

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 11

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

Modern A.I.

The Age of Deep Learning

Deep Learning




Deep Learning

Training multilayered networks

1989: Yann LeCun's LeNet-1

2006: Geoffrey Hinton's Deep Belief Networks


2007: Yoshua Bengio's Stacked Autoencoders



Deep Learning





2006: Hinton's Deep Belief Networks



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

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

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




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


2011: Yoshua Bengio's ReLu



Activation Functions



Deep Learning



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



Trivia: The Stars of Deep Learning

Jitendra Malik: India

Dong Yu: China

Oriol Vinyals: Spain

Ian Goodfellow: Canada

Karen Simonyan: Armenia

Andrew Zisserman: Britain

Christian Szegedy: Germany

Aja Huang: China

Kaiming He: China

Jian Sun: China

Andrej Karpathy: Slovakia/ Canada


Pieter Abbeel: Belgium

Ronan Collobert: France

Yangqing Jia: China

Rajat Monga: India

Richard Socher: Germany



Trivia: The Stars of Deep Learning

Elon Musk: South Africa

Jen-Hsun Huang: Taiwan



Trivia: The Women of Deep Learning

Hava Siegelmann, Univ of Massachusetts

Fei-Fei Li, Stanford Univ

Daphne Koller, Calico Labs

Cynthia Breazeal, Jibo

Andrea Frome, Clarifai

Rana el Kaliouby, Affectiva

Hua Wu, Baidu

Angelica Lim, Simon Fraser Univ

Daniela Rus, MIT

Jane Wang, Google DeepMind

Rama Akkiraju, IBM Watson


Suzanne Gildert, KindredAI

Raia Hadsell, DeepMind

Raquel Urtasun, Univ of Toronto

Leslie Kaelbling, MIT

Jeannette Bohg, Max Planck Inst



The Homeland of Deep Learning



Deep Learning

2010: The ImageNet challenge (ILSVRC)

IMGENET



Large Scale Visual Recognition
Challenge 2010



Winner team:



Deep Learning

Annus mirabilis of Deep Learning: 2012

Deep Learning →


Error rates on the ILSVRC-2012 competition

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


much bigger gap

ImageNet: Image Classification Task

Classification Error (%)



Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton




Deep Learning

AlexNet: 8 layers



Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton



Deep Learning

- Andrew Ng's Google Brain (2012): 1.7 billion connections (and 16,000 processors) learn to recognize cats in YouTube videos




2010s



ICDAR
OFFLINE CHINESE
HANDWRITING
RECOGNITION
CONTEST


OCT 2013:

NEAR-HUMAN
PERFORMANCE



JÜRGEN SCHMIDHUBER, IDSIA

USI/SUPSI



Istituto
Dalle Molle
di studi
sull'intelligenza
artificiale

IDSIA

2010s

- 2015: Jun Sun's team at Fujitsu's Chinese laboratories surpasses human performance in handwritten Chinese character recognition



Fujitsu Achieves 96.7% Recognition Rate for Handwritten Chinese Characters Using AI That Mimics the Human Brain

First time ever to be more accurate than human recognition, according to conference


Fujitsu R&D Center Co., Ltd., Fujitsu Laboratories Ltd.

Beijing, China and Kawasaki, Japan, September 17, 2015

Deep Learning


Region-based CNNs (Jitendra Malik's team, 2013)

- Object detection



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

Deep Learning

Annus mirabilis of Deep Learning: 2012

Why did it take until 2012?

Too many parameters to be learned from too few labeled examples

Processing speed wasn't there before Google paid for it

Training datasets weren't big enough earlier


*Thus to the Eastern wealth through Storms we go,
But now, the Cape once doubled, fear no more:
A constant Trade-wind will securely blow,*



*And gently lay us on the Spicy shore.
(John Dryden, Annus Mirabilis, 1667)*

The real heroes of Deep Learning

Nvidia's GPUs



*Deep Learning
is born*

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.




Nvidia Tesla P100 for deep learning

Nvidia's CEO, Jen-Hsun Huang, delivers the first DGX-1 to Elon Musk's OpenAI (2016)


The real heroes of Deep Learning

Google Tensor Processing Unit (2016)

Intel Nervana (2017)



Intel unveils new family of AI chips to take on Nvidia's GPUs



Oct 17, 2017, 2:00pm EDT

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)
- 2000: 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)

Unsung Heroes: The datasets

Videos:

Human3.6m (2014, Romanian Academy)

YouTube-8M (2016, Google)

Kinetics-600 (2017, Oxford)

Human-machine interaction:

SEMAINE (2007, Queen's University Belfast)

Cam3D (2011, Cambridge)

MAHNOB-HCI (2011, ICL)

Facial expressions:

EAGER (2013, Binghamton Univ)

FaceWarehouse (2014, Zhejiang Univ)

CelebA (2015, Hong Kong)

DCU (2014, Dublin City Univ)

Unsung Heroes: The datasets

Scenes:

LSUN (2015, Princeton University)

Human interactions:

DCU (2014, Dublin City Univ)

Machine translation:

UM (2014, Macau Univ)

Tvsub & Mvsub (2016, Dublin City Univ)

Reading comprehension:

SNLI (2015, Stanford)

SQuAD (2016, Stanford)

MARCO (2016, Microsoft)

WebText (2018, OpenAI)

Sentiment analysis:

LMRD (2011, Stanford)

SST (2013, Stanford)


Unsung Heroes: The Platforms

- Open-source platforms for deep learning
 - Torch (Ronan Collobert @ IDIAP, Switzerland, 2002)
 - Theano (Bengio's group @ Univ of Montreal, Canada, 2010)
 - Caffe (Yangqing Jia @ UC Berkeley, 2013)
 - TensorFlow (Rajat Monga @ Google, 2015)



Deep Learning

- Scales well with memory/data/computation
- Solves the representation learning problem
- State-of-the-art for images, audio, language, ...





R. Gershick

2015

Deep Learning

Deep Learning mimics the workings of the brain: the audiovisual cortex works in multiple hierarchical stages




[Gallant & Van Essen]

[picture from Simon Thorpe]

Artificial Intelligence

Genealogy of Intelligent Machines




Progress in Neural Networks

- Physics: Hopfield's recurrent neural networks mimic annealing
- Neuroscience: Fukuyama's convolutional neural networks mimic the cat's visual system
- Rediscovery of old mathematics (mainly optimization and control methods)

2010s

2014: Microsoft's Skype Translator, a real-time spoken language translation system



2014: IBM's TrueNorth processor

Lawrence Livermore National Laboratory and IBM Collaborate to Build Brain-Inspired Supercomputer

LIVERMORE, Calif. and ARMONK, N.Y. - 29 Mar 2016

The 2010s

- Widespread adoption of image recognition
 - Facebook's DeepFace (2015)
 - Deep Face (2014) claims 97% accuracy... but mysteriously disappears
 - Microsoft's CaptionBot (2016)



CaptionBot



Deep Face

2014

claim to have close to human-level performance



2010s

- Facebook (2010): face recognition
- FindFace (2016): identify the pictures of strangers
- Face++: pay with your face



The 2010s

- Face-reading algorithms: detecting human emotion (eg Emotient)

Apple buys artificial intelligence software that can tell people's emotions from their faces

EMOTIENT™



The 2010s

- Gadgets for
 - Speech Recognition
 - Image Recognition



State of the Art in Deep Learning

- Perceptron
- CNN = Convolutional Neural Networks (face recognition, object identification, disease detection)
- RNN/ LSTM = Recurrent Neural Networks (speech recognition, translation, caption generation)
- GAN = Generative Adversarial Networks (image generation, style transfer)
- DRL = Deep Reinforcement Learning (robot training, game playing)

Evolution of Neural Networks

Perceptrons

1960s

Reinforcement
Learning

1970s

Recurrent
Neural
Networks

1980s

Convolutional
Neural
Networks

1990s

Deep
Learning

2000s

Generative
Adversarial
Networks


2010s

Foundation
Models


2020s

Evolution of Neural Networks


1960's



1980's



2000's



Not much has changed...

Table of Contents

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

1. Very Deep Learning
2. Reinforcement Learning
3. Recurrent Neural Nets
4. Generative Adversarial Networks
5. Recursive Cortical Networks
6. Capsule Nets
7. Development Platforms
8. Automatic Machine Learning
9. Explainable Deep Networks
10. Machine Creativity

Very Deep Learning

Most influential architectures:



- AlexNet (2012)
- R-CNN (2013)
- VGGNet or VGG-16 (2014)
- GoogleNet or Inception (2014)
- DeepFace (2014)
- ResNet (2015)
- Conditional networks (2016)
- Inception v4 and Inception-ResNet (2016)
- SqueezeNet (2016)
- ResNeXt (2017)

Very Deep Learning

Most influential architectures:

- AlexNet (2012)

ImageNet Classification with Deep Convolutional Neural Networks 2012




Very Deep Learning

Most influential architectures:

- R-CNN (2013) – Ross Girshick and Jitendra Malik

Deep Learning for Object Detection within Images

R-CNN (2013)



Ross Girshick



Jitendra Malik

Berkeley
UNIVERSITY OF CALIFORNIA

Very Deep Learning


Most influential architectures:

- VGGNet or VGG-16 (2014)
Karen Simonyan & Andrew Zisserman

Published as a conference paper at ICLR 2015

VERY DEEP CONVOLUTIONAL NETWORKS FOR LARGE-SCALE IMAGE RECOGNITION

Karen Simonyan* & Andrew Zisserman*
Visual Geometry Group, University of Oxford




Very Deep Learning

Most influential architectures:


- GoogleNet or Inception (2014) – Christian Szegedy
 - The first layers is a “depthwise separable convolution” (Laurent Sifre , 2013)

Christian Szegedy



GoogLeNet architecture with nine inception modules


<http://joelouismarino.github.io>



Standard way to model a neuron

$$f(x) = \tanh(x) \quad \text{or} \quad f(x) = (1 + e^{-x})^{-1}$$


Very slow to train



Non-saturating nonlinearity (RELU)

$$f(x) = \max(0, x)$$

Quick to train




Very Deep Learning

Most influential architectures:

•GoogLeNet (Inception)

9 **Inception** modules



Very Deep Learning

Most influential architectures:

- DeepFace (2014) – Lior Wolf

DeepFace: Closing the Gap to Human-Level Performance in Face Verification

Yaniv Taigman


Ming Yang

Marc' Aurelio Ranzato

Lior Wolf

Facebook AI Research
Menlo Park, CA, USA

Tel Aviv University
Tel Aviv, Israel



Outline of the *DeepFace* architecture.

Very Deep Learning

Most influential architectures:

- Network-in-Network/ NiN (2014) – Shuicheng Yan




Figure 2: The overall structure of Network In Network

Very Deep Learning

Most influential architectures:

- ResNet (2015) – Kaiming He, Jian Sun – 152 layers
- Another “network-in-network architectures” like Inception

Dec 2015

Deep Residual Learning for Image Recognition

Kaiming He




Xiangyu Zhang



Microsoft

Shaoqing Ren

Jian Sun



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.

Very Deep Learning


Most influential architectures:

- ResNet with identity mappings (2016) – Kaiming He – 1001 layers

$$y_l = h(x_l) + \mathcal{F}(x_l, \mathcal{W}_l),$$

$$x_{l+1} = f(y_l),$$


f is ReLU



identity mapping:

$$h(x_l) = x_l$$

f is also an identity mapping: $x_{l+1} \equiv y_l$



$$x_L = x_l + \sum_{i=l}^{L-1} \mathcal{F}(x_i, \mathcal{W}_i)$$

for any deeper unit L
and any shallower unit l

Very Deep Learning

2016: Microsoft's conditional networks – Antonio Criminisi

Decision Forests, Convolutional Networks and the Models in-Between

Microsoft Research Technical Report 2015-58

Yani Ioannou¹

Duncan Robertson²

Darko Zikic²

Peter Kontschieder²


Jamie Shotton²

Matthew Brown³



Antonio Criminisi²

Our automatically-optimized conditional architecture is ~ 5 times faster and ~ 6 times smaller than NiN, with same accuracy.



Very Deep Learning

Object detection (improving R-CNN):

- SPPnet (2014)
- Fast RCNN (2014)

Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition ECCV 2014



fully-connected layers (fc_6, fc_7)





Figure 3: A network structure with a **spatial pyramid pooling layer**. Here 256 is the filter number of the $conv_5$ layer, and $conv_5$ is the last convolutional layer.

Sep 2015

Fast R-CNN

Ross Girshick
Microsoft Research



Very Deep Learning

Object detection (improving R-CNN):

- Faster RCNN (2015)
- Yolo (2016)

Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks

Advances in Neural Information Processing Systems 28 (NIPS 2015)

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun



May 2016

You Only Look Once: Unified, Real-Time Object Detection

Joseph Redmon*, Santosh Divvala*[†], Ross Girshick[‡], Ali Farhadi*[†]
University of Washington*, Allen Institute for AI[†], Facebook AI Research[‡]






Figure 2: Faster R-CNN is a single, unified network for object detection. The RPN module serves as the 'attention' of this unified network.

Very Deep Learning

Batch Normalization (Sergey Ioffe and Christian Szegedy at Google, 2015)

A way to accelerate the training of a deep network



Batch Normalization

Sergey Ioffe Christian Szegedy



Input: Values of x over a mini-batch: $\mathcal{B} = \{x_{1...m}\}$;

Parameters to be learned: γ, β

Output: $\{y_i = \text{BN}_{\gamma, \beta}(x_i)\}$

$$\mu_{\mathcal{B}} \leftarrow \frac{1}{m} \sum_{i=1}^m x_i \quad // \text{ mini-batch mean}$$

$$\sigma_{\mathcal{B}}^2 \leftarrow \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2 \quad // \text{ mini-batch variance}$$

$$\hat{x}_i \leftarrow \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}} \quad // \text{ normalize}$$

$$y_i \leftarrow \gamma \hat{x}_i + \beta \equiv \text{BN}_{\gamma, \beta}(x_i) \quad // \text{ scale and shift}$$

Christian Szegedy




Aug 2016


Most influential architectures:

- Inception-v4 (2016)
- Inception-ResNet (2016)

Inception-ResNet



Inception-v4



Very Deep Learning

Most influential architectures:

- Inception v1 - Introduced inception blocks and one “depthwise separable convolution”
- Inception v2 - Added batch normalization (Dec 2015)
- Inception v3 - Factorized the inception blocks further
- Inception v4 - Adds residual connections (2016)



Very Deep Learning

Most influential architectures:

- SqueezeNet (2016) – Kurt Keutzer

SQUEEZENET: ALEXNET-LEVEL ACCURACY WITH
50X FEWER PARAMETERS AND <0.5MB MODEL SIZE

Kurt Keutzer



Forrest Iandola



Song Han




Nov 2016

Very Deep Learning

Most influential architectures:

- Xception (2016, Francois Chollet) replaces all Inception modules with “depthwise separable convolutions”



Very Deep Learning

Most influential architectures:

- ResNeXt (2017) – Ross Girshick, Kaiming He

Apr 2017

Aggregated Residual Transformations for Deep Neural Networks

Saining Xie¹ Ross Girshick² Piotr Dollár² Zhuowen Tu¹ Kaiming He²
¹UC San Diego ²Facebook AI Research



Aggregated Transformation

$$\mathcal{F}(x) = \sum_{i=1}^C \mathcal{T}_i(x),$$

C=Cardinality



Figure 1. Left: A block of ResNet [14]. Right: A block of ResNeXt with cardinality = 32, with roughly the same complexity. A layer is shown as (# in channels, filter size, # out channels).

Very Deep Learning

Most influential architectures:

- R-net (2017) – Microsoft China
 - Bidirectional RNN for reading comprehension (Mikolov et al., 2010) (Wang & Jiang, 2016)



Furu Wei




Figure 1: R-NET structure overview.

GeekWire

January 15, 2018

Microsoft and Alibaba AI programs beat humans in Stanford reading comprehension test for 1st time

Very Deep Learning

Most influential architectures:

- FractalNet (Chicago, 2017)
- PolyNet (Hong Kong, 2017)

Published as a conference paper at ICLR 2017

FRACTALNET:

ULTRA-DEEP NEURAL NETWORKS WITHOUT RESIDUALS

Gustav Larsson
University of Chicago

Michael Maire
TTI Chicago

Gregory Shakhnarovich
TTI Chicago




Figure 1: **Fractal architecture.** Left: A simple expansion rule generates a fractal architecture with

17 Jul 2017

PolyNet: A Pursuit of Structural Diversity in Very Deep Networks

Xingcheng Zhang*

Zhizhong Li*

Chen Change Loy

Dahua Lin



香港中文大學
The Chinese University of Hong Kong





Figure 4: Examples of PolyInception structures.

Very Deep Learning

Most influential architectures:

- Tsung-wei Ke's Multigrid networks (2017)

Multigrid Neural Architectures



Very Deep Learning

Most influential architectures:

- Gao Wang's MSDNet (2018)



Published as a conference paper at ICLR 2018

MULTI-SCALE DENSE NETWORKS FOR RESOURCE EFFICIENT IMAGE CLASSIFICATION

Gao Huang **Danlu Chen** **Tianhong Li** **Felix Wu**
Cornell University Fudan University Tsinghua University Cornell University
Laurens van der Maaten **Kilian Weinberger**
Facebook AI Research Cornell University

Densely Connected Convolutional Networks

Gao Huang* Zhuang Liu* Laurens van der Maaten Kilian Q. Weinberger
Cornell University Tsinghua University Facebook AI Research Cornell University



Very Deep Learning

The Power of Depth

Ronen Eldan and Ohad Shamir prove that "depth can be exponentially more valuable than width" (2016).

The Power of Depth for Feedforward Neural Networks

May 2016



Ronen Eldan

Weizmann Institute of Science

Ohad Shamir

Weizmann Institute of Science

all known activation functions, including rectified linear units, sigmoids and thresholds, and formally demonstrates that depth – even if increased by 1 – can be exponentially more valuable than width for standard feedforward neural networks. Moreover, compared to related results in the context of Boolean





Huang's Law

"Nvidia's GPUs today are 25 times faster than 5 years ago" (Nvidia CEO Jensen Huang, April 2018)

2012: It took six days on two Nvidia GTX 580s to train AlexNet

2018: It takes 18 minutes with the Nvidia DGX-2



2 Apr 2018

Move Over, Moore's Law: Make Way for Huang's Law

Graphics processors are on a supercharged development path that eclipses Moore's Law, says Nvidia's Jensen Huang

Huang's Law


2016: Nvidia's Volta

2016: Nvidia's DGX-1: 8 GPUs + software to train neural networks



2016

Nvidia Tesla P100 for deep learning



640 TENSOR CORES

An Exponential Leap in
Performance

Volta delivers over 100 Teraflops per second

Huang's Law


Google's TPU (programmed via TensorFlow and offered via Google Cloud)

Internal search ranking model training (2018):

~9 hours on 1/4 pod vs. ~132 hours on 275 high end CPU machines
(14.2 times faster)

Internal image model training:

~22 hours on 1/4 pod vs. ~216 hours on previous production setup
(9 times faster)



A.I. for Makers

Let's start with the easy part: how can a maker build A.I. systems?

Luckily, this is increasingly easy.

A.I. for Makers

Arduino & Raspberry

- Google TensorFlow available for Raspberry Pi
- Dobot's Arduino A.I. Suite



Raspberry Pi



Arduino Artificial Intelligence Suite

LEARNING THROUGH PLAYING

Entry-level Learning Kit for Artificial Intelligence

[Watch Video](#)

[Contact Us](#)



A.I. for Makers

Google Coral

**GOOGLE LAUNCHES AI
PLATFORM THAT LOOKS
REMARKABLY LIKE A
RASPBERRY PI**

by: **Brian Benchoff**

80 Comments

f t 8+

March 5, 2019



Google Coral Dev Board

A.I. for Makers

Nvidia Jetson




AI FOR MAKERS – LEARN
WITH JETBOT

NVIDIA WEBINAR

Nvidia announces \$99 AI
computer for developers,
makers, and researchers

Mar 18, 2019,

The Jetson Nano



A.I. for Makers

HiKey
Sipeed
Horned Sungem
Intel

HiKey 970 Development Board



A.I. for Makers

Free resources

- Google's AI Education site <https://ai.google/education>
- Free online book "Neural Networks and Deep Learning"
<http://neuralnetworksanddeeplearning.com/>
- Free Online Course "Creative Applications of Deep Learning with TensorFlow"
<https://www.classcentral.com/course/kadenze-creative-applications-of-deep-learning-with-tensorflow-6679>
- David Silver's course on Reinforcement Learning on YouTube
<https://www.youtube.com/watch?v=2pWv7GOvuf0>

A.I. for Makers

Udemy's 14-hour video course
“Deep Learning with Python”



Udemy


Complete Guide to TensorFlow for Deep Learning with Python

Learn how to use Google's Deep Learning Framework - TensorFlow with Python! Solve problems with cutting edge techniques!

BESTSELLER ★★★★★ 4.5 (10,593 ratings) 58,684 students enrolled

Created by Jose Portilla Last updated 11/2018

English English [Auto-generated], French [Auto-generated], 9 more



A.I. for Makers

The tools are the easy part... the problem is what to do with the tools...


You can do simple things like robots that say “hello” and recognize objects...

... but A.I. wants to do much more

What is A.I. studying now?

Very Deep Learning

ImageNet 2012-15



Binarized Networks

Binarized Networks (2015)

BinaryConnect: Training Deep Neural Networks with binary weights during propagations

Matthieu Courbariaux
École Polytechnique de Montréal

Yoshua Bengio
Université de Montréal, CIFAR Senior Fellow



Binary weights, i.e., weights which are constrained to only two possible values (e.g. -1 or 1), would bring great benefits to specialized DL hardware by replacing many multiply-accumulate op-

Ternary Networks

Ternary Networks (2016)

Nov 2016

Ternary weight networks


Fengfu Li and Bo Zhang

Institute of Applied Math., AMSS, CAS
Beijing, China

Bin Liu

Moshanghua Tech Co., Ltd.
Beijing, China

A. Trends in DNN accuracies on ImageNet




*Ternary ResNet is the first *sparse and very low precision* (2bit) DNN to offer comparable accuracy to ResNet-152

CMU +
Intel
study
(2016)

Reinforcement Learning

Types of Machine Learning



Reinforcement Learning

Supervised learning: learning from a dataset of positive instances (e.g., object recognition).

Unsupervised learning: clustering objects together by similarity (e.g., concept formation)

Reinforcement learning: improving performance by direct interaction with the environment

Reinforcement Learning

Edward Thorndike (1911): "The greater the satisfaction or discomfort, the greater the strengthening or weakening of the bond."



Ivan Pavlov (1926): "The magnitude and timing of the conditioned response changes as a result of the contingency between the conditioned stimulus and the unconditioned stimulus".



Burrhus Skinner (1938): "The only differences between the behavior of rat and man lie in the field of verbal behavior."



Donal Hebb (~1946): Neurons that fire together wire together



Reinforcement Learning

Claude Shannon (1949): use an evaluation function to help a computer learn how to play chess



Marvin Minsky's thesis (1954)



Richard Bellman (1957): dynamic programming/ optimal control



Arthur Samuel (1959): checkers program



Donald Michie (1961): MENACE program



Harry Klopff (1972): neurons actively seek "excitatory" signal and avoid "inhibitory" signals.



Reinforcement Learning

Ian Witten (1977): actor-critic method



John Holland (1977): classifier systems



Richard Sutton (1981): temporal-difference method



Bart Kosko (1986): differential Hebbian learning



Chris Watkins (1989): Q Learning



Reinforcement Learning

Gerald Tesauro (1992): reinforcement learning to play backgammon



Wulfram Gerstner (1996): neuronal dynamics



David Fogel (1996): evolutionary algorithm to learn checkers



...


DeepMind (2013): deep reinforcement learning to play Atari games



DeepMind (2016): deep reinforcement learning for continuous action

Reinforcement Learning

- Unsupervised learning
- A computational approach to goal-directed learning from interaction between an active decision-making agent and its environment
- Learning what to do so as to maximize a reward
- The four pillars of reinforcement learning: a policy, a reward function, a value function, and a model of the environment
- Learning by self-play



Harry Klopff:
“The Hedonistic
Neuron” (1982)



Andrew Barto
Richard Sutton


Reinforcement Learning

- Policy-based RL. Search directly for the policy achieving maximum future reward
- Value-based RL. Estimate the maximum value achievable under any policy
- Model-based RL. Build a transition model of the environment

Reinforcement Learning

- Markov Decision Process:

- Set of states, S
- Set of actions, A
- Reward function, R
- Policy, π
- Value, V



- Action (A) transitions from start state to end state (S) and is assigned a (positive or negative) reward (R)
 - Policy: set of actions
 - Value: set of rewards
- Goal: choose policy that maximizes value

Reinforcement Learning

- Markov Decision Process:
 - Richard Bellman's "value iteration" method (1957)
 - Ronald Howard's "policy iteration" method 1960)

Reinforcement Learning

- Q-learning
- Sarsa
- TD-learning

Reinforcement Learning

Deep Reinforcement Learning

- Applying DL to RL
- Use a deep network to represent value function and/or policy and/or model
- Optimise the value function and/or policy and/or model
- Combines CNN and RL



Deep Reinforcement Learning

Deep Reinforcement Learning

- Deep Q-Networks (DQN) provide a stable solution to deep value-based RL (Volodymyr Mnih, 2013)



Reinforcement Learning in Atari



Playing Atari with Deep Reinforcement Learning


Volodymyr Mnih Koray Kavukcuoglu David Silver Alex Graves Ioannis Antonoglou
Daan Wierstra Martin Riedmiller
DeepMind Technologies

Figure 1. The DQN algorithm is composed of three main components, the Q-network ($Q(s, a; \theta)$) that defines the behavior policy, the target Q-network ($Q(s, a; \theta^-)$) that is used to generate target Q values for the DQN loss term and the replay memory that the agent uses to sample random transitions for training the Q-network.
Source: Arun Nair (2015)

The Limitations of Reinforcement Learning

Andrey Kurenkov

<https://thegradiant.pub/why-rl-is-flawed/>



DQN fails at Montezuma Revenge

DQN

Best linear learner

Deep Reinforcement Learning

Arun Nair's GORILA (2015)


First massively distributed architecture for deep reinforcement learning

A lot faster implementation of DQN

Massively Parallel Methods for Deep Reinforcement Learning

Arun Nair, Praveen Srinivasan, Sam Blackwell, Cagdas Alceick, Rory Fearon, Alessandro De Maria, Vedavyas Panneershelvam, Mustafa Suleyman, Charles Beattie, Stig Petersen, Shane Legg, Volodymyr Mnih, Koray Kavukcuoglu, David Silver

Google DeepMind, London



Source: Arun Nair (2015)

Jul 2015

Four main components:
parallel actors that generate new behavior; parallel learners that are trained from stored experience; a distributed neural network to represent the value function or behavior policy; and a distributed store of experience.

Reinforcement Learning

- Computer go/weiqi
 - 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

2016: Google/DeepMind's AlphaGo beats the weiqi champion Se-dol Lee



 INDEPENDENT

Google AlphaGo computer beats professional at 'world's most complex board game' Go

Milestone in AI research likened to defeat of world chess champion Garry Kasparov in 1997 by IBM's Deep Blue computer



The New York Times

Google's Computer Program Beats Lee Se-dol in Go Tournament

By CHOE SANG-HUN MARCH 15, 2016




Reinforcement Learning

2017: AlphaGo beats the Go champion Ke Jie



Reinforcement Learning


AlphaGo (2016)



Reinforcement Learning

AlphaGo Overview

based on: Silver, D. et al. Nature Vol 529, 2016
copyright: Bob van den Hoek, 2016




© Bob van den Hoek


Reinforcement Learning

AlphaGo architecture Training policy and value networks

Mastering the game of Go with deep neural networks and tree search



neural network architecture



Reinforcement Learning

AlphaGo architecture

Planning with an environment model and Monte Carlo Tree Search





Figure 3 | Monte Carlo tree search in AlphaGo. **a**, Each simulation traverses the tree by selecting the edge with maximum action value Q , plus a bonus $u(P)$ that depends on a stored prior probability P for that edge. **b**, The leaf node may be expanded; the new node is processed once by the policy network p_σ and the output probabilities are stored as prior probabilities P for each action. **c**, At the end of a simulation, the leaf node

is evaluated in two ways: using the value network v_θ ; and by running a rollout to the end of the game with the fast rollout policy p_π , then computing the winner with function r . **d**, Action values Q are updated to track the mean value of all evaluations $r(\cdot)$ and $v_\theta(\cdot)$ in the subtree below that action.

Deep Reinforcement Learning

Fanuc/ Preferred Networks robot (2015)

Toyota's self-teaching car (2016)



Deep Reinforcement Learning

AlphaGo Zero (2017)




nature
International journal of science

April 2017

Mastering the game of Go without human knowledge

David Silver✉, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou,
Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian
Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van
den Driessche, Thore Graepel & Demis Hassal

Nature **550**, 354–359 (19 October 2017)




Deep Reinforcement Learning

AlphaGo Zero (2017)

Playing go without human knowledge



NATURE | VOL 550 | 19 OCTOBER 2017



Deep Reinforcement Learning

AlphaGo Zero (2017)

Playing go without human knowledge




Figure 2 | MCTS in AlphaGo Zero



Deep Reinforcement Learning

AlphaGo Zero (2017)

Playing go without human knowledge




Performance of AlphaGo Zero



Deep Reinforcement Learning

AlphaGo Zero (2017)

- Learns the rules of go in three hours
- Consumes 12 times less power than AlphaGo



3 hours

AlphaGo Zero plays like a human beginner, forgoing long term strategy to focus on greedily capturing as many stones as possible.




70 hours

AlphaGo Zero plays at super-human level. The game is disciplined and involves multiple challenges across the board.

Deep Reinforcement Learning

AlphaGo Zero (2017)

- Combines policy and value into just one neural network
- Monte-Carlo tree search with policy-iteration



Deep Reinforcement Learning

AlphaZero (2017)

- Learns multiple games (eg chess)

Science

Science 07 Dec 2018:

Vol. 362, Issue 6419, pp. 1140-1144



A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play


David Silver^{1,2*†}, Thomas Hubert^{1*}, Julian Schrittwieser^{1*}, Ioannis Antonoglou¹, Matthew Lai¹, Arthur Guez¹, Marc Lanctot¹, Laurent Sifre¹, Dhharshan Kumaran¹, Thore Graepel¹, Timothy Lillicrap¹, Karen Simonyan¹, Demis Hassabis^{1†}

The game of chess is the longest-studied domain in the history of artificial intelligence. The strongest programs are based on a combination of sophisticated search techniques, domain-specific adaptations, and handcrafted evaluation functions that have been refined by human experts over several decades. By contrast, the AlphaGo Zero program recently achieved superhuman performance in the game of Go by reinforcement learning from self-play. In this paper, we generalize this approach into a single AlphaZero algorithm that can achieve superhuman performance in many challenging games. Starting from random play and given no domain knowledge except the game rules, AlphaZero convincingly defeated a world champion program in the games of chess and shogi (Japanese chess), as well as Go.

Deep Reinforcement Learning

AlphaZero (2017)


– Hassabis at MIT (2019)



Deep Reinforcement Learning

AlphaZero (2017)

– Hassabis at MIT (2019)




Deep Reinforcement Learning

Tools

2016 OpenAI Gym, a toolkit for research on reinforcement learning,


2017 OpenAI Roboschool, a software environment to create real-world simulations for training robots.



Deep Reinforcement Learning

Methods

John Schulman's proximal policy optimization (PPO, 2017)



$$L^{CLIP}(\theta) = \hat{E}_t [\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1 - \varepsilon, 1 + \varepsilon)\hat{A}_t)]$$

- θ is the policy parameter
- $\hat{E}_t$ denotes the empirical expectation over timesteps
- r_t is the ratio of the probability under the new and old policies
- $\hat{A}_t$ is the estimated advantage at time t
- ε is a hyperparameter, usually 0.1 or 0.2

Deep Reinforcement Learning

Methods

Deep Mind's Rainbow (2017)



Rainbow: Combining Improvements in Deep Reinforcement Learning

Matteo Hessel
Will Dabney


Joseph Modayil
Dan Horgan

Hado van Hasselt
Bilal Piot

Tom Schaul
Mohammad Azar

Georg Ostrovski
David Silver

Oct 2017



Deep Reinforcement Learning

Challenging Environments

PPO to train a robot in Roboschool environments

E.g. an agent tries to reach a target (the pink sphere), learning to walk, run, turn, recover from being hit by objects, and get up again when knocked over.



Deep Reinforcement Learning



Challenging Environments

PPO to train a robot in Roboschool environments

Agents trained with PPO develop flexible movement policies that let them improvise turns and tilts as they head towards a target location


Controllable, Complicated Robots



Deep Reinforcement Learning

Challenging Environments


PPO to teach complicated, simulated robots to walk, like the 'Atlas' model from Boston Dynamics: the model has 30 distinct joints, versus 17 for the bipedal robot



Deep Reinforcement Learning

Reinforcement Learning: Issues

1. Reinforcement learning is highly unnatural: humans don't need thousands of examples



Deep Reinforcement Learning

Reinforcement Learning: Issues

2. Simpler algorithms can solve the same problems

- Jeff Clune's team at **Uber AI Labs** ("Deep Neuroevolution: Genetic Algorithms Are a Competitive Alternative for Training Deep Neural Networks for Reinforcement Learning", 2017),
- Tim Salimans at **OpenAI** ("Evolution Strategies as a Scalable Alternative to Reinforcement Learning", 2017)
- Aravind Rajeswaran at the **University of Washington** ("Towards Generalization and Simplicity in Continuous Control", 2017)
- Horia Mania and Aurelia Guy at **UC Berkeley** ("Simple Random Search Provides a Competitive Approach to Reinforcement Learning", 2018)


Deep Reinforcement Learning

3. Designing the correct reward function is extremely difficult in the real world
4. Balancing exploration and exploitation ("the multi-armed bandit problem"): DQN could play Atari games so well but not Montezuma Revenge, a game that requires more exploration.

Atari 2600 Games:



Pong,



Breakout,



Space Invaders,



Seaquest, Beam Rider



Deep Reinforcement Learning

Go belongs to the simplest category of A.I. problems:

- deterministic,
- fully observable and known,
- discrete
- static


2018: OpenAI Five learns from scratch the multiplayer videogame Dota 2 (neither fully observable/known nor discrete nor static)

Note: OpenAI Five uses proximal policy optimization

2018: and then applies to robotic hand Dactyl to learn to manipulate cubes

Deep Reinforcement Learning

2018: ... and then applies to robotic hand Dactyl to learn to manipulate cubes



An AI-driven robot hand spent a hundred years teaching itself to rotate a cube


July 30, 2018

A reinforcement-learning algorithm allows Dactyl to learn physical tasks by practicing them in a virtual-reality environment.



Deep Reinforcement Learning

OpenAI Five (2018)



Deep Reinforcement Learning

OpenAI Five (2018)

OpenAI Five trains by playing **180 years** worth of games against itself every day, running on **256 GPUs** and **128,000 CPU cores**, using a separate LSTM for each hero

OPENAI FIVE  OpenAI	
CPU	128,000 <u>preemptible</u> CPU cores on GCP
GPU	256 P100 GPUs on GCP
Experience collected	~180 years per day (~900 years per day counting each hero separately)

Deep Reinforcement Learning

OpenAI Five (2018)

OpenAI Five contest

OpenAI Five

AUG
05

OpenAI Five
Benchmark

2018



Deep Reinforcement Learning

OpenAI Five (2018)
but...

Emergent Tech ▶ Artificial Intelligence

23 Aug 2018

**OpenAI bots smashed in their first
clash against human Dota 2 pros**

AI can react faster than humans, but don't play well
enough to beat the masters yet

Deep Reinforcement Learning

October 2019

OpenAI's robotic arm solves the Rubik's Cube using OpenAI Five's reinforcement learning algorithm + Josh Tobin's Automatic Domain Randomization (ADR) which endlessly generates progressively more difficult environments in simulation. This frees us from having an accurate model of the real world (2017)

ADR generates endlessly more complicated simulation environments which help train the robotic arm.

Similarly Rui Wang's Paired Open-Ended Trailblazer (POET) at Uber endlessly generates increasingly complex learning environments (2019)

Deep Reinforcement Learning

OpenAI's robotic arm solves the Rubik's Cube (2019)

Josh Tobin: Automatic Domain Randomization - ADR (2017)



Solving Rubik's Cube
with a Robot Hand

OCTOBER 16, 2019 • 8 MINUTE READ

Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World

Josh Tobin¹, Rachel Fong², Alex Ray², Jonas Schneider², Wojciech Zaremba², Pieter Abbeel³

Mar 2017



OpenAI

¹OpenAI and UC Berkeley

²OpenAI,

³OpenAI, UC Berkeley



Berkeley
UNIVERSITY OF CALIFORNIA



Deep Reinforcement Learning

Microsoft investment in OpenAI (2019)

**Microsoft to invest
\$1 billion in OpenAI**



Sam Altman, CEO of OpenAI (left), and Microsoft CEO Satya Nadella

Deep Reinforcement Learning

Rui Wang: Paired Open-Ended Trailblazer – POET (Uber, 2019)




**Paired Open-Ended Trailblazer (POET): Endlessly
Generating Increasingly Complex and Diverse
Learning Environments and Their Solutions**


UBER

Feb 2019

Rui Wang Joel Lehman Jeff Clune* Kenneth O. Stanley*
Uber AI Labs
San Francisco, CA 94103




(a) ES-/POET-generated agents attempting gaps




Deep Reinforcement Learning

DeepMind AlphaStar (2018)




Deep Reinforcement Learning

DeepMind AlphaStar (2018)



Deep Reinforcement Learning

DeepMind AlphaStar (2018)



Deep Reinforcement Learning

DeepMind AlphaStar (2018)

AlphaFold
Winner of CASP13 protein folding competition
Structures: Ground truth, Predicted

T0954 / 6CVZ T0965 / 6D2V T0955 / 5W9F



Predicting most accurate structure for **25 out of 43 proteins**
(compared to three for runner-up)



CENTER FOR
**Brains
Minds+
Machines**

March 20, 2019

The Power of
Self-Learning Systems

Demis Hassabis
DeepMind

Deep Reinforcement Learning

DeepMind AlphaStar (2018)

Many key challenges remain

Unsupervised Learning

Memory and one-shot learning

Imagination-based Planning with
Generative Models

Learning Abstract Concepts

Transfer Learning

Language understanding



CENTER FOR
**Brains
Minds+
Machines**

March 20, 2019

The Power of
Self-Learning Systems

Demis Hassabis

DeepMind

Deep Reinforcement Learning

DeepMind AlphaStar plays StarCraft (2019)

nature

10 October 2019

Grandmaster level in StarCraft II using multi-agent reinforcement learning

<https://doi.org/10.1038/s41586-019-1724-z>

Received: 30 August 2019

Accepted:

Published online: 30 October 2019

Oriol Vinyals^{1,2,*}, Igor Babuschkin^{1,2}, Wojciech M. Czarnecki^{1,2}, Michaël Mathieu^{1,2}, Andrew Dudzik^{1,2}, Junyoung Chung^{1,2}, David H. Choi^{1,2}, Richard Powell^{1,2}, Timo Ewalds^{1,2}, Petko Georgiev^{1,2}, Junhyuk Oh^{1,2}, Dan Horgan^{1,2}, Manuel Kroiss^{1,2}, Ivo Danihelka^{1,2}, Aja Huang^{1,2}, Laurent Sifre^{1,2}, Trevor Cai^{1,2}, John P. Agapiou^{1,2}, Max Jaderberg¹, Alexander S. Vezhnevets¹, Rémi Leblond¹, Tobias Pohlen¹, Valentin Dalibard¹, David Budden¹, Yury Sulsky¹, James Molloy¹, Tom L. Paine¹, Caglar Gulcehre¹, Ziyu Wang¹, Tobias Pfaff¹, Yuhuai Wu¹, Roman Ring¹, Dani Yogatama¹, Dario Wünsch², Katrina McKinney¹, Oliver Smith¹, Tom Schaul¹, Timothy Lillicrap¹, Koray Kavukcuoglu¹, Demis Hassabis¹, Chris Apps^{1,2} & David Silver^{1,2,*}

THE VERGE

Oct 30, 2019

DeepMind's StarCraft 2 AI is now better than 99.8 percent of all human players



live from DeepMind in London

Deep Reinforcement Learning

Robot training

DQN works well for videogames, which have a discrete action space, but it cannot be used for the continuous spaces of robot movement.

DDPG for continuous action spaces (DeepMind, 2016)

Published as a conference paper at ICLR 2016

CONTINUOUS CONTROL WITH DEEP REINFORCEMENT LEARNING

Timothy P. Lillicrap*, Jonathan J. Hunt*, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver & Daan Wierstra
Google Deepmind



Lillicrap



Hunt



David Silver & Daan Wierstra

Update critic by minimizing the loss: $L = \frac{1}{N} \sum_i (y_i - Q(s_i, a_i | \theta^Q))^2$
Update the actor policy using the sampled policy gradient:

$$\nabla_{\theta^\mu} J \approx \frac{1}{N} \sum_i \nabla_a Q(s, a | \theta^Q)|_{s=s_i, a=\mu(s_i)} \nabla_{\theta^\mu} \mu(s | \theta^\mu)|_{s_i}$$

Deep Reinforcement Learning

Robot training

Train the visual system and the motor system at the same time
The visual system adapts to the motor system, i.e. to the goal
Turning policy search into supervised learning



Journal of Machine Learning Research 17 (2016) 1-40

End-to-End Training of Deep Visuomotor Policies

Sergey Levine[†]
Chelsea Finn[†]
Trevor Darrell
Pieter Abbeel

Division of Computer Science
University of California
Berkeley, CA 94720-1776, USA

[†]These authors contributed equally.



Deep Reinforcement Learning


Robot vision and manipulation

Kalashnikov & Levine (UC Berkeley + Google, 2018)

2nd Conference on Robot Learning (CoRL 2018), Zürich, Switzerland.

Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation

Dmitry Kalashnikov¹, Alex Irpan¹, Peter Pastor², Julian Ibarz¹,
Alexander Herzog², Eric Jang¹, Deirdre Quillen³, Ethan Holly¹,
Mrinal Kalakrishnan², Vincent Vanhoucke¹, Sergey Levine^{1,3}



Deep Reinforcement Learning

Robot training


Reinforcement Learning in Robotics: A Survey

Jens Kober^{*†}



(a) OBELIX robot

J. Andrew Bagnell[‡]



(b) Zebra Zero robot

Jan Peters^{§¶}



Figure 5: Boston Dynamics



(c) Autonomous helicopter



(d) Sarcos humanoid

Deep Reinforcement Learning

Chemical industry: optimizing chemical reactions

Optimizing Chemical Reactions with Deep Reinforcement Learning


Zhenpeng Zhou,[†] Xiaocheng Li,[‡] and Richard N. Zare^{*,†}

[†]Department of Chemistry, Stanford University, Stanford, California 94305,

[‡]Department of Management Science and Engineering, Stanford University,



Scheme 1. Visualization of the DRO Model Unrolled over Three Time Steps^a





Deep Reinforcement Learning

Resource management (like Skynet in “Terminator”!)

Resource Management with Deep Reinforcement Learning

HotNets-XV, November 09-10, 2016, Atlanta, GA, USA

Hongzi Mao^{*}, Mohammad Alizadeh^{*}, Ishai Menache[†], Srikanth Kandula[†]
Massachusetts Institute of Technology^{*}, Microsoft Research[†]



Deep Reinforcement Learning

Multi-agent reinforcement learning

E.g. OpenAI Five: Dota 2 AI agents are trained to coordinate with each other to compete against humans

UC Berkeley: Rlib (2018)



Scaling Multi-Agent Reinforcement Learning

Eric Liang and Richard Liaw

Dec 12, 2018


RLlib: Abstractions for Distributed Reinforcement Learning

29 Jun 2018

Eric Liang, Richard Liaw, Philipp Moritz, Robert Nishihara, Roy Fox, Ken Goldberg, Joseph E. Gonzalez, Michael I. Jordan, Ion Stoica

RLlib: Scalable Reinforcement Learning

RLlib is an open-source library for reinforcement learning that offers both a collection of reference algorithms and scalable primitives for composing new ones.



Deep Reinforcement Learning

Multi-agent reinforcement learning

- Most successes of RL are in single-agent scenarios
- Multi-agent (multiple players interacting)
 - imperfect information
 - joint action space
 - large state space
 - delayed credit assignment

Note: the probability of taking a gradient step in the correct direction decreases

exponentially with the number of agents (Lowe et al., 2017)


Note: the “LazyAgent Problem” (Sunehag et al., 2017)

Deep Reinforcement Learning

Multi-agent reinforcement learning

– Challenge: SC2LE1 (StarCraft II Learning Environment)

- a multi-agent reinforcement-learning environment based on the StarCraft II video game (DeepMind, Oriol Vinyals, 2017)



Deep Reinforcement Learning

Multi-agent reinforcement learning

Decentralized Actor Centralized Critic
(Ryan Lowe & Yi Wu, OpenAI, 2017)

**Multi-Agent Actor-Critic for Mixed
Cooperative-Competitive Environments**

16 Jan 2018

Ryan Lowe*
McGill University
OpenAI

Yi Wu*
UC Berkeley

Aviv Tamar
UC Berkeley

Jean Harb
McGill University
OpenAI

Pieter Abbeel
UC Berkeley
OpenAI

Igor Mordatch
OpenAI








Figure 2: Illustrations of the experimental environment and some tasks we consider, including a) *Cooperative Communication* b) *Predator-Prey* c) *Cooperative Navigation* d) *Physical Deception*.

Deep Reinforcement Learning

Multi-agent reinforcement learning

Counterfactual Multi-agent or COMA
(Jakob Foerster & Gregory Farquhar,
Oxford Univ, 2017)


QMIX (Oxford, 2018)



Information Flow in COMA

QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning

Tabish Rashid^{*1} Mikayel Samvelyan^{*2} Christian Schroeder de Witt¹
Gregory Farquhar¹ Jakob Foerster¹ Shimon Whiteson¹



6 Jun 2018




Figure 2. (a) Mixing network structure. In red are the hypernetworks that produce the weights and biases for mixing network layers shown in blue. (b) The overall QMIX architecture. (c) Agent network structure. Best viewed in colour.

Deep Reinforcement Learning

Multi-agent reinforcement learning

Social Influence as Intrinsic Motivation (MIT
+ DeepMind, 2019)

8 Feb 2019

Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning

Natasha Jaques^{1,2} Angeliki Lazaridou² Edward Hughes² Caglar Gulcehre² Pedro A. Ortega² DJ Strouse³
Joel Z. Leibo² Nando de Freitas²




Figure 6: The Model of Other Agents (MOA) architecture learns both an RL policy π_e , and a supervised model that predicts the actions of other agents, a_{t+1} .





Figure 3: The communication model has two heads, which learn the environment policy, π_e , and a policy for emitting communication symbols, π_m . Other agents' communication messages m_{t-1} are input to the LSTM.

Deep Reinforcement Learning

DRL for continuous action spaces
(e.g. robot behavior)

D4PG (Timothy Lillicrap,
DeepMind, 2018)




Published as a conference paper at ICLR 2018

DISTRIBUTED DISTRIBUTIONAL DETERMINISTIC POLICY GRADIENTS

Gabriel Barth-Maron*, Matthew W. Hoffman*, David Budden, Will Dabney,
Dan Horgan, Dhruva TB, Alistair Muldal, Nicolas Heess, Timothy Lillicrap
DeepMind



Figure 9: Control Suite domains used for benchmarking. *Top*: acrobot, cartpole, cheetah, finger, fish, hopper. *Bottom*: humanoid, manipulator, pendulum, reacher, swimmer6, swimmer15, walker.



Deep Reinforcement Learning

Model-based DRL

PETS (Kurtland Chua, UC Berkeley, 2019)





Figure 1: Our method (PE-TS): **Model:** Our probabilistic ensemble (PE) dynamics model is shown as an ensemble of two bootstraps (bootstrap disagreement far from data captures epistemic uncertainty: our subjective uncertainty due to a lack of data), each a probabilistic neural network that captures aleatoric uncertainty (inherent variance of the observed data). **Propagation:** Our trajectory sampling (TS) propagation technique uses our dynamics model to re-sample each particle (with associated bootstrap) according to its probabilistic prediction at each point in time, up until horizon T . **Planning:** At each time step, our MPC algorithm computes an optimal action sequence, applies the first action in the sequence, and repeats until the task-horizon.







Deep Reinforcement Learning

Model-based DRL

PETS (Kurtland Chua, UC Berkeley, 2019)



Figure 2: Tasks evaluated.



Deep Reinforcement Learning

Model-based DRL

Google Brain's SimPLe (Finn, Levine, etc 2019)





Google Brain

Mar 2019

Model Based Reinforcement Learning for Atari

Lukasz Kaiser^{*1} Mohammad Babaeizadeh^{*23} Piotr Miłoś^{*45} Błażej Osiniński^{*453} Roy H Campbell²
Konrad Czechowski⁴ Dumitru Erhan¹ Chelsea Finn¹ Piotr Kozakowski⁴ Sergey Levine¹ Ryan Sepassi¹
George Tucker¹ Henryk Michalewski⁴⁵



Deep Reinforcement Learning

Model-based DRL

Google Brain's SimPLe (Finn, Levine, etc, 2019)




Google Brain

Mar 2019

Model Based Reinforcement Learning for Atari

Lukasz Kaiser^{*1} Mohammad Babaeizadeh^{*2,3} Piotr Miłoś^{*4,5} Błażej Osiniński^{*4,5,3} Roy H Campbell²
Konrad Czechowski⁴ Dumitru Erhan¹ Chelsea Finn¹ Piotr Kozakowski⁴ Sergey Levine¹ Ryan Sepassi¹
George Tucker¹ Henryk Michalewski^{4,5}



Deep Reinforcement Learning

Model-based DRL

Performance of Simple




Figure 3: Comparison with Rainbow. Each bar illustrates the number of interactions with environment required by Rainbow to achieve the same score as our method (Simple).




Figure 4: Comparison with PPO

Deep Reinforcement Learning

Model-based DRL


DeepMind's Deep Planning Network (PlaNet, Danijar Hafner & Timothy Lillicrap, 2019)

Learning Latent Dynamics for Planning from Pixels




UNIVERSITY OF
TORONTO

Deterministic




(a) Deterministic model (RNN)

Stochastic




(b) Stochastic model (SSM)




(c) Recurrent state-space model (RSSM)


Deterministic
and
stochastic



(a) Standard variational bound



(b) Observation overshooting





(c) Latent overshooting

Deep Reinforcement Learning

Model-based DRL

DeepMind's PlaNet (2019)



Training time:

Method	Modality	Episodes	Cartpole Balance	Cartpole Swingup	Finger Spin	Cheetah Run	Ball in cup Catch	Walker Walk
A3C	proprioceptive	100,000	952	558	129	214	105	311
D4PG	pixels	100,000	993	862	985	524	980	968
PlaNet (ours)	pixels	2,000	986	831	744	650	914	890
CEM + true simulator	simulator state	0	998	850	825	656	993	994
Data efficiency gain PlaNet over D4PG (factor)			100	180	16	50+	20	11



UNIVERSITY OF
TORONTO

Deep Reinforcement Learning

Model-based DRL

DeepMind's PlaNet (2019)



UNIVERSITY OF
TORONTO



Deep Reinforcement Learning

DRL & Evolutionary algorithms

Shauharda Khadka & Kagan Tumer:
Collaborative Evolutionary
Reinforcement Learning (2019)



6 May 2019

Collaborative Evolutionary Reinforcement Learning

Shauharda Khadka^{1,2} Somdeb Majumdar¹ Tarek Nassar¹ Zach Dwiell¹ Evren Tumer¹ Santiago Miret¹
Yinyin Liu¹ Kagan Tumer²

¹ Intel AI Lab ² Oregon State University.




Figure 1. High level schematic of CERL. A portfolio of policy gradient learners operate in parallel to neuroevolution

Deep Reinforcement Learning

Distributed DRL:


- R2D2 (DeepMind, 2018): model-free RL
- MuZero (DeepMind, 2019): model-based RL
- Agent 57 (DeepMind, 2020), based on NGU, which is based on R2D2

Agent57: Outperforming the Atari Human Benchmark

Adrià Puigdomènech Badia^{*1} Bilal Piot^{*1} Steven Kapturowski^{*1} Pablo Sprechmann^{*1} Alex Vitvitskyi¹ — Daniel Guo¹ Charles Blundell¹



Mar 2020



Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model

Julian Schrittwieser,^{1*} Ioannis Antonoglou,^{1,2*} Thomas Hubert,^{1*} Karen Simonyan,¹ Laurent Sifre,¹ Simon Schmitt,¹ Arthur Guez,¹ Edward Lockhart,¹ Demis Hassabis,¹ Thore Graepel,^{1,2} Timothy Lillicrap,¹ David Silver^{1,2*}

¹DeepMind, 6 Pancras Square, London N1C 4AG.

²University College London, Gower Street, London WC1E 6BT.


*These authors contributed equally to this work.

Agent	Median	Mean	Env. Frames	Training Time
Ape-X [18]	434.1%	1695.6%	22.8B	5 days
R2D2 [21]	1920.6%	4024.9%	37.5B	5 days
MuZero	2041.1%	4999.2%	20.0B	12 hours
IMPALA [9]	191.8%	957.6%	200M	—
Rainbow [17]	231.1%	—	200M	10 days
UNREAL ^a [19]	250% ^a	880% ^a	250M	—
LASER [36]	431%	—	200M	—
MuZero Reanalyze	731.1%	2168.9%	200M	12 hours

Table 1: Comparison of **MuZero** against previous agents in Atari.

Deep Reinforcement Learning

Distributed DRL:



Deep Reinforcement Learning

Distributed DRL:





Image Credit: DeepMind

Recurrent Neural Nets

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



Recurrent Neural Nets

- Using RNNs to generate sequences
 - Alex Graves (2014)

"Generating Sequences With
Recurrent Neural Networks"
(2014)



UNIVERSITY OF
TORONTO

Alex Graves



<http://www.cs.toronto.edu/~graves/handwriting.cgi>

Type a message into the text box, and the network will try to write it out longhand


Text --- up to 100 characters, lower case letters work best

piero scaruffi typed this text

piero scaruffi typed this text
piero scaruffi typed this text
piero scaruffi typed this text

Recurrent Neural Nets

- Using RNNs for machine translation



Google Translate now provides live translation of Japanese text

Posted 10 hours ago by Darrell Etherington (@etherington)





Italian - detected ▾



English ▾



L'intervento di Zappa nella melodia di consumo tradizionale comincia dall'esagerazione, fino all'exasperazione, degli elementi piu` correvi, tipo i coretti fatiscanti 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 conduce da un tema al suo opposto senza intoppi, discontinuita` o fratture armoniche, con la coerenza assurda che e` solo dei pazzi e dei geni. |


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 that is devious in the mad reel of sound events, in the perfect twist that leads from a theme to its opposite, smooth, discontinuous or harmonious fractures, with the absurd coherence that is just crazy and genes.

Recurrent Neural Nets

- Using RNNs for scene analysis (Oriol Vinyals)

Show and Tell: A Neural Image Caption Generator

CVPR2015



Recurrent Neural Nets





- Scene analysis (Karpathy & Fei-fei Li)

Deep Visual-Semantic Alignments for Generating Image Descriptions

Andrej Karpathy

Li Fei-Fei


Department of Computer Science, Stanford University



Recurrent Neural Nets

2014: Image captioning

Fei-Fei Li's on algorithm to annotate photos




Recurrent Neural Nets

2014: Google's Scene Analysis system (Vinyals)



Google Research Blog



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
 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>

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.

Recurrent Neural Nets

Using RNNs and variational autoencoders to generate images (DeepMind, 2015)

DRAW: A Recurrent Neural Network For Image Generation

Karol Gregor
Ivo Danihelka
Alex Graves
Danilo Jimenez Rezende
Daan Wierstra
Google DeepMind

May 2015



Deep 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

Variational Autoencoders

2013: Max Welling and Diederik Kingma

2014: Danilo Rezende, Shakir Mohamed,
Daan Wierstra

Application: generate images that are similar
to the ones in the training dataset



Variational Autoencoders

Continuous-time neural networks

- Barak Pearlmutter (1995)
- Ricky Chen's Neural ordinary differential equations (2018)

15 Jan 2019

Neural Ordinary Differential Equations

Ricky T. Q. Chen*, Yulia Rubanova*, Jesse Bettencourt*, David Duvenaud
University of Toronto, Vector Institute



UNIVERSITY OF
TORONTO





Figure 6: Computation graph of the latent ODE model.

Generative Adversarial Networks

Ian Goodfellow (2014)

Alex Radford (2015)



Ian Goodfellow

the game between the generator G and the discriminator D is the minimax objective

$$\min_G \max_D \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_r} [\log(D(\mathbf{x}))] + \mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_g} [\log(1 - D(\tilde{\mathbf{x}}))],$$

Generative Adversarial Networks

Loss functions

losslandscape.com / Javier Ideami

Loss function

$$Loss(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(x_i, y_i; \theta)$$



Depends on θ



High dimensionality

Average over the dataset

Mean square error

$$\min_{\theta} \sum_i \|f(x_i; \theta) - y_i\|^2$$

Minimizes difference between output of network and label

$$\theta^* = \arg \min L(\theta)$$

$f(x, w)$



Non-linear



Non-convex loss function

Loss functions

Regression / Mean square error

$$\min_{\theta} \sum_i \|f(x_i; \theta) - y_i\|^2$$

Penalizes predictions that deviate more from the right values

Classification / Cross entropy loss

$$-(y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i))$$

Penalizes predictions that are confident but wrong

Original GAN

$$E_x[\log(D(x))] + E_z[\log(1 - D(G(z)))]$$

The generator tries to minimize it and the discriminator to maximize it (minimax loss)

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

Conditional GAN (Mehdi Mirza @ Univ Montreal, 2014)

Vanilla GAN

Vanilla GAN


(Goodfellow, et al., 2014)



Discriminator Looks at Latent Variables


Conditional GAN

(Mirza & Osindero, 2014)



Bidirectional GAN

(Donahue, et al., 2016; Dumoulin, et al., 2016)



Generative Adversarial Networks


GANs of 2016

- InfoGAN (Xi Chen @ UC Berkeley, 2016)
- Semi-supervised GAN (Augusto Odena @ Google, 2016)
- Auxiliary Classifier GAN (Augusto Odena @ Google)

Discriminator Predicts Latent Variables


Semi-Supervised GAN

(Odena, 2016; Salimans, et al., 2016)




InfoGAN

(Chen, et al., 2016)



Auxiliary Classifier GAN

(Odena, et al., 2016)




Generative Adversarial Networks

- More than 100 variants of GANs are introduced

Vanilla GAN

Vanilla GAN


(Goodfellow, et al., 2014)



Discriminator Looks at Latent Variables


Conditional GAN

(Mirza & Osindero, 2014)



Bidirectional GAN

(Donahue, et al., 2016; Dumoulin, et al., 2016)




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



Bernt Schiele **Honglak Lee**

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.



MAX-PLANCK-GESELLSCHAFT

Generative Adversarial Networks

Caption generation

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



MAX-PLANCK-GESELLSCHAFT

Generative Adversarial Networks

Video generation

Vondrick & Torralba (MIT)



Vondrick Torralba




Generating Videos with Scene Dynamics

Carl Vondrick
MIT

Hamed Pirsiavash
University of Maryland Baltimore County

Antonio Torralba
MIT



Generative Adversarial Networks


Video generation

Clark & Simonyan (DeepMind, 2019)



ADVERSARIAL VIDEO GENERATION ON COMPLEX DATASETS

Sep 2019

Aidan Clark
DeepMind
Karen Simonyan
DeepMind




Jeff Donahue
DeepMind




a set of four-second synthesized video clips

Generative Adversarial Networks

Video Generation



a set of four-second synthesized video clips



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.

Generative Adversarial Networks

- Photorealistic faces of fake celebrities (Jaakko Lehtinen, 2017)



Figure 5: 1024×1024 images generated using the CELEBA-HQ dataset.



Published as a conference paper at ICLR 2018

PROGRESSIVE GROWING OF GANS FOR IMPROVED
QUALITY, STABILITY, AND VARIATION

Tero Karras
NVIDIA

Timo Aila
NVIDIA

Samuli Laine
NVIDIA

Jaakko Lehtinen
NVIDIA and Aalto University



Figure 10: Top: Our CELEBA-HQ results


Generative Adversarial Networks

- Photorealistic faces of fake celebrities (Tero Karras, 2018)



NVIDIA

A Style-Based Generator Architecture for Generative Adversarial Networks



Generative Adversarial Networks

- WGAN-GP (Aaron Courville, 2017)

Dec 2017

Improved Training of Wasserstein GANs



Université 
de Montréal

Ishaan Gulrajani¹, Faruk Ahmed¹, Martin Arjovsky², Vincent Dumoulin¹, Aaron Courville^{1,3}

¹ Montreal Institute for Learning Algorithms




Generative Adversarial Networks

- BEGAN (Google Brain, 2017)



Boundary Equilibrium Generative Adversarial Networks

David Berthelot, Thomas Schumm, Luke Metz



Generative Adversarial Networks

- SAGAN (Rutgers, 2018)

May 2018

Self-Attention Generative Adversarial Networks

Han Zhang*
Rutgers University

Ian Goodfellow
Google Brain

Dimitris Metaxas
Rutgers University

Augustus Odena
Google Brain





Figure 2: The proposed self-attention mechanism. The $\otimes$ denotes matrix multiplication. The softmax operation is performed on each row.



Generative Adversarial Networks

- BigGAN (DeepMind, 2018): details generated based on long-range dependencies (using Xiaolong Wang's self-attention)

Sep 2018

LARGE SCALE GAN TRAINING FOR
HIGH FIDELITY NATURAL IMAGE SYNTHESIS

Andrew Brock*
Heriot-Watt University



Jeff Donahue
DeepMind



Karen Simonyan
DeepMind



Generative Adversarial Networks


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



(a)Generated Images



Generative Adversarial Networks



Nvidia GauGAN (2018)



Generative Adversarial Networks

Video-to-video

- Dongdong Chen: COVST (Microsoft Research Asia, 2017)



Generative Adversarial Networks

Video-to-video

- Ting-Chun Wang
(Nvidia, 2018)







Figure 1: Generating a photorealistic video from an input segmentation map video on Cityscapes. Top left: input. Top right: pix2pixHD. Bottom left: COVST. Bottom right: vid2vid (ours).

Generative Adversarial Networks

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



Leonid Afremov



Neural Network

Generative Adversarial Networks

Multi-style transfer



Xinyuan Chen (Jiao Tong Univ, 2019)



Apr 2019

Gated-GAN: Adversarial Gated Networks for Multi-Collection Style Transfer

Xinyuan Chen, Chang Xu, Xiaokang Yang, Li Song, and Dacheng Tao.



Generative Adversarial Networks

Artificially intelligent painters invent new styles of art



CAN: Creative Adversarial Networks, Generating "Art" by Learning About Styles

and Deviating from Style Norms | Ahmed Elgammal, Bingchen Liu, Mohamed Elhoseiny, Marian Mazzone 21 Jun 2017

Generative Adversarial Networks

- Cornell Univ's style-transfer network (2017)



Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization



Xun Huang

Serge Belongie

Department of Computer Science & Cornell Tech, Cornell University



Style

Content

Ours


Gatys et al.

Generative Adversarial Networks

- Nvidia's face generator (2018)

A Style-Based Generator Architecture for Generative Adversarial Networks

Dec 2018




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

- Open-source platforms for deep learning
 - Pytorch (Soumith Chintala @ Facebook, 2016)
 - Keras (Francois Chollet @ Google, 2015) which (in 2018) ships out-of-the-box with 5 CNNs pre-trained on the ImageNet dataset
 - VGG16
 - VGG19
 - ResNet50
 - Inception V3
 - Xception



Platforms

- Open-source platforms for deep learning
 - MXNet (UW, NYU, NUS, MIT, CMU, 2015)
 - ONNX (Facebook and Microsoft, 2017)
 - Gluon (Microsoft and Amazon, 2017)

 **DMLC** for Scalable and
Reliable Machine Learning

**MXNet: A Flexible and Efficient Machine Learning
Library for Heterogeneous Distributed Systems**



DMLC UW

2015

Tianqi Chen, Mu Li*, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang,
U. Washington CMU Stanford NUS TuSimple NYU
Tianjun Xiao, Bing Xu, Chiyuan Zhang, Zheng Zhang
Microsoft U. Alberta MIT NYU Shanghai



Gluon Open Source AI
Framework — Amazon and
Microsoft Partner for AI in the
Cloud

October 12, 2017

Platforms

- Open-source platforms of 2018
 - Fast.ai library for training of neural nets
 - Facebook: Detectron for object detection
 - Google: Dopamine for fast prototyping of reinforcement learning algorithms
 - Nvidia: vid2vid for video-to-video synthesis



facebook research

Detectron



Platforms

- Google TensorFlow
- Apple CoreML

Google.ai

Tools for everyone



Platforms



**Machine Learning for
Developers.**

Nexosis makes it easy to start building
machine learning applications.

Hardware Platforms

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




Aids for Machine Learning

- Open-source platforms for deep learning
 - CUDA Convnet
 - cuDNN
 - Deeplearning4j
 - PyBrain
 - PyLearn2
 - SINGA
 - ...

Aids for Machine Learning

- Facebook's FBLearner Flow and Asimo, that allow to optimize neural networks (i.e. automate the job of neural-network engineers).



The screenshot displays the FBLearner Flow web interface. At the top, there is a navigation bar with the FBLearner logo, links to 'Workflow Library', 'Models/Features', and 'Help', and a search bar. Below the navigation bar, the main content area features a 'Launch New Run' button and a 'Compare' button. A search bar is also present, labeled 'Search by workflow, tag, owner, name, or ID'. The interface shows a list of runs under the 'Custom' tab. The table below lists the runs with columns for ID, Owner, Workflow, Name, Progress, Start Time, Tags, Log Loss, and AUC.

ID	Owner	Workflow	Name	Progress	Start Time	Tags	Log Loss	AUC
1047165	Mahaveer Jain	Parameter Sweep Example	-		9/9, 9:06pm	london-demo	-	-
1047298	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.35		9/9, 9:19pm	-	0.00105	0.95524
1047297	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.25		9/9, 9:19pm	-	0.00107	0.95776
1047296	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.3		9/9, 9:19pm	-	0.00104	0.95719
1047295	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.1		9/9, 9:19pm	-	0.00122	0.95871
1047294	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.2		9/9, 9:19pm	-	0.00109	0.95796
1047293	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.15		9/9, 9:19pm	-	0.00115	0.95887
1047292	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.4		9/9, 9:19pm	-	0.00106	0.95355
1047291	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.45		9/9, 9:19pm	-	0.00110	0.95293
1037778	Jason Briceno	Parameter Sweep Example	-		9/8, 2:30pm	-	-	-
950428	Li Zhang	Parameter Sweep Example	-		8/21, 2:40pm	-	-	-
900673	Jiawei Chen	Parameter Sweep Example	-		8/8, 9:11pm	-	-	-
832281	Giri Rajaram	Parameter Sweep Example	-		7/24, 12:56pm	-	-	-
832027	Giri Rajaram	Parameter Sweep Example	-		7/24, 12:34pm	-	-	-

Displaying results 1 - 8 out of 8 matches.

Aids for Machine Learning

- MOA <http://moa.cs.waikato.ac.nz/>
- Neurokernel <http://neurokernel.github.io/>
- NuPic <https://github.com/numenta/nupic>
- Orange <http://orange.biolab.si/>
- RapidMiner <http://rapidminer.com>
- Scikit-learn <http://scikit-learn.org/stable/>
- Spark <http://spark.apache.org/mllib/>
- TunedIT <http://tunedit.org/>
- Vahara <https://github.com/thedatachef/varaha>
- Viv <http://viv.ai/>
- Vowpal Wabbit https://github.com/JohnLangford/vowpal_wabbit/wiki
- Weka <http://www.cs.waikato.ac.nz/ml/weka/>
- Dlib <http://dlib.net/ml.html>
- MADLib <http://madlib.net/>
- Mahout <http://mahout.apache.org/>
- MCMLL <http://mcml.sourceforge.net/>
- MLC++ <http://www.sgi.com/tech/mlc/>
- mloss <http://mloss.org/software/>
- mlpack <http://mlpack.org/>
- Shogun <http://www.shogun-toolbox.org/>
- Stan <http://mc-stan.org/>

Aids for Machine Learning

- Not open-source:
 - Ayasdi <http://www.ayasdi.com/>
 - BigML <https://bigml.com/>
 - H2O <http://h2o.ai>
 - IBM Watson
 - Matlab <http://uk.mathworks.com/solutions/machine-learning/>
 - Nutonian <http://www.nutonian.com/>
 - Prediction.io <http://prediction.io/>
 - Rocketfuel <http://rocketfuel.com/>
 - Skytree <http://www.skytree.net/>
 - Trifacta <http://www.trifacta.com/>
 - Wolfram Alpha <http://www.wolframalpha.com/>
 - Wise.io <http://www.wise.io/>
 - Yhat <https://yhathq.com/>

Info about Machine Learning

- Conferences
 - Intl Joint Conference on Artificial Intelligence (IJCAI)
 - ICML icml.cc/
 - NIPS
 - Machine Learning Conference mlconf.com
 - AAAI aaai.org
 - Machine Intelligence Summit re-work.co
 - Artificial Intelligence and Applications (AIAPP)
 - AI Summit theaisummit.com
 - The AI Conference aiconference.ticketleap.com
 - The O'Reilly Artificial Intelligence Conference
 - Intl Joint Conference on Computational Intelligence ijcci.org

Info about Machine Learning

- Journals

- ML Journal <http://www.springer.com/computer/ai/journal/10994>
- JMLR <http://jmlr.org/papers/>
- Pattern Recognition <http://www.jpr.org/index.php/jpr>
- arXiv <http://arxiv.org/list/stat.ML/recent>
- gitXiv <http://gitxiv.com>

- Blogs

- FastML <http://fastml.com/>
- Chris Olah <http://colah.github.io/>
- Andrej Karpathy <http://karpathy.github.io>
- DeepLearning.net <http://deeplearning.net/>

Info about Machine Learning


- Books
 - Russell & Norvig: Artificial Intelligence: A Modern Approach 3rd Edition, (2009)
 - Witten, Ian et al: Data Mining, Practical Machine Learning Tools and Techniques, 3rd ed (2011)
 - Abu-Mostafa, Yaser et al: Learning from Data (2012)
 - Murphy, Kevin: Machine Learning: A Probabilistic Perspective (2012)
 - Rickhert and Coelho: Building Machine Learning Systems with Python (2013)
 - Buduma, Nikhil – Fundamentals of Deep Learning (2015)
 - Gibson and Patterson: Deep Learning: A Practitioner's Approach (2015)
 - Maters, Timothy: Deep Belief Nets in C++ and CUDA C (2015)
 - Bengio, Yoshua et al – Deep Learning (2016)

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



 Google Cloud
Cloud AutoML^{BETA}



Automatic Machine Learning

Texas A&M University (2018): Auto-Keras

- automatically search for the best architecture and hyperparameters



Automatic Machine Learning

Amazon SageMaker for automatic model tuning (2017)

- BUILD
 - Collect & prepare training data
 - Data labeling & pre-built notebooks for common problems
 - Choose & optimize your ML algorithm
 - Model & algorithm marketplace
 - Built-in, high-performance algorithms
- TRAIN
 - Setup & manage environments for training
 - One-click training on the highest performing infrastructure
 - Train & tune model
- DEPLOY
 - Deploy model in production
 - Scale & manage the production environment
 - Fully managed with auto-scaling for 75% less



Automatic Machine Learning

Google MorphNet (2019) takes an existing neural network as input and produces a new neural network that is smaller and faster



The latest news from Google AI

MorphNet: Towards Faster and Smaller Neural Networks

Wednesday, April 17, 2019

XAI (Explainable AI)

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

XAI

- Carlos Guestrin
- Wojciech Samek
- DARPA's XAI

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


“Why Should I Trust You?”

Explaining the Predictions of Any Classifier

Marco Tulio Ribeiro
University of Washington

Sameer Singh
University of Washington

Carlos Guestrin
University of Washington




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

XAI

Zeiler & Fergus (2013)



XAI

- Naftali Tishby

“Optimal Deep Learning” and the Information Bottleneck method

ICRI-CI retreat, Haifa, May 2015



Naftali Tishby

Noga Zaslavsky



School of Engineering and Computer Science

The Edmond & Lily Safra Center for Brain Sciences

Hebrew University, Jerusalem, Israel



XAI

- Google Brain (2018)

The Building Blocks of Interpretability



March 6, 2018



XAI

- MIT-IBM (2018)

GAN Dissection: Visualizing and Understanding Generative Adversarial Networks

David Bau^{1,2}, Jun-Yan Zhu¹, Hendrik Strobelt^{2,3}, Bolei Zhou⁴, Joshua B. Tenenbaum¹,
William T. Freeman¹, Antonio Torralba^{1,2}

¹MIT CSAIL, ²MIT-IBM Watson AI Lab, ³IBM Research, ⁴The Chinese University of Hong Kong



The [#GANpaint app](#) works by directly activating and deactivating sets of neurons in a deep network trained to generate images. Each button on the left ("door", "brick", etc) corresponds to a set of 20 neurons. The app demonstrates that, by learning to draw, the network also learns about objects such as trees and doors and rooftops. By switching neurons directly, you can observe the structure of the visual world that the network has learned to model.


XAI

- Been Kim at Google Brain: Testing with Concept Activation Vectors – TCAV (2018)

Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV)

Jun 2018

Been Kim Martin Wattenberg Justin Gilmer Carrie Cai James Wexler
Fernanda Viegas Rory Sayres



Software 2.0


Deep Learning represents a fundamental paradigm shift in software engineering

"Neural networks are not just another classifier, they represent the beginning of a fundamental shift in how we write software. They are Software 2.0." (Andrej Karpathy, 2017)

Big Data has changed the way we can solve problems

Software 2.0 can learn from the data


The programmer's job: create and maintain the data (NOT trivial!)



Andrej Karpathy
Director of AI at Tesla

Nov 11, 2017

Software 2.0



Software 2.0

- Google's Transformer (2017) changes the way neural-network architectures are designed
- Models becoming commodities
- Focus on the data-flow, not on the model

- Google Snorkel



snorkel

[Http://snorkel.org](http://snorkel.org)

- Uber Ludwig

LUDWIG

- Apple Overton (Chris Re)

- Bootleg



BOOTLEG

- SambaNova

SambaNova[®]
SYSTEMS




Christopher Ré

Deep Learning



Warning: Deep Learning is NOT the only kind of A.I., it is NOT always the best, it is NEVER precise, it is NOT a brain

For example:






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

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.

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.



441

Danger of A.I.

An Artificial Intelligence Developed Its Own Non-Human Language

Researchers shut down AI that invented its own language

Facebook engineers panic, pull plug on AI after bots develop their own language

'Terminator' Come To Life? – Facebook Shuts Down Artificial Intelligence After It Developed Its Own Language

Danger of A.I.

The Atlantic

How the Enlightenment Ends

Philosophically, intellectually—in every way—human society is unprepared for the rise of artificial intelligence.

HENRY A. KISSINGER JUNE 2018 ISSUE



Yann LeCun

15 May · 🌐

Kissinger falls prey to a very commonly-held erroneous belief, following the AlphaZero results: "On its own, in just a few hours of self-play, it achieved a level of skill that took human beings 1,500 years to attain. Only the basic rules of the game were provided to AlphaZero. Neither human beings nor human-generated data were part of its process of self-learning. If AlphaZero was able to achieve this mastery so rapidly, where will AI be in five years?"

Danger of A.I.

Allen Institute

Asilomar Conference (2017)

AI for Good Global Summit (2017)



BENEFICIAL AI 2017

ASILOMAR AI PRINCIPLES



Danger of A.I.



Electronic Frontier Foundation

<https://www.eff.org/ai/metrics>



Atari 2600 Bank Heist



The Real Danger

- Software to create alternative realities
- Face morphing
 - Project VoCo (Adobe, 2016)
 - Face2Face (Stanford, 2017)
- Voice cloning
 - CandyVoice
 - VivoText
 - Festvox (CMU)

Face Morphing

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



Being John Malkovich
ECCV 2010



Face Morphing

Matthias Niessner's Face2Face (2016): real-time face capture and reenactment



Face Morphing

Matthias Nießner

Stanford 2017

Face2Face

AMAZING NEW VIDEO TWEAKING PROGRAM LETS USERS EDIT A PERSON'S LIP MOVEMENTS IN REAL TIME

By Chloe Olewitz — March 21, 2016 1:08 PM



1



228



Subscribe



Share

Subscribe on



YouTube

Face2Face: Real-time Face Capture and Reenactment of RGB Videos (CVPR 2016 Oral)

Real-time Reenactment




Reenactment Result

<http://www.digitaltrends.com/cool-tech/make-videos-say-what-you-want-face2face/>

Face Morphing

Deepfakes/ Fakeapp (2018)



Face Morphing

Deep video portraits (Christian Theobalt, 2018)

Deep Video Portraits

HYEONGWOO KIM, Max Planck Institute for Informatics, Germany

PABLO GARRIDO, Technicolor, France

AYUSH TEWARI and WEIPENG XU, Max Planck Institute for Informatics, Germany

JUSTUS THIES and MATTHIAS NIESSNER, Technical University of Munich, Germany

PATRICK PÉREZ, Technicolor, France

CHRISTIAN RICHARDT, University of Bath, United Kingdom


MICHAEL ZOLLHÖFER, Stanford University, United States of America

CHRISTIAN THEOBALT, Max Planck Institute for Informatics, Germany



MAX-PLANCK-GESELLSCHAFT

May 2018



Voice 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 Morphing

THE VERGE

Adobe is working on an audio app that lets you add words someone never said

by Nick Statt | @nickstatt | Nov 3, 2016, 6:30pm EDT



Voice Morphing

- Voice cloning



We create virtual speech
with the voices you love

**CUSTOMIZE A TEXT-TO-SPEECH
SYNTHETIC VOICE IS VERY EASY!**

Real-time imitation of various voices from
Microsoft Azure TTS (Text-To-Speech) voice

festvox

[CMU Speech Software](#) | [CMU Speech Group](#) |

The Festvox project aims to make the building of new synthetic voices more systemic and better documented, making it possible for anyone to build a new voice.

Voice Morphing

Google WaveNet (2016)



Voice Morphing

Supasorn Suwajanakorn (2017): given the audio, create the video of a speech

Synthesizing Obama: Learning Lip Sync from Audio

Supasorn Suwajanakorn

SIGGRAPH 2017



Output Obama Video

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

Paul Debevec



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


How 'Rogue One' Brought Back Familiar Faces



Scene Understanding

Jiajun Wu (MIT, 2017): Neural scene de-rendering

Neural Scene De-rendering



Scene Understanding

Lukasz Romaszko (University of Edinburgh, 2017)

Vision-as-Inverse-Graphics:

Obtaining a Rich 3D Explanation of a Scene from a Single Image

Lukasz Romaszko¹ Christopher K.I. Williams^{1,2} Pol Moreno¹ Pushmeet Kohli^{3,*}



Input image



Predicted scene



Different view



Input image



Predicted scene





Scene Understanding

DeepMind's Generative Query Network – GQN (2018)

Ali Eslami and Danilo Rezende: understand the layout of a room



Ali Eslami Danilo Rezende



Scene Understanding

Roei Herzig and Moshiko Raboh (Tel Aviv University, 2018)

Mapping Images to Scene Graphs with Permutation-Invariant Structured Prediction



Roei Herzig*
Tel Aviv University

Gal Chechik
Bar-Ilan University, NVIDIA Research


Amir Globerson
Tel Aviv University

Moshiko Raboh*
Tel Aviv University


Jonathan Berant
Tel Aviv University, AI2




(a) Input image




(b) Ground-truth



(c) Baseline



(d) GPI: Linguistic



(e) Attention for neighbors of tree

Attention Score, Label
0.88, window
0.04, building
0.03, building
0.02, bird
0.02, cat
0.005, tail
0.004, tree

(f) Attention per entities

Attention Score, Label
0.69, bird
0.16, window
0.097, tail
0.039, cat
0.005, building
0.005, building
0.003, tree

Variational Inference

David Duvenaud
(University of
Toronto):
Neural Ordinary
Differential
Equations (2018)

22 Oct 2018

Neural Ordinary Differential Equations

Ricky T. Q. Chen*, Yulia Rubanova*, Jesse Bettencourt*, David Duvenaud
University of Toronto, Vector Institute



Algorithm 1 Reverse-mode derivative of an ODE initial value problem

Input: dynamics parameters θ , start time t_0 , stop time t_1 , final state $\mathbf{z}(t_1)$, loss gradient $\partial L / \partial \mathbf{z}(t_1)$

$$\frac{\partial L}{\partial t_1} = \frac{\partial L}{\partial \mathbf{z}(t_1)}^\top f(\mathbf{z}(t_1), t_1, \theta) \quad \triangleright \text{Compute gradient w.r.t. } t_1$$

$$s_0 = [\mathbf{z}(t_1), \frac{\partial L}{\partial \mathbf{z}(t_1)}, \mathbf{0}, -\frac{\partial L}{\partial t_1}] \quad \triangleright \text{Define initial augmented state}$$

def aug_dynamics($[\mathbf{z}(t), \mathbf{a}(t), -, -], t, \theta$): $\triangleright$ Define dynamics on augmented state

$$\text{return } [f(\mathbf{z}(t), t, \theta), -\mathbf{a}(t)^\top \frac{\partial f}{\partial \mathbf{z}}, -\mathbf{a}(t)^\top \frac{\partial f}{\partial \theta}, -\mathbf{a}(t)^\top \frac{\partial f}{\partial t}] \quad \triangleright \text{Concatenate time-derivatives}$$

$$[\mathbf{z}(t_0), \frac{\partial L}{\partial \mathbf{z}(t_0)}, \frac{\partial L}{\partial \theta}, \frac{\partial L}{\partial t_0}] = \text{ODESolve}(s_0, \text{aug_dynamics}, t_1, t_0, \theta) \quad \triangleright \text{Solve reverse-time ODE}$$

return $\frac{\partial L}{\partial \mathbf{z}(t_0)}, \frac{\partial L}{\partial \theta}, \frac{\partial L}{\partial t_0}, \frac{\partial L}{\partial t_1}$ $\triangleright$ Return all gradients




Figure 6: Computation graph of the latent ODE model.

What is really new since 1989?

- OpenAI's CLIP: contrastive pre-training

CLIP: Connecting Text and Images

January 5, 2021

Learning Transferable Visual Models From Natural Language Supervision


Alec Radford*¹ Jong Wook Kim*




OpenAI



(1) Contrastive pre-training





(2) Create dataset classifier from label text




What is really new since 1989?

- OpenAI's DALL-E



What is really new since 1989?


- OpenAI's GLIDE: diffusion



Generative models:

What is really new since 1989?

- OpenAI's DALL-E2




DALL-E (2021)

OpenAI's DALL-E (2021)

OpenAI



DALL-E



vibrant portrait painting of Salvador Dalí with a robotic half face



a shiba inu wearing a beret and black turtleneck



a close up of a hand/palm with leaves growing from it



an espresso machine that makes coffee from human souls, animation



panda mad scientist mixing sparkling chemicals, animation



a dog's head depicted as an explosion of a nebula



a dolphin in an astronaut suit on saturn, animation



a propaganda poster depicting a cat dressed as french emperor



a teddy bear on a skateboard in times square

When to use Deep Learning


- When to use Deep Learning
 - Large dataset
 - Clean/balanced data
 - Modeling time-series data
 - Statistical methods don't work
- When NOT to use Deep Learning
 - Small dataset
 - Data not balanced
 - The goal is insight
 - Your boss told you to do it because it's popular

When to use Deep Learning

- A good example: NASA + Google (2017): discover patterns hidden in a large dataset of astronomical observations

CHRISTOPHER J. SHALLUE¹ & ANDREW VANDERBURG^{2,3} in press
¹ Google Brain, 1600 Amphitheatre Parkway, Mountain View,
CA 94043
² Department of Astronomy, The University of Texas at
Austin, 2515 Speedway, Stop C1400, Austin, TX 78712
³ Harvard-Smithsonian Center for Astrophysics, 60 Garden
St., Cambridge, MA 02138


December 2017




Deep Learning



Warning: Deep Learning is NOT the only kind of A.I., it is NOT always the best, it is NEVER precise, it is NOT a brain





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; Zur Eye, 2017; Ozlo, 2017)
 - Yahoo (LookFlow, 2013; Incredible Labs, 2014)
 - Twitter (WhetLab, 2015; Magic Pony, 2016)
 - Salesforce (TempoAI, 2015; MetaMind, 2016; PredictionIO
 - Samsung (Viv Labs, 2016)
 - Intel (Nervana and ItseeX in 2016)
 - General Electric (Wise.io, 2017)



Business

Venture Funding Into US Artificial Intelligence, Machine Learning, And Related Startups


2008 through 2017. Dollar volume based on deals of known size; round counts are for all deals.



crunchbase news

Business



- Venture Scanner November 2017 and 2019



Business

- A.I. startups by country

Cognite Ventures 2017 A.I. STARTUPS



Business

- DeepMind (2010, acquired by Google in 2014)
- Vicarious
- Sentient
- ...



Robotics

Inside Vicarious, the Secretive AI Startup Bringing Imagination to Computers

By reinventing the neural network, the company hopes to help computers make the leap from processing words and symbols to comprehending the real world.

by Will Knight May 19, 2016

≡ Forbes / Tech

Interview With The CEO Of The World's Highest Funded AI Company

FEB 29, 2016 @ 09:02 AM

[Sentient Technologies](#)

\$143

million in venture investment,

276
276

Business

- Startups
 - AlchemyAPI
 - Clarifai
 - Ersatz Labs
 - Memkite
 - Nervana
 - Skymind
 - Wise.io
 - Saffron
 - Narrative Science
 - CrowdAI
 - ...


Mostly acquired by
big corporations

Business

- Investment in AI startups in 2014: \$230 million
- Investment in AI startups in 2015: \$128 million


Artificial Intelligence, Real Money

Total venture capital money for pure AI startups, by year



Data: CB Insights

Total Artificial Intelligence Funding by Year



[https://www.venturescanner.com/blog/tags/artificial intelligence](https://www.venturescanner.com/blog/tags/artificial%20intelligence)

What is counted as A.I.?

Business

Corporate Participation in Artificial Intelligence By Quarter
2010-2015



Business

- Between 2010 and 2015 over 45 companies and corporate VC arms have invested in AI startups (top of the list: Bloomberg, Samsung, Rakuten)
- Investment in AI startups in 2014: \$230 million
- Investment in AI startups in 2015: \$128 million

Business





Figure 1: Journal articles mentioning “deep learning” or “deep neural network”

Business

Artificial Intelligence VC Funding by Country

 **Venture Scanner**

Number of AI startups



Source: CBI Insights

Data from April 2017

Business

- Element
- Gradient
- Andre Ng

Element AI, a platform for companies to build AI solutions, raises \$102M



Posted Jun 14, 2017



siliconrepublic 17 AUG 2017

AI visionary Andrew Ng is raising \$150m for a new fund

Google launches Gradient Ventures, an AI-focused venture fund

JULY 11, 2017 2:45 PM



Business

- But most A.I. startups were still losing money in 2020...
- DeepMind's losses: \$154 million in 2016, \$341 million in 2017, \$572 million in 2018, \$649 million in 2019...



DeepMind A.I. unit lost \$649 million last year and had a \$1.5 billion debt waived by Alphabet

PUBLISHED THU, DEC 17 2020-7:11 AM EST UPDATED THU, DEC 17 2020-9:21 AM EST

The Forbes logo, with the word "Forbes" in a blue serif font.

2019, 11:14am EDT

DeepMind Losses Soared To \$570 Million In 2018

Centers of A.I. Research

Univ of Toronto: Geoffrey Hinton, Brendan Frey, Radford Neal

Univ of Montreal: Yoshua Bengio, Pascal Vincent, Aaron Courville, Roland Memisevic

New York Univ: Yann Lecun, Rob Fergus

Stanford Univ: Sebastian Thrun, Andrew Ng, Christopher Manning, Fei-fei Li

UC Berkeley: Bruno Olshausen, Trevor Darrell, Pieter Abbeel, Alexei Efros, Michael Jordan

University of Washington: Pedro Domingos, Ali Farhadi

Carnegie Mellon: Chris Dyer, Tom Mitchell, James Kuffner

Northwestern Univ: Ronald Williams

MIT: Tomaso Poggio, Cynthia Breazeal, Joshua Tenenbaum, Daniel Rus, Antonio Torralba

CalTech: Pietro Perona

Centers of A.I. Research

Oxford Univ: Nando de Freitas, Andrew Zisserman

Cambridge Univ: Carl-Edward Rasmussen

Univ of London: Christopher Watkins

IDSIA (Switzerland): Jurgen Schmidhuber, Dan Ciresan

IDIAP (Switzerland): Ronan Collobert

EPFL (Switzerland): Boi Faltings

Univ of Amsterdam: Max Welling

Tel Aviv Univ: Lior Wolf – Technion: Ran El-Yaniv

National Univ of Singapore: Shuicheng Yan

Australia: Marcus Hutter

Centers of A.I. Research

New York Univ: Yann Lecun, Rob Fergus

Stanford Univ: Sebastian Thrun, Andrew Ng, Christopher Manning, Fei-fei Li

UC Berkeley: Bruno Olshausen, Trevor Darrell, Pieter Abbeel, Alexei Efros,
Michael Jordan, Jitendra Malik, Sergey Levine

University of Washington: Pedro Domingos

Carnegie Mellon: Chris Dyer, Tom Mitchell

Northwestern Univ: Ronald Williams

MIT: Antonio Torralba

Centers of A.I. Research

OpenAI: Ilya Sutskever, Ian Goodfellow, Durk Kingma

Google: Sebastian Thrun, Geoffrey Hinton, Samy Bengio, Ilya Sutskever, Ian Goodfellow, Oriol Vinyals, Quoc Le, Christian Szegedy,

Ross Girshick, Kaiming He, and Ronan Collobert. OpenAI hired Ilya Sutskever. Megvii hired Jian Sun, Tencent hired Dong Yu.

DeepMind: Alex Graves, Karol Gregor, Koray Kavukcuoglu, Andriy Mnih, Shane Legg, Christian Szegedy, Volodymyr Mnih, Daan Wierstra, Matthew Lai

Facebook: Yann Lecun, Rob Fergus, Jason Weston, Leon Bottou, Ronan Collobert, Tomas Mikolov, Ross Girshick, Kaiming He

Twitter: Hugo Larochelle

Microsoft: Li Deng, Dong Yu, Antonio Criminisi

Tesla: Andrej Karpathy

Salesforce: Richard Socher

Indico: Alec Radford

Nvidia: Ming-yu Liu, Raquel Urtasun

A.I. in China

- Microsoft chatbot Xiaoce: 100 million users
- News recommendation: Toutiao
- Voice recognition: iFlytek
- Face recognition: Face++
- Several A.I. unicorns
- Cambricon AI chip



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,

究了硅谷

A.I. in China

- Cambricon AI chip (2016)



Chinese AI Chip Maker Cambricon Aims For 1 Billion Devices In Three Years

YIMIAN WU

November 7, 2017

Cambricon
寒武纪科技

Artificial Intelligence in China

- Tencent Dream Writer (2015)
- Toutiao Xiaomingbot (2016)


AM ASIA MEDIA
INTERNATIONAL

Oct 11, 2015

CHINA: DREAMWRITER IS A DREAM-
KILLER FOR JOURNALISTS



Toutiao has seen
unmatched revenue
growth in its first 4 years



Futurism

The Future of Writing?
China's AI Reporter
Published 450 Articles

A.I. in China

- The B.A.T.
 - Alibaba: ET City Brain (traffic optimization based on visual recognition), ET Medical Brain (medical image analysis)
 - Baidu: open software platform Apollo for self-driving cars (the “Android” of self-driving cars)
 - Tencent has data about social media (Wechat), payments (Wechat Pay) and games



A.I. in China

- A simple definition of A.I.: computational mathematics

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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


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

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 paper 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

- 2015: Ministry of Public Security launches a project to build the world's most powerful facial recognition system with the power to identify any one of China's 1.3 billion citizens within three seconds
- SeetaTech, a start-up established by several researchers from the Institute of Computing Technology at the Chinese Academy of Sciences in Beijing
- 2017: SeetaTech's face recognition deployed at Public Security Departments in at least 15 provinces



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)

79.083

86.450

Joint Laboratory of HIT and iFLYTEK

Artificial Intelligence in China

- October 2017: The robot Xiaoyi (Tsinghua University & iFlyTek) passes the medical licensing examination
- October 2017: Megvii beats Facebook and Google at Microsoft COCO object recognition challenge
- 2018: iFlytech offline speech translator



CHINADAILY 中国日报网
COM.CN

2017-11-10

Chinese robot becomes world's first machine to pass medical exam



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

Artificial Intelligence in China

Baidu

- 2017: Xiaoyu Zaijia/ Little Fish voice-controlled family robot powered by Baidu's DuerOS AI platform
- 2017: Project Apollo for autonomous vehicle technology (with Microsoft and others)
- 2017: Partnership with Nvidia
- 2017: Acquisition of smart-home startup Raven Tech
- 2017: Acquisition of computer-vision startup xPerception



Artificial Intelligence in China

- Anuhi A.I. hospital (2017)



Artificial Intelligence in China

- Medical robot (2017)



Chinese AI-powered robot Xiaoyi took the country's medical licensing examinations
it got a score of 456 points, 96 points above the required



A.I. in China



Monday, Aug 28, 2017

China to add AI courses in primary education



The next breakthrough

- National projects that changed the world
 - Apollo Program (1963)
 - Arpanet (1969)
 - Human Genome Project (1990)



The next breakthrough

- National projects
 - USA: BRAIN Initiative (2013)
 - EU: Human Brain Project (2013)
 - Switzerland: Blue Brain (2015)



Henry Markram



The next breakthrough

- National projects
 - Machine Intelligence from Cortical Networks (Microns) to map a cubic millimeter of cortex (about 10 million synapses for 100,000 neurons)



Machine Intelligence from Cortical Networks (MICrONS)



Christof Koch



ANDREAS TOLIAS

Baylor College of Medicine



David Cox

HARVARD
UNIVERSITY

Next...

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

