Deep Learning

with Artificial Neural Networks

Jon Krohn
jon@untapt.com

Chief Data Scientist at untapt

Metis — January 17th, 2017
(slides available at jonkrohn.com/talks)

Outline

Outline

Outline

Outline

Outline

Jeanne Calment

Jeanne Calment

(1875-1997 — i.e., 122 years)

Life in the Year 2138

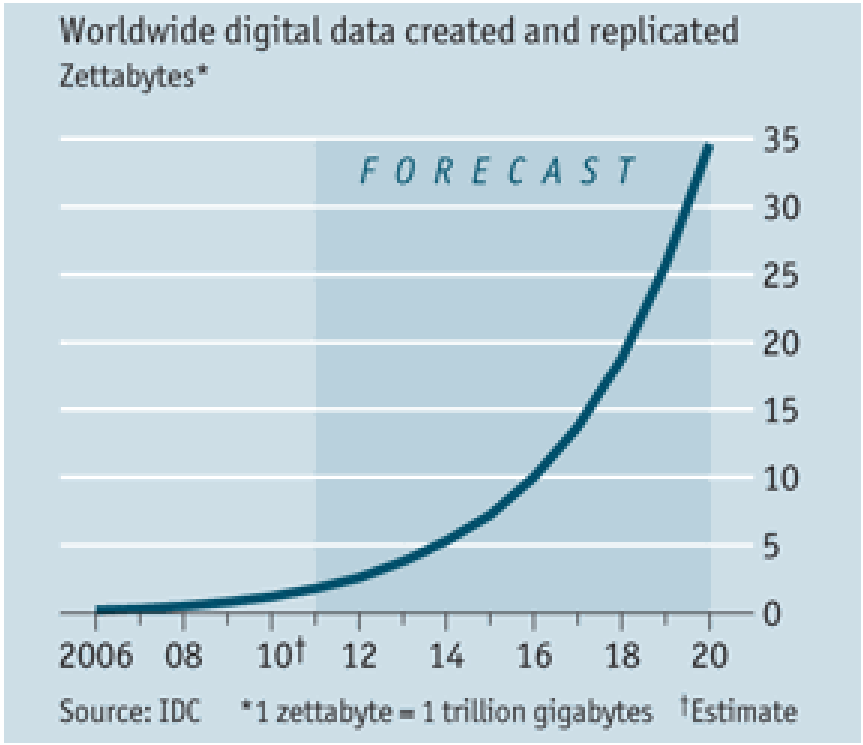
Worldwide digital data created and replicated
Zettabytes*
FORECAST
35
30
25
20
15
10
5
0
2006
08
10†
12
14
16
18
20
Source: IDC
*1 zettabyte = 1 trillion gigabytes
†Estimate

Outline

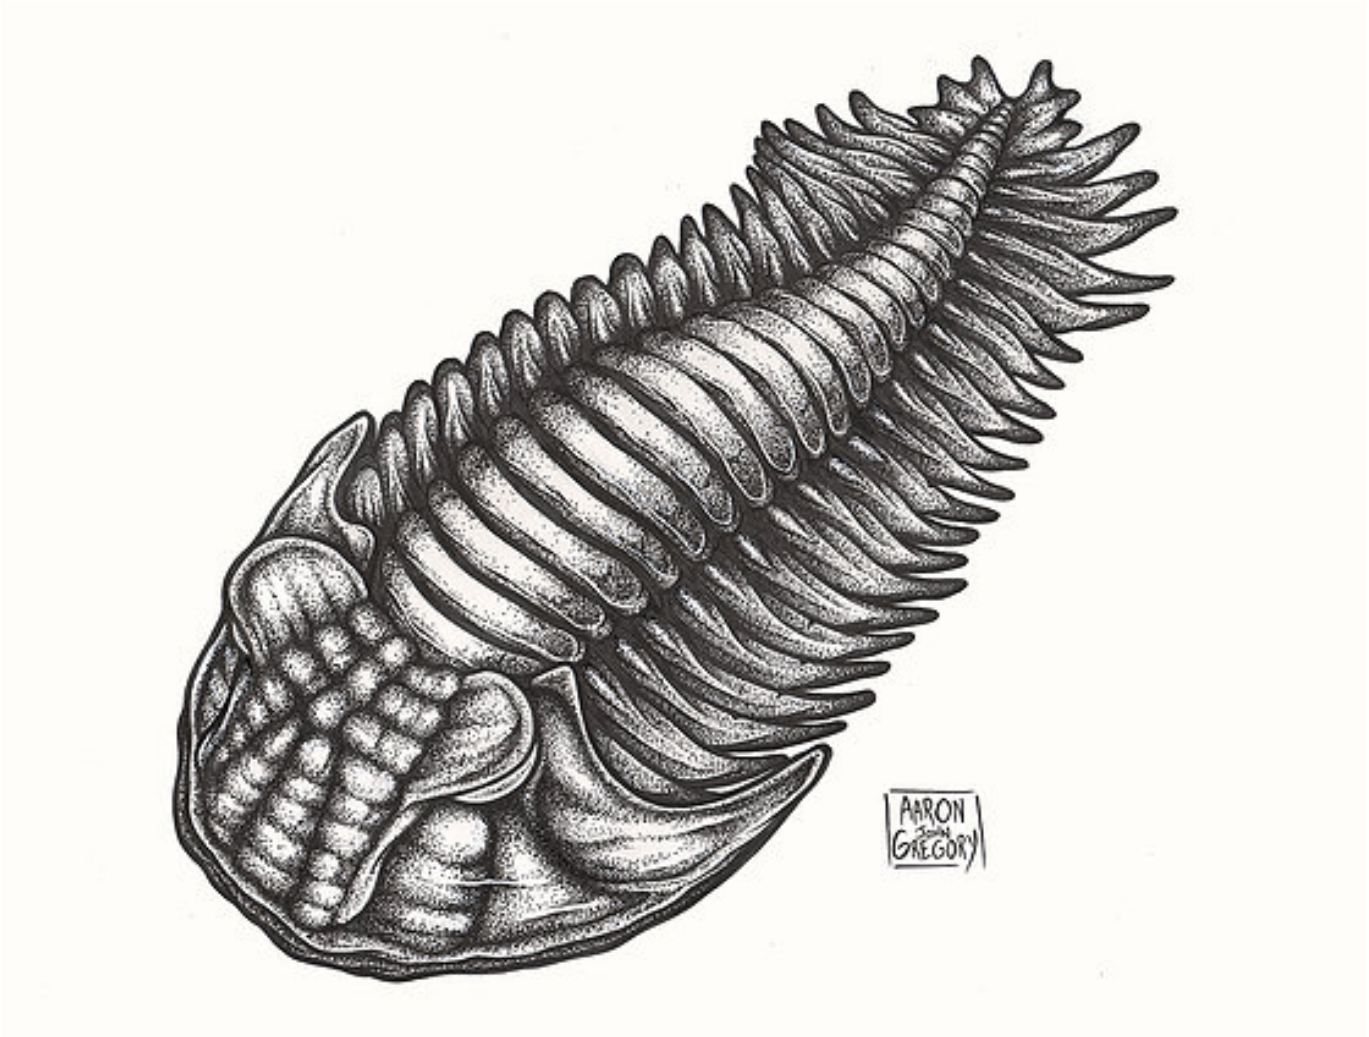
AARON
GREGORY

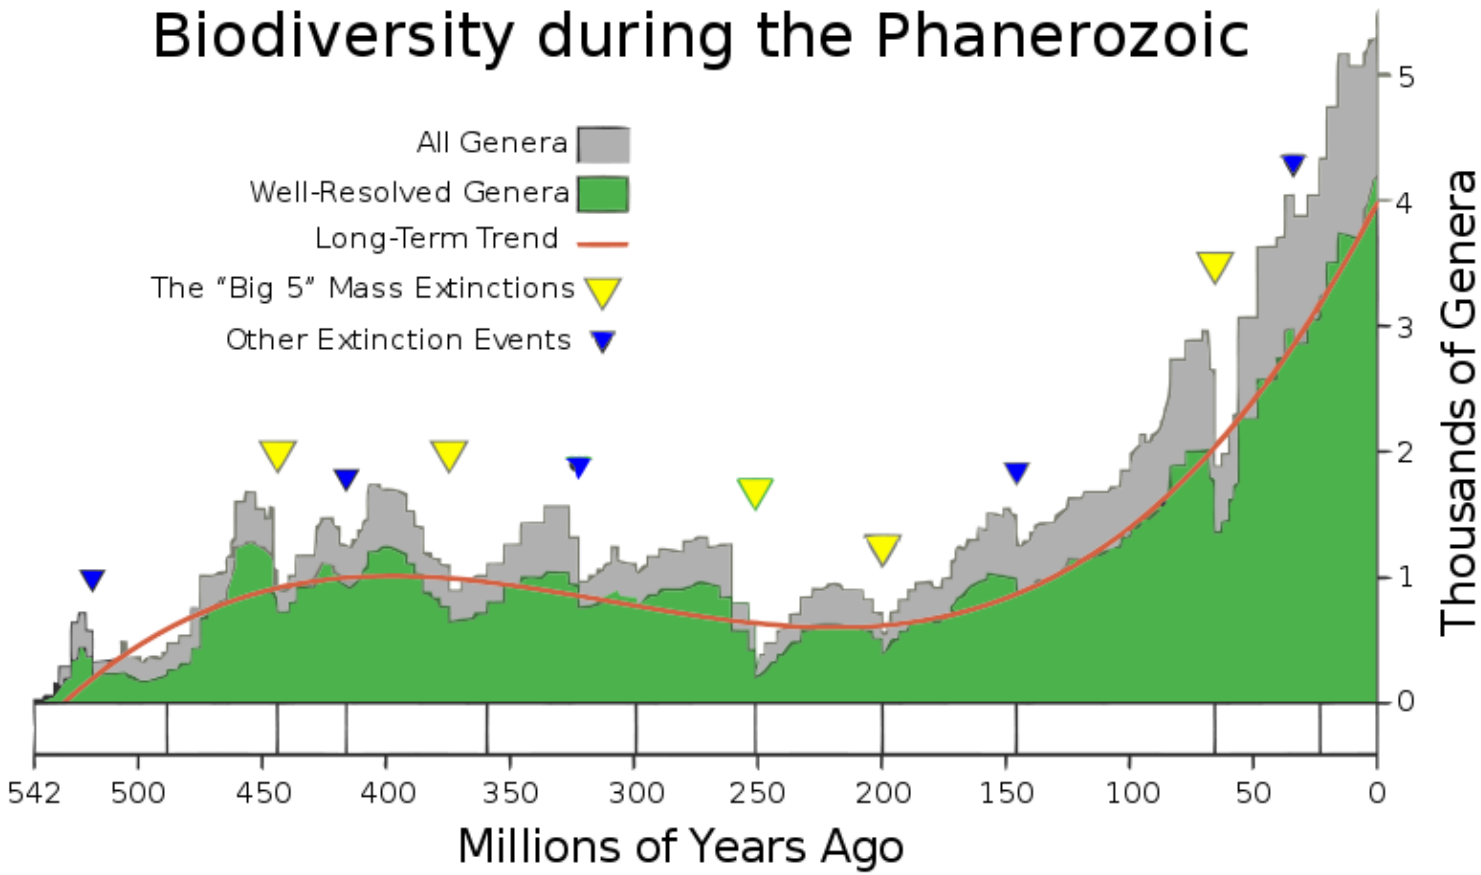
Biodiversity during the Phanerozoic
All Genera
Well-Resolved Genera
Long-Term Trend
The "Big 5" Mass Extinctions
Other Extinction Events
542
500
450
400
350
300
250
200
150
100
50
0
Millions of Years Ago
0
1
2
3
4
5
Thousands of Genera

Hubel & Wiesel (1959)

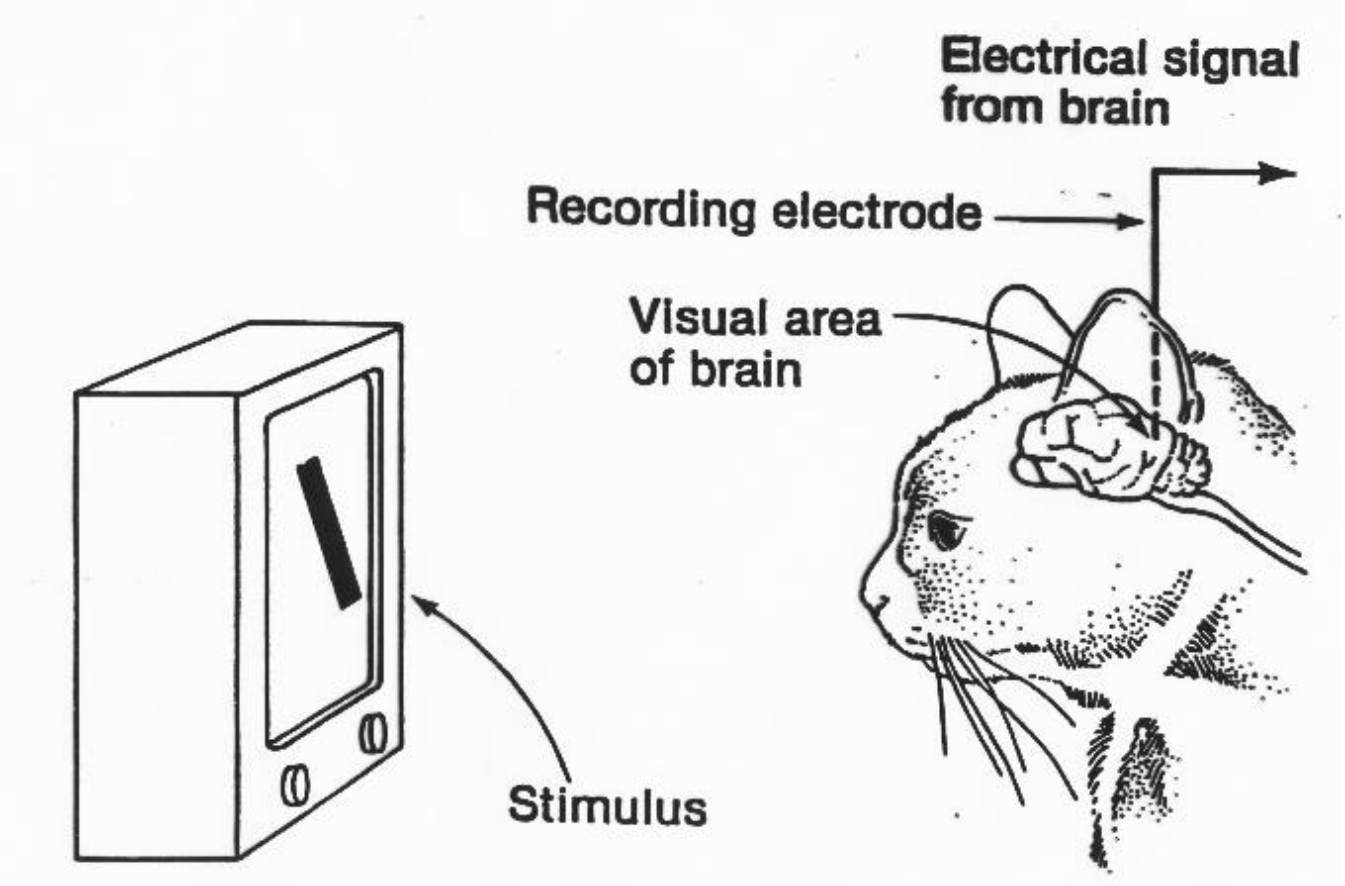

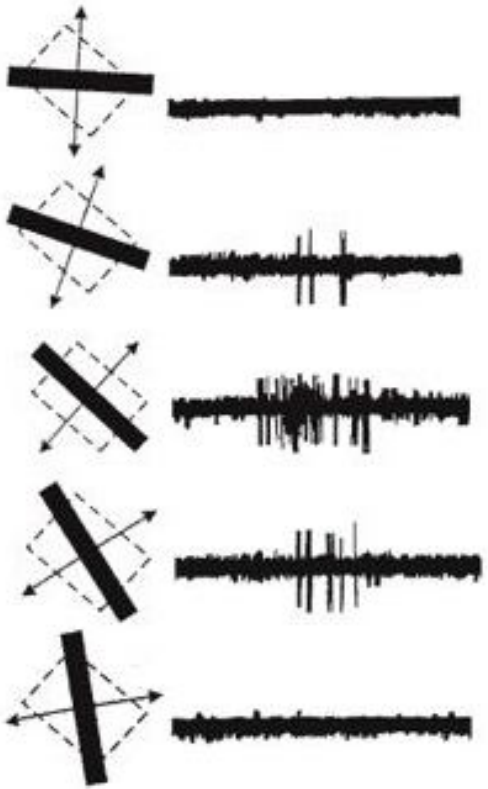

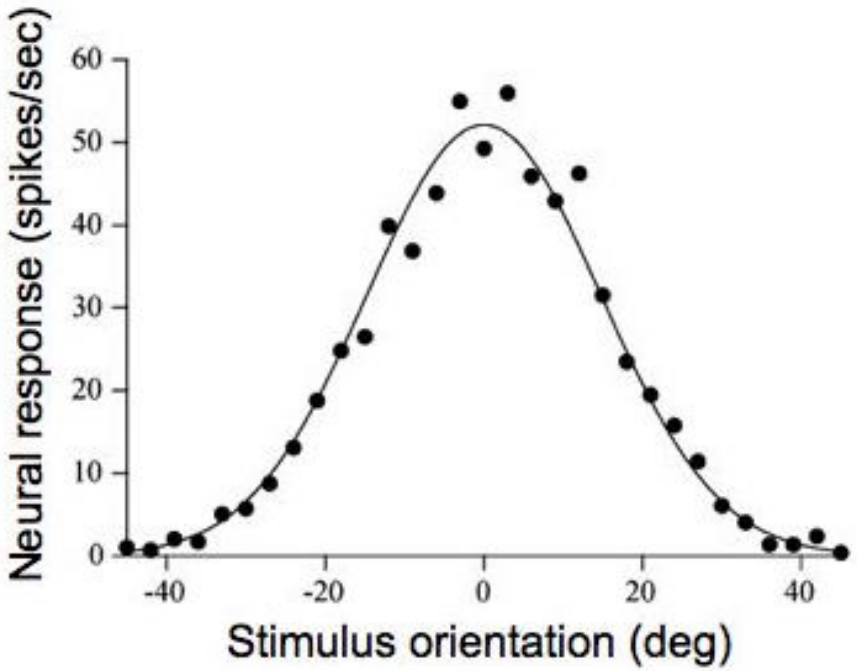

Hubel & Wiesel, 1968

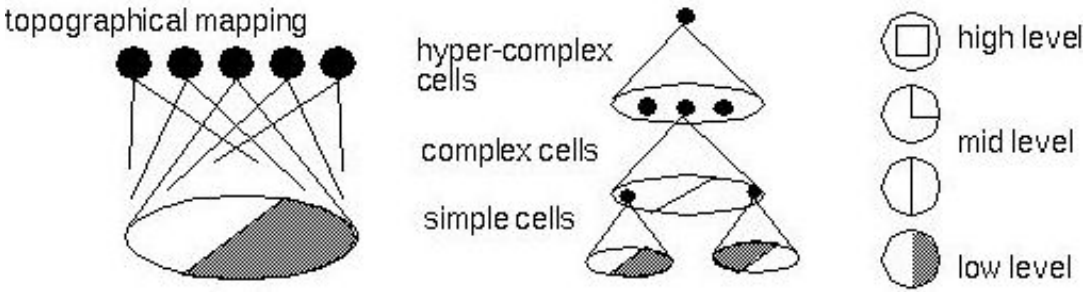

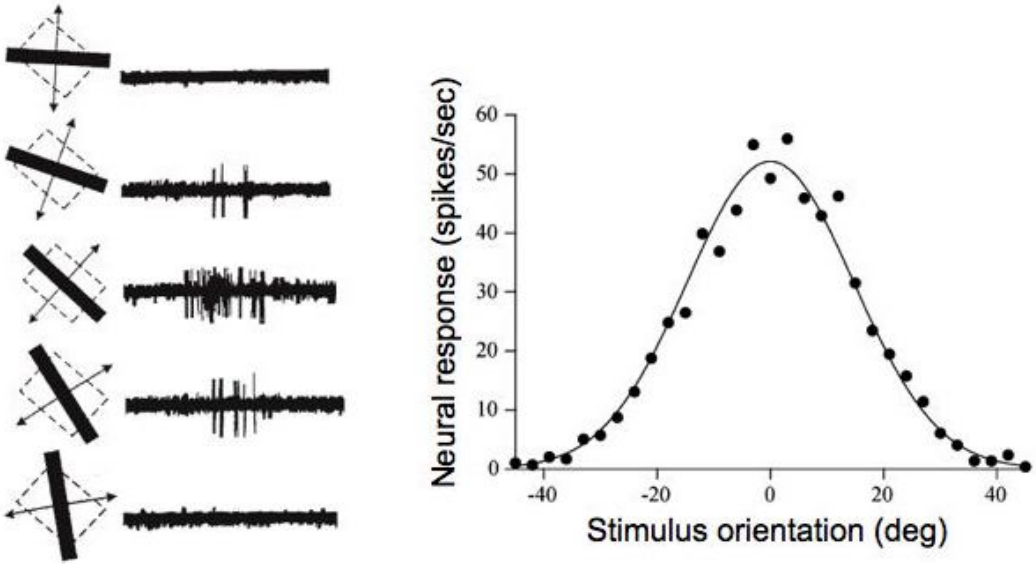

Hubel & Wiesel, 1968

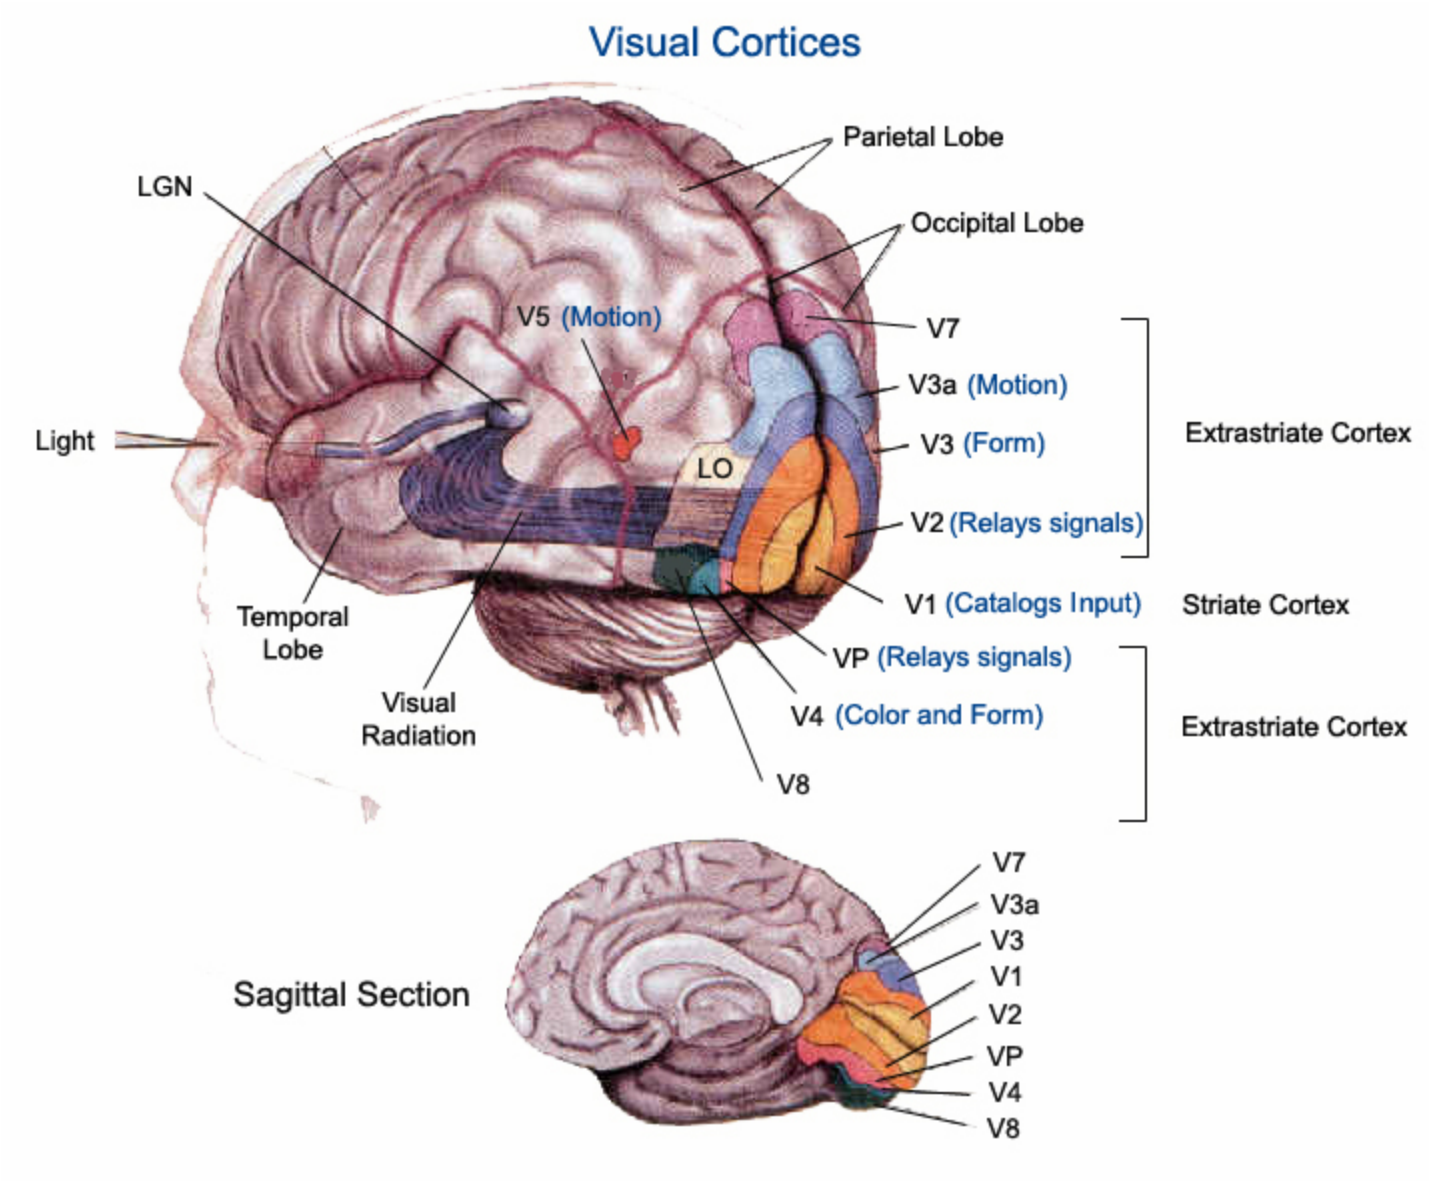
Visual Cortices
Parietal Lobe
LGN
Occipital Lobe
V5 (Motion)
V7
V3a (Motion)
Light
V3 (Form)
LO
Extrastriate Cortex
V2 (Relays signals)
Temporal Lobe
V1 (Catalogs Input)
Striate Cortex
VP (Relays signals)
Visual Radiation
V4 (Color and Form)
Extrastriate Cortex
V8
Sagittal Section
V7
V3a
V3
V1
V2
VP
V4
V8

Camera Obscura

da Vinci (15th Century)

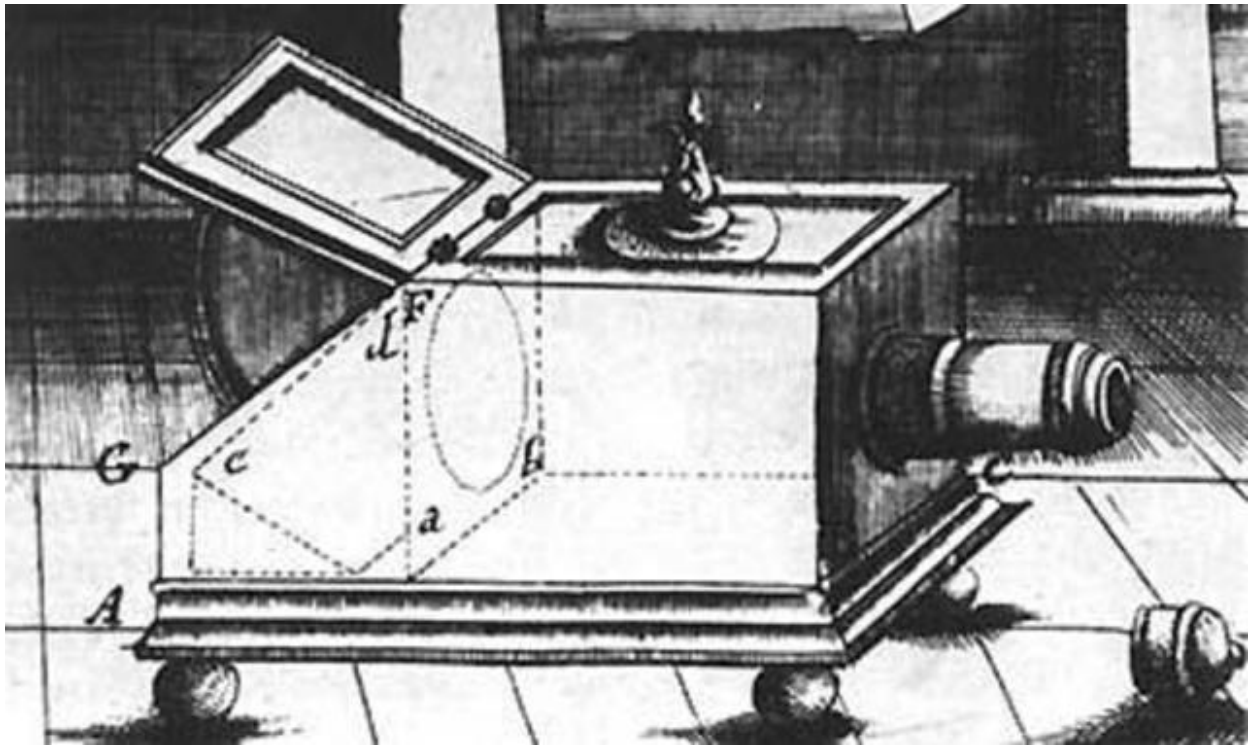

Block World

Larry Roberts (1965)

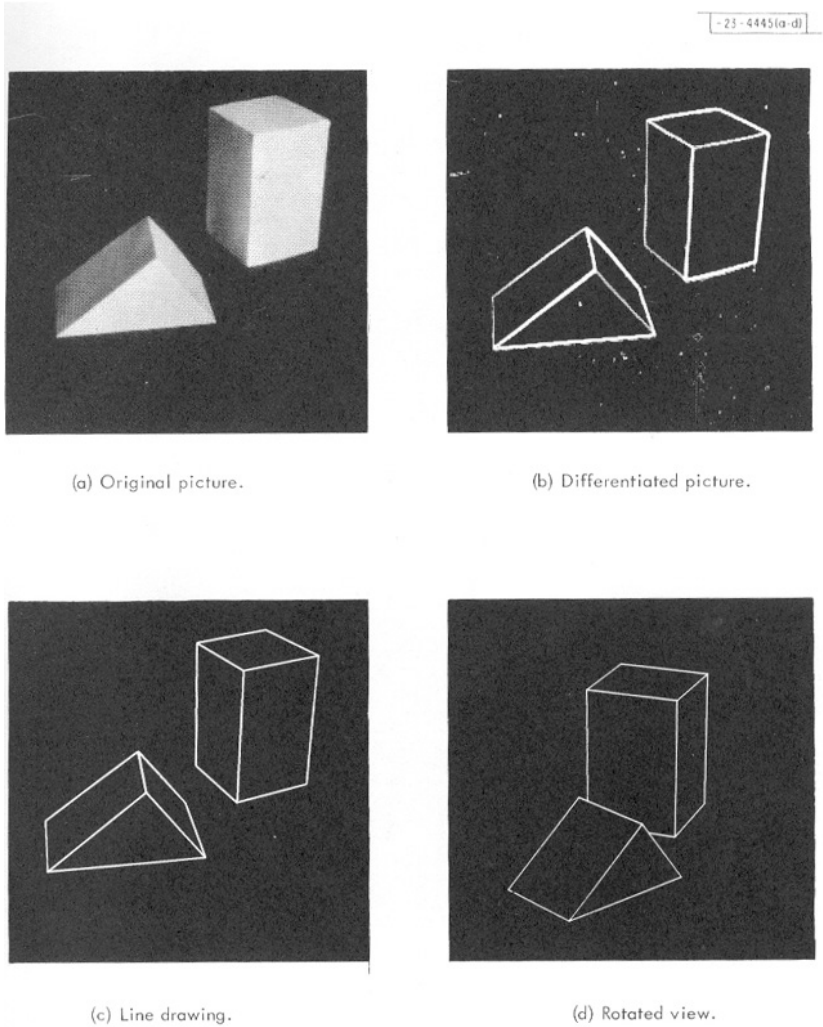

(a) Original picture.

(b) Differentiated picture.

(c) Line drawing.

(d) Rotated view.

Viola & Jones (2001)

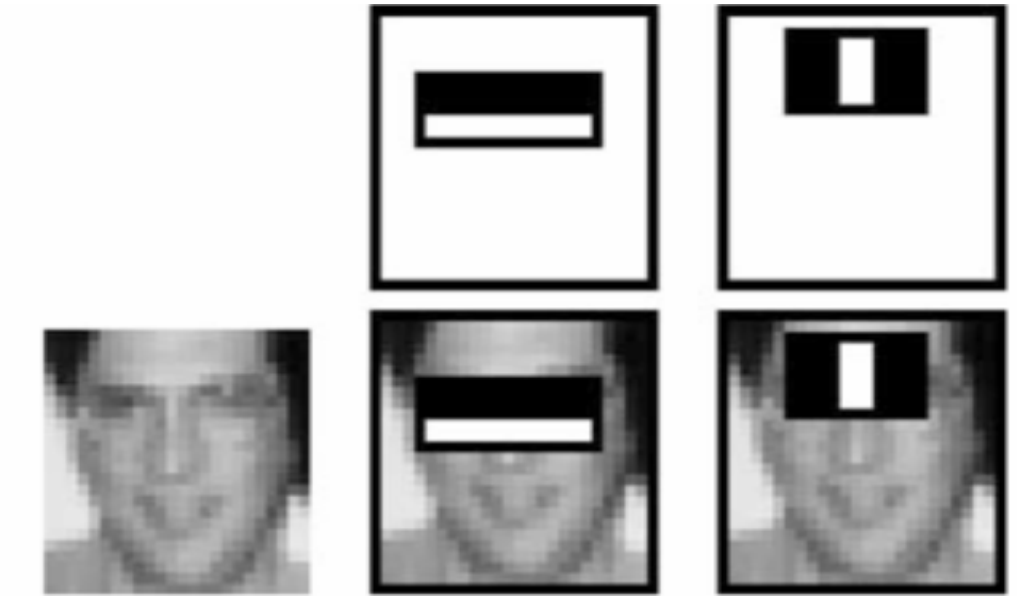

Neurocognitron

Fukushima (1980)

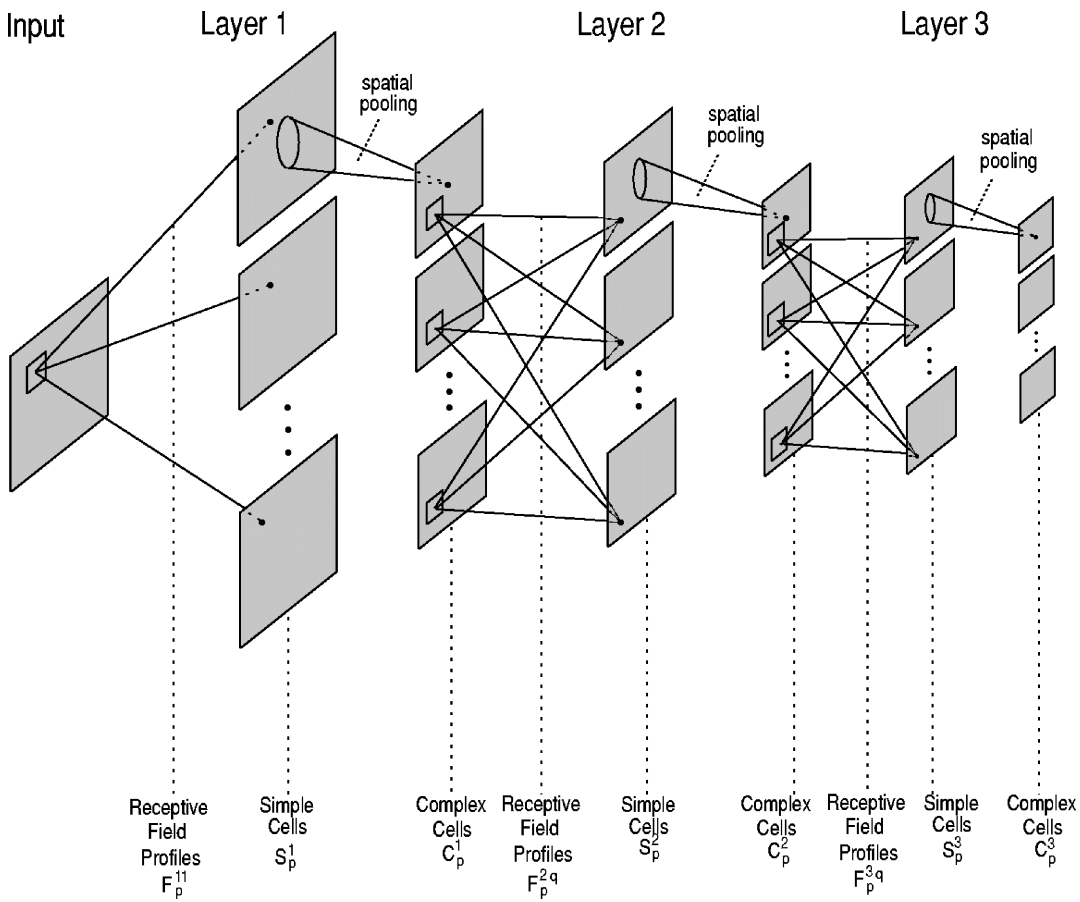

MNIST & LeNet-5

LeCun et al. (1998)

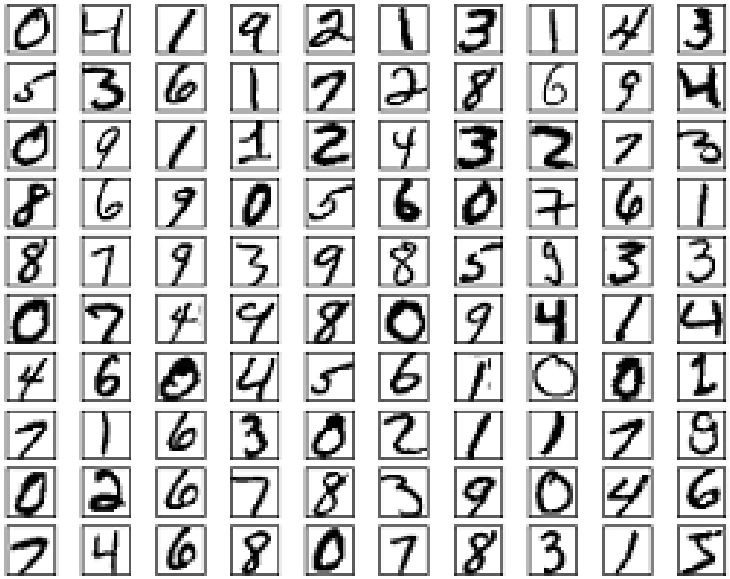

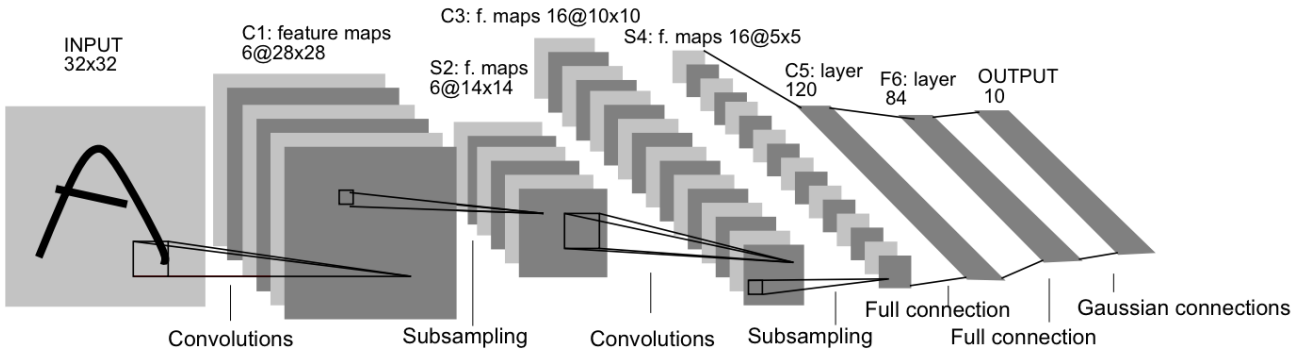

Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Boutou, Bengio & Haffner (1998)

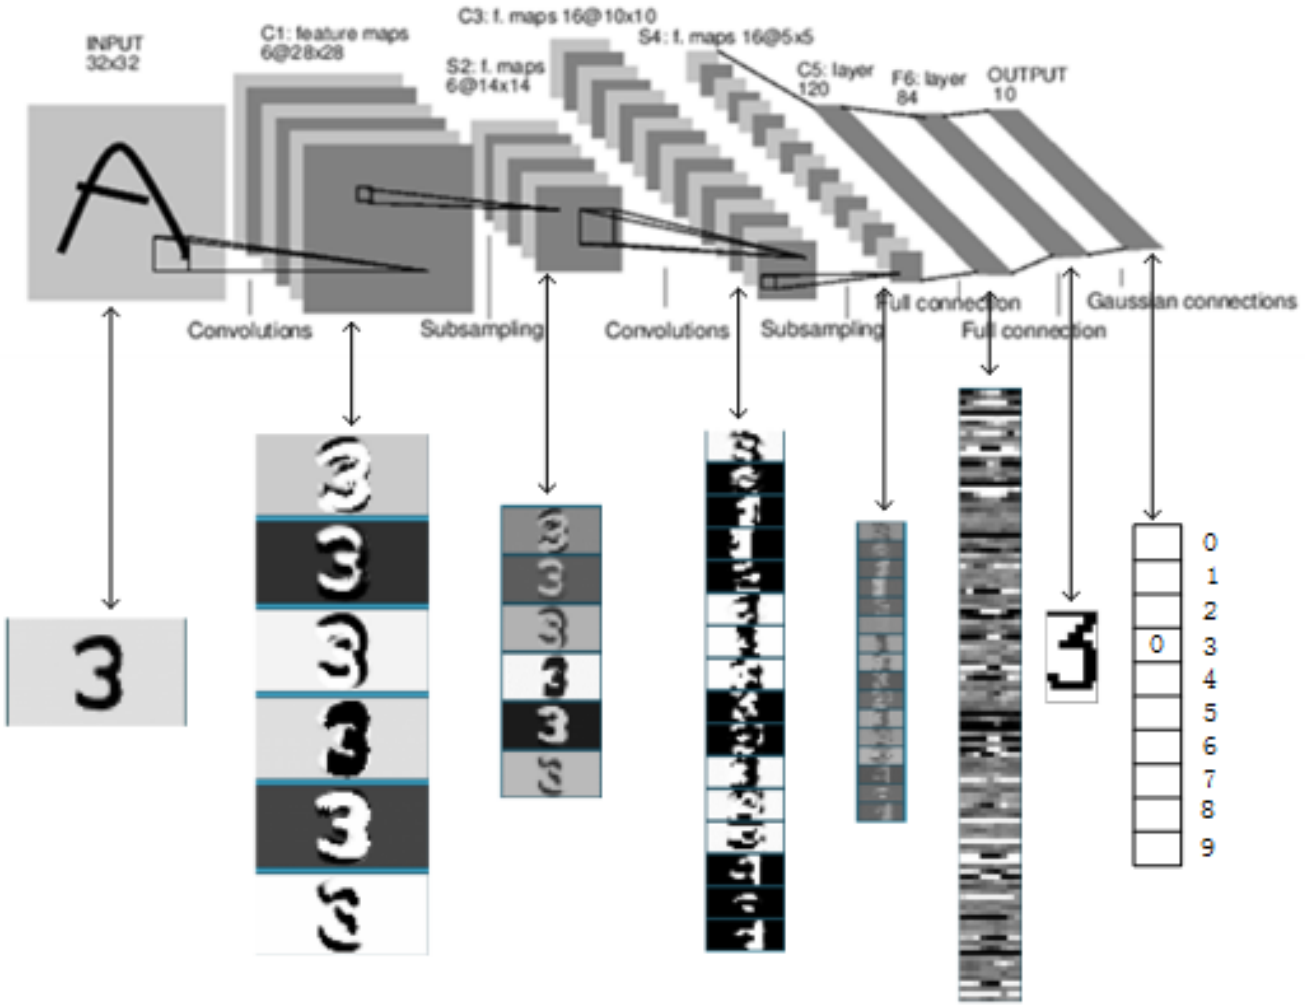

ImageNet

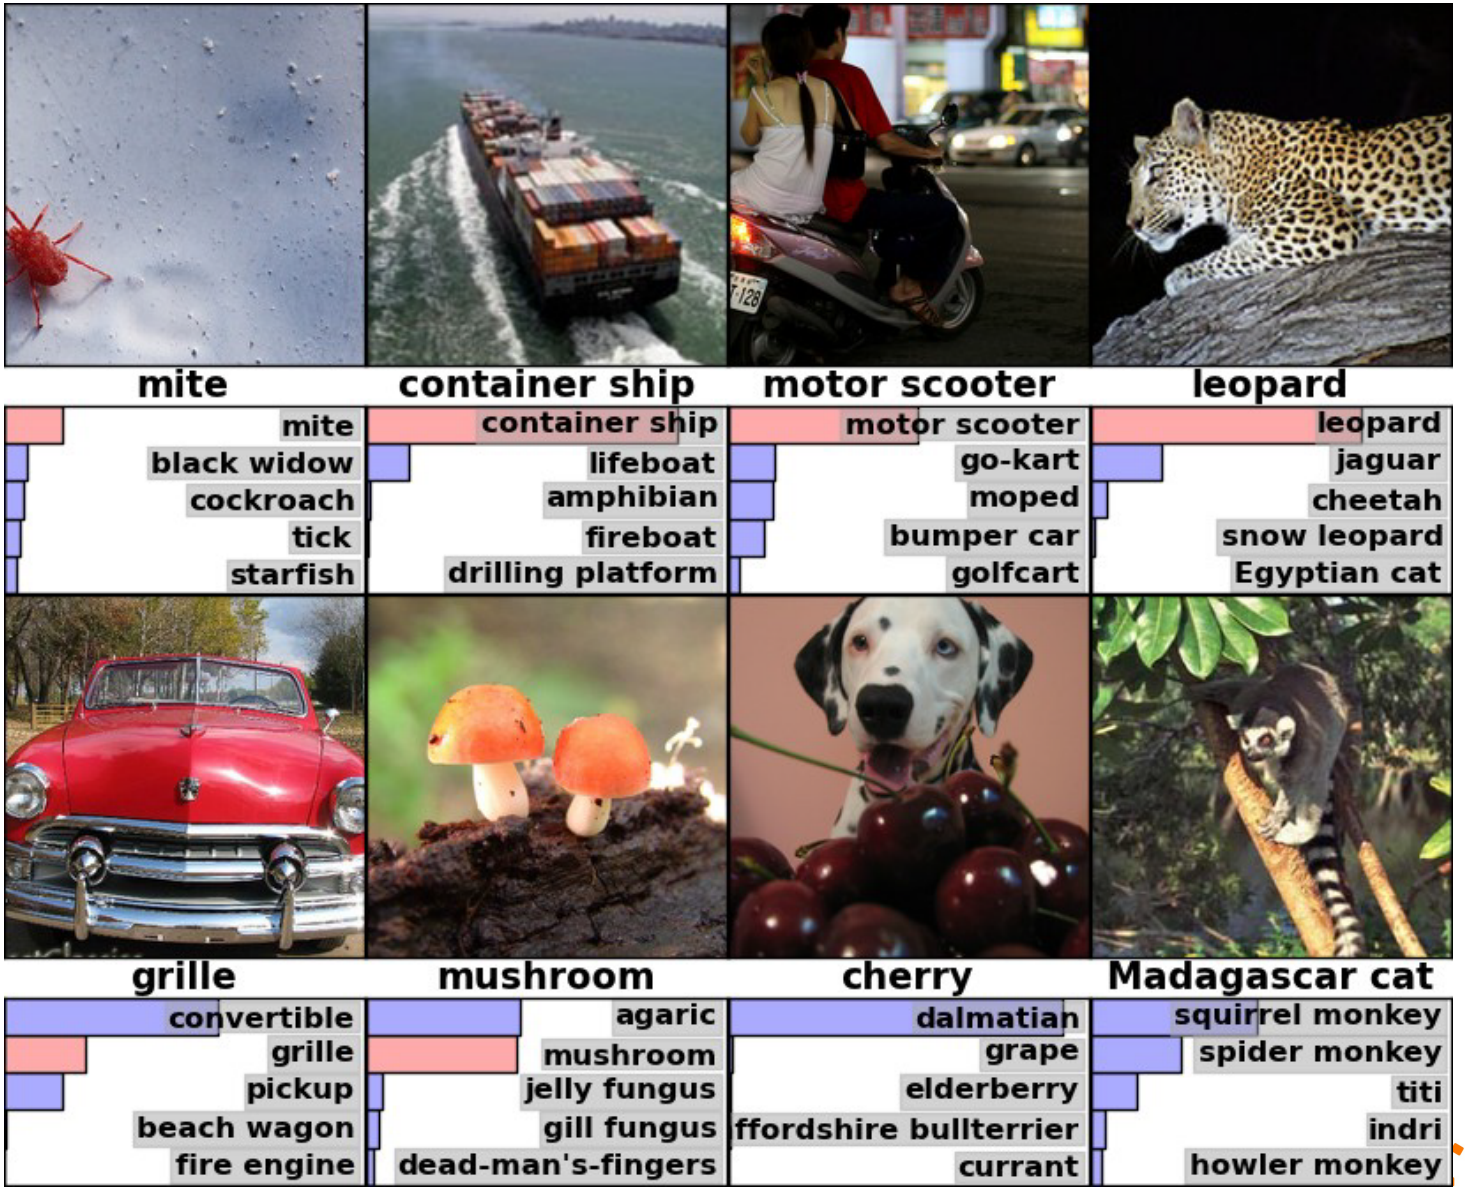
mite
mite
black widow
cockroach
tick
starfish
container ship
container ship
lifeboat
amphibian
fireboat
drilling platform
motor scooter
motor scooter
go-kart
moped
bumper car
golfcart
leopard
leopard
jaguar
cheetah
snow leopard
Egyptian cat
grille
convertible
grille
pickup
beach wagon
fire engine
mushroom
agaric
mushroom
jelly fungus
gill fungus
dead-man's-fingers
cherry
dalmatian
grape
elderberry
ffordshire bullterrier
currant
Madagascar cat
squirrel monkey
spider monkey
titi
indri
howler monkey

ImageNet Classification Error

Krizhevsky, Sutskever & Hinton (2012)

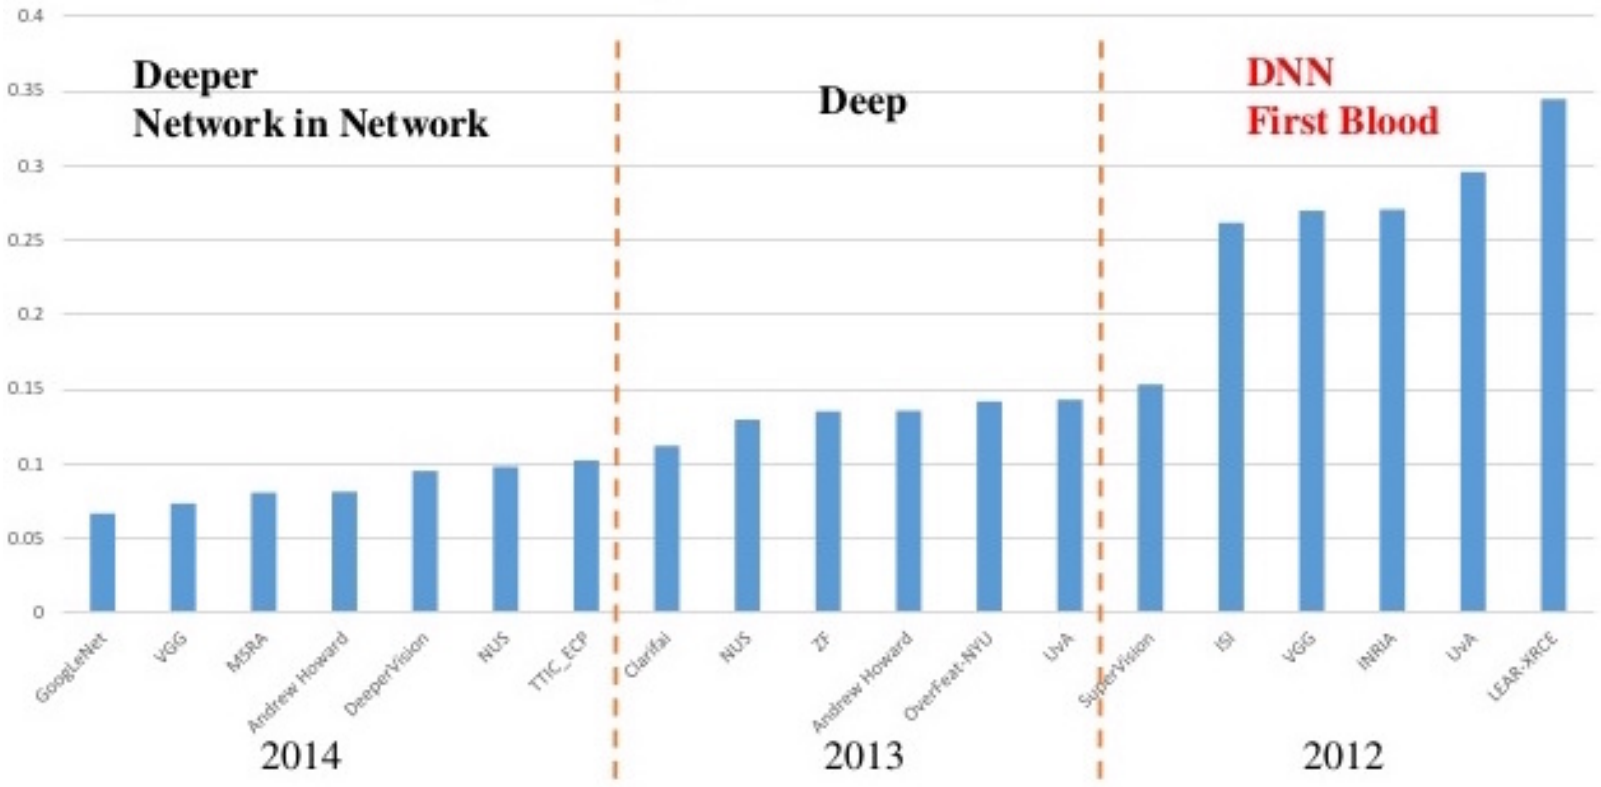

Krizhevsky et al. (2012)

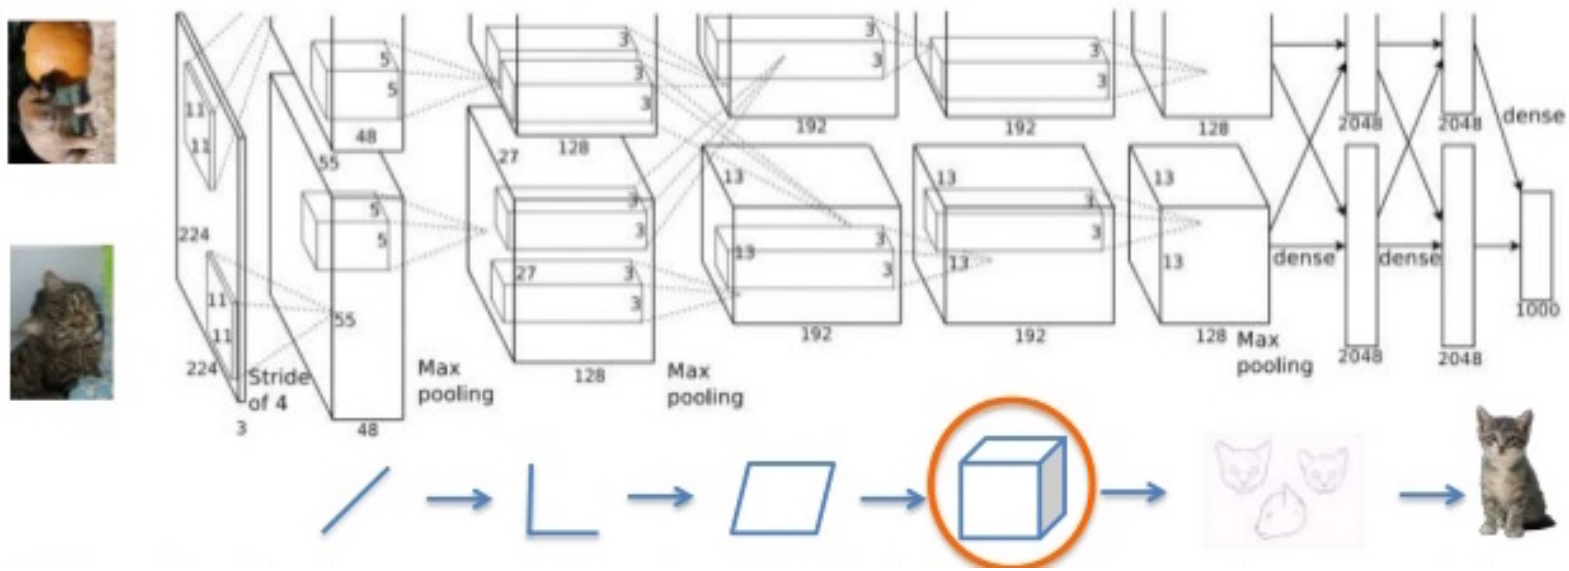

Outline


mark zuckerberg
Mark Zuckerberg
11,814 like this. · Business P...

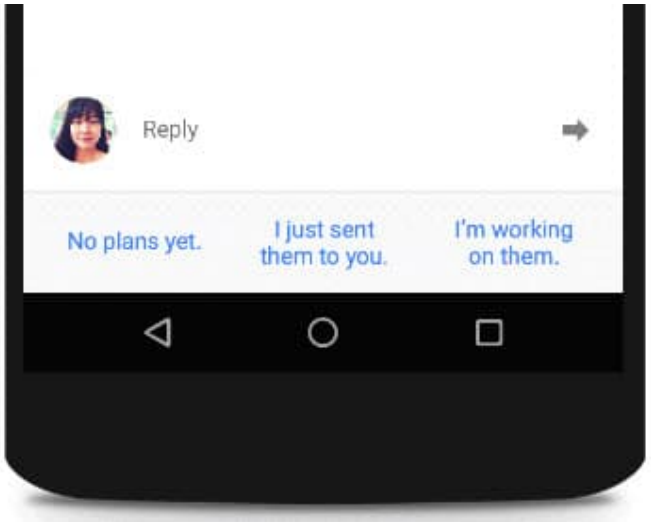
Reply
No plans yet.
I just sent them to you.
I'm working on them.


"Remind me about this when I get home"
tap to edit
OK, I'll remind you when you arrive at home:
Reminder
11 Hot Beach Escapes - Jetsetter
Arriving: Home
Remove


mark zuckerberg
Mark Zuckerberg
11,814 like this. · Business P...

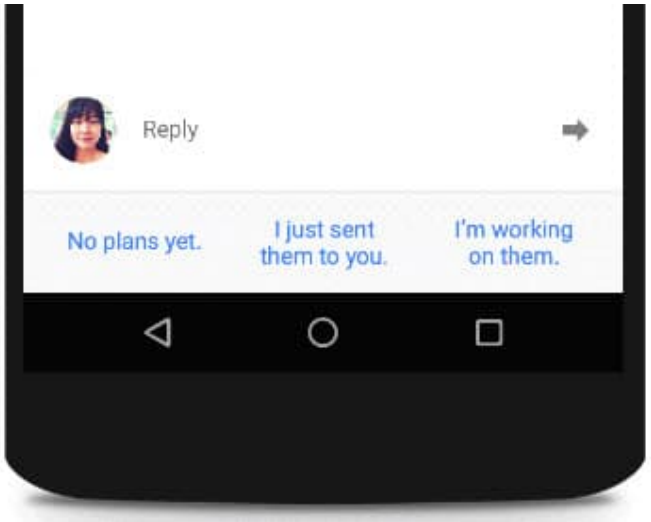
Reply
No plans yet.
I just sent them to you.
I'm working on them.


"Remind me about this when I get home"
tap to edit
OK, I'll remind you when you arrive at home:
Reminder
11 Hot Beach Escapes - Jetsetter
Arriving: Home
Remove


mark zuckerberg
Mark Zuckerberg
11,814 like this. · Business P...

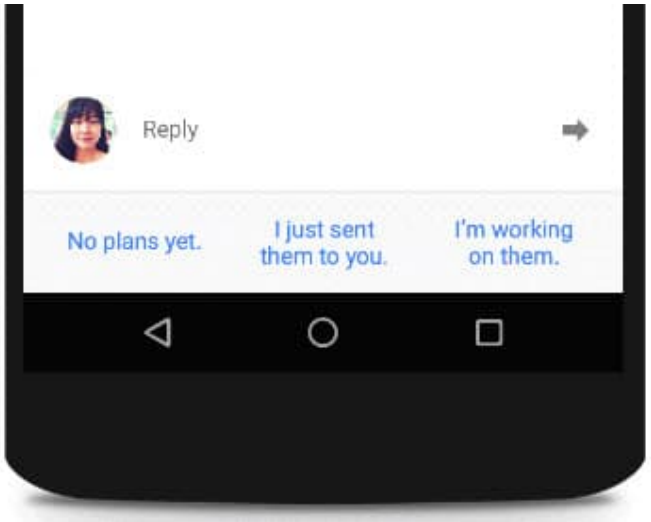
Reply
No plans yet.
I just sent them to you.
I'm working on them.


"Remind me about this when I get home"
tap to edit
OK, I'll remind you when you arrive at home:
Reminder
11 Hot Beach Escapes - Jetsetter
Arriving: Home
Remove

Deep Learning in Late 2016

[Image to Lyrics and Music]

[Daddy's Car]

[Sunspring]

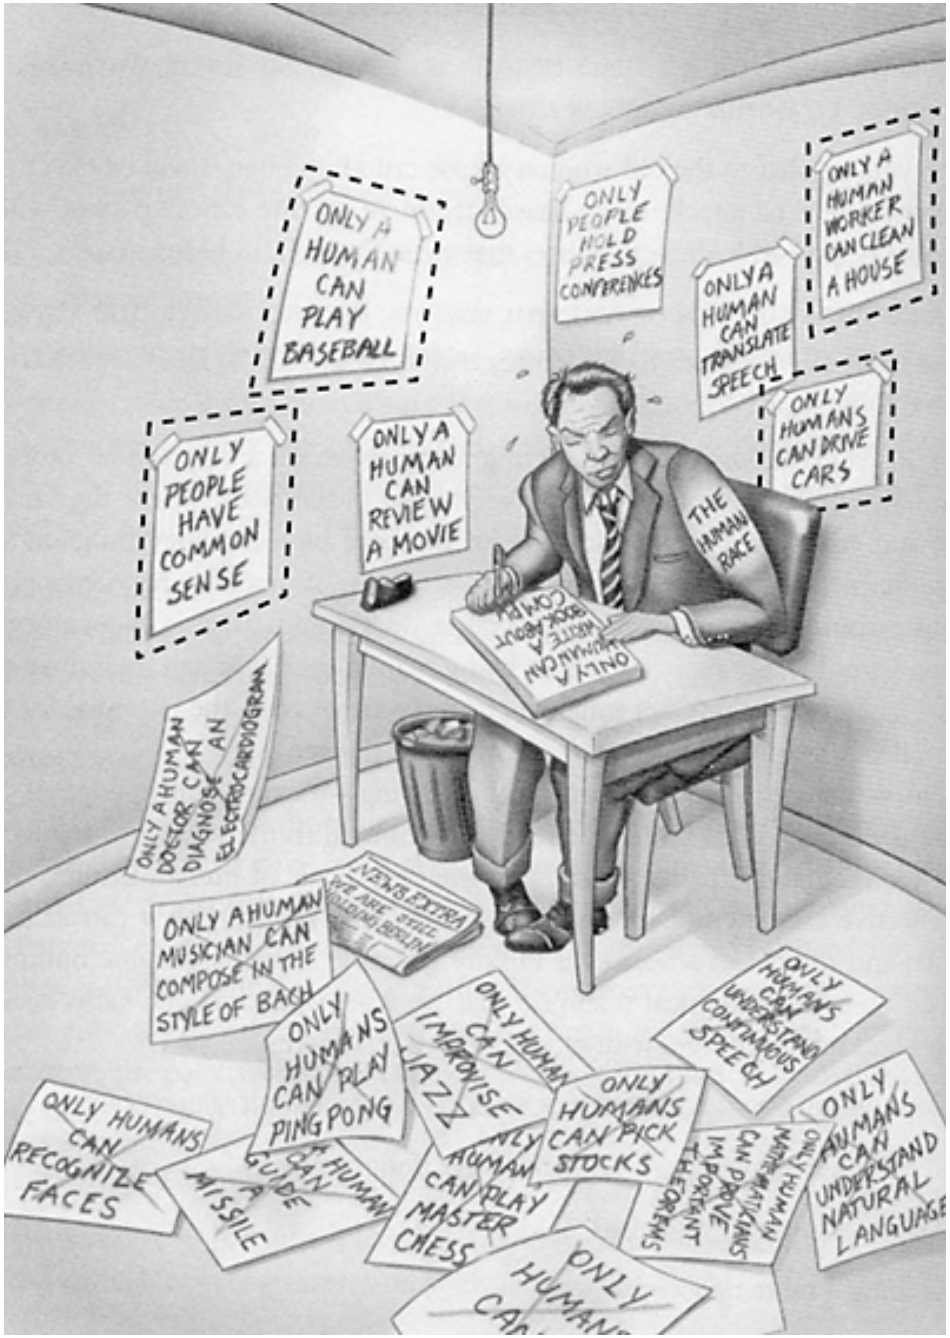
ONLY A HUMAN CAN PLAY BASEBALL
ONLY PEOPLE HOLD PRESS CONFERENCES
ONLY A HUMAN WORKER CAN CLEAN A HOUSE
ONLY A HUMAN CAN TRANSLATE SPEECH
ONLY PEOPLE HAVE COMMON SENSE
ONLY A HUMAN CAN REVIEW A MOVIE
ONLY HUMANS CAN DRIVE CARS
THE HUMAN RACE
ONLY A HUMAN DOCTOR CAN DIAGNOSE AN ELECTROCARDIOGRAM
ONLY A HUMAN MUSICIAN CAN COMPOSE IN THE STYLE OF BACH
NEWS EXTRA
ONLY HUMANS CAN PLAY PING PONG
ONLY HUMANS CAN IMPROVISE JAZZ
ONLY HUMANS CAN UNDERSTAND CONTINUOUS SPEECH
ONLY HUMANS CAN RECOGNIZE FACES
A HUMAN CAN GUIDE A MISSILE
ONLY HUMANS CAN PICK STOCKS
A HUMAN CAN PLAY MASTER CHESS
ONLY HUMANS CAN UNDERSTAND NATURAL LANGUAGE

Outline

Biological Neuron Morphology

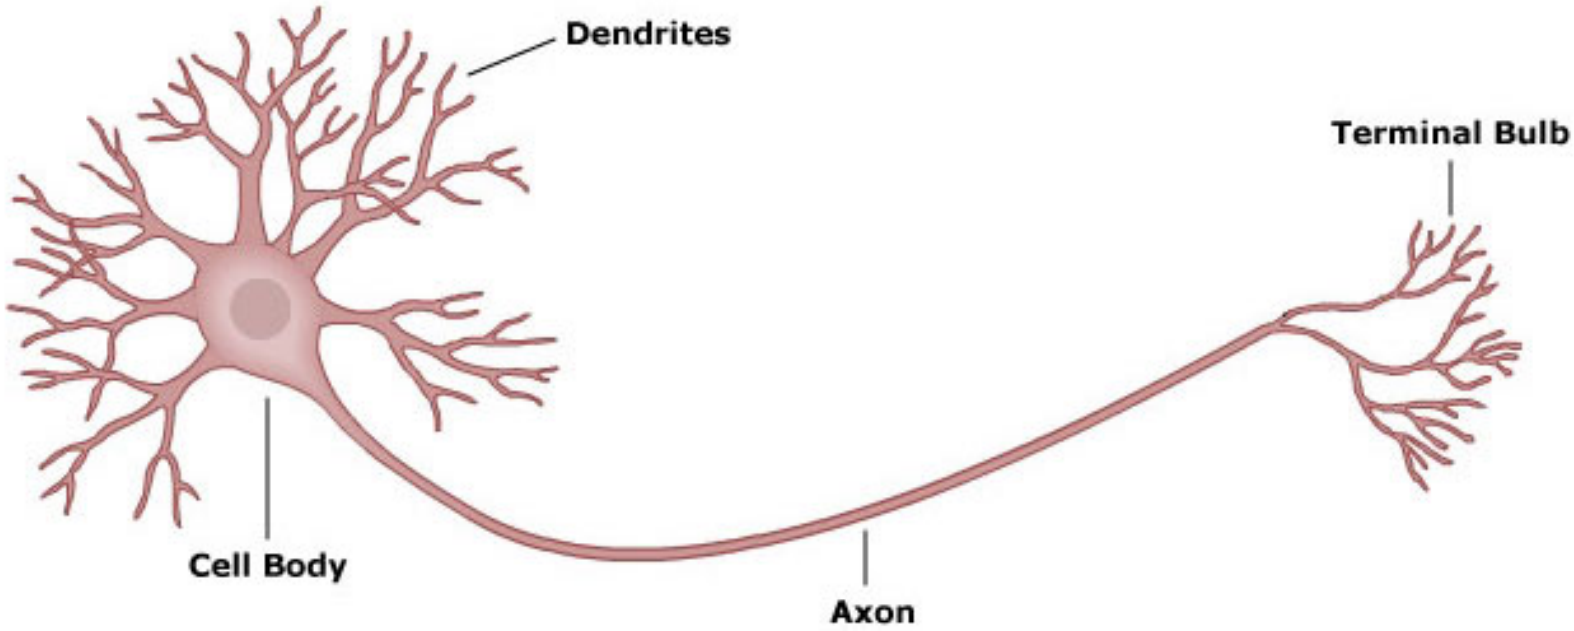

Perceptron

Rosenblatt (1957)

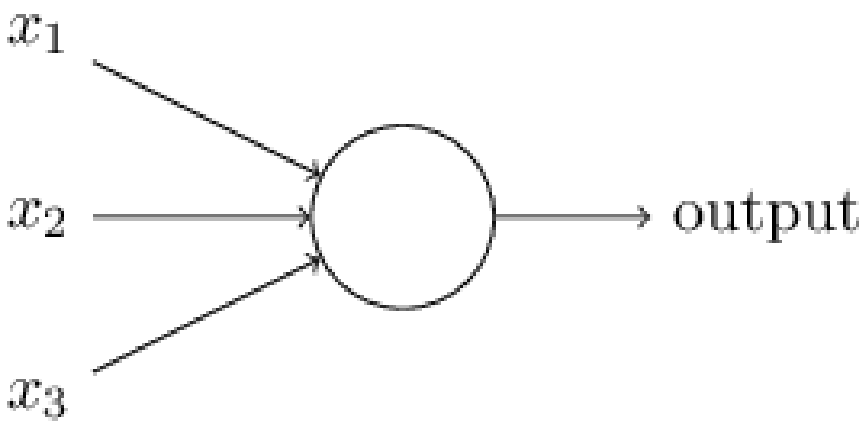

$$\text{output} = \begin{cases} 0 & \text{if } \sum_j w_j x_j \leq \text{threshold} \\ 1 & \text{if } \sum_j w_j x_j > \text{threshold} \end{cases}$$

Biological Neuron Physiology

The *Binary* Action Potential

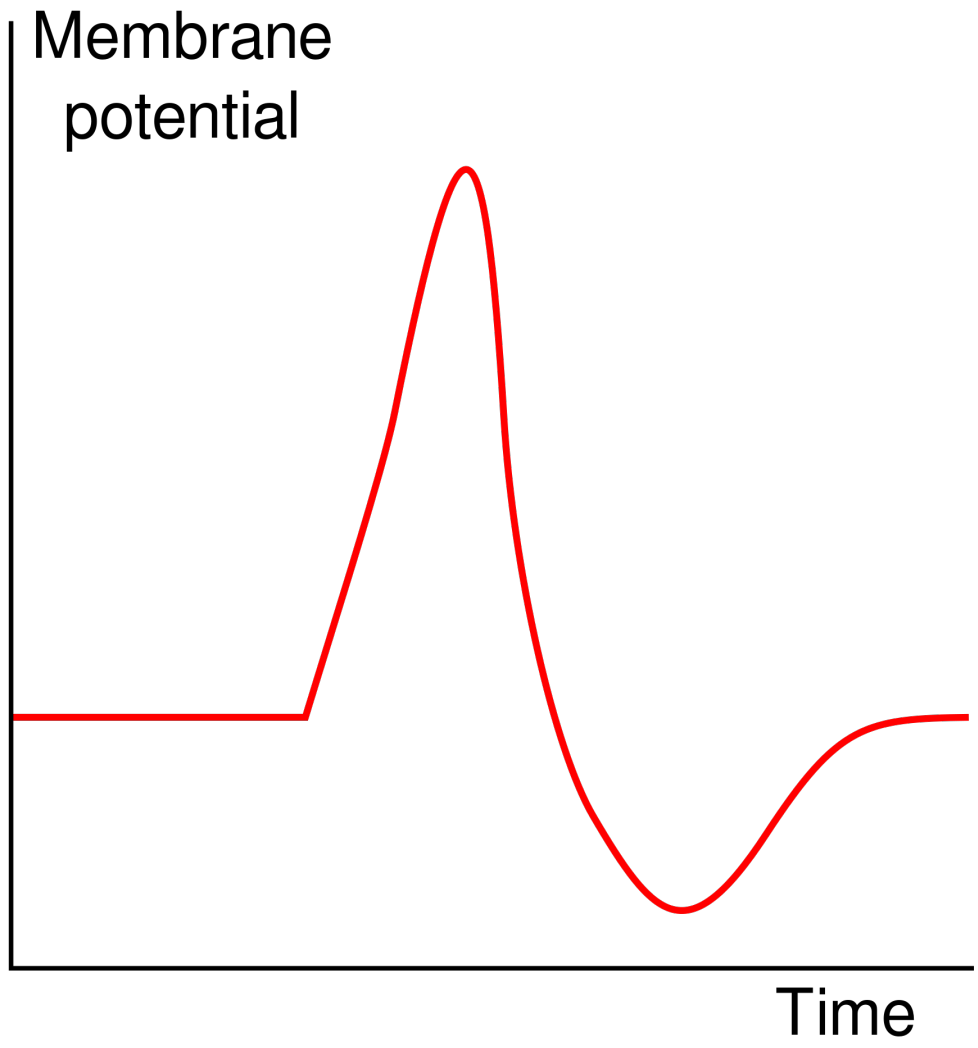

Perceptron

Rosenblatt (1957)

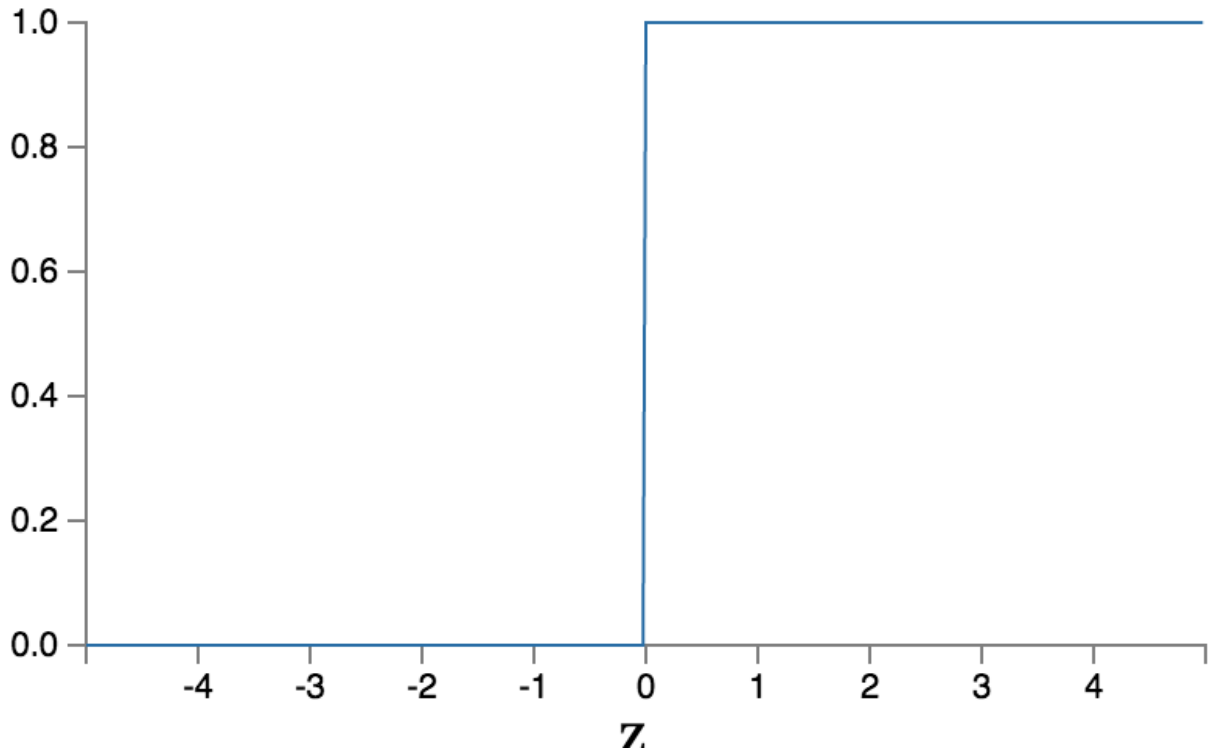

Multi-Layer Perceptron

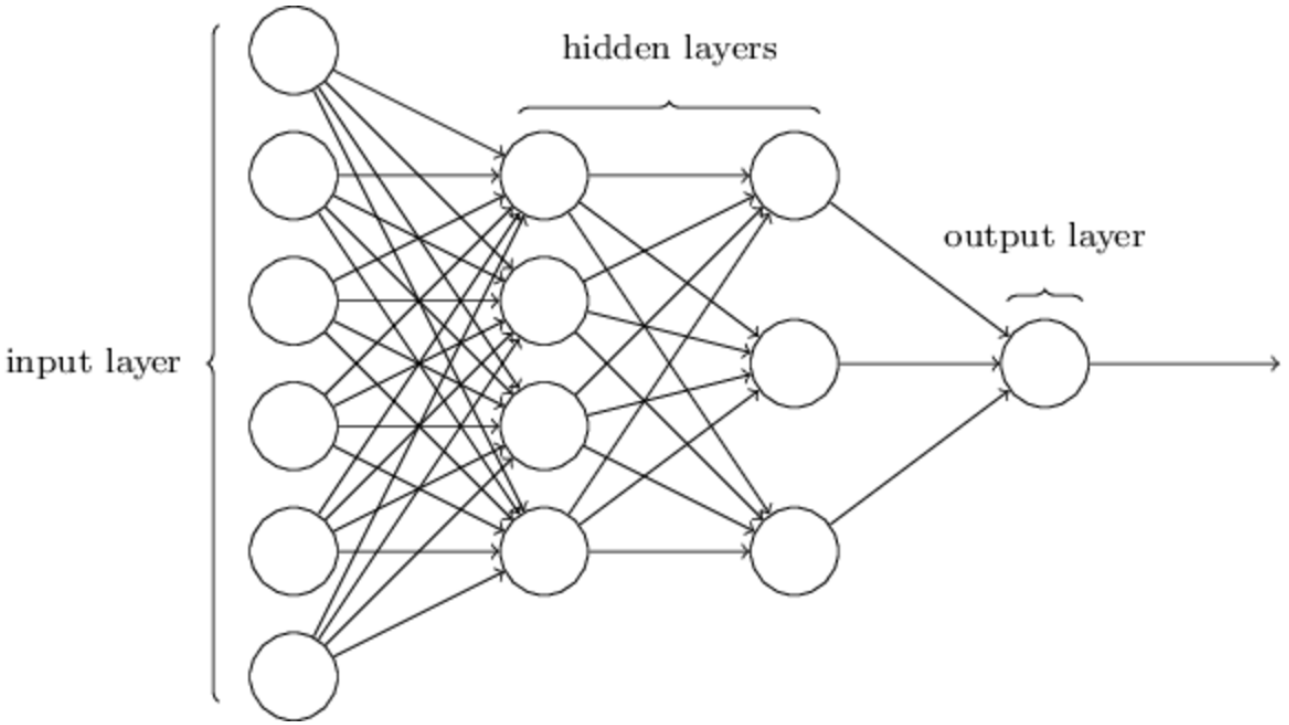

Multi-Layer Perceptron

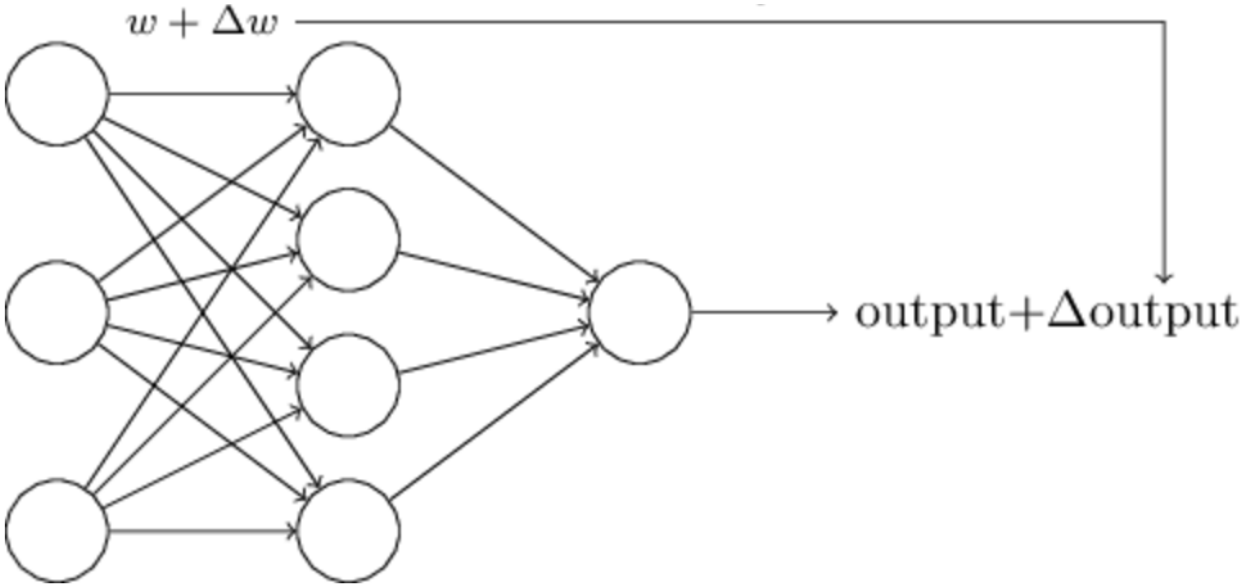

Sigmoid Neuron

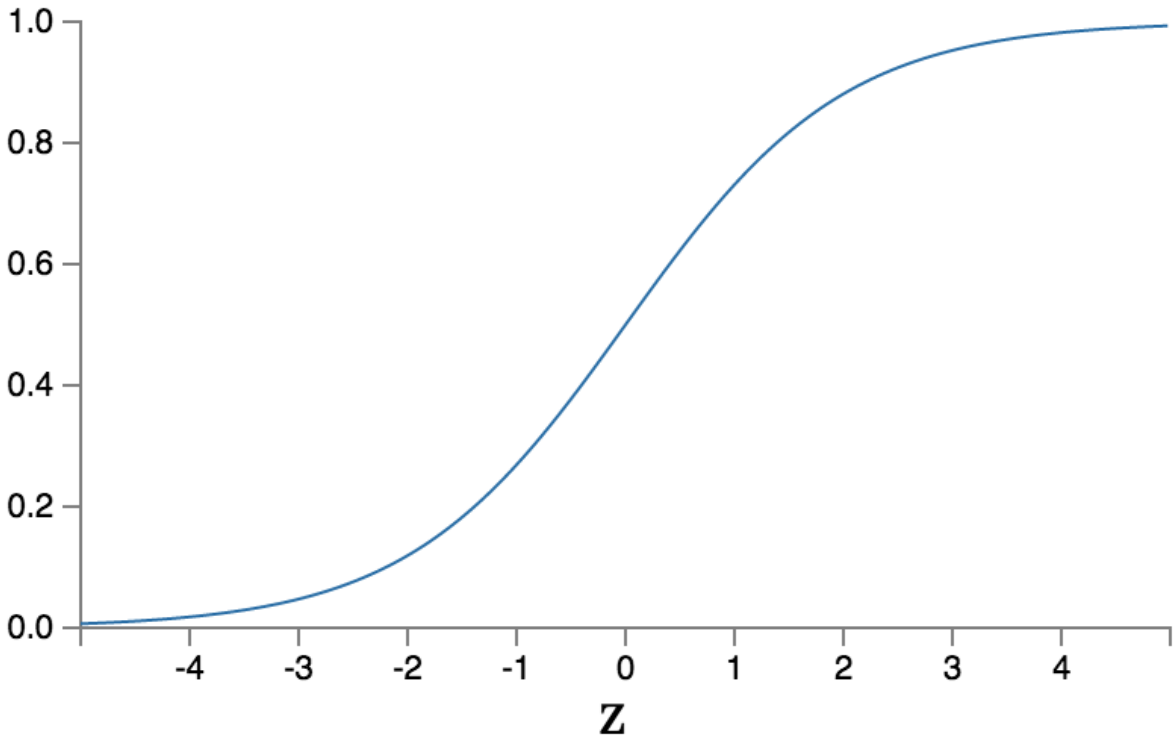

$$\frac{1}{1 + \exp(-\sum_j w_j x_j - b)}$$

tanh Neuron

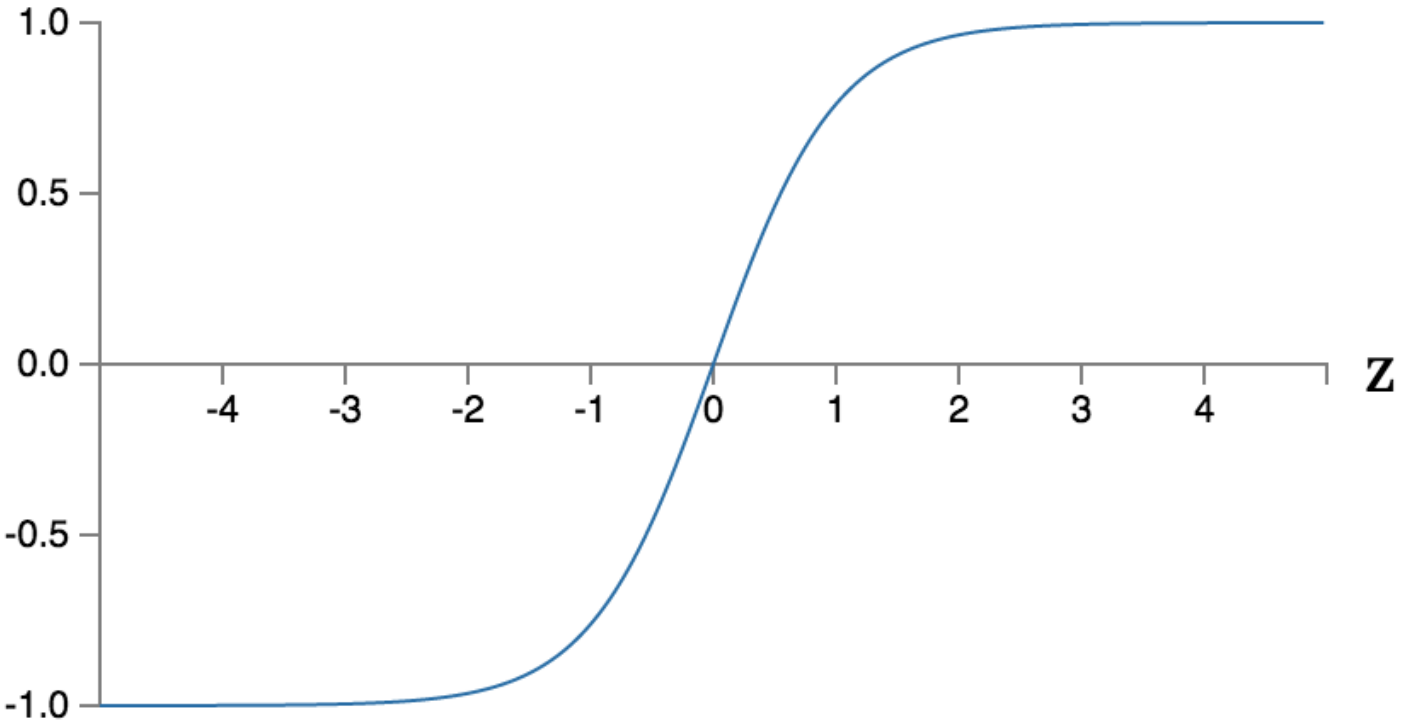

$$\sigma(z) = \frac{1 + \tanh(z/2)}{2}$$

ReLU: Rectified Linear Units

Nair & Hinton (2010); Maas, Hannun & Ng (2014)

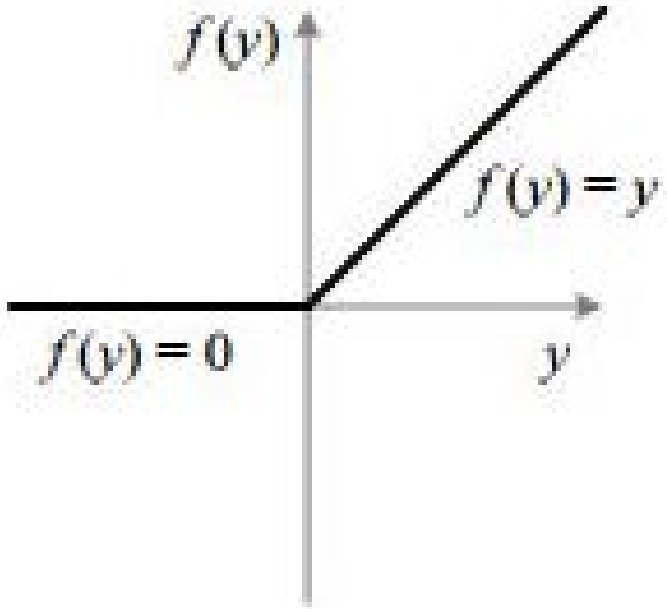

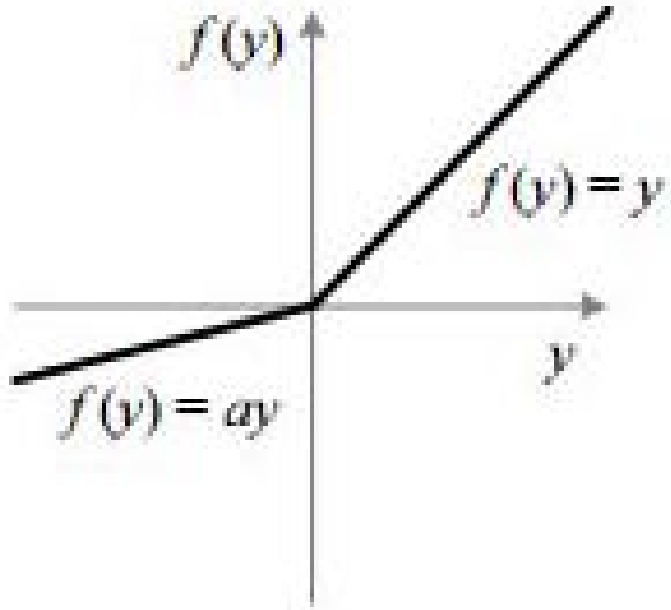

Outline

MNIST

Handwritten Digits

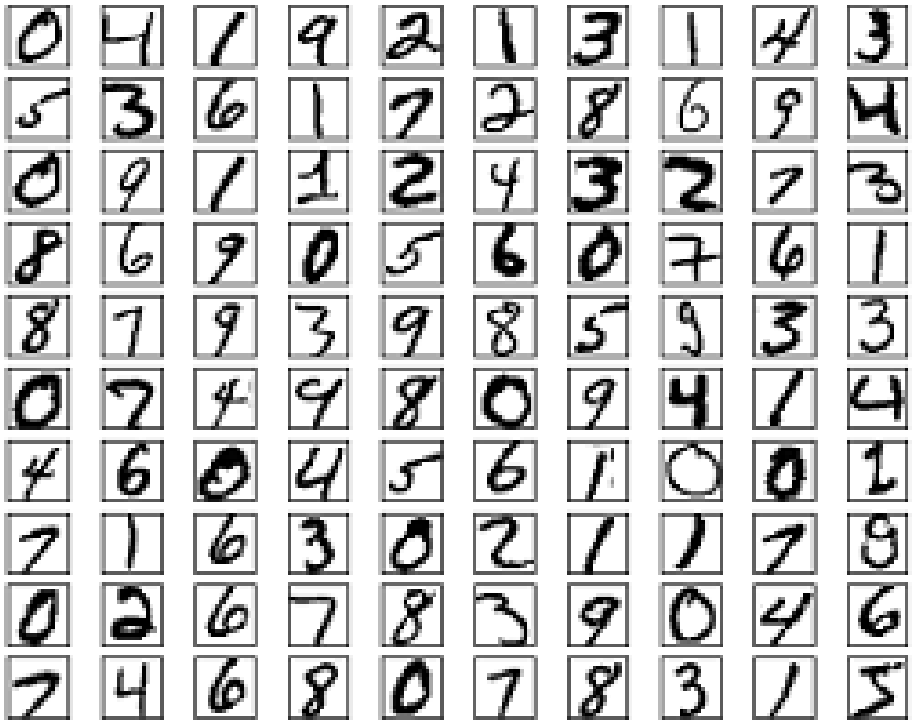

Fully-Connected Neural Net

Single Hidden Layer

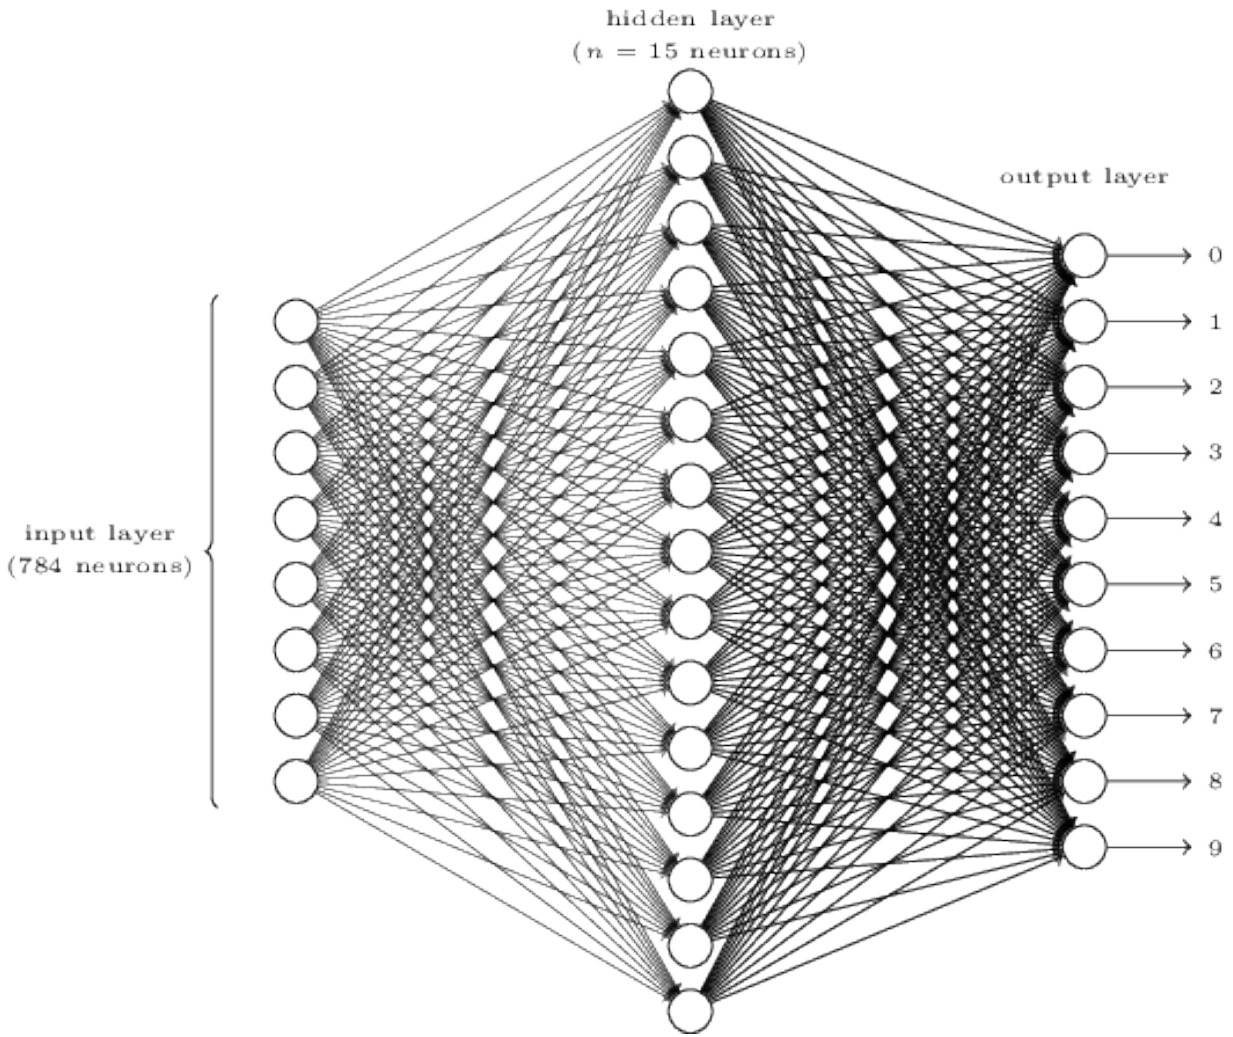

Popular Libraries

Never pay for software

- Theano
- Torch
- Caffe
- TensorFlow [demo]

Higher-Level APIs:

- TFLearn
- Keras

Popular Libraries

Never pay for software

- Theano
- Torch
- Caffe
- TensorFlow [demo]

Higher-Level APIs:

- TFLearn
- Keras

Popular Libraries

Never pay for software

- Theano
- Torch
- Caffe
- TensorFlow [demo]

Higher-Level APIs:

- TFLearn
- Keras

Popular Libraries

Never pay for software

- Theano
- Torch
- Caffe
- TensorFlow [demo]

Higher-Level APIs:

- TFLearn
- Keras

Popular Libraries

Never pay for software

- Theano
- Torch
- Caffe
- TensorFlow [demo]

Higher-Level APIs:

- TFLearn
- Keras

Popular Libraries

Never pay for software

- Theano
- Torch
- Caffe
- TensorFlow [demo]

Higher-Level APIs:

- TFLearn
- Keras

Outline

Deep Fully-Connected Net

3 (or more) Hidden Layers

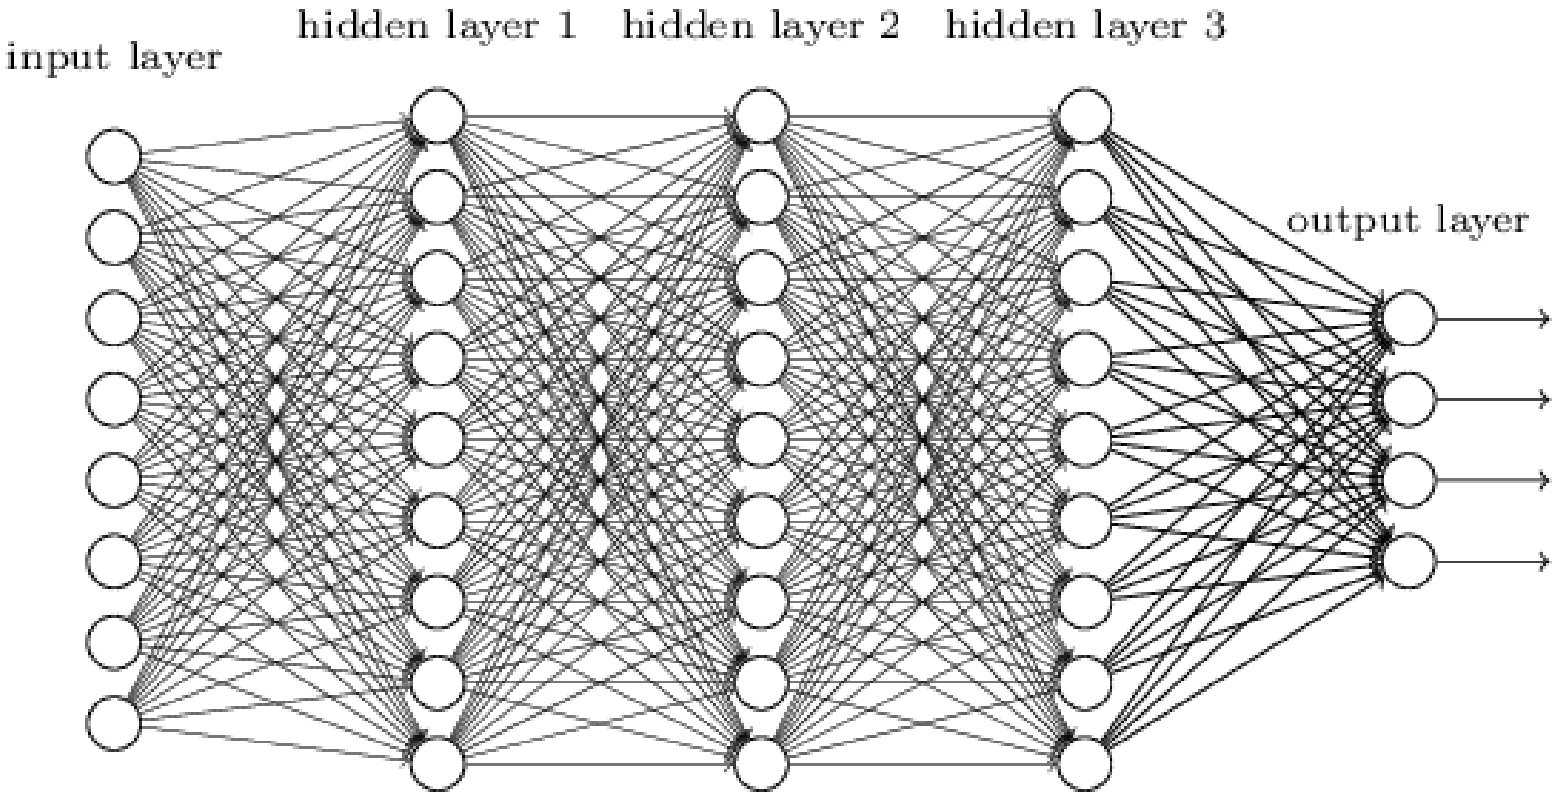

Synaptic Pruning

(Stochastic) Gradient Descent

Adam = AdaGrad + RMSprop

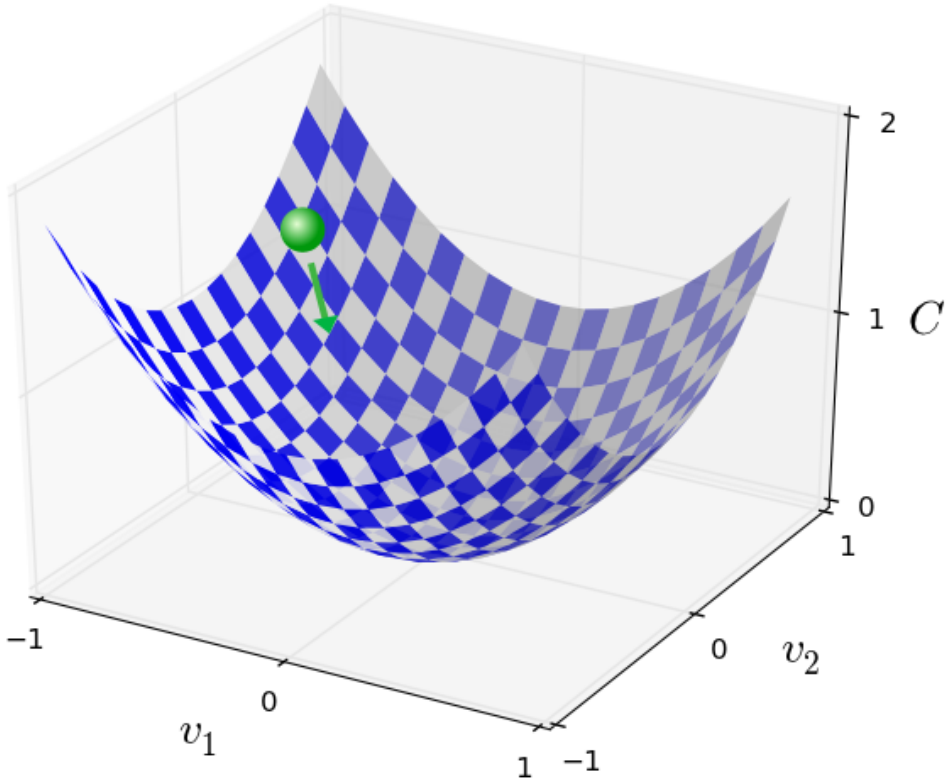

Backpropagation

computes error & gradient of cost function

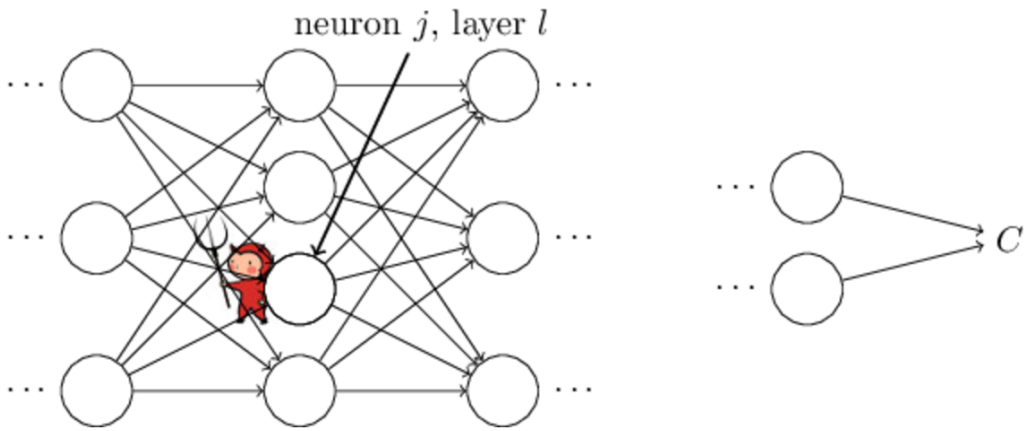

$$\delta^L = \nabla_a C \odot \sigma'(z^L) \qquad \text{(BP1)}$$

$$\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l) \qquad \text{(BP2)}$$

$$\frac{\partial C}{\partial b^l_j} = \delta^l_j \qquad \text{(BP3)}$$

$$\frac{\partial C}{\partial w^l_{j\,k}} = a^{l-1}_k \delta^l_j \qquad \text{(BP4)}$$

Overfitting

...and avoiding it

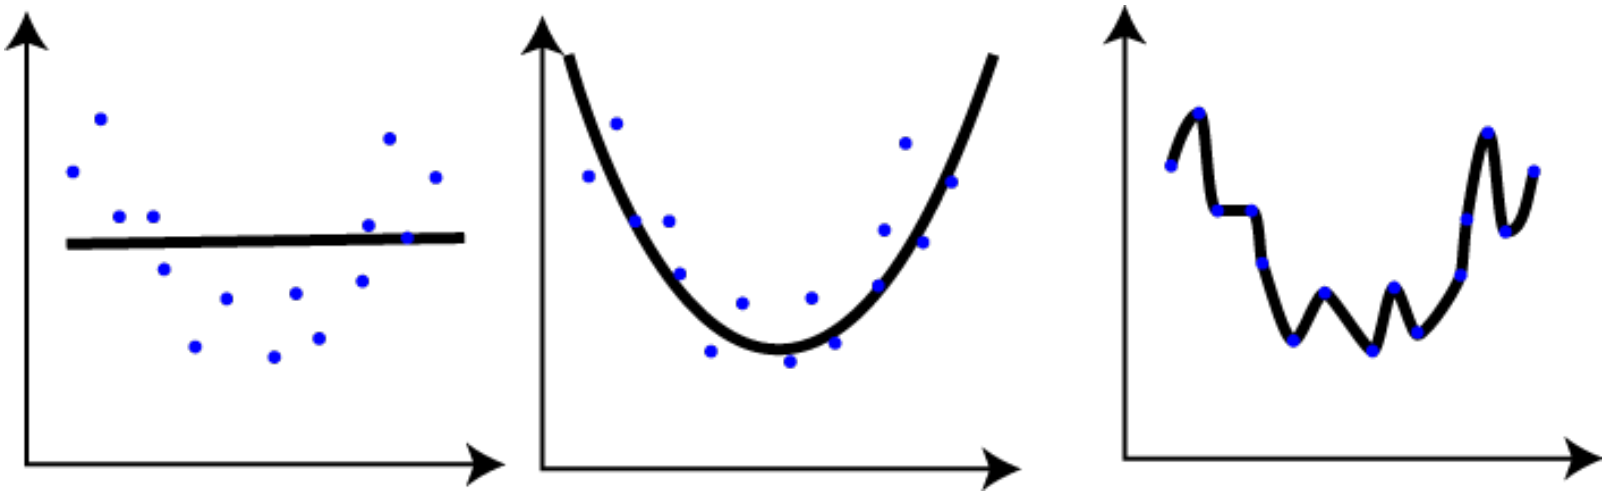

- L1/L2 regularization
- dropout
- artificial data set expansion

Overfitting

...and avoiding it

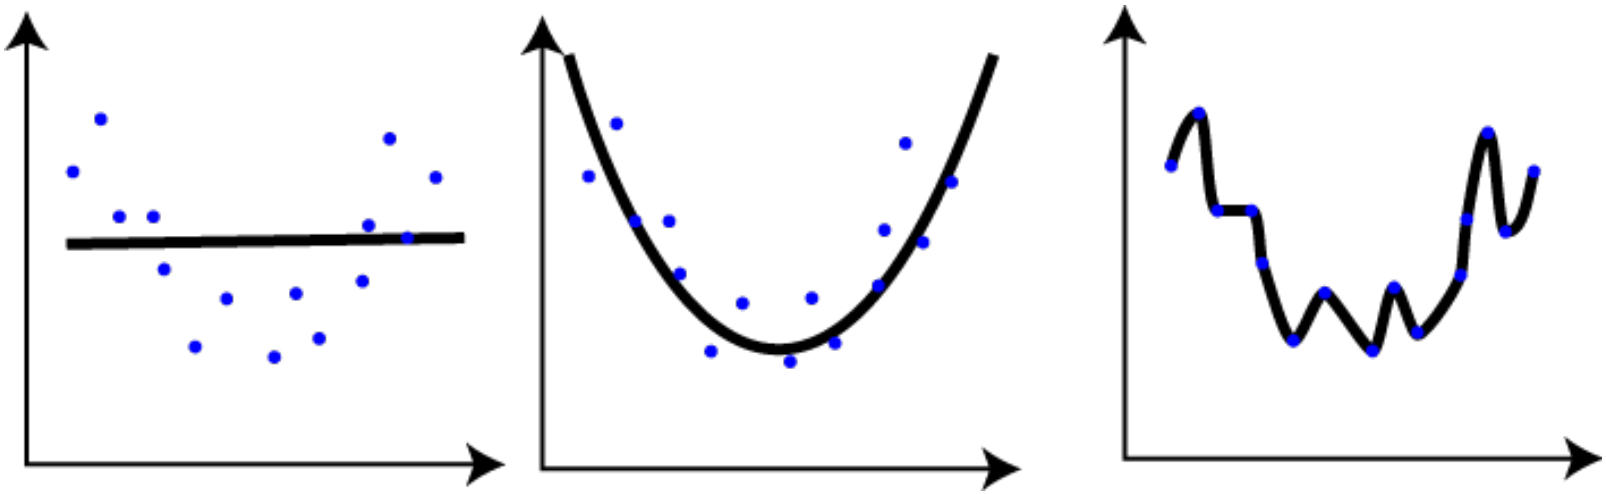

- L1/L2 regularization
- dropout
- artificial data set expansion

Overfitting

...and avoiding it

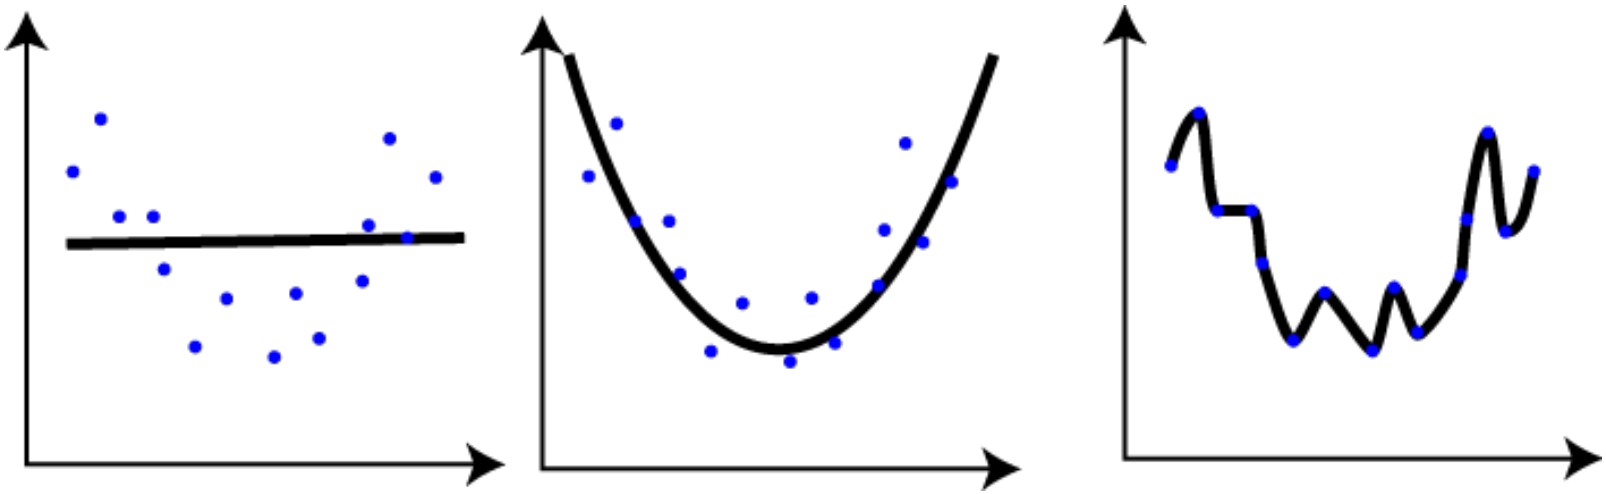

- L1/L2 regularization
- dropout
- artificial data set expansion

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Improving Neural Networks

Attribute & Hyperparameter Tuning

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η, variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

Universality

Solve Any Continuous Function (Nielsen, 2015)

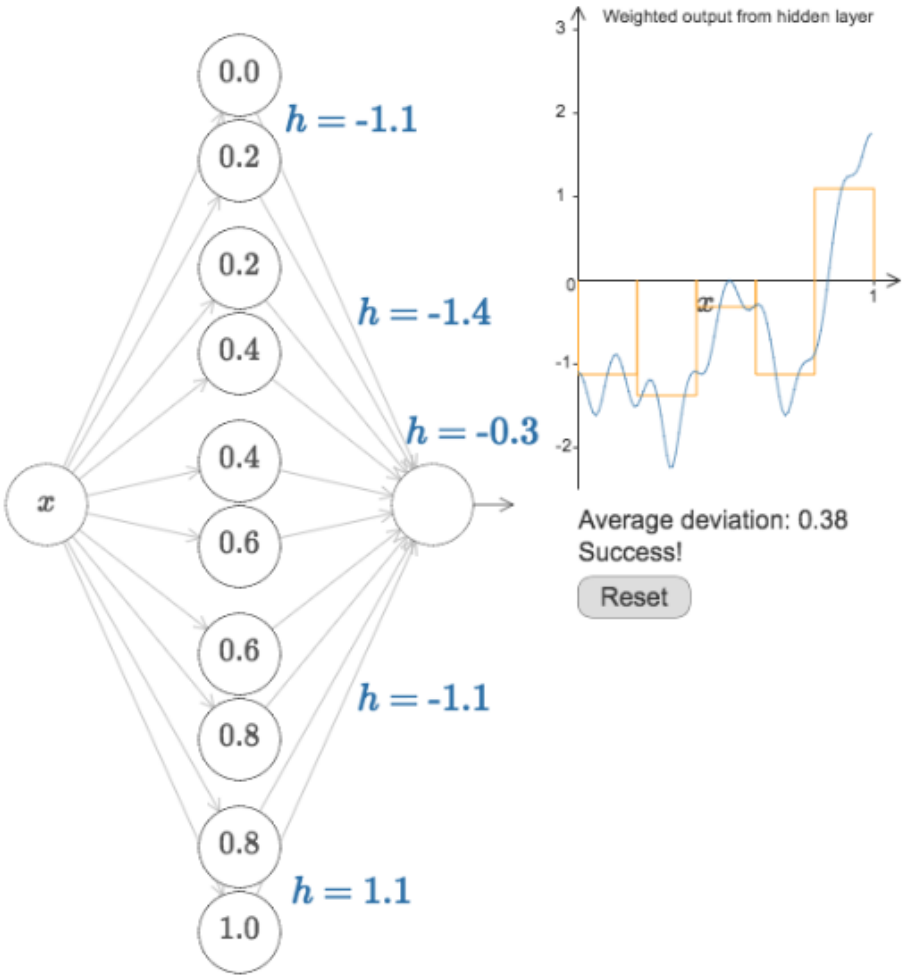

Unstable Gradient

Typically *Vanishes* (but can *Explode*)

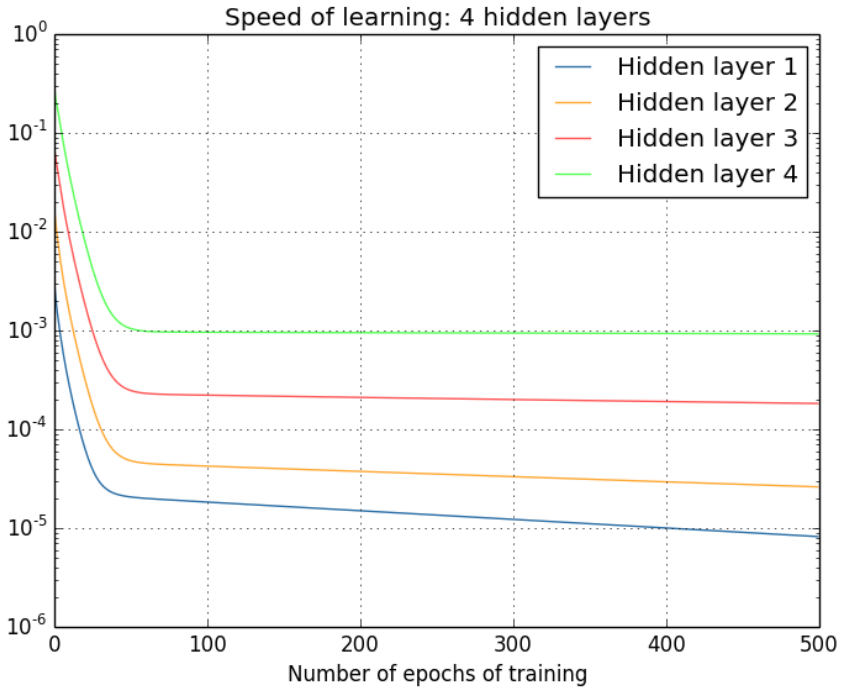

$$\frac{\partial C}{\partial b_1} = \sigma'(z_1) \times w_2 \times \sigma'(z_2) \times w_3 \times \sigma'(z_3) \times w_4 \times \sigma'(z_4) \times \frac{\partial C}{\partial a_4}$$

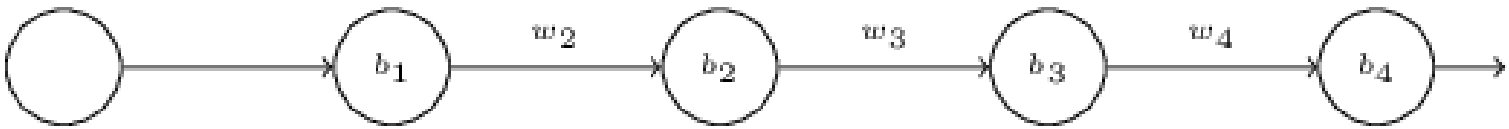

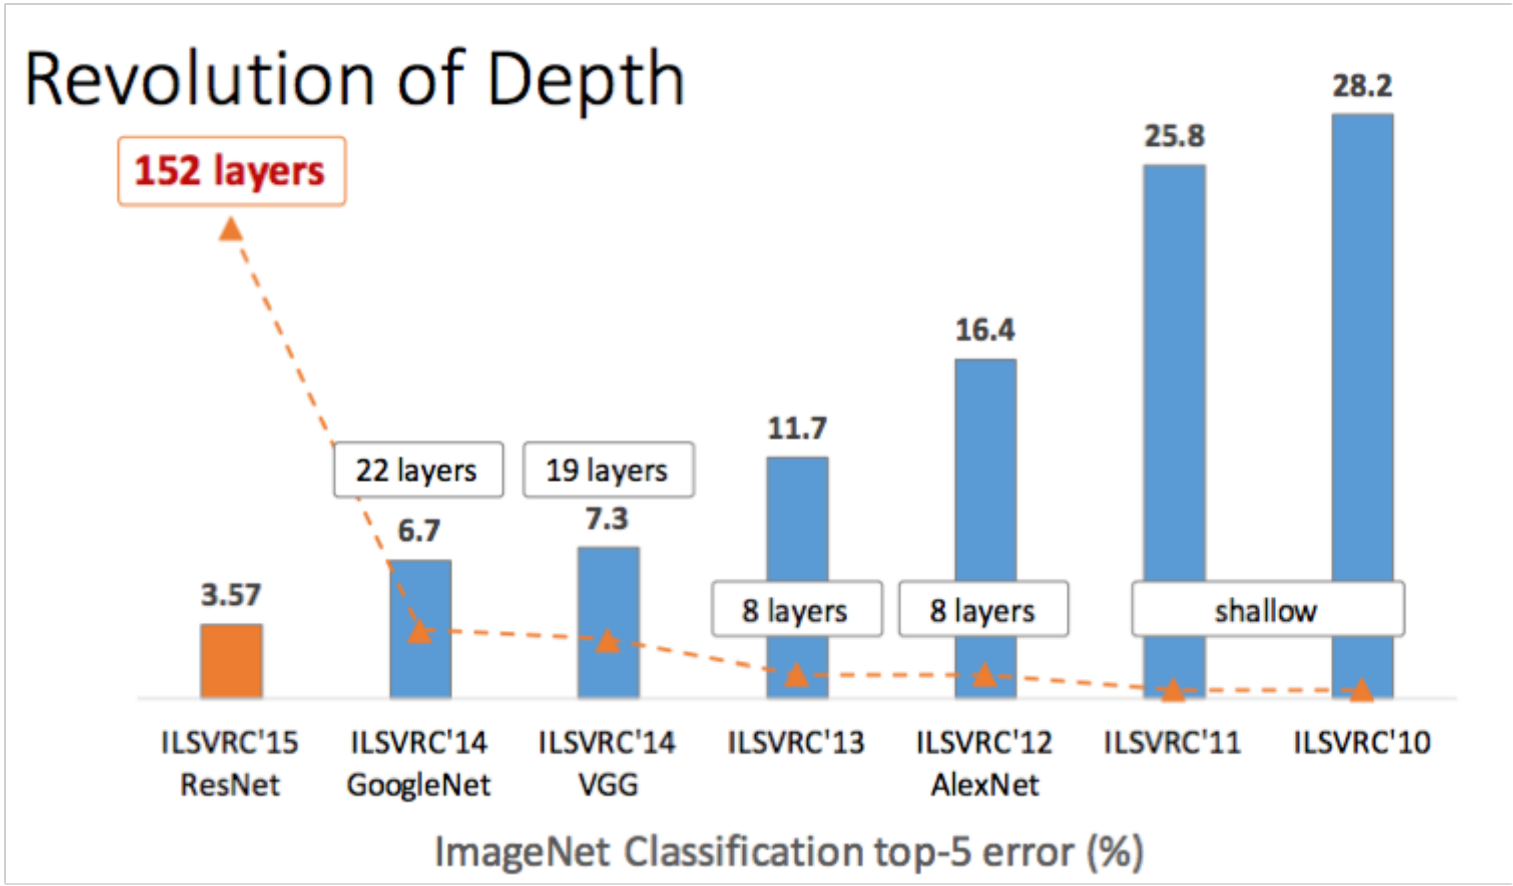

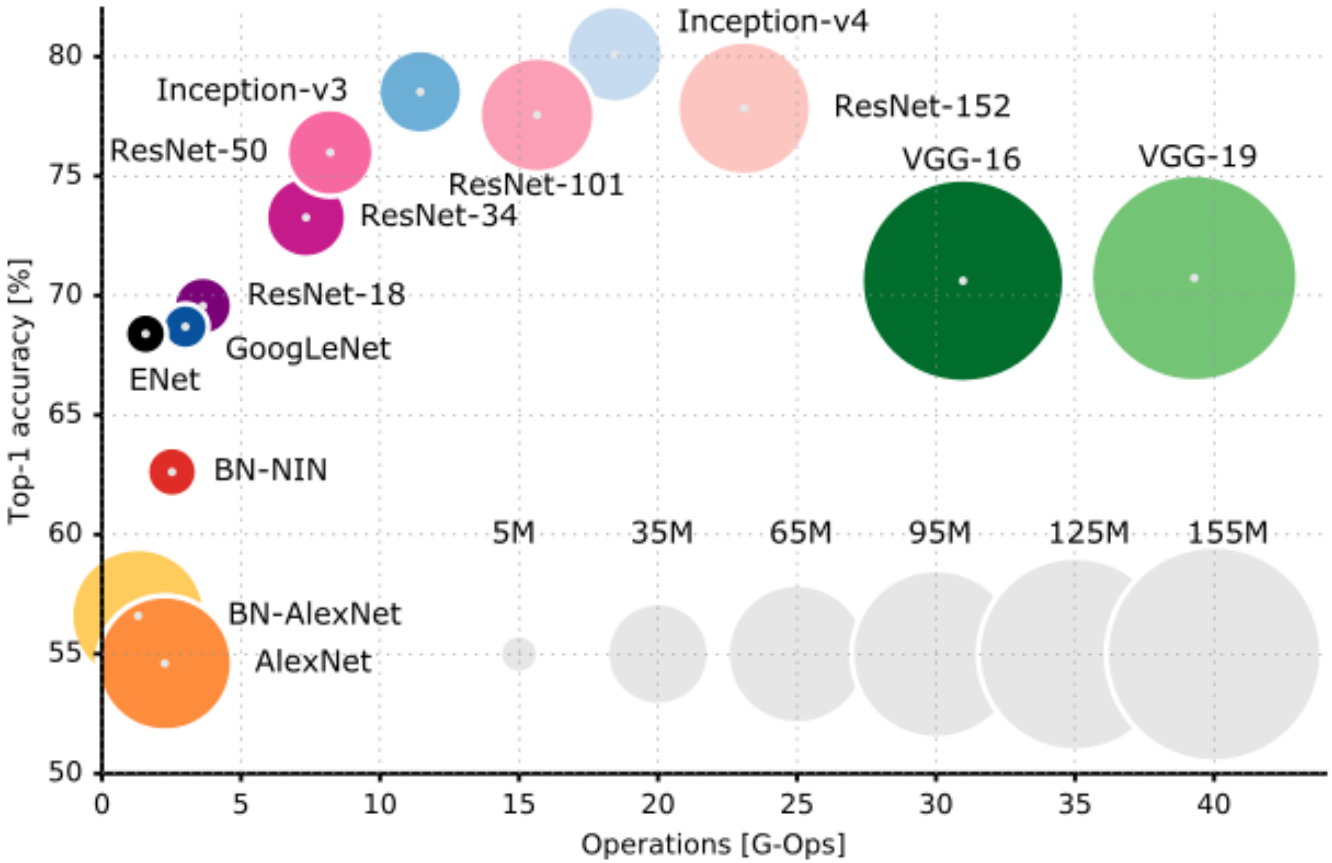
Inception-v4
Inception-v3
ResNet-152
ResNet-50
ResNet-101
VGG-16
VGG-19
ResNet-34
ResNet-18
GoogLeNet
ENet
BN-NIN
5M
35M
65M
95M
125M
155M
BN-AlexNet
AlexNet
Top-1 accuracy [%]
80
75
70
65
60
55
50
0
5
10
15
20
25
30
35
40
Operations [G-Ops]

Outline

Hubel & Wiesel (1959)

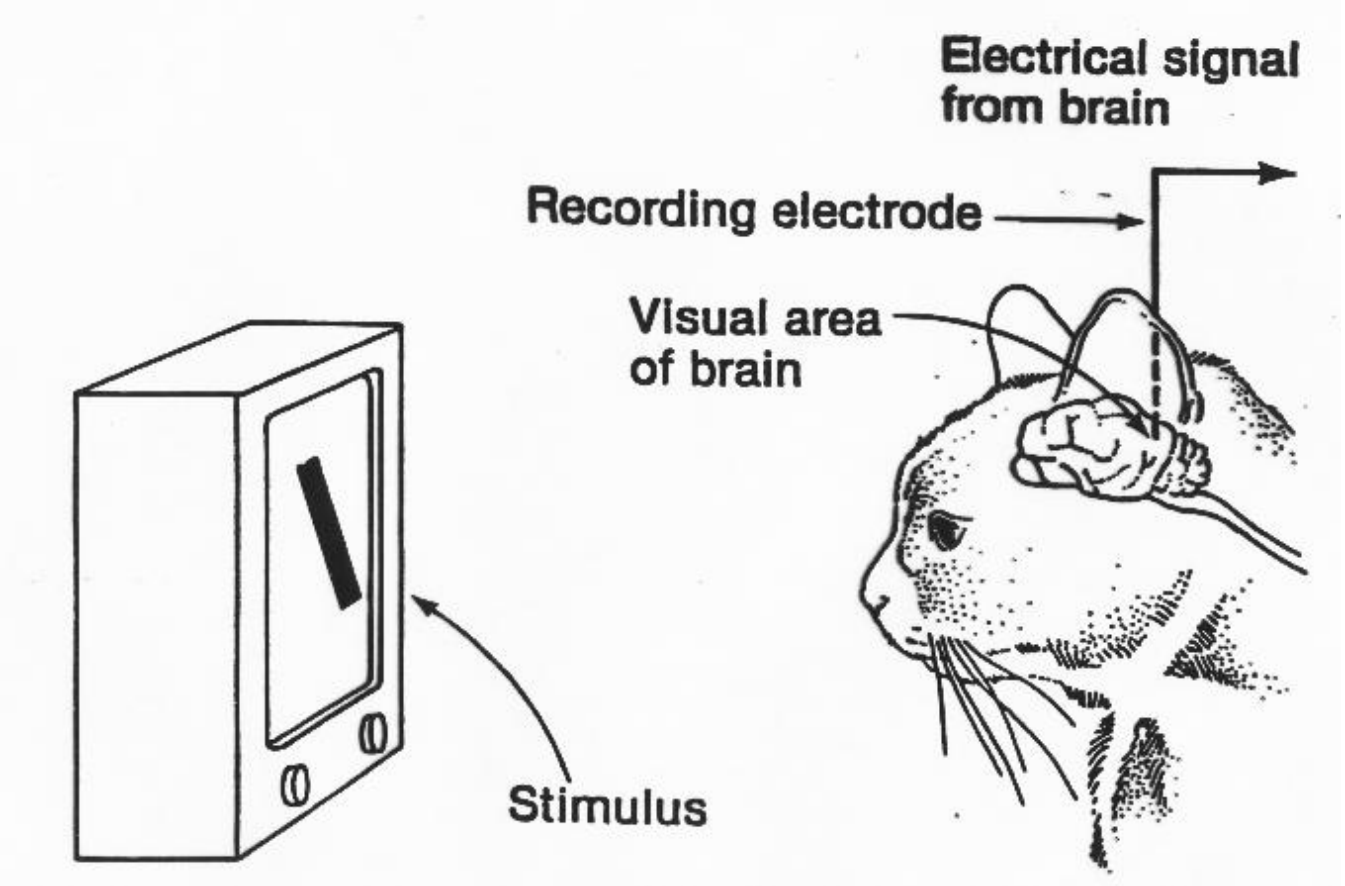

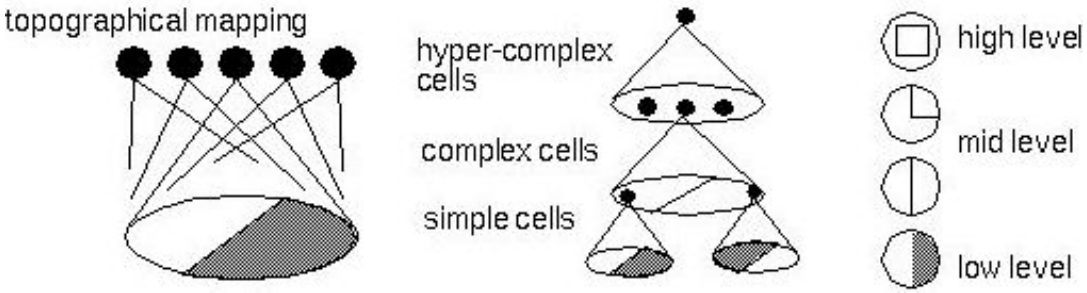

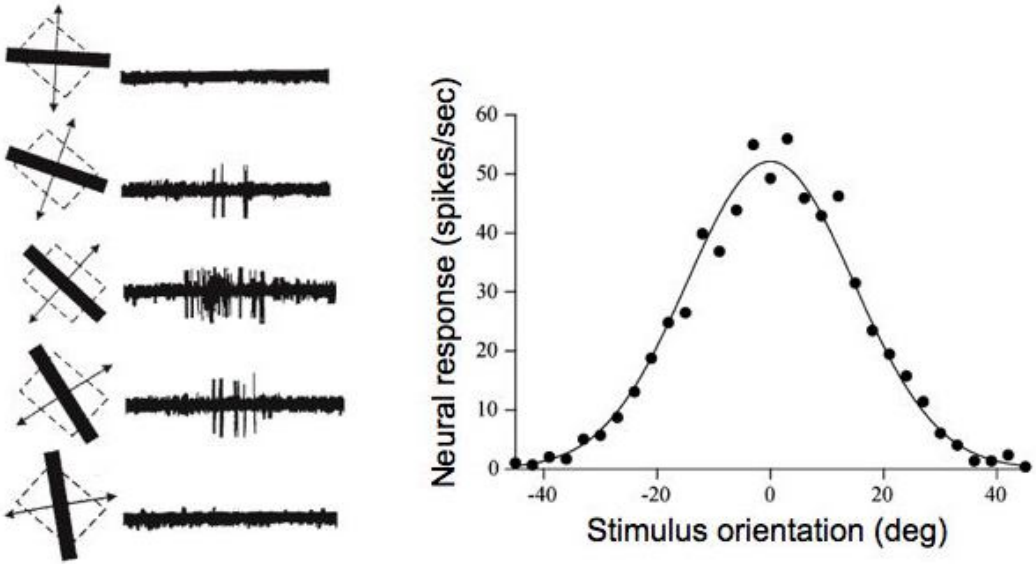

Hubel & Wiesel, 1968

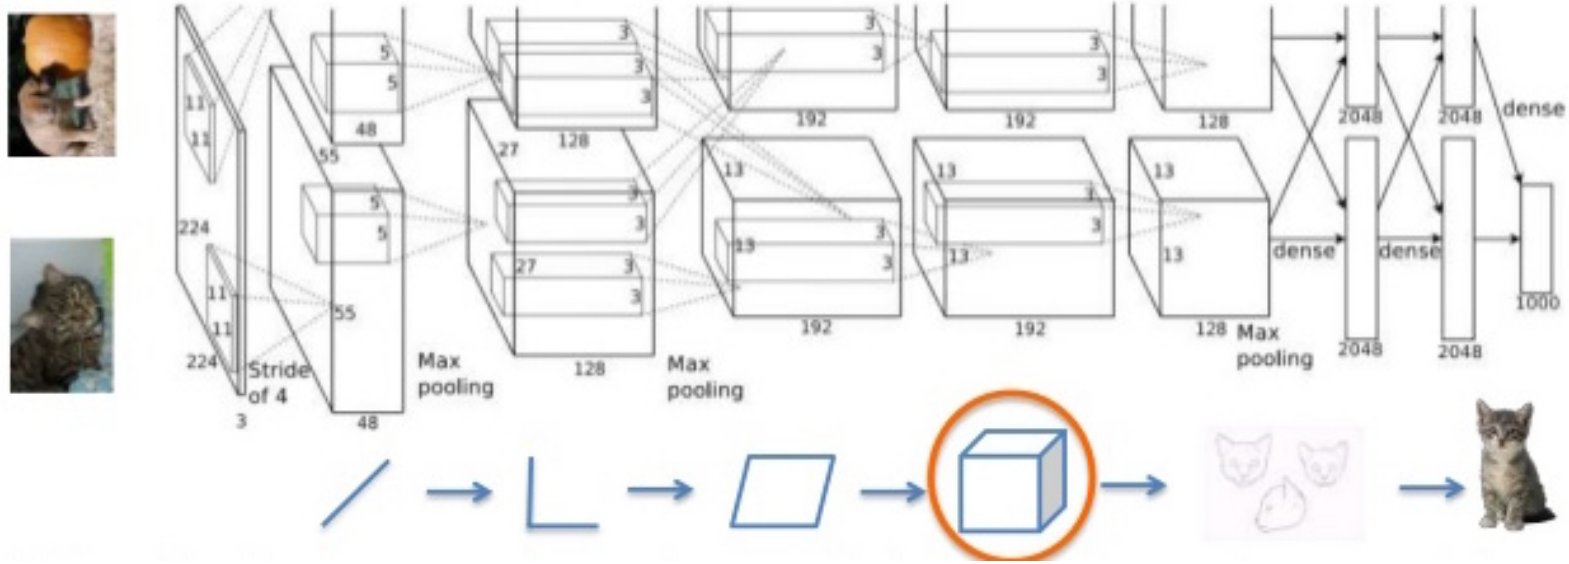

*De*ConvNet

Yosinski et al. (2015)

[Deep Visualization Toolbox]

“2.5-dimension” CT Scans

Roth et al. (2015)

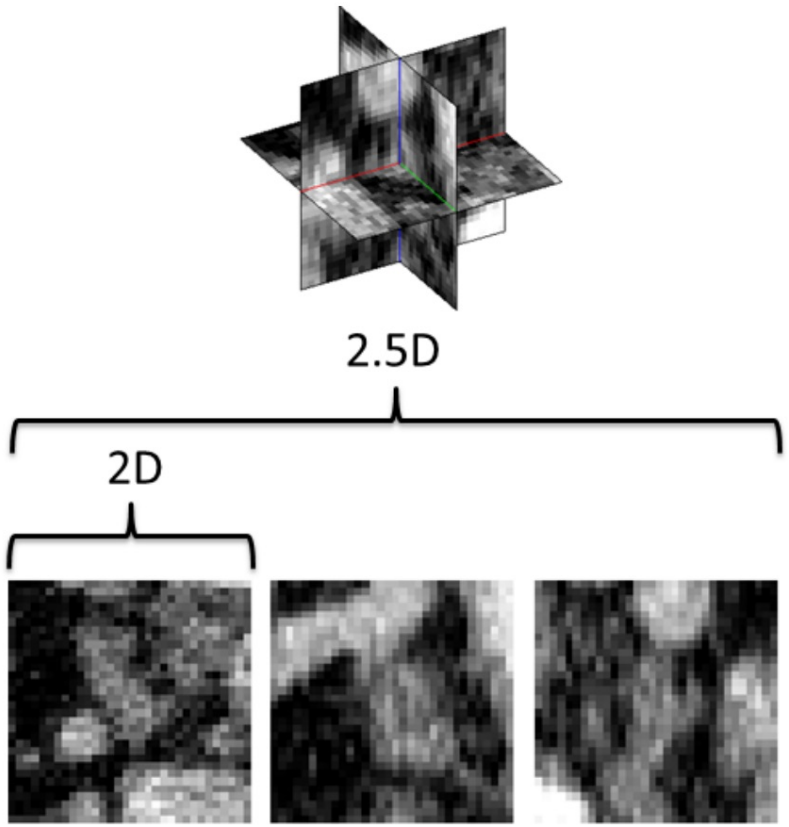

Computer-Aided Detection

Shin et al. (2016); Roth et al. (2016)

Experimental Results (~100% sensitivity but ~40 FPs/patient at candidate generation step; then 3-fold CV with data augmentation)

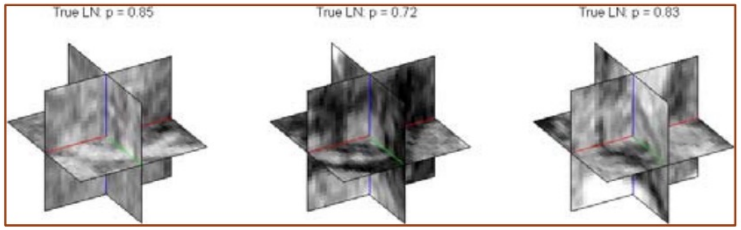

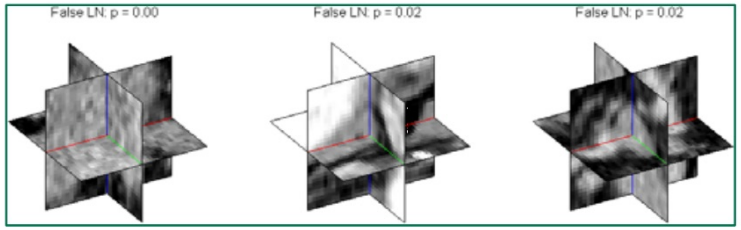

- Mediastinum

71% @ 3 FPs (was 55%)

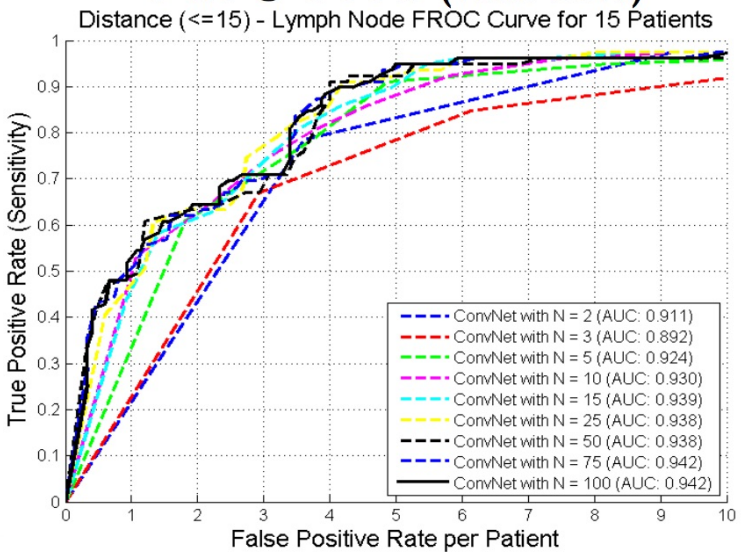

- Abdomen

83% @ 3 FPs (was 30%)

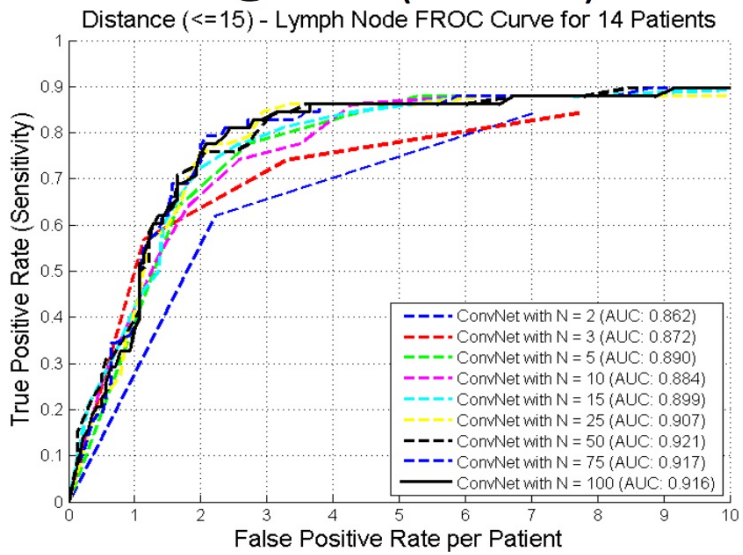

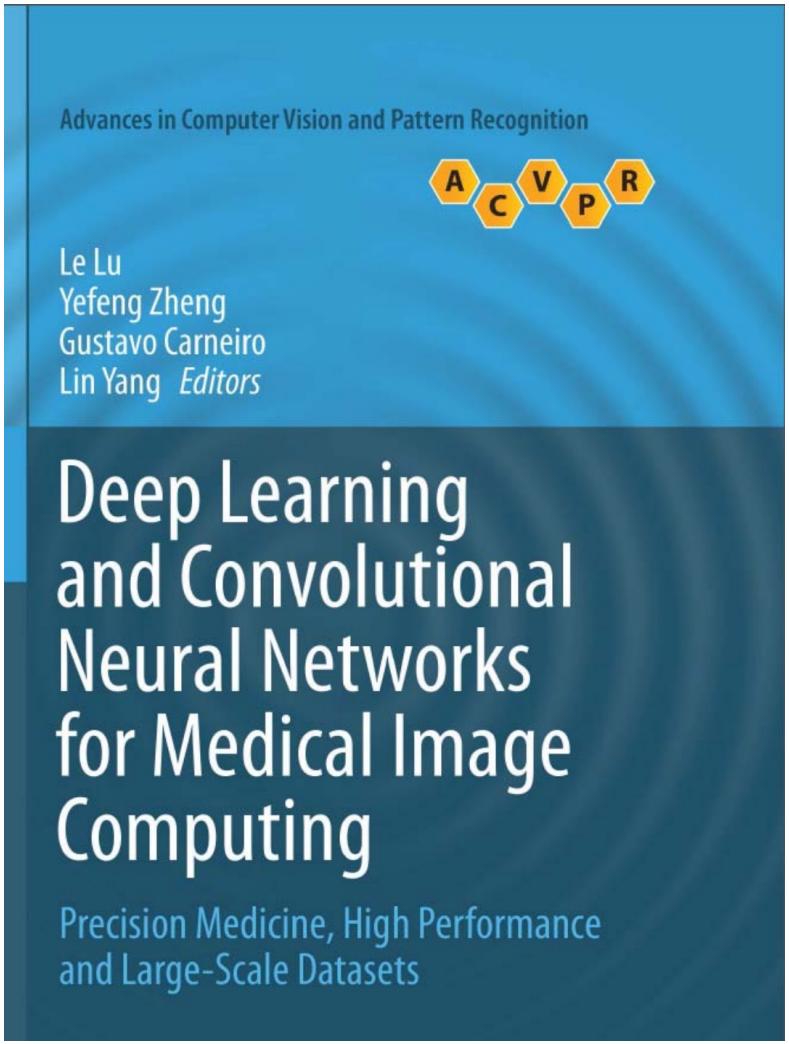

DAILY NEWS 22 November 2016

Google's DeepMind agrees new deal to share NHS patient data

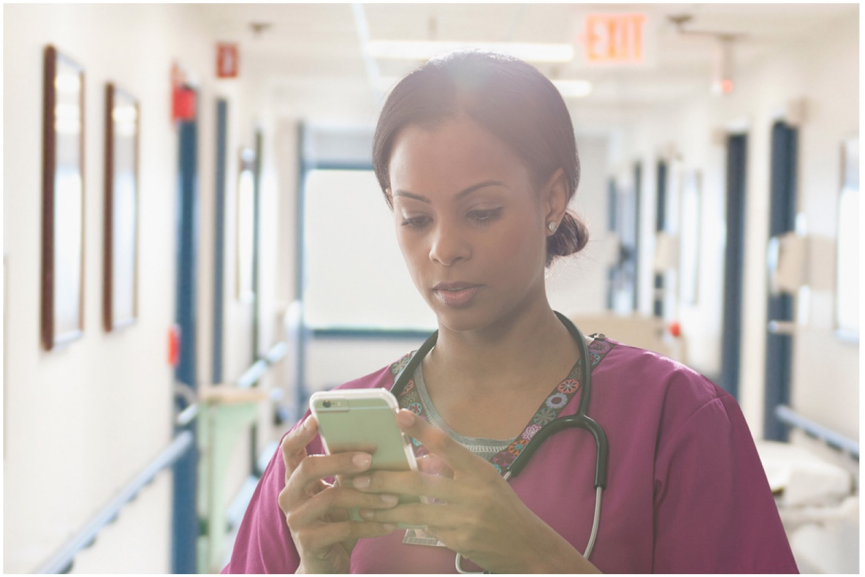

Phone alerts could save lives
Jose Luis Pelaez Inc/Getty

By **Victoria Turk**

Google's DeepMind has announced a five-year agreement with a UK National Health Service (NHS) trust that will give it access to patient data to develop and deploy its healthcare app, Streams.

Outline

Long Short-Term Memory

Hochreiter & Schmidhuber (1997)

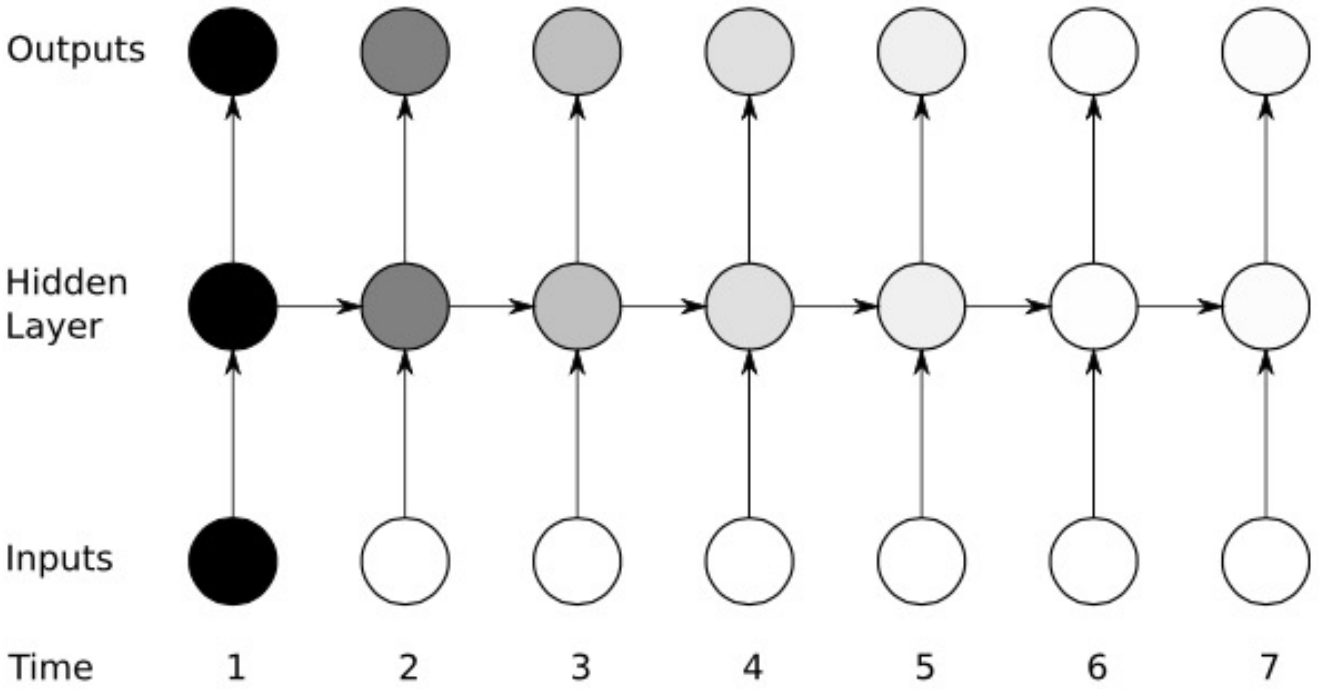

Word2Vec

Mikolov, Sutskever, Chen, Corrado & Dean (2013)

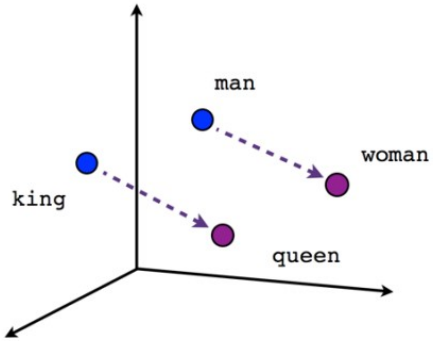

Male-Female

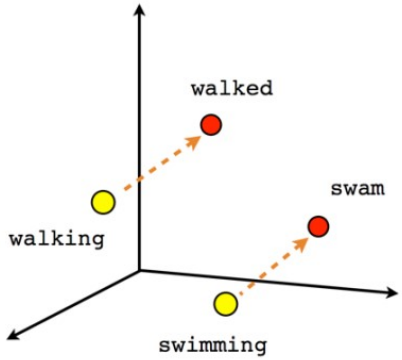

Verb tense

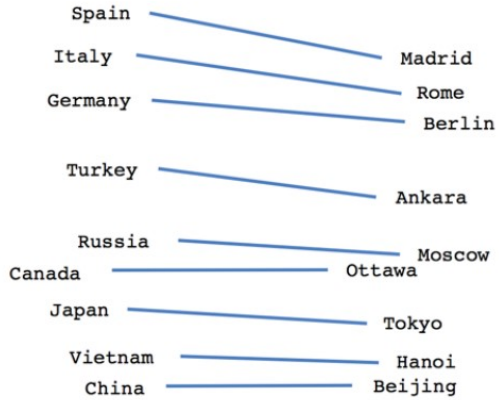

Country-Capital

t-SNE

Hinton & van der Maaten (2008)

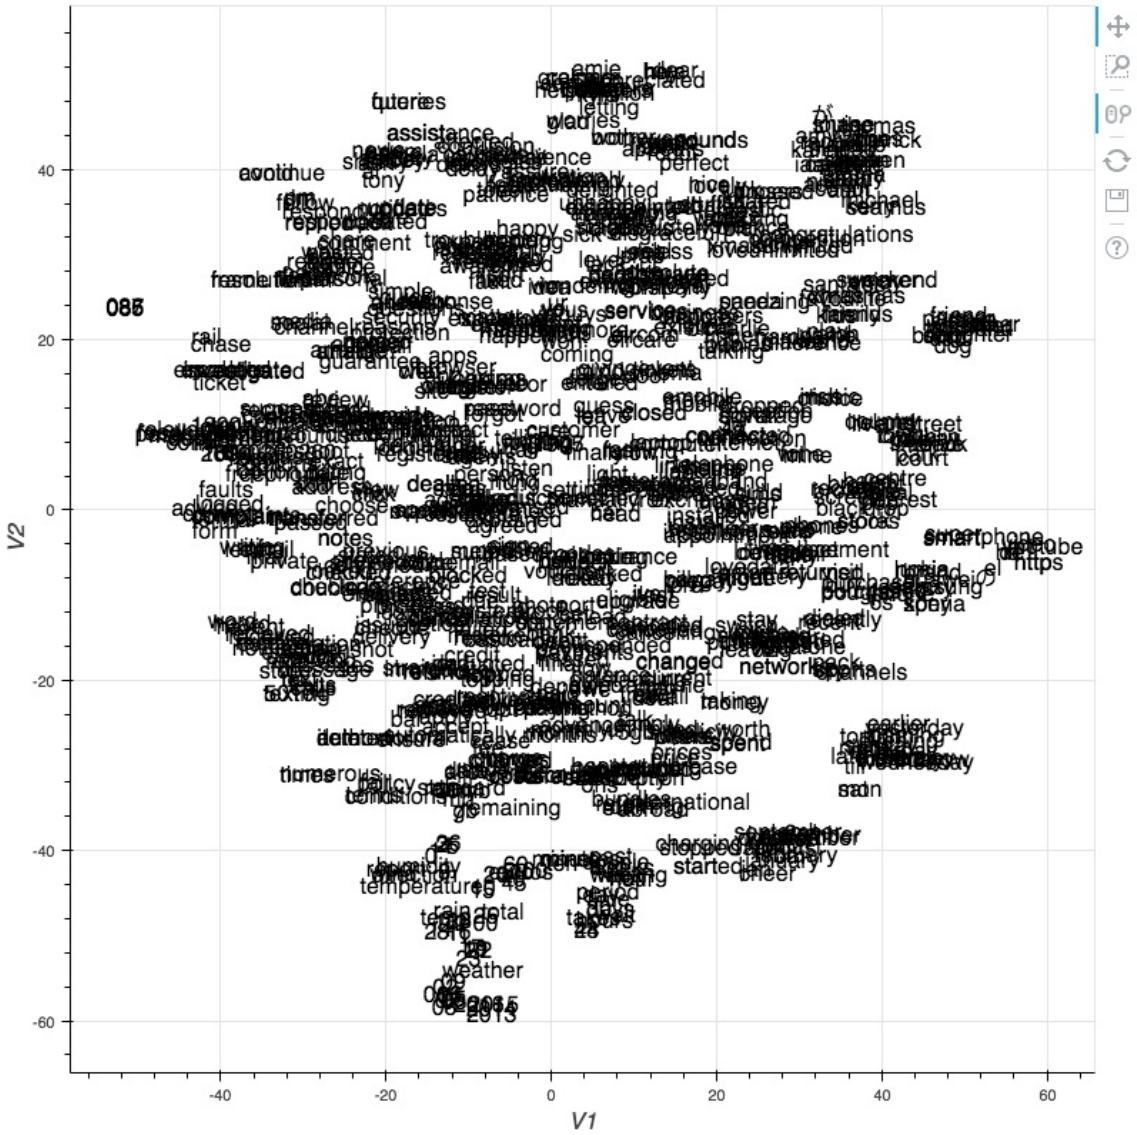

Word2Vec + t-SNE

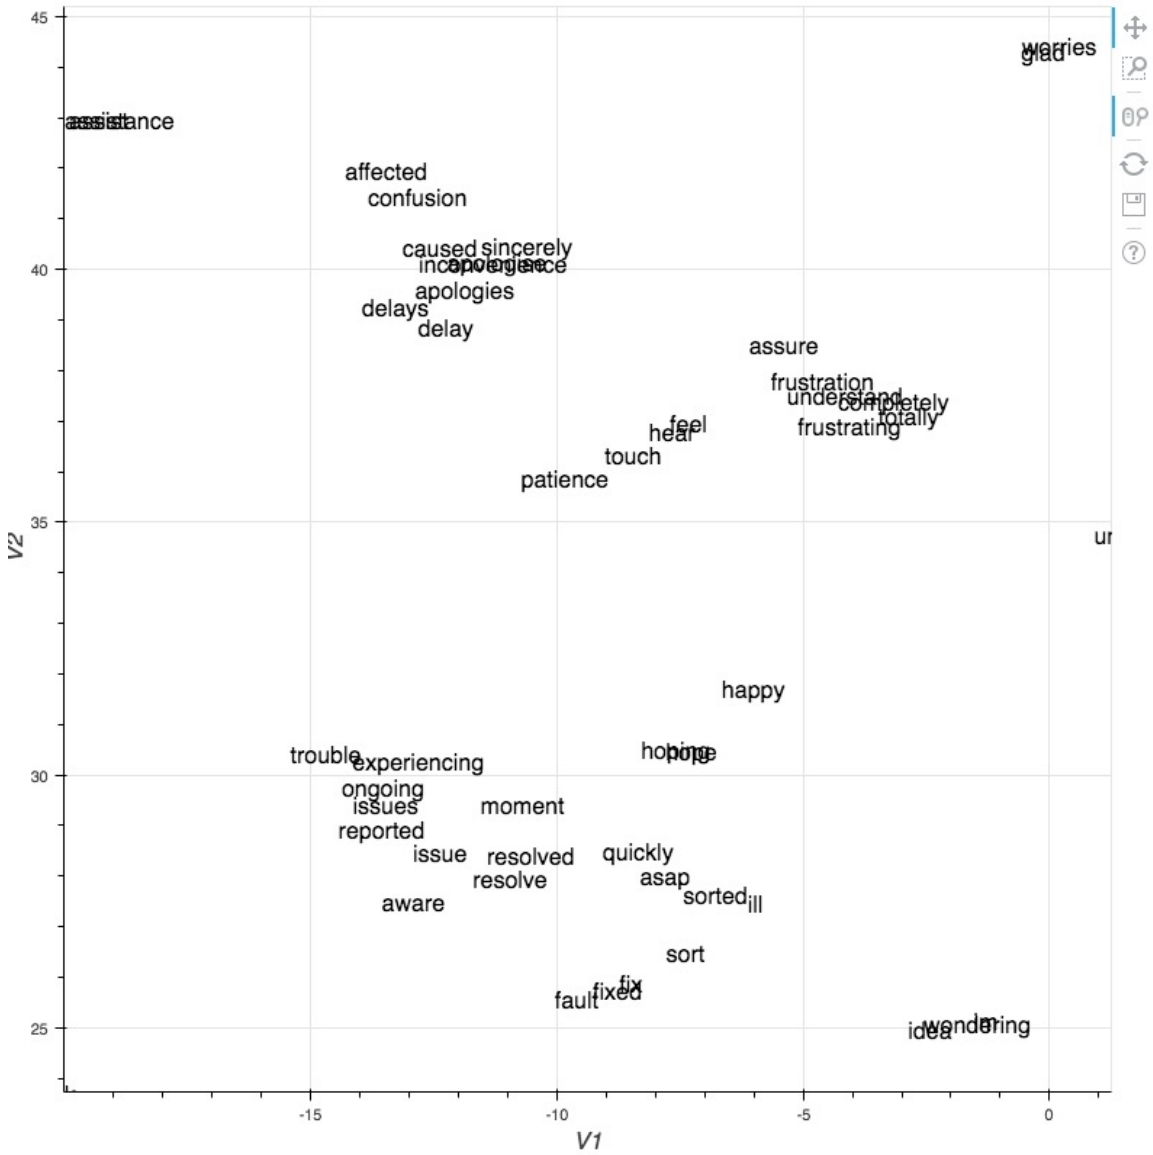

Word2Vec + t-SNE

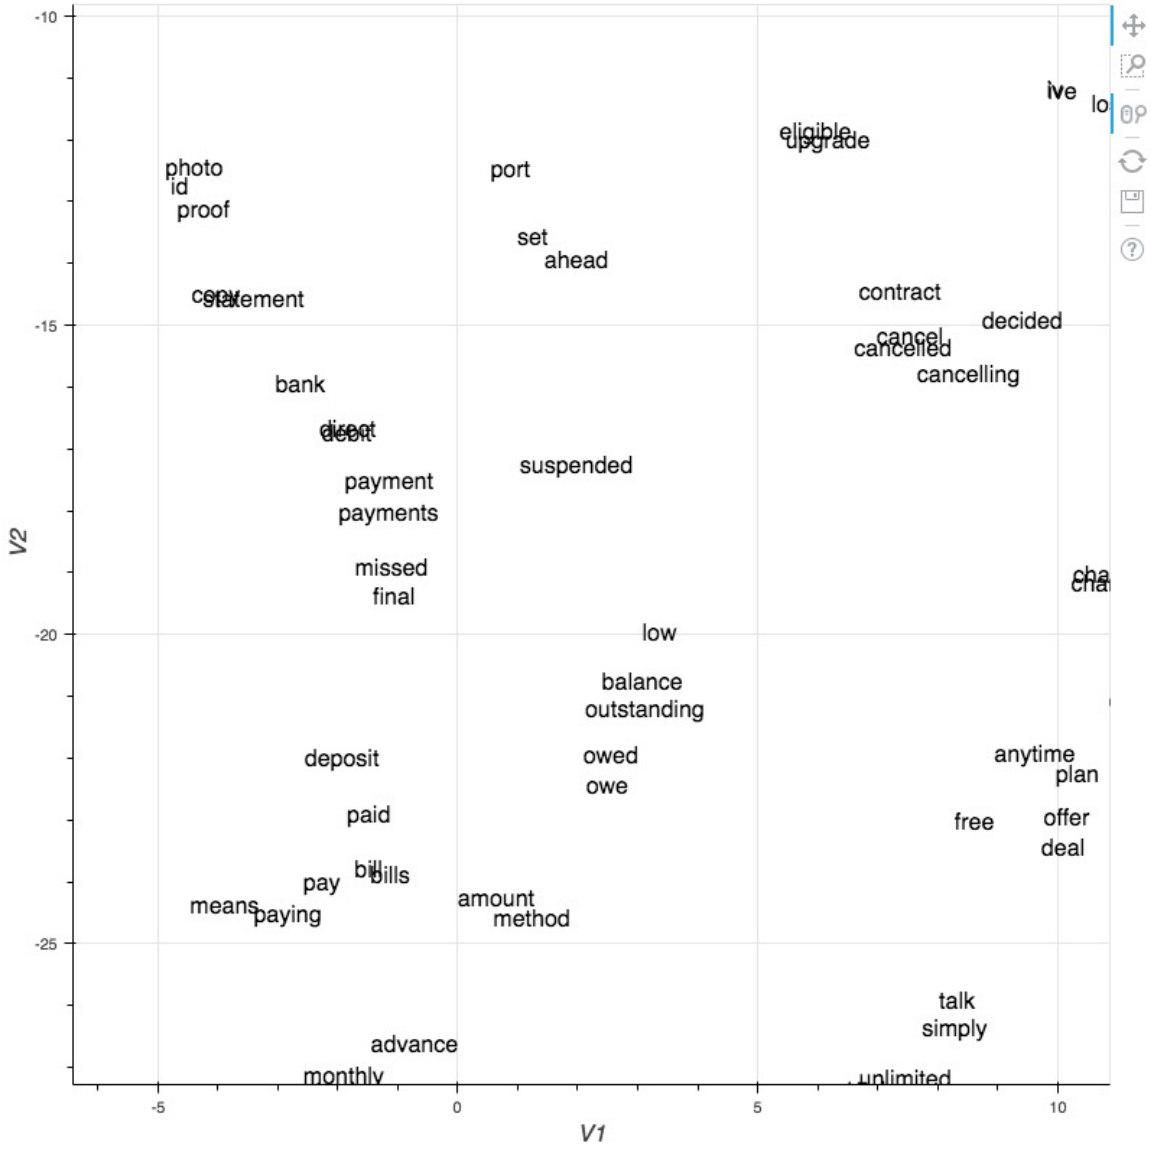

Word2Vec + t-SNE

```
model.most_similar(positive=['angular'])
```

```
[('angularjs', 0.9534549117088318),
 ('backbonejs', 0.9315043687820435),
 ('ember', 0.905410647392273),
 ('emberjs', 0.9029799103736877),
 ('reactjs', 0.896049439907074),
 ('requirejs', 0.8759748339653015),
 ('coffeescript', 0.8645504713058472),
 ('bootstrap', 0.8554328083992004),
 ('nodejs', 0.8515532612800598),
 ('backbone', 0.8443130254745483)]
```

```
model.most_similar(positive=['managed'])
```

```
[('oversaw', 0.8659406900405884),
 ('directed', 0.8491166234016418),
 ('supervised', 0.8058902621269226),
 ('coordinated', 0.7858685851097107),
 ('led', 0.7539615035057068),
 ('orchestrated', 0.7211644649505615),
 ('supported', 0.7198437452316284),
 ('comanaged', 0.6774874925613403),
 ('encompassing', 0.6726169586181641),
 ('administered', 0.6706464886665344)]
```

Outline

untapt

Digital Recruitment Platform

untapt

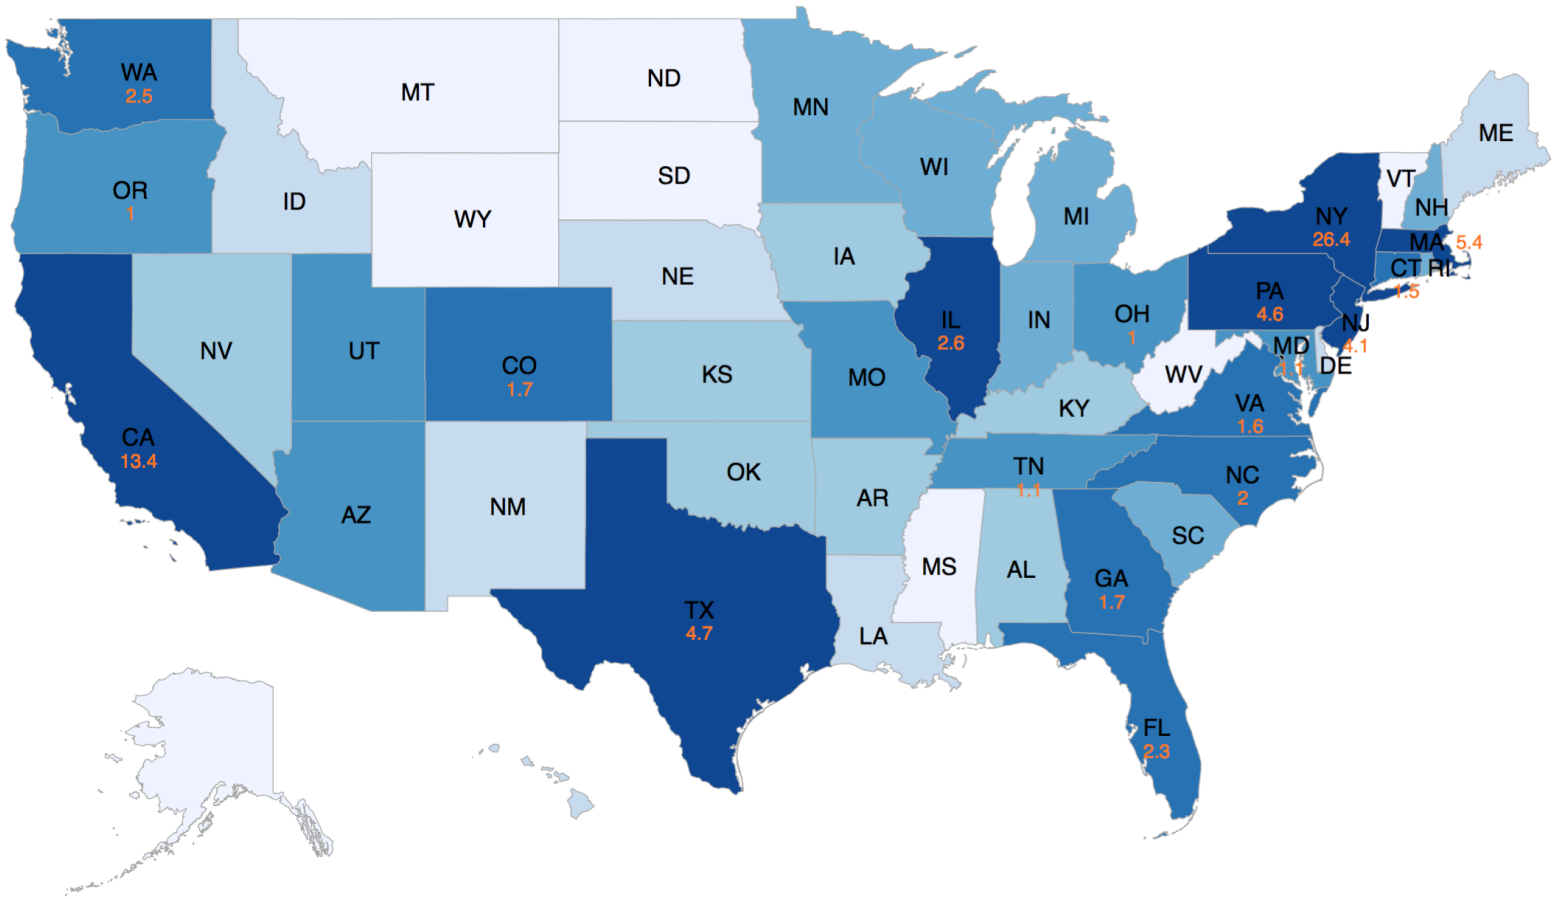
WA 2.5
OR 1
CA 13.4
ID
NV
MT
WY
UT
AZ
CO 1.7
NM
ND
SD
NE
KS
OK
TX 4.7
MN
IA
MO
AR
LA
WI
IL 2.6
MI
IN
OH 1
KY
TN 1.1
MS
AL
GA 1.7
FL 2.3
SC
NC 2
VA 1.6
WV
PA 4.6
NY 26.4
MD 1.1
DE
NJ 4.1
CT
RI 1.5
MA 5.4
VT
NH
ME

untapt

Candidate-Side Feedback

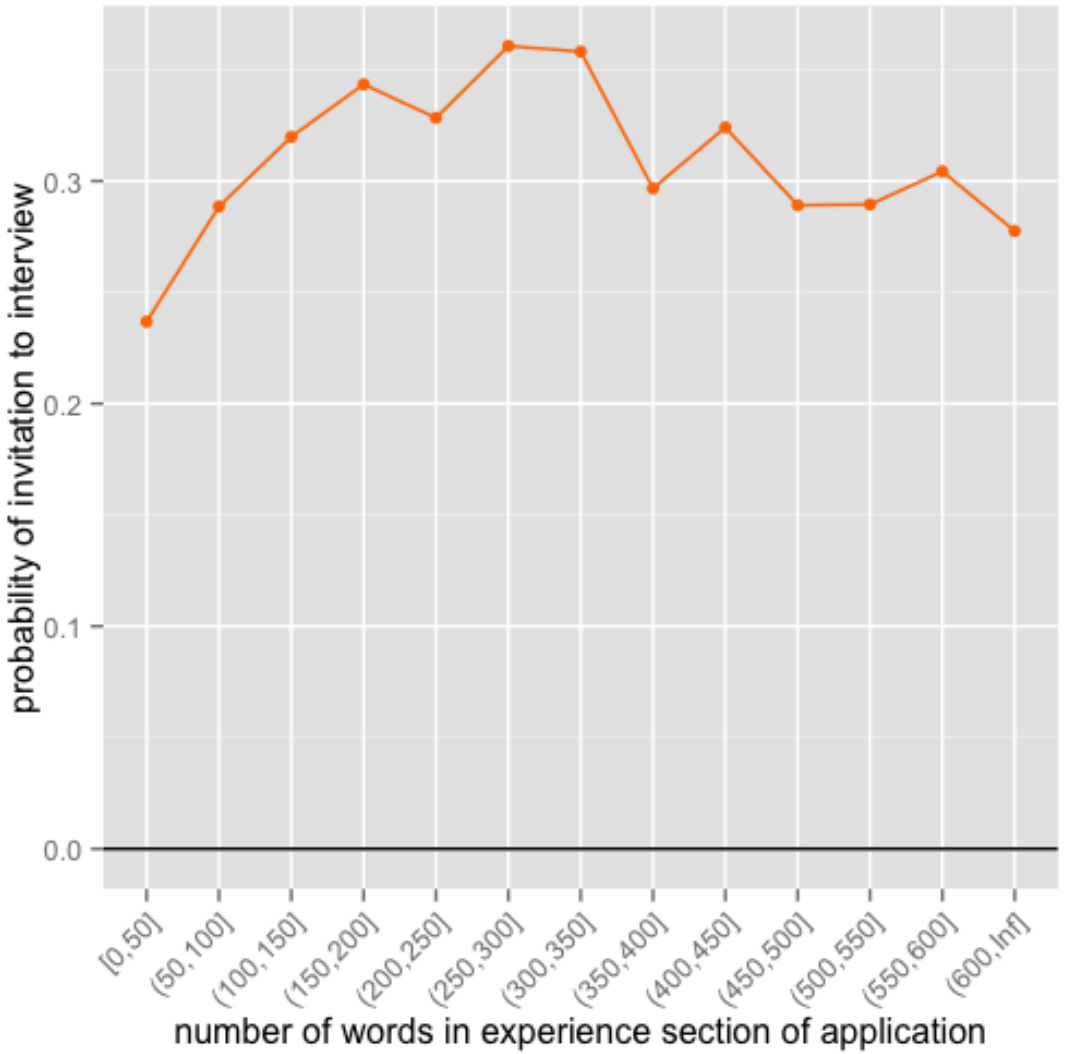

untapt

Client-Side Feedback

untapt

Multi-Stage Bayesian Regression

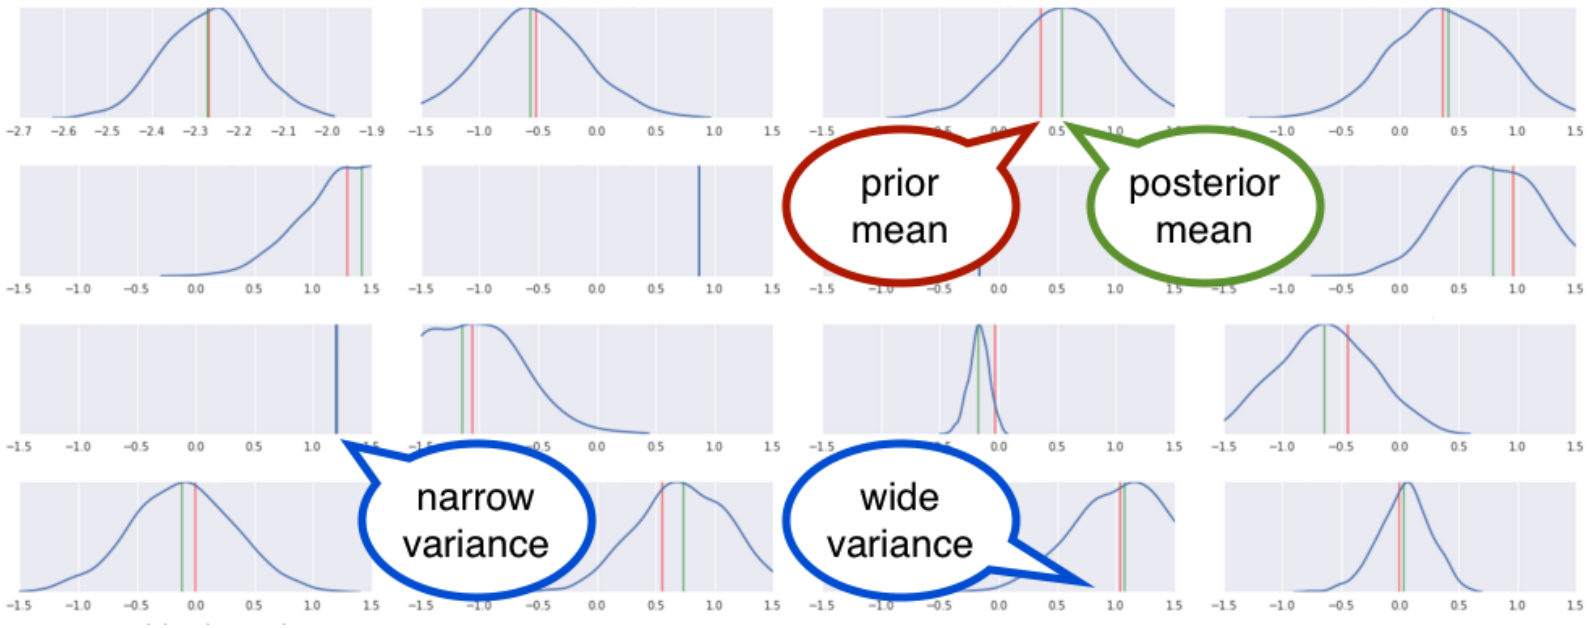

Krohn, Rives-Corbett & Donner (2016)

untapt

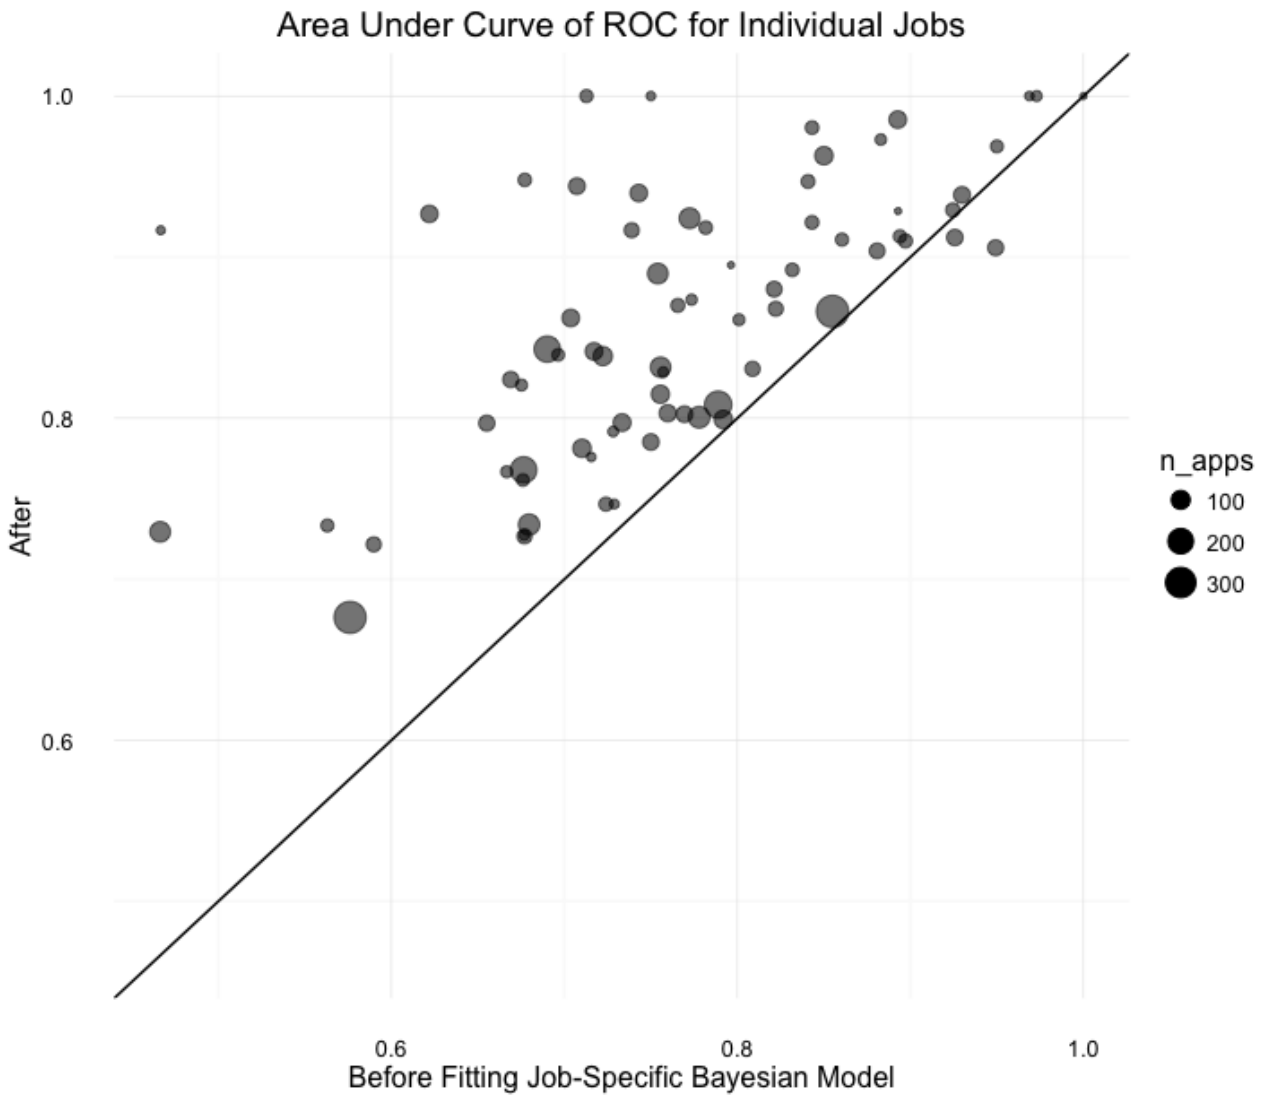

Krohn, Rives-Corbett & Donner (2016)

untapt

Ensemble with Deep Neural Net

```
Give me one bullet-point from your resume:
>> • Sat around all day checking my Facebook feed
I predict a 0.0% chance of interview
```

```
Give me one bullet-point from your resume:
>> • Developed trading applications in Python
I predict a 24.6% chance of interview
```

```
Give me one bullet-point from your resume:
>> • Developed python solution for Monte Carlo risk calculation using numpy,
 scipy and pandas, with a Javascript frontend in AngularJS and React
I predict a 98.1% chance of interview
```

deep-orange.untapt.com

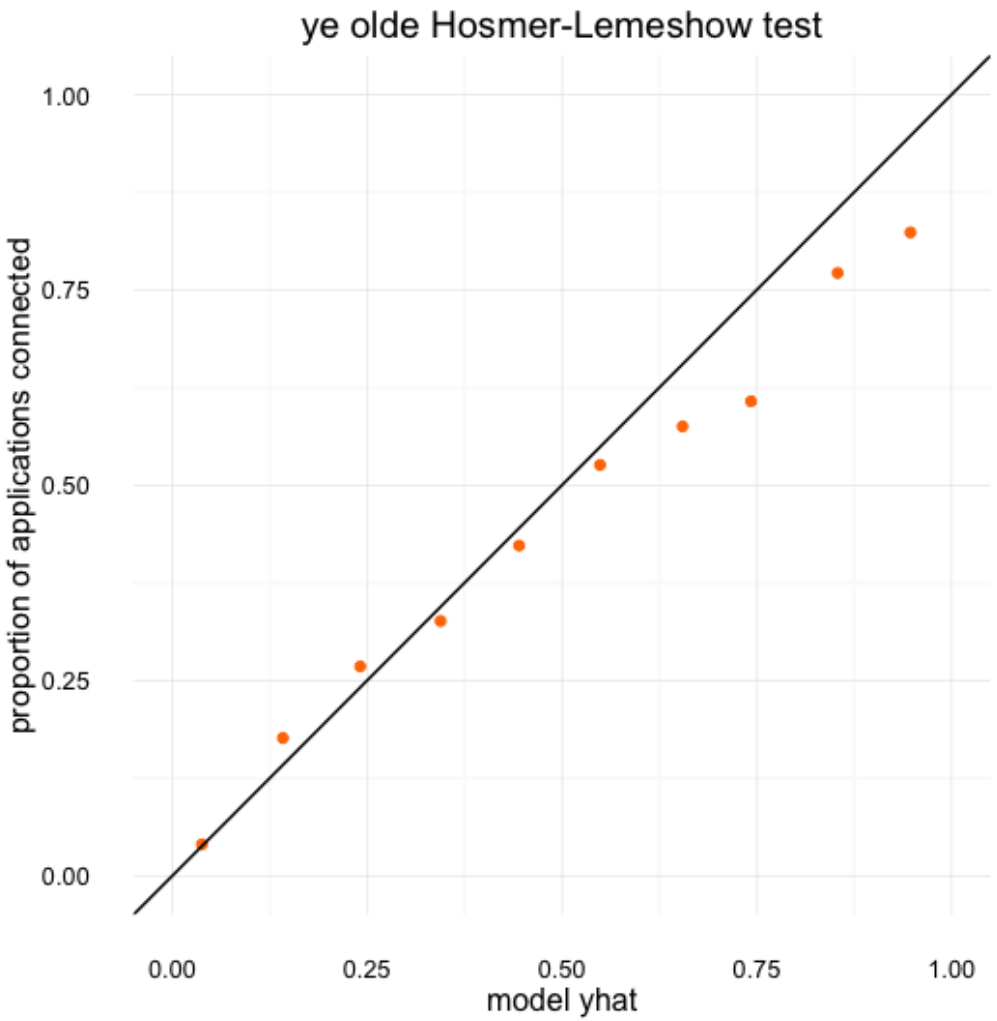
ye olde Hosmer-Lemeshow test
proportion of applications connected
1.00
0.75
0.50
0.25
0.00
0.00
0.25
0.50
0.75
1.00
model yhat


U

Outline

AlphaGO

Silver et al. (2016)

Deep Q-Learning

Mnih et al. (2015)

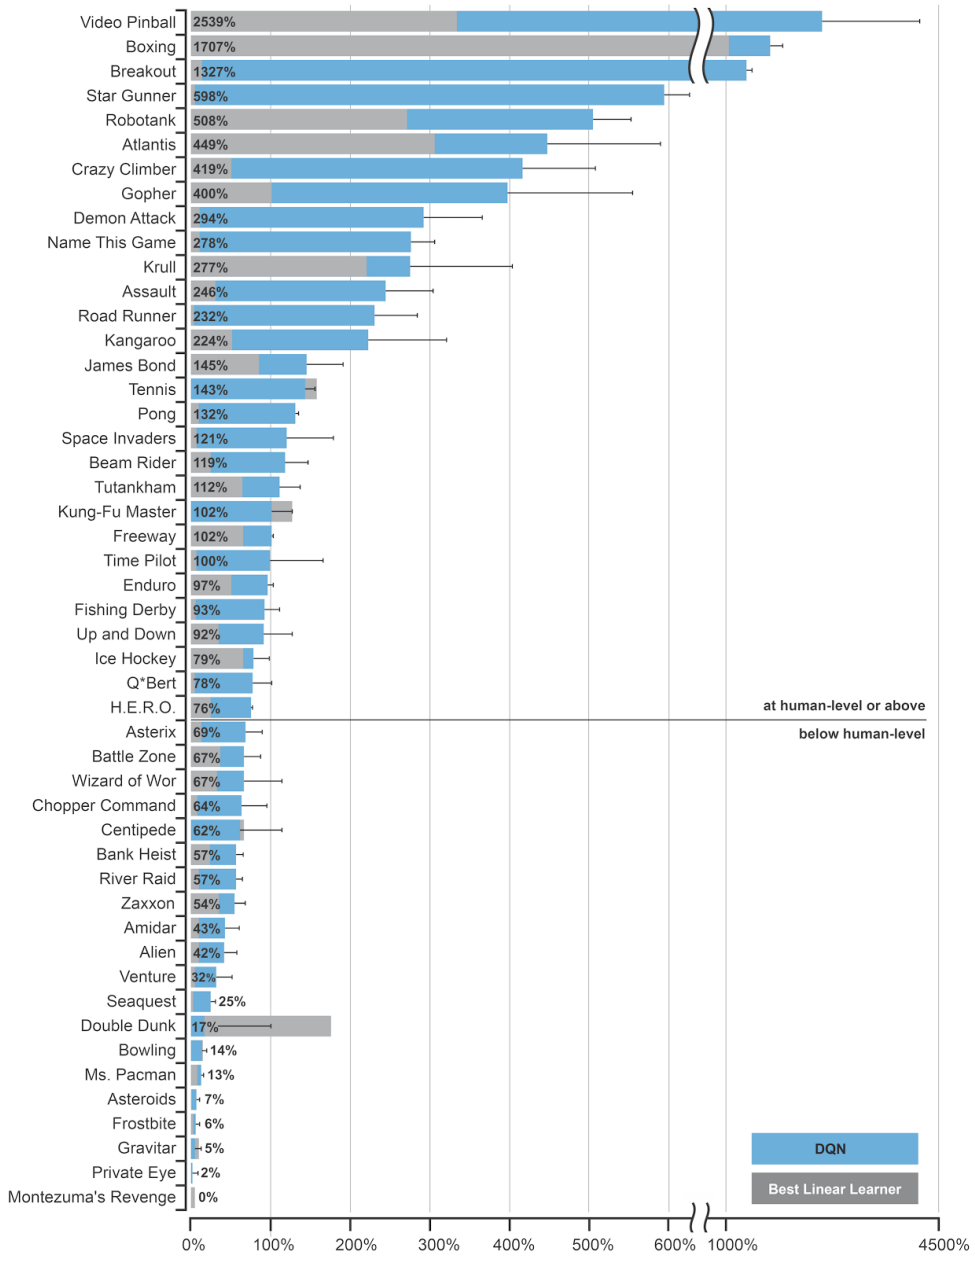

[Atari Games]

Deep Q-Learning

Mnih et al. (2015)

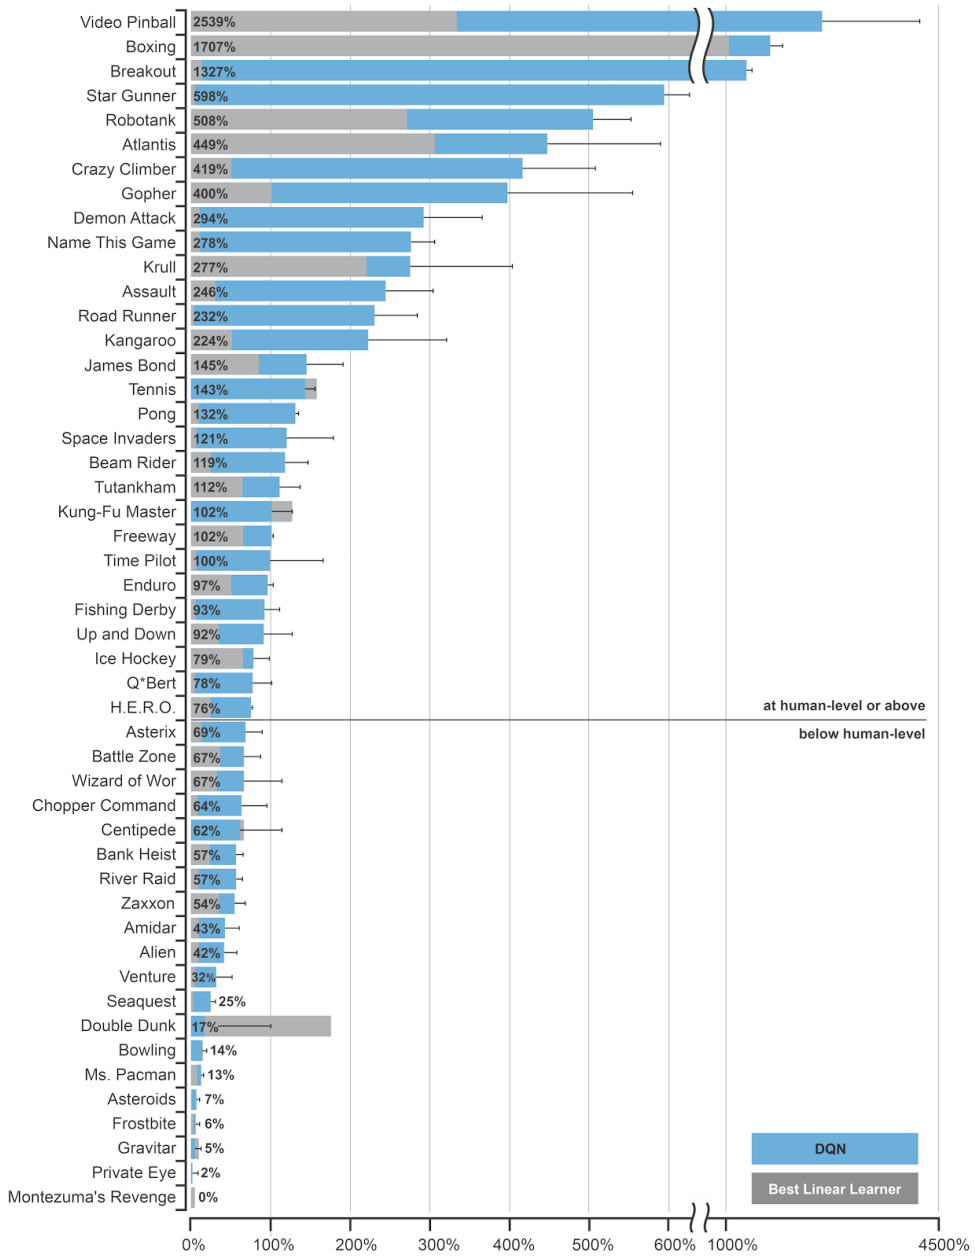

[Atari Games]

Deep Q-Learning

Mnih et al. (2015)

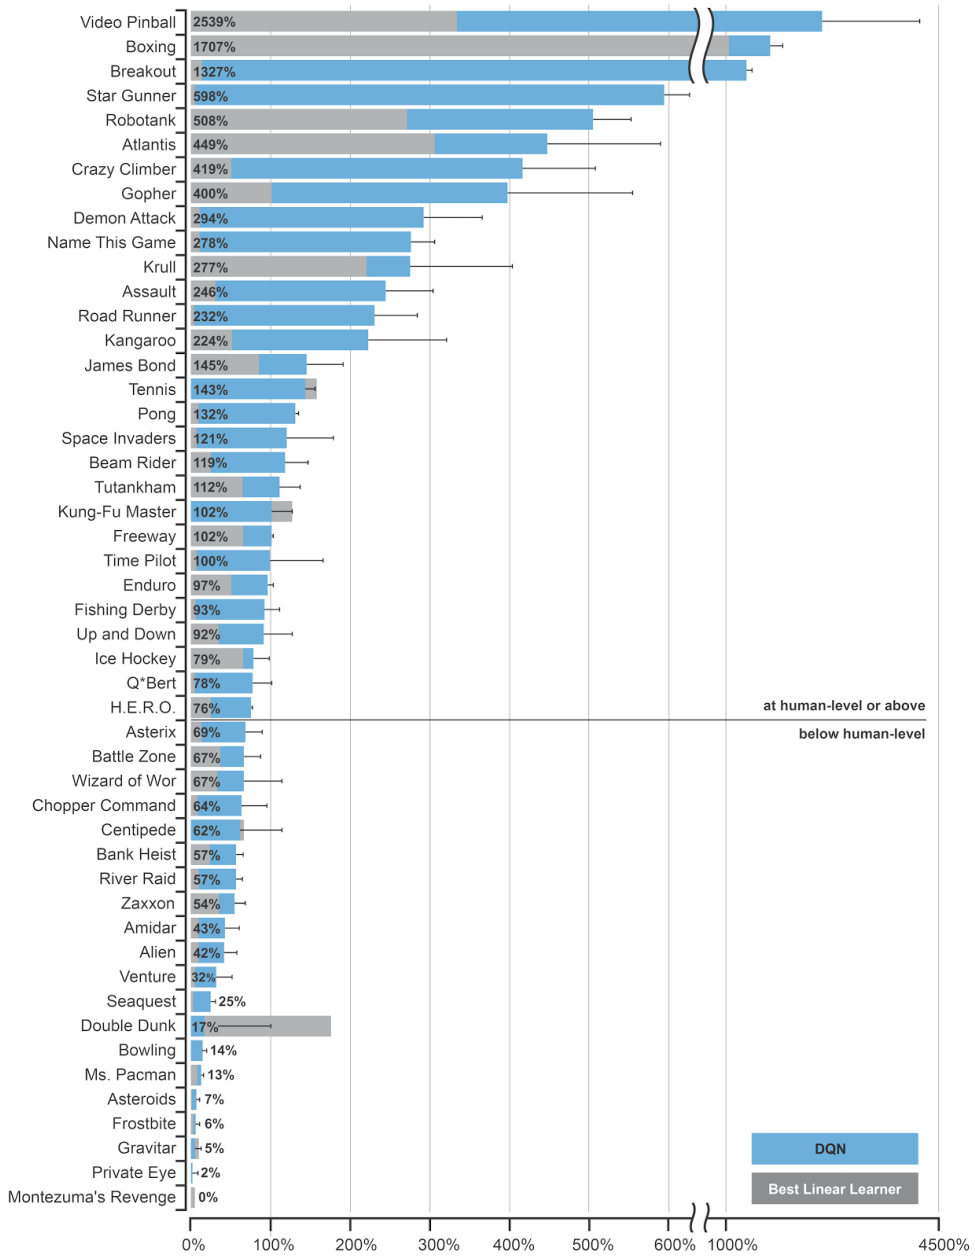

[Atari Games]

[OpenAI Universe]

[Google DeepMind Lab]

Outline

Hardware

- local machine
- build your own cluster
- AWS
- GPU(s) / TPU(s)

Hardware

- local machine
- build your own cluster
- AWS
- GPU(s) / TPU(s)

Hardware

- local machine
- build your own cluster
- AWS
- GPU(s) / TPU(s)

Hardware

- local machine
- build your own cluster
- AWS
- GPU(s) / TPU(s)

Resources for Human Learning

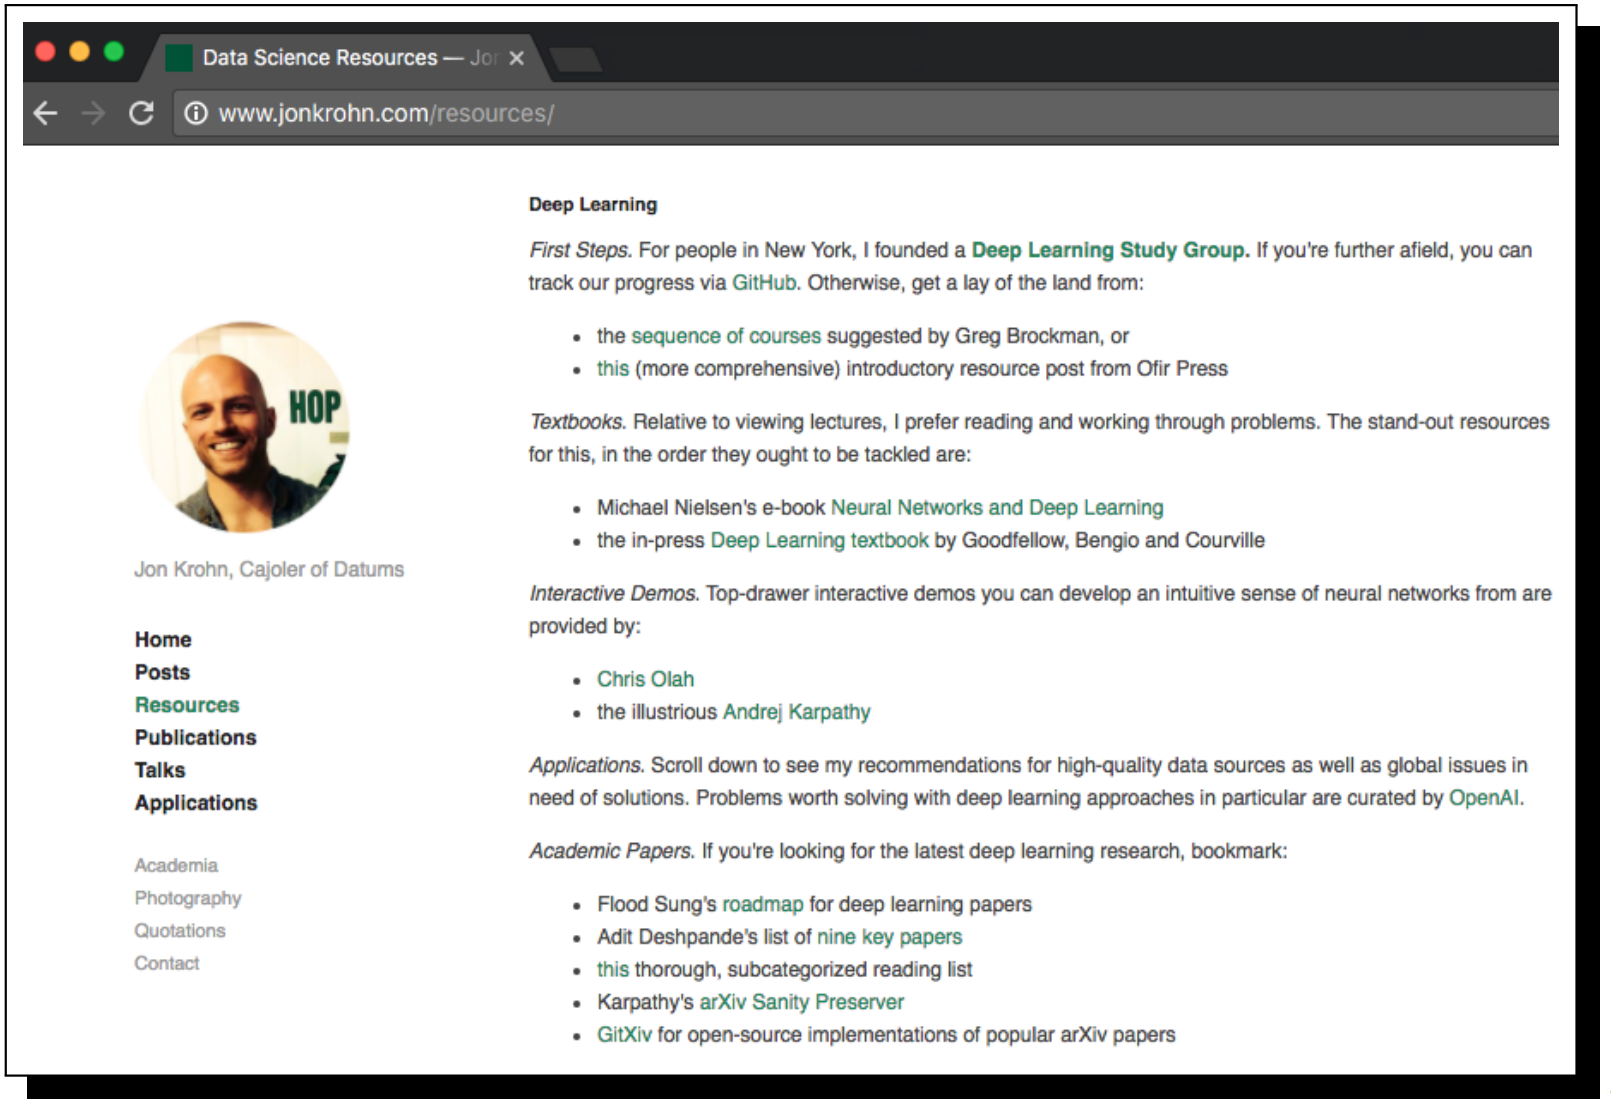

Transfer Learning

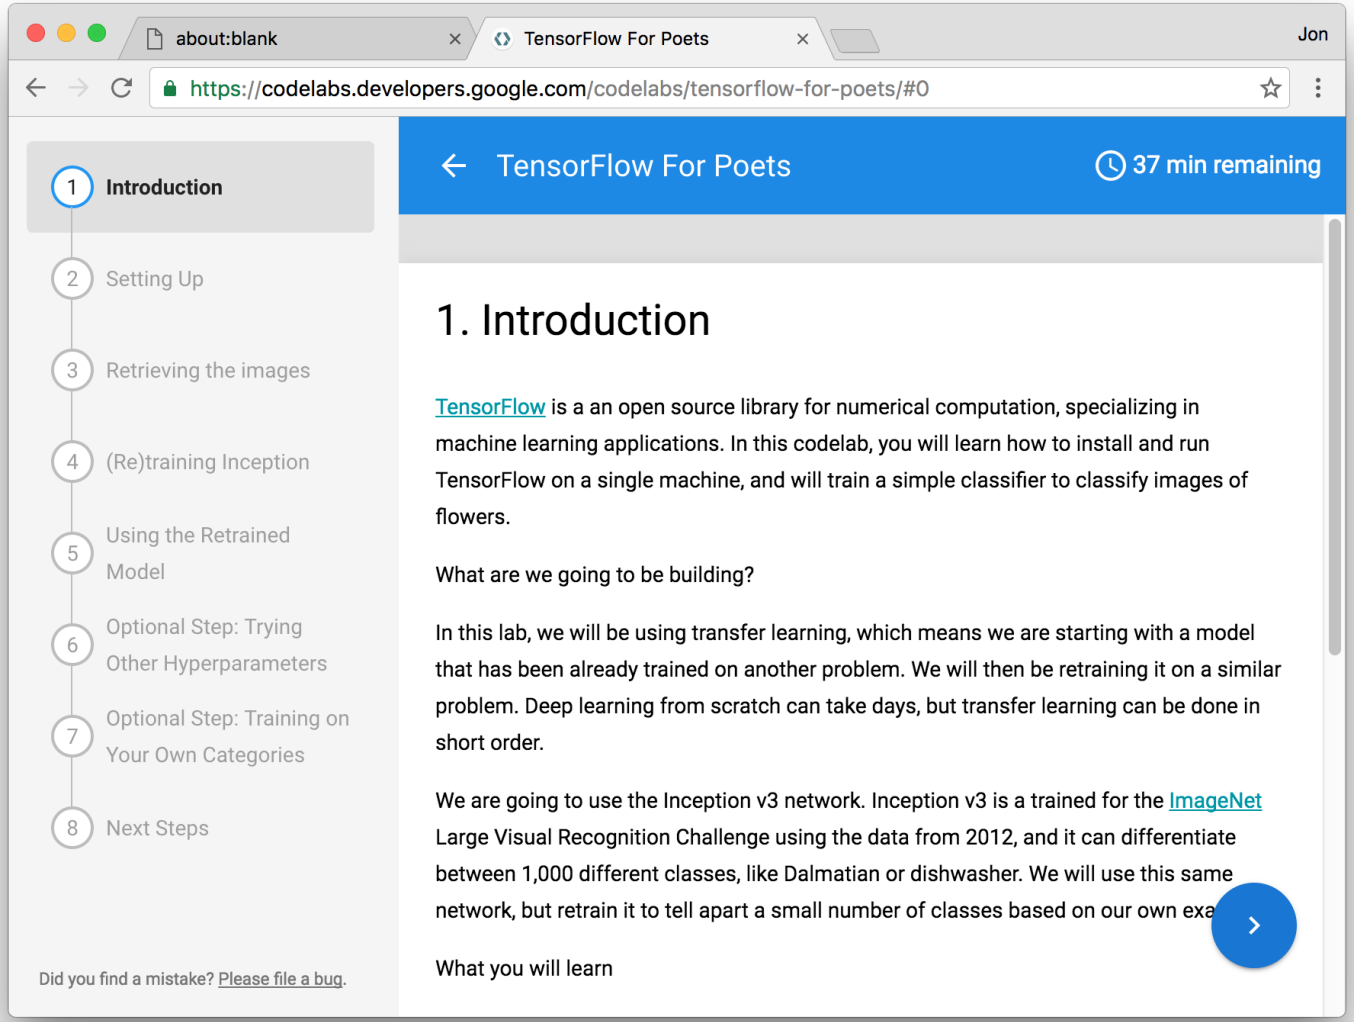

Jeanne Calment

Jeanne Calment

(1875-1997 — i.e., 122 years)

Life in the Year 2138

Thiel & Masters (2014)

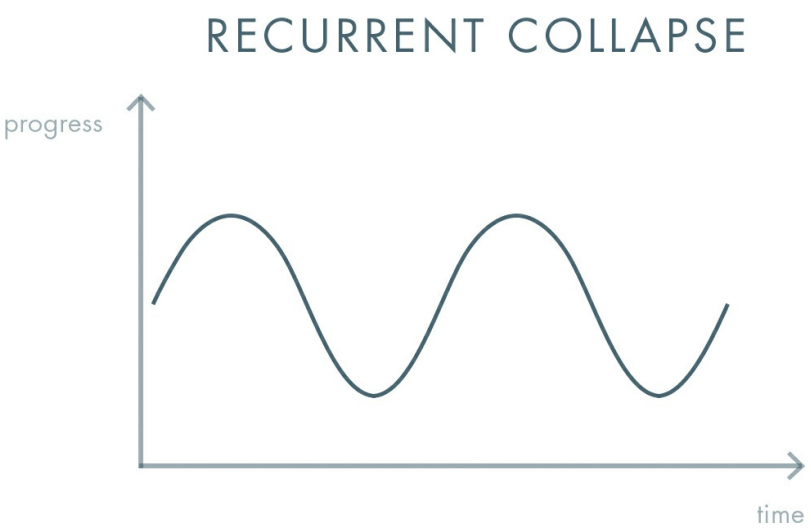

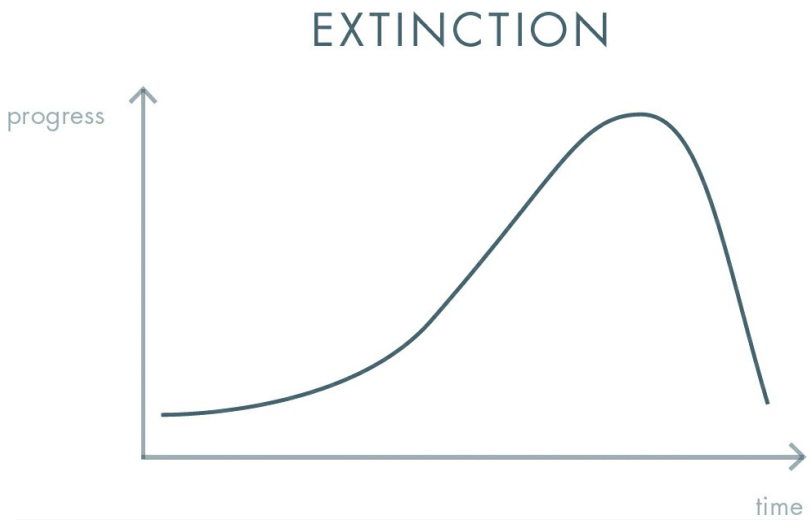

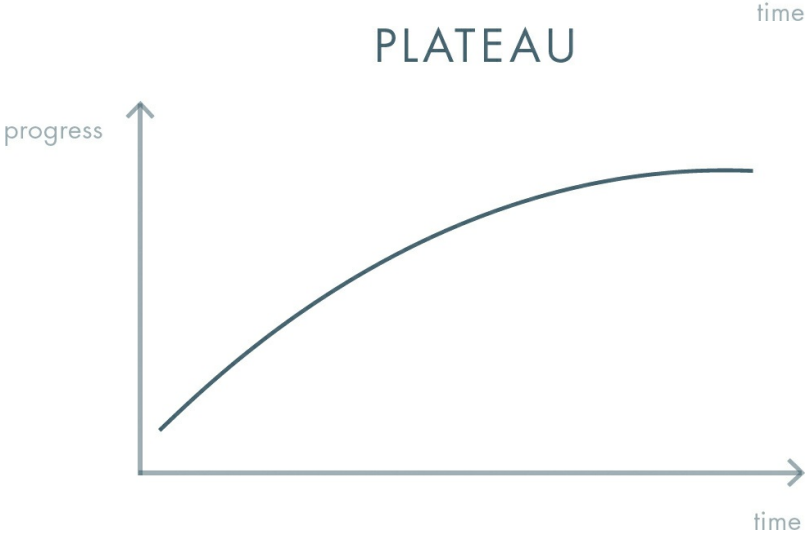

Thiel & Masters (2014)

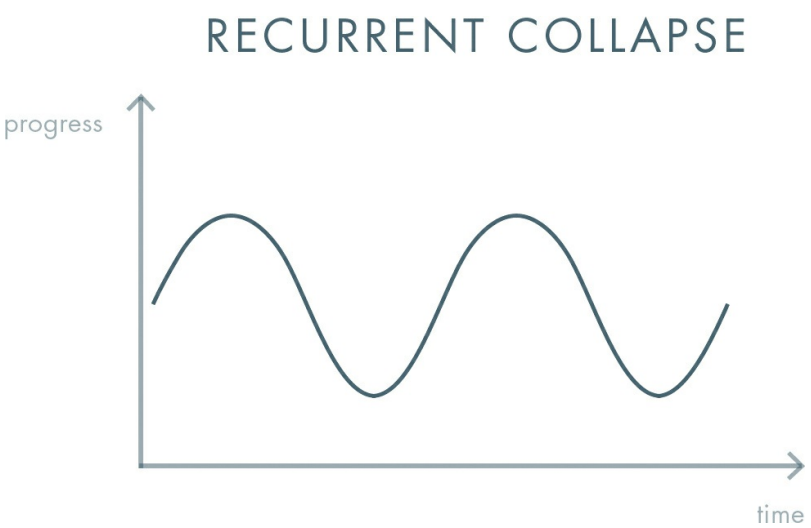

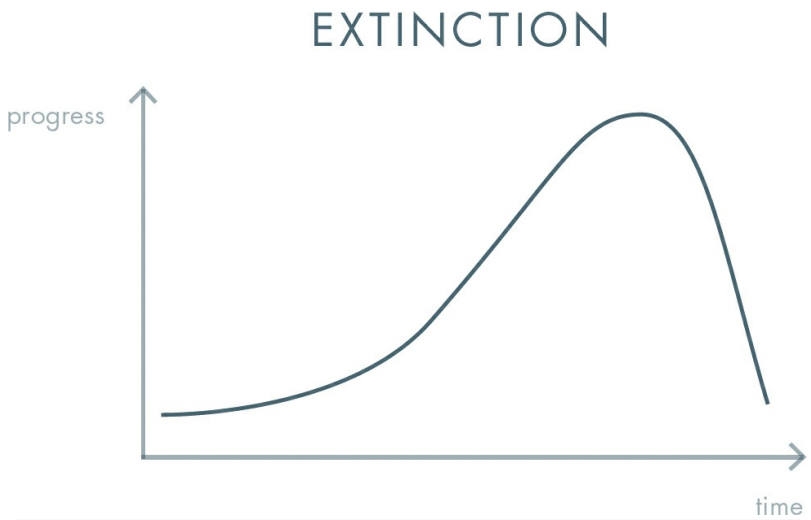

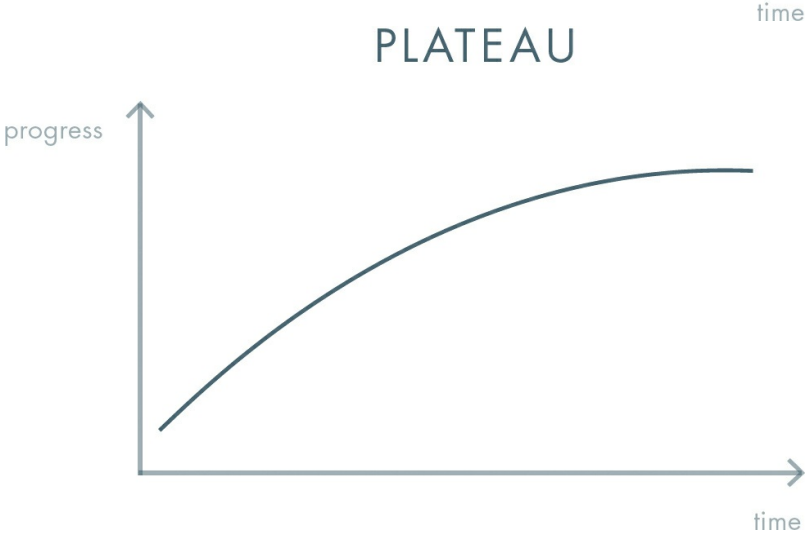

Thiel & Masters (2014)

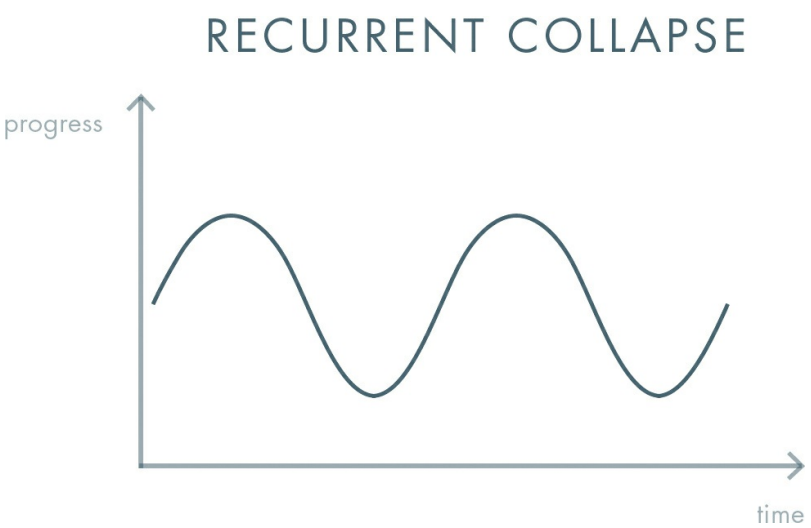

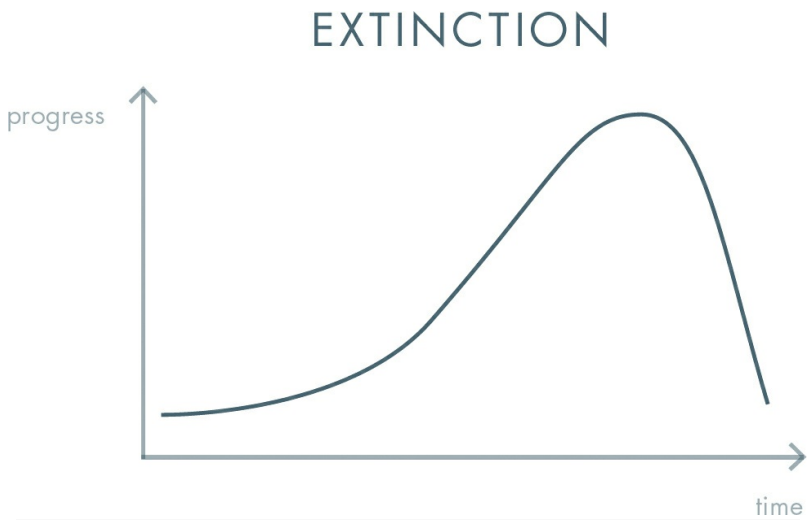

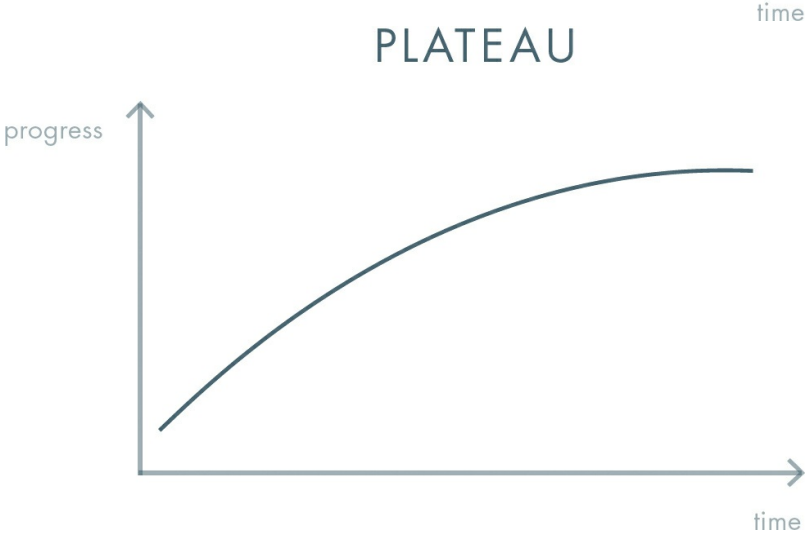

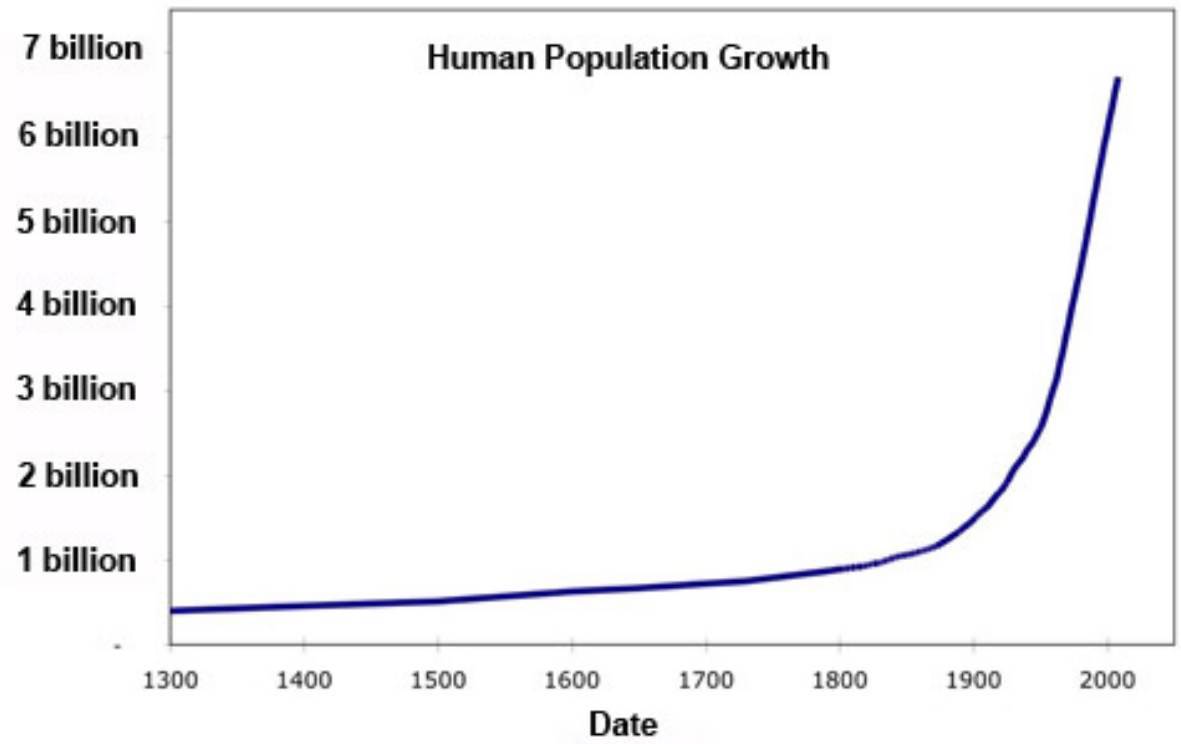
Human Population Growth
7 billion
6 billion
5 billion
4 billion
3 billion
2 billion
1 billion
1300
1400
1500
1600
1700
1800
1900
2000
Date

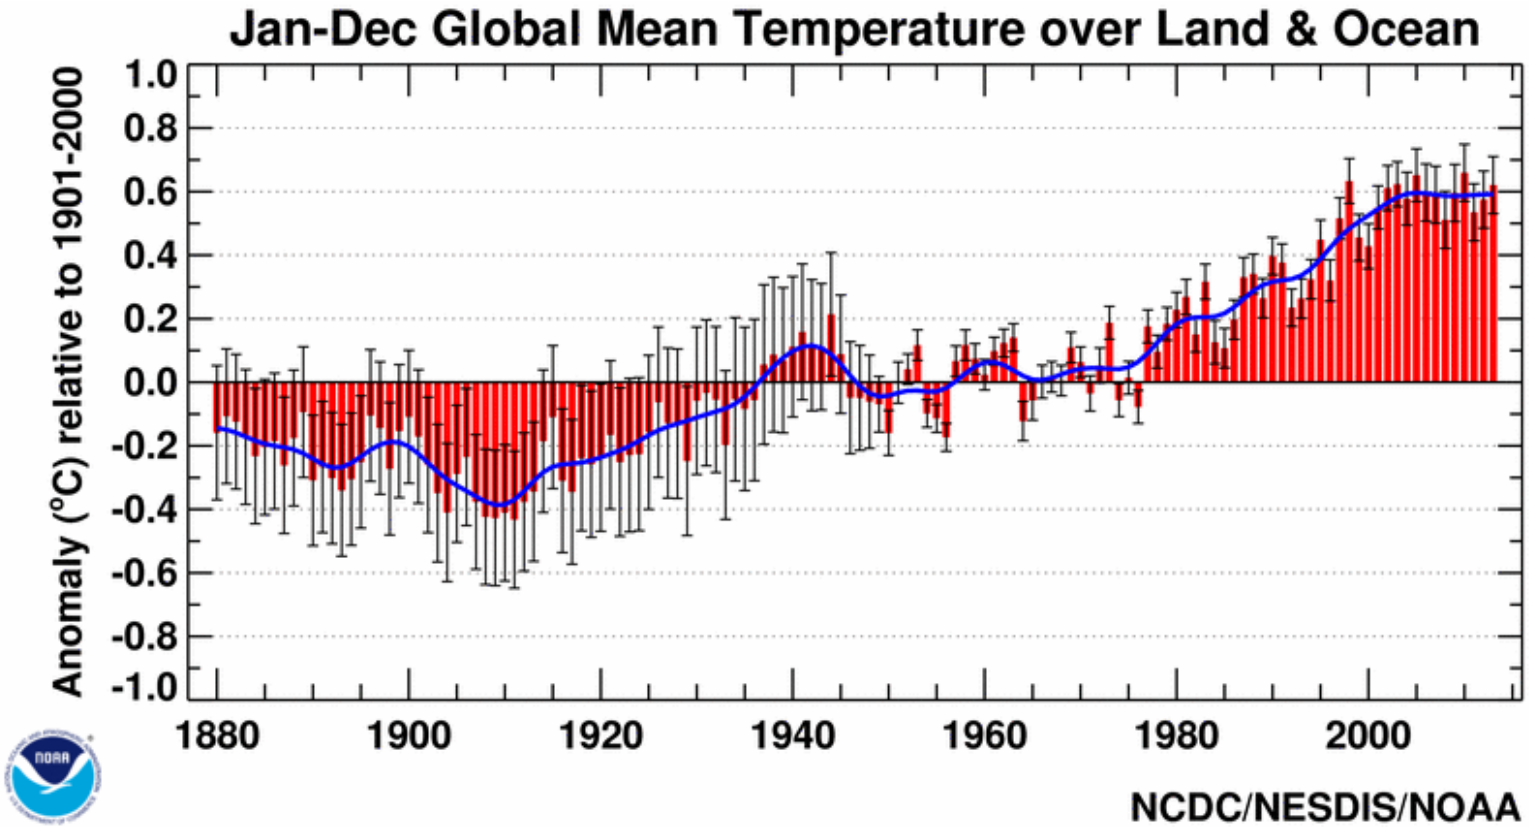
Jan-Dec Global Mean Temperature over Land & Ocean
Anomaly (°C) relative to 1901-2000
1.0
0.8
0.6
0.4
0.2
0.0
-0.2
-0.4
-0.6
-0.8
-1.0
1880
1900
1920
1940
1960
1980
2000
NCDC/NESDIS/NOAA

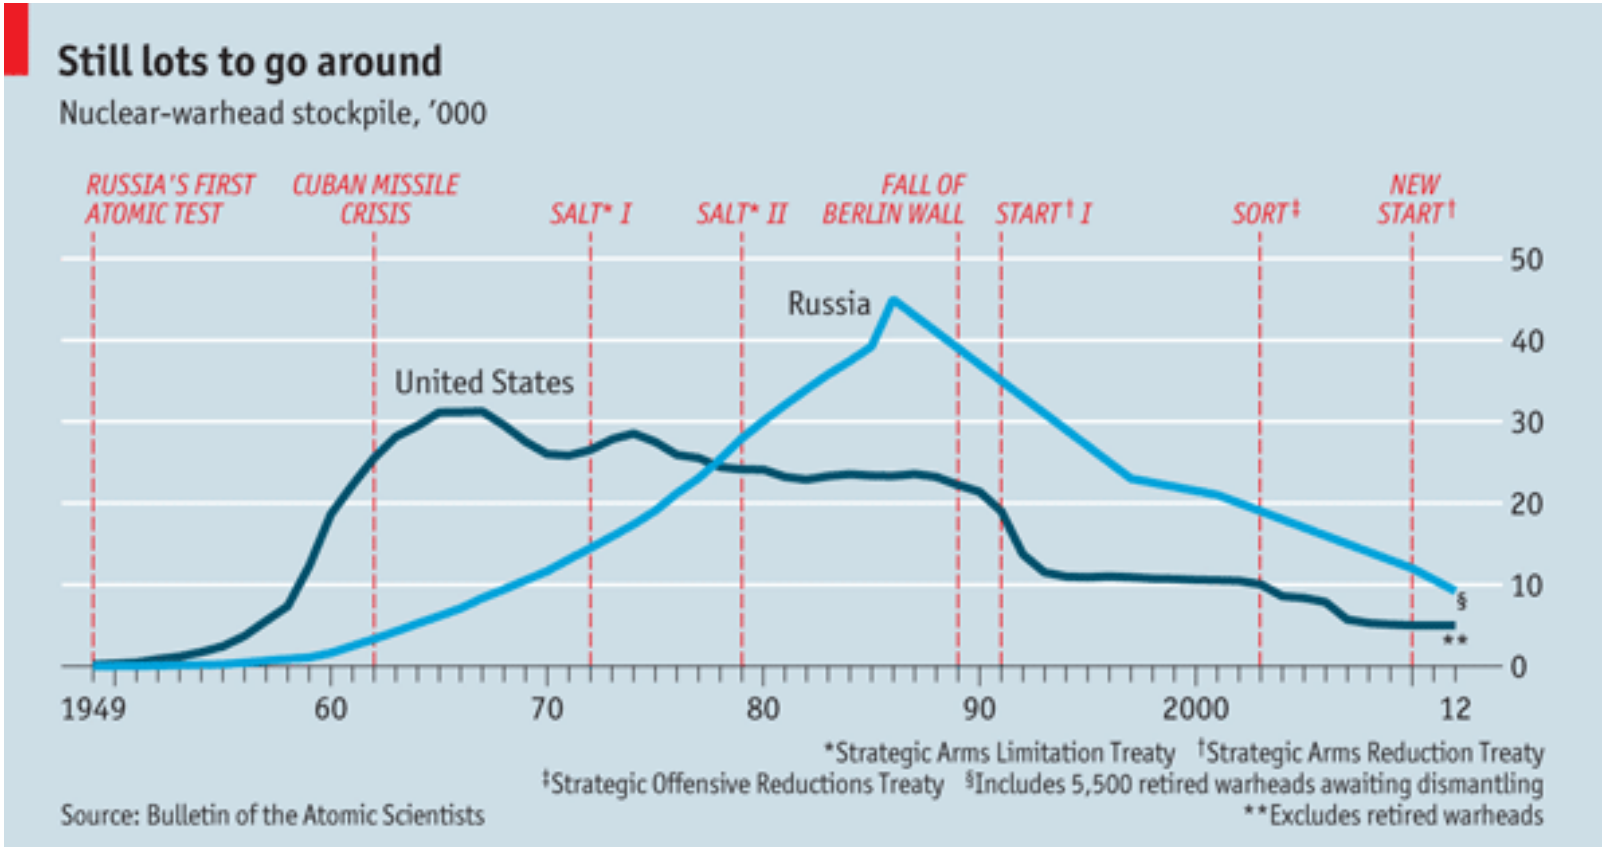
Still lots to go around
Nuclear-warhead stockpile, '000
RUSSIA'S FIRST ATOMIC TEST
CUBAN MISSILE CRISIS
SALT* I
SALT* II
FALL OF BERLIN WALL
START† I
SORT‡
NEW START†
Russia
United States
50
40
30
20
10
0
§
**
1949
60
70
80
90
2000
12
*Strategic Arms Limitation Treaty †Strategic Arms Reduction Treaty
‡Strategic Offensive Reductions Treaty §Includes 5,500 retired warheads awaiting dismantling
**Excludes retired warheads
Source: Bulletin of the Atomic Scientists

Pinker & Mack (2014)

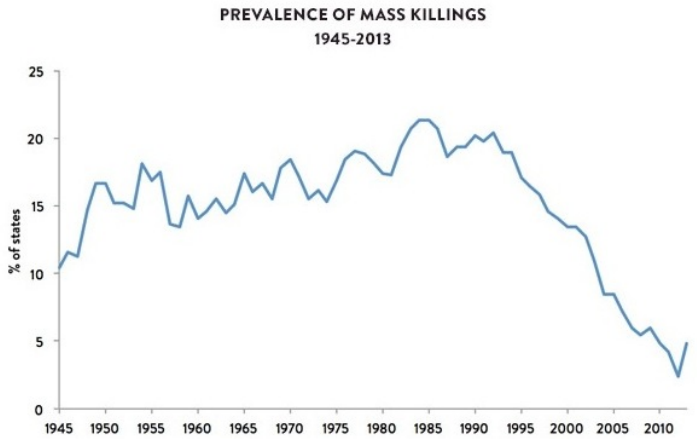

Source: Early Warning Project, https://cpgearlywarning.wordpress.com/; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Feb. 1, 2008). Available at SSRN: http://ssrn.com/abstract=1703426 or http://dx.doi.org/10.2139/ssrn.1703426

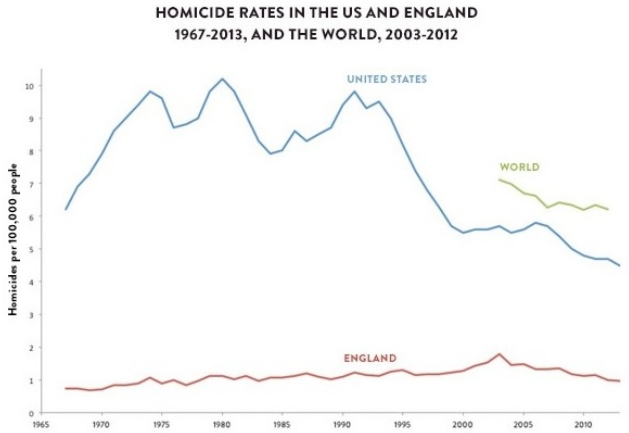

Sources: U.S., FBI Uniform Crime Reports. England (including Wales): U.K. Office for National Statistics. World: U.N. Office on Drugs and Crime, reported in U.N. Economic and Social Council's "World crime trends and emerging issues and responses in the field of crime prevention and social justice," Feb. 12, 2014, Figure 1. The percentages were converted to homicide rates by setting the 2012 rate at 6.2, the figure reported in the UNODC Global Study on Homicide 2013, Page 12.

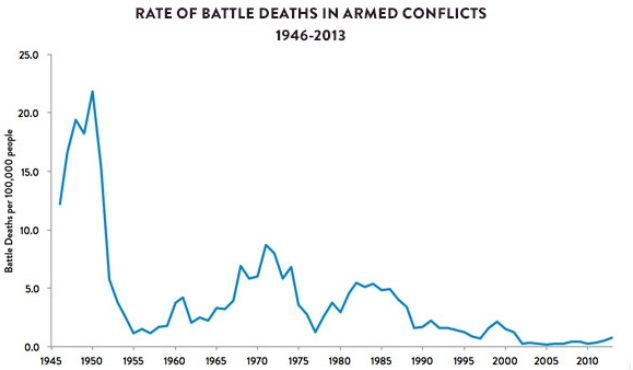

Source: Uppsala Conflict Data Program Battle-Related Deaths Dataset v.5-2014, 1989-2013, "Best Estimates." Missing estimate for 2013 for the Syrian civil war calculated from a revised UCDP "Low Estimate" and a "High Estimate" from the Peace Research Institute Oslo (both obtained in consultation with Erik Melander of UCDP). The "High Estimate" is for May 2013–April 2014, rather than calendar year 2013, and comes from Page 9 of M. Price, A. Gohdes, & P. Ball, "Updated Statistical Analysis of Documentation of Killings in the Syrian Arab Republic," Human Rights Data Analysis Group, http://www.ohchr.org/Documents/Countries/SY/HRDAGUpdatedReportAug2014.pdf. World population figures from U.S. Census Bureau, http://www.census.gov/population/international/data/worldpop/table_population.php

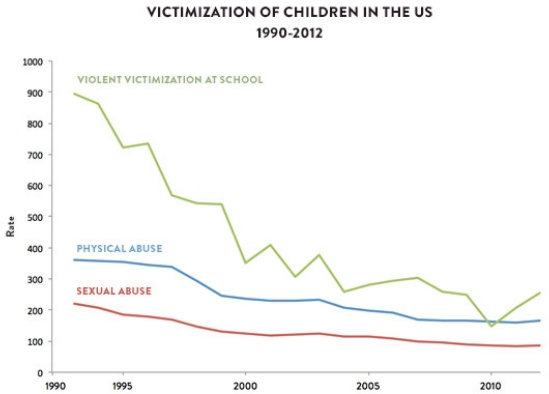

Rates for physical and sexual abuse are per 100,000 children younger than 18. Rates for violent victimization at school are per 10,000 children age 12–17.

Sources: Physical and sexual abuse, National Child Abuse and Neglect Data System, analyzed by David Finkelhor, 2014, in "Trends in Child Welfare," Presentation at the Carsey Institute Policy Series, March 31, 2014. Victimization at school, Bureau of Justice Statistics (bjs.gov), using the National Crime Victimization Survey Victimization Analysis Tool.

Pinker & Mack (2014)

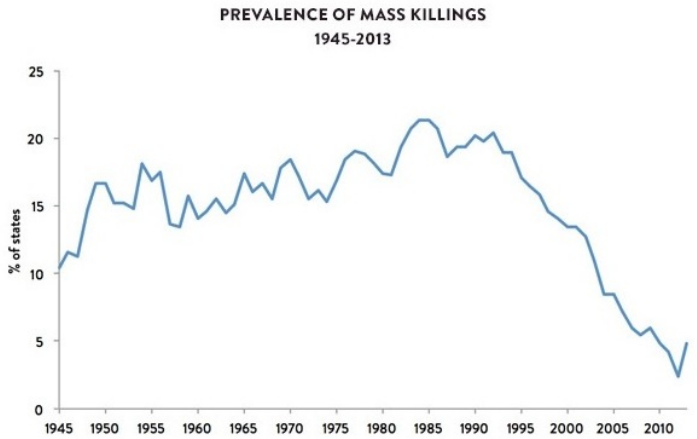

Source: Early Warning Project, https://cpgearlywarning.wordpress.com/; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Feb. 1, 2008). Available at SSRN: http://ssrn.com/abstract=1703426 or http://dx.doi.org/10.2139/ssrn.1703426

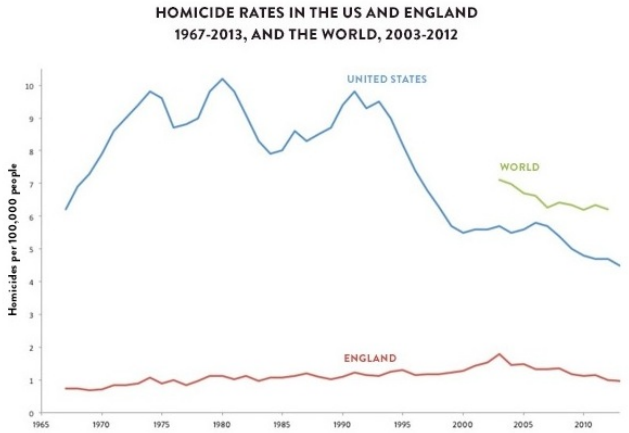

Sources: U.S., FBI Uniform Crime Reports. England (including Wales): U.K. Office for National Statistics. World: U.N. Office on Drugs and Crime, reported in U.N. Economic and Social Council's "World crime trends and emerging issues and responses in the field of crime prevention and social justice," Feb. 12, 2014, Figure 1. The percentages were converted to homicide rates by setting the 2012 rate at 6.2, the figure reported in the UNODC Global Study on Homicide 2013, Page 12.

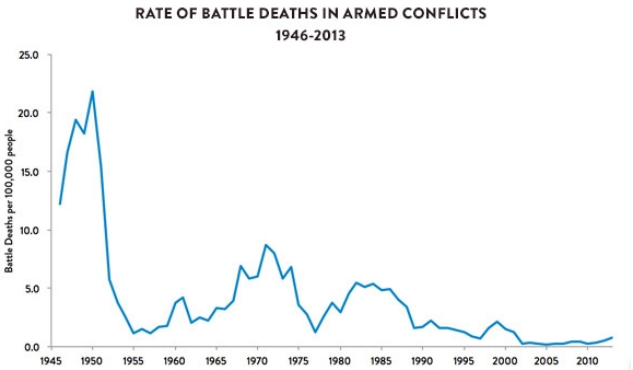

Source: Uppsala Conflict Data Program Battle-Related Deaths Dataset v.5-2014, 1989-2013, "Best Estimates." Missing estimate for 2013 for the Syrian civil war calculated from a revised UCDP "Low Estimate" and a "High Estimate" from the Peace Research Institute Oslo (both obtained in consultation with Erik Melander of UCDP). The "High Estimate" is for May 2013–April 2014, rather than calendar year 2013, and comes from Page 9 of M. Price, A. Gohdes, & P. Ball, "Updated Statistical Analysis of Documentation of Killings in the Syrian Arab Republic," Human Rights Data Analysis Group, http://www.ohchr.org/Documents/Countries/SY/HRDAGUpdatedReportAug2014.pdf. World population figures from U.S. Census Bureau, http://www.census.gov/population/international/data/worldpop/table_population.php

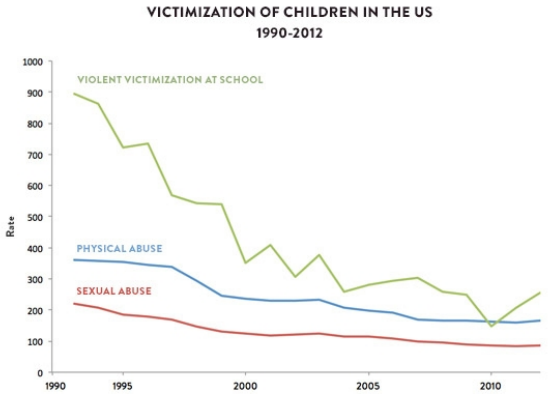

Rates for physical and sexual abuse are per 100,000 children younger than 18. Rates for violent victimization at school are per 10,000 children age 12–17.

Sources: Physical and sexual abuse, National Child Abuse and Neglect Data System, analyzed by David Finkelhor, 2014, in "Trends in Child Welfare," Presentation at the Carsey Institute Policy Series, March 31, 2014. Victimization at school, Bureau of Justice Statistics (bjs.gov), using the National Crime Victimization Survey Victimization Analysis Tool.

Pinker & Mack (2014)

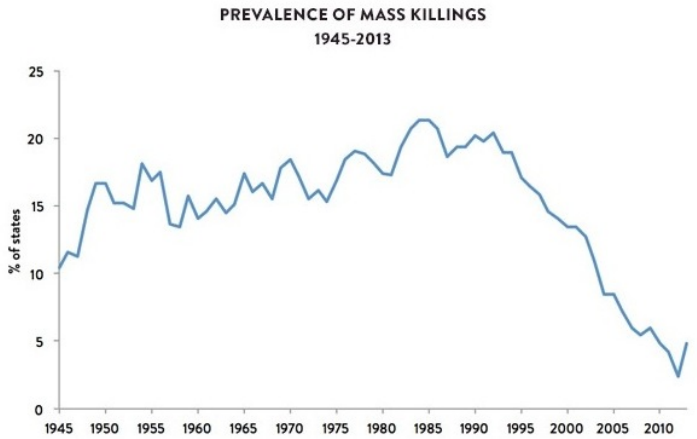

Source: Early Warning Project, https://cpgearlywarning.wordpress.com/; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Feb. 1, 2008). Available at SSRN: http://ssrn.com/abstract=1703426 or http://dx.doi.org/10.2139/ssrn.1703426

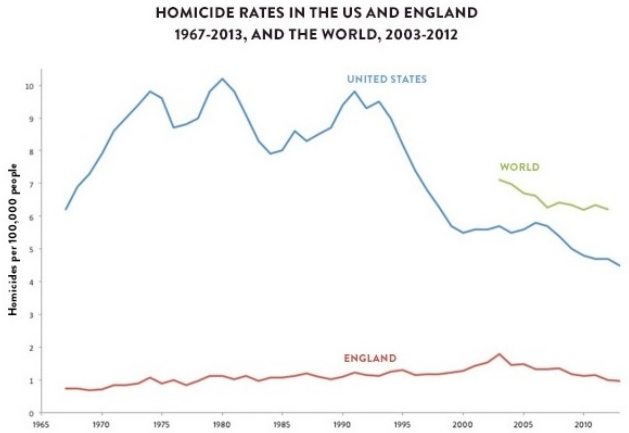

Sources: U.S., FBI Uniform Crime Reports. England (including Wales): U.K. Office for National Statistics. World: U.N. Office on Drugs and Crime, reported in U.N. Economic and Social Council's "World crime trends and emerging issues and responses in the field of crime prevention and social justice," Feb. 12, 2014, Figure 1. The percentages were converted to homicide rates by setting the 2012 rate at 6.2, the figure reported in the UNODC Global Study on Homicide 2013, Page 12.

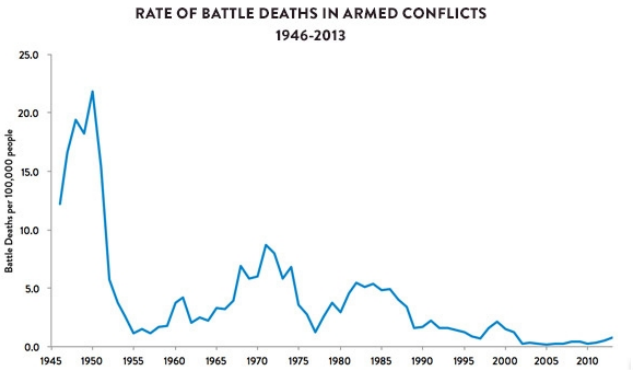

Source: Uppsala Conflict Data Program Battle-Related Deaths Dataset v.5-2014, 1989-2013, "Best Estimates." Missing estimate for 2013 for the Syrian civil war calculated from a revised UCDP "Low Estimate" and a "High Estimate" from the Peace Research Institute Oslo (both obtained in consultation with Erik Melander of UCDP). The "High Estimate" is for May 2013–April 2014, rather than calendar year 2013, and comes from Page 9 of M. Price, A. Gohdes, & P. Ball, "Updated Statistical Analysis of Documentation of Killings in the Syrian Arab Republic," Human Rights Data Analysis Group, http://www.ohchr.org/Documents/Countries/SY/HRDAGUpdatedReportAug2014.pdf. World population figures from U.S. Census Bureau, http://www.census.gov/population/international/data/worldpop/table_population.php

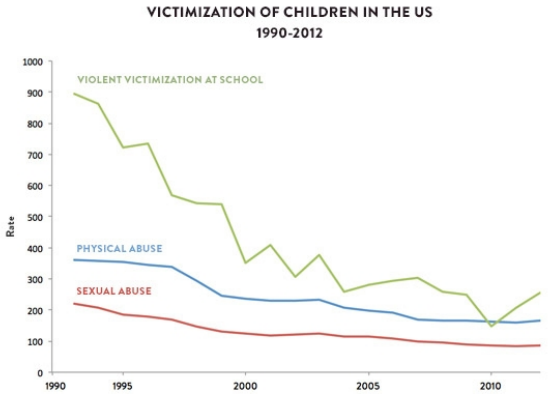

Rates for physical and sexual abuse are per 100,000 children younger than 18. Rates for violent victimization at school are per 10,000 children age 12–17.

Sources: Physical and sexual abuse, National Child Abuse and Neglect Data System, analyzed by David Finkelhor, 2014, in "Trends in Child Welfare," Presentation at the Carsey Institute Policy Series, March 31, 2014. Victimization at school, Bureau of Justice Statistics (bjs.gov), using the National Crime Victimization Survey Victimization Analysis Tool.

World population living in extreme poverty, 1820 to 2015

Our World in Data

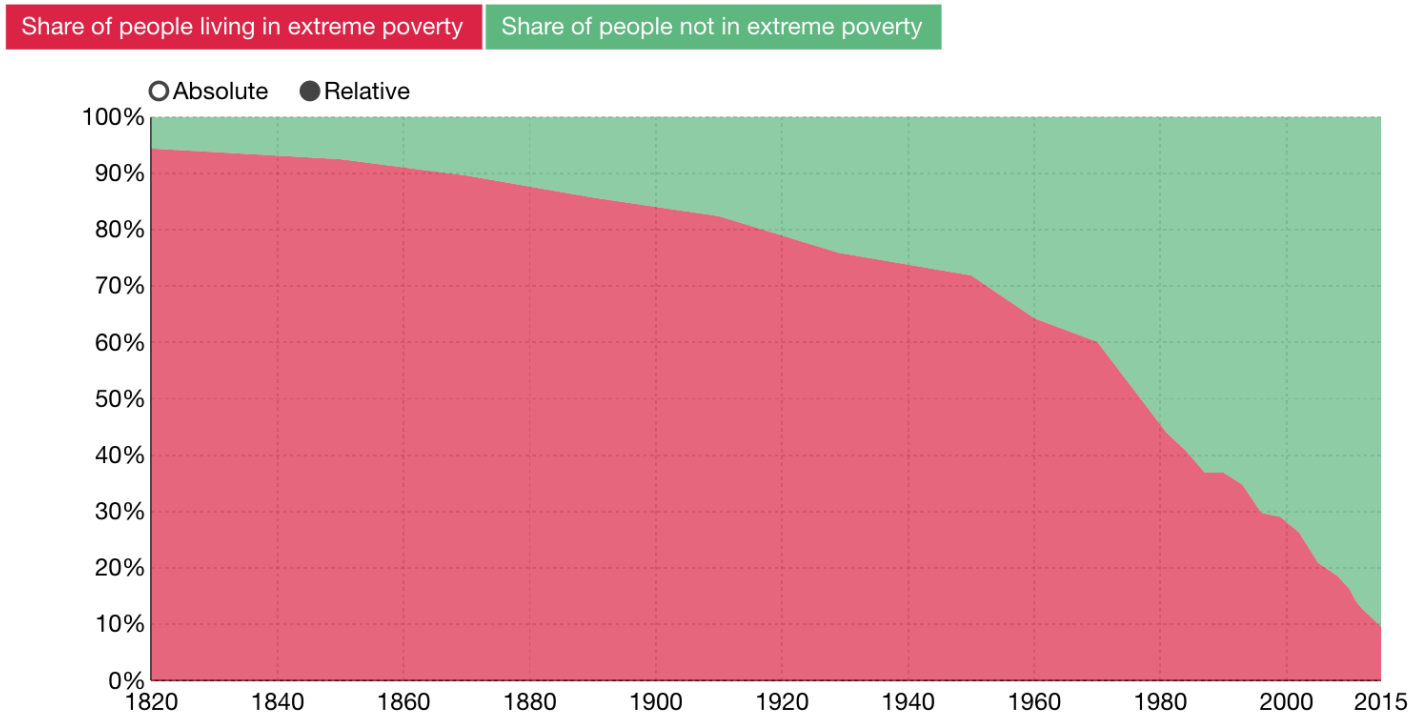

Data source: World Poverty in absolute numbers (Max Roser based on World Bank and Bourguignon and Morrisson (2002))

Literate and illiterate world population, 1800 to 2014

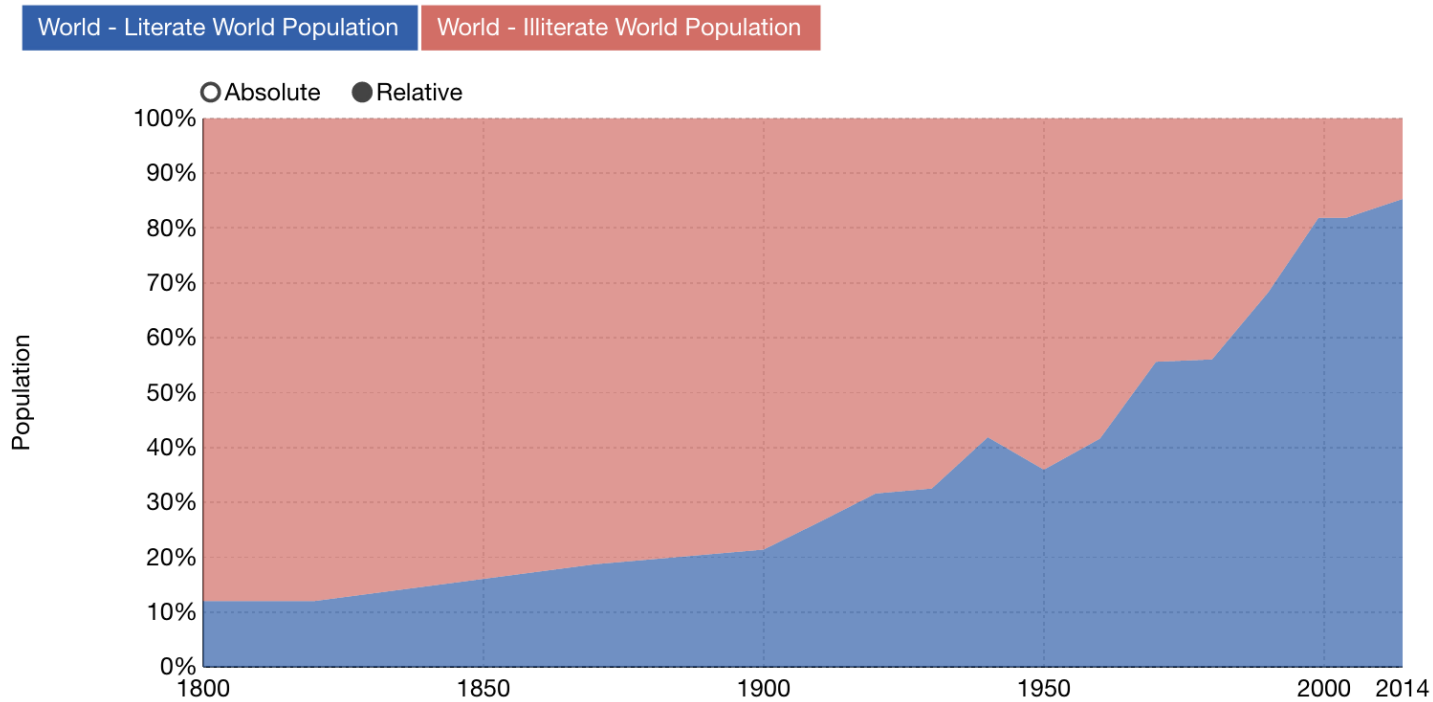

Data source: Literate World Population (Our World In Data based on OECD and UNESCO)

Number of world citizens living under different political regimes

The Polity IV score captures the type of political regime for each country on a range from -10 (full autocracy) to +10 (full democracy). Regimes that fall into the middle of this spectrum are called anocracies.

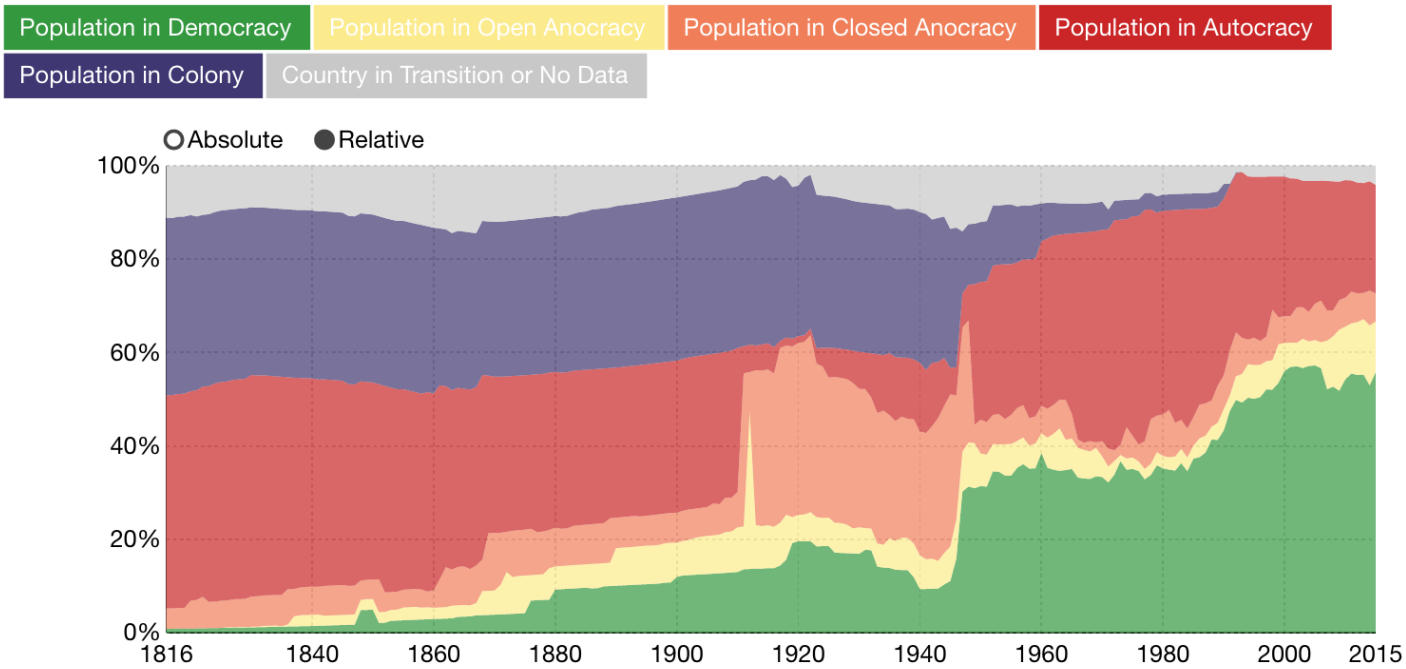

Data source: World Population by Political Regime they live in (by Our World In Data)

Global child mortality, 1800 to 2015

Share of the world population dying and surviving the first 5 years of life.

Our World in Data

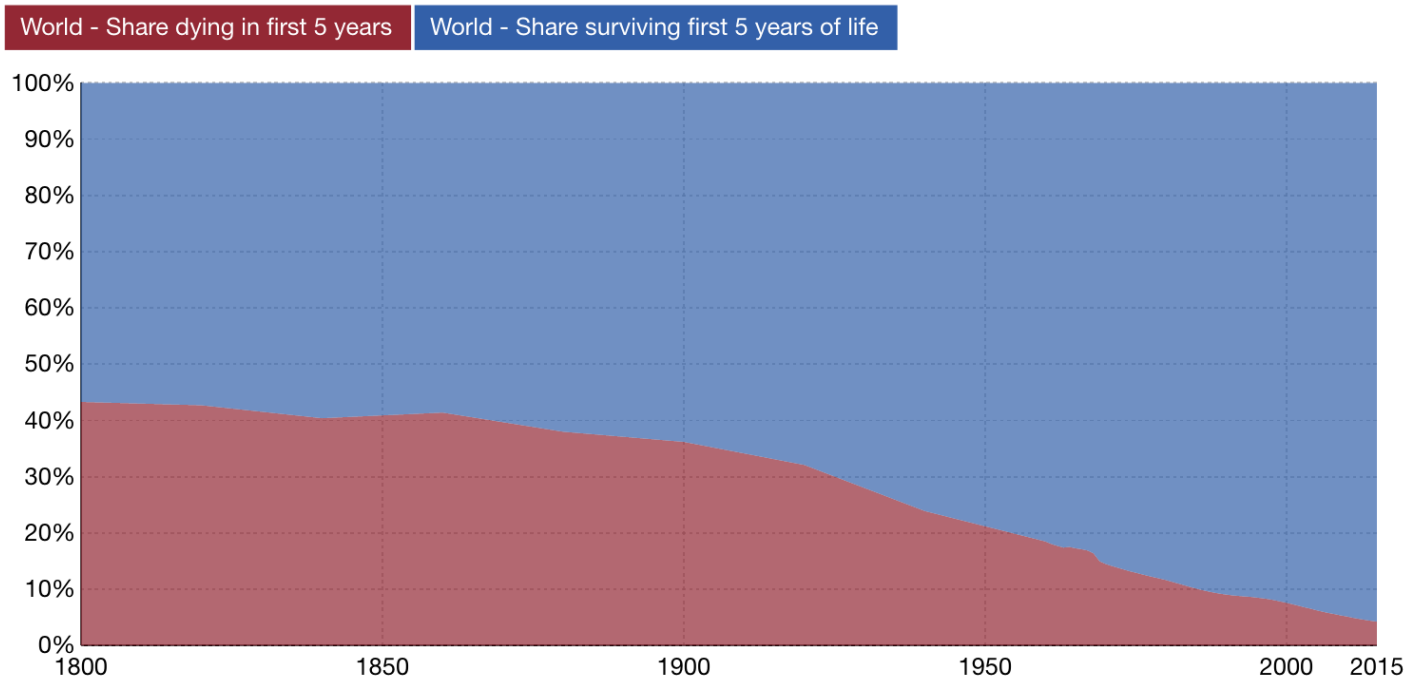

Data source: Global child mortality (since 1800) based on Gapminder and World Bank

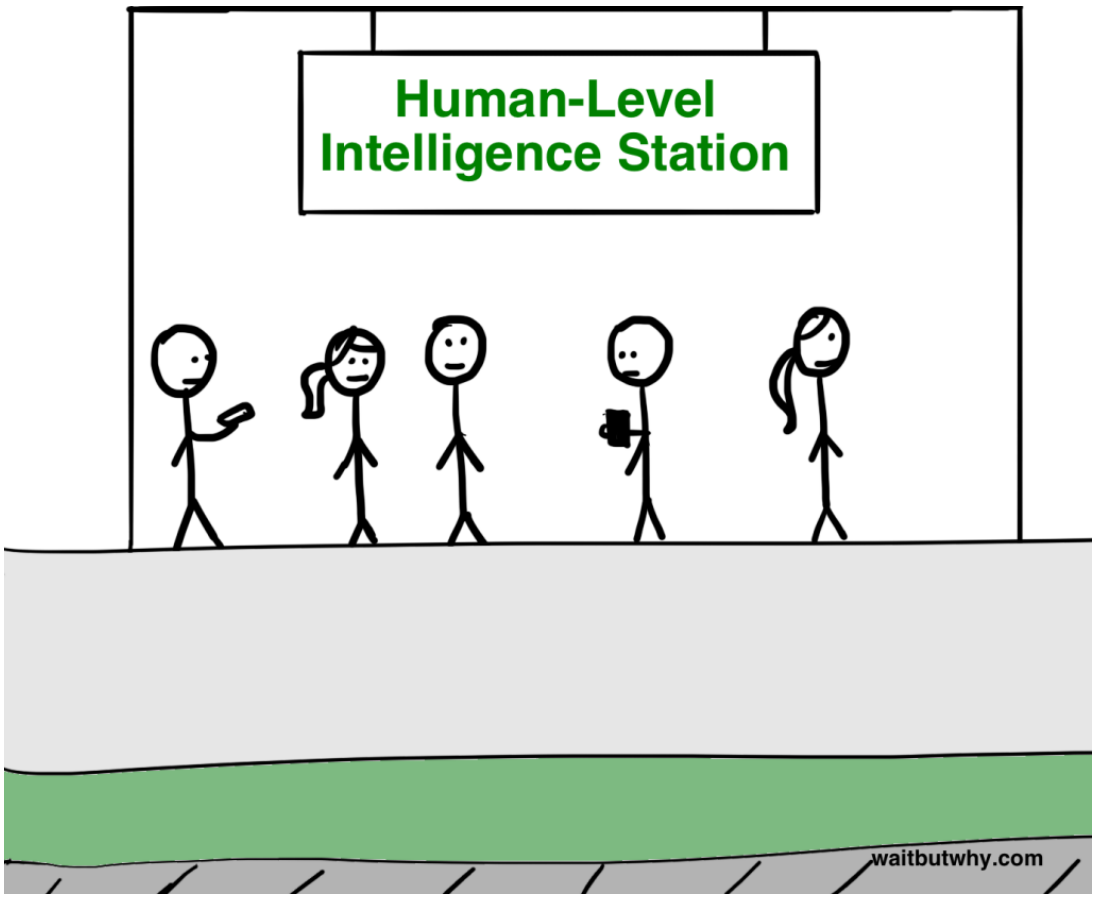

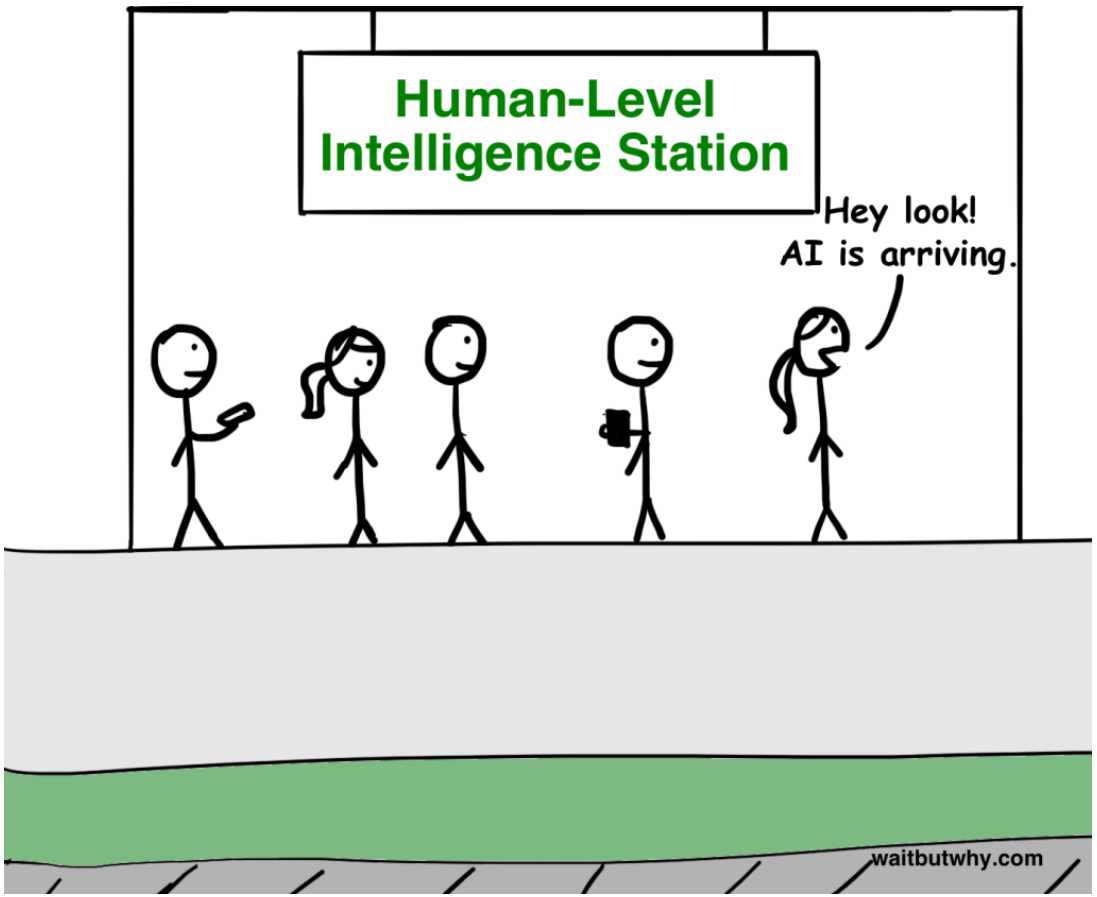
Human-Level
Intelligence Station
Hey look!
AI is arriving.
waitbutwhy.com

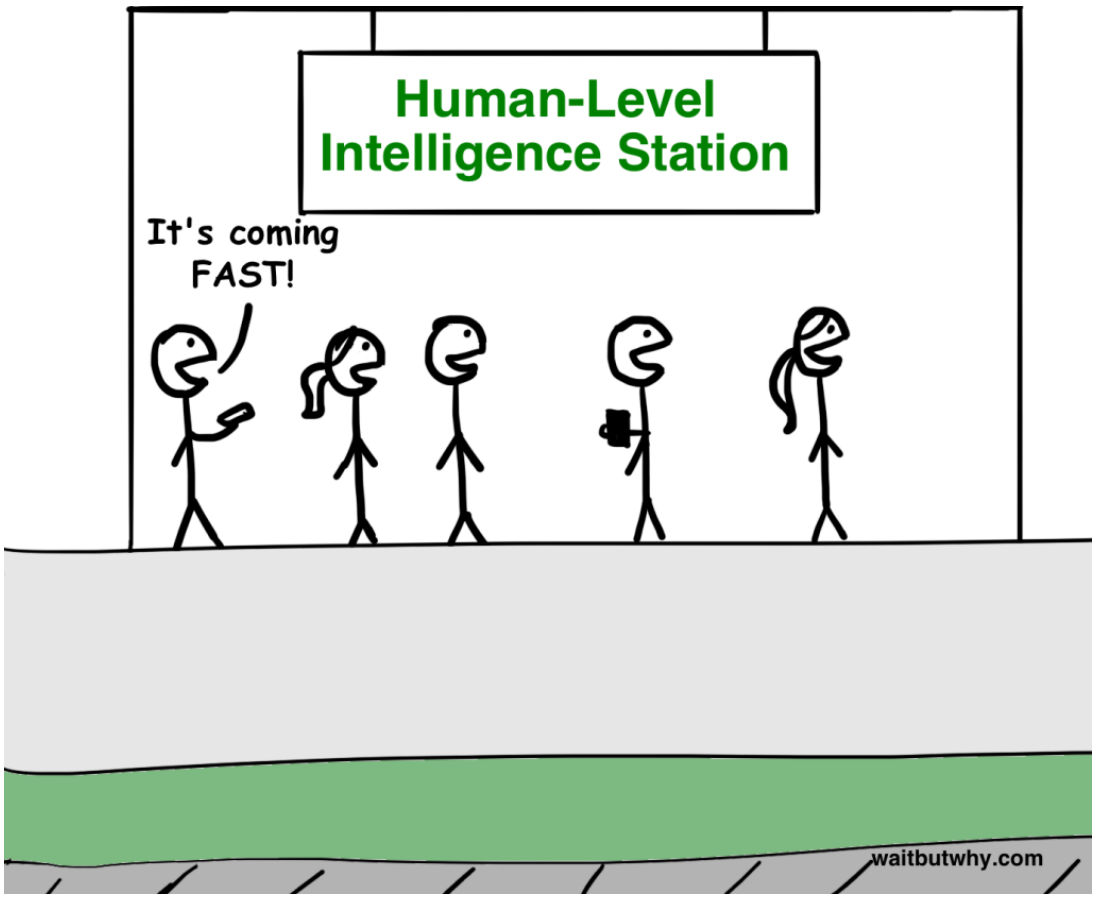
Human-Level
Intelligence Station
It's coming
FAST!
waitbutwhy.com

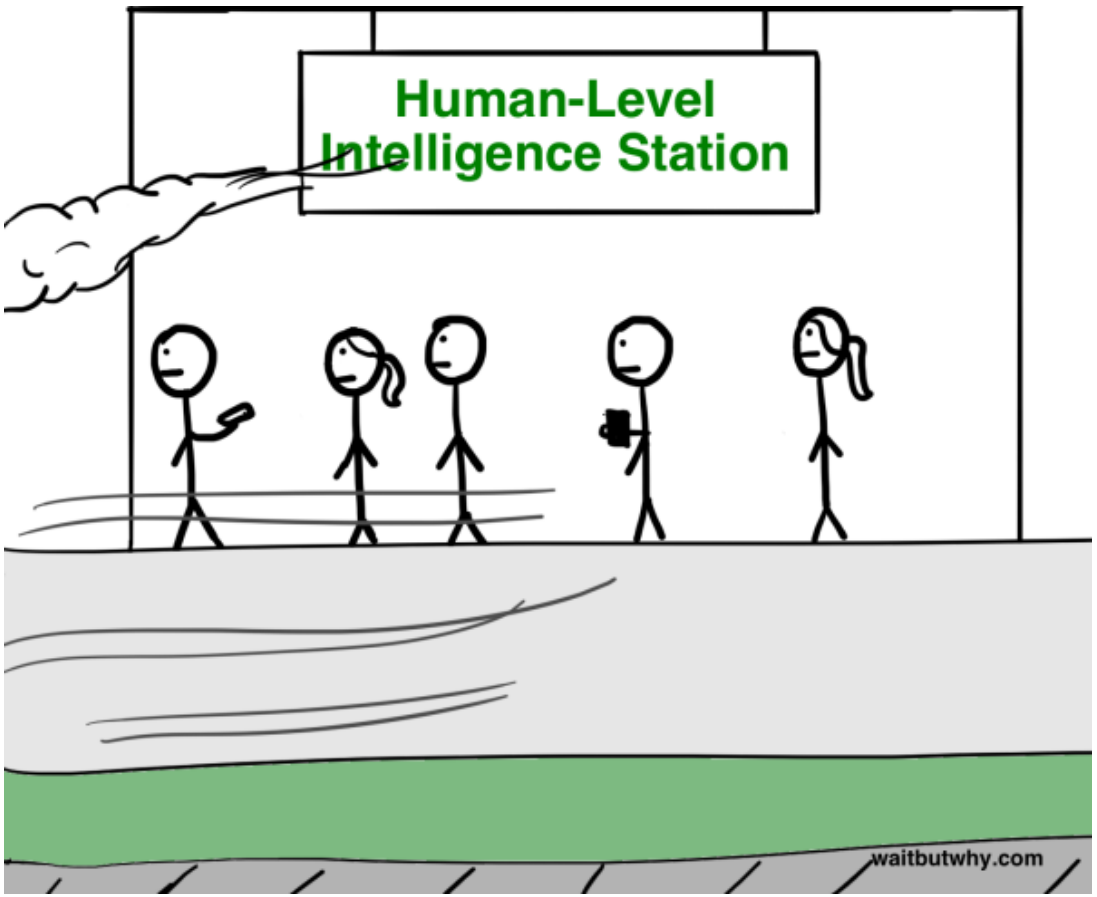
Human-Level
Intelligence Station
waitbutwhy.com

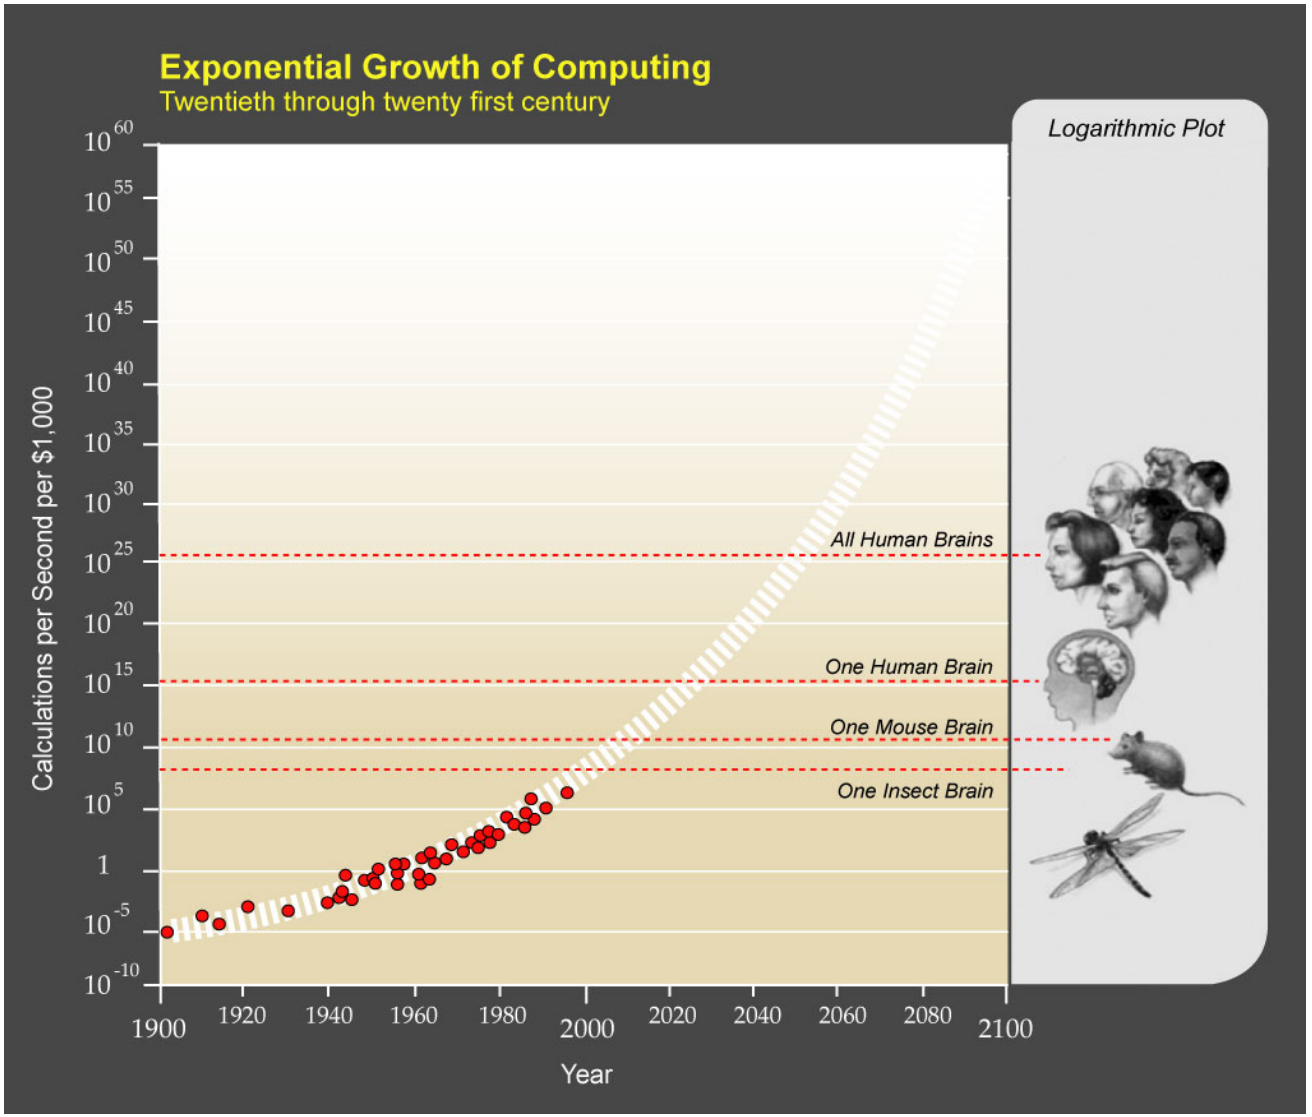
Exponential Growth of Computing
Twentieth through twenty first century
Logarithmic Plot
Calculations per Second per $1,000
10^{60}
10^{55}
10^{50}
10^{45}
10^{40}
10^{35}
10^{30}
10^{25}
10^{20}
10^{15}
10^{10}
10^{5}
1
10^{-5}
10^{-10}
All Human Brains
One Human Brain
One Mouse Brain
One Insect Brain
1900
1920
1940
1960
1980
2000
2020
2040
2060
2080
2100
Year

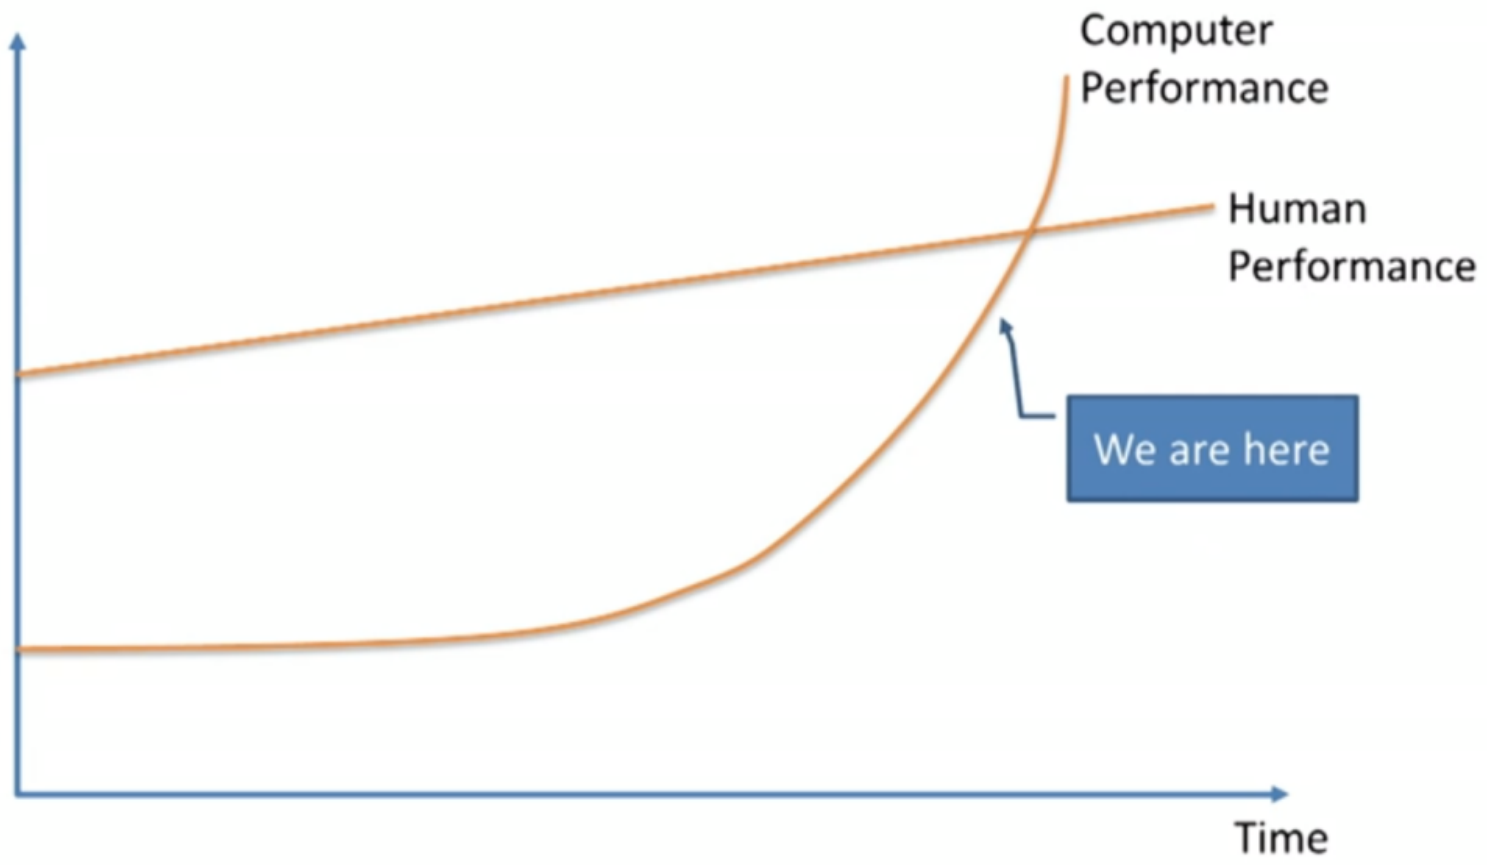
Computer Performance
Human Performance
We are here
Time

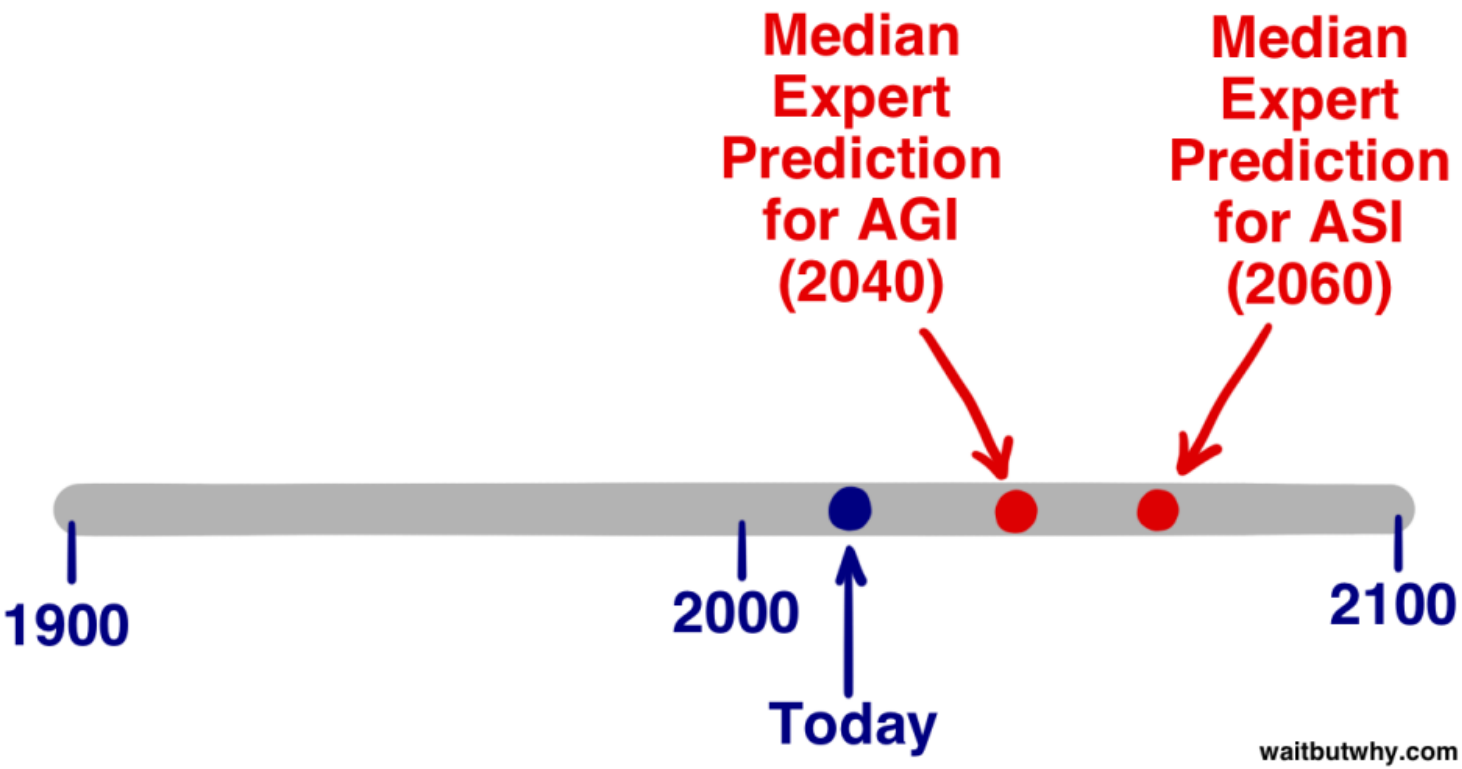
Median Expert Prediction for AGI (2040)
Median Expert Prediction for ASI (2060)
1900
2000
Today
2100
waitbutwhy.com

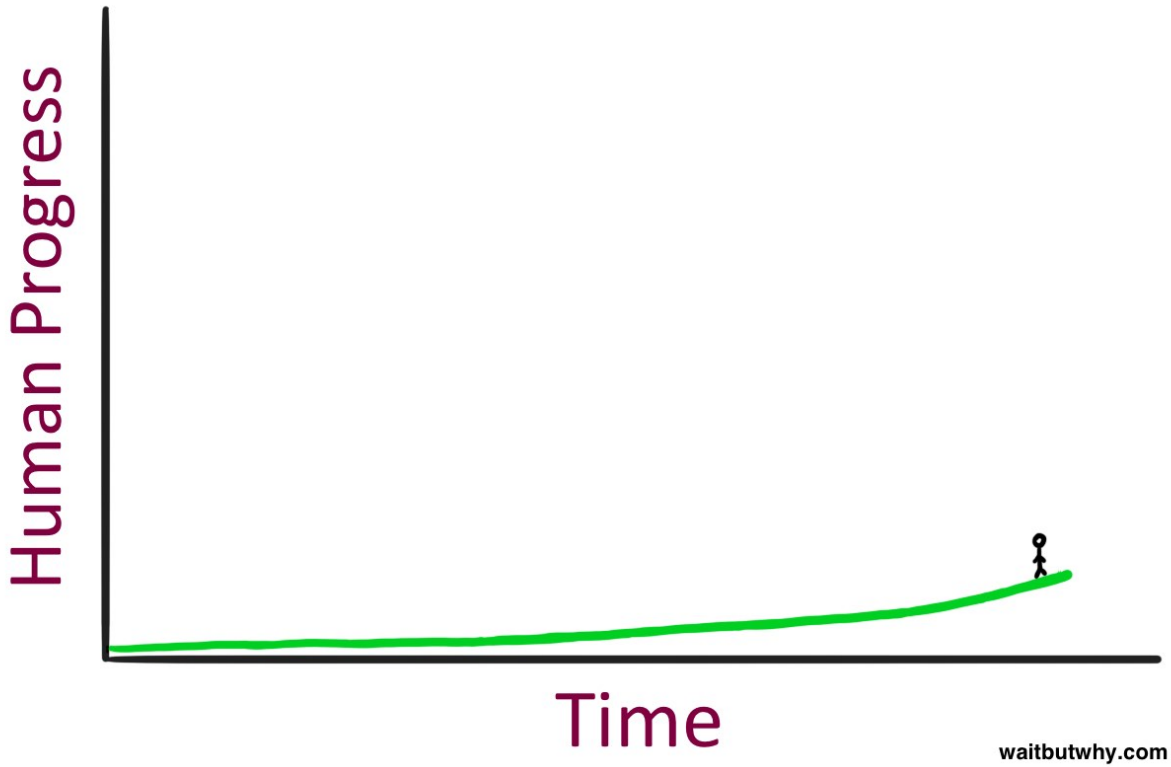
Human Progress
Time
waitbutwhy.com

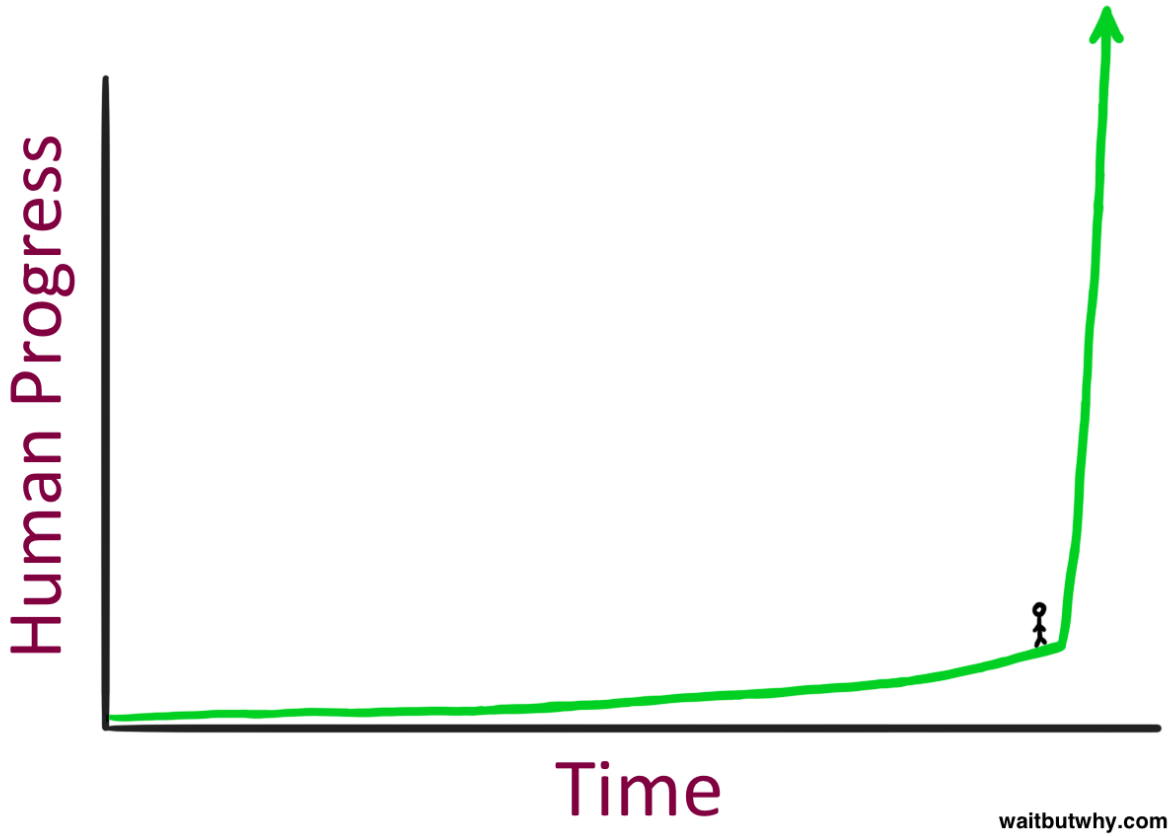
Human Progress
Time
waitbutwhy.com