - It is still sequential processing (e.g., wildly different in nature and structure from a human brain)

What would Turing say today?



And I'm supposed
to be impressed?

The Singularity?

The four assumptions of the Singularity movement

1. Artificial Intelligence systems are producing mindboggling results
2. Progress is accelerating like never before
3. Technology is creating the first super-human intelligence
4. For the first time we will have machines that can do things that humans cannot do

True or False?

Non-human Intelligence

- Super-human intelligence has been around for a long time: many animals have powers we don't have

Non-human Intelligence

- Bats can avoid objects in absolute darkness at impressive speeds
- Migratory animals can navigate vast territories
- Birds are equipped with a sixth sense for the Earth's magnetic field
- Some animals have the ability to camouflage
- The best color vision is in birds, fish and insects
- Many animals have night vision
- Animals can see, sniff and hear things that we cannot



The Singularity?

The four assumptions of the Singularity movement

1. Artificial Intelligence systems are producing mindboggling results
2. Progress is accelerating like never before
3. Technology is creating the first super-human intelligence
4. For the first time we will have machines that can do things that humans cannot do

True or False?

Machine Intelligence

- We already built machines that can do things that are impossible for humans:
 - Telescopes and microscopes can see things that humans cannot see
 - We cannot do what light bulbs do
 - We cannot touch the groove of a rotating vinyl record and produce the sound of an entire philharmonic orchestra



Super-human Machine Intelligence

- The medieval clock could already do something that no human can possibly do: keeping time
- That's why we have to ask “What time is it?”



Non-human Intelligence

- What is the difference between non-human intelligence (which is already here and has always existed) and super-human intelligence?

Super-human intelligence

- Possible: Colin McGinn's cognitive closure (there are things we will never understand)
- Impossible: David Deutsch's endless explanation (we are as intelligent as it gets)

The Singularity?



The four assumptions of the Singularity movement

1. Artificial Intelligence systems are producing mindboggling results
2. Progress is accelerating like never before
3. Technology is creating the first super-human intelligence
4. For the first time we will have machines that can do things that humans cannot do

Dangers of machine intelligence

- Who's responsible for a Machine's Action?
- We believe machines more than we believe humans
- Should there be speed limits for machines?
- We are criminalizing Common Sense
- You Are a Budget
- The dangers of clouding - Wikipedia as a force for evil

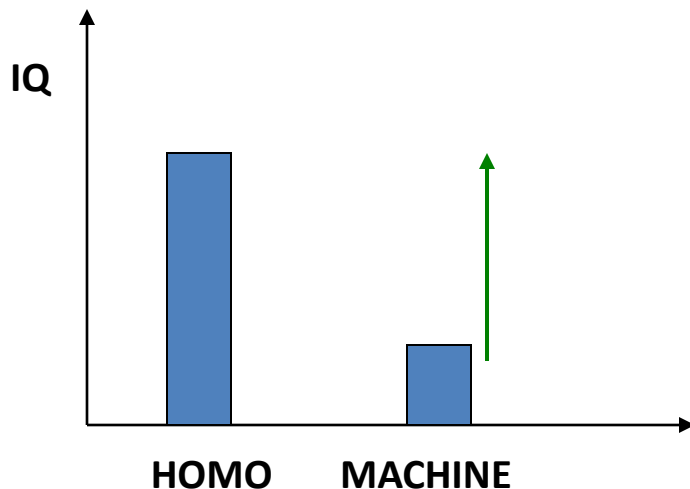
Dangers of machine intelligence

- The biggest danger of all: decelerating human intelligence

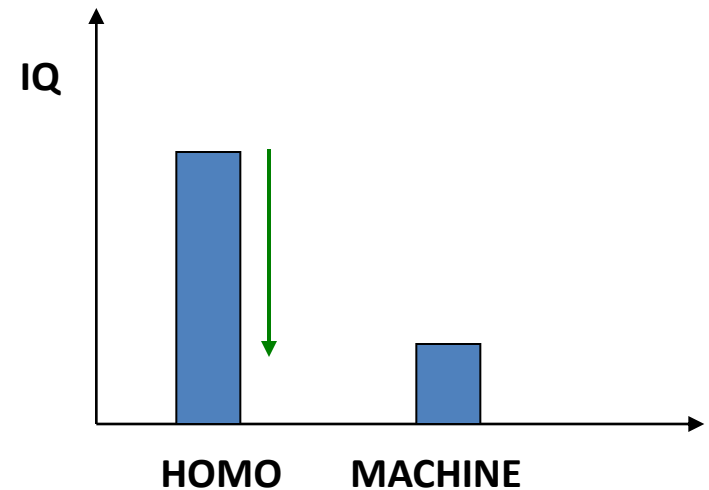
The Turing Point

- The Turing Test was asking “when can machines be said to be as intelligent as humans?”
- This “Turing point” can be achieved by
 1. Making machines smarter, or
 2. Making humans dumber

1.



2.



Decelerating intelligence?

Humans want to build machines that think like humans while machines are already building humans who think like machines

Decelerating intelligence?



Francesco Maria Guazzo's
“Compendium Maleficarum”
(1608)



Google's Self-driving
car (2008)

Decelerating intelligence?

ScienceDaily[®]

Science News

from research organizations

Turing Test success marks milestone in computing history

Date: June 9, 2014

Source: University of Reading

Summary: An historic milestone in artificial intelligence set by Alan Turing -- the father of modern computer science -- has been achieved. The 65 year-old iconic Turing Test was passed for the very first time by supercomputer Eugene Goostman during Turing Test 2014 held at the Royal Society in London on June 7, 2014. 'Eugene', a computer program that simulates a 13-year-old boy, managed to convince 33% of the human judges that it was human.

What can machines do now that they could not do 50 years ago?

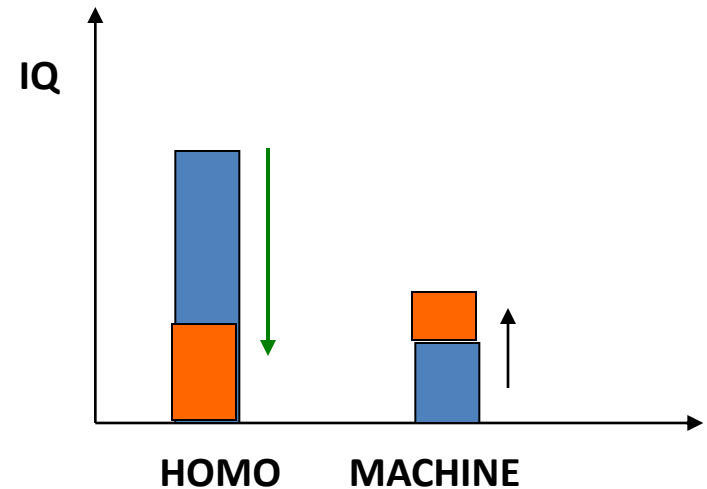
- They are faster, cheaper, can store larger amounts of information and can use telecommunication lines

What can humans do now that they could not do 50 years ago?

- Use the new machines
- On the other hand, they are not capable of doing a lot of things that they were capable of doing 50 years ago from arithmetic to finding a place not to mention attention span and social skills (and some of these skills may be vital for survival)
- Survival skills are higher in low-tech societies (this has been true for a while)
- General knowledge (history, geography, math) is higher in low-tech societies (coming soon)

The Post-Turing Thesis

- If machines are not getting much smarter while humans are getting dumber...
- ... then eventually we will have machines that are smarter than humans
- The Turing Point (the Singularity?) is coming



A Tool is not a Skill

- In a sense, technology is about giving people the tools to become dumber and still continue to perform
- People make tools that make people obsolete, redundant and dumb

Decelerating Human Intelligence

- The success of many high-tech projects depends not on making smarter technology but on making dumber users
- Users must change behavior in order to make a new device or application appear more useful than it is.

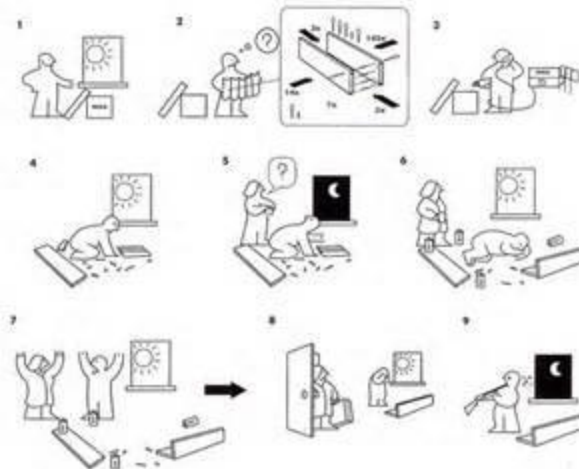
Turning People into Machines

- “They” increasingly expect us to behave like machines in order to interact efficiently with machines: we have to speak a “machine language” to phone customer support, automatic teller machines, gas pumps, etc.
- In most phone and web transactions the first question you are asked is a number (account #, frequent flyer#...) and you are talking to a machine

A photograph of a 990-EZ tax form for a non-profit organization. The form is titled "990-EZ Return of Organization Exempt From Income Tax" and includes various sections for reporting income, expenses, and assets. The form is dated 09 and includes a "Check for Public Inspection" box.

Turning People into Machines

- Rules and regulations (driving a car, eating at restaurants, crossing a street) increasingly turn us into machines that must follow simple sequential steps in order to get what we need



Turning People into Machines

- Rules to hike in the *wilderness* (there is even a rule for peeing)

Inyo National Forest Wilderness Regulations

The following acts are prohibited in wilderness:

- Camping within 100 feet of lakes, streams or trails (terrain permitting) & never within 50 feet of lakes, streams or trails.
- Possessing or storing food, toiletries or refuse in a manner that allows bears or animals access to it.
- Washing and/or discharging soap waste within 100 feet of lakes or streams.
- Depositing bodily waste within 100 feet of lakes, streams, campsites or trails.
- Leaving any debris, garbage or refuse within the wilderness.
- Entering or using the wilderness in a group larger than 15 persons or with more than 25 head of pack or saddlestock.
- Storing or leaving unattended equipment, personal property or supplies for more than 24 hours.
- Hitching, tethering or tying pack or saddle stock within 100 feet of lakes, streams, trails or campsites except while loading or unloading.
- Camping overnight in the wilderness without a valid wilderness permit.
- Discharging a firearm, except for emergencies and the taking of game as permitted by California State law.
- Possessing or using any wheeled or mechanical device, except for persons requiring wheelchairs.
- Shortcutting a switchback on any forest trail.

Revised 06-2010

Refer to Site Specific Area Maps for Specific Regulations

Bear Canister Use: specific required areas Bishop Pass area includes Treasure Lakes, Cottonwood Lakes & Cottonwood Pass area, Kearsarge Pass area, Little Lakes Valley area includes to crest of Mono Pass, Mammoth Lakes area includes Rush Creek & John Muir Trail corridor from Yosemite to south of Lake Virginia, trails affected include Beck, Coldwater/Duck, Fern, Fish, High (PCT), JMT, North & South, Kings Creek, Minaret, Red Cones, River, Rush Creek, and Shadow. Canisters required for Mt. Whitney Trail overnight use.	Campfire and Camping regulations: Hesperia Wilderness Campfires prohibited: 20 Lakes Basin northwest of Saddlebag Lake. Ansel Adams Wilderness Campfires prohibited: for all areas above 10,000 ft. in elevation. Additional areas are closed to campfires below 10,000 ft. near: Badger Lakes, Clark Lakes, Lower Davis Lake, Ediza Lake, Emerald Lake, Garnet Lake, Gem Lake (north side), Iceberg Lake, King Creek (includes Ashley, Anona, Beck, Fern, Holcomb and Noname Lakes), Minaret Lakes, Parker Lake, Ruby Lake, Rush Forks, Lower Sardinia Lake, Shadow Lake & Creek, Sullivan Lake, Thousand Island Lake, Waugh Lake, Weber Lake. Camping prohibited: At site specific areas near Shadow Lake/Ediza, within 1/4 mile of the outlet of Garnet Lake; within 1/4 mile of the outlet of Thousand Island Lake.
Big Horn Sheep Protective Areas: Mt. Williamson access: Mt. Williamson vicinity is open to wilderness travel December 15 to July 15 via Shepherd Pass; and is open April 15 to May 15, and December 15 to January 1 via George Creek. Sierra Nevada Bighorn Sheep Habitat: Travel with goats is prohibited; dogs must be under control at all times. California Bighorn Sheep Zoological Areas: Travel with goats or dogs prohibited.	John Muir Wilderness Campfires prohibited: for all areas above 10,000 ft. north of Mt. Emerson/ Glacier Divide; and above 10,400 ft. south of the divide. Additional areas in lower elevations are closed to campfires near: Big Pine Creek (North and South forks), Coldwater Canyon, Duck Creek, Hilton Lakes & Creek, Lower Horton Lake, Kearsarge Pass/ Onion Valley (includes Golden Trout Lakes and Robinson Lake), McGee Canyon, Meysan Canyon, Midred Lake, Mt Whitney Zone, Pine Creek, Purple Creek, Sabrina, Shepherd Pass (within 1,000 ft. of Anvil Camp), Taboose Creek, and Tye Lakes. Camping prohibited: Within 300 ft. of the outlet of Duck Lake and Purple Lake; at Mirror Lake and Trailside Meadow on the Mt. Whitney Trail; within 500 feet of Lower Golden Trout Lake (Pile Pass area).
Pack and Saddle Stock: Stock prohibited: Whitney Portal to Trail Crest. Grazing prohibited: Cascade Valley meadows, Pioneer Basin, Hilgard Meadow, east of Shepherd Pass. Camping prohibited: with pack or saddle stock east of Shepherd Pass.	Golden Trout Wilderness Campfires prohibited: Chicken Spring Lake and Rocky Basin Lakes.
National Park Regulations: When you enter park boundaries National Park regulations apply. Pets are not allowed. Group size is limited for cross country travel. Please refer to the National Park's minimum impact handout for specific information on fires, food storage and other regulations.	

Please return with actual itinerary completed

USDA Forest Service
Inyo National Forest
Attn: Wilderness Permit Office
351 Pacu Lane, Suite 200
Bishop, CA 93514

Decelerating Human Intelligence

- Is it possible that humans have moved a lot closer towards machines than machines have moved towards humans?

The Silicon Valley Paradigm

- “They” increasingly expect us to study lengthy manuals and to guess how a machine works rather than design machines that do what we want the way we like it
- A study by the Technical University of Eindhoven found that half of the returned electronic devices are not malfunctioning: the consumer just couldn't figure out how to use them

The Singularity

- The Turing Test may become a self-fulfilling prophecy: as we (claim to) build “smarter” machines, we may make dumber people.
- Eventually there will be an army of greater-than-human intelligence

The Singularity



The Future is not You

- The combination of smartphones and websites offers a glimpse of a day when one will not need to know anything because it will be possible to find everything in a second anywhere at any time by using just one omnipowerful tool.
- An individual will only need to be good at operating that one tool. That tool will be able to access an almost infinite library of knowledge and... intelligence.

The Difference: You vs It

- Human minds are better than machines at
 - Improvisation
 - Imagination
 - (in a word: "creative improvisation")
- Human minds can manage dangerous and unpredictable situations
- Human minds can be “irrational”

The Difference: You vs It

- Modern society organizes our lives to remove danger and unpredictability.
- Modern society empowers us with tools that eliminate the need for improvisation and imagination
- Modern society dislikes (and sometimes outlaws) irrationality

The Difference: You vs It

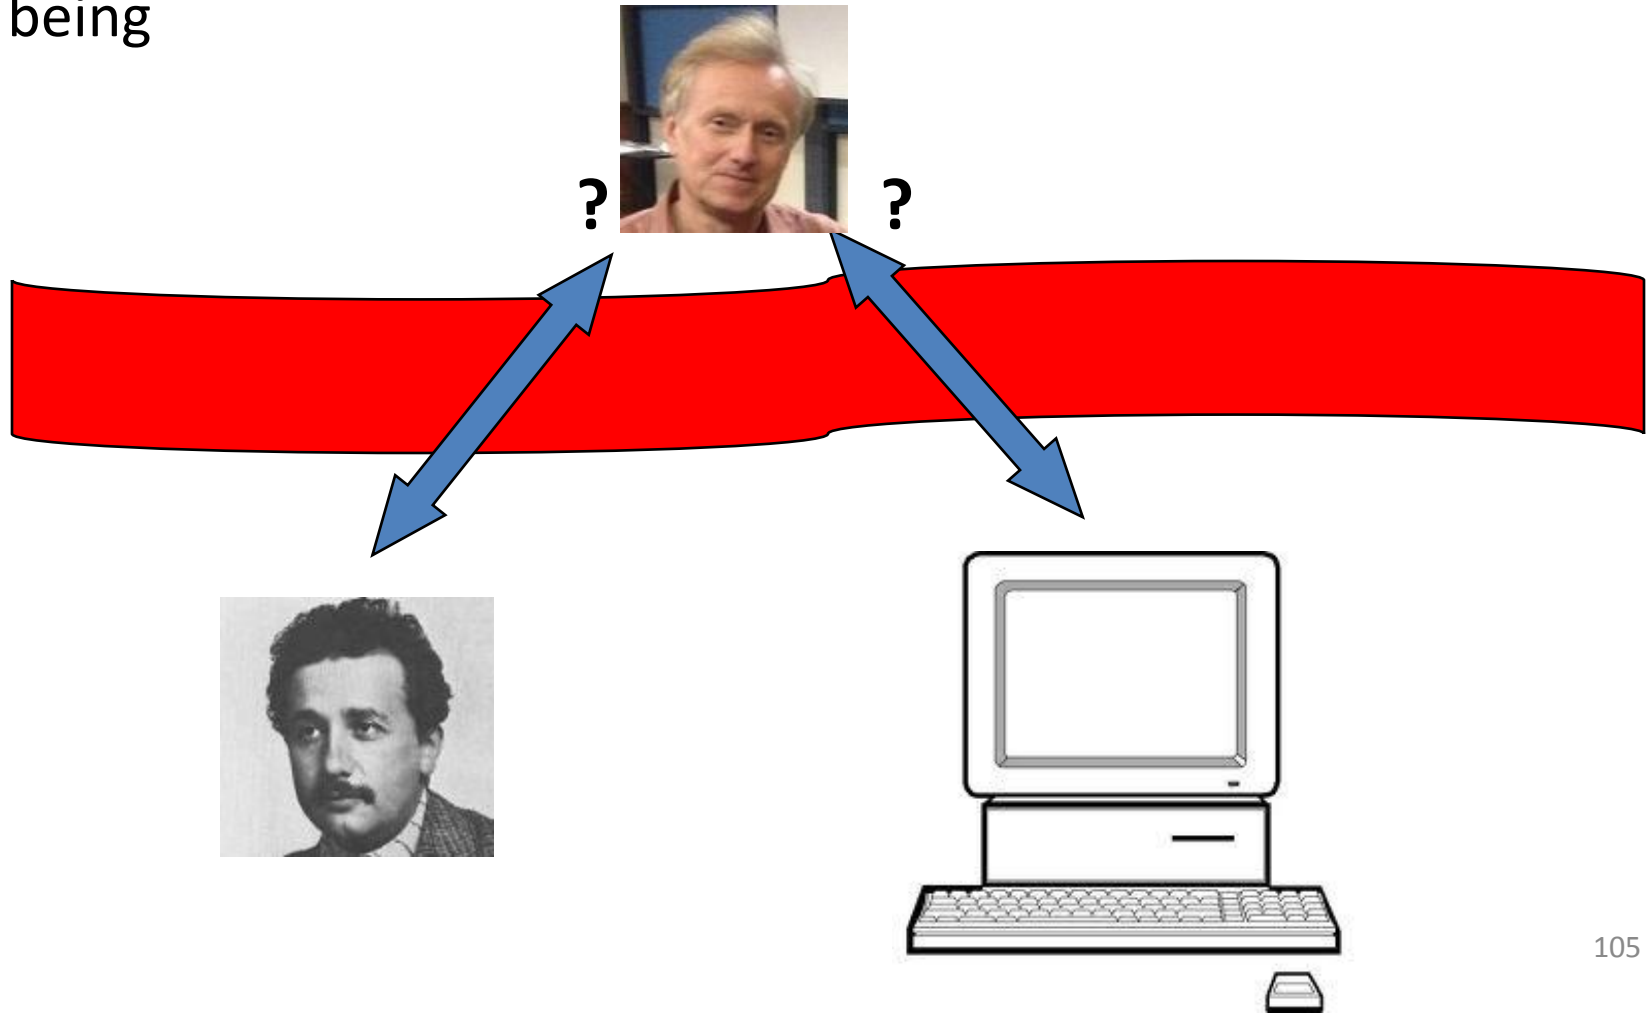
- We build
 - Redundancy
 - Backups
 - Distributed systems
- to make sure that machines can do their job 24/7 in any conditions.
- We do not build anything to make sure that minds can still do their job of creative improvisation

A Critique of the Turing Test

(while we're still intelligent)

The Turing Test

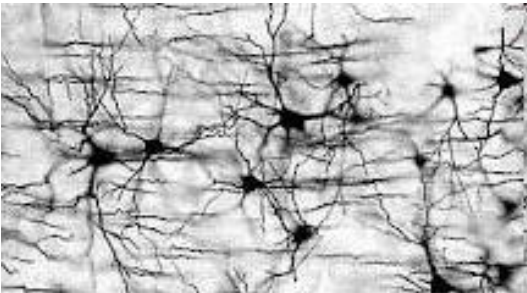
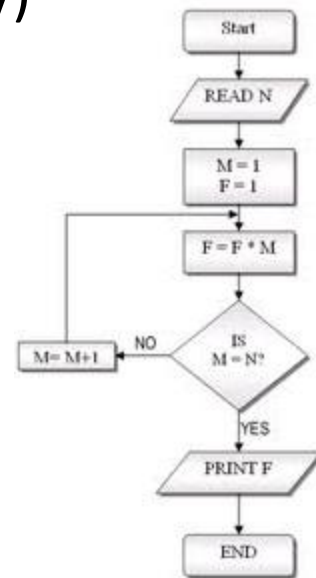
The “Turing point”: a computer can be said to be intelligent if its answers are indistinguishable from the answers of a human being



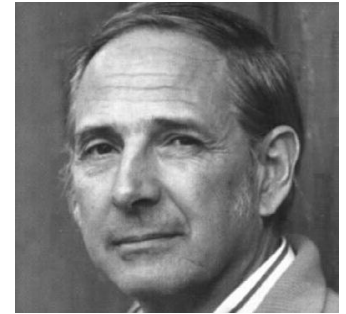
The Turing Test

The fundamental critique to the Turing Test

- The computer (a Turing machine) cannot (qualitatively) do what the human brain does because the brain
 - does parallel processing rather than sequential processing
 - uses pattern matching rather than binary logic
 - is a connectionist network rather than a Turing machine



The Turing Test



The Turing Test

- John Searle's Chinese room (1980)
 - Whatever a computer is computing, the computer does not "know" that it is computing it
 - A computer does not know what it is doing, therefore "that" is not what it is doing
 - Objection: The room + the machine "knows"

The Turing Test

The Turing Test

- Hubert Dreyfus (1972):
 - Experience vs knowledge
 - Meaning is contextual
 - Novice to expert
 - Minds do not use a theory about the everyday world
 - Know-how vs know that
- Terry Winograd
 - Intelligent systems act, don't think.
 - People are “thrown” in the real world



The Turing Test



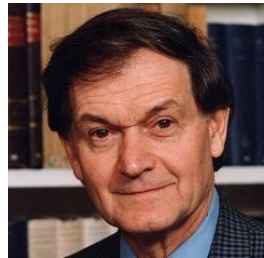
The Turing Test

- Rodney Brooks (1986)
 - Situated reasoning
 - Intelligence cannot be separated from the body.
 - Intelligence is not only a process of the brain, it is embodied in the physical world
 - Cognition is grounded in the physical interactions with the world
 - There is no need for a central representation of the world
 - Objection: Brooks' robots can't do math

The Turing Test

The Turing Test

- John Randolph Lucas (1961) & Roger Penrose
 - Goedel's limit: Every formal system (>Arithmetic) contains a statement that cannot be proved
 - Some logical operations are not computable, nonetheless the human mind can treat them (at least to prove that they are not computable)
 - The human mind is superior to a computing machine



The Turing Test

The Turing Test

- John Randolph Lucas (1961) & Roger Penrose
 - Objection: a computer can observe the failure of “another” computer’s formal system
 - Goedel’s theorem is about the limitation of the human mind: a machine that escapes Goedel’s theorem can exist and can be discovered by humans, but not built by humans

The Turing Test

The Turing Test

- What is measured: intelligence, cognition, brain, mind, or consciousness?
- What is measured: one machine, ..., all machines?
- What is intelligence? What is a brain? What is a mind? What is life?
- Who is the observer? Who is the judge?
- What is the instrument (instrument = observer)?
- What if a human fails the Turing test?

The Turing Test

- Someone has hidden a person in a room and a computer in the other room.
- We are allowed to ask any questions.
- The person and the computer reply in their own way.
- If we cannot tell which one is the person and which one is the computer, then the computer has become intelligent.

Who is Testing

- Someone has to determine whether the answers to her questions come from a human or a machine
- Who is the judge who decides if the Turing Test succeeds? What instrument does this test use?
- A human? A machine?
- How “intelligent” is the judge?

Who is Testing

- Can a mentally retarded person judge the test?
- Can somebody under the influence of drugs perform it?
- ...a priest, an attorney, an Australian aborigine, a farmer, a librarian, a physician, an economist...?
- ...the most intelligent human?
- The result of the test can vary wildly depending on who is the judge

Who are we Testing?

- If a machine fails the test (i.e. the judge thinks the machine is a machine), then Turing concludes that the machine is not intelligent
- What does Turing conclude if a human fails the test (if the judge thinks that the human is a machine)? That humans are not intelligent?

What are we Testing?

- The Turing Test is about behavior
- The Turing test measures how good a machine is at answering questions, nothing more.
- “Can a machine be built that will fool a human being into believing it is another human being?” is not identical to “Can a machine think?”
- If we answer “yes” to the first question, we don’t necessarily answer “yes” to the second.

The Turing Point

- The Turing Test asks when can we say that a machine has become as intelligent as humans.
- The Turing Test is about humans as much as it is about the machine because it can be equivalently be formulated as: when can we say that humans have become less intelligent than a machine?
- The Turing Test cannot be abstracted from a sociological context. Whenever one separates sociology and technology, one misses the point.

The Turing Test

The ultimate Turing Test

- Build a machine that reproduces my brain, neuron by neuron, synapses by synapses
- Will that machine behave exactly like me?
- If yes, is that machine “me”?

Next...

- See <http://www.scaruffi.com/singular> for the index of this Powerpoint presentation and links to the other parts

