

Lecture 02

A Brief Summary of the History of Neural Networks and Deep Learning

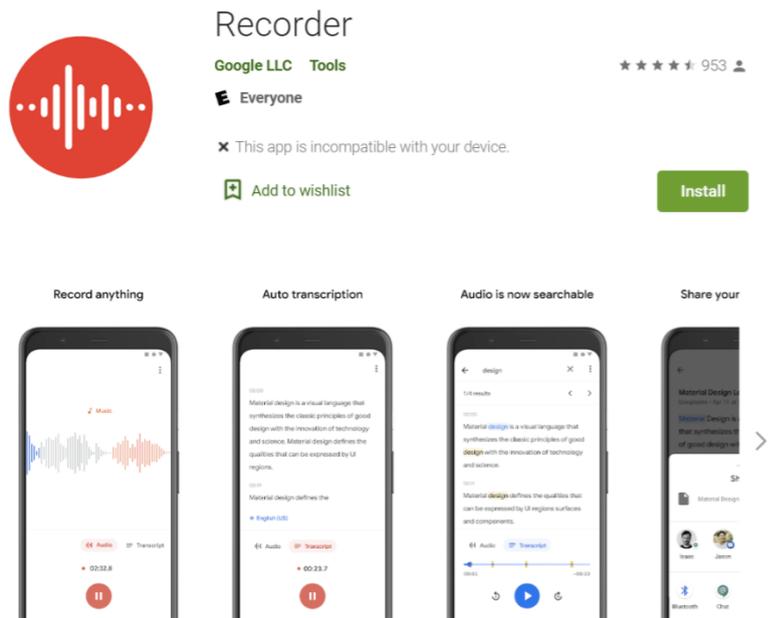
STAT 453: Introduction to Deep Learning and Generative Models
Spring 2020

Sebastian Raschka

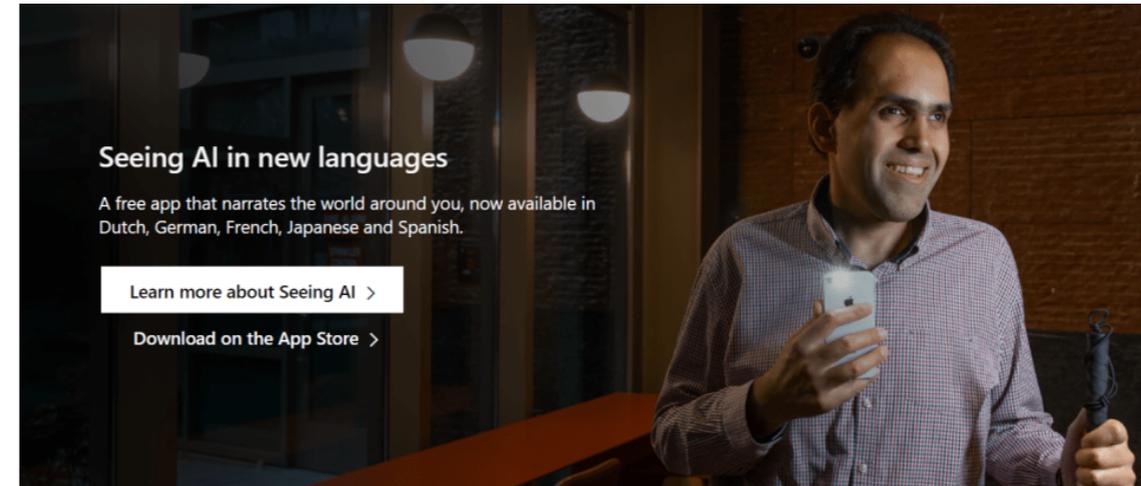
<http://stat.wisc.edu/~sraschka/teaching/stat453-ss2020/>

新年快乐！ 身体健康，工作顺利，一帆风顺，万事如意

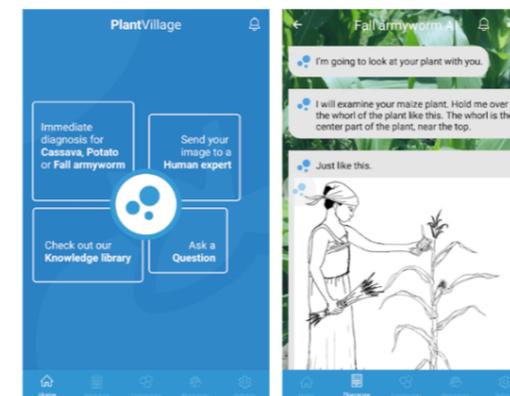
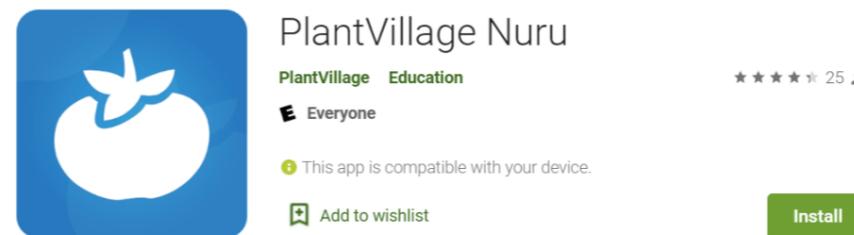
2019 in Review: 10 Brilliant AI Apps



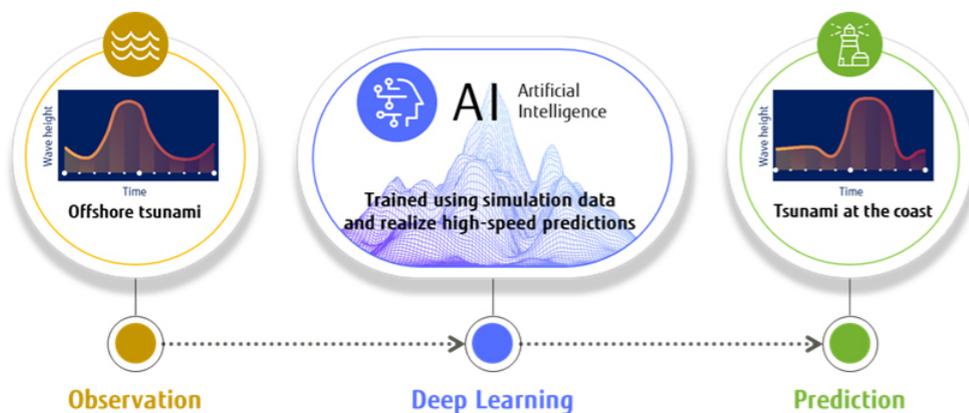
The app uses AI-powered speech recognition to automatically transcribe recordings of lectures, meetings, interviews, etc.



A tap of the finger on a touchscreen image produces audio descriptions of objects in the image and the spatial relationship between them.



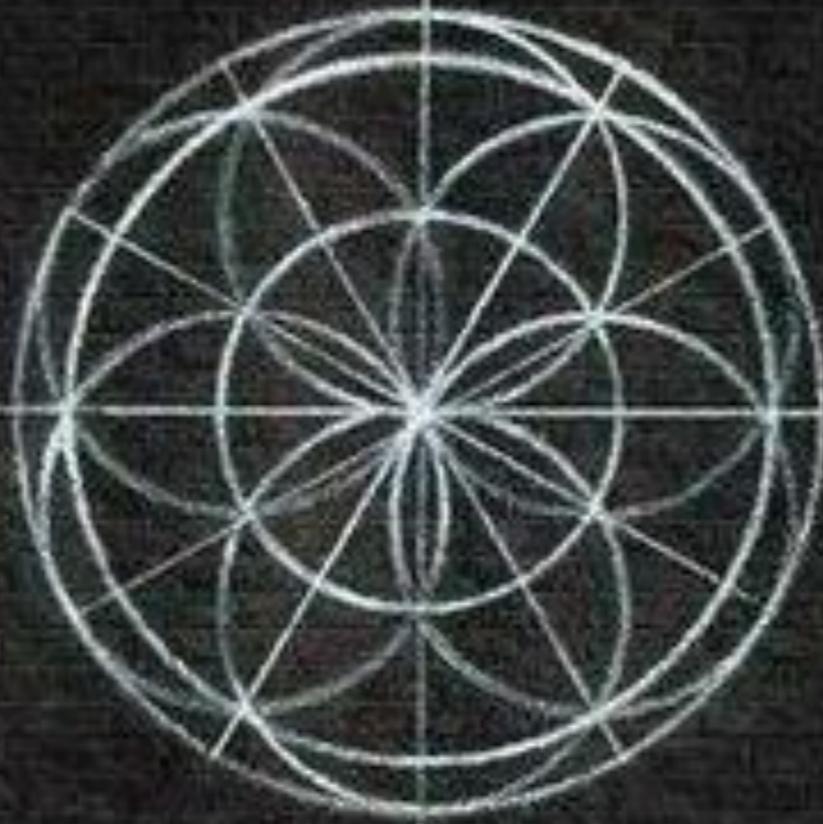
Predicting near-term crop yields for farmers in Africa could help protect popular crops from threats posed for example by climate change.



evacuation app for smartphones designed to help vulnerable individuals make informed and timely evacuation decisions in case of fire or flood

Source: <https://medium.com/syncedreview/2019-in-review-10-brilliant-ai-apps-60d690976ccb>

In today's lecture, we will ...



"Learn how to see. Realize that everything connects to everything else."

Leonardo da Vinci

Image source: <https://me.me/i/learn-how-to-see-realize-that-everything-connects-to-everything-8d07da5f06654b1eacd8c0459a886149>

Overview

1/5 -- Artificial neurons: the Perceptron

2/5 -- Multilayer perceptrons

3/5 -- Deep learning

4/5 -- DL literature and resources

5/5 -- Current trends

The Perceptron

(Will be covered in greater detail next lecture)

1/5 -- Artificial neurons: the Perceptron

2/5 -- Multilayer perceptrons

3/5 -- Deep learning

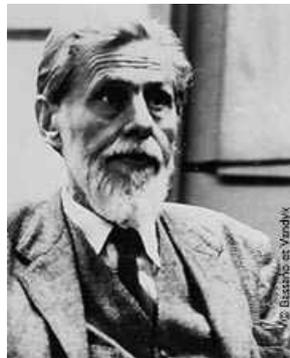
4/5 -- DL literature

5/5 -- Current trends

Neural Networks and Deep Learning -- A Timeline

McCulloch & Pitt's neuron model (1943)

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4), 115-133.



Warren McCulloch



Walter Pitts

WARREN S. MCCULLOCH AND WALTER PITTS 129

$Pr_1(z_1)$, and $C_{mn}(z_1)$ belong to it, where C_{mn} denotes the property of being congruent to m modulo n , $m < n$.

3. The set K has no further members

Then every member of K is realizable.

For, if $Pr_1(z_1)$ is realizable, nervous nets for which

$$N_i(z_1) \equiv Pr_1(z_1) \cdot SN_i(z_1)$$

$$N_i(z_1) \equiv Pr_1(z_1) \vee SN_i(z_1)$$

are the expressions of equation (4), realize $(z_1)z_1 \cdot Pr_1(z_1)$ and $(E z_1)z_1 \cdot Pr_1(z_1)$ respectively; and a simple circuit, c_1, c_2, \dots, c_n , of n links, each sufficient to excite the next, gives an expression

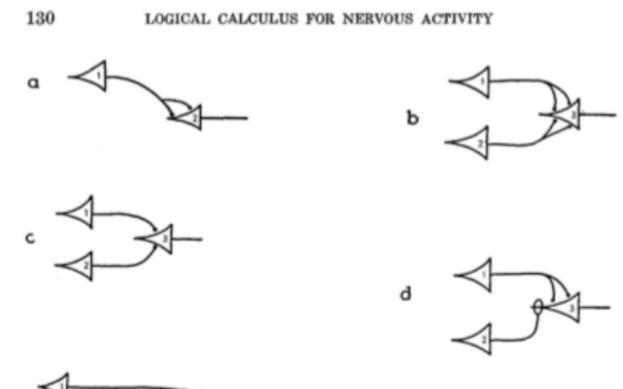
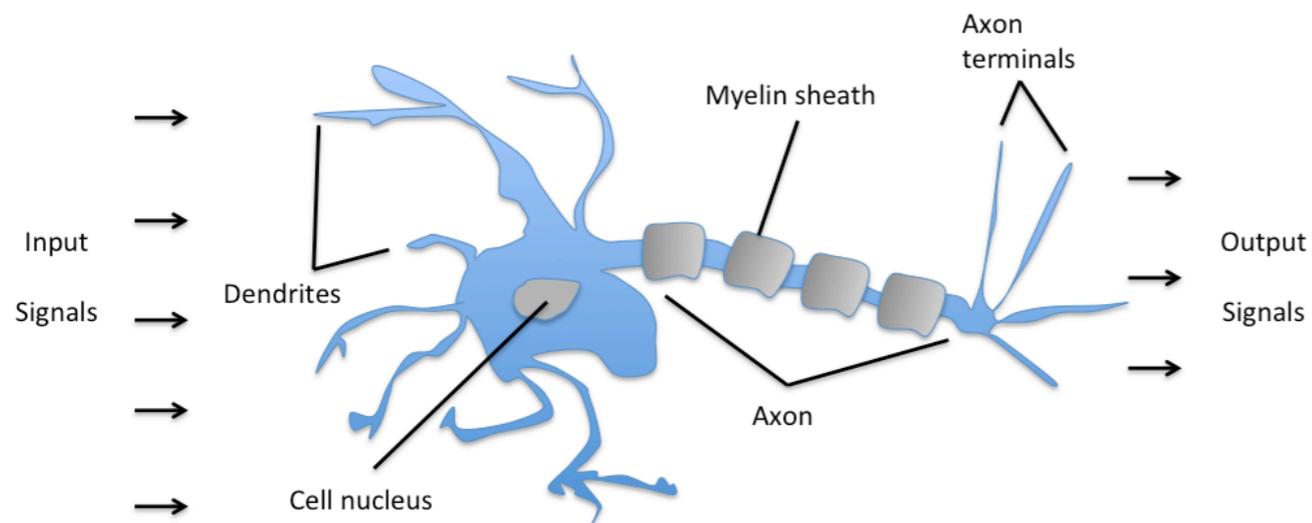


Image Source: <https://www.i-programmer.info/babbages-bag/325-mcculloch-pitts-neural-networks.html>



Schematic of a biological neuron.

Mathematical formulation of a biological neuron, could solve AND, OR, NOT problems

Neural Networks and Deep Learning -- A Timeline

Frank Rosenblatt's Perceptron (1957)

A learning algorithm for the neuron model

Rosenblatt, F. (1957). *The perceptron, a perceiving and recognizing automaton Project Para*. Cornell Aeronautical Laboratory.

Inspired by

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4), 115-133.

Hebb, D. O. (1949). *The organization of behavior. A neuropsychological theory*.

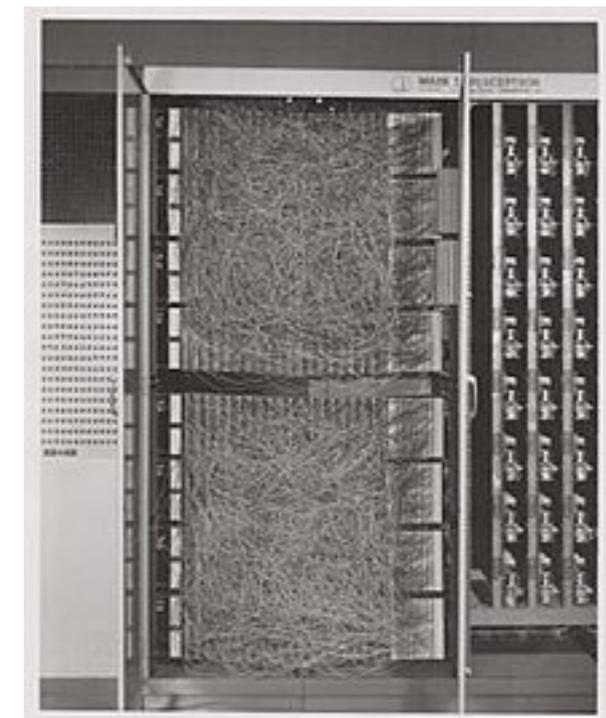
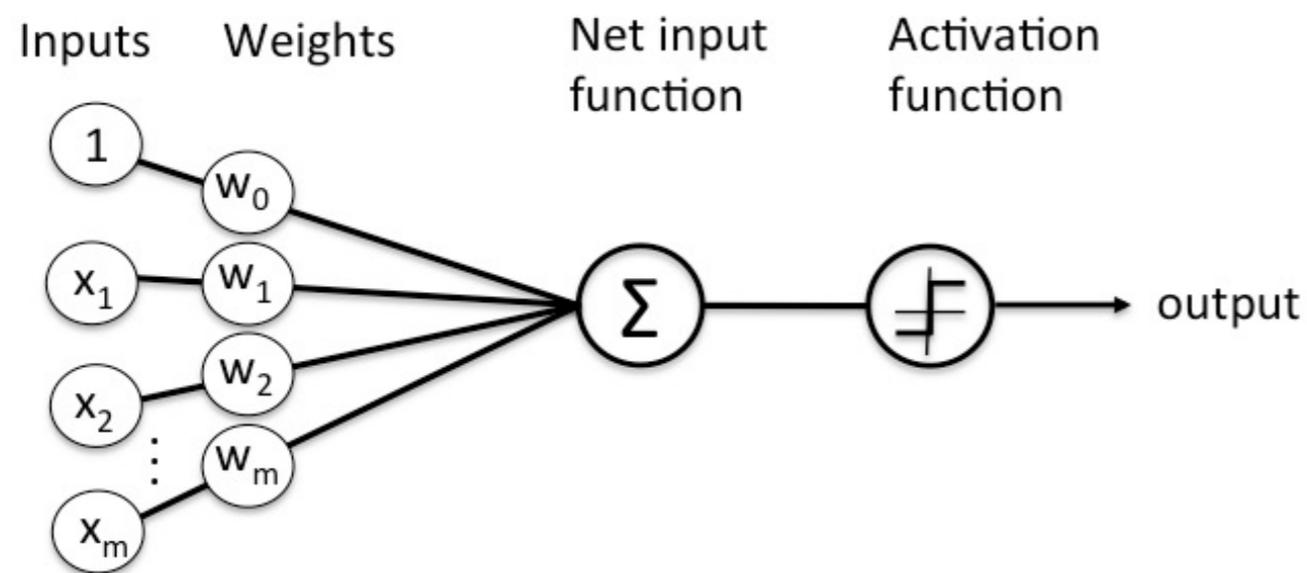


Image source:
https://en.wikipedia.org/wiki/Perceptron#/media/File:Mark_I_perceptron.jpeg

Neural Networks and Deep Learning -- A Timeline

Widrow and Hoff's ADALINE (1960)

A nicely differentiable neuron model

Widrow, B., & Hoff, M. E. (1960). *Adaptive switching circuits* (No. TR-1553-1). Stanford Univ Ca Stanford Electronics Labs.

Widrow, B. (1960). *Adaptive "adaline" Neuron Using Chemical "memistors."*

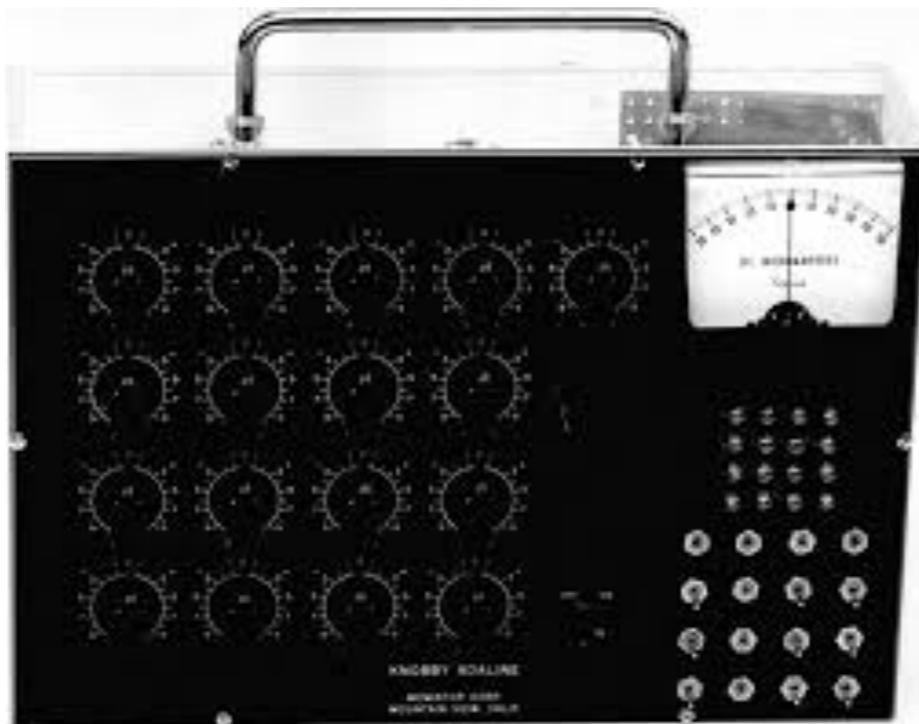
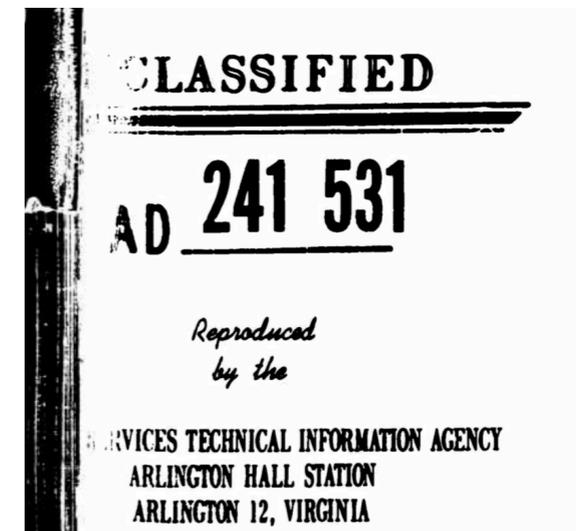


Image source: https://www.researchgate.net/profile/Alexander_Magoun2/publication/265789430/figure/fig2/AS:392335251787780@1470551421849/ADALINE-An-adaptive-linear-neuron-Manually-adapted-synapses-Designed-and-built-by-Ted.png



THIS REPORT HAS BEEN DELIMITED
AND CLEARED FOR PUBLIC RELEASE
UNDER DOD DIRECTIVE 5200.20 AND
NO RESTRICTIONS ARE IMPOSED UPON
ITS USE AND DISCLOSURE.

DISTRIBUTION STATEMENT A

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED.

Widrow and Hoff's ADALINE (1960)

A "single layer neural network"

A nicely differentiable neuron model

Widrow, B., & Hoff, M. E. (1960). *Adaptive switching circuits* (No. TR-1553-1). Stanford Univ Ca Stanford Electronics Labs.

Widrow, B. (1960). *Adaptive "adaline" Neuron Using Chemical "memistors."*

Does the figure below remind you of something you know from other statistics classes?

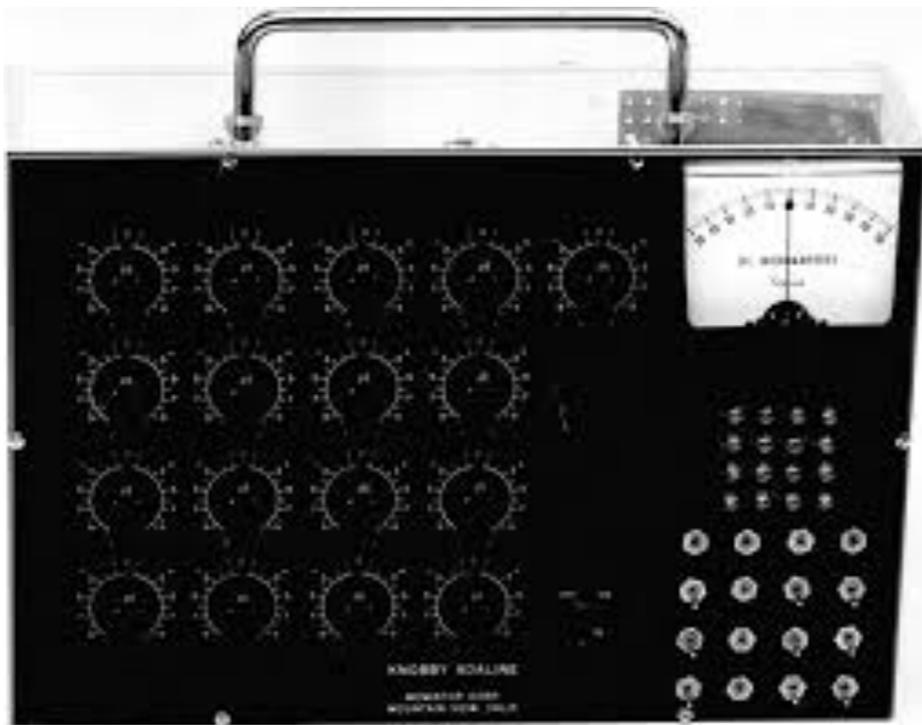
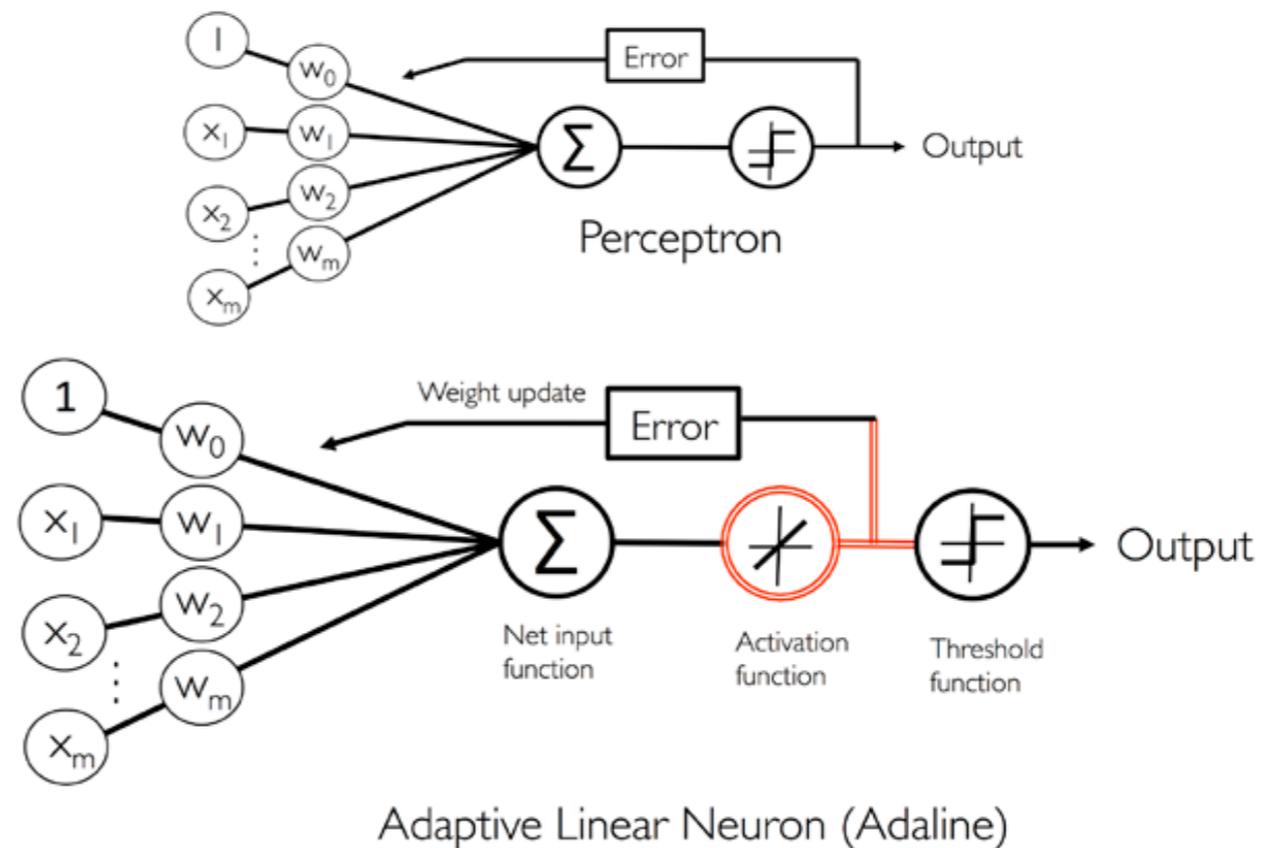


Image source: https://www.researchgate.net/profile/Alexander_Magoun2/publication/265789430/figure/fig2/AS:392335251787780@1470551421849/ADALINE-An-adaptive-linear-neuron-Manually-adapted-synapses-Designed-and-built-by-Ted.png



Source: Raschka and Mirjalily (2019). *Python Machine Learning, 3rd Edition*

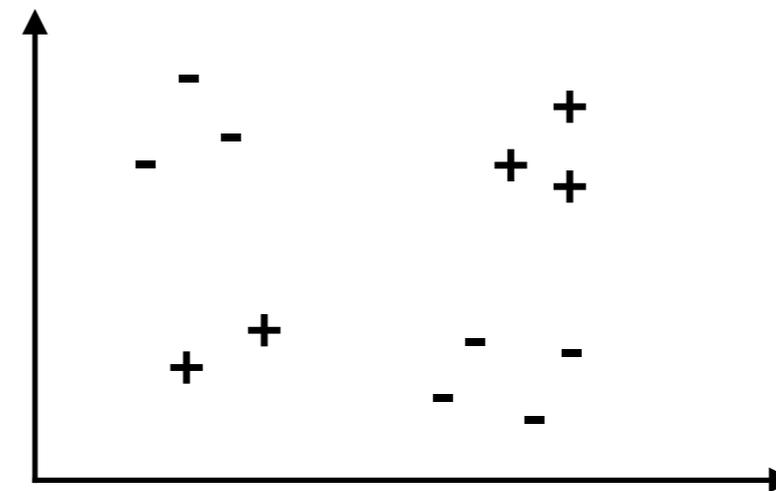
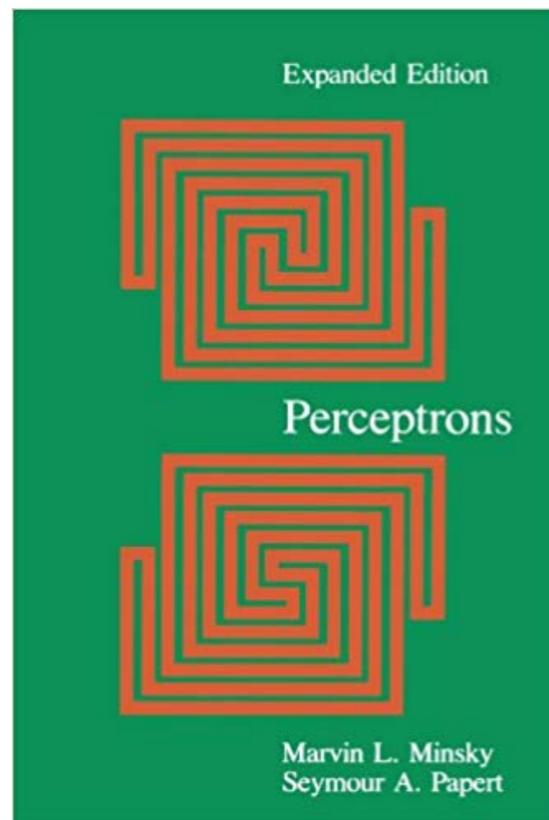
Neural Networks and Deep Learning -- A Timeline

Minsky and Pappert (1969): "Perceptrons" book

=> Problem: Perceptrons (and ADALINEs) could not solve XOR problems!

=> Neurons, a dead end? Start of the first "AI Winter"

Minsky, M. L., & Papert, S. A. (1969). Perceptrons: an introduction to computational geometry. MA: MIT Press, Cambridge.



Training neural networks with multiple layers

(Lecture 09)

1/5 -- Artificial neurons: the Perceptron

2/5 -- Multilayer perceptrons

3/5 -- Deep learning

4/5 -- DL literature

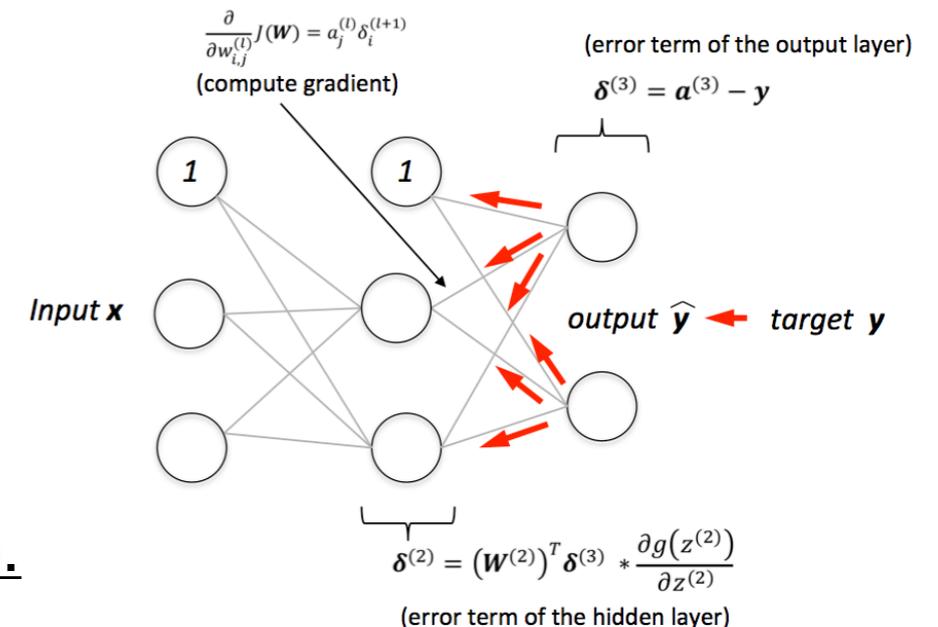
5/5 -- Current trends

Neural Networks and Deep Learning -- A Timeline

- Solution to the XOR problem: hidden layers and non-linear activation functions
- New problem: Hard to train
- Solution: Backpropagation

Note that backpropagation has been independently formulated many times ...

<http://people.idsia.ch/~juergen/who-invented-backpropagation.>



Rumelhart and Hinton (1986) formulated it independently and then showed that it really works (and formed the basis of all consequent neural network and DL progress):

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533.

Also, it was later shown that "neural nets" are universal function approximators

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5), 359-366.

Neural Networks and Deep Learning -- A Timeline

Rumelhart and Hinton (1986) formulated it independently and then showed that it really works (and formed the basis of all consequent neural network and DL progress):

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533.

You just get the output layer to reproduce the input layer, and then you don't need a separate teaching signal. Then the hidden units are representing some code for the input.

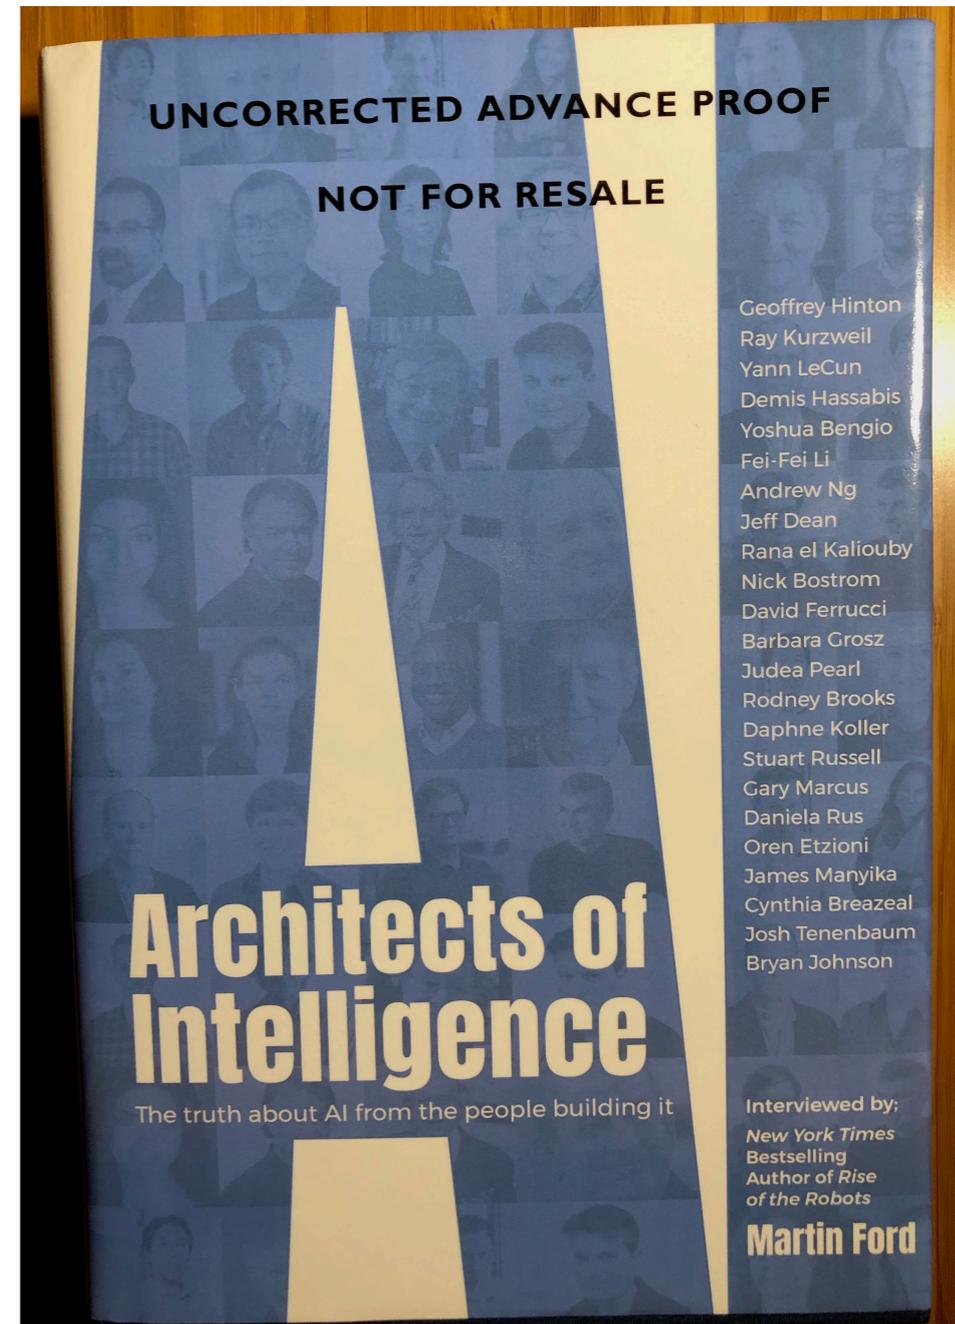
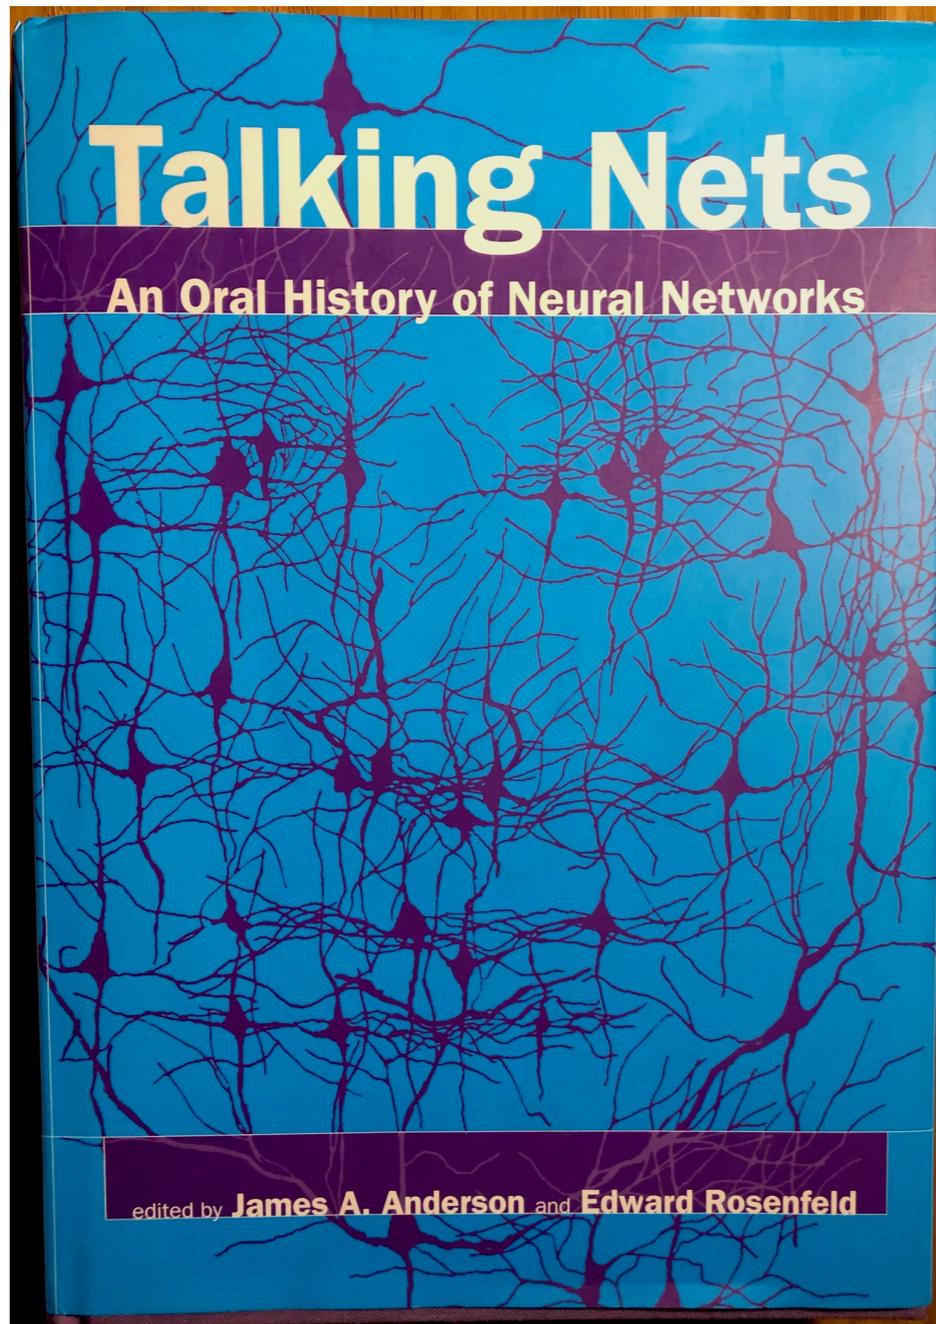
In late 1985, I actually had a deal with Dave Rumelhart that I would write a short paper about backpropagation, which was his idea, and he would write a short paper about autoencoders, which was my idea. It was always better to have someone who didn't come up with the idea write the paper because he could say more clearly what was important.

So I wrote the short paper about backpropagation, which was the *Nature* paper that came out in 1986, but Dave still hasn't written the short paper about autoencoders. I'm still waiting.

What he did do was tell Dave Zinser about the idea of autoencoders and

– *Geoffrey Hinton in Talking Nets - An Oral History of Neural Networks, pg. 380*

Interview with Experts: Suggestions for Pleasure Reading



Deep Learning

Representation learning with advanced architectures with many layers & algorithmic improvements for better computational efficiency and convergence

(Lecture 10 and onwards)

1/5 -- Artificial neurons: the Perceptron

2/5 -- Multilayer perceptrons

3/5 -- Deep learning

4/5 -- DL literature

5/5 -- Current trends

Neural Networks and Deep Learning -- A Timeline

Shortly after followed a breakthrough in image recognition using some clever enhancements to

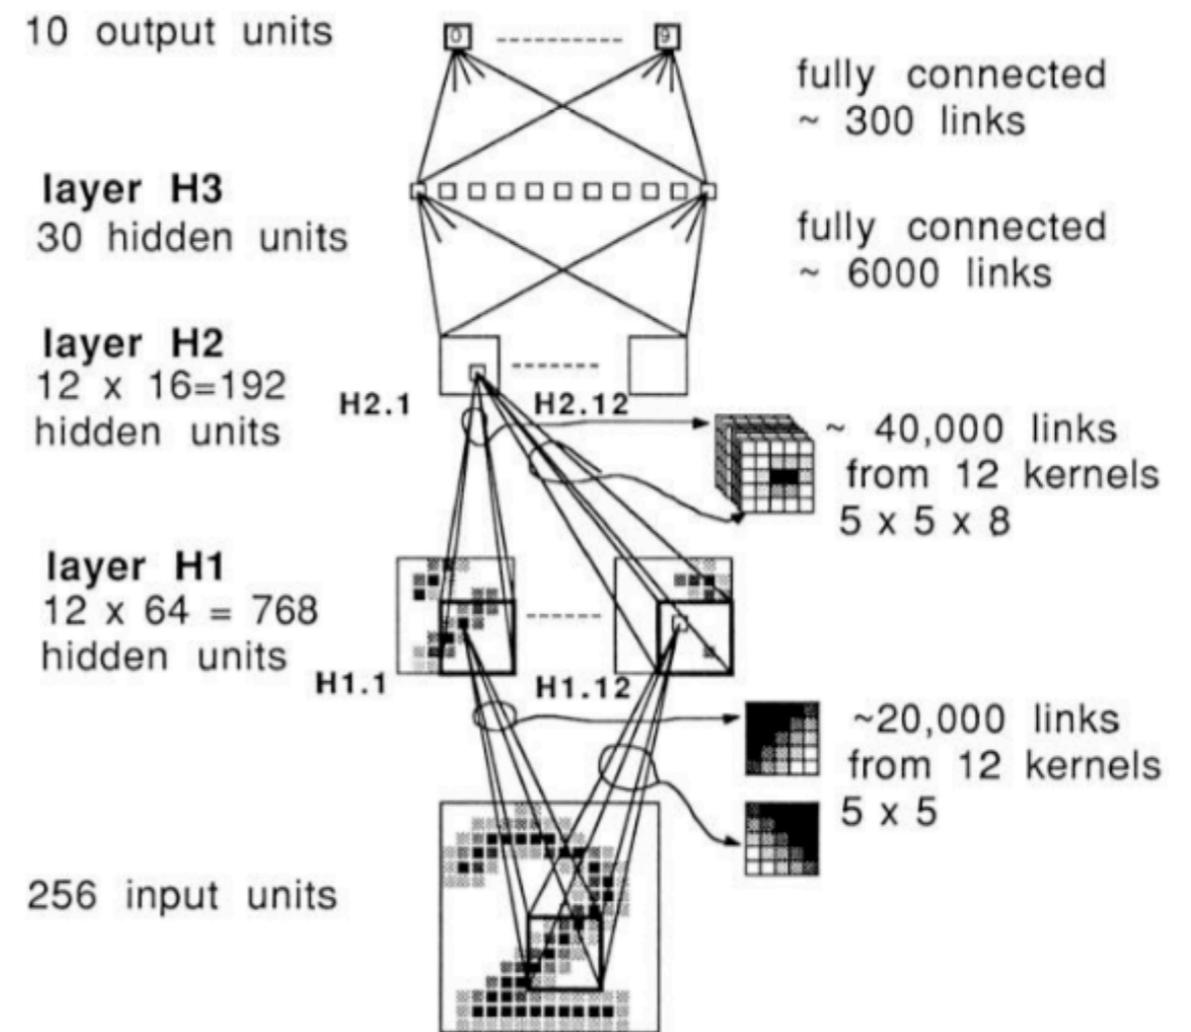
- make training more efficient
- extract local features (and better capture feature dependency)

by

- Weight sharing
- Pooling

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., & Jackel, L. D. (1989). Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4), 541-551.

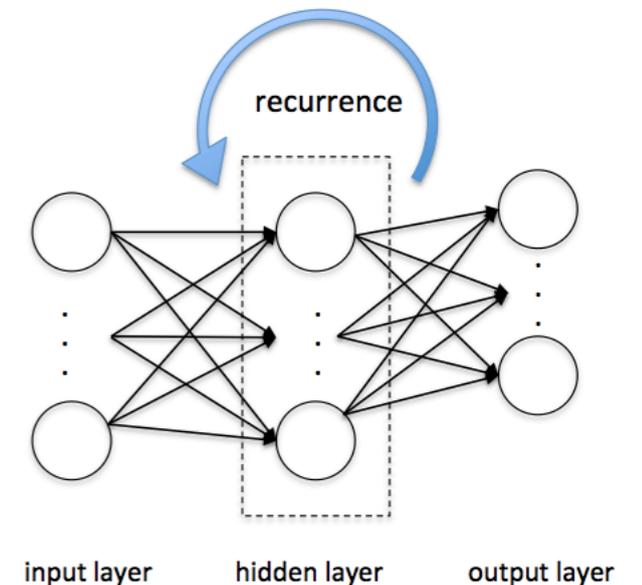
(the origin of Convolutional Neural Networks and MNIST)



Neural Networks and Deep Learning -- A Timeline

Recurrent Neural Networks and Backpropagation through time

Some time created in the 1980's based on Rumelhart's work



New problems: vanishing and exploding gradients!

Schmidhuber, Jürgen (1993). *Habilitation thesis: System modeling and optimization*. Page 150 ff demonstrates credit assignment across the equivalent of 1,200 layers in an unfolded RNN.

Solution: LSTMs (still popular and commonly used)

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735-1780.

Neural Networks and Deep Learning -- A Timeline

- 2nd "AI Winter" in the late 1990's and 2000's
- Probably due to popularity of Support Vector Machines and Random Forests
- Also, neural networks were still expensive to train, until GPUs came into play

Oh, K. S., & Jung, K. (2004). GPU implementation of neural networks. *Pattern Recognition*, 37(6), 1311-1314.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097-1105).



Specifications	Intel® Core™ i7-6900K Processor Extreme Ed.	NVIDIA GeForce® GTX™ 1080 Ti
Base Clock Frequency	3.2 GHz	< 1.5 GHz
Cores	8	3584
Memory Bandwidth	64 GB/s	484 GB/s
Floating-Point Calculations	409 GFLOPS	11300 GFLOPS
Cost	~ \$1000.00	~ \$700.00

(data from ~2017)

Image source: <https://www.amax.com/blog/?p=907>

When did Deep Learning Become Really Popular?

That was the view of people in computer vision until 2012. Most people in computer vision thought this stuff was crazy, even though Yann LeCun sometimes got systems working better than the best computer vision systems, they still thought this stuff was crazy, it wasn't the right way to do vision. They even rejected papers by Yann, even though they worked better than the best computer vision systems on particular problems, because the referees thought it was the wrong way to do things. That's a lovely example of scientists saying, "We've already decided what the answer has to look like, and anything that doesn't look like the answer we believe in is of no interest."

In the end, science won out, and two of my students won a big public competition, and they won it dramatically. They got almost half the error rate of the best computer vision systems, and they were using mainly techniques developed in Yann LeCun's lab but mixed in with a few of our own techniques as well.

MARTIN FORD: This was the ImageNet competition?

GEOFFREY HINTON: Yes, and what happened then was what should happen in science. One method that people used to think of as complete nonsense had now worked much better than the method they believed in, and within two years, they all switched. So, for things like object classification, nobody would dream of trying to do it without using a neural network now.

(Excerpt from "Architects of Intelligence")

Neural Networks and Deep Learning -- A Timeline

- Many enhancements were developed to make neural networks perform better and solve new problems

Rectified Linear Units

Nair, V., & Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)* (pp. 807-814).

BatchNorm

Ioffe, S., & Szegedy, C. (2015, June). Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *International Conference on Machine Learning* (pp. 448-456).

Dropout

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929-1958.

GANs

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672-2680).

(All covered in this course)

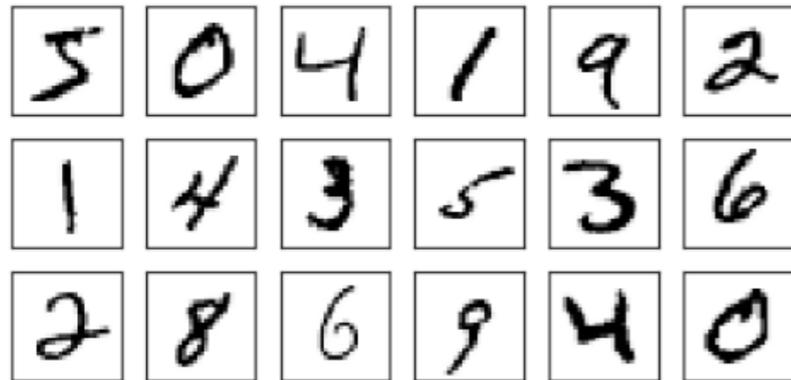
& many more

About the term "Deep Learning" ...

Representation learning is a set of methods that allows a machine to be fed with raw data and to automatically discover the representations needed for detection or classification. Deep-learning methods are representation-learning methods with multiple levels of representation [...]

-- *LeCun, Y., Bengio, Y., & Hinton, G. (2015).
Deep learning. Nature, 521(7553), 436.*

Evolution of Benchmark Datasets for Computer Vision



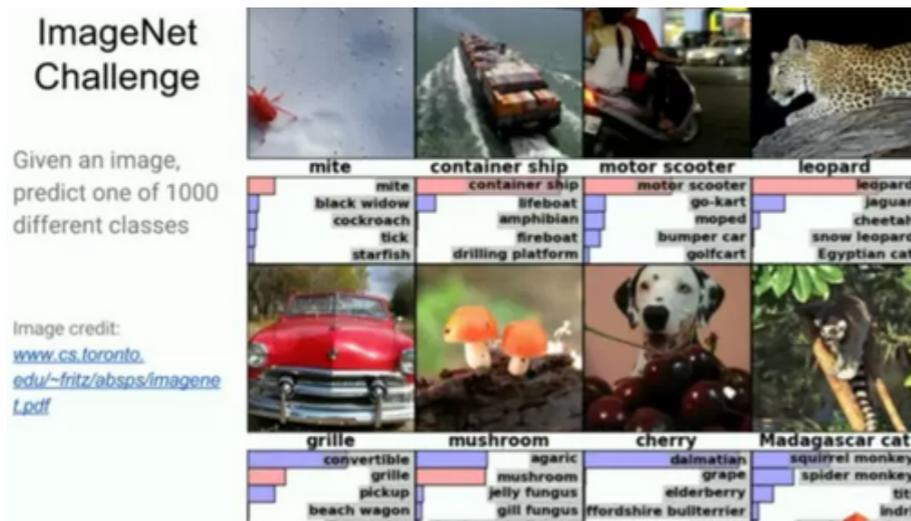
MNIST (1998)

- 60,000 examples, 10 classes
- features: 28x28x1
- <http://yann.lecun.com/exdb/mnist/>



CIFAR-10/CIFAR-100 (2009)

- 60,000 examples, 10 or 100 classes
- features: 32x32x3,
- <https://www.cs.toronto.edu/~kriz/cifar.html>



ImageNet (~2010)

- ~14 million images
- features: full resolution
- <http://www.image-net.org>

Deep Learning Literature

Deep learning is a fast moving field: How to stay up to date with current trends

(Lecture 10 and onwards)

1/5 -- Artificial neurons: the Perceptron

2/5 -- Multilayer perceptrons

3/5 -- Deep learning

4/5 -- DL literature

5/5 -- Current trends

Publishing Culture

In contrast to most other fields:

- Mostly Open Access
- Mostly Conferences (double-blind peer review, competitive ~25% acceptance rate)
- Usually preprints online

Top Conferences for Machine Learning & Arti. Intelligence

Ranking is based on *Conference H5-index* ≥ 12 provided by Google Scholar Metrics

<input type="checkbox"/> Show Due only		All Categories		All Countries		<input type="text" value="Search by keyword"/>	
Hindex	Publisher	Conference Details					
1	158  IEEE	CVPR : IEEE Conference on Computer Vision and Pattern Recognition, CVPR Jun 15, 2019 - Jun 21, 2019 - Long Beach , United States http://cvpr2019.thecvf.com/					
2	101  Neural Information Processing Systems Foundation	NIPS : Neural Information Processing Systems (NIPS) Dec 3, 2018 - Dec 6, 2018 - Palais des Congrès de Montréal , Canada https://nips.cc/					
3	98  Springer	ECCV : European Conference on Computer Vision Sep 8, 2018 - Sep 14, 2018 - Munich , Germany https://eccv2018.org/					
4	91  IMLS	ICML : International Conference on Machine Learning (ICML) Jul 10, 2018 - Jul 15, 2018 - Stockholm , Sweden https://icml.cc/					
5	89  IEEE	ICCV : IEEE International Conference on Computer Vision Oct 27, 2019 - Nov 3, 2019 - Seoul , South Korea http://iccv2019.thecvf.com/ Deadline : Mon 22 Apr 2019					
10	73  Association for Computing Machinery	SIGKDD : ACM SIGKDD International Conference on Knowledge discovery and data mining Aug 19, 2018 - Aug 23, 2018 - London , United Kingdom http://www.kdd.org/kdd2018/					
16	67  Association for Computational Linguistics	ACL : Meeting of the Association for Computational Linguistics (ACL) Aug 28, 2019 - Sep 2, 2019 - Florence , Italy http://www.acl2019.org Deadline : Thu 04 Apr 2019					
22	59  Association for Computing Machinery	SIGMOD : ACM SIGMOD International Conference on Management of Data Jun 30, 2019 - Jul 15, 2019 - Amsterdam , Netherlands http://www.sigmod.org/					

Source: <http://www.guide2research.com/topconf/machine-learning>



Journal of Machine Learning Research

[Home Page](#)

[Papers](#)

[Submissions](#)

[News](#)

[Editorial Board](#)

[Announcements](#)

[Proceedings](#)

[Open Source Software](#)

[Search](#)

The Journal of Machine Learning Research (JMLR) provides an international forum for the electronic and paper publication of high-quality scholarly articles in all areas of machine learning. All published papers are freely available online.

JMLR has a commitment to rigorous yet rapid reviewing. Final versions are [published electronically](#) (ISSN 1533-7928) immediately upon receipt. Until the end of 2004, paper volumes (ISSN 1532-4435) were published 8 times annually and sold to libraries and individuals by the MIT Press. Paper volumes (ISSN 1532-4435) are now published and sold by [Microtome Publishing](#).

News

- [2019.01.20: Volume 19 completed; Volume 20 began.](#)
- [2018.08.28: Volume 18 completed; Volume 19 began.](#)
- [2018.04.16: Changes in JMLR leadership team.](#)
- [2016.12.01: Special topic on Learning from Electronic Health Data](#)
- [2015.09.01: Special issue in Memory of Alexey Chervonenkis.](#)

© JMLR 2019.

<http://www.jmlr.org>



NEWS

[Comments Welcome on New Site](#)

[Help Support JAIR](#)

[New Special Track on Deep Learning, Knowledge Representation, and Reasoning](#)

[New Special Track on AI & Society](#)

<https://www.jair.org/>

Recent Articles

nature machine intelligence

<https://www.nature.com/natmachintell/volumes/2/issues/1>

Volume 2 Issue 1, January 2020



Tree explainer.

Machine learning models based on trees are popular non-linear models. But are trees easy to interpret? In real-world situations, models must be accurate and interpretable, so that humans can understand how the model uses input features to make predictions. In a paper in this issue, Scott Lundberg and colleagues propose TreeExplainer, a general method to explain the individual decisions of a tree-based model in terms of input contributions. The utility... [show more](#)

Image: Scott Lundberg. Cover Design: Karen Moore.

ICML | 2019

Thirty-sixth International Conference on Machine Learning

Dates

Schedule

Calls

Long Beach

ICML 2019 Call for Papers

The 36th International Conference on Machine Learning (ICML 2019) will be held in Long Beach, CA, USA from June 10th to June 15th, 2019. The conference will consist of one day of tutorials (June 10), followed by three days of main conference sessions (June 11-13), followed by two days of workshops (June 14-15). We invite submissions of papers on all topics related to machine learning for the

Year (2019)

Help

policy, and the organizers have the right to reject such submissions, and remove them from the proceedings.

There are several exceptions to this rule:

1. Submission is permitted of a short version of a paper that has been submitted to a journal, but will not be published in that journal on or before June 2019. Authors must declare such dual-submissions either through the CMT submission form, or via email to the program chairs (icml2019pc@gmail.com). It is the author's responsibility to make sure that the journal in question allows dual concurrent submissions to conferences.
2. Submission is permitted for papers presented or to be presented at conferences or workshops without proceedings (e.g., ICML or NIPS workshops), or with only abstracts published.
3. Submission is permitted for papers that are available as a technical report (or similar, e.g., in arXiv). In this case we suggest the authors not cite the report, so as to preserve anonymity.

Finally, note that previously published papers with substantial overlap written by the authors must be cited in such a way so as to preserve author anonymity. Differences relative to these earlier papers must be explained in the text of the submission. For example, (This work develops [our earlier work], which showed that).

Reviewing Criteria

Accepted papers must contain significant novel results. Results can be either theoretical or empirical. Results will be judged on the degree to which they have been objectively established and their potential for scientific and technological impact. Reproducibility of results and easy availability of code will be taken into account in the decision-making process.

Most ML/DL papers are also available as preprints



Cornell University

We gratefully acknowledge support from the Simons Foundation and member institutions.

Logged in as **rasbt** | [My Account](#) | [Logout](#)

arXiv.org

Search or Article ID All fields

[\(Help | Advanced search\)](#)

Open access to 1,490,045 e-prints in Physics, Mathematics, Computer Science, Quantitative Biology, Quantitative Finance, Statistics, Electrical Engineering and Systems Science, and Economics

Subject search and browse:

14 Jan 2019: [The annual update from the arXiv team is now available](#)

3 Jan 2019: [Holiday schedule announced for 21 January](#)

5 Sept 2018: [arXiv looks to the future with move to Cornell CIS](#)

See cumulative ["What's New"](#) pages. Read [robots beware](#) before attempting any automated download

<https://arxiv.org>

- The "computer science -- machine learning (cs.LG)" category that is currently moderated by Tom Dietterich, Luis Lamb, and me receives **~150 new submissions (papers) each day**

(careful & critical though, note that not all preprints were later peer-reviewed / were accepted for publication)

Nice Preprint Recommender System by Andrej Karpathy

The screenshot shows the top navigation bar of the Arxiv Sanity Preserver website. The header is dark red and contains the text "Arxiv Sanity Preserver" in white, followed by a subtitle: "Built in spare time by @karpathy to accelerate research. Serving last 63398 papers from cs.[CV|CL|LG|AI|NE]/stat.ML". On the right side of the header, there are two buttons: "rasbt" and "log out". A diagonal banner in the top right corner says "Fork me on GitHub". Below the header is a search bar with a magnifying glass icon. Underneath the search bar is a row of navigation tabs: "most recent", "top recent", "top hype", "friends", "discussions", "recommended", and "library". The "top recent" tab is currently selected. Below the tabs is a button that says "Only show v1". At the bottom of the visible area, there is a light blue bar with the text "Showing most recent Arxiv papers:".

<http://www.arxiv-sanity.com>

<https://scholar.google.com>



Google Scholar

Articles Case law

Recommended articles

Facial Gender Classification—Analysis using Convolutional Neural Networks

B Lee, SZ Gilani, GM Hassan, A Mian - 2019 Digital Image Computing: Techniques ..., 2019

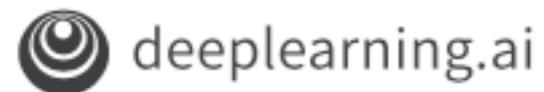
Generating Anime from Real Human Image with Adversarial Training

S Roy - 2019 1st International Conference on Advances in ..., 2019

[See all recommendations](#)

Stand on the shoulders of giants

Newsletters and Podcasts



<https://www.deeplearning.ai/thebatch>
by Andrew Ng

Import AI Newsletter

<https://us13.campaign-archive.com>
by Jack Clark

Artificial Intelligence: AI Podcast by Lex Fridman

<https://lexfridman.com/ai/>

Current Trends

1/5 -- Artificial neurons: the Perceptron

2/5 -- Multilayer perceptrons

3/5 -- Deep learning

4/5 -- DL literature

5/5 -- Current trends

Current Trends and Research Directions

- Applications across fields and in industry
- Engineering new tricks
- Developing specialized hardware
- Developing theory and understanding

Applications

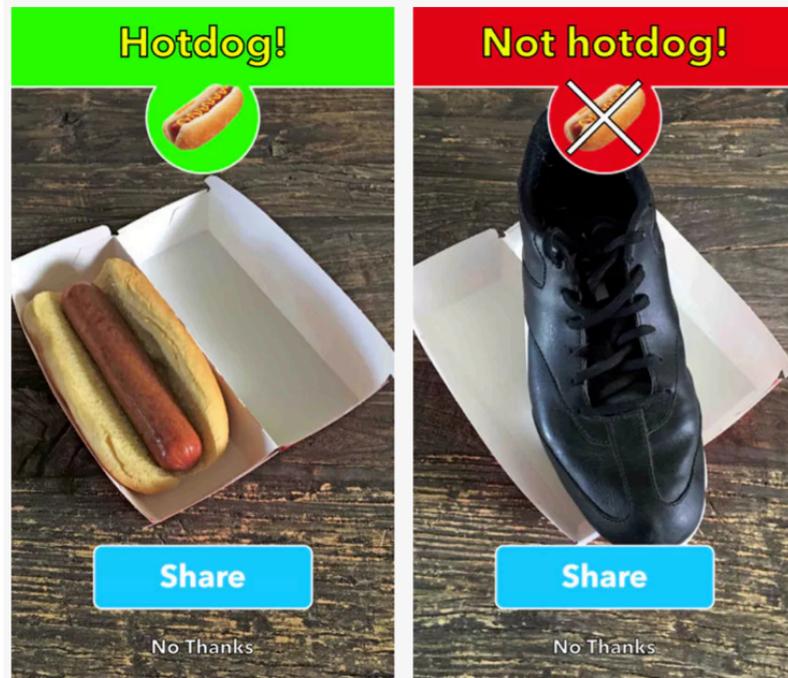


Photo: SeeFood Technologies, Inc.

<https://www.theverge.com/tldr/2017/5/14/15639784/hbo-silicon-valley-not-hotdog-app-download>

Dermatologist-level classification of skin cancer with deep neural networks

Andre Esteva^{1*}, Brett Kuprel^{1*}, Roberto A. Nova^{2,3}, Justin Ko², Susan M. Swetter^{2,4}, Helen M. Blau⁵ & Sebastian Thrun⁶

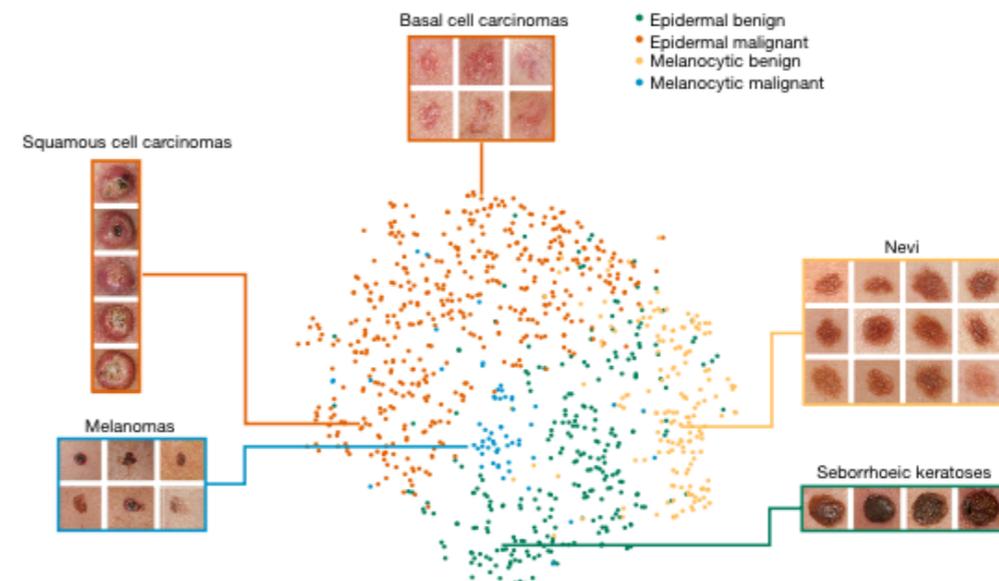


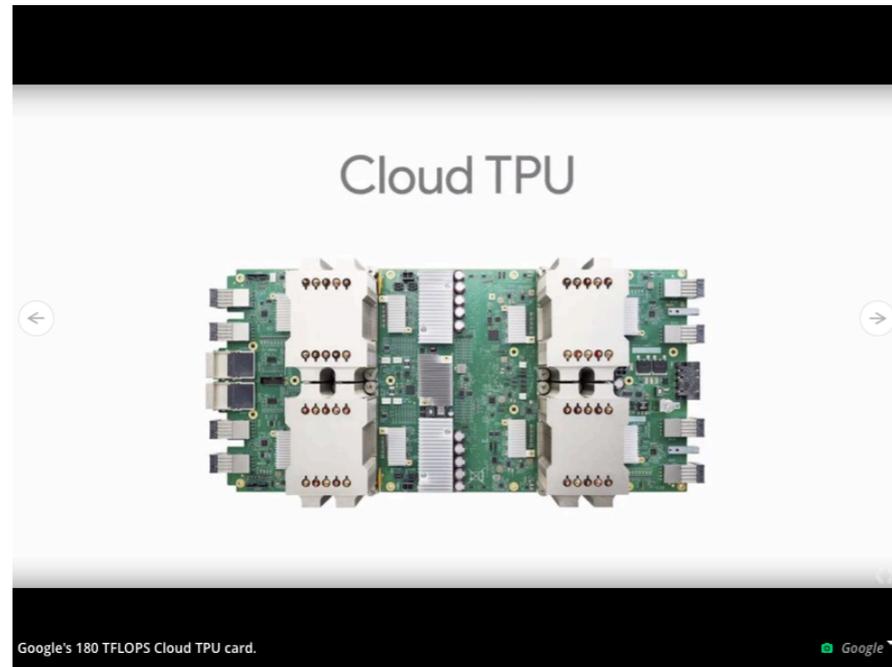
Figure 4 | t-SNE visualization of the last hidden layer representations in the CNN for four disease classes. Here we show the CNN's internal representation of four important disease classes by applying t-SNE (932 images). Coloured point clouds represent the different disease categories, showing how the algorithm clusters the diseases. Insets show images corresponding to various points. Images reprinted with permission

<https://www.nature.com/articles/nature21056.epdf>



<https://ai.googleblog.com/2015/07/how-google-translate-squeezes-deep.html>

Developing Specialized Hardware



<https://arstechnica.com/gadgets/2018/07/the-ai-revolution-has-spawned-a-new-chips-arms-race/>

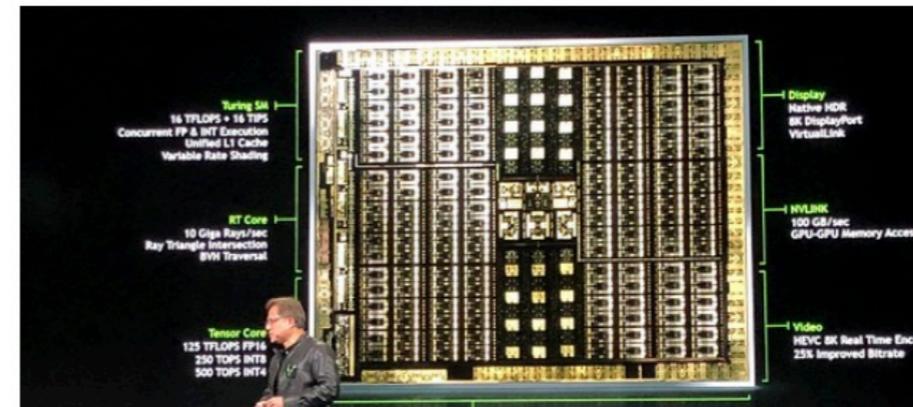
Opinion: New Nvidia chip extends the company's lead in graphics, artificial intelligence

By Ryan Shrout

Published: Aug 14, 2018 2:35 p.m. ET



The only question that remains: How big is Nvidia's advantage over its rivals?



<https://www.marketwatch.com/story/new-nvidia-chip-extends-the-companys-lead-in-graphics-artificial-intelligence-2018-08-14>



<https://developer.arm.com/products/processors/machine-learning/arm-ml-processor>

TECHNOLOGY NEWS

NOVEMBER 28, 2018 / 2:59 PM / 2 MONTHS AGO

Amazon launches machine learning chip, taking on Nvidia, Intel

<https://www.reuters.com/article/us-amazon-com-nvidia/amazon-launches-machine-learning-chip-taking-on-nvidia-intel-idUSKCN1NX2PY>

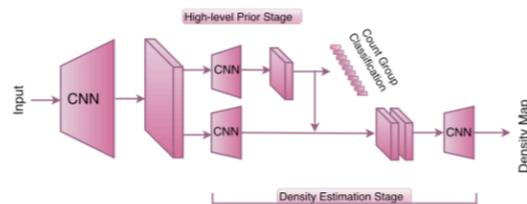
Engineering New Tricks

CNN-based Cascaded Multi-task Learning of High-level Prior and Density Estimation for Crowd Counting

Vishwanath A. Sindagi Vishal M. Patel
 Department of Electrical and Computer Engineering, Rutgers University
 94 Brett Road, Piscataway, NJ, 08854, USA
 vishwanath.sindagi@rutgers.edu, vishal.m.patel@rutgers.edu

Abstract

Estimating crowd count in densely crowded scenes is an extremely challenging task due to non-uniform scale variations. In this paper, we propose a novel end-to-end cascaded network of CNNs to jointly learn crowd count classification and density map estimation. Classifying crowd into various groups is tantamount to coarsely es-



<https://arxiv.org/pdf/1707.09605.pdf>

Group Normalization

Yuxin Wu Kaiming He
 Facebook AI Research (FAIR)
 {yuxinwu, kaiminghe}@fb.com

Abstract

Batch Normalization (BN) is a milestone technique in the development of deep learning, enabling various networks to train. However, normalizing along the batch dimension introduces problems — BN's error increases rapidly when batch size becomes smaller, caused by inaccurate batch statistics estimation. This limits BN's usage for training server models and transferring features to computer vision tasks including detection, segmentation, and video, which require small batches constrained by memory consumption. In this paper, we present Group Normalization (GN) as a simple alternative to BN. GN divides the channels into groups and computes within each group the mean and variance.

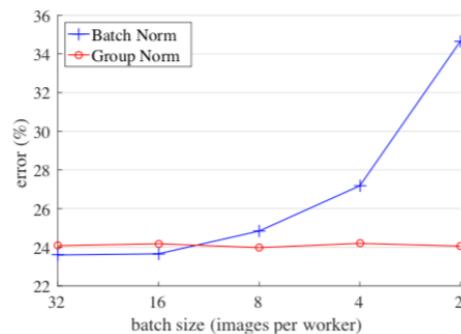


Figure 1. ImageNet classification error vs. batch sizes. This is

<https://arxiv.org/pdf/1803.08494.pdf>

Cyclical Learning Rates for Training Neural Networks

Leslie N. Smith
 U.S. Naval Research Laboratory, Code 5514
 4555 Overlook Ave., SW., Washington, D.C. 20375
 leslie.smith@nrl.navy.mil

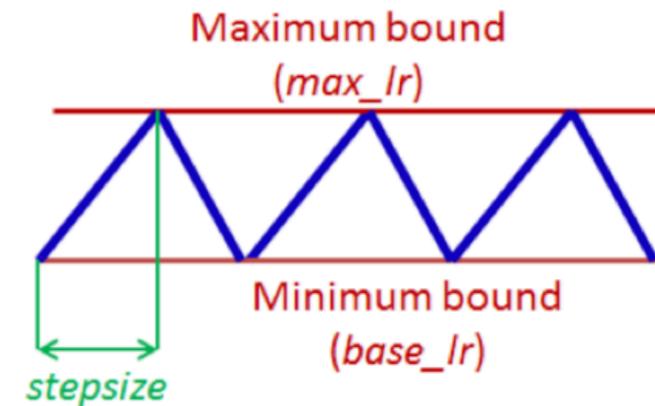


Figure 2. Triangular learning rate policy. The blue lines represent learning rate values changing between bounds. The input parameter *stepsize* is the number of iterations in half a cycle.

<https://arxiv.org/pdf/1506.01186.pdf>

Developing Theory and Understanding

Opening the Black Box of Deep Neural Networks via Information

Ravid Shwartz-Ziv, Naftali Tishby

(Submitted on 2 Mar 2017 (v1), last revised 29 Apr 2017 (this version, v3))

Despite their great success, there is still no comprehensive theoretical understanding of learning with Deep Neural Networks (DNNs) or their inner organization. Previous work proposed to analyze DNNs in the *Information Plane*; i.e., the plane of the Mutual Information values that each layer preserves on the input and output variables. They suggested that the goal of the network is to optimize the Information Bottleneck (IB) tradeoff between compression and prediction, successively, for each layer.

In this work we follow up on this idea and demonstrate the effectiveness of the Information-Plane visualization of DNNs. Our main results are: (i) most of the training epochs in standard DL are spent on *compression* of the input to efficient representation and not on fitting the training labels. (ii) The representation compression phase begins when the training errors becomes small and the Stochastic Gradient Decent (SGD) epochs change from a fast drift to smaller training error into a stochastic relaxation, or random diffusion, constrained by the training error value. (iii) The converged layers lie on or very close to the Information Bottleneck (IB) theoretical bound, and the maps from the input to any hidden layer and from this hidden layer to the output satisfy the IB self-consistent equations. This generalization through noise mechanism is unique to Deep Neural Networks and absent in one layer networks. (iv) The training time is dramatically reduced when adding more hidden layers. Thus the main advantage of the hidden layers is computational. This can be explained by the reduced relaxation time, as this it scales super-linearly (exponentially for simple diffusion) with the information compression from the previous layer.

Geometric Understanding of Deep Learning

Na Lei, Zhongxuan Luo, Shing-Tung Yau, David Xianfeng Gu

(Submitted on 26 May 2018 (v1), last revised 31 May 2018 (this version, v2))

Deep learning is the mainstream technique for many machine learning tasks, including image recognition, machine translation, speech recognition, and so on. It has outperformed conventional methods in various fields and achieved great successes. Unfortunately, the understanding on how it works remains unclear. It has the central importance to lay down the theoretic foundation for deep learning.

In this work, we give a geometric view to understand deep learning: we show that the fundamental principle attributing to the success is the manifold structure in data, namely natural high dimensional data concentrates close to a low-dimensional manifold, deep learning learns the manifold and the probability distribution on it.

We further introduce the concepts of rectified linear complexity for deep neural network measuring its learning capability, rectified linear complexity of an embedding manifold describing the difficulty to be learned. Then we show for any deep neural network with fixed architecture, there exists a manifold that cannot be learned by the network. Finally, we propose to apply optimal mass transportation theory to control the probability distribution in the latent space.

• • •

Some of The Currently "Hot" Research Topics

Semi-supervised learning

(Self-supervised learning and computer vision by Jeremy Howard; https://www.fast.ai/2020/01/13/self_supervised/)

Graph neural networks

(A gentle introduction to graph neural networks by Huang Kung-Hsiang <https://towardsdatascience.com/a-gentle-introduction-to-graph-neural-network-basics-deepwalk-and-graphsage-db5d540d50b3>)

Bias & Fairness

(A tutorial on fairness in machine learning by Ziyuan Zhong, <https://towardsdatascience.com/a-tutorial-on-fairness-in-machine-learning-3ff8ba1040cb>)

...

Reading Material For This Lecture

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.

REVIEW

doi:10.1038/nature14539

Deep learning

Yann LeCun^{1,2}, Yoshua Bengio³ & Geoffrey Hinton^{4,5}

Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other domains such as drug discovery and genomics. Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer. Deep convolutional nets have brought about breakthroughs in processing images, video, speech and audio, whereas recurrent nets have shone light on sequential data such as text and speech.

Machine-learning technology powers many aspects of modern society: from web searches to content filtering on social networks to recommendations on e-commerce websites, and it is increasingly present in consumer products such as cameras and smartphones. Machine-learning systems are used to identify objects in images, transcribe speech into text, match news items, posts or products with users' interests, and select relevant results of search. Increasingly, these applications make use of a class of techniques called deep learning.

Conventional machine-learning techniques were limited in their ability to process natural data in their raw form. For decades, constructing a pattern-recognition or machine-learning system required careful engineering and considerable domain expertise to design a feature extractor that transformed the raw data (such as the pixel values of an image) into a suitable internal representation or feature vector from which the learning subsystem, often a classifier, could detect or classify patterns in the input.

intricate structures in high-dimensional data and is therefore applicable to many domains of science, business and government. In addition to beating records in image recognition¹⁻⁴ and speech recognition⁵⁻⁷, it has beaten other machine-learning techniques at predicting the activity of potential drug molecules⁸, analysing particle accelerator data^{3,18}, reconstructing brain circuits¹¹, and predicting the effects of mutations in non-coding DNA on gene expression and disease^{12,13}. Perhaps more surprisingly, deep learning has produced extremely promising results for various tasks in natural language understanding¹⁴, particularly topic classification, sentiment analysis, question answering¹⁵ and language translation^{16,17}.

We think that deep learning will have many more successes in the near future because it requires very little engineering by hand, so it can easily take advantage of increases in the amount of available computation and data. New learning algorithms and architectures that are currently being developed for deep neural networks will only accelerate this progress.

- <https://www.nature.com/articles/nature14539>
- or <http://www.csri.utoronto.ca/~hinton/absps/NatureDeepReview.pdf>

Downloading research papers on and off campus

Important: Homework for next lecture (ungraded)

As preparation for next lecture

1)

- **Installing and using Python:**

https://github.com/rasbt/stat479-machine-learning-fs18/blob/master/03_python/03_python_notes.pdf

2)

- **Scientific Python and NumPy:**

https://github.com/rasbt/stat479-machine-learning-fs18/blob/master/04_scipython/04_scipython_notes.pdf

- or (better), download & run the interactive Jupyter Notebook version:

https://github.com/rasbt/stat479-machine-learning-fs18/blob/master/04_scipython/04_scipython_notes.ipynb

Next Lecture:

The Perceptron