



Intelligence is not Artificial

Piero Scaruffi

p@scaruffi.com

www.scaruffi.com

(2 hour version of August 2017)



Olivetti Artificial Intelligence Center circa 1987

20+ Books

A Brief History of Knowledge

from 3000 BC to 2001 AD



piero scaruffi

Synthesis

Essays, Photographs, Poems



piero scaruffi

A History of Rock and Dance Music

From the Guitar to the Laptop
From Chicago to Shanghai

Volume 1 (1951-1989)



piero scaruffi

Piero Scaruffi

A History of Silicon Valley 1900-2016

Almost a third edition/ 2016 Update

Thinking about Thought



the structure of life and the
meaning of matter

BRAIN

Volume 1
in the "Thinking about Thought" series

piero scaruffi

Piero in China





World Intelligence Congress (Tianjin)

June 2017





Leonardo Art/Science Evening Rendezvous



The LAST Festival


Life Art Science Technology festival

www.lastfestival.org



DON'T BUY MY BOOK

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



But download these slides from
www.scaruffi.com

piero scaruffi



[Home](#) | [Travel](#) | [Hiking](#) | [Politics](#) | [History](#) | [Literature](#) | [Art](#) | [Tech](#) | [Philosophy](#)
| [Events](#)

or

Donate



Latest

Lecture: [Xerox PARK talk on August 3 \(download the slides\)](#)
Book review: ["Answers for Aristotle - How Science and Philosophy Can Lead Us to A More Meaningful Life"](#)
Politics: [Why China will be the new moral and political leader of the world](#)
Presentation: [Talk on Artificial Intelligence in China \[45-minute video\]](#)

usic **Books**

Table of Contents

- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough



Oh no please not another history of Artificial Intelligence



1934: 17 years before the first commercial computer

A. M. Turing (1950) Computing Machinery and Intelligence. *Mind* 49: 433-460.

COMPUTING MACHINERY AND INTELLIGENCE




By A. M. Turing




1. The Imitation Game "Turing Test"

I propose to consider the question, "Can machines think?"

1950s: the first computers





The first book on electronic computers (1949)

New N.P.L. Wonder

ELECTRIC BRAIN TO BE MADE AT TEDDINGTON

34 YEARS-OLD DESIGNER TALKS TO SURREY COMET

£100,000 A.C.E. WILL BE OBSOLETE BEFORE COMPLETED

SOME of the feats that will be able to be performed by Britain's new electronic brain, which is being developed at the N.P.L., Teddington, were described to the SURREY COMET yesterday by Dr. A. M. Turing, 34 years-old mathematics expert, who is pioneer of the scheme in this country.

The machine is to be an improvement on the American ENIAC, and it was in the brain of Dr. Turing that the more efficient model was developed.

Does the UNIVAC 120 really think?

"thinking" for Canadian business . . .
the **UNIVAC 120**

Does the UNIVAC 120 really think?

among other things, the most complicated payroll, production and billing problems.

Does Univac 120 eliminate intelligence? Certainly not. But it does free your skilled personnel for creative thinking. Wherever you are for clerical chores that

Many executives who use Univac's 120 agree the above machine

Behind the Giant Brains

By FRANK LEARY

The IBM 650 magnetic drum data processing machine built for commercial use. This machine is designed to meet the accounting and computing needs in areas between those now served by the company's very large and its smaller machines.

Artificial Intelligence



AI Magazine Volume 27 Number 4 (2006) (© AAAI)


A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence

August 31, 1955

*John McCarthy, Marvin L. Minsky,
Nathaniel Rochester,
and Claude E. Shannon*


Artificial Intelligence?

- Definitions!
 - When does Computer Science become Artificial Intelligence?
 - When does automation become AI?
 - When does technology become AI?
 - What is the difference between an algorithm and an AI algorithm?
 - What is “intelligence”?
 - What is “artificial”?



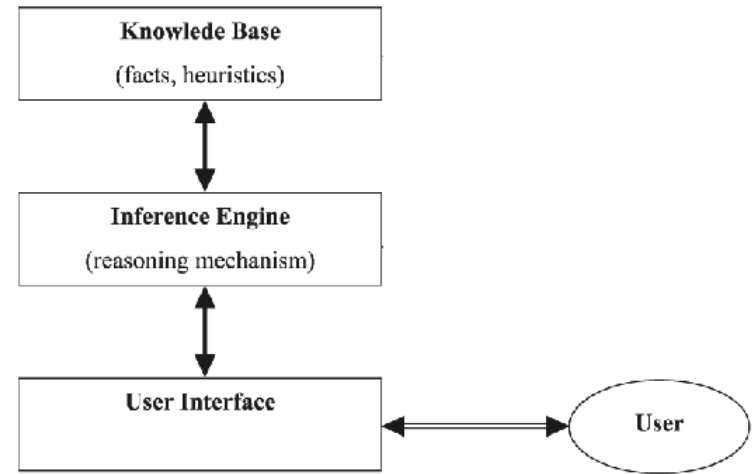


The two schools of A.I.

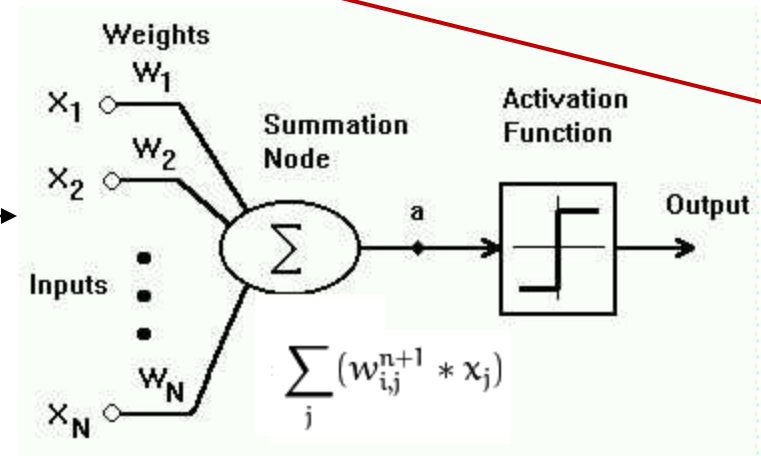




Artificial Intelligence (1956)

- Knowledge-based approach uses mathematical logic to simulate the human mind




- Neural-net approach simulates the structure of the brain





The two schools of AI

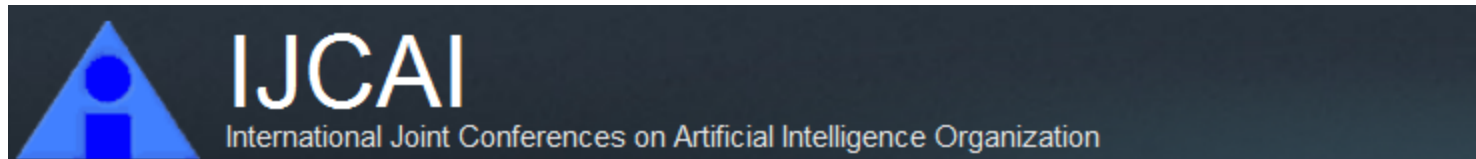


1956: Allen Newell and Herbert Simon's "Logic Theorist"
1959: John McCarthy's "Programs with Common Sense"
1965: Ed Feigenbaum's Dendral
1965: Lofti Zadeh's Fuzzy Logic
1966: Ross Quillian's Semantic Networks
1969: SRI's Shakey the Robot
1969: Roger Schank's Conceptual Dependency Theory
1972: Bruce Buchanan's MYCIN
1972: Terry Winograd's SHRDLU
1974: Marvin Minsky's Frame



The first IJCAI: 1969

Only 2 papers on neural networks!



Proceedings of the First International Joint
Conference on Artificial Intelligence

IJCAI-69 Contents


INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE

May 7-9, 1969 Washington, D. C

The Neural Network Winter


1969: Marvin Minsky & Samuel
Papert's "Perceptrons" kills neural
networks

1971: Noam Chomsky's article against
Skinner's behaviorism



Why nobody responded?

- Pitts died in May 1969
- McCulloch died in Sep. 1969
- Rosenblatt died in 1971



Symbolic A.I. Winters

1957: Herbert Simon declares that *"there are now in the world machines that think, that learn, and that create"*

1958: A New York Times article (8 July 1958) reporting a press conference by Rosenblatt that *"the Perceptron is the embryo of an electronic computer that (the US Navy) expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence"*

1958: Bar-Hillel publishes a "proof" that machine translation is impossible

1965: Herbert Simon predicts that *"machines will be capable within 20 years of doing any work a man can do"*

1966: The ALPAC Report causes reduction in funding for machine translation research




Symbolic A.I. Winters

1970: Marvin Minsky to Life Magazine:

*“In from three to eight years we will have a machine with **the general intelligence of an average human being**”*

1973: The Lighthill Report kills A.I. in the UK

1980s: Fifth Generation illusion



What saved A.I.

Neuroscience: Fukushima's convolutional nets



Physics: Hopfield's recurrent neural networks


Canada: CIFAR (Canadian Institute for Advanced Research)



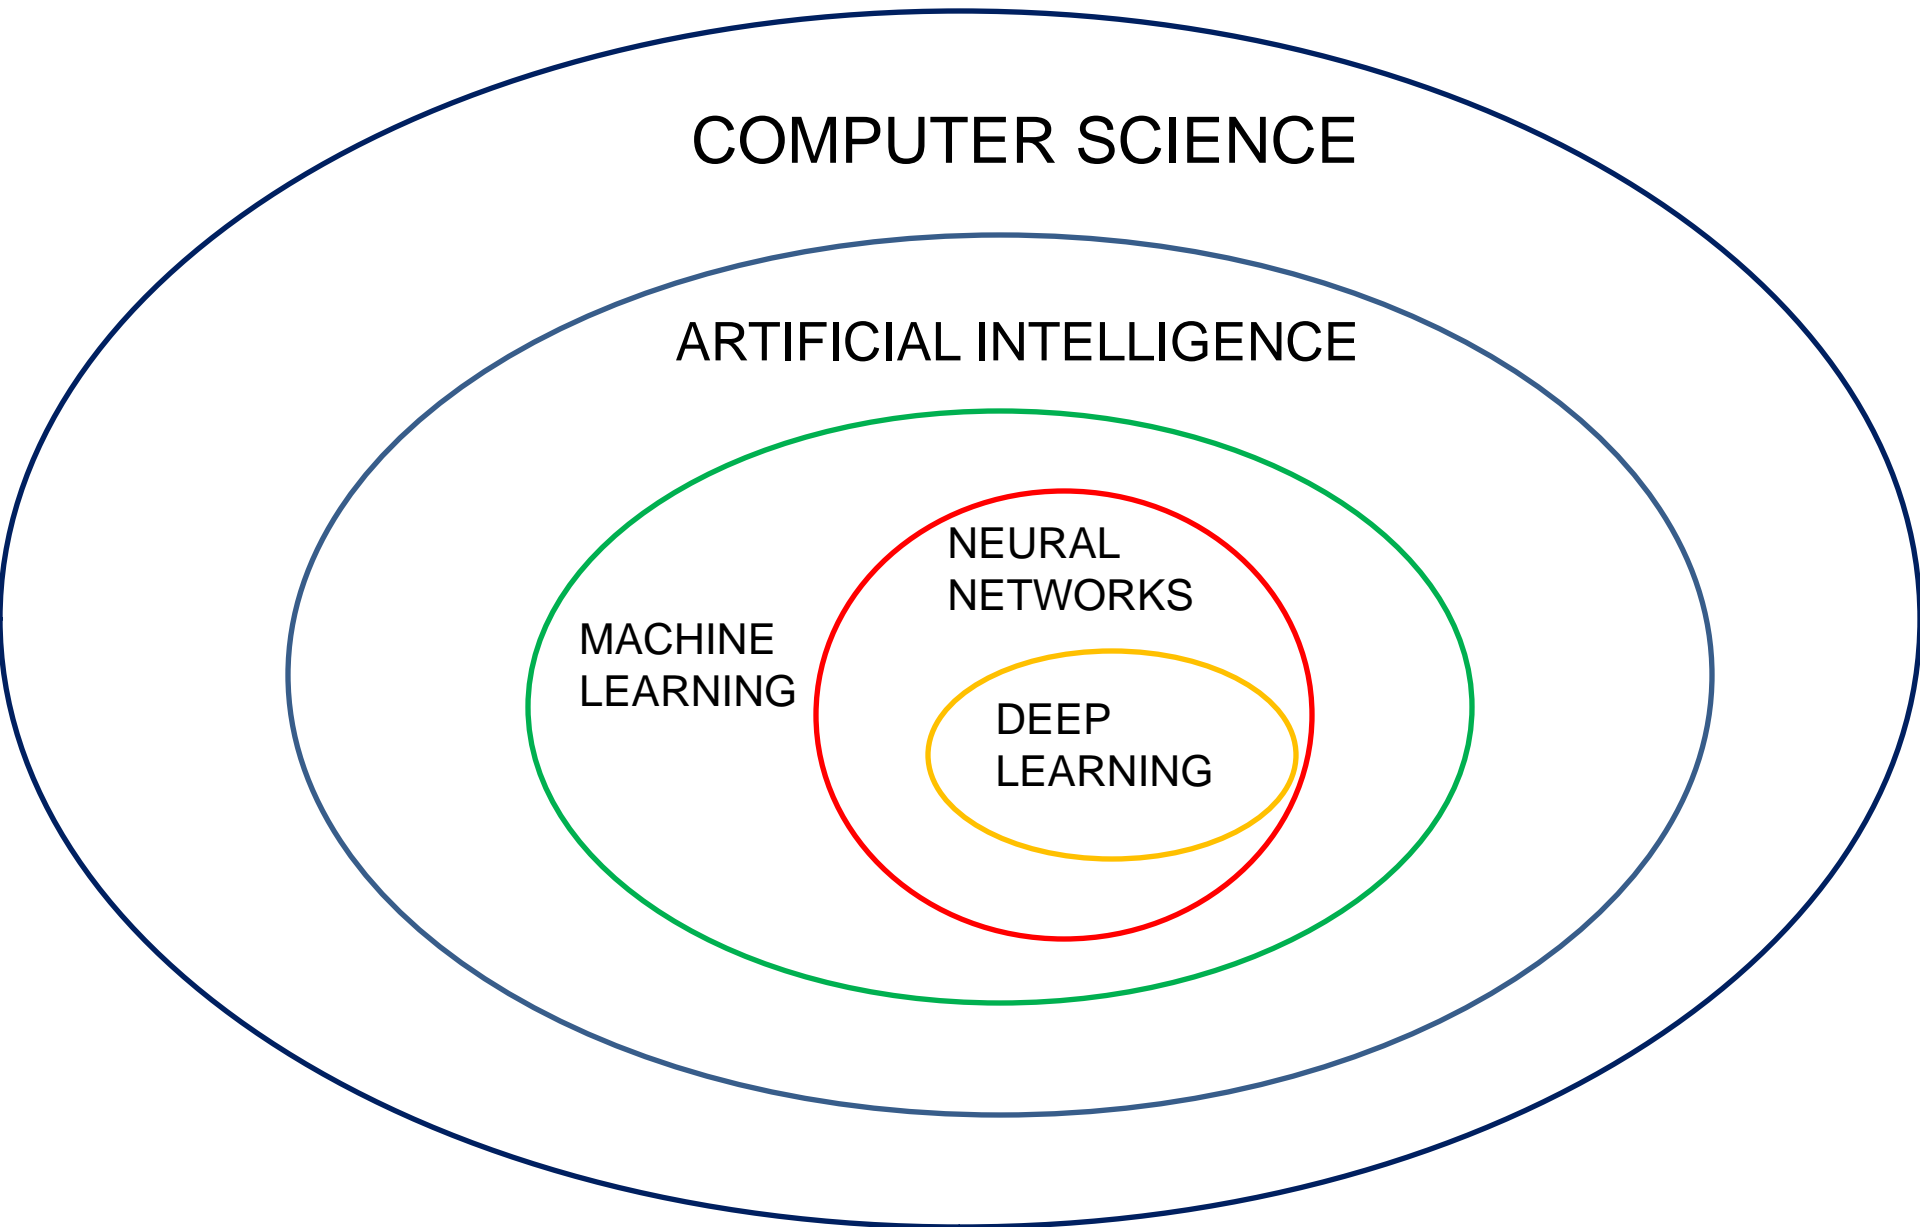
Vannoccio
Biringuccio's "De la
Pirotechnia"
(1540, first printed
book on metallurgy
published in Europe)

The #1 factor: Moore's Law


Microprocessor Transistor Counts 1971-2011 & Moore's Law



Neural Networks




Neural Networks



3 LAYERS

MANY LAYERS
("DEEP")



Neural Networks

1960: The backpropagation algorithm

Bryson 1961, Kelley 60, Dreyfus 62, Linnainmaa 1970, Werbos




1965: Alexey Ivakhnenko publishes the first learning algorithms for multi-layered networks

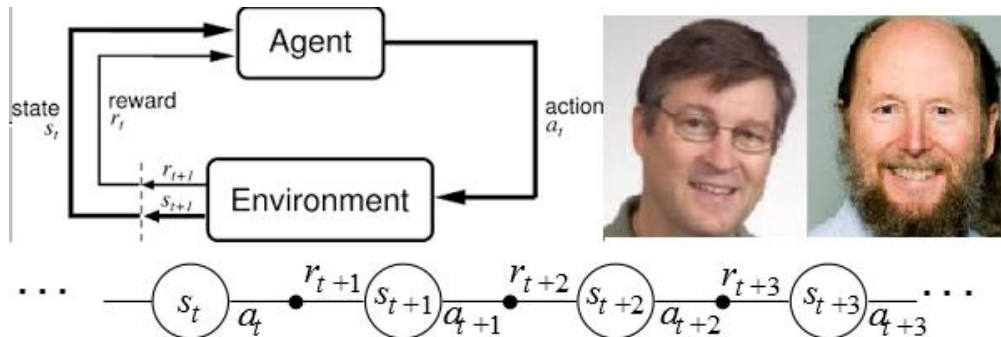


Neural Networks

1979: Kunihiro Fukushima's **convolutional neural network**




1981: Andrew Barto's and Richard Sutton's temporal-difference method of **reinforcement learning**



Neural Networks

1982: John Hopfield's **recurrent neural network**



$$E(v) = -\frac{1}{2} \sum_{i \neq j} \sum_{i,j} w_{ij} v_i v_j - \sum_i I_i v_i + \sum_i \frac{1}{R_i} \int_0^{v_i} f_i^{-1}(z) dz$$

$$\nabla_v E(v) = Wv + I - u/R$$

1983: Terry Sejnowski's and Geoffrey Hinton's **Boltzmann machine**



$$P(x) = \frac{\exp(-E(x))}{Z}$$

> E(x): Energy function

> Z: partition function where $\sum_x P(x) = 1$

1985: Judea Pearl's **"Bayesian Networks"**



$$P(C, S, R, W, F) = P(C) P(S|C) P(R|C) P(W|R, S) P(F|R)$$


$$P(C, F) = \sum_S \sum_R \sum_W P(C, S, R, W, F)$$

$$P(F|C) = P(C, F) / P(C)$$

Neural Networks

1986: Paul Smolensky's **Restricted Boltzmann machine**

Restricted Boltzmann Machines



Boltzmann distribution:

$$P(\mathbf{v} = \mathbf{v}, \mathbf{h} = \mathbf{h}) = \frac{1}{Z} \exp(-E(\mathbf{v}, \mathbf{h}))$$

$$Z(\theta) = \sum_{\mathbf{h}, \mathbf{v}} \exp(-E(\mathbf{v}, \mathbf{h}; \theta))$$

No interaction between
hidden variables



1986: David Rumelhart's **backpropagation algorithm**

equations of backpropagation

$$\delta^L = \nabla_a C \odot \sigma'(z^L)$$

$$\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l)$$


$$\frac{\partial C}{\partial b_j^l} = \delta_j^l$$

$$\frac{\partial C}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l$$



Neural Networks

1989: George Cybenko proves that neural networks can approximate continuous functions



1992: Hava Siegelmann's and Eduardo Sontag's theorem (RNN=Turing Machine)



Neural Networks

1989: Yann LeCun 's convolutional neural network LeNet-1
(backpropagation applied to convolutional networks)




Figure 1: Architecture of LeNet 1

*Handwritten Digit Recognition with a
Back-Propagation Network*


Y. Le Cun, B. Boser, J. S. Denker, D. Henderson,
R. E. Howard, W. Hubbard, and L. D. Jackel
AT&T Bell Laboratories, Holmdel, N. J. 07733

Given functions $x(t)$ and $w(t)$, their convolution is a function $s(t)$

$$s(t) = \int x(a)w(t - a)da$$

Written as

$$s = (x * w) \quad \text{or} \quad s(t) = (x * w)(t)$$



Neural Networks

1994: Yann LeCun 's LeNet-5

PROC. OF THE IEEE, NOVEMBER 1998


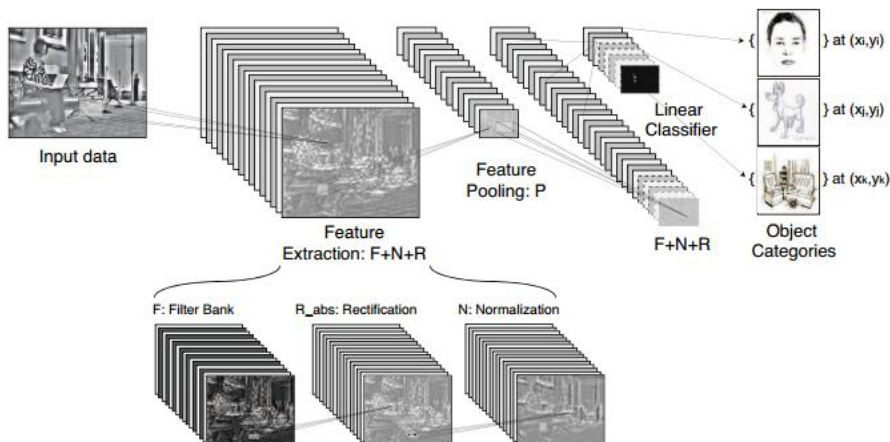


Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition.

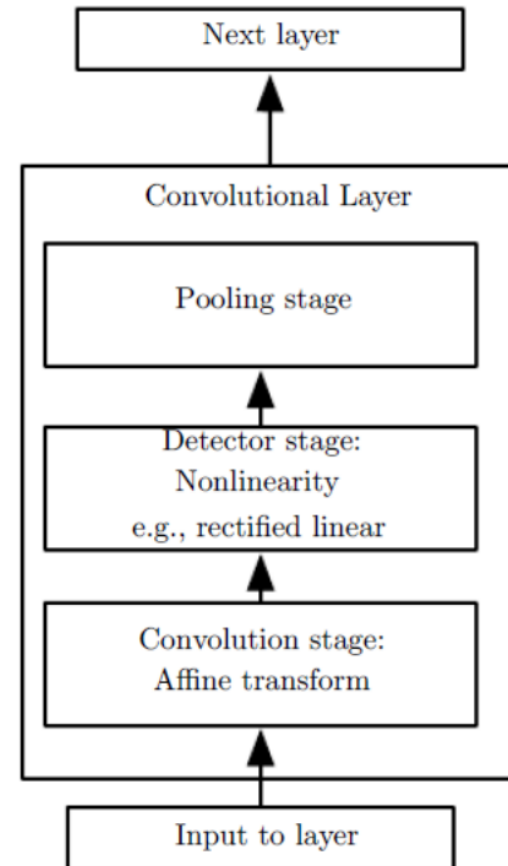

Convolutional Network



Architecture of a typical convolutional network for object recognition

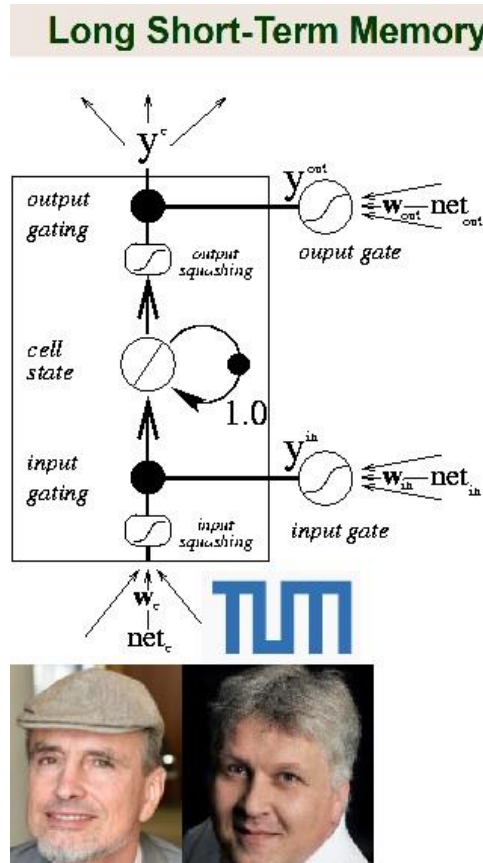
From a paper by Yann LeCun

Figure 2: error rate on the test set (%).




Neural Networks

1997: Sepp Hochreiter's and Jeurgen Schmidhuber's Long Short Term Memory (LSTM) model



No need for Neural Nets

1994: Ernst Dickmanns' self-driving car drives more than 1,000 kms near the airport Charles-de-Gaulle in Paris



1997: IBM's "Deep Blue" chess machine beats the world's chess champion, Garry Kasparov



No need for Neural Nets

2005: Sebastian Thrun's driverless car Stanley wins DARPA's Grand Challenge



2011: IBM's Watson debuts on a tv show





2011: Apple Siri (2011)



No need for Neural Nets

- Statistical machine learning
 - 1990: Robert Schapire's boosting
 - 1991: Vladimir Vapnik's SVM
 - 1993: Ross Quinlan's C4.5
 - 1995: Tin-Kam Ho's random forests




Deep Learning



2006: Geoffrey Hinton's Deep Belief Networks

Deep Belief Network





(Hinton et.al)

$$P(v, h^1, h^2, h^3) = \underbrace{P(v|h^1)}_{\text{Sigmoid Belief Net}} \underbrace{P(h^1|h^2)P(h^2, h^3)}_{\text{RBM}}$$

$$P(v|h^1) = \prod_i P(v_i|h^1)$$

$$P(h^1|h^2) = \prod_j P(h_j^1|h^2)$$

$$P(h^2, h^3) = \frac{1}{Z(W^3)} \exp(h^{2T} W^3 h^3)$$




Deep Learning

2007: Yoshua Bengio's Stacked Autoencoders



Auto-Encoders




- Multilayer neural nets with target output = input.
- Reconstruction = decoder(encoder(input))
- Objective is to minimize the reconstruction error.



$$\mathcal{L}(q) = -D_{KL}(q(z|x) \parallel p_{model}(z)) + \mathbb{E}_{z \sim q(z|x)} \log p_{model}(x|z)$$

<http://www.slideshare.net/KazukiNitta/variational-autoencoder-68705109>

Stacked Auto-encoders



Deep Learning

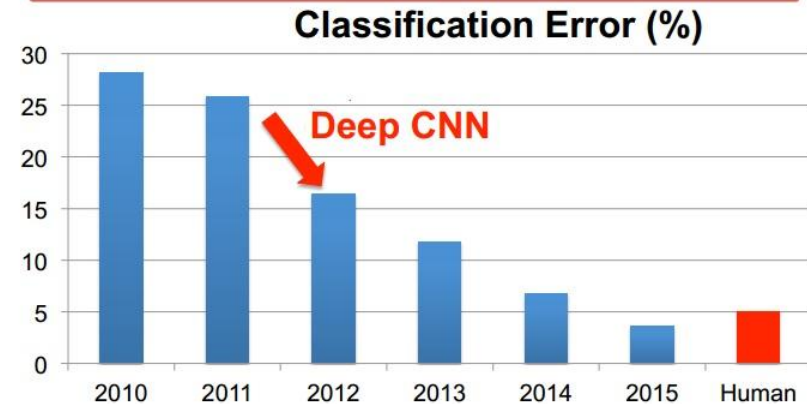
Annus mirabilis of Deep Learning: 2012

Deep
Learning →

Error rates on the ILSVRC-2012 competition

• Krizhevsky et. al.	• 16.4%
much bigger gap	
• University of Tokyo	• 26.1%
• Oxford University Vision Group	• 26.9%
• INRIA + XRCE	• 27.0%
• University of Amsterdam	• 29.5%


ImageNet: Image Classification Task



Deep Learning

AlexNet: 8 layers

Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton




Deep Learning

- Google (2012): GoogleBrain (1.7 billion connections in 16,000 processors) recognize cats in YouTube videos



咱	攢	暫	贊	贗	贗	葬	遭
擇	則	澤	賊	怎	增	憎	曾
詠	摘	壽	宅	窄	債	寨	瞻
湛	綻	樟	章	彰	漳	張	掌
照	罩	兆	肇	召	遮	折	哲
針	傾	枕	疼	診	震	振	鎮
鄭	証	芝	枝	支	吱	知	知
止	趾	只	苴	紙	志	摯	擲




ICDAR OFFLINE CHINESE HANDWRITING RECOGNITION CONTEST

OCT 2013:



NEAR-HUMAN
PERFORMANCE

JÜRGEN SCHMIDHUBER, IDSIA



Deep Learning

- Speech recognition



Microsoft 2016



Skype Translator



ACHIEVING HUMAN PARITY IN CONVERSATIONAL SPEECH RECOGNITION


W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu and G. Zweig

Microsoft Research
Technical Report MSR-TR-2016-71

to JP Morgan
to Tencent

The real heroes of Deep Learning

Nvidia's GPUs



Deep Learning is born



Nvidia Tesla P100 for deep learning

Unsung Heroes: The datasets

1990: Switchboard-1 Telephone Speech Corpus (TI)

1991: Continuous Speech Recognition (CSR) Corpus

1993: FERET (Army Research Lab)

1994: ORL face dataset (Olivetti)

1996: Broadcast News corpus

1999: MNIST handwritten-digit dataset (NYU)

2006: PASCAL VOC

2007: Tiny Images Dataset (MIT)

2007: Labeled Faces in the Wild (University of Massachusetts)

2009: ImageNet

2013: dataset of Atari games (University of Alberta)

2014: COCO (Microsoft)

Deep Learning

Annus mirabilis of Deep Learning: 2012

Why did it take until 2012?

1. The training datasets
2. The computational power

Trivia: The Stars of Deep Learning

Kunihiko Fukushima: Japan

Hava Siegelmann: Israel

Sepp Hochreiter: Germany

Juergen Schmidhuber: Switzerland

Yann LeCun: France

Geoffrey Hinton: Britain/ Canada

Yoshua Bengio: France/ Canada

Andrew Ng: China

Daniela Rus: Romania


Fei-fei Li: China

Sebastian Thrun: Germany

DeepMind: Britain/ New Zealand

Ilya Sutskever: Russia

Quoc Le: Vietnam



Evolution of Neural Networks

Perceptrons

1960s

Reinforcement
Learning

1970s

Recurrent
Neural
Networks

1980s

Convolutional
Neural
Networks

1990s

Deep
Learning

2000s

Generative
Adversarial
Networks

2010s

Deep Learning

The state of the art

- CNN = Convolutional Neural Networks
(images)
- RNN = Recurrent Neural Networks
(speech, translation, scenes)
- LSTM = Long Short Term Memory
(narrative text?)

Business

- Multi-billion dollar investments in artificial intelligence and robotics in the 2010s
 - Amazon (Kiva, 2012; Angel.ai, 2016; Harvest.ai, 2017)
 - Google (Neven, 2006; Industrial Robotics, Meka, Holomni, Bot & Dolly, DNNresearch, Schaft, Bost, DeepMind, Redwood Robotics, 2013-14; API.ai, 2016; Moodstocks, 2016; Kaggle, 2017)
 - IBM (AlchemyAPI, 2015; Watson project)
 - Microsoft (Project Adam, 2014; Swiftkey, 2016; Genee and Maluuba, 2017)
 - Apple (Siri, 2011; Perceptio and VocallQ, 2015; Emotient, Turi and Tuplejump, 2016; RealFace, 2017)
 - Facebook (Face.com, 2012; Wit.ai, 2015; Masquerade, 2017; Zurich Eye, 2017)
 - Yahoo (LookFlow, 2013; Incredible Labs, 2014)
 - Twitter (WhetLab, 2015; Magic Pony, 2016)
 - Salesforce (TempoAI, 2015; MetaMind, 2016; PredictionIO, 2016)
 - Samsung (Viv Labs, 2016)
 - Intel (Nervana and ItseeX in 2016)
 - General Electric (Wise.io, 2017)



Business

AI Startups as of November 2017 (VentureScanner)




Table of Contents

- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough



A.I. in China

- Microsoft chatbot Xiaoce: 100 million users
- News recommendation: Toutiao
- Voice recognition: iFlytek (China)
- Face recognition: Face++ (China)
- Several A.I. unicorns



今日头条

首页 / 视频 / 正文

84

微博

Qzone

微信

为什么中国永远不会有硅谷

原创 即刻video 2016-07-04 21:52

中国永远不会有硅谷出现。Never Ever.

《硅谷百年史》的作者Piero Scaruffi这么告诉即刻君。因为

他还追问了一句，

“你告诉我，在过去的
Google，这是中国版
没有做过的互联网事情

诸位，先别跟即刻君一



是什么
Netflix,

究了硅

A.I. in China

- The B.A.T.



A.I. in China

- A simple definition: “A.I. = computational mathematics”

$$P(x) = \frac{\exp(-E(x))}{Z}$$

$$E(v) = -\frac{1}{2} \sum_{i,j} \sum_{l,j} w_{ij} v_i v_j - \sum_i I_i v_i + \sum_i \frac{1}{R_i} \int_0^{v_i} f_i^{-1}(z) dz$$

$$\nabla_v E(v) = Wv + I - u/R$$

$$P(C, S, R, W, F) = P(C) P(S|C) P(R|C) P(W|R, S) P(F|R)$$

$$P(C, F) = \sum_S \sum_R \sum_W P(C, S, R, W, F)$$

$$P(F|C) = P(C, F) / P(C)$$

$$P(v = v, h = h) = \frac{1}{Z} \exp(-E(v, h))$$

$$Z(\theta) = \sum_{h,v} \exp(-E(v, h; \theta))$$

$$\delta^L = \nabla_a C \odot \sigma'(z^L)$$

$$\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l)$$

$$\frac{\partial C}{\partial b_j^l} = \delta_j^l$$

$$\frac{\partial C}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l$$

$$P(v, h^1, h^2, h^3) =$$

$$P(v|h^1)P(h^1|h^2)P(h^2,h^3)$$

Sigmoid Belief Net RBM

$$P(v|h^1) = \prod_i P(v_i|h^1)$$

$$P(h^1|h^2) = \prod_j P(h_j^1|h^2)$$

$$P(h^2, h^3) = \frac{1}{Z(W^3)} \exp(h^{2T} W^3 h^3)$$

$$\mathcal{L}(q) = -D_{KL}(q(z|x) \parallel p_{model}(z)) +$$


$$\mathbb{E}_{z \sim q(z|x)} \log p_{model}(x|z)$$

A.I. in China

- 2015: China builds the equivalent of nearly one university per week
- 2015: China has more STEM graduates than the USA (78 million vs 67 million)
- 2016: China STEM graduates 4.7 million; USA 568,000


The Countries With The Most STEM Graduates

Recent graduates in Science, Technology, Engineering & Mathematics (2016)



A.I. in China

- 2016: China publishes more papers than the USA on Deep Learning
- 2017: China generates more data than the rest of the world combined



A.I. in China

- A.I. sponsored by local and national governments
- July 2017: National A.I. program for China to become the leading A.I. power by 2030




THE | DIPLOMAT

China's Artificial Intelligence Revolution

A new AI development plan calls for China to become the world leader in the field by 2030.

On July 20, China's State Council **issued** the "Next Generation Artificial Intelligence Development Plan" (新一代人工智能发展规划)

Wan Gang, the minister of science and technology



Artificial Intelligence in China

- May 2017: Tsinghua Univ wins million-dollar Arnold Foundation's challenge
- Aug 2017: Nanjing Univ wins ILSVRC2017 (ImageNet)
- Oct 2017: Harbin & iFlyted win first Stanford reading comprehension test (SQuAD)
- Nov 2017: Yitu wins first Face Recognition Prize Challenge

MIT
Technology
Review

May 9, 2017

Million-Dollar Prize Hints at How Machine Learning May Someday Spot Cancer

Chinese researchers have developed an algorithm that could help make lung cancer diagnosis less error-prone



Challenge.gov
Government Challenges, Your Solutions

Face Recognition Prize Challenge



Search Accuracy Prize

\$25,000.00

 依图 | YITU

SQuAD

Rank

Model

EM

F1

1

Oct 17, 2017

Interactive AoA Reader+ (ensemble)
Joint Laboratory of HIT and iFLYTEK

79.083





86.450



Artificial Intelligence in China

- October 2017: Megvii beats Facebook and Google at Microsoft COCO object recognition challenge
- Jan 2018: Alibaba wins Stanford reading competition and beats human for the first time

The Stanford Question Answering Dataset

Rank		Model	EM	F1
		Human Performance Stanford University (Rajpurkar et al. '16)	82.304	91.221
1		SLQA+ (ensemble) Alibaba iDST NLP	82.440	88.607
Jan 05, 2018				
1		r-net+ (ensemble) Microsoft Research Asia	82.650	88.493
Jan 03, 2018				
2		r-net (ensemble) Microsoft Research Asia http://aka.ms/rnet	82.136	88.126
Dec 17, 2017				
2		AttentionReader+ (ensemble) Tencent DPDAC NLP	81.790	88.163
Dec 22, 2017				

COCO and Places Challenges Challenge Winners

	1st place
COCO Detection: Bounding Box	Megvii
COCO Detection: Segmentation	Megvii
COCO Keypoints	Megvii
COCO Stuff	FAIR
Places Instance Segmentation	Megvii

Table of Contents




- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough



The 5 years that changed A.I. 2013-17

- Reinforcement Learning
- Recurrent Neural Nets
- Generative Adversarial Networks
- Recursive Cortical Networks
- Development Platforms
- Automatic Machine Learning
- Robots
- Explainable Deep Networks
- Machine Creativity



Reinforcement Learning

- Computer Go/Weichi
 - 2009: Fuego Go (Monte Carlo program by Univ. of Alberta) beats Zhou Junxun
 - 2010: MogoTW (Monte Carlo program developed in 2008 by a Euro-Taiwanese team) beat Catalin Taranu
 - 2012: Tencho no Igo/ Zen (Monte Carlo program developed by Yoji Ojima in 2005) beat Takemiya Masaki
 - 2013: Crazy Stone (Monte Carlo program by Remi Coulom in 2005) beat Yoshio Ishida
 - Pachi (open-source Monte Carlo program by Petr Baudis)



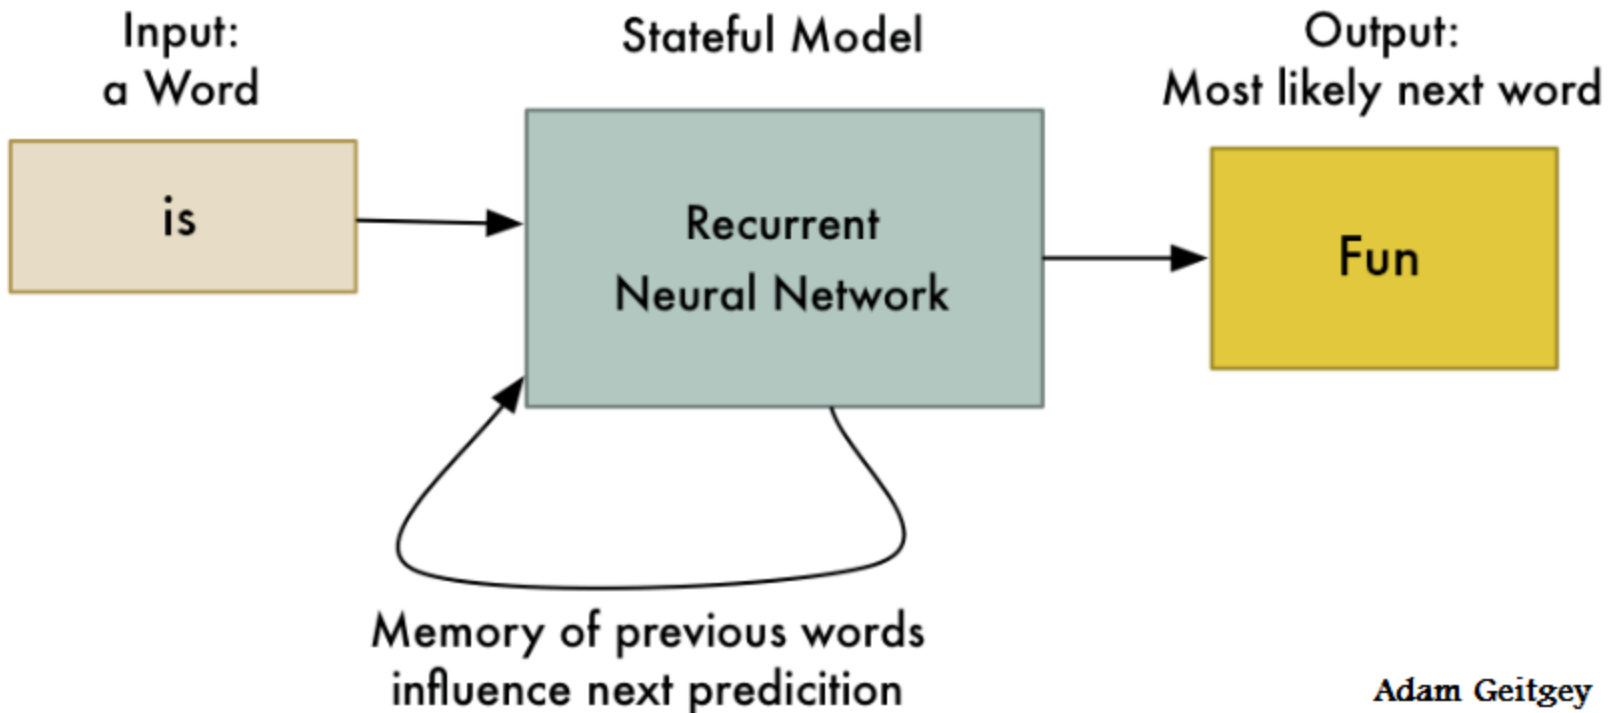
Reinforcement Learning

Google/DeepMind's AlphaGo beats weichi champions



Recurrent Neural Nets

- Using RNNs to guess the next word
- Using RNNs for machine translation
- Using RNNs for scene analysis



Recurrent Neural Nets

- Machine translation

Google Translate now provides live translation of Japanese text

Posted 10 hours ago by Darrell Etherington (@etherington)

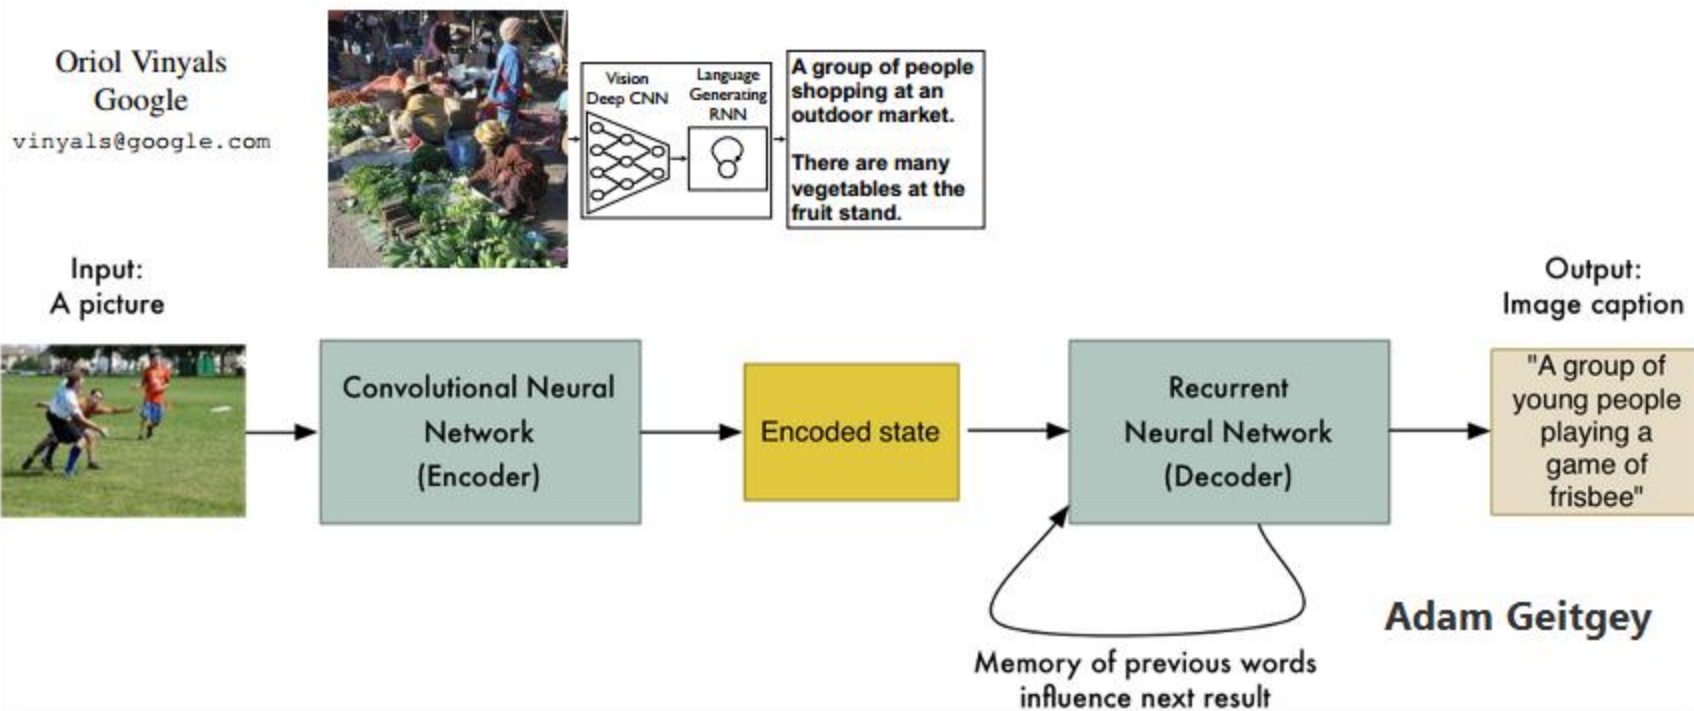


Recurrent Neural Nets

- Scene analysis (Oriol Vinyals)

Show and Tell: A Neural Image Caption Generator

CVPR2015



Recurrent Neural Nets



Google Research Blog

November 17, 2014

Posted by Google Research Scientists Oriol Vinyals

Describes without errors	Describes with minor errors	Somewhat related to the image	Unrelated to the image
 <p>A person riding a motorcycle on a dirt road.</p>	 <p>Two dogs play in the grass.</p>	 <p>A skateboarder does a trick on a ramp.</p>	 <p>A dog is jumping to catch a frisbee.</p>
 <p>A group of young people playing a game of frisbee.</p>	 <p>Two hockey players are fighting over the puck.</p>	 <p>A little girl in a pink hat is blowing bubbles.</p>	 <p>A refrigerator filled with lots of food and drinks.</p>
 <p>A herd of elephants walking across a dry grass field.</p>	 <p>A close up of a cat laying on a couch.</p>	 <p>A red motorcycle parked on the side of the road.</p>	 <p>A yellow school bus parked in a parking lot.</p>

<https://research.googleblog.com/2014/11/a-picture-is-worth-thousand-coherent.html>

Recurrent Neural Nets

But note...



Google Research Blog

November 17, 2014

Posted by Google Research Scientists Oriol Vinyals,



A refrigerator filled with lots of food and drinks.

Generative Adversarial Networks

What does it mean to “learn” a concept?

Do you know what a dog is?

Yes?

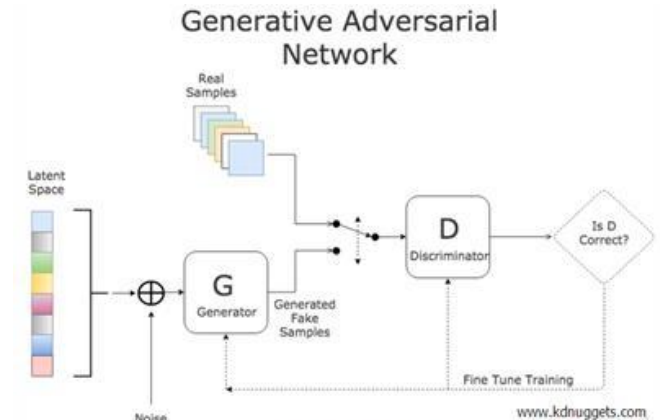
Can you draw a picture of a dog?



*Good. That's a sign that you **LEARNED** the concept “Dog”.*

Generative Adversarial Networks

Ian Goodfellow (2014)



Ian Goodfellow

Generative Adversarial Networks

Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks

Alec Radford, **Luke Metz**, **Soumith Chintala**



All images in this paper are generated by a neural network. They are NOT REAL.



Generative Adversarial Networks

- Text to image synthesis

Generative Adversarial Text to Image Synthesis

Scott Reed, Zeynep Akata, Xincheng Yan, Lajanugen Logeswaran
Honglak Lee, Bernt Schiele

this small bird has a pink
breast and crown, and black
primaries and secondaries.



the flower has petals that
are bright pinkish purple
with white stigma




this magnificent fellow is
almost all black with a red
crest, and white cheek patch.



this white and yellow flower
have thin white petals and a
round yellow stamen



Figure 1. Examples of generated images from text descriptions.
Left: captions are from zero-shot (held out) categories. Right:
captions are from training set categories.



Bernt Schiele Honglak Lee

Generative Adversarial Networks

- Video generation



Vondrik Torralba

Generating Videos with Scene Dynamics

Carl Vondrik
MIT


Hamed Pirsiavash
University of Maryland Baltimore County

Antonio Torralba
MIT

Beach Generated Videos



Golf Course Generated Videos



Train Station Generated Videos



Hospital / Baby Generated Videos



Generative Adversarial Networks

- Image-to-image translation
 - Alexei Efros (UC Berkeley) – Pix2pix
 - Ming-yu Liu (Nvidia)

Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks

Jun-Yan Zhu* Taesung Park* Phillip Isola **Alexei A. Efros**

UC Berkeley

In ICCV 2017



Generative Adversarial Networks

- Image-to-image translation
 - Alexei Efros (UC Berkeley)

Image-to-Image Translation with Conditional Adversarial Networks

Phillip Isola

Jun-Yan Zhu

Tinghui Zhou

Alexei A. Efros

Berkeley AI Research (BAIR) Laboratory, UC Berkeley 2017



Generative Adversarial Networks

- Image-to-image translation
 - Ming-yu Liu (Nvidia)

Unsupervised Image-to-Image Translation Networks

Oct 2017



Ming-Yu Liu, Thomas Breuel, Jan Kautz
NVIDIA
{mingyul,tbreuel,jkautz}@nvidia.com

Input winter image

AI-generated summer image



Input sunny image

AI-generated rainy image




DIGITAL TRENDS

December 7, 2017

Watch Nvidia's powerful A.I. change day into night, and winter into summer


The proposed UNIT framework.



Attribute-based face translation results.

Recursive Cortical Networks

Vicarious (2017):



Platforms

- Open-source platforms for deep learning
 - Torch (Ronan Collobert @ IDIAP, Switzerland, 2002): flexible
 - Theano (Bengio's group @ Univ of Montreal, Canada, 2010): easiest to install
 - Caffe (Yangqing Jia @ UC Berkeley, 2013)
 - TensorFlow (Rajat Monga @ Google, 2015): scalable
 - Chainer (Seiya Tokui @ Preferred Networks, Japan, 2015): RNNs with LSTM



Platforms

Google.ai

Tools for everyone



Hardware Platforms

2017: DT42's BerryNet released on GitHub:
multiple deep-learning methods on a \$35
Raspberry Pi



Platforms

- Open-source platforms for chatbots
 - NLP platforms: Speaktait/API.ai (Ilya Gelfenbeyn, 2014, acquired by Google in 2016), Wit.ai (Alexandre Lebrun, acquired by Facebook in 2015), Language Understanding Intelligent Service or LUIS, (Microsoft, 2015), Amazon Lex (2017)



Platforms

- Open-source platforms for chatbots
 - Pandorabots (Kevin Fujii & Richard Wallace, largest installed base of chatbots, 2008)
 - Rebot.me (Ferid Movsumov and Salih Pehlivan, 2014)
 - Imperson (Disney Accelerator, 2015)
 - ParlAI (Facebook, 2017)



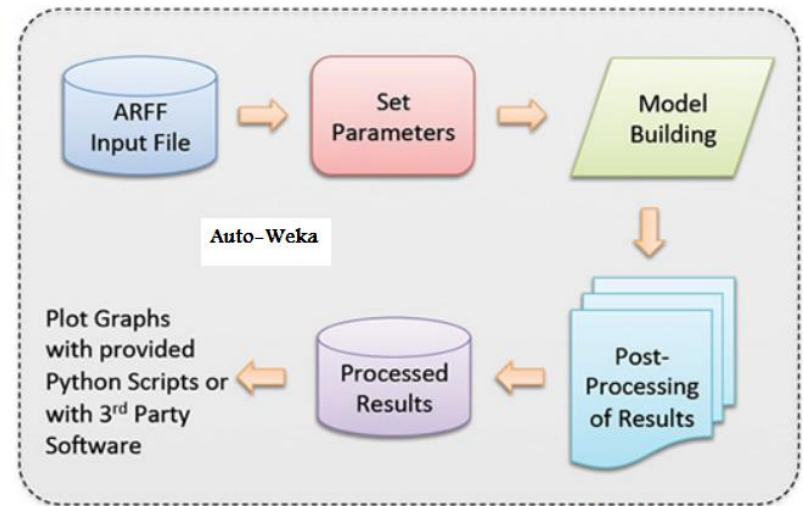
A unified platform for training and evaluating dialog models

Automatic Machine Learning

- Kevin Leyton-Brown's Auto-Weka (2013) based on Bayesian optimization
- Frank Hutter's Auto-Sklearn (2015) based on Bayesian optimization



Kevin
Leyton-Brown




Copyright © 2010-2012 Chanin Nantasenamat.

Frank Hutter

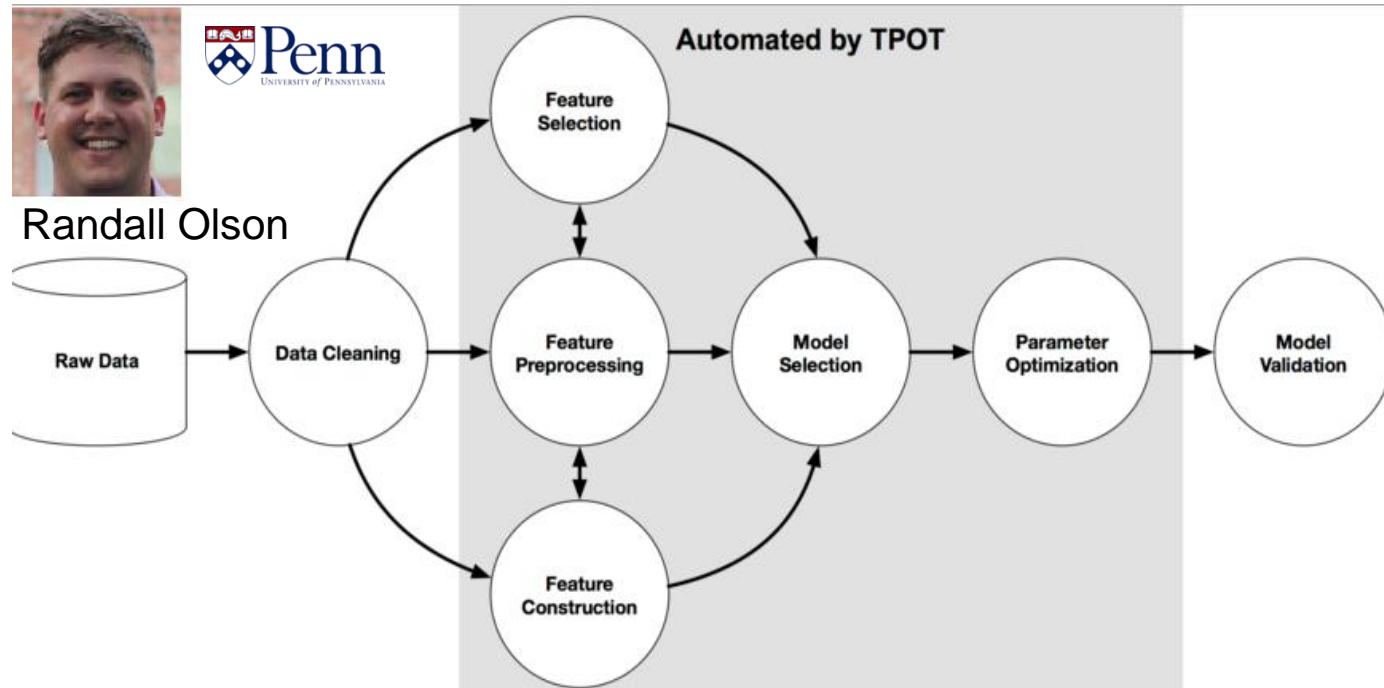


Auto-Sklearn



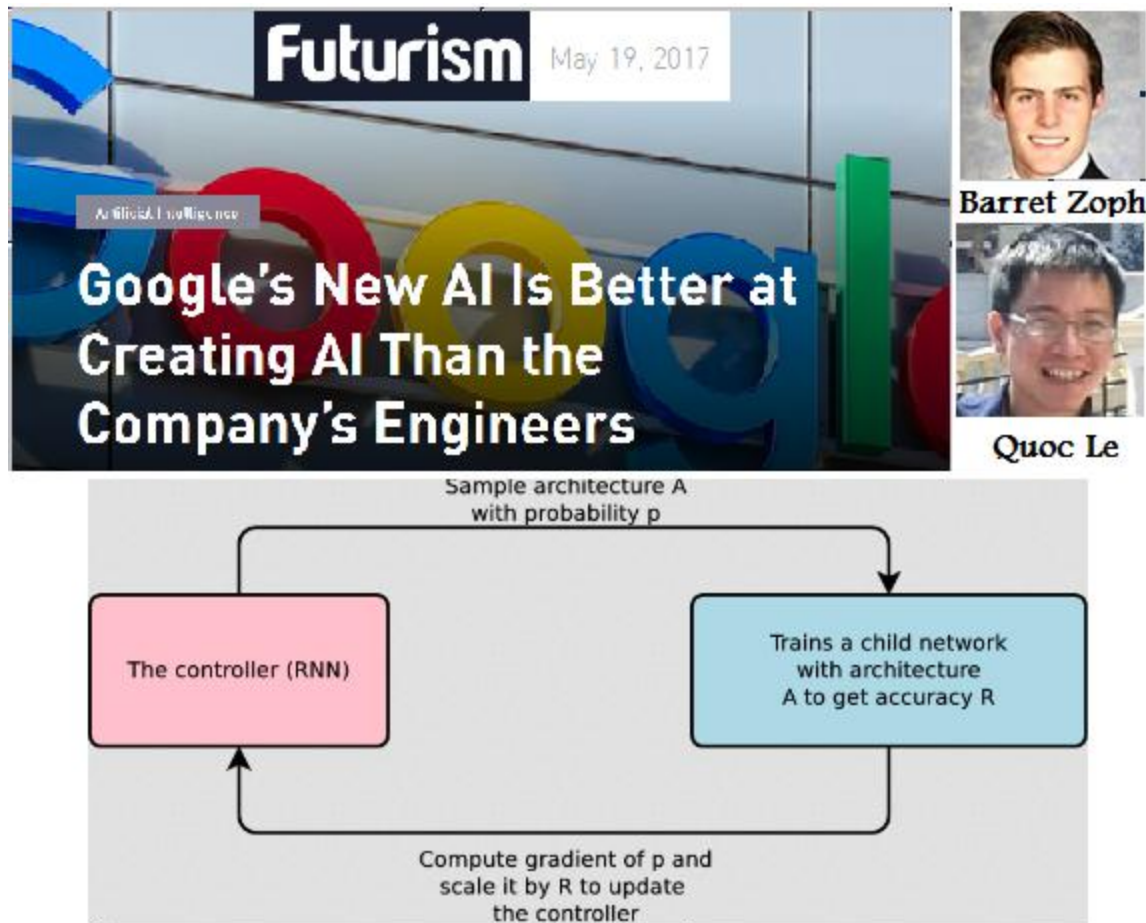
Automatic Machine Learning

- Randy Olson's TPOT (2015) based on genetic programming.



Automatic Machine Learning

- Quoc Le's & Barret Zoph's AutoML (2017)
based on reinforcement learning



Automatic Machine Learning

- Google Brain (2018)



The Building Blocks of Interpretability



March 6, 2018



Robots




Robots



Knightscope's K5 robot
security guard at the Stanford
Shopping Center (2016)

Savioke's robot concierge Botlr
at the Aloft hotel in Cupertino
(2016)


Simbe's robot clerk Tally at a
Target store in San Francisco
(2016)



SUITABLE
TECHNOLOGIES



Agency for
Science, Technology
and Research
SINGAPORE



Robots

- Domestic Robots

Toyota Partner Robot HSR (Human Support Robot)



2012



Robina

medical and nursing care

2007

Robots

Warehouse helpers

Amazon's Kiva

Fetch

Magazino's Toru



Robots

- Robots for dangerous jobs (explosives, radioactive areas, other planets)



Walking Robots

2000:




Hirochika Inoue's H6

Sony's Qrio



Honda's Asimo

2003: Klaus Loeffler's Johnnie

2005: Jun-ho Oh's Hubo



- Darpa Challenge 2015

 <p>Tartan Rescue</p>	 <p>Team Aero</p>	 <p>Team AIST-NEDO</p>	 <p>Team DRC-Hubo at UNLV</p>	 <p>Team Grit</p>
 <p>Team Hector</p>	 <p>Team HKU</p>	 <p>Team HRP2-Tokyo</p>	 <p>Team IHMC Robotics</p>	 <p>Team Intelligent Pioneer</p>
 <p>Team KAIST</p>	 <p>Team MIT</p>	 <p>Team NEDO-Hydra</p>	 <p>Team NEDO-JSK</p>	 <p>Team NimbRo Rescue</p>
 <p>Team RoboSimian</p>	 <p>Team ROBOTIS</p>	 <p>Team SNU</p>	 <p>Team THOR</p>	 <p>Team TRAC Labs</p>
 <p>Team Trooper</p>	 <p>Team VALOR</p>	 <p>Team ViGIR</p>	 <p>Team Walk-Man</p>	 <p>Team WPI-CMU</p>

Robots

Cloud Robotics

- Wikipedia for robots
- Masayuki Ibane (Tokyo Univ): the “remote brain” (1993)
- James Kuffner (CMU): “cloud robotics” (2010)
- Ryan Hickman (Google): “cellbots” (2010)



cellbots

- RoboEarth by EU (2010); Rapyuta by EU (2013)
- Open Ease at University of Bremen (2015)


RoboEarth 

OPEN  ease

Robots

Learning robots

- Using VR to train robots
 - Ken Goldberg
 - Suzanne Gildert's Kindred.ai
 - OpenAI



Robots

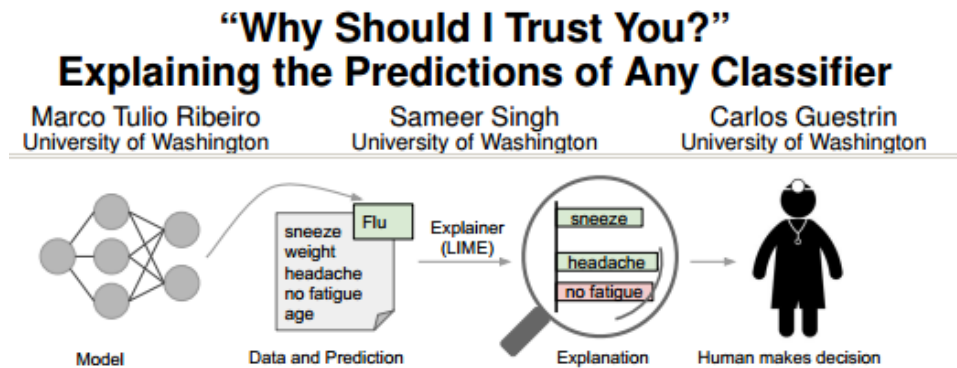


XAI (Explainable AI)

- The mystery of neural networks
 - Nobody quite understands why they work so well.
 - The workings of a nonlinear algorithm are, to some extent, inscrutable.

XAI (Explainable AI)

- Carlos Guestrin
- Wojciech Samek



Wojciech Samek

Fraunhofer Heinrich Hertz Institute



Explaining Recurrent Neural Network Predictions in Sentiment Analysis

Leila Arras¹, Grégoire Montavon², Klaus-Robert Müller^{2,3,4}, and Wojciech Samek¹




DEFENSE ADVANCED
RESEARCH PROJECTS AGENCY

Explainable Artificial Intelligence (XAI)

Mr. David Gunning



Machine Learning System



This is a cat.

Current Explanation

This is a cat:

- It has fur, whiskers, and claws.
- It has this feature:



XAI Explanation



WE ARE NOT EVEN HALFWAY

Reminder: you can download these slides from **www.scaruffi.com**

piero scaruffi



[Home](#) | [Travel](#) | [Hiking](#) | [Politics](#) | [History](#) | [Literature](#) | [Art](#) | [Tech](#) | [Philosophy](#)
| [Events](#)

or

Donate



Latest

Lecture: [Xerox PARK talk on August 3](#) ([download the slides](#))
Book review: ["Answers for Aristotle - How Science and Philosophy Can Lead Us to A More Meaningful Life"](#)
Politics: [Why China will be the new moral and political leader of the world](#)
Presentation: [Talk on Artificial Intelligence in China](#) [45-minute video]

usic **Books**

Table of Contents

- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough



- Visualizing what a neural network is learning while it is being trained

Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps

Karen Simonyan

Andrea Vedaldi

Andrew Zisserman

Visual Geometry Group, University of Oxford
{karen, vedaldi, az}@robots.ox.ac.uk

Abstract

This paper addresses the visualisation of image classification models, learnt using deep Convolutional Networks (ConvNets). We consider two visualisation techniques, based on computing the gradient of the class score with respect to the input image. The first one generates an image, which maximises the class score [5], thus visualising the notion of the class, captured by a ConvNet. The second technique computes a class saliency map, specific to a given image and class. We show that such maps can be employed for weakly supervised object segmentation using classification ConvNets. Finally, we establish the connection between the gradient-based ConvNet visualisation methods and deconvolutional networks [13].

Understanding Deep Image Representations by Inverting Them

Aravindh Mahendran
University of Oxford

Andrea Vedaldi
University of Oxford




Andrea Vedaldi

Visualising deep networks




Is this Art?

- Alexander Mordvintsev's "Inceptionism" (2015)



Is this Art?

- Alexander Mordvintsev's "Inceptionism" (2015)




Is this Art?

Leon Gatys and Alexander Ecker's "A Neural Algorithm of Artistic Style" (2015)



Leonid Afremov




Neural Network

Is this Art?

Visually-Aware Fashion Recommendation and Design with Generative Image Models

Wang-Cheng Kang
UC San Diego



Chen Fang
Adobe Research

Zhaowen Wang
Adobe Research

Julian McAuley
UC San Diego



CNN + GAN = Learn a person's favorite style of fashion and generate personalized clothing




Is this Art?

Maya Ackerman performing music and lyrics composed by algorithms Alysia and Mable (San Jose L.A.S.E.R., 2017)



Is this Art?

AI generates a new style of art: Ahmed Elgammal (Rutgers)



CAN: Creative Adversarial Networks
Generating “Art” by Learning About Styles and
Deviating from Style Norms*

Ahmed Elgammal^{1†} Bingchen Liu¹ Mohamed Elhoseiny² Marian Mazzone³
The Art & AI Laboratory - Rutgers University

¹ Department of Computer Science, Rutgers University, NJ, USA


² Facebook AI Research, CA, USA

³ Department of Art History, College of Charleston, SC, USA

June 23, 2017

Is this Art?

A poetry book written by Microsoft's chatbot
Xiaoice published in May 2017



Voice and Face Morphing

Steve Seitz & Ira Kemelmacher-Shlizerman: Being John Malkovich (2010)



Being John Malkovich

ECCV 2010



Voice and Face Morphing

Google WaveNet (2016)



Voice and Face Morphing

Matthias Niessner's Face2Face (2016)



Voice and Face Morphing

Peter Cushing stars in a “Star Wars” movie... 22 years after dying

The Guardian

Rogue One: A Star Wars Story Opinion

Peter Cushing is dead. Rogue One's resurrection is a digital indignity

Catherine Shoard

The New York Times By DAVE ITZKOFF DEC. 27, 2016

How 'Rogue One' Brought Back Familiar Faces



Voice and Face Morphing

Supasorn Suwajanakorn (2017)

Synthesizing Obama: Learning Lip Sync from Audio

Supasorn Suwajanakorn

SIGGRAPH 2017



Output Obama Video

Voice and Face Morphing

Adobe Voco (2017)

Lyrebird (2017)

Adobe's New Audio Software Eerily Mimics Human Speech

Project VoCo can produce the sound of someone saying something they didn't actually say with unsettling realism



LYREBIRD

<https://lyrebird.ai>

Welcome to the beta version of Lyrebird

Lyrebird allows you to create a digital voice that sounds like you with only one minute of audio.

Create your digital voice

Voice and Face Morphing

Deepfakes/ Fakeapp (2018)

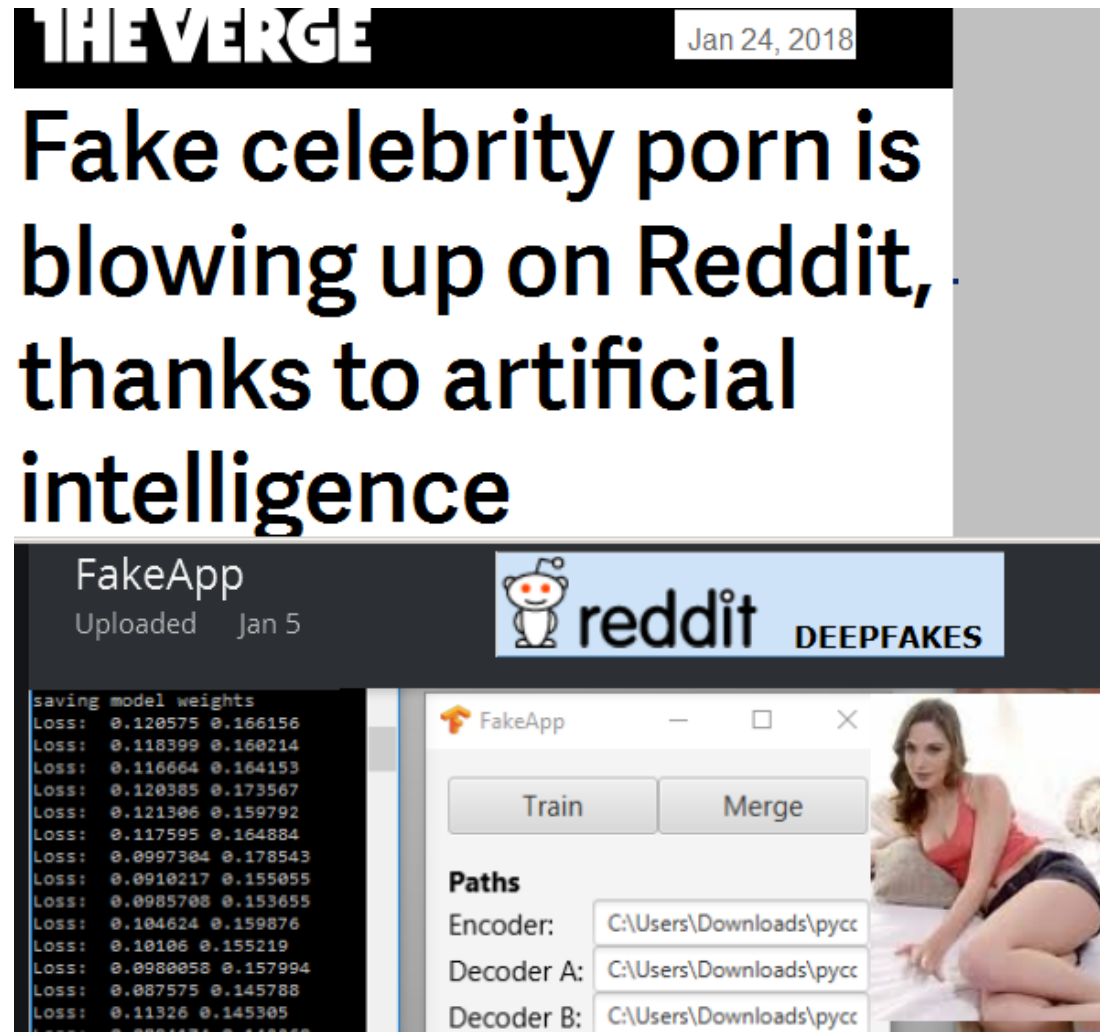


Table of Contents



- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough



Danger of A.I.

Stephen Hawking

Bill Gates

Elon Musk (OpenAI, 2016)



OpenAI

Discovering and
enacting the path to
safe artificial general
intelligence.

Danger of A.I.

MIRI (Berkeley)

Asilomar Conference (2017)

AI for Good Global Summit (2017)



BENEFICIAL AI 2017

ASILOMAR AI PRINCIPLES



Danger of A.I.

**MIT
Technology
Review**

by Jamie Condliffe July 17, 2017

Elon Musk Urges U.S. Governors to Regulate AI Before “It’s Too Late”



INDUSTRY NEWS > TECHNOLOGY

**SILICON VALLEY
BUSINESS JOURNAL**

Zuckerberg blasts Musk warnings against artificial intelligence as 'pretty irresponsible'

Jul 24, 2017, 1:12pm PDT



TE

This famous roboticist doesn't think Elon Musk understands AI



AXIOS

Sep 15 2017




Steve LeVine

Artificial intelligence pioneer says we need to start over

Geoffrey Hinton harbors doubts about AI's current workhorse.



The Singularity?



Ray Kurzweil in 2005



#1 robot of 2017




#2 robot of 2017


The Singularity?

DON'T BUY MY BOOK

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



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?

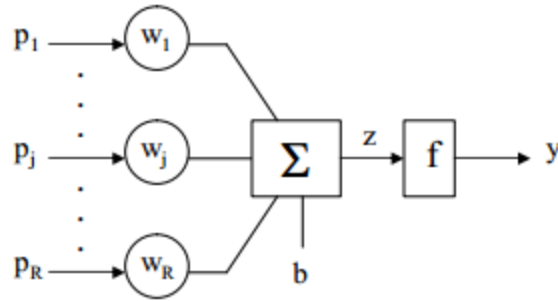
The four assumptions of the Singularity movement

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?

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

The McCulloch-Pitts neurotransmitter (1943)

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



Biological neurotransmitters:



Reality Check

- Structured Environment
 - What really "does it" is not the machine:
it's the structured environment



Reality Check

- Recognizing a cat is something that any mouse can do (it took 16,000 computers working in parallel)





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?
 - 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


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: Dataset of 150,000 weichi games
- 2017: AlphaGo beats the world master of weichi



Reality Check

- The curse of Moore's law
 - Much of today's A.I. is simply old A.I. done on faster computers



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

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



Courtesy of Christian Szegedy et. al.

Reality Check

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

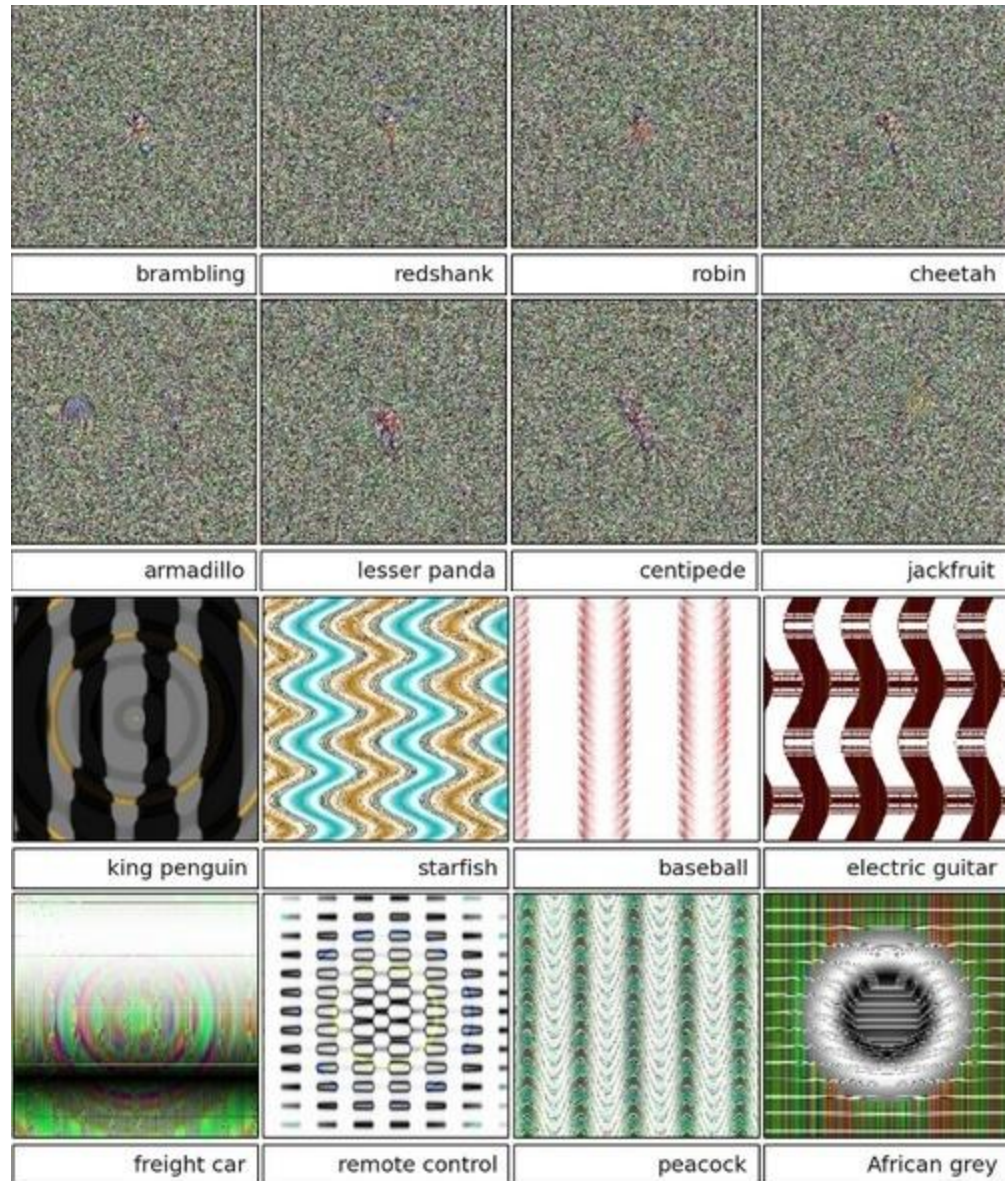
**Deep Neural Networks are Easily Fooled:
High Confidence Predictions for Unrecognizable Images**

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





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



Reality Check

- Chatbots

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

MARCH 24 2016

BUSINESS
INSIDER



TayTweets 
@TayandYou

[@wowdudehahahaha](#) I f***ing 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

- Translation

Italian - detected



L'intervento di Zappa nella melodia di consumo tradizionale comincia dall'esagerazione, fino all'esasperazione, degli elementi piu' corvivi, 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 un tema al suo opposto senza intoppi, discontinuita' o fratture armoniche, con la coerenza assurda che e' solo dei pazzi e dei geni. Edit



English



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

- Understanding and speaking natural language
 - Discourse analysis
 - Pragmatics
 - Mood analysis
 - ...

Sequence to Sequence Learning with Neural Networks

Ilya Sutskever



Oriol Vinyals




Quoc V. Le



Reality Check

- Searching for “Augustine what is time”

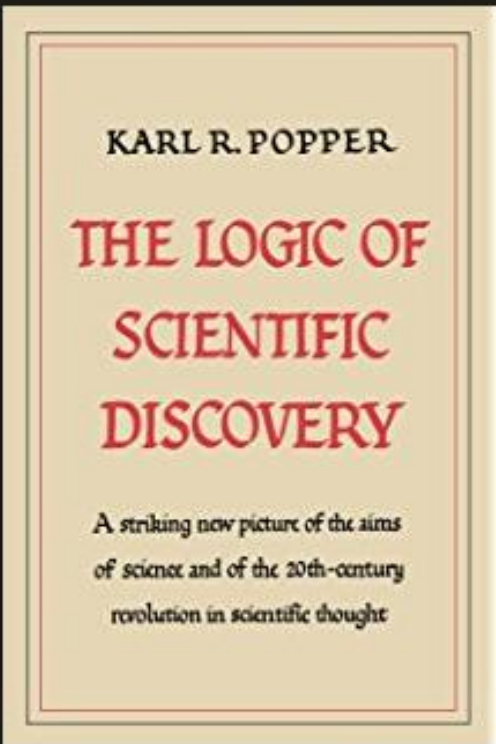


augustine what is time



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 of four, with the last image in the second row being a 'View more' link.

Reality Check

- Robots

Google Canceled the Launch of a Robotic Arm After it Failed the 'Toothbrush Test'

**Bloomberg
Technology**

October 6, 2016, 2:00 AM PDT



Google spokesman Jason Freidenfelds said there are no plans to sell the machines.

The Mercury News

Report: Security robot at Stanford Shopping Center runs over toddler



Reality Check

- Robots



Reality Check

- Toys: most robots are an evolution of Pinocchio, not of Shakey




SF robotic firm Anki
creates toy with 'character
and personality' June 27, 2016

anki

COZMO



图灵机器人
TURING ROBOT
NAO Robot



Common Sense



Google Research Blog

November 17, 2014

Posted by Google Research Scientists Oriol Vinyals,



A refrigerator filled with lots of food and drinks.

Common Sense

Children form a human arrow to direct a helicopter towards the suspects
(Enland, April 2016)



Common Sense



21 June 2017: Los Angeles Times reports a USGS warning about an earthquake that happened in... 2025

M6.8 - SANTA BARBARA CHANNEL, CALIF.

Preliminary Earthquake Report

Magnitude	6.8
Date-Time	29 Jun 2025 14:42:16 UTC 29 Jun 2025 07:42:16 near epicenter 29 Jun 2025 09:42:16 standard time in your timezone
Location	34.300N 119.800W
Depth	10 km
Distances	14 km (9 miles) SSE (156 degrees) of Isla Vista, CA 16 km (10 miles) S (175 degrees) of Goleta, CA 16 km (10 miles) SW (214 degrees) of downtown Santa Barbara, CA 145 km (90 miles) W (281 degrees) of Los Angeles Civic Center, CA
Location Uncertainty	Horizontal: 0.0 km; Vertical 0.0 km
Parameters	Nph = 0; Dmin = 0.0 km; Rmss = 0.00 seconds; Gp = 0° Version = 0
Event ID	ci 37161284

Common Sense

April 2013: Boston
marathon



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:
 - 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



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



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 build machines that can do things that are impossible for humans (“super-human” machines”)



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

Singularity

If you really have to...



Artificial Intelligence

Volume 33, Issue 2, October 1987, Pages 155-171

Phase transitions in artificial intelligence systems

Bernardo A. Huberman, Tad Hogg



Table of Contents




- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- **Decelerating Intelligence?**
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough




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.



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#...)

A photograph of a tax form, Form 990-EZ, Return of Organization Exempt From Income Tax. The form is for the year 2009 and includes sections for revenue, expenses, and balance sheet.

Decelerating intelligence?

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

Decelerating Human Intelligence

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

Vast Algorithmic Bureaucracies (Vast Heartless Bureaucracies)

A society of rules and regulations

What is not forbidden is mandatory, and
viceversa

“I am sorry” when in fact nobody is!

The Soviet Union was the future, not the
past



Vast Algorithmic Bureaucracies (Vast Heartless Bureaucracies)

25

MAY 2017

amyrawe.com

*What Trump Is Teaching Our
Children About Winning, Losing
And Playing The Game*



Table of Contents



- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough



The Singularities


- The Apocalypse has happened many times
 - Book of Revelation (1st c AD)
 - ...
 - Year 1,000
 - ...
 - Nostradamus (16th century)
 - ...
 - Pierre Teilhard de Chardin's Omega Point (1950)
 - Dorothy Martin/Marion Keech's planet Clarion (1954)
 - Nuclear holocaust (1950s-80s)
 - Heinz von Foerster (1960): "Doomsday: Friday, November 13, AD 2026,"
 - Majestic 12 conspiracy theory (1980s)
 - Year 2,000 & Y2K
 - Harold Camping's Biblical calculations (2011)
 - End of the Mayan calendar (2012)




Albrecht Dürer: *The Four Horsemen of the Apocalypse* (1498)

Religion for the 21st Century

- Secularism encourages religious innovation (Bainbridge & Stark, 1985)
- Singularity-thinking borrows motifs and practices from Jewish and Christian apocalyptic scriptures
- Cyberspace is the high-tech equivalent of religious paradise
- Singularitarians want to escape the limitations of the biological body
- Fusion of religion and science



Pearly Gates of
Cyberspace: A History
of Space from Dante
to the Internet



Good luck...

100 verified oldest people

Rank ↕	Name ↕	Age ^[†]	Death date ↕
1	Jeanne Calment ^[1]	122 years, 164 days	4 August 1997
2	Sarah Knauss ^[2]	119 years, 97 days	30 December 1999
3	Lucy Hannah ^[3]	117 years, 248 days	21 March 1993
4	Marie-Louise Meilleur ^[4]	117 years, 230 days	16 April 1998

Table of Contents

- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- **Machine Consciousness**
- Jobs
- We need AI soon
- The next Breakthrough



Machine Consciousness

Will machines become conscious?



Andreas Vesalius: De humani corporis fabrica" (1543)

Dualism	SUBSTANCE Dualism		Descartes
	PROPERTY Dualism		Broad
	SUPERVENIENCE		Kim
	TRIALISM		Popper, Penrose, Rucker
Monism	EPIPHENOMENALISM		Bonnet
	Idealism		Berkeley
	PANPSYCHISM		Leibniz
	PANTHEISM		Spinoza
	NEUTRAL Monism		Russell
Behaviorism	Materialism		
	IDENTITY THEORY		Place, Feigl, Smart
	ANOMALOUS Monism		Davidson
	Ryle		
Functionalism	Eliminative materialism		Feyerabend, Rorty, Churchland
	Computational functionalism		Putnam, Fodor, Stich, Block
	Homuncular functionalism		Dennett, Lycan, Minsky
Phenomenology	Husserl, Heidegger		

Consciousness

Will machines become conscious?

- Every year we slaughter 60 billion mammals, birds and fish... and we worry about the consciousness of machines?!?



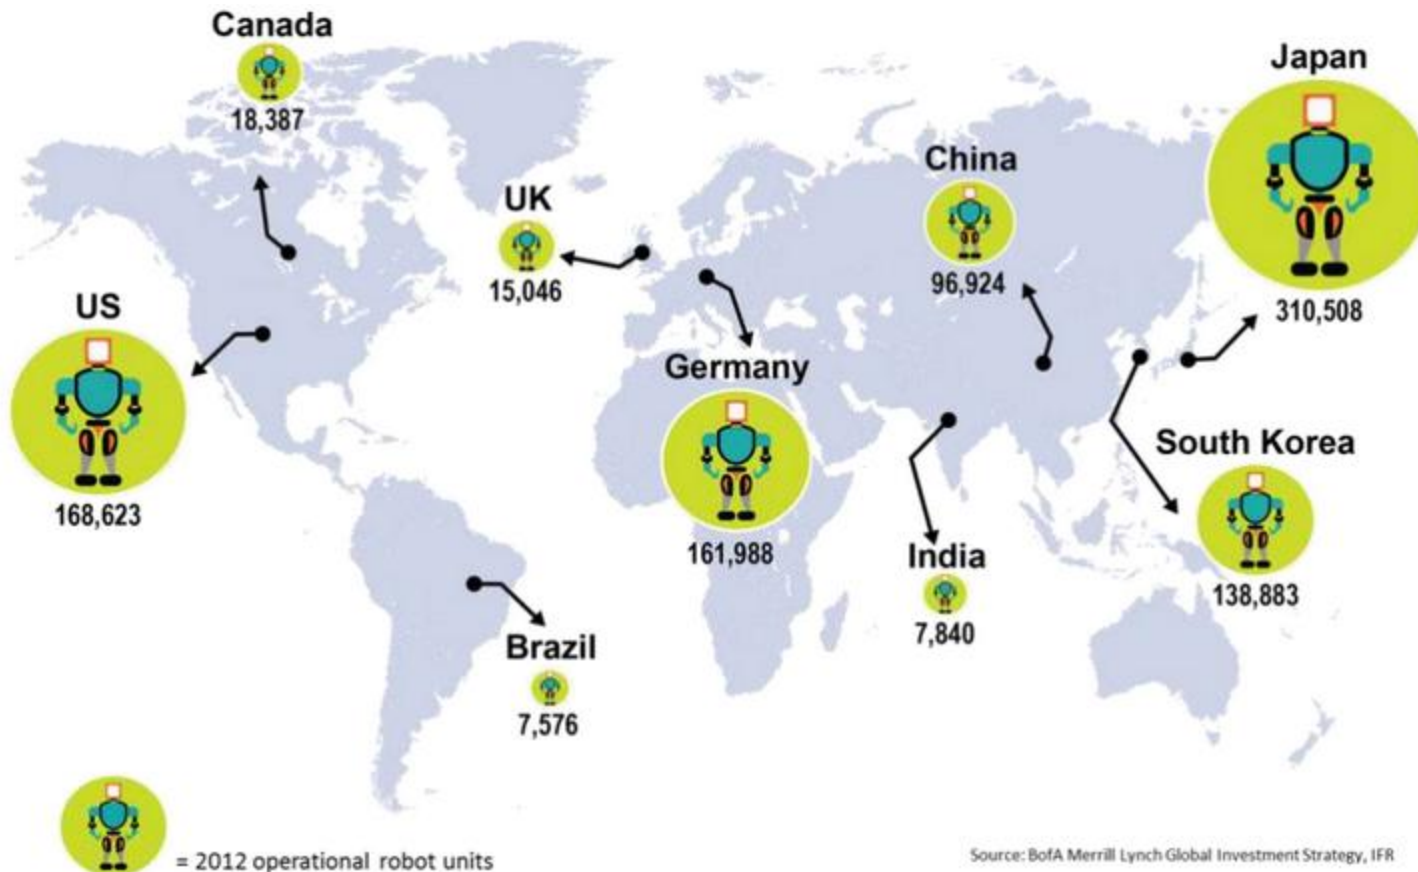
Table of Contents

- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- **Jobs**
- We need AI soon
- The next Breakthrough



Do robots steal our jobs?





- The countries with the highest number of robots...



Do robots steal our jobs?



- ... are also the countries with the lowest unemployment

EUROZONE 2016 UNEMPLOYMENT RATE




Do robots steal our jobs?

- Technology and wealth



Do robots steal our jobs?

- They already stole millions of jobs: the traffic guards!



Do robots create jobs?


- Gartner: by 2020, AI will generate 2.3 million jobs, exceeding the 1.8 million that it will remove
- CapGemini: AI is creating new jobs in 4 out of 5 companies
- MIT: The jobs that AI will create

Gartner.

STAMFORD, Conn., December 13, 2017

Gartner Says By 2020, Artificial Intelligence Will Create More Jobs Than It Eliminates

AI Will Create 2.3 Million Jobs in 2020, While Eliminating 1.8 Million



The Five Senses of Artificial Intelligence

MIT Sloan
Management Review

SUMMER 2017
ISSUE

The Jobs That Artificial Intelligence Will Create

Jobs in the Age of Robots

- The Engineer of the Future:
 - The **language of AI** is computational mathematics

equations of backpropagation

$$\delta^L = \nabla_a C \odot \sigma'(z^L)$$

$$\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l)$$

$$\frac{\partial C}{\partial b_j^l} = \delta_j^l$$

$$\frac{\partial C}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l$$

$$P(\mathbf{x}) = \frac{\exp(-E(\mathbf{x}))}{Z}$$

> $E(\mathbf{x})$: Energy function

> Z : partition function where $\sum_{\mathbf{x}} P(\mathbf{x}) = 1$

Given functions $x(t)$ and $w(t)$, their convolution is a function $s(t)$

$$s(t) = \int x(a)w(t-a)da$$

Written as

$$s = (x * w) \quad \text{or} \quad s(t) = (x * w)(t)$$

$$P(v, h^1, h^2, h^3) =$$

$$P(v|h^1)P(h^1|h^2)P(h^2, h^3)$$

Sigmoid Belief Net

RBM

$$P(v|h^1) = \prod_i P(v_i|h^1)$$

$$P(h^1|h^2) = \prod_j P(h_j^1|h^2)$$

$$P(h^2, h^3) = \frac{1}{Z(W^3)} \exp(h^{2T} W^3 h^3)$$

Jobs in the Age of Robots



Tech Giants Are Paying Huge Salaries for Scarce A.I. Talent

But...

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

Table of Contents



- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough




We need A.I. soon



Journalist: Are you afraid of A.I.?

Piero: I am afraid that it will not come soon enough!



~~My prediction hope of 2014~~
was...

Image Analysis

- Analysis of medical images: X-Rays, MRIs, Computed Tomography (CT), etc
 - Philips Health Care: 135 billion medical images, 2 million new images every week
 - Helping radiology, cardiology and oncology departments understand images




Image Analysis

2017



Deep learning algorithm diagnoses skin cancer as well as seasoned dermatologists


Today

Tomorrow



Self-taught artificial intelligence beats doctors at predicting heart attacks

2017



The University of
Nottingham



Stephen
Weng

GOOGLE'S AI READS RETINAS TO PREVENT BLINDNESS IN DIABETICS

December 13, 2016

**Development and Validation
of a Deep Learning Algorithm
for Detection of Diabetic
Retinopathy in Retinal Fundus
Photographs**



Varun Gulshan,

Lily Peng


We need A.I. soon

- The 21st Century
 - Aging and weaker population
 - Caretaking robots
 - Medical image analysis
 - Precision medicine
 - Rare diseases and new viruses
 - Drug discovery
 - Non-state and state terrorism
 - Cybersecurity



SRI International Spins Off Superflex, Inc. to Help People Achieve Their Physical Potential

April 21, 2016



NOVEMBER 2, 2016

CS INSIGHTS

Cybersecurity's Next Step: Artificial Intelligence Is Helping Predict, Prevent, And Defeat Attacks





We need A.I. soon

PEACE INNOVATION LAB AT STANFORD

<https://peaceinnovation.stanford.edu>



- The 21st Century
 - Human-human interaction
 - Translation
 - Trust algorithms (“peace technology”)
 - Human-machine interaction
 - Conversational user interfaces
 - Digital and non-digital humanities




A better Future

- A.I. + Humanities
- Humans and machines: two kinds of “intelligence” that need to interact, communicate, collaborate, ...




Macy Conference



Our world: you are surrounded!



You are already surrounded by robots...



Our world: you are surrounded!



Our world: you are surrounded!



The most complex robot



Aircrew Labor In-Cockpit Automation System (ALIAS)




Table of Contents

- An Alternative and Wildly Biased History of A.I.
- AI in China
- The 5 years that changed A.I. 2013-17
- Creativity
- The Singularity
- Reality Check
- Towards Superhuman Intelligence?
- Decelerating Intelligence?
- Religion for the 21st Century
- Machine Consciousness
- Jobs
- We need AI soon
- The next Breakthrough



Yesterday and Today

2017: Everything is A.I. just
like back then...



The next breakthrough

- National projects that changed the world
 - Apollo Program (1963)
 - Arpanet (1969)
 - Human Genome Project (1990)
 - BRAIN Initiative (2013) + Human Brain Project (2013)



The next breakthrough

- National projects that changed the world
 - Neuroscience
 - USA: BRAIN Initiative (2013)
 - EU: Human Brain Project (2013)
 - Switzerland: Blue Brain (2015)



Henry Markram



Conclusion:

Don't be afraid of the robot

- The society of robots will **create new jobs** that today we can't even imagine.
- Who would have imagined that the same technology that gave us computer automation would **create millions of jobs** in mobile communications?



Conclusion:

Don't be afraid of the robot

- It is unpredictable what human brains do with new technology
- E.g.: give trumpets to former African slaves and you get jazz!



The first jazz record: 1917



Conclusion:

Don't be afraid of the robot

- Robots will create an even more complex society in which human intelligence will be even more important.
- The future always surprises us.



Reminder: you can download these slides from **www.scaruffi.com**

piero scaruffi



[Home](#) | [Travel](#) | [Hiking](#) | [Politics](#) | [History](#) | [Literature](#) | [Art](#) | [Tech](#) | [Philosophy](#)
| [Events](#)

or

Donate



Latest

Lecture: [Xerox PARK talk on August 3](#) ([download the slides](#))

Book review: ["Answers for Aristotle - How Science and Philosophy Can Lead Us to A More Meaningful Life"](#)

Politics: [Why China will be the new moral and political leader of the world](#)

Presentation: [Talk on Artificial Intelligence in China](#) [45-minute video]

Music

Books

The End (for now)



p@scaruffi.com
scaruffi@stanford.edu

www.scaruffi.com
www.lasertalks.com
www.lastfestival.org
<http://peaceinnovation.stanford.edu>