

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 8

- 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 Future of Artificial Intelligence

Updated to...



The Future of A.I.

We have plenty of humans who like to predict the future of humans.

When will we have an A.I. program to predict the future of A.I. programs?

Behavior Recognition

- SRI Intl (2013): Holistic assessment of behavior
 - Evaluating facial expression
 - Evaluating body posture
 - Evaluating voice

Holistic Assessment of Behavior – Multimodal Sensing



Behavior Recognition

- Sergio Escalera (Barcelona):
Looking at People (LAP) contests
 - Affective Computing
 - Social Signal Processing
 - Human Behavior Analysis
 - Social Robotics
- 2011/12 Gesture Challenge
- 2013 Multi-modal Challenge
- 2014-16 Looking At People Challenges



Behavior Designer

- The Behavior Designer: AI not as a replacement for people but instead as a useful companion for altering social behavior (Hirokazu Shirado and Nicholas A. Christakis 2017)



Locally noisy autonomous agents improve global human coordination in network experiments


Hirokazu Shirado & Nicholas A. Christakis

Nature 545, 370–374 (18 May 2017)


Received 29 December 2016



Information-based System



Knowledge-based System



Human intelligence is weird...

- Susan Langer: our brains are symbol-producing machines
- We don't just “do” things: we invent complex rituals to “mean” things
- E.g., the tea ceremony
- What symbols does a neural network learn?



The Age of Invisible Robots

- What do these technologies have in common?
 - Internet of Things
 - Biotech
 - Wearables
 - 3D Printing
 - Social Media
 - Sharing Economy
 - Fintech
 - ...
- They produce data. Billions of data.
- 44 zettabytes of data by 2020 (IDC)

The Age of Invisible Robots

- The **Data Economy**: is your business really your business?
- E.g.: credit card companies are **data companies** not financial companies (they know what you buy)
- Social media are data companies not social networks (they know what you
- The Convergence of Big Data and AI
 - *Deep Learning needs large datasets*
 - *Large datasets need deep learning*

The Age of Invisible Robots

- Deep Learning = Big Data Intelligence
 - Traditional AI = Pattern recognition (image, speech, ...)
 - AI + Big Data = Profile recognition (behavior, desires, emotions...)
 - What you buy + what you do + who are your friends + ... = **your digital behavior**
 - Your digital behavior = **Your mind** (desires, beliefs, needs)

The Age of Invisible Robots

- The machine
 - knows more about you than your best friends know
 - has more data about everybody else than you do
- so the machine can use *knowledge about humans* to understand *who you are* and use *your data* to understand how *to manipulate your mind*
- *The invisible robot is spying on you all the time!*

The Age of Invisible Robots

- The invisible robot can understand intimate facts about your life
- The invisible robot can "design" your behavior

The Age of Invisible Robots

- Invisible robots are the future
- Examples
 - **Politics**: a political candidate can customize the political message to maximize your attention
 - **Staffing**: the invisible robot can select the best person for the job
 - **Justice**: the invisible robot can calculate the probability that a person will commit the same crime again and calculate the optimal sentence accordingly
 - ...

The Age of Invisible Robots

- Sounds scary?
- It gets **worse**: thousands of invisible bots will collaborate on this project



The Age of Invisible Robots

- The **positive** side
 - An algorithm is not biased: not racist, not sexist, etc
 - You like people who like you and you like information that confirms your opinion
 - An algorithm doesn't like anybody in particular
 - So... A.I. can build a more **fair** society

The Age of Invisible Robots

- The Engineer of the Future:
 - The Behavior Designer
 - You need to understand the **language of A.I.**, not just the language of humans

The Age of Invisible 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(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

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

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

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

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

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

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

The Age of Invisible Robots

- You want to be the one who shapes AI...
- ...not the one who is shaped by AI



“The **best**
way to predict
the **⚡FUTURE⚡**
is to
INVENT IT.”



—Alan Kay

Common Sense



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
 A person riding a motorcycle on a dirt road.	 Two dogs play in the grass.	 A skateboarder does a trick on a ramp.	 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 Blog



Google Research Blog

November 17, 2014


Posted by Google Research Scientists Oriol Vinyals,



A refrigerator filled with lots of food and drinks.

Common Sense

The search engine puts the dots precisely where the interesting information is.



The image shows a Google search interface. The search bar contains the text "yearly stem graduates". Below the search bar, there are tabs for "All", "News", "Shopping", "Images", "Videos", "More", and "Settings". The "All" tab is selected. Below the tabs, it says "About 582,000 results (0.44 seconds)". There are two search results displayed. The first result is titled "Data check: U.S. producing more STEM graduates even without ..." and includes a URL "www.sciencemag.org/.../data-check-us-producing-more-stem-graduates-even-without...". The second result is titled "STEM Education Data and Trends - National Science Foundation" and includes a URL "https://www.nsf.gov/nsb/sei/edTool/data/college-14.html". Both results mention "degrees" and "annually". Red circles are drawn around the word "degrees" in the first result and the word "annually" in the second result.

Google

yearly stem graduates

All News Shopping Images Videos More Settings

About 582,000 results (0.44 seconds)

Data check: U.S. producing more STEM graduates even without ...
www.sciencemag.org/.../data-check-us-producing-more-stem-graduates-even-without...
Jun 30, 2014 - Nation on path to awarding 1 million additional technical **degrees** by ... of **STEM** bachelor's **degrees** awarded **annually** by U.S. institutions.

STEM Education Data and Trends - National Science Foundation
<https://www.nsf.gov/nsb/sei/edTool/data/college-14.html>
Who earns bachelor's **degrees** in science and engineering? ... The number of S&E bachelor's awarded **annually** rose steadily from 398,602 in 2000 to ...

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

IMPORTANT - YOU SHOULD READ THIS NOTICE CAREFULLY

Dear Mrs Lindsay Durdle,

This is a default notice served under section 87(1) of the Consumer Credit Act 1974. Your account has an outstanding balance of £3,240.72.

Provision of Agreement Breached

You are in breach of condition 15.4(c) of your agreement with PayPal Credit as we have received notice that you are deceased. In accordance with condition 15.4(c), we are entitled to close your account, terminate your agreement and demand repayment of the full amount outstanding.

Nature of Breach

This breach is not capable of remedy. We therefore intend to take some or all of the following actions after 20 JULY 2018.

1. Act to limit or restrict your ability to access your credit limit;
2. With effect from the date shown, terminate your agreement and require the immediate repayment of the whole amount outstanding under the agreement;
3. Take legal proceedings against you and/or instruct a third party debt collection agency to begin further collections activity to recover the whole amount outstanding under the agreement; and/or
4. Notify your executors or representatives of our claim as a creditor.

We are not required to give any further notice before taking such action, however we do understand that you may be experiencing financial difficulties and are eager to help.

Continued

Copyright © 1999 to 2018 PayPal. All rights reserved.

PayPal (Europe) S.à r.l. et Cie, S.C.A. Société en Commandite par Actions Registered Office:

22-24 Boulevard Royal, L-2449 Luxembourg RCS Luxembourg B 118 349

Paypal bot (2018)
sends a letter to a
dead woman

Common Sense

- What happens if you don't have common sense...

300-pound mall robot runs over toddler - Jul. 14, 2016 - CNN Money
money.cnn.com/2016/07/14/technology/robot-stanford-mall/ ▼

Jul 14, 2016 - **Knightscope**, which makes the robot, said the machine veered to the left to

300-Pound Security Robot Runs Over Toddler At California Shopping
www.huffingtonpost.com/.../security-robot-toddler us 57863670e4b03f... ▼

Security Robot Runs Over Toddler at California Mall [Updated ...
www.roboticstrends.com/.../security_robot_runs_over_toddler_at_califor... ▼

Crime-fighting robot hits, rolls over child at Silicon Valley mall
www.latimes.com/.../la-me-ln-crimefighting-robot-hurts-child-bay-area-... ▼

Security robot accidentally attacks child | Fusion
fusion.net/story/324731/security-robot-runs-over-child/ ▼



Common Sense

- What happens if you don't have common sense...

YAHOO!
NEWS

BGR News 11 July 2018

PayPal just demanded payment from a dead woman



You are in breach of condition 15-4(c) of your agreement with Paypal Credit as we have received notice that you are deceased. In accordance with condition

Common Sense

- What happens if you don't have common sense...

STAT

July 25, 2018

EXCLUSIVE

IBM's Watson supercomputer recommended 'unsafe and incorrect' cancer treatments, internal documents show



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



MARCH 19, 2018

Self-driving Uber car kills Arizona woman crossing street



MARCH 30, 2018

Tesla says crashed vehicle had been on autopilot prior to accident



Common Sense

Children form a human arrow to direct a helicopter towards the suspects



Common Sense

April 2013: Boston
marathon attacks



Common Sense

Can you tell what is wrong with this picture?



Common Sense


- Marriage with Children episode
 - Wife reads from a magazine: “Whom would you like to spend the evening with? A. Your wife – B. ...”
 - Husband interrupts and shouts: “B!”
 - Why is it funny?



Common Sense

- What happens if you have common sense...

VanGogh and Nietzsche went mad in the same year, 1888, one year after Emile Berliner invented the gramophone (that records sounds) and in the same year in which Kodak introduced the first consumer camera (that records images),



Common Sense

- Knowledge-based systems (“expert” systems) did NOT fail
- Do we need deep learning to solve these problems?
 - If you are an American and Americans are humans, are you a human?
 - What is $2+2$?

Common Sense



Common Sense



Common Sense






Common Sense

- What human neural networks do better than artificial neural networks: reasoning.
- Neural networks are good at guessing the cause from the effect, not the effect from the cause
- Google's DeepMath (2016)
- Microsoft's DeepCoder (2017)
- Integration of neural networks and search



Common Sense

- Yejin Choi (University of Washington, 2019):
ATOMIC, a graph of commonsense knowledge,
and COMET, a generator of such graph



Common Sense

- Yejin Choi (University of Washington, 2019): COMET



COMET: Commonsense Transformers for Automatic Knowledge Graph Construction

ACL 2019

Antoine Bosselut

Hannah Rashkin

Maarten Sap

Chaitanya Malaviya

Asli Çelikyilmaz

Me

Common Sense


- DARPA (2018)

Defense Advanced Research Projects Agency News And Events

10/11/2018

Teaching Machines Common Sense Reasoning

DARPA program seeks to articulate and encode humans' basic background knowledge for intelligent systems



Machine Learning

- The success stories/ Why we are happy
 - **Supervised Learning:**
 - image recognition
 - speech recognition
 - machine translation
 - **Reinforcement Learning:**
 - game-playing
 - Unsupervised Learning: ----

Machine Learning

- **Supervised Learning:**

- AlexNet (2012)
- VGGNet or VGG-16 (2014)
- GoogleNet or Inception (2014)
- DeepFace (2014)
- ResNet (2015)
- Inception-ResNet (2016)
- ResNeXt (2017)
- ...



Speech Recognition



Image Recognition

Machine Learning

- The Methods/ Why we are **NOT** happy:
 - Supervised Learning: requires a lot of data - **UNNATURAL**
 - Reinforcement Learning: requires a lot of self-play - **UNNATURAL**
 - Unsupervised Learning: **the way humans and animals learn!**

*A.I. is based on big data, right?
H.I. (human intelligence) is not!
Animal intelligence is not.*

The Future of Machine Learning

Towards real “learning”:

- Meta-learning
- Learning by Demonstration
- Transfer Learning
- Curiosity-driven Exploration
- Multitask Learning
- Self-supervised Learning

The Future of Machine Learning

- Meta-learning
 - Schmidhuber : "Simple Principles of Metalearning" (1996)
 - Nicolas Schweighofer and Kenji Doya (2001)



PERGAMON

Neural Networks letter

Meta-learning in Reinforcement Learning

Nicolas Schweighofer^{a,*}, Kenji Doya^{a,b,1}

^a*CREST, Japan Science and Technology Corporation, ATR, Human Information Science Laboratories, 2-2-2 Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0288, Japan*

^b*ATR Human Information Science Laboratories, 2-2-2 Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0288, Japan*

Received 6 September 2002; accepted 10 October 2002

Neural Networks 16 (2003) 5–9

The Future of Machine Learning

- Meta-learning
 - Yan Duan 's RL Square (2016)



RL²: FAST REINFORCEMENT LEARNING VIA SLOW REINFORCEMENT LEARNING

Yan Duan^{†‡}, John Schulman^{†‡}, Xi Chen^{†‡}, Peter L. Bartlett[†], Ilya Sutskever[‡], Pieter Abbeel^{†‡}

[†] UC Berkeley, Department of Electrical Engineering and Computer Science

[‡] OpenAI



10 Nov 2016



The Future of Machine Learning

- Meta-learning
 - Jane Wang's "deep meta-reinforcement learning" (2017)

LEARNING TO REINFORCEMENT LEARN

JX Wang¹, Z Kurth-Nelson¹, D Tirumala¹, H Soyer¹, JZ Leibo¹,
R Munos¹, C Blundell¹, D Kumaran^{1,3}, M Botvinick^{1,2}

¹DeepMind, London, UK

²Gatsby Computational Neuroscience Unit, UCL, London, UK

³Institute of Cognitive Neuroscience, UCL, London, UK



23 Jan 2017



nature
neuroscience

14 May 2018

Nature Neuroscience **21**, 860–868 (2018)

Prefrontal cortex as a meta-reinforcement learning system

Jane X. Wang, Zeb Kurth-Nelson, Dharshan Kumaran, Dhruva Tirumala,
Hubert Soyer, Joel Z. Leibo, Demis Hassabis & Matthew Botvinick




The Future of Machine Learning

- Meta-learning
 - Chelsea Finn 's "model-agnostic meta-learning" MAML (2017)




Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks

Chelsea Finn¹ Pieter Abbeel^{1,2} Sergey Levine¹



MAML on HalfCheetah.



MAML on Ant.



18 Jul 2017

The Future of Machine Learning

- Meta-learning
 - Alex Nichol and John Schulman's Reptile (2018)

On First-Order Meta-Learning Algorithms

Alex Nichol and Joshua Achiam and John Schulman
OpenAI



OpenAI

4 Apr 2018

The Future of Machine Learning

- Learning by Demonstration
 - Kuniyoshi, Inaba & Inoue in Japan (1989)
 - Donald Michie (1990)
 - Claude Sammut (1992)
 - Stefan Schaal's "Is Imitation Humanoid Robots?" (1999)

Trends in Cognitive Sciences 3:233-242. (1999).

Is Imitation Learning the Route to Humanoid Robots?




Figure 3: Conceptual sketch of an imitation learning system.

The Future of Machine Learning

- Learning by Demonstration
 - Dean Pomerleau's ALVINN to follow the lane (1988)



- Pieter Abbeel and Andrew Ng (2004)

Apprenticeship Learning via Inverse Reinforcement Learning

Pieter Abbeel

Andrew Y. Ng

Computer Science Department, Stanford University, Stanford, CA 94305, USA




Figure 5. Screenshot of driving simulator.

The Future of Machine Learning

- Learning by Demonstration
 - Tianhe Yu and Chelsea Finn: Domain-Adaptive Meta-Learning (DAML) that trains a deep network with many videos of human and robot movements performed for different tasks (2018)

One-Shot Imitation from Observing Humans via Domain-Adaptive Meta-Learning

Tianhe Yu*, Chelsea Finn*,



Pieter Abbeel, Sergey Levine
University of California, Berkeley

5 Feb 2018



The Future of Machine Learning

- Learning by Demonstration
 - Tianhe Yu and Chelsea Finn: Domain-Adaptive Meta-Learning (DAML) that trains a deep network with many videos of human and robot movements performed for different tasks (2018)




Fig. 4. Illustration of the policy architecture. The policy consists of a sequence of five convolutional (conv) layers, followed by a spatial soft-argmax and fully-connected (fc) layers.


The Future of Machine Learning

- Learning by Demonstration
 - Aravind Rajeswaran (Univ of Washington, 2018)

Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations



Aravind Rajeswaran^{1*}, Vikash Kumar^{1,2*}, Abhishek Gupta³, Giulia Vezzani⁴,
John Schulman², Emanuel Todorov¹, Sergey Levine³



26 Jun 2018

The Future of Machine Learning

- Learning by Demonstration
 - Oriol Vinyals (DeepMind, 2016)



Matching Networks for One Shot Learning




Figure 1: Matching Networks architecture

The Future of Machine Learning

- Learning by Demonstration
 - Byron Boots (Georgia Tech, 2018)



Agile Autonomous Driving using End-to-End Deep Imitation Learning

Yunpeng Pan^{*‡}, Ching-An Cheng^{*}, Kamil Saigol^{*}, Keuntaek Lee[†], Xinyan Yan^{*},
Evangelos A. Theodorou^{*}, and Byron Boots^{*}

^{*}Institute for Robotics and Intelligent Machines, [†]School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, Georgia 30332-0250

[‡]JD.com American Technologies Corporation, Mountain View, California 94043



May 2018



Fig. 1: The high-speed off-road driving task.

The Future of Machine Learning

- Learning by Demonstration + Reinforcement
 - Nando de Freitas (DeepMind, 2018)

Reinforcement and Imitation Learning for Diverse Visuomotor Skills

Yuke Zhu[†] Ziyu Wang[‡] Josh Merel[‡] Andrei Rusu[‡] Tom Erez[‡] Serkan Cabi[‡]
Saran Tunyasuvunakool[†] János Kramár[‡] Raia Hadsell[‡] Nando de Freitas[‡] Nicolas Heess[‡]

[†]Computer Science Department, Stanford University, USA

[‡]DeepMind, London, UK



May 2018

Collecting
demonstrations

3D motion controller



Training in
simulation

physics engine



Running on
real robot

real environment



The Future of Machine Learning


- Learning by Demonstration + Reinforcement
 - Aravind Rajeswara & Vikash Kumar (UC Berkeley, 2018)

Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations

Aravind Rajeswaran^{1*}, Vikash Kumar^{1,2*}, Abhishek Gupta³, Giulia Vezzani⁴,
John Schulman², Emanuel Todorov¹, Sergey Levine³



Jun 2018



The Future of Machine Learning

- Learning by Demonstration
 - Andi Palan & Dorsa Sadigh (Stanford, 2019)

Stanford researchers teach robots what humans want



Andi Palan



Dorsa Sadigh

Robotics: Science and Systems 2019
Freiburg im Breisgau, June 22-26, 2019

Learning Reward Functions by Integrating Human Demonstrations and Preferences

Malayandi Palan, Gleb Shevchuk, Nicholas C. Landolfi, Dorsa Sadigh





Fig. 2. The full DemPref framework. The human user provides demonstrations, which are used to learn a prior over reward functions. Then, we actively generate two trajectories that, together with the human's demonstration, maximize the expected information gained from querying the human with them. We then query the human for her ranking of the trajectories. We use that information to update the distribution over reward functions and use the chosen trajectory in place of the demonstration in the next query. We repeat this process of generating trajectories and querying the human until convergence.

The Future of Machine Learning

- Transfer learning: transfer knowledge across tasks
 - Pioneers: Lorian Pratt, Sebastian Thrun, Rich Caruana
 - Schmidhuber ("Evolutionary Principles in Self-referential Learning", 1987)
 - Barto ("Intrinsically Motivated Reinforcement Learning", 2004)
 - Demis Hassabis (DeepMind) talk (2018)



The Future of Machine Learning

- Transfer learning: transfer knowledge across tasks
 - Exploration based on "intrinsic motivation"
 - DeepMind (2016): plays Montezuma Revenge
 - OpenAI (2017): EMAML and ERL2




OpenAI

Some Considerations on Learning to Explore via Meta-Reinforcement Learning

Bradly C. Stadie¹ Ge Yang² Rein Houthoofd³ Xi Chen⁴

Yan Duan⁴ Yuhuai Wu⁵ Pieter Abbeel^{4,6} Ilya Sutskever³



Unifying Count-Based Exploration and Intrinsic Motivation

Marc G. Bellemare
bellemare@google.com


Sriram Srinivasan
srsrinivasan@google.com

Georg Ostrovski
ostrovski@google.com

Tom Schaul
schaul@google.com

David Saxton
saxton@google.com

Rémi Munos
munos@google.com



The Future of Machine Learning


- Transfer learning: transfer knowledge across tasks
 - Curiosity-driven exploration
 - DeepMind (2018)



Kickstarting Deep Reinforcement Learning

Simon Schmitt^{*1} Jonathan J. Hudson^{*1} Augustin Zidek^{*1}
Simon Osindero¹ Carl Doersch¹ Wojciech M. Czarnecki¹ Joel Z. Leibo¹ Heinrich Kuttler¹
Andrew Zisserman¹ Karen Simonyan¹ S. M. Ali Eslami¹

10 Mar 2018




The Future of Machine Learning

- Transfer learning: transfer knowledge across tasks
 - Curiosity-driven exploration
 - Abhishek Gupta (2018)



Meta-Reinforcement Learning of Structured Exploration Strategies

Abhishek Gupta¹ Russell Mendonca¹ YuXuan Liu¹ Pieter Abbeel^{1,2} Sergey Levine¹



20 Feb 2018

The Future of Machine Learning


- Transfer learning: transfer knowledge across tasks
 - using natural language descriptions of the environment to help guide policy transfer



Deep Transfer in Reinforcement Learning by Language Grounding



Karthik Narasimhan, Regina Barzilay and Tommi Jaakkola
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology



The Future of Machine Learning

- Curiosity-driven Exploration
 - Lerrel Pinto (CMU, 2016)





Jul 2016

The Curious Robot: Learning Visual Representations via Physical Interactions

Lerrel Pinto, Dhiraj Gandhi, Yuanfeng Han, Yong-Lae Park, Abhinav Gupta

The Robotics Institute, Carnegie Mellon University



Physical Interaction Data



Conv Layer1 Filters



Conv3 Neuron Activations



Conv5 Neuron Activations

Learned Visual Representation

Fig. 8. Our shared convolutional architecture for four different tasks.

The Future of Machine Learning

- Curiosity-driven Exploration
 - Joshua Achiam (2017): surprise-based intrinsic motivation
 - Deepak Pathak (2017): “Intrinsic Curiosity Model”, a self-supervised reinforcement learning system that uses curiosity as a natural reward signal to enable the agent to explore its environment and learn skills for use later

Mar 2017

SURPRISE-BASED INTRINSIC MOTIVATION FOR DEEP REINFORCEMENT LEARNING

Joshua Achiam & Shankar Sastry
Department of Electrical Engineering and Computer Science
UC Berkeley



Curiosity-driven Exploration by Self-supervised Prediction

34 in International Conference on Machine Learning, Sydney, Australia, 2017

Deepak Pathak¹ Pulkit Agrawal¹ Alexei A. Efros¹ Trevor Darrell¹



Figure 1. Discovering how to play Super Mario Bros



Our curious agent learns to move along the corridors without any extrinsic rewards



The Future of Machine Learning

- Curiosity-driven Exploration
 - E.g. application: Piaget-like learning agent

Stanford EE Computer Systems Colloquium

4:30 PM, Wednesday, December 5, 2018

NEC Auditorium, Gates Computer Science Building Room B3

Leela: a Semantic Intelligent Agent

David Henkel-Wallace



About the talk:

Leela is a semantic artificially intelligent agent modeled on the theories of Jean Piaget. She builds increasingly abstract semantic models of the world from her experiences of exploration, play, and experimentation. As an agent she is able to formulate, execute, and explain her own plans.

The Future of Machine Learning

- Multitask Learning
 - Sebastian Thrun's "Lifelong Robot Learning" (1995): how to enable a robot to continuously learn as it collects new experiences



The Future of Machine Learning

- Multitask Learning
 - Bharath Ramsundar (Stanford + Google, 2015)

Massively Multitask Networks for Drug Discovery

Bharath Ramsundar^{*,†,°}

Steven Kearnes^{*,†}

Patrick Riley[°]

Dale Webster[°]

David Konerding[°]

Vijay Pande[†]

(*Equal contribution, †Stanford University, °Google Inc.)



6 Feb 2015




Figure 1. Multitask neural network.

The Future of Machine Learning

- Multitask Learning
 - Ishan Misra and Abhinav Shrivastava (CMU, 2016): "cross-stitch units", units that look for the best shared representations for multitask learning

Apr 2016

Cross-stitch Networks for Multi-task Learning


Ishan Misra*

Abhinav Shrivastava*

Abhinav Gupta

Martial Hebert

The Robotics Institute, Carnegie Mellon University



The Future of Machine Learning

- Multitask Learning
 - Fei Wang (China, 2017): Residual Attention Network

Residual Attention Network for Image Classification

Fei Wang¹, Mengqing Jiang², Chen Qian¹, Shuo Yang³, Cheng Li¹,
Honggang Zhang⁴, Xiaogang Wang³, Xiaoou Tang³

¹SenseTime Group Limited, ²Tsinghua University,

³The Chinese University of Hong Kong, ⁴Beijing University of Posts and Telecommunications

Apr 2017

The Future of Machine Learning

- Multitask Learning
 - Shikun Liu (Imperial College London, 2018): Multi-Task Attention Network (MTAN), a shared network with shared features and with task-specific attention modules

28 Mar 2018

End-to-End Multi-Task Learning with Attention



Shikun Liu Edward Johns Andrew Davison



Imperial College
London



The Future of Machine Learning

- Multitask Learning
 - Sebastian Ruder (National University of Ireland, 2018): "sluice networks", a framework for learning loosely-related tasks, a generalization of cross-stitch networks

Nov 2018




Latent Multi-task Architecture Learning

Sebastian Ruder^{1,2}, Joachim Bingel³, Isabelle Augenstein³, Anders Søgaard³

¹Insight Research Centre, National University of Ireland, Galway

²Aylien Ltd., Dublin, Ireland

³Department of Computer Science, University of Copenhagen, Denmark



The Future of Machine Learning

- Multitask Learning
 - Andrei Rusu (DeepMind, 2016): Progressive Nets



Sep 2016

Progressive Neural Networks DeepMind

Andrei A. Rusu*, Neil C. Rabinowitz*, Guillaume Desjardins, Hubert Soyer,
James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, Raia Hadsell

* These authors contributed equally to this work

Google DeepMind



(c) Atari games

The Future of Machine Learning


- Multitask Learning
 - Iasonas Kokkinos (UCL, 2017): UberNet

UberNet: Training a Universal Convolutional Neural Network for Low-, Mid-, and High-Level Vision using Diverse Datasets and Limited Memory





Iasonas Kokkinos

University College London & Facebook Artificial Intelligence Research



UberNet architecture



The Future of Machine Learning

- Multitask Learning
 - Aravind Rajeswaran (Univ of Washington, 2018)

Towards Generalization and Simplicity in Continuous Control



Aravind Rajeswaran* Kendall Lowrey* Emanuel Todorov Sham Kakade
University of Washington Seattle

20 Mar 2018




Figure 2: Hopper completes a get-up sequence before moving to its normal forward walking behavior. The getup sequence is learned along side the forward hopping in the modified task setting.

Conditional Computing

- Conditional Computation for faster networks
 - Bengio's conditional computation (2013)



Deep Learning of Representations: Looking Forward

Yoshua Bengio

Department of Computer Science and Operations Research
Université de Montréal, Canada




Figure 1: Conditional computation architecture: gater path on the left produces sparse output which are multiplied elementwise by the expert outputs, hidden units H_i .

Conditional Computing

- Conditional Computation for faster networks
 - Google's Sparsely-gated Mixture of Experts (2017)



March 2018

ICLR 2017

OUTRAGEOUSLY LARGE NEURAL NETWORKS: THE SPARSELY-GATED MIXTURE-OF-EXPERTS LAYER

Noam Shazeer¹, Azalia Mirhoseini^{*†1}, Krzysztof Maziarsz^{*2}, Andy Davis¹, Quoc Le¹, Geoffrey Hinton¹ and Jeff Dean¹

¹Google Brain, {noam,azalia,andydavis,qvl,geoffhinton,jeff}@google.com

²Jagiellonian University, Cracow, krzysztof.maziarsz@student.uj.edu.pl





Figure 1: A Mixture of Experts (MoE) layer embedded within a recurrent language model.

Recursive Cortical Networks

Vicarious (2017):



What Robots Need to Understand

(What is going to happen next)

- Roozbeh Mottaghi (Allen Institute, 2016): **effect of forces on objects** (what happens if you push an object beyond the edge of a table)
- Jacob Walker (Carnegie Mellon, 2014): **effect of movement of objects** (where can the car go?)
- Michael Ryoo (ETRI in Korea, 2011): **effect of people's movement** (what happens next when a person pulls out a gun?)

What Robots Need to Understand


- Roozbeh Mottaghi (Allen Institute, 2016): **effect of forces on objects** (what happens if you push an object beyond the edge of a table)

2016

“What happens if...”

Learning to Predict the Effect of Forces in Images


Roozbeh Mottaghi, Mohammad Rastegari, Abhinav Gupta, Ali Farhadi



t = 0



t = 1



t = 2

What Robots Need to Understand

Patch to the Future: Unsupervised Visual Prediction

Jacob Walker, Abhinav Gupta, and Martial Hebert
Robotics Institute, Carnegie Mellon University

- Jacob Walker (CMU, 2014): **effect of movement of objects** (where can the car go?)




Figure 1. Consider the scene shown in image (a). Our data-driven approach uses a large collection of videos to predict the likely future of an agent in the scene. The heatmap (b) shows the likely locations the car can visit in the future (along with a few possible trajectories.) (c) shows the hallucination of car moving straight and (d) shows the hallucination of the car turning left.

What Robots Need to Understand

- Michael Ryoo (ETRI in Korea, 2011): **effect of people's movement** (what happens next when a person pulls out a gun?)

2011

Human Activity Prediction: Early Recognition of Ongoing Activities from Streaming Videos

M. S. Ryoo

Electronics and Telecommunications Research Institute, Daejeon, Korea

A. Human activity *classification* problem

Problem: What is this activity?

Observation:



B. Human activity *prediction* problem

Problem: What is this activity?

Observation:



Self-supervised Learning

Self-Supervised Learning, Yann LeCun, Facebook AI Research

After a brief presentation of the state of the art in deep learning, some promising principles & methods for self-supervised learning will be discussed.



Wednesday, October 24, 2018

3:30pm-4:30pm


Carpenter 013

**Google's Amazing Self-Supervised
Computer Vision Model can Track Objects
in any Video!**

JUNE 28, 2018

Applications

- Robotic surgery
- Autonomous driving



Prediction Networks

“Predictive coding” from the neuroscience literature

Discard the predictable, learn from the unpredictable (from your mistakes of prediction)

Srinivasan, Laughlin and Dubs




(1982): predictive coding

Felleman and VanEssen (1991): the cortex is layered and hierarchical

Rao and Ballard (1999): the brain is a prediction machine

Rasmus Palm (2012): the "predictive" autoencoder

Chalasani & Principe (2013): Deep Predictive Coding Networks



Unsupervised Learning

Why we need unsupervised learning

Supervised learning is difficult for videos because videos are much higher-dimensional entities than single images

Where unsupervised learning can work

Luckily videos contain a lot of structure (spatial and temporal regularities) that can be used to train the network

Self-supervised learning!



Self-supervised Learning

Nets trained to predict the next frame in a video

- Marc'Aurelio Ranzato (Facebook, 2014)
- Bill Lotter (Harvard Univ, 2015)
- Xiaolong Wang (CMU, 2015)
- Nitish Srivastava (Univ. of Toronto, 2015)
- Phillip Isola (UC Berkeley, 2016)
- Carl Vondrick (MIT, 2016)



UNIVERSITY OF
TORONTO



Berkeley
UNIVERSITY OF CALIFORNIA



Self-supervised Learning

Other pioneers

- Pulkit Agrawal: KittiNet (UC Berkeley, 2015)
- Carl Doersch: spatial correlations (UC Berkeley, 2015)



Self-supervised Learning

Bill Lotter (Harvard Univ, 2015): PredNet

- prediction of future frames in a video sequence
- unsupervised learning for learning the structure of the world
- the network learns internal representations
- implicit learning of scene structure



DEEP PREDICTIVE CODING NETWORKS FOR VIDEO PREDICTION AND UNSUPERVISED LEARNING


William Lotter, Gabriel Kreiman & David Cox
Harvard University
Cambridge, MA 02215, USA

ICLR 2017



Self-supervised Learning

PredNet



Actual

1

Predicted



Self-supervised Learning


Xiaolong Wang (CMU, 2015): Visual tracking of object (follow the object while the video is rolling)




Unsupervised Learning of Visual Representations using Videos

Xiaolong Wang, Abhinav Gupta
Robotics Institute, Carnegie Mellon University

Oct 2015



(a) Unsupervised Tracking in Videos




(b) Siamese-triplet Network

(c) Ranking Objective

Self-supervised Learning

Phillip Isola (UC Berkeley, 2016): even learns to group pictures by theme




Self-supervised Learning

Carl Vondrick (MIT, 2016): predicting human actions

Anticipating Visual Representations from Unlabeled Video

Carl Vondrick Hamed Pirsiavash[†] Antonio Torralba
Massachusetts Institute of Technology [†]University of Maryland, Baltimore County



Nov 2016



Self-supervised Learning


Pulkit Agrawal (UC Berkeley, 2015): learn from the camera of a self-driving car

Pre-training + Fine-tuning



Self-supervised Learning

Mehdi Noroozi and Paolo Favaro (Univ. of Bern, 2017): pre-training (jigsaw puzzle) + fine-tuning (classification task)



Aug 2017

Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles


Mehdi Noroozi and Paolo Favaro

Institute for Informatiks
University of Bern

Self-supervised Learning

PVM (Filip Piekiewicz, UC San Diego, 2016):

- Not based on deep learning
- Sven Behnke and Raul Rojas (1998): both horizontal (lateral) and vertical (feedback and feedforward) loops
- Rodney Douglas and Kevan Martin (2007): connectivity is mostly local
- Learns from video tracking



Unsupervised Learning from Continuous Video in a Scalable Predictive Recurrent Network


Filip Piekiewicz* Patryk Laurent Csaba Petre
filip@piekiewicz.info laurent@braincorporation.com csaba.petre@gmail.com

Micah Richert Dmitry Fisher
richert@braincorporation.com fisher@braincorporation.com

Todd L. Hylton
October 3, 2016

Self-supervised Learning

PVM



	Deep Learning	Predictive Vision Model
Information Flow	Feedforward ⁹	Extensive Feedback
Computation	End-to-end error propagation	Local prediction
Spatio-Temporal Context	Rarely used ¹⁰	Top-Down/Integrated
Data	Shuffled frames/vectors	Continuous video
Advantages of Neuromorphic Impl.	Small	Large ¹¹
Learning	Offline, primarily Supervised	Online, primarily Unsupervised
Example Tasks	Human-centric Classification	Visually guided behavior such as object tracking

Self-supervised Learning

Chelsea Finn (UC Berkeley, 2016): Learn from visual tracking and “imagine” possible futures

Unsupervised Learning for Physical Interaction through Video Prediction


17 Oct 2016

Chelsea Finn
UC Berkeley



Ian Goodfellow
OpenAI

Sergey Levine
Google Brain



Architecture of the CDNA model

Self-supervised Learning

Deepak Pathak (UC Berkeley, 2017):

Apr 2017

Learning Features by Watching Objects Move

Deepak Pathak^{1,2,*}, Ross Girshick¹, Piotr Dollár¹, Trevor Darrell², and Bharath Hariharan¹

¹Facebook AI Research (FAIR)

²University of California, Berkeley






Figure 2. Overview of our approach. We use motion cues to segment objects in videos *without any supervision*. We then train a ConvNet to predict these segmentations from *static frames*, i.e. without any motion cues. We then transfer the learned representation to other recognition tasks.

Self-supervised Learning




Olivia Wiles (Oxford, 2018): self-supervised networks trained only using videos

- Facial Attributes-Net (FAb-Net): a self-for learning a facial attribute about pose and expression
- X2Face for face puppeteering




Self-supervised Learning

Carl Vondrick (Google, 2018): networks for visual tracking trained using large amounts of unlabeled video



Reference Frame Future Frame (gray) Predicted Color



Self-supervised Learning

Chuang Gan (MIT, 2018): geometry as auxiliary supervision for the self-supervised learning of video representations




Method	YUPENN Acc (%)	Maryland Acc (%)
Spacetime [10]	67.7	86.0
Orientation [6]	43.1	80.7
Object Patch [37]	48.46	70.47
Seq Ver. [24]	51.53	76.67
Geometry (Ours)	69.23	86.90

Method	UCF101 Acc (%)	HMDB51 Acc (%)
DrLim [13]	38.4	13.4
TempCoh [25]	45.4	15.9
Object Patch [37]	42.7	15.6
Seq Ver. [24]	50.9	19.8
OPN [20]	56.3	22.1
Geometry (Ours)	54.1	22.6
Bidirectional Geometry (Ours)	55.1	23.3

37 = Wang (CMU)
20 = Lee (UC Merced)
24 = Misra (Facebook)
10 = Feichtenhofer (Austria)

Self-supervised Learning

Andrew Owens (UC Berkeley, 2018):
multisensory representation of video





Self-supervised Learning

Dahun Kim (South Korea, 2018): three-dimensional representation of video

Self-Supervised Video Representation Learning with Space-Time Cubic Puzzles


Dahun Kim, Donghyeon Cho, In So Kweon
Dept. of Electrical Engineering, KAIST, Daejeon, Korea



Method	Backbone	UCF101(%)	HMDB51(%)
Random initialization	3D ResNet-18	42.4	17.1
Random initialization	AlexNet	38.4	13.4
Temporal Coherency (Mobahi, Collobert, and Weston 2009)	AlexNet	45.4	15.9
Object Patch (Wang and Gupta 2015)	AlexNet	42.7	15.6
Sequence Verification (Misra, Zitnick, and Hebert 2016)	AlexNet	50.9	19.8
OPN (Lee et al. 2017)	AlexNet	56.3	22.1
Geometry (Gan et al. 2018)	AlexNet	54.1	22.6
Time Arrow (Wei et al. 2018)	AlexNet	55.3	-
Video Generation (Vondrick, Pirsiavash, and Torralba 2016)	C3D	52.1	-
3D ST-puzzle (ours)	C3D	60.6	28.3
	3D ResNet-10	63.4	30.8
	3D ResNet-18	65.8	33.7

Self-supervised Learning

Self-supervised learning forces the network to learn a higher-level representation that can then be used for other recognition tasks



Mehdi Noroozi¹

Self-supervised Learning

2016: the year of Deep Reinforcement Learning (e.g. AlphaGo)

2017: the year of Generative Adversarial Networks (e.g. Pix2pix)

2018: the year of Self-supervised Learning

Small-data algorithms

Remedies when there aren't enough data to train a neural network

- Few-shot learning (like OpenAI's GPT)
- Self-supervised learning
- Transfer learning
- Synthetic data generation by GANs


Unsupervised Learning

Applications: Darktrace (2017)

MIT
Technology
Review

The rare form of machine learning that can spot hackers who have already broken in

Darktrace's unsupervised-learning models sound the alarm before intruders can cause serious damage.



Unsupervised Learning


Missing Data Methods: Alexei Efros (UC Berkeley)

Unsupervised Methods: **Missing Data Problems**

Unsupervised Visual Representation Learning by Context Prediction


Carl Doersch, Abhinav Gupta, Alexei A. Efros

(Submitted on 19 May 2015 (v1), last revised 16 Jan 2016 (this version, v3))



$$X = \begin{pmatrix} \text{bird head} \\ \text{bird neck} \end{pmatrix}; Y = 3$$

Figure 7: The visualization shows how well the learned features work.



Context Encoders: Feature Learning by Inpainting

Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, Alexei A. Efros

(Submitted on 25 Apr 2016 (v1), last revised 21 Nov 2016 (this version, v2))

Unsupervised Methods: **Missing Data Problems**


Colorful Image Colorization

Rizang Zhang, Phillip Isola, Alexei A. Efros
(rich.zhang, isola, efros}@eecs.berkeley.edu

University of California, Berkeley




Fig. 1. Example input grayscale photos and output colorization from our algorithm. These examples are cases where our model works especially well. Please visit <http://richzhang.github.io/colorization/> to see the full range of results and to try our model and tools. Best viewed in color (please!).




These "prediction" techniques create better representations,
"force" the system to create better representations

Unsupervised Learning

Local aggregation (Daniel Yamins, Stanford)





Performance increases not just on object categorization but also many other visual tasks ... suggesting general improvement in representation.




Capsule Nets

Geoff Hinton (2017)
Problem of convnets:

Proposal of 2011:



Face



Face

Implementation of 2017:




image 3.0: CapsNet Architecture

$$\mathbf{v}_j = \frac{\|\mathbf{s}_j\|^2}{1 + \|\mathbf{s}_j\|^2} \frac{\mathbf{s}_j}{\|\mathbf{s}_j\|}$$


image 3.1: Novel Squashing Function



Dynamic Routing

Bruno Olshausen (1993)

Dynamic routing
(Olshausen, Anderson, Van Essen 1993)



World Model

David Ha & Jürgen Schmidhuber (2018)



- V: a variational autoencoder
- M: an LSTM recurrent neural network with 256 hidden units
- C: a densely connected neural network

At each time step, our agent receives an **observation** from the environment.



World Model

The **Vision Model (V)** encodes the high-dimensional observation into a low-dimensional latent vector.

The **Memory RNN (M)** integrates the historical codes to create a representation that can predict future states.

A small **Controller (C)** uses the representations from both **V** and **M** to select good actions.

The agent performs **actions** that go back and affect the environment.



Concept Formation

Igor Mordatch (OpenAI, 2018)

6 Nov 2018


Concept Learning with Energy-Based Models



Igor Mordatch
OpenAI
San Francisco, CA



OpenAI



6 Nov 2019


Implicit Generation and Modeling with Energy-Based Models



Yilun Du *
MIT CSAIL

Igor Mordatch
Google Brain

Work done at OpenAI



Concept Formation

Vicarious (2018)

**Behavior is Everything – Towards Representing Concepts
with Sensorimotor Contingencies**

**Nicholas Hay, Michael Stark, Alexander Schlegel, Carter Wendelken,
Dennis Park, Éric Purdy, Tom Silver, D. Scott Phoenix, and Dileep George**
Vicarious AI, San Francisco, CA, USA



<i>concept</i>		<i>CSP expressions</i>
<i>cl_self-left-of-two-objects_vs_one</i>	c	$\exists f, x, y. \text{floor}(f) \wedge \text{on-top}(x, f) \wedge \text{on-top}(y, f) \wedge \text{on-top}(\text{self}, f)$
	g	$\exists f, x, y. \text{blue-floor}(f) \wedge \text{white}(x) \wedge \text{white}(y) \wedge \text{small-blob}(x) \wedge \text{small-blob}(y)$ $\wedge \text{on-top}(x, f) \wedge \text{on-top}(y, f) \wedge \text{on-top}(\text{self}, f) \wedge \neg \text{touches}(x, y)$ $\wedge \text{dir-right-of}(x, \text{self}) \wedge \text{dir-right-of}(y, \text{self})$
	g	$\exists f, x. \text{blue-floor}(f) \wedge \text{white}(x) \wedge \text{small-blob}(x) \wedge \text{on-top}(x, f)$ $\wedge \text{on-top}(\text{self}, f) \wedge \text{dir-right-of}(x, \text{self})$
	g	
<i>ba_self-in-container-from-between-potential-containers</i>	c	$\exists x. \text{container}(x) \wedge \text{inside-supported}(\text{self}, x)$
	g	$\exists f, x, y. \text{blue-floor}(f) \wedge \text{white}(x) \wedge \text{white}(y) \wedge \text{on-top}(x, f) \wedge \text{on-top}(y, f)$ $\wedge \text{container}(x) \wedge \text{noncontainer}(y) \wedge \text{h-between}(\text{self}, x, y)$
	g	$\exists f, x, y. \text{blue-floor}(f) \wedge \text{white}(x) \wedge \text{white}(y) \wedge \text{on-top}(x, f) \wedge \text{on-top}(y, f)$ $\wedge \text{container}(x) \wedge \text{noncontainer}(y) \wedge \text{h-between}(\text{self}, y, x)$
	g	

Concept Formation


SCAN (DeepMind, 2018)

6 Jun 2018

Published as a conference paper at ICLR 2018

SCAN: LEARNING HIERARCHICAL COMPOSITIONAL VISUAL CONCEPTS

Irina Higgins, Nicolas Sonnerat, Loic Matthey, Arka Pal,
Christopher P Burgess, Matko Bošnjak, Murray Shanahan,
Matthew Botvinick, Demis Hassabis, Alexander Lerchner
DeepMind, London, UK



The next breakthroughs

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

Sequence to Sequence Learning with Neural Networks


Ilya Sutskever



Oriol Vinyals



Quoc V. Le



The next breakthroughs

- Open-source toolkits
 - Stanford Core NLP suite
 - Natural Language Toolkit (nltk.org)
 - Apache's OpenNLP
 - Google SyntaxNet
 - Matthew Honnibal's Spacy.io



The next breakthroughs

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



The next breakthroughs



The next breakthroughs


- National projects that changed the world
 - BRAIN Initiative: Iarpa's Microns (David Markovitz)



The next breakthroughs

- New mathematical models
 - Salvatore Torquato 's hyperuniformity (2003)
 - From bird vision...
 - ...to the large-scale structure of the universe
 - Etaphase (Salvatore Torquato) to manufacture "hyperuniform" photonic circuits, devices that transmit data via light rather than electrons

The colored dots below correspond to the arrangement of green, blue, red, violet and double-type (black) cone photoreceptors in a chicken's retina.



#1 most frequently asked question

- ***Will robots be conscious?***
 - The food industry slaughters 60 billion farmed animals (mammals, birds and fish) every year: why in heaven are we concerned for the consciousness of robots when we are not concerned for the consciousness of mammals and birds, whose brain is so similar to ours?


#1 most frequently asked question

- ***Will machines become conscious?***
- Judea Pearl's causal calculus at UCLA ("A Probabilistic Calculus of Actions", 1994)
- Giulio Tononi ("Consciousness as Integrated Information - A Provisional Manifesto", 2008) at the University of Wisconsin: the amount of integrated information in a neural network increases as larger and larger groups of neurons form and become independent self-standing units.
- At each level of integration there is a probability of error in predicting the future state (the effect of a cause is unpredictable to a lower and lower degree)



#1 most frequently asked question

- ***Will machines become conscious?***
- The state transitions can be mathematically modeled using Markov chains
- There is a maximum scale at which system states reach the greatest causal power, and this could be the state at which desires and beliefs manifest themselves.
- This is the state at which predicting future states becomes the most reliable.
- If this is consciousness, then yes machines will become conscious.



Φ

#1 most frequently asked question

- ***Will machines become conscious?***
- If this is consciousness, then yes machines will become conscious...
- ... just like the 60 billion mammals, birds and fish) that we slaughter every year



Beware of...

- Smart appliances that try to understand your habits and customize your experience



Beware of...

- Chatbots
- Intelligent agents



Wei Zexi's parents (2016)



Two Google Homes are arguing on Twitch

<https://www.twitch.tv/seebotschat>



Beware of...

- Driver-less cars = brain-less cars

Ford to ship self-driving cars without steering wheels, brake or gas pedals by 2021



August 2016

Tesla driver dies in fatal crash while using autopilot mode



The Age and AI and...

ROGER SCHANK

©2017 Roger Schank

**THE FRAUDULENT CLAIMS MADE BY
IBM ABOUT WATSON AND AI.**

They are not doing "cognitive computing" no
matter how many times they say they are

The Age of A.I. and..

- Fake news



Protecting from A.I.

- Future of Humanity Institute (Oxford)
- Center for Science and Existential Risk (Cambridge)
- Future of Life Institute (MIT)
- Machine Intelligence Research Institute (San Francisco)
- OpenAI (San Francisco)

Machine Ethics

- Eliezer Yudkowsky: Friendly A.I. (2001)
- Steve Omohundro: Transhuman Morality (2008)
- Wendell Wallach, Colin Allen, and Iva Smit: artificial moral agents



Creating Friendly AI 1.0:
The Analysis and Design of
Benevolent Goal Architectures

Eliezer Yudkowsky



The principle of complementarity in A.I.

- For every activity that gets automated, someone will automate the counter-activity
- Eg: If someone creates an automated recruiting assistant to recruit the best engineer, someone will create an automated resume writer to write the best resume that fools the recruiting assistant
- Eg: If someone creates a travel assistant that picks the best combination of hotels, car rentals, trains, etc, someone will write a travel-agency assistant to offer a package that fools the travel assistant



Niels Bohr

Jobs

**MIT
Technology
Review**

October 2017

**Automation could hit
India particularly hard
because much of its
high-tech economy
involves relatively
routine work that is
prime for computers to
take over.**

Jobs

- Boom of jobs
- Very competitive talent market
- Skyrocketing salaries in A.I.
- Not enough A.I. engineers in the pipeline
- Millions of job openings being created within the next decade
- The age of hyper-employment (negative unemployment)?
- The strategic good is not software (intelligent or not) but data

ITIF | INFORMATION TECHNOLOGY
& INNOVATION FOUNDATION

**False Alarmism: Technological
Disruption and the U.S. Labor Market,
1850–2015**

[Robert D. Atkinson](#) and [John Wu](#) May 8, 2017

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(x): Energy function

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

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

Jobs in the Age of Robots



DEEP
LEARNING
INSTITUTE

MIT
Technology
Review

August 8, 2017

**Andrew Ng's Next Trick:
Training a Million AI
Experts**



Jobs

- New kinds of jobs
- The Chief A.I. Officer



Harvard
Business
Review

LEADERSHIP

Hiring Your First Chief AI Officer


by Andrew Ng

NOVEMBER 11, 2016



Jobs

- The Chief A.I. Officer (CAIO)
 - Good technical understanding of AI and Big Data: what can A.I. really do?
 - Communication and management skills
 - Charisma to attract AI talent in a very competitive job market
 - Someone who **doesn't tweet** all the time but actually understands the issue



Jobs

- Why the protagonists of deep learning tend to be foreign-born?



Donald J. Trump

@McDonaldJTrump



 Follow

My grandparents didn't come to America all the way from Germany to see it get taken over by immigrants. Not on my watch.

[#TrumpForPresident](#)

RETWEETS

103

LIKES


115



6:31 PM - 6 Sep 2015

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?

- What will be the effect of hyper-technology on employment? Hyper-employment.


Do robots steal our jobs?

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



Jobs

- Maybe it is time to rethink the economy: why not universal basic income?



The AI Brain Drain

- Google: Andrew Ng (2011 from Stanford), Geoff Hinton's team (2013 from Toronto)
- Facebook: Yann LeCun (NYU)
- Microsoft: Geoffrey Zweig (Berkeley)
- OpenAI: Sutskever (Toronto via Google)
- JP Morgan: Afsheen Afshar (Stanford), Geoffrey Zweig (Berkeley via Microsoft)
- Tesla: Karpathy

The AI Brain Drain

- Good or bad?
- 1950s: Semiconductor Brain Drain as leading scientists joined industrial R&D labs
- One of them was William Shockley (Beckman Instruments, 1956)
- who in turn hired some of the best graduate students in the nation
- who in turn founded Fairchild and hired the best graduate students in the nation
- who in turn created Silicon Valley as it is today.

What A.I. will not give us

- Free humans to pursue arts and sciences
- Accelerate new discoveries in science and tech
- Increased efficiencies
- Unemployment
- Machine ethics

What A.I. will give us

- More automation
- Faster, cheaper, more accurate services
- A whole new class of jobs
- DIY customer support
- More rules and regulations to allow dumb machines to work in the human world
- A dumber human brain
- Bigger inequality gap and bigger gap between developed and underdeveloped countries
- Fake news
- Cyber-terrorism
- Cyborgs

Don't be afraid of the robot

- AI systems "don't have the intentionality, really, even of an insect" (Rodney Brooks)

You Should Be More Scared of Smartphones Than Robots

Gill Pratt, one of the top roboticists in the world, says people should not fear robots because they have legs. He says you should be more worried about all the data you carry around on your cellphone, and who has access to it.

Don't be afraid of the robot


- We need AI soon.
- 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?



The world we live in...



You are already surrounded by robots...



The world we live in...



The world we live in...



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.



The robots are coming!

Journalist: Are you afraid of A.I.?

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



Some Conclusions

A.I. is not magic

A.I. is just computational math applied to automation

The border between A.I. and plain Computer Science is blurred

When does Computer Science become A.I.? It depends on whether the term is fashionable or not.

The theory of the singularity belongs to the realm of religion, not science

Automation will cause hyper-employment

The End (for now)

*"A man provided with paper, pencil, and rubber
(and subject to strict discipline) is in effect a
universal machine"*



(Alan Turing, 1948)

*"Computers are useless: they can only give
you answers"*

(Pablo Picasso)

Contact: p@scaruffi.com

Index: <http://www.scaruffi.com/singular>



Next...

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

