

CS 190I

Deep Learning

Introduction

Lei Li (leili@cs)

UCSB

Acknowledgement: Slides borrowed from Bhiksha Raj's 11485 and Mu Li & Alex Smola's 157 courses on Deep Learning, with modification

About the Course

- This course focus on a sub-field of machine learning -- Deep Learning, with moderate introduction to general learning concepts and methods.
- This course will teach models, algorithms, and implementation practice
- If you want to take a broader ML course, please elect 165B instead.

Your Instructor

Lei Li

Co-Director, [UCSB NLP Group](#)

Assistant Professor

[Computer Science Department](#)

[Univeristy of California Santa Barbara](#)

Research area: natural language processing, machine learning, data mining.

Topics:

- Machine translation, speech translation, multilingual NLP.
- Text generation and summarization.
- Reasoning and question answering.
- Information extraction.
- AI for drug discovery
- Green and Efficient ML
- Time series mining and prediction
- Probabilistic inference, Bayesian sampling methods

Career path: CMU -> UCB -> Baidu -> Bytedance (Tiktok) -> UCSB

TA

- Krushna Chirag Shah (Office Hour: Tuesday 2-3pm, Trailer 936)
- Zoey Song (Office Hour: Thursday 5-6pm, HH 2014)
- Danqing Wang (Office Hour: Wednesday 11-12am, HH 2014)

Prerequisite

- You should have taken the following courses:
 - Calculus: Math 3A, 3B, 6A
 - Integration and derivative
 - Calculate gradients for multiple variables
 - Linear algebra: Math 4A, 4B
 - Vector, Matrix, norm, linear independence
 - Probability: Pstat 120A & 120B
 - Bayes Rule, likelihood, MLE
 - Algorithm & coding: CS 130A & 130B
 - Python, numpy, notebook

Logistics

- Course website:
 - <https://www.cs.ucsb.edu/~leili/course/dl23w/>
- Text
 - Dive into Deep Learning, Aston Zhang, Zachary Lipton, Mu Li, Alexander Smola.
(available [online](#))
 - You are **required** to read the chapters of the book for each lecture listed in the syllabus.
 - (Optional) Mathematics for Machine Learning.

Lecture

- Required to attend
- M/W 2pm-3:15pm, CHEM 1171
- If you have legitimate reason to be absent (approval from DSP or Student Health or Department), please email me and TA
- In-class quiz/poll/discussion at random times
 - Please respond to all during the class. We will mark your participation (but not the correctness)
 - 5% in-class quiz
 - 5% for forum discussion
- Final exam: 30%

Discussion Forum

- Ed platform
 - <https://edstem.org/us/join/DEfrn>
 - Post questions or discussion on topics related to course material, assignments
 - Message can be private if only send to instructor & TA
 - We will use the same platform for in-class quiz

Homework

- 3 Writing Assignments
 - 10% each
- 2 Machine Problems
 - 15% each
 - Building real DL models
 - Compete on realtime leaderboards on Kaggle
 - please start as early as possible
- Submission on Gradescope
 - **You should already be added, let me know if not.**
- Deadline: midnight on the due date
- Late days
 - a total of 3 days (for all HW, based on days), no penalty.
 - Solution submitted after late days will be graded 0.

Academic integrity is absolutely required

- Allowed:
 - Discussion of lecture and textbook materials
 - Discussion of how to approach assignments, what techniques to consider, what textbook or lecture material is relevant
- Not allowed:
 - Sharing ideas in the form of code, pseudocode, or solutions
 - Turning in someone else's work as your own, even with that person's permission.
 - Allowing someone else to turn in your work as his or her own.
 - Turning in work without proper acknowledgment of the sources of the content (including ideas) contained within the work.
- We will use software to detect plagiarism.
 - It will detect even if change of variables

Recitation Sessions

- Fridays:
 - 9-10am, PHELP 1444
 - 10-11am, PHELP 1440
 - 11-12, PHELP 1440
- Encourage to attend
- TA will cover
 - Background materials
 - Additional explanation and examples
 - Coding assistance
 - Q/A

Computing Resources

- UCSB supercomputing center
 - <http://csc.cnsi.ucsb.edu/acct>
- ECE ML computing lab
 - instruction to use will be sent later
- CS Computing Lab
- Google Colab
 - <https://colab.research.google.com/>
 - Free

Discussion

- Why you are taking the course?

Deep Learning

- Deep Learning have become one of *the* main approaches to AI
- They have been successfully applied to various pattern recognition, prediction, and analysis problems
- In many problems they have established the state of the art
 - Often exceeding previous benchmarks by large margins
 - Sometimes solving problems you couldn't solve using earlier ML methods

Breakthroughs with Deep Learning

www.technewsworld.com/story/84013.html

TECHNEWSWORLD EMERGING TECH

Computing Internet IT Mobile Tech Reviews Security Technology Tech Blog Reader Services

Microsoft AI Beats Humans at Speech Recognition

By Richard Adhikari
Oct 20, 2016 11:40 AM PT

Print Email




Image: Adobe Stock

Microsoft's Artificial Intelligence and Research Unit earlier this week reported that its speech recognition technology had surpassed the performance of human transcriptionists.

5
25
45
11
0
104

Facebook Twitter LinkedIn Google+ StumbleUpon RSS

Most Popular Newsletters News Alerts

How do you feel about Black Friday and Cyber Monday?

- They're great -- I get a lot of bargains!
- The deals are too spread out -- I'd prefer just one day.
- They're a fun way to kick off the holiday season.
- I don't like the commercialization of Thanksgiving Day.
- They're crucial for the retail industry and the economy.
- The deals typically aren't that good.

Vote to See Results

E-Commerce Times

Black Friday Shoppers Hungry for New Experiences, New Tech

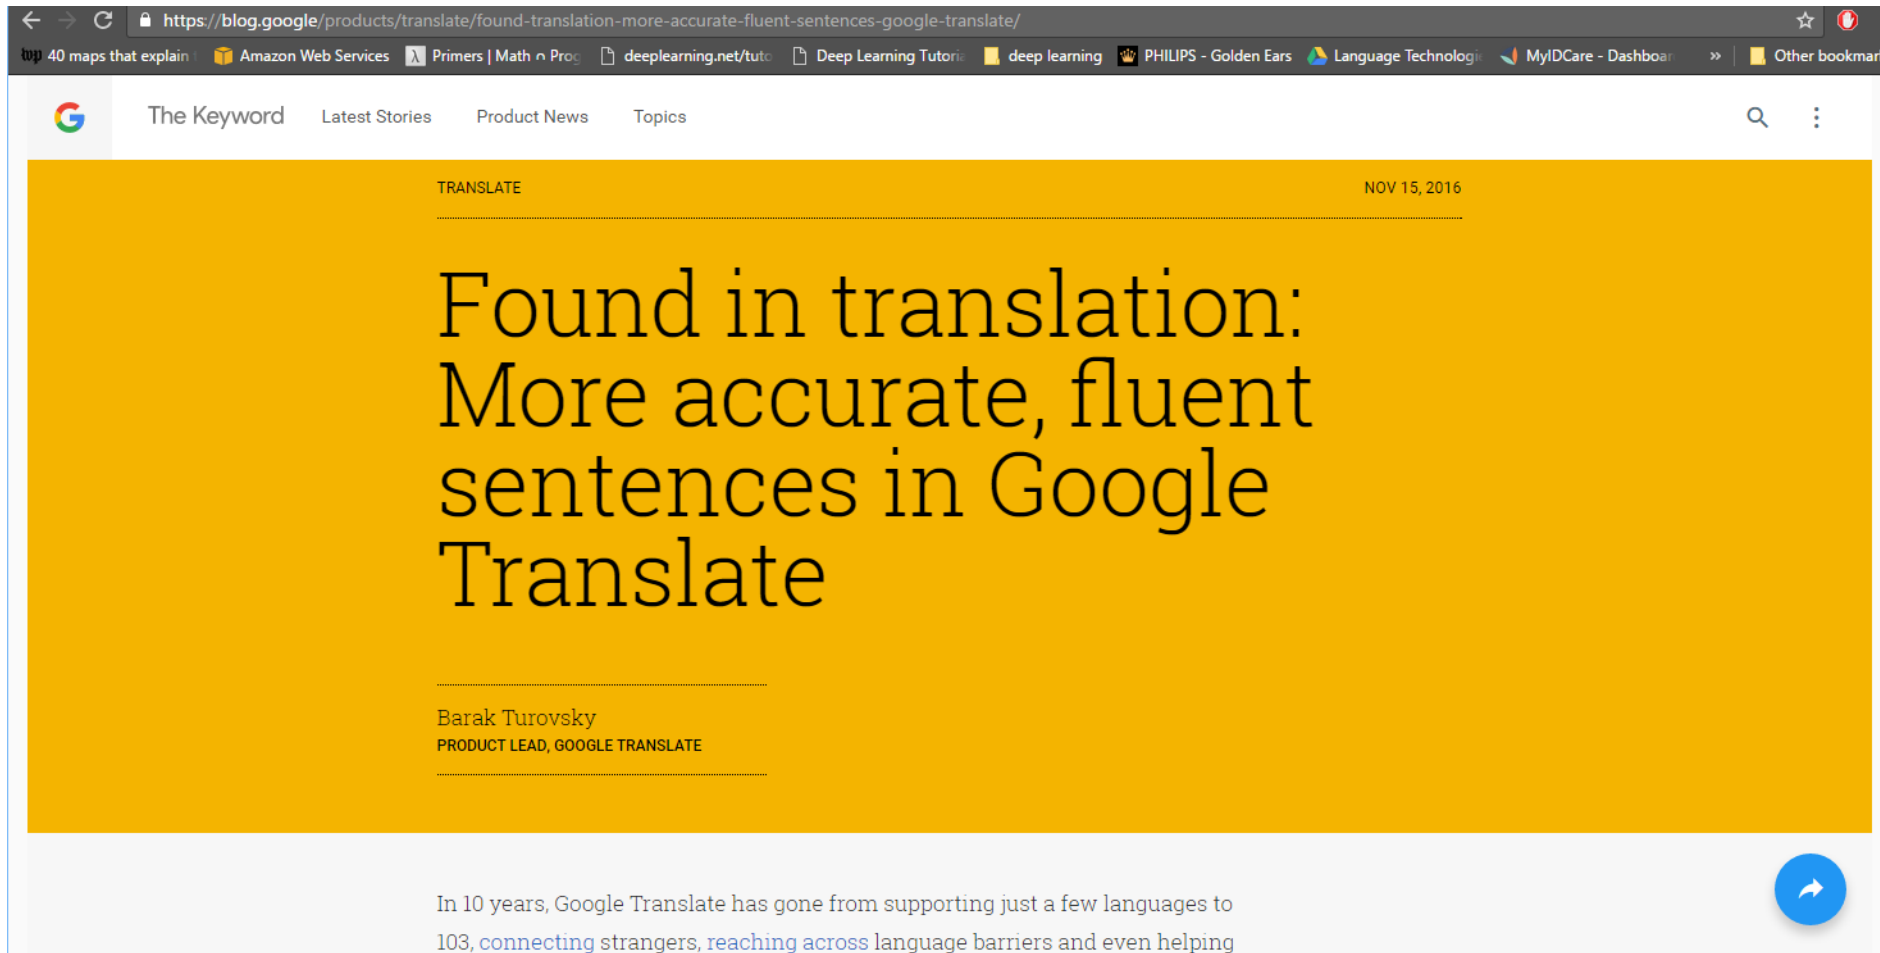
Pay TV's Newest Innovation: Giving Users Control

Apple Celebrates Itself in \$300 Coffee Table Tome

AWS Enjoys Top Perch in IaaS, PaaS Markets

US Comptroller Gears Up for Blockchain and

Breakthrough with Deep Learning



The image shows a screenshot of a web browser displaying a Google blog post. The browser's address bar shows the URL: <https://blog.google/products/translate/translation-more-accurate-fluent-sentences-google-translate/>. The browser's tab bar includes several open tabs, such as '40 maps that explain', 'Amazon Web Services', 'Primers | Math n Prog', 'deeplearning.net/tuto', 'Deep Learning Tutori', 'deep learning', 'PHILIPS - Golden Ears', 'Language Technologi', 'MyIDCare - Dashboar', and 'Other bookmark'. The Google logo is visible in the top left corner of the page, followed by navigation links: 'The Keyword', 'Latest Stories', 'Product News', and 'Topics'. A search icon and a menu icon are in the top right. The main content area has a yellow background. At the top left of this area, it says 'TRANSLATE' and at the top right, 'NOV 15, 2016'. The title of the post is 'Found in translation: More accurate, fluent sentences in Google Translate'. Below the title, the author is listed as 'Barak Turovsky' with the title 'PRODUCT LEAD, GOOGLE TRANSLATE'. At the bottom of the page, there is a light blue banner with the text: 'In 10 years, Google Translate has gone from supporting just a few languages to 103, connecting strangers, reaching across language barriers and even helping'. A blue circular share icon is in the bottom right corner of the banner.

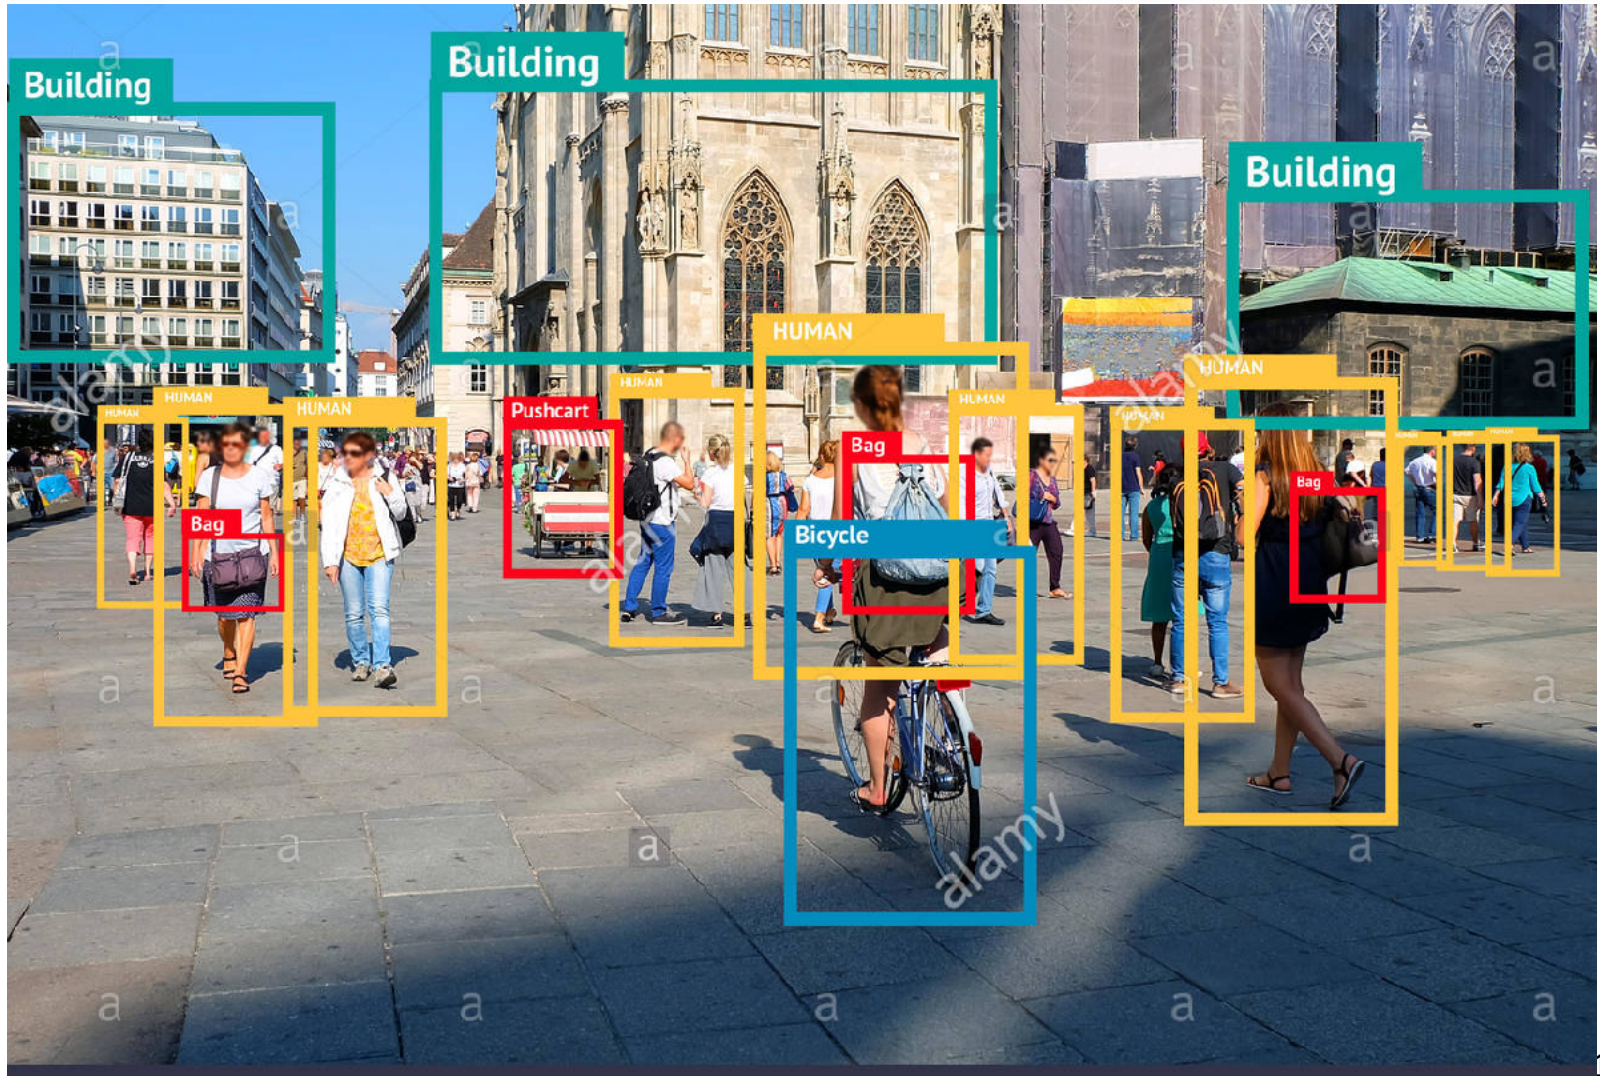
TRANSLATE NOV 15, 2016

Found in translation: More accurate, fluent sentences in Google Translate

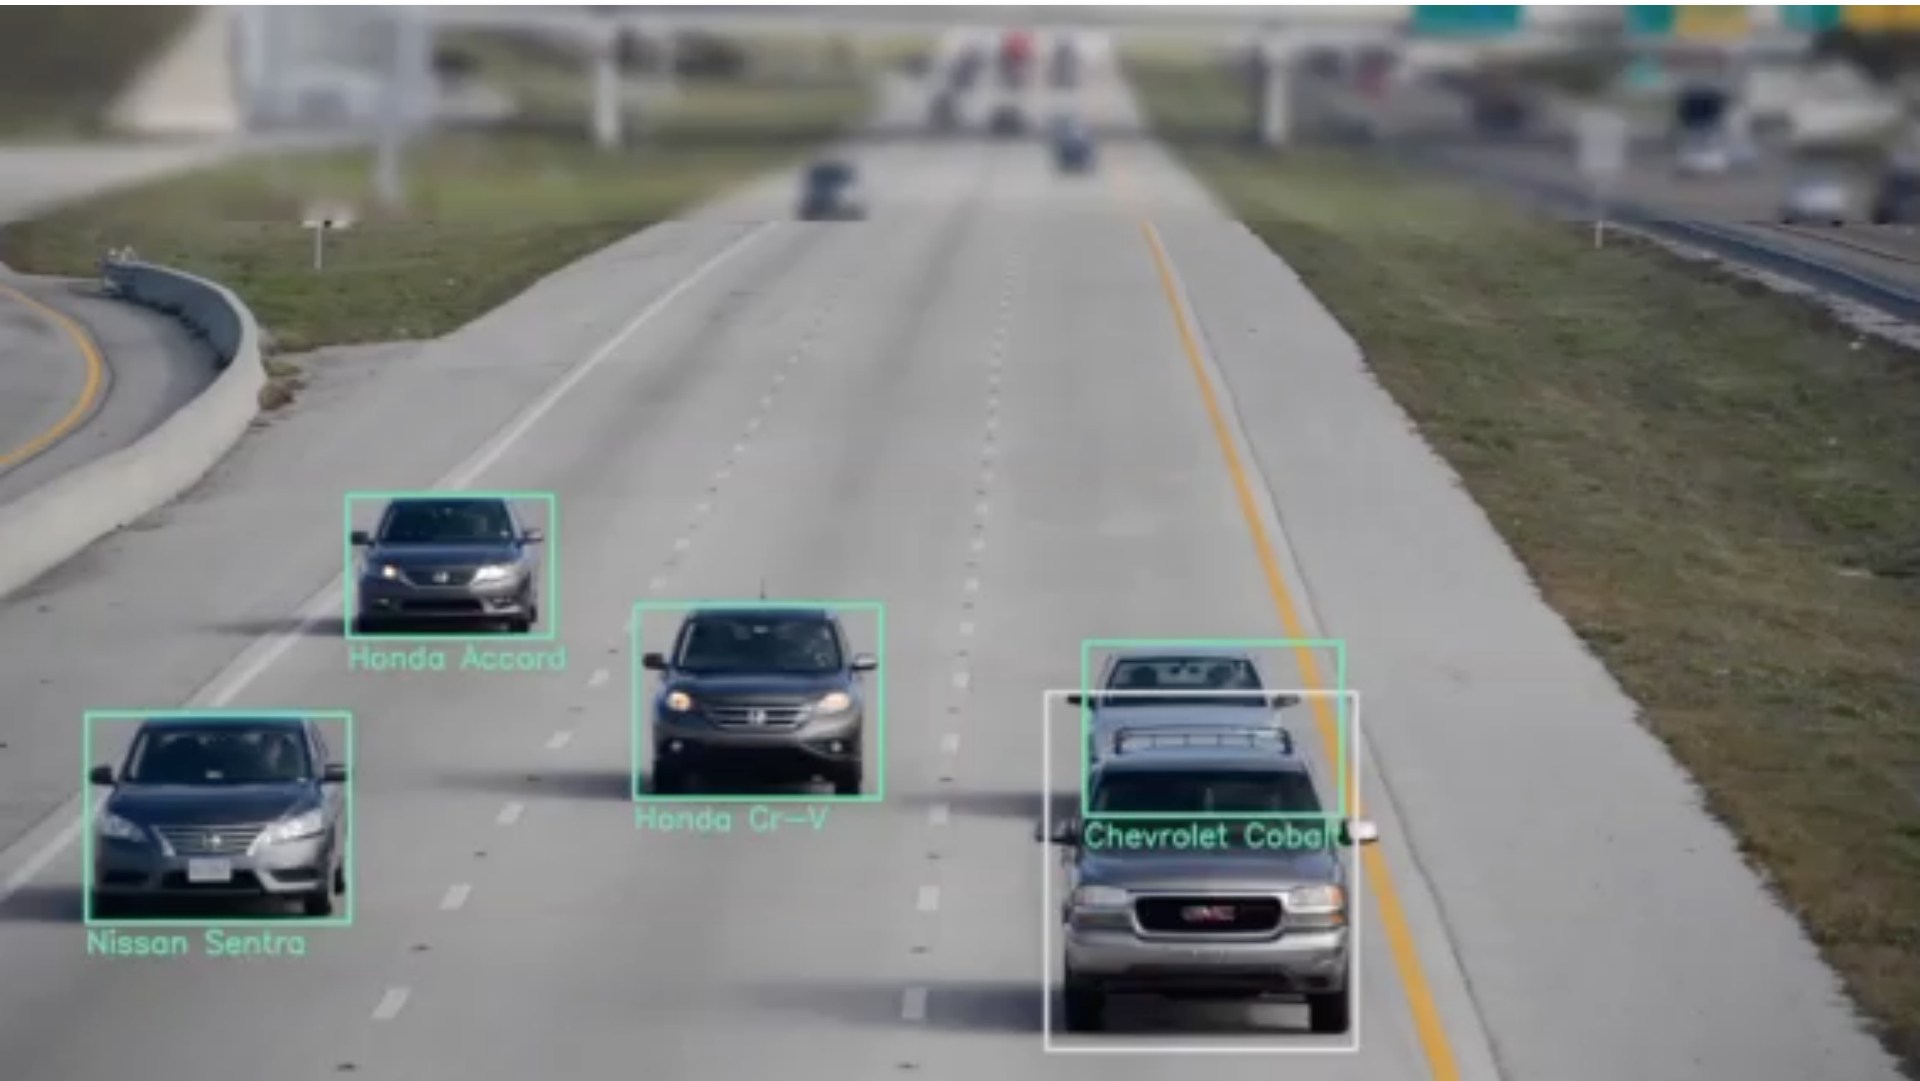
Barak Turovsky
PRODUCT LEAD, GOOGLE TRANSLATE

In 10 years, Google Translate has gone from supporting just a few languages to 103, connecting strangers, reaching across language barriers and even helping

Image segmentation and Object recognition



Autonomous Driving



Achieving Master Level in GO

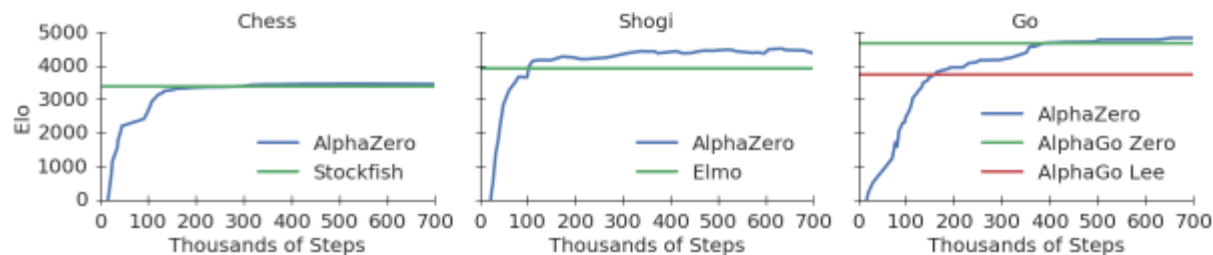


Figure 1: Training *AlphaZero* for 700,000 steps. Elo ratings were computed from evaluation games between different players when given one second per move. **a** Performance of *AlphaZero* in chess, compared to 2016 TCEC world-champion program *Stockfish*. **b** Performance of *AlphaZero* in shogi, compared to 2017 CSA world-champion program *Elmo*. **c** Performance of *AlphaZero* in Go, compared to *AlphaGo Lee* and *AlphaGo Zero* (20 block / 3 day) (29).

Image Captioning

Human captions from the training set



Automatically captioned



DL is Transforming the Industries

- Search engine, with neural ranking
- Recommendation systems, Youtube, Tiktok, Facebook
- Speech Recognition
- Machine Translation
- Healthcare
- Finance
- ...

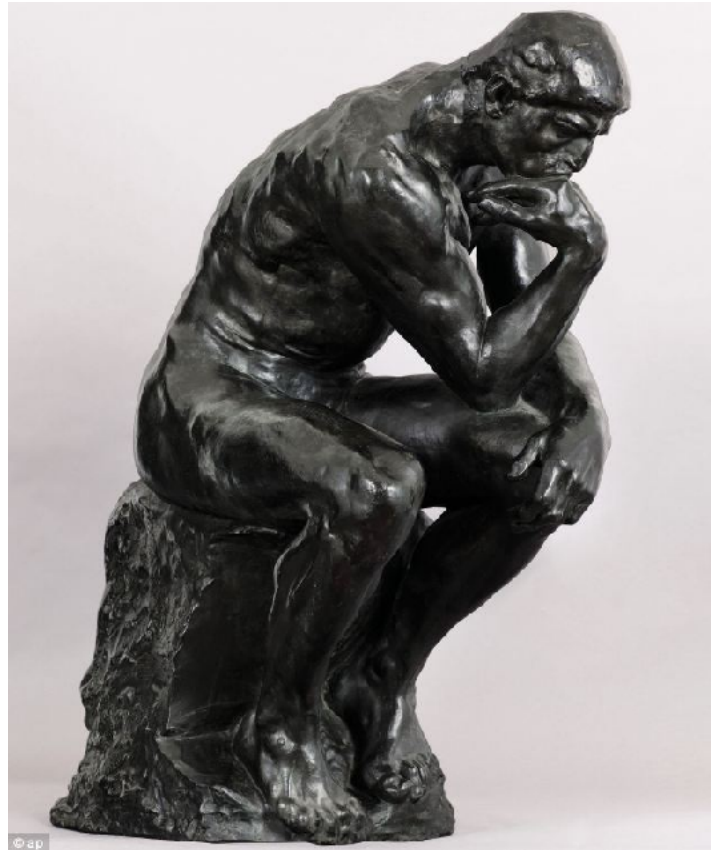
So what are neural networks??

- It begins with this..



So what are neural networks??

- Or even earlier.. with this..



*"The Thinker!"
by Augustin Rodin*

The magical capacity of humans

- Humans can
 - Learn
 - Solve problems
 - Recognize patterns
 - Create
 - Cogitate
 - ...
- Worthy of emulation
- But how do humans “work“?



Dante!

Cognition and the brain..

- “If the brain was simple enough to be understood - we would be too simple to understand it!”
 - Marvin Minsky

Early Models of Human Cognition

- Associationism
 - Humans learn through association
- 400BC-1900AD: Plato, David Hume, Ivan Pavlov..



- **But where are the associations stored??**
- **And how?**

Observation: *The Brain*

- Mid 1800s: The brain is a mass of interconnected neurons



Brain: Interconnected Neurons

- Many neurons connect *in* to each neuron
- Each neuron connects *out* to many neurons



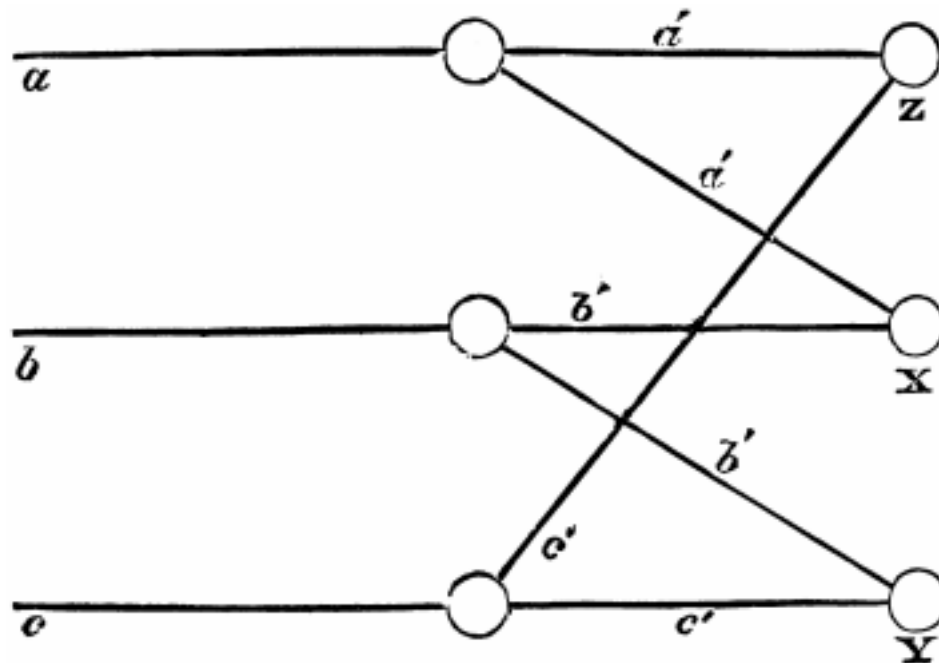
Enter *Connectionism*

- Alexander Bain, philosopher, psychologist, mathematician, logician, linguist, professor
- **1873: The information is in the *connections***
 - *Mind and body* (1873)



Bain's Idea 1: Neural Groupings

- Neurons excite and stimulate each other
- Different combinations of inputs can result in different outputs

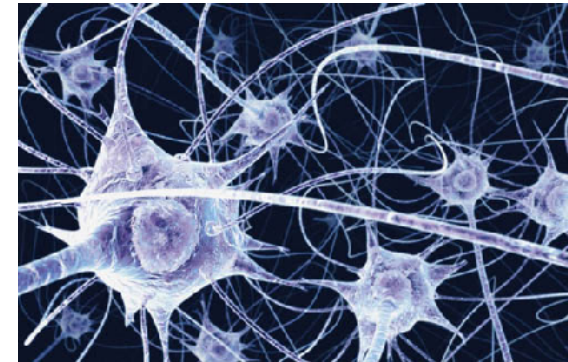


Bain's Idea 2: Making Memories

- “when two impressions concur, or closely succeed one another, the nerve-currents find some bridge or place of continuity, better or worse, according to the abundance of nerve-matter available for the transition.”
- Predicts “Hebbian” learning (three quarters of a century before Hebb!)

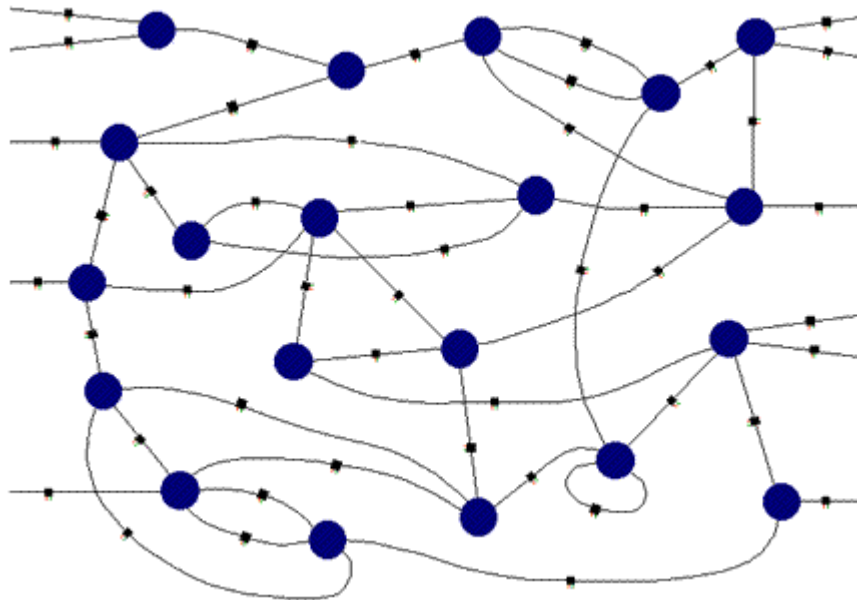
Connectionism lives on..

- The human brain is a connectionist machine
 - Bain, A. (1873). *Mind and body. The theories of their relation*. London: Henry King.
 - Ferrier, D. (1876). *The Functions of the Brain*. London: Smith, Elder and Co
- Neurons connect to other neurons. The processing/capacity of the brain is a function of these connections
- Connectionist machines emulate this structure



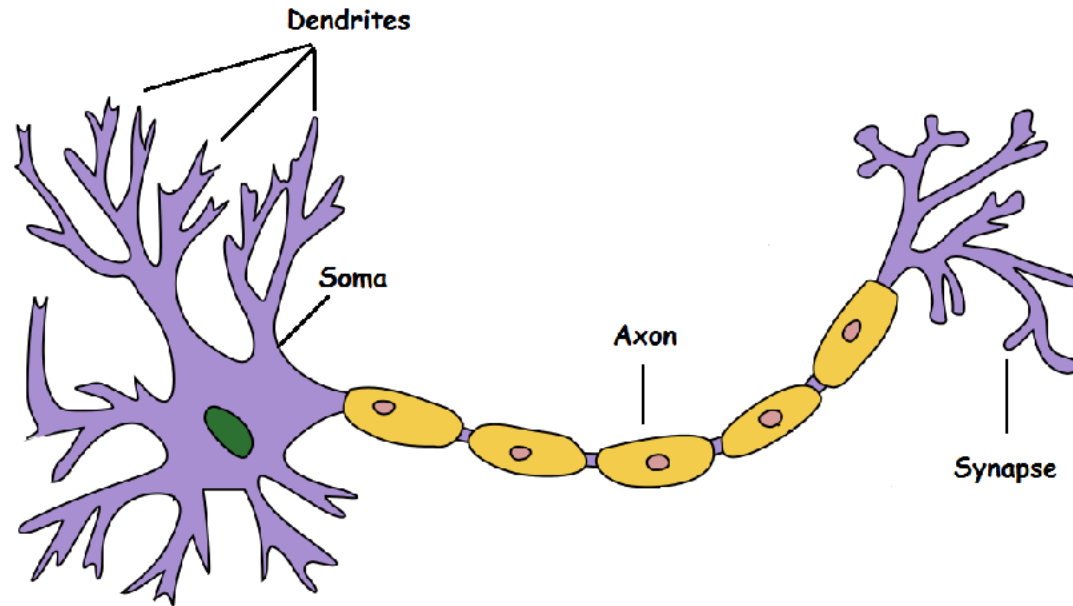
Connectionist Machines

- Network of processing elements
 - All world knowledge is stored in the *connections* between the elements
- *But what are the individual elements?*



Modelling the brain

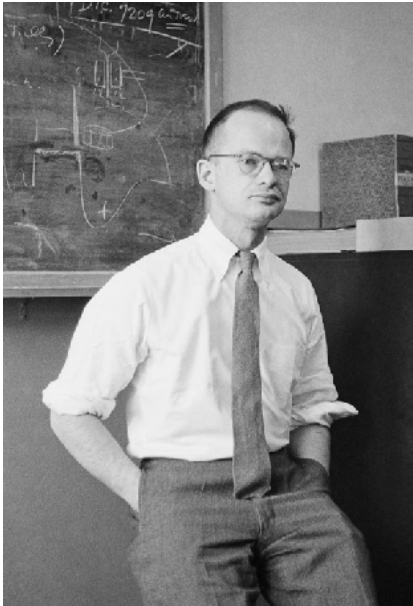
- A neuron:



- Signals come in through the dendrites into the Soma
- A signal goes out via the axon to other neurons
 - Only one axon per neuron
- Factoid that may only interest me: Neurons do not undergo cell division
 - Neurogenesis occurs from neuronal stem cells, and is minimal after birth

McCulloch and Pitts

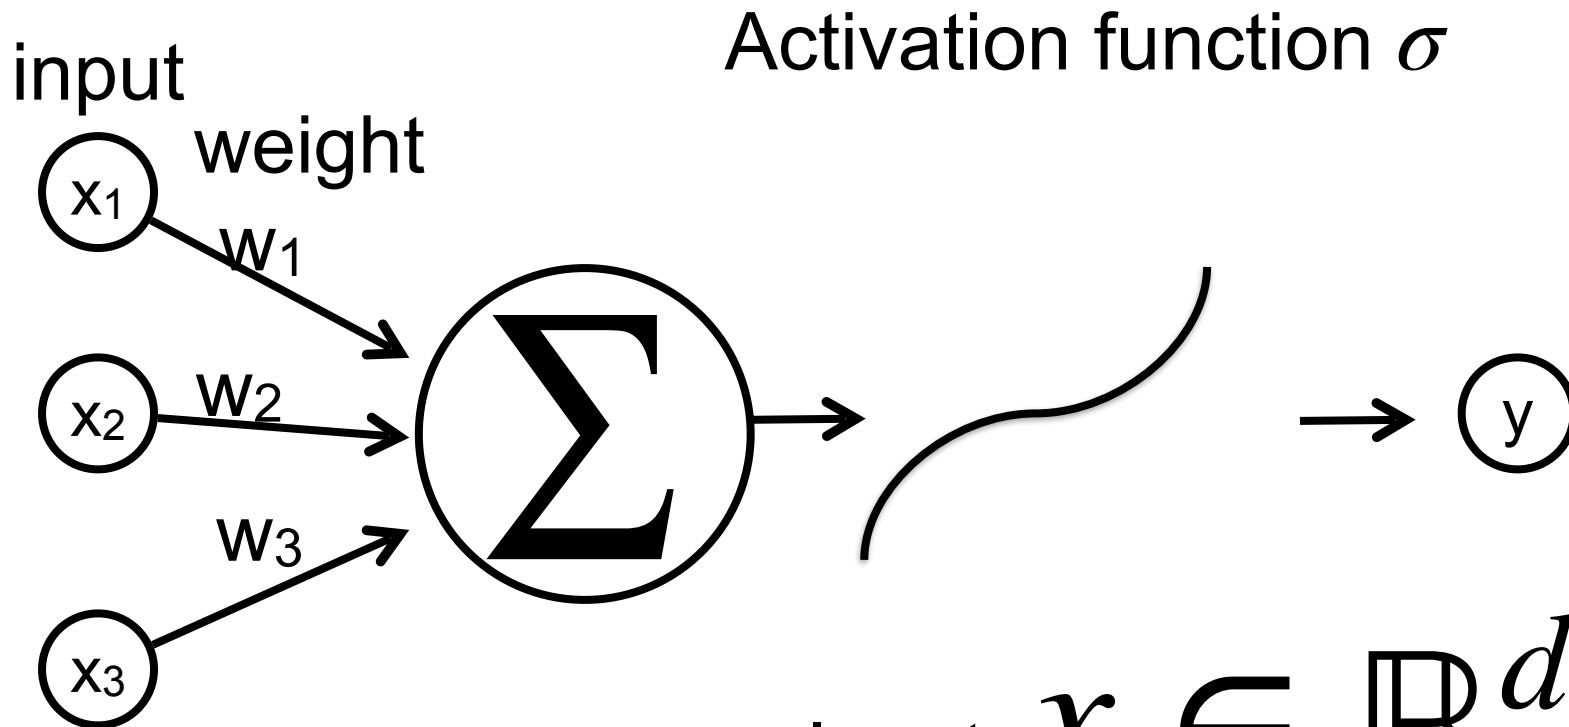
- The Doctor and the Hobo..
 - Warren McCulloch: Neurophysiologist
 - Walter Pitts: Homeless wannabe logician who arrived at his door



The McCulloch and Pitts model

- A mathematical model of a neuron
 - McCulloch, W.S. & Pitts, W.H. (1943). A Logical Calculus of the Ideas Immanent in Nervous Activity, Bulletin of Mathematical Biophysics, 5:115-137, 1943
 - Pitts was only 20 years old at this time

A single Artificial Neuron



Input: $x \in \mathbb{R}^d$

Weight:

$w \in \mathbb{R}^d, b \in \mathbb{R}$

Criticisms

- They claimed that their nets
 - Should be able to compute a small class of functions
 - Also, if tape is provided their nets can compute a richer class of functions.
 - Additionally they will be equivalent to Turing machines
 - Dubious claim that they're Turing complete
 - They didn't prove any results themselves
- Didn't provide a learning mechanism..

Donald Hebb



Novelist, farmer,
hobo, schoolteacher
psychologist

- “Organization of behavior”, 1949
- A learning mechanism:
 - “When an axon of cell *A* is near enough to excite a cell *B* and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that *A*'s efficiency, as one of the cells firing *B*, is increased.”
 - As *A* repeatedly excites *B*, its *ability* to excite *B* improves
 - *Neurons that fire together wire together*

Hebbian Learning

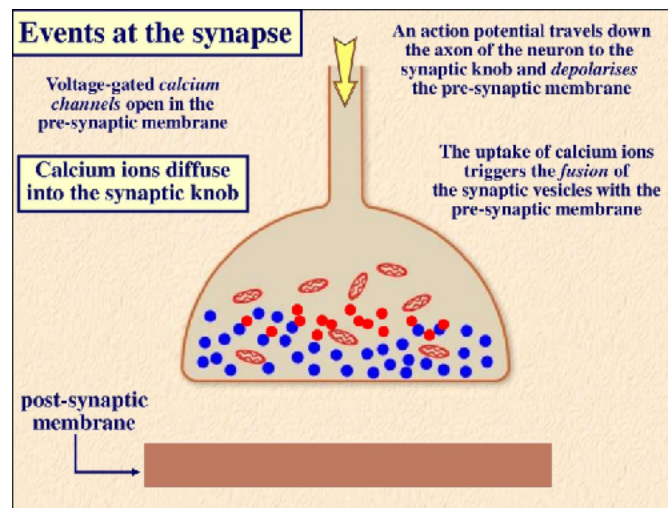
- If neuron x repeatedly triggers neuron y , the synaptic knob connecting x to y gets larger
- In a mathematical model:

$$w_{xy} = w_{xy} + \eta xy$$

- Weight of the connection from input neuron x to output neuron y
- This simple formula is actually the basis of many learning algorithms in ML

Axonal connection from neuron X

Dendrite of neuron Y



Hebbian Learning

- **Fundamentally unstable**

- Stronger connections will enforce themselves
- No notion of “competition”
- No *reduction* in weights
- Learning is unbounded

- Number of later modifications, allowing for weight normalization, forgetting etc.

- E.g. Generalized Hebbian learning, aka Sanger’s rule

$$w_{ij} = w_{ij} + \eta y_j \left(x_i - \sum_{k=1}^j w_{ik} y_k \right)$$

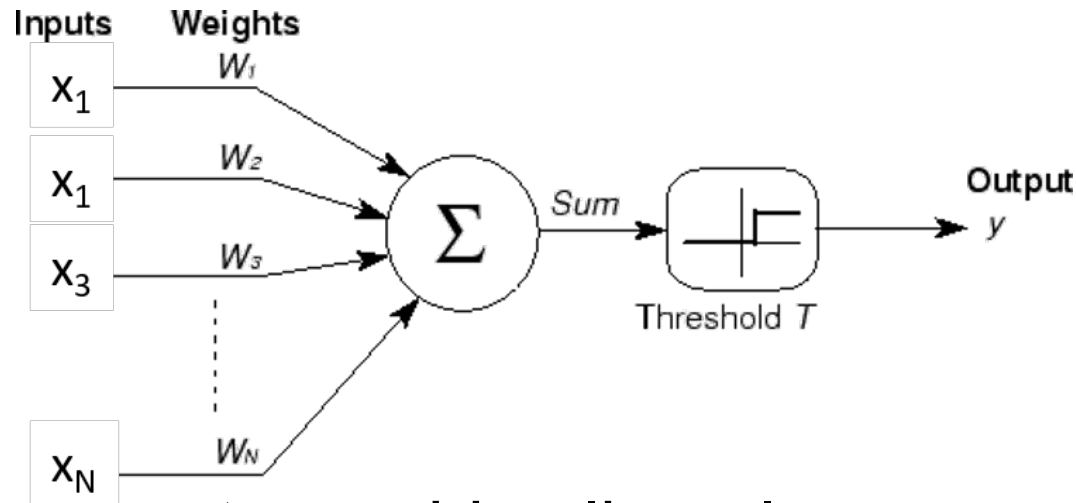
- The contribution of an input is incrementally *distributed* over multiple outputs..

A better model

- Frank Rosenblatt
 - Psychologist, Logician
 - Inventor of the solution to everything, aka the Perceptron (1958)



Perceptron: Simplified model

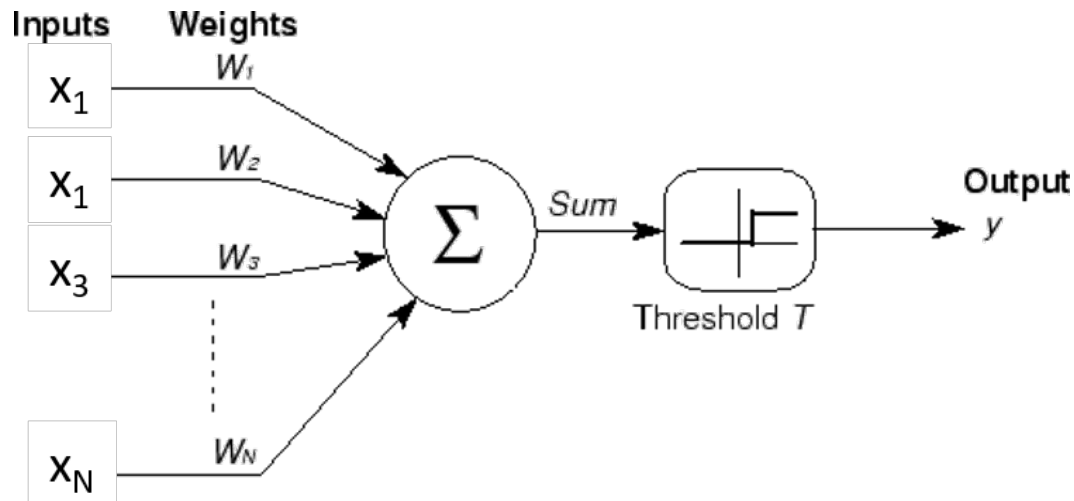


- Number of inputs combine linearly
 - Threshold logic: Fire if combined input exceeds threshold

$$Y = \begin{cases} 1 & \text{if } \sum_i w_i x_i - T \geq 0 \\ 0 & \text{else} \end{cases}$$

The Universal Model

- Originally assumed could represent *any* Boolean circuit and perform any logic
 - “*the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence,*” New York Times (8 July) 1958
 - “*Frankenstein Monster Designed by Navy That Thinks,*” Tulsa, Oklahoma Times 1958



Also provided a learning algorithm

- Boolean tasks
- Update the weights whenever the perceptron output is wrong
 - Update the weight by the product of the input and the *error* between the desired and actual outputs
- Proved convergence for linearly separable classes

$$\mathbf{w} = \mathbf{w} + \eta (d(\mathbf{x}) - y(\mathbf{x})) \mathbf{x}$$

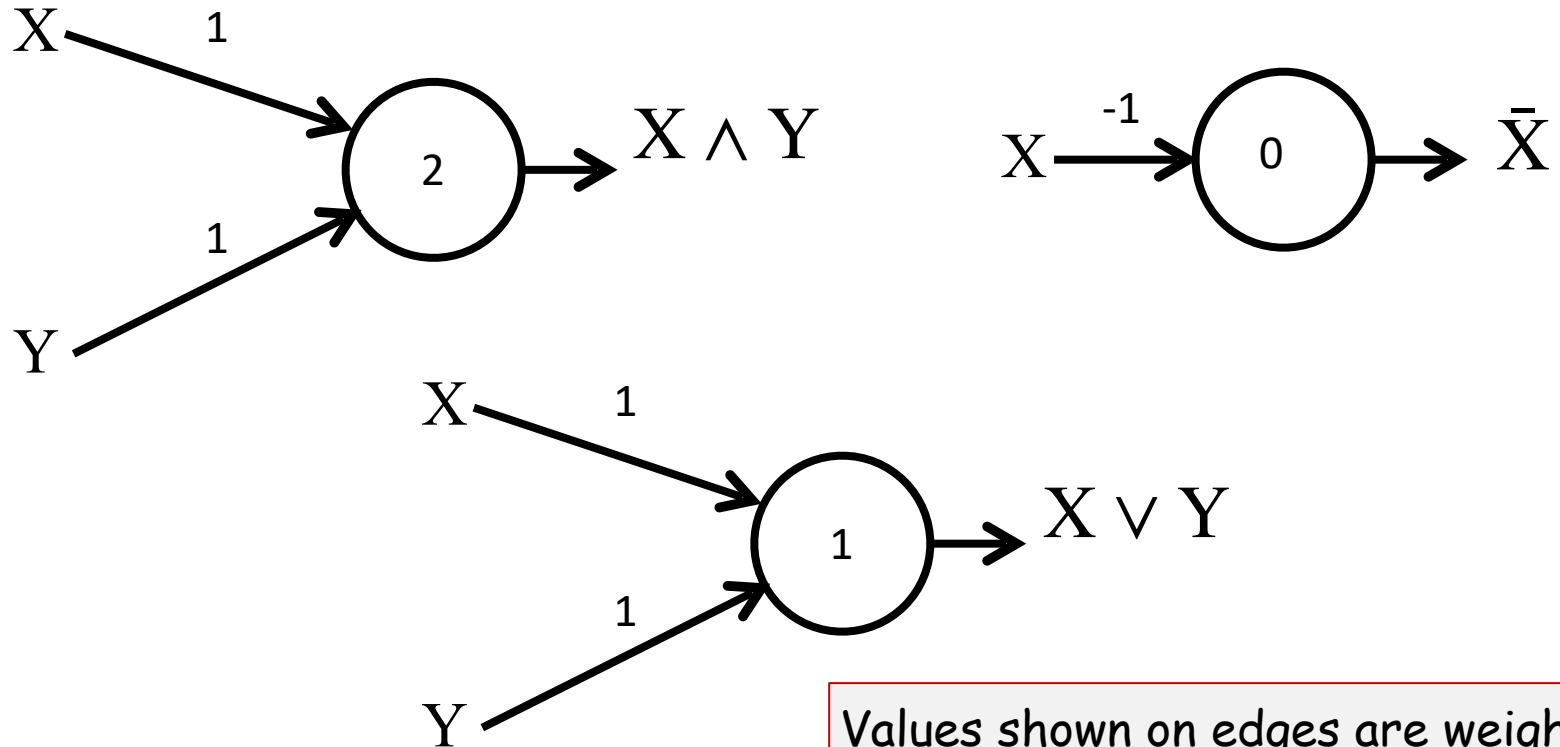
Sequential Learning:

$d(x)$ is the desired output in response to input \mathbf{x}

$y(x)$ is the actual output in response to \mathbf{x}

Perceptron

- Easily shown to mimic any Boolean gate

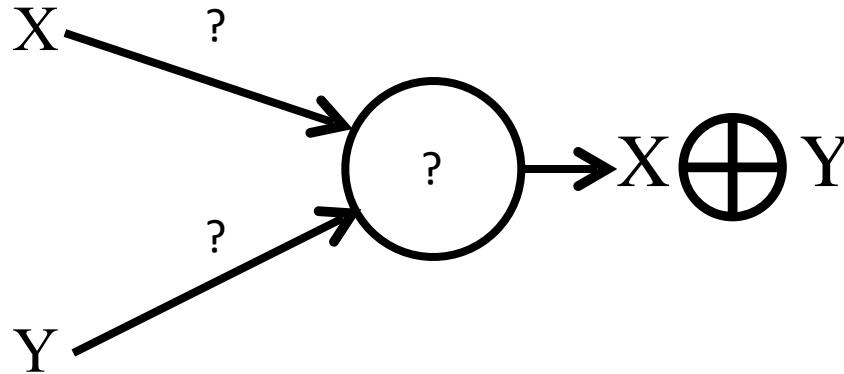


- But...

Values shown on edges are weights,
numbers in the circles are thresholds

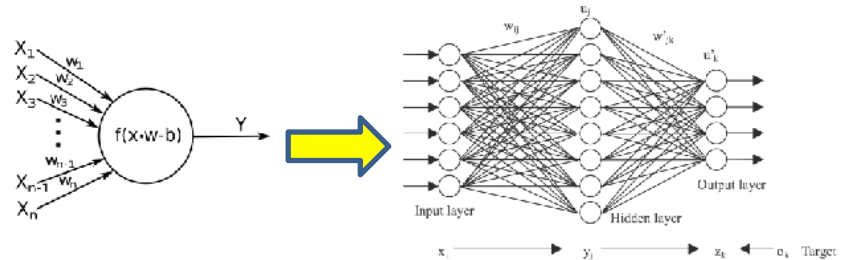
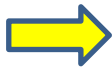
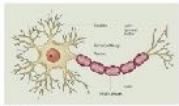
Perceptron

- Minsky and Papert, 1968
No solution for XOR!
Not universal!



A single neuron is not enough

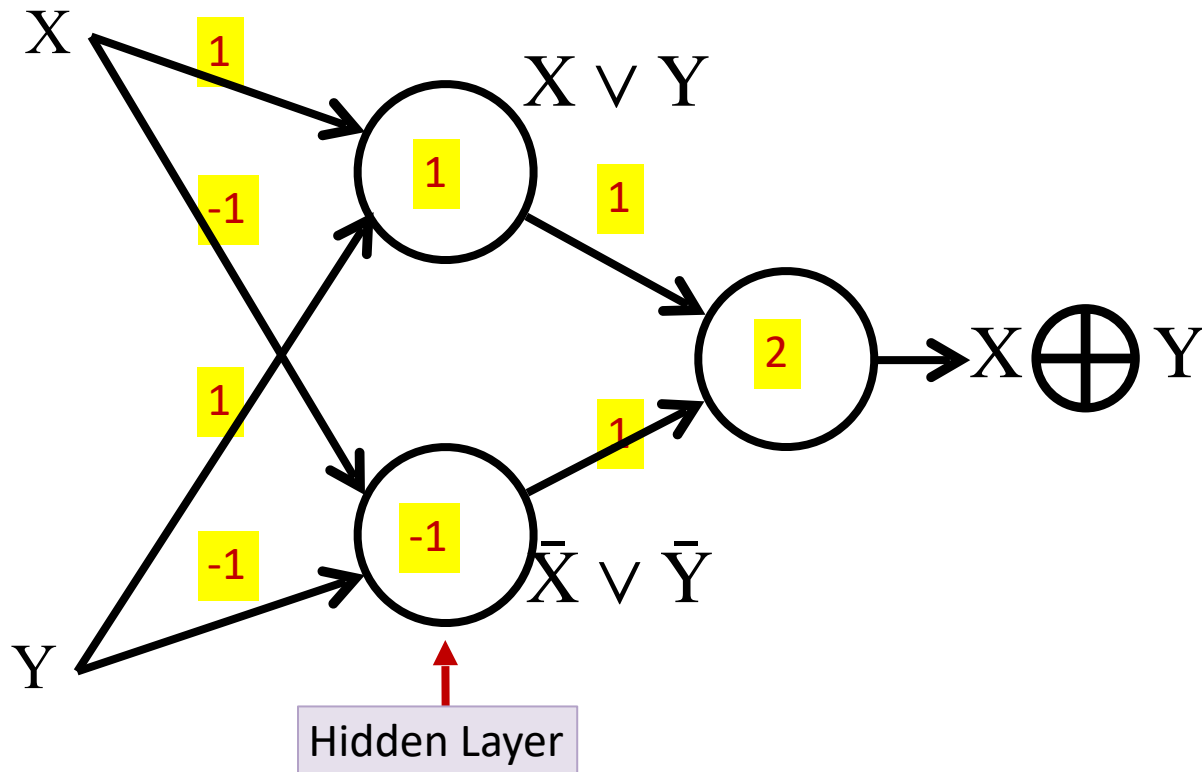
- Individual elements are weak computational elements
 - Marvin Minsky and Seymour Papert, 1969, *Perceptrons: An Introduction to Computational Geometry*
- *Networked* elements are required



Multi-layer Perceptron!

- XOR

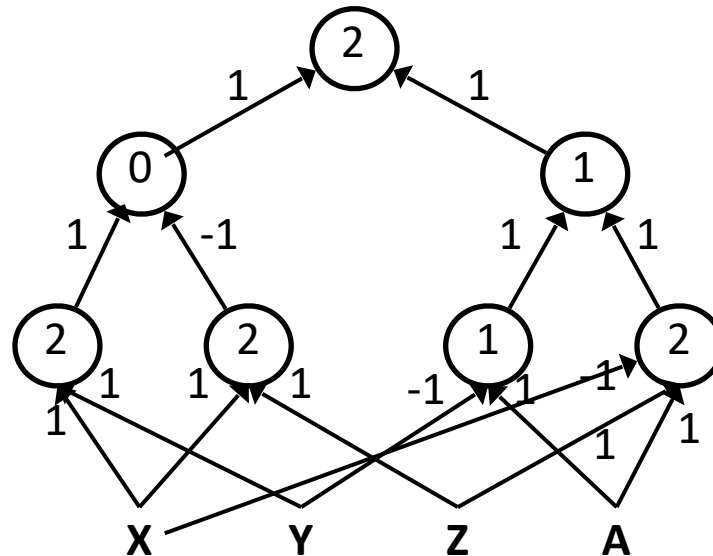
- The first layer is a “hidden” layer
- Also originally suggested by Minsky and Papert 1968



A more generic model

- A “multi-layer” perceptron
- Can compose arbitrarily complicated Boolean functions!
 - In cognitive terms: Can compute arbitrary Boolean functions over sensory input
 - More on this in the next class

$$((A\bar{X}Z) | (A\bar{Y}))((X Y) | (X\bar{Z}))$$



Story so far

- Neural networks began as computational models of the brain
- Neural network models are *connectionist machines*
 - They comprise networks of neural units
- McCulloch and Pitt model: Neurons as Boolean threshold units
 - Models the brain as performing propositional logic
 - But no learning rule
- Hebb's learning rule: Neurons that fire together wire together
 - Unstable
- Rosenblatt's perceptron : A variant of the McCulloch and Pitt neuron with a provably convergent learning rule
 - But individual perceptrons are limited in their capacity (Minsky and Papert)
- Multi-layer perceptrons can model arbitrarily complex Boolean functions

Next Up

- What is Machine learning
- Linear Models
- More on neural networks as universal approximators
 - And the issue of depth in networks
 - How to train neural network from data