

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 6

- 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

The Singularity?



Singularity

- Jack Good: "Speculations Concerning the First Ultraintelligent Machine" (1965)
- Masahiro Mori: "The Buddha in the Robot" (1974)
- Hans Moravec: "Today's Computers, Intelligent Machines and Our Future" (1978)
- Marvin Minsky : "Will Robots Inherit the Earth" (1994)
- Ray Kurzweil: "The Singularity is Near" (2005)



The Singularity?



Wells Cathedral Clock of the 14th century, a machine that can do something that no human can do: keeping time



Ray Kurzweil at the Singularity University

Arakawa's 1966 solution to the weather forecast, the "mission impossible" of the early computers

Computational Design for Long-Term Numerical Integration of the Equations of Fluid Motion: Two-Dimensional Incompressible Flow. Part I¹

Akio Arakawa
University of California, Los Angeles

The integral constraints on quadratic quantities of physical importance, such as conservation of mean kinetic energy and mean square vorticity, will not be maintained in finite difference analogues of the equation of motion for two-dimensional incompressible flow, unless the finite difference Jacobian expression for the advection term is restricted to a form which properly represents the interaction between grid points, as derived in this paper. It is shown that the derived form of the finite difference Jacobian prevents nonlinear computational instability and thereby permits long-term numerical integrations. © 1966 Academic Press

INTRODUCTION

A major difficulty, which has blocked progress in long-term numerical integration of the equations of fluid motion, has been nonlinear computational instability of the finite difference analogues of the governing differential equations. The existence and cause of this instability was first called to our attention by Phillips [1, 2].

The instability can be illustrated by integration of the vorticity equation for two-dimensional incompressible flow,

$$\frac{\partial \zeta}{\partial t} = J(\zeta, \psi), \quad (2)$$

or

$$\frac{\partial \nabla^2 \psi}{\partial t} = J(\nabla^2 \psi, \psi), \quad (3)$$

where J is the Jacobian operator with respect to the rectangular coordinates, x and y , in the plane.

When the Jacobian in this equation is replaced by space-differences of the usual form,

$$\begin{aligned} \bar{\partial}_{i,j}(\zeta, \psi) = \frac{1}{4d^2} [(\zeta_{i+1,j} - \zeta_{i-1,j})(\psi_{i,j+1} - \psi_{i,j-1}) \\ - (\zeta_{i,j+1} - \zeta_{i,j-1})(\psi_{i+1,j} - \psi_{i-1,j})], \end{aligned} \quad (4)$$

where i is the finite-difference grid index in x , j is the index in y , and d is the grid interval, and the equation is integrated over some tens of time steps, using an ordinary time-centered differencing scheme, it is found that the solution begins to show a characteristic structure termed "stretch-

The Singularity?

- Singularity: a super-human form of intelligence
 - The pessimists: those who think that it will kill us all
 - The optimists: those who think that it will give us immortality



The Singularity?

Ray Kurzweil's predictions

- Infinite life extension: *“Medical technology will be more than a thousand times more advanced than it is today... every new year of research guaranteeing at least one more year of life expectancy”** (2022)
- Precise computer simulations of all regions of the human brain (2027)
- Small computers will have the same processing power as human brains (2029)
- 2030s: Mind uploading - humans become software-based
- 2045: The Singularity



* recently postponed to 2040

The Singularity?

2014: Deep Knowledge Ventures (Hong Kong) appoints an algorithm to its board of directors

16 May 2014 Last updated at 06:29 ET



Algorithm appointed board director



Would board meetings be improved by sitting next to an algorithm?

A venture capital firm has appointed a computer algorithm to its board of directors.

[Related Stories](#)

The Singularity?

- 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
 - Harold Camping's Biblical calculations (2011)
 - End of the Mayan calendar (2012)



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

The Singularity?

Journalist: Are you afraid of A.I.?

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



The Singularity?

The five dogmas of the Singularity movement

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


The Singularity?

Five arguments against the Singularity

1. Reality Check
2. Accelerating Progress?
3. Non-human Intelligence
4. Human Intelligence
5. A Critique of the Turing Test

Reality Check

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



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

NVIDIA GPUs help Microsoft build record-breaking Image Recognition System

Dec 21, 2015 | 39 views | 10 Likes | 0 Comments | [in](#) [f](#) [t](#)

Thanks to the use of GPUs, Microsoft researchers achieved record results on ImageNet, a prestigious image-recognition benchmark. Compared to last year, Microsoft's system cut the top-5 error rate by half, correctly classifying images within 1,000 pre-defined categories more than 96 percent of the time. Their system uses a 152-layer neural network, which is nearly five times deeper than the state of the art.

Reality Check

- Who needs intelligence anymore?
 - Actually, you can find the answer to most questions by simply using a search engine: no need to think, no need for intelligence

How To Fix a Cracked iPhone Screen - YouTube




<https://www.youtube.com/watch?v=3t678W5>

Dec 10, 2015 - Uploaded by HowToBasic

Today I show you how to easily fix any cracked tablet screen for free. This technique works ...

How To Fix a Water Damaged iPhone - YouTube



<https://www.youtube.com/watch?v=D5XrHkL>

Jan 29, 2016 - Uploaded by HowToBasic

Today I show you how to quickly fix any water damaged iPhone. This fast repair technique works with all ...

How To Repair a Door - YouTube



<https://www.youtube.com/watch?v=BCFuE1t>

Apr 7, 2016 - Uploaded by HowToBasic

Today I show you how to easily repair a broken door. This technique works incredibly well on all types of c

Reality Check

- The curse of Moore's law
 - What progress in deep learning will we have if Moore's Law stops working?
 - A major revolution started in 2016

Computing

Intel: Chips Will Have to Sacrifice Speed Gains for Energy Savings

A major technological shift is needed in the next few years if computer chips are to keep improving.

by Katherine Bourzac February 5, 2016

Intel says chips to become slower but more energy efficient

Martin Anderson


Fri 5 Feb 2016 10:35am



William Holt, Intel


Reality Check

- The curse of Moore's law




Reality Check

- Are we moving in the right direction?



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



IMAGENET



Reality Check

- Examples of searches in 2015

The image is a collage of four screenshots illustrating a search process in 2015. A red line traces the path from the initial search on Frys.com to the Google Images results.

Top Left: Frys.com
The Frys.com website header is visible, including the logo, navigation links like "Cool Stuff We Sell", "Things We Do for You", "myFrys", and "Welcome Guest!", and a sidebar with categories such as "Computers(55)", "Electronics(8)", and "Movies, TV Shows & Music(4)".

Top Right: Google Search
A Google search bar shows the query "augustine what is time". Below the search bar, the "Images" tab is selected, displaying a grid of four images related to the search.

Bottom Left: Frys.com Search Results
A screenshot of the Frys.com search results page for the query "widescreen". The results show "All Products (70)" and "Your Store: Palo Alto, CA (48)". A "Search Within Results" box is highlighted, and a "Refine" button is visible. Below the search box, a suggestion is displayed: "Did you mean widescreen, widescreen or widescreen?".

Bottom Right: Google Images Results
Four images are shown in a grid, representing the results of the Google search. The images include: a group of three men standing behind a counter in a pizza shop; a large pepperoni pizza on a tray; a close-up of two slices of pizza on a plate; and a man in a yellow shirt working behind a pizza counter.

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 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
 - How would you call a human being who needs to make 20,000 bigger effort than you to do less than what you can do? More intelligent or more stupid?
 - 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
 - What AlphaGo failed to do: simulating human intuition
 - A program that was succeeding at simulating human intuition: Giraffe, a “***deep reinforcement learning***” engine that taught itself to play chess like a master in 72 hours (September 2015)
 - Giraffe was designed by just one person and it ran on a the humble university computer

TECHNOLOGY

**ARTIFICIAL INTELLIGENCE
IS TAKING COMPUTER
CHESS BEYOND BRUTE
FORCE**

TO INTERNATIONAL MASTER STATUS

By Dave Gershgorin Posted September 16, 2015

Reality Check

- DeepMind's AlphaGo
 - It was designed by Matthew Lai
 - Lai was hired by DeepMind in January 2016



- MOTHERBOARD -

The Chess Engine that Died So AlphaGo Could Live

WRITTEN BY ROLLIN BISHOP

March 14, 2016 // 08:00 AM EST

Reality Check

- DeepMind's AlphaGo
 - What we really learned from AlphaGo (and many other A.I. projects):
 - That Google is very good at publicizing itself and at eliminating competitors



A.I.



Drinks


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

- DeepMind's AlphaGo
 - Limitations of image recognition
 - 2013 (Google + New York Univ + UC Berkeley): tiny perturbations alter the way a neural network



Courtesy of [Christian Szegedy et. al.](#)

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

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 to very 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: AlphaGo's neural network was trained with 150,000 weichi games



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

- Image recognition of 2015 (Google Photos)



Reality Check

- Image recognition of 2016 (Microsoft CaptionBot)

CaptionBot



Microsoft

I am not really confident, but I think it's a bus that is sitting in front of a building.



Common Sense



Google Research Blog

November 17, 2014

Posted by Google Research Scientists Oriol Vinyals

Describes without errors



A person riding a motorcycle on a dirt road.

Describes with minor errors



Two dogs play in the grass.

Somewhat related to the image



A skateboarder does a trick on a ramp.

Unrelated to the image



A dog is jumping to catch a frisbee.



A group of young people playing a game of frisbee.



Two hockey players are fighting over the puck.



A little girl in a pink hat is blowing bubbles.



A refrigerator filled with lots of food and drinks.



A herd of elephants walking across a dry grass field.



A close up of a cat laying on a couch.



A red motorcycle parked on the side of the road.



A yellow school bus parked in a parking lot.

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

Common Sense



Google Research Blog

November 17, 2014


Posted by Google Research Scientists Oriol Vinyals,



A refrigerator filled with lots of food and drinks.

Reality Check

- A remote-controlled toy is NOT a step toward superhuman intelligence
- Human-looking automata that mimic human behavior have been built since ancient times
- A human being is NOT a toy (yet)



Reality Check

- The brain of the roundworm (a few hundred neurons connected by a few thousand synapses) is still smarter than the smartest neural network ever built.



Reality Check

- An easy science
 - Artificial Intelligence is not subjected to the same scrutiny as other sciences
 - Its success stories are largely unproven

Reality Check

- 60 years later it is not machines that learned to understand human language but humans who got used to speak like machines in order to be understood by automated customer support (and mostly not even speak it but simply press keys)

Reality Check

- What “automation” really means...
 - The jobs that have been automated are repetitive and trivial.
 - And in most cases the automation of those jobs has required the user/customer to accept a lower (not higher) quality of service.
 - The more automation around you, the more you (you) are forced to behave like a machine to interact with machines

Reality Check

- Intelligent Behavior from Structured Environments



Reality Check

- Structuring the Environment
 - We structure the chaos of nature because it makes it easier to survive and thrive in it
 - The more we structure the environment, the easier for extremely dumb people and machines to survive and thrive in it.
 - It is easy to build a machine that has to operate in a highly structured environment
 - What really "does it" is not the machine: it's the structured environment

Reality Check

- Semantics
 - It is not intelligent to talk about intelligent machines: whatever they do is not what we do, and, therefore, is neither "intelligent" nor "stupid" (attributes invented to define human behavior)
 - We apply to machines many words invented for humans simply because we don't have a vocabulary for the states of machines

Reality Check

- Semantics
 - Memory is reconstructive
 - Data storage is not “memory”
 - Exponentially increasing data storage does not mean better memory
 - What is “computer speed”?
 - Who is faster at picking a cherry from a tree, the fastest computer in the world or you?

Reality Check

- Where A.I. is truly successful...
 - Most machine intelligence is being employed to couple real-time customization and machine learning in order to understand who you are and tailor situations in real time that will prompt you to buy some products (custom advertising)
 - *"The best minds of my generation are thinking about how to make people click ads"* (former Facebook research scientist Jeff Hammerbacher in 2012)
 - So far A.I. has not created better doctors or engineers, but better traveling salesmen

Reality Check

- Can you please stop improving my experience?



Updating Skype

Just a moment, we're improving your Skype experience...



Reality Check

- Can you please stop beeping?



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



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

- The self-driving car vs... the self-flying plane!



Artificial General Intelligence

- Task-specific vs General-purpose Intelligence
- Originally, A.I. was looking for general-purpose intelligence
- Today's A.I. is looking for task-specific intelligence (recognizing a cat, driving a car)

Artificial General Intelligence

- How to simulate an average human (not just one human task) - the “logic theorist” solution (1960s): create a system that can perform reasoning on knowledge and infer the correct behavior for any situation
- How to simulate an average human (not just one human task) - the brute force solution (2000s): create one specific program/robot for each of the millions of possible situations, and then millions of their variants

Artificial General Intelligence

- The Multiplication of Appliances and Artificial Intelligence by Enumeration
 - We have machines that dispense money (ATMs), machines that wash clothes (washing machines), machines that control the temperature of a room (thermostats), and machines that control the speed of a car (cruise controls).
 - We can build machines for all the other tasks and then collectively call them “equal” to humans

Artificial General Intelligence

- The enumeration problem: which human functions qualify as "intelligent"?
 - There are very human functions that people don't normally associate with "intelligence". They just happen to be things that human bodies do.
 - Do we really want machines that fall asleep or urinate?
 - We swing arms when we walk, but we don't consider "swinging arms while walking" a necessary feature of intelligent beings.


Tips for better A.I.

1. IBM's Watson of 2013 consumes 85,000 Watts compared with the human brain's 20 Watts.
2. The brain is an analog device, not digital
3. What we need: a machine that has only a limited knowledge of all the chess games ever played and is allowed to run only so many logical steps before making a move and can still consistently beat the world champion of chess.
4. Memory is not storage

Tips for better A.I.

- What conditions may foster a breakthrough: it is not the abundance of a resource (such as computing power or information) that triggers a major paradigm shift but the scarcity of a resource.

Why Futurists always get it Wrong



1907: A futurist predicts the world of "wireless telegraphy"




2007: Apple ships the iPhone



Today's world of "wireless telegraphy"

Why Futurists always get it Wrong



L.A. Times of 3 April 1988 predicts life in 2013



L.A. in 2013

Next...

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

