AI and Radiology

Bartek Papież

Big Data Institute

Rise of Artificial Intelligence (AI)


BBC Sign in Home News Sport Weather iPlayer

NEWS

Home | Cost of Living | War in Ukraine | Coronavirus | Climate | UK | World | Business | Politics | Tech | Sci

Health

AI 'outperforms' doctors diagnosing breast cancer

2 January 2020

https://www.bbc.co.uk/news/health-50857759

Rise of Artificial Intelligence (AI)

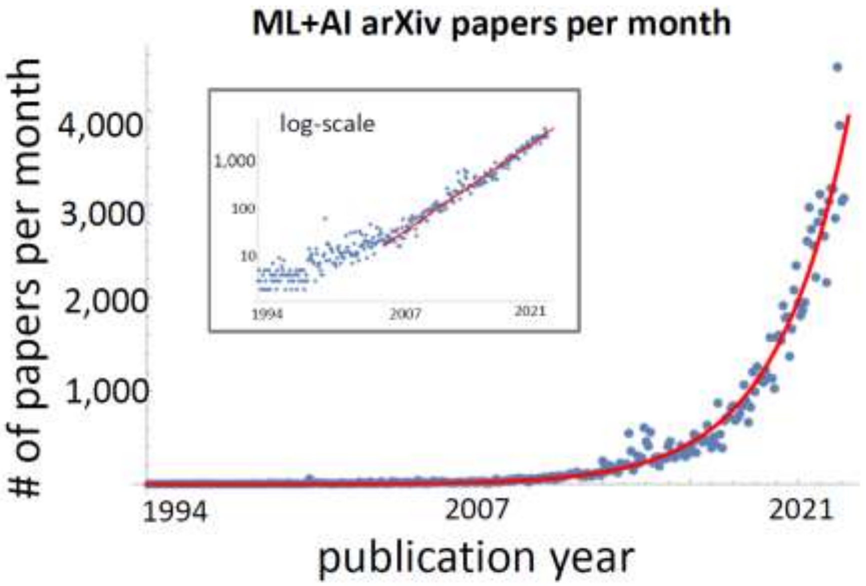

Figure 1. **Number of papers published per months in the arXiv categories of AI grow exponentially.** The doubling rate of papers per months is roughly 23 months, which might lead to problems for publishing in these fields, at some point. The categories are `cs.AI`, `cs.LG`, `cs.NE`, and `stat.ML`.

Krenn, M. et al. "Predicting the Future of AI with AI: High-quality link prediction in an exponentially growing knowledge network." *arXiv preprint arXiv:2210.00881* (2022).

Word cloud generated from the titles of the all accepted papers for Conference on Medical Image Understanding and Analysis since 1997

Machine Learning vs Deep Learning

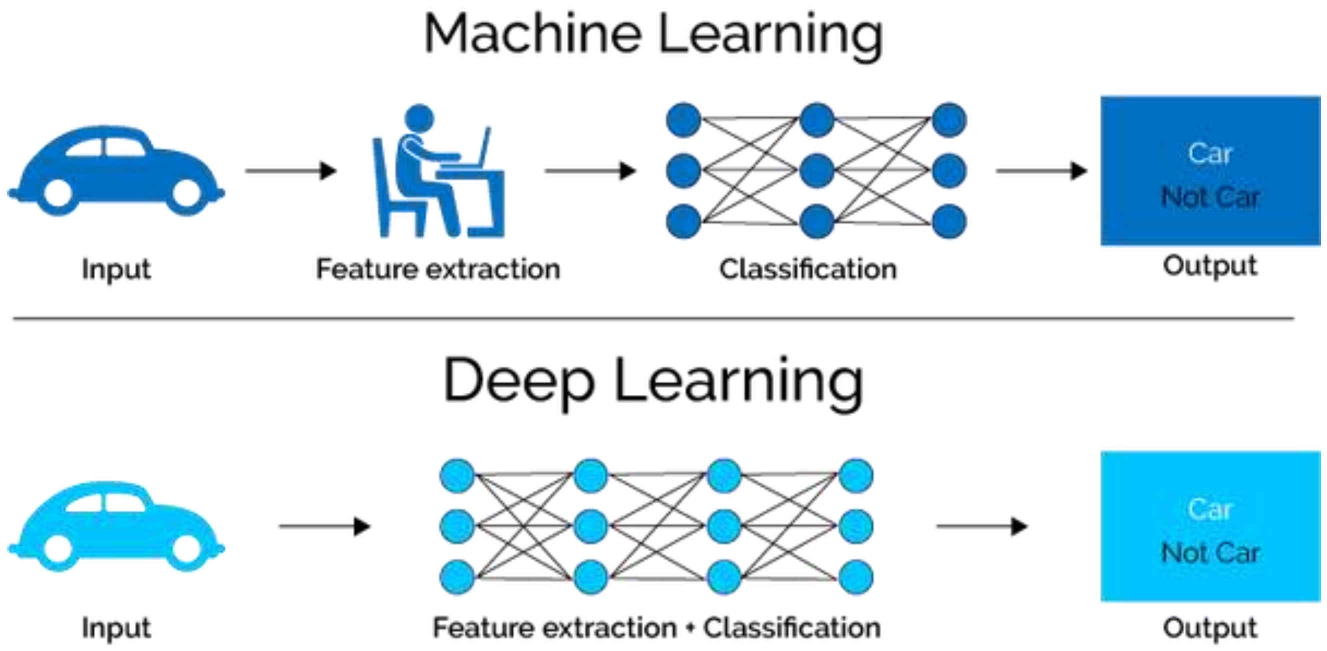

Picture from https://www.quora.com/What-is-the-difference-between-deep-learning-and-usual-machine-learning

Machine Learning vs Deep Learning

- Hand-crafting features

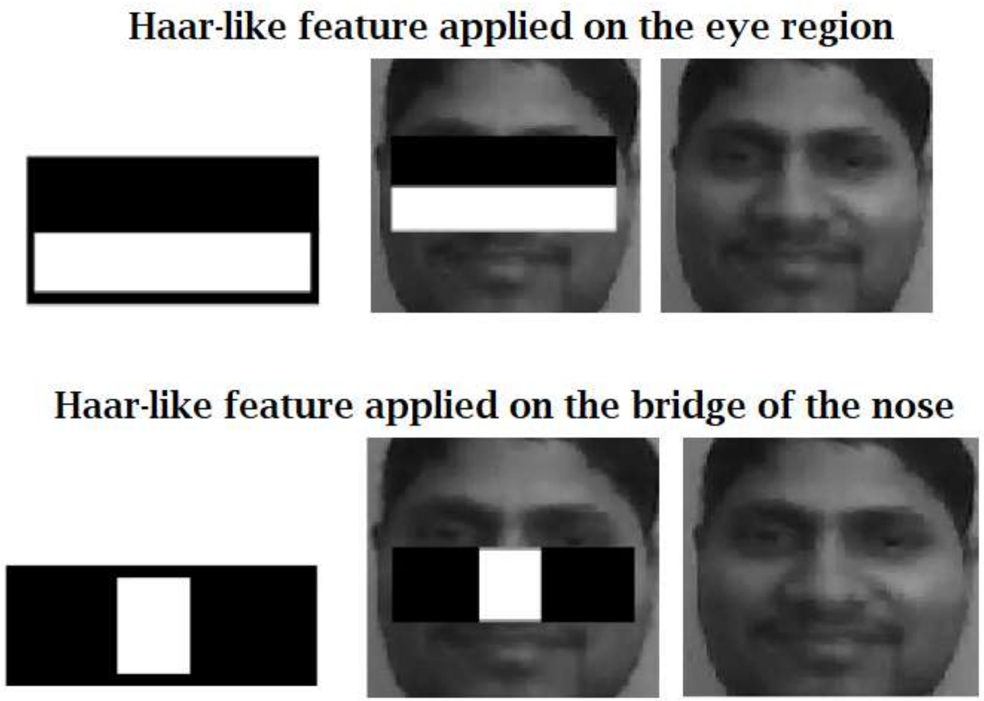

Viola, Paul, and Michael Jones. "Rapid object detection using a boosted cascade of simple features." Proceedings of the 2001 IEEE computer society conference on computer vision and pattern recognition. CVPR 2001. Vol. 1. Ieee, 2001.

- Learning features

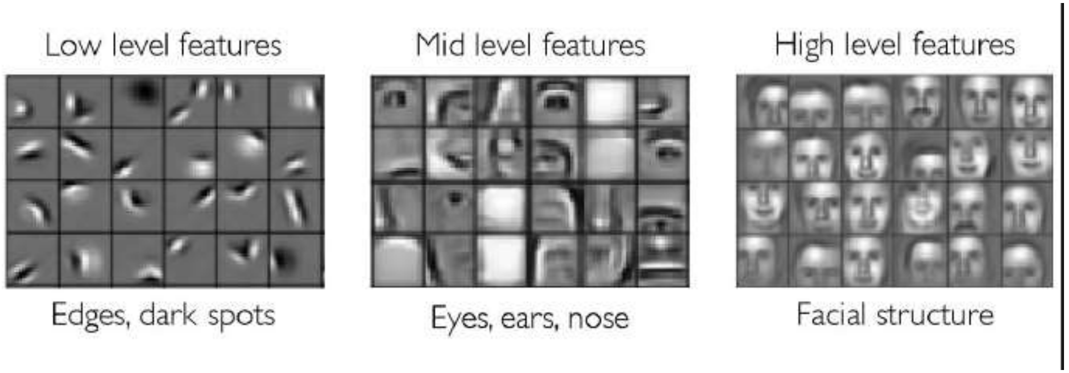

Hubel and Wiesel, 1959

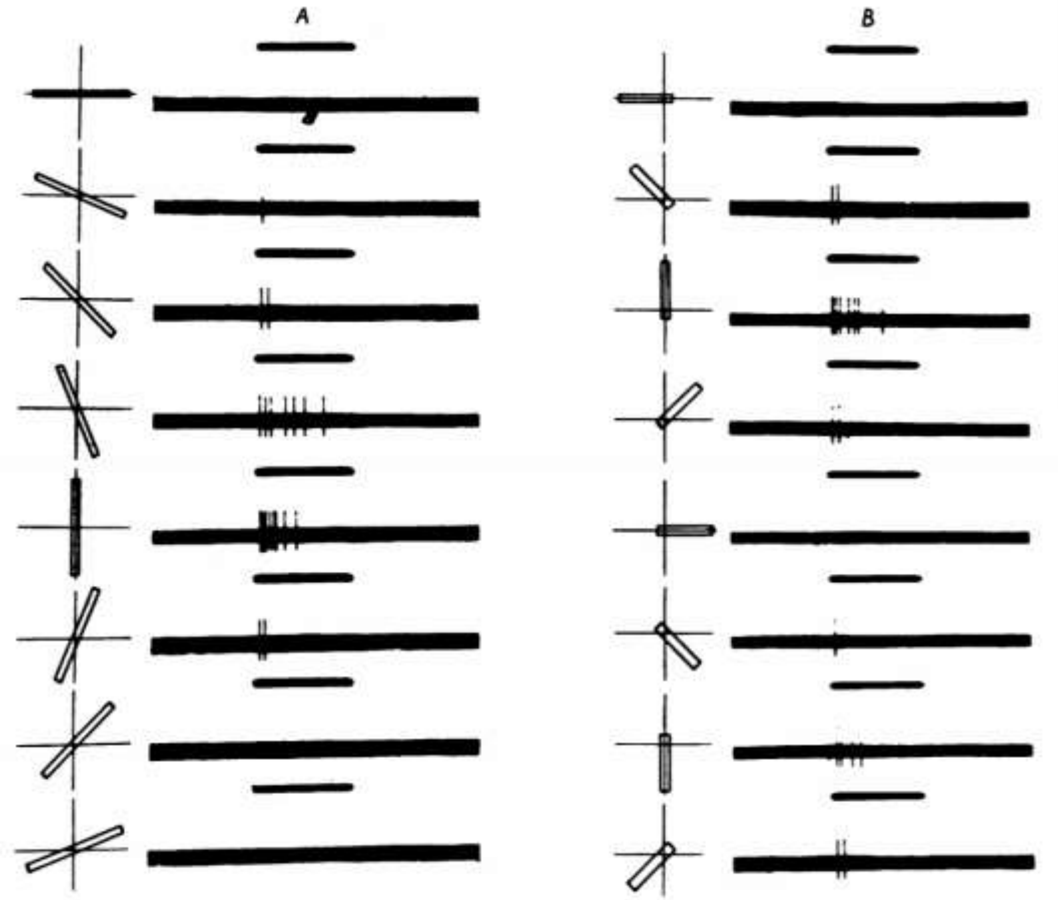

RECEPTIVE FIELDS OF SINGLE NEURONES IN THE CAT'S STRIATE CORTEX

BY D. H. HUBEL* AND T. N. WIESEL*

From the Wilmer Institute, The Johns Hopkins Hospital and University, Baltimore, Maryland, U.S.A.

(*Received* 22 *April* 1959)

Video on Hubel and Wiesel Cat Experiment:

https://www.youtube.com/watch?v=IOHayh06LJ4

Perceptron - biology

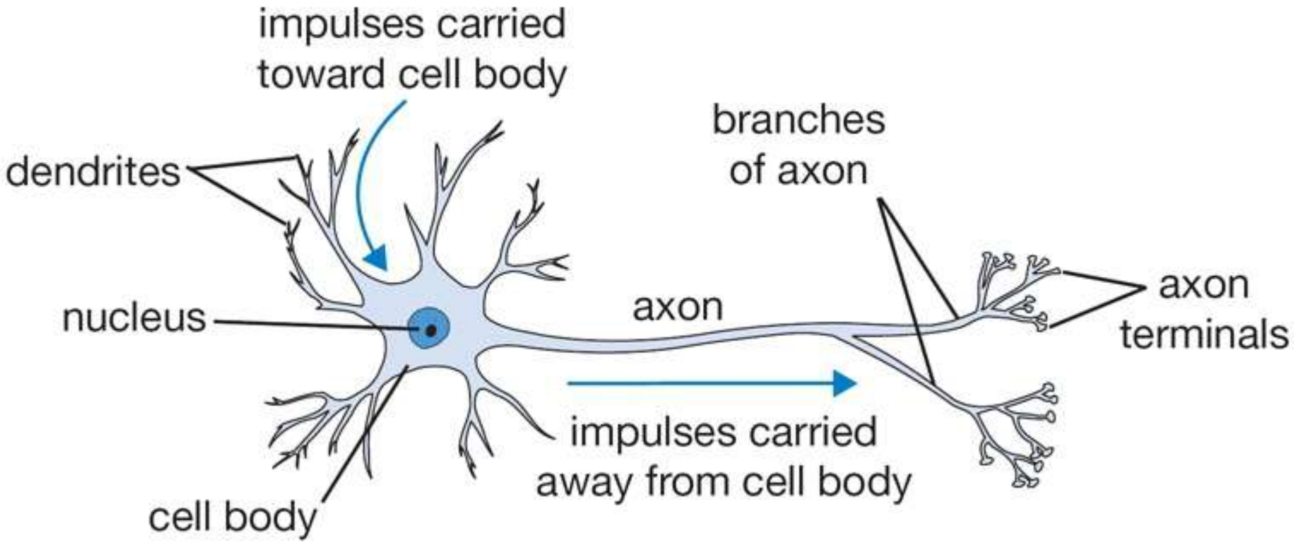

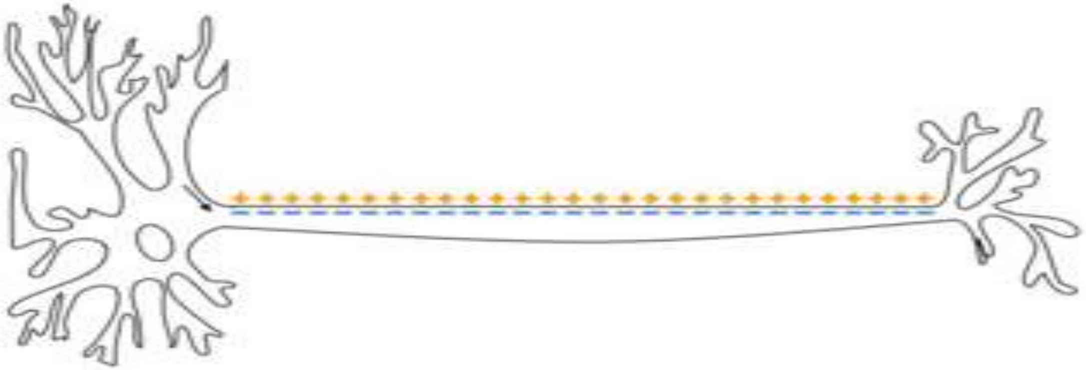

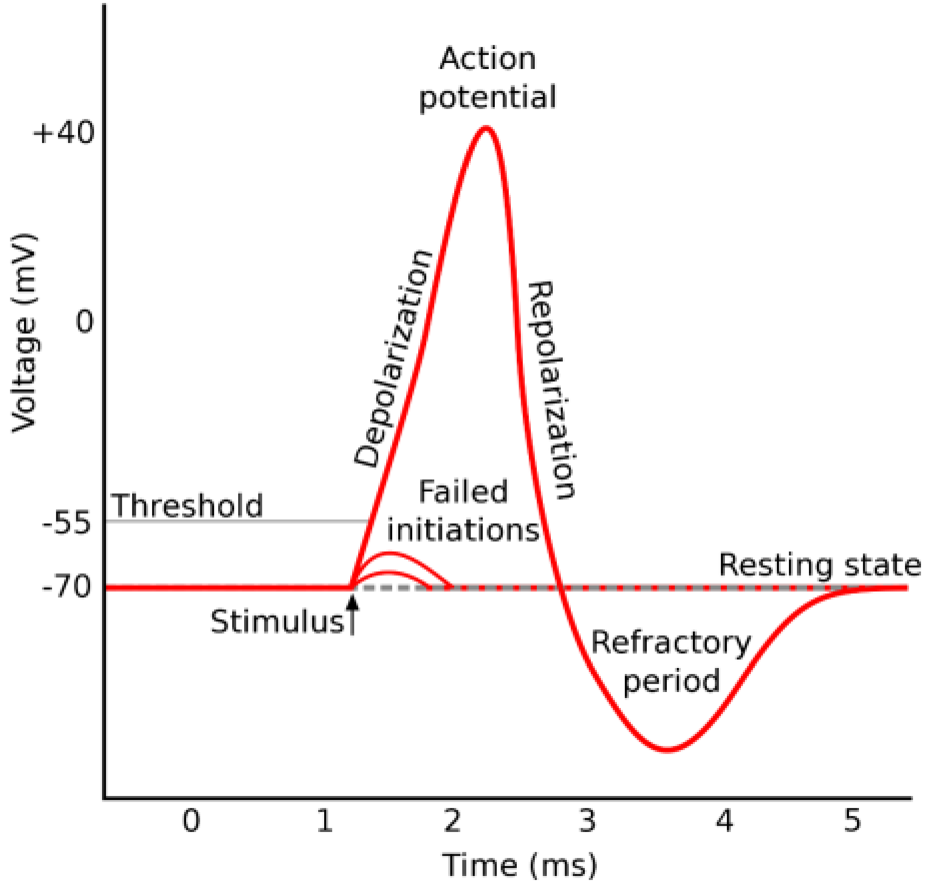

Frank Rosenblatt, 1957: Perceptron

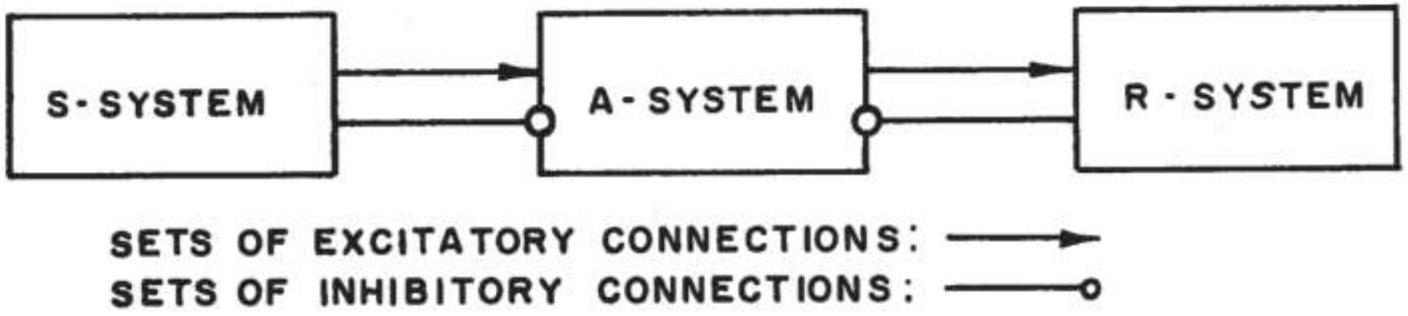

FIGURE I
GENERAL ORGANIZATION OF THE PERCEPTRON

$$f(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{w} \cdot \mathbf{x} + b > 0, \\ 0 & \text{otherwise} \end{cases}$$

Neural Networks - Perceptron

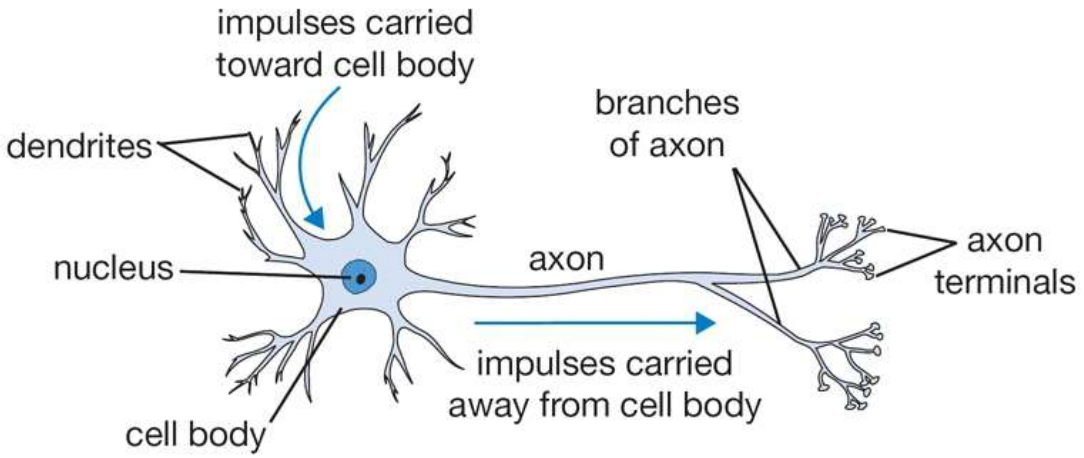

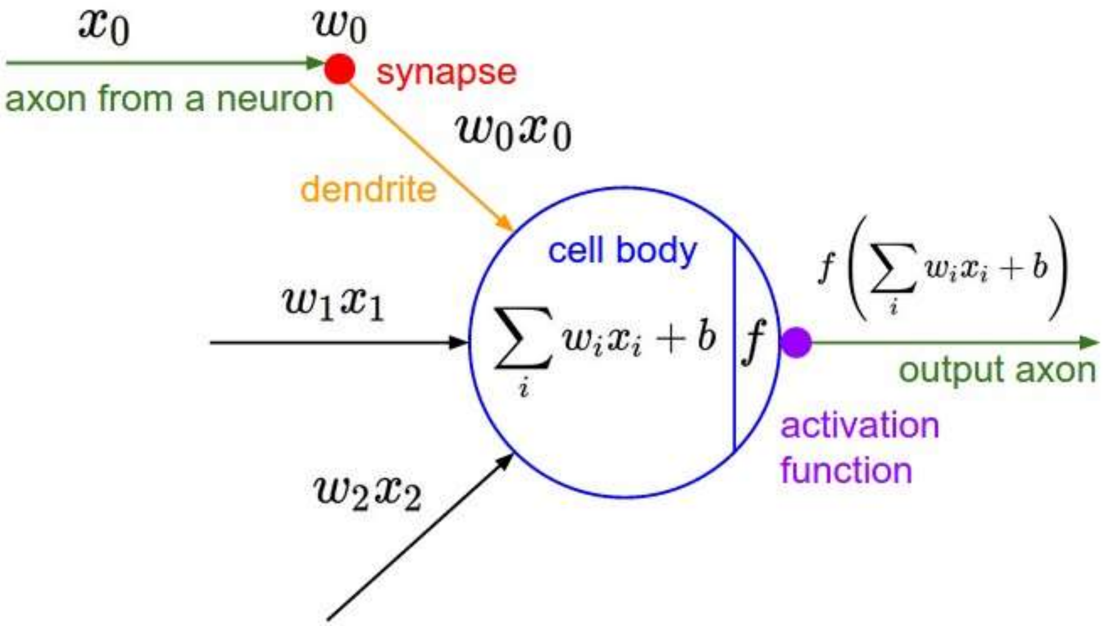

Anatomy of Perceptron

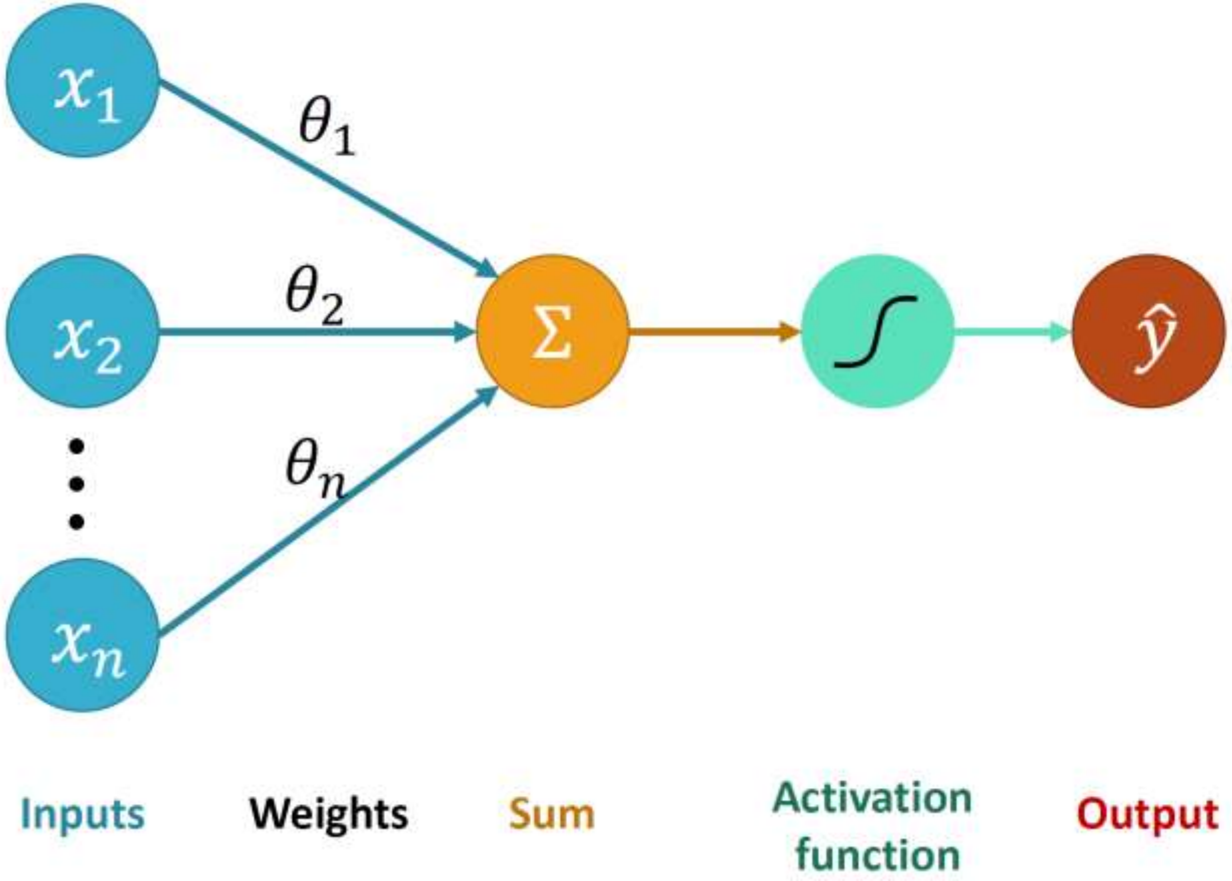

$$\underbrace{\hat{y}}_{\text{Output (prediction)}} = \underbrace{g}_{\text{Non-linear activation function}}\overbrace{\left(\sum_{i=1}^{n} x_i \quad \theta_i\right)}^{\text{Linear combination of inputs}}$$

Anatomy of Perceptron

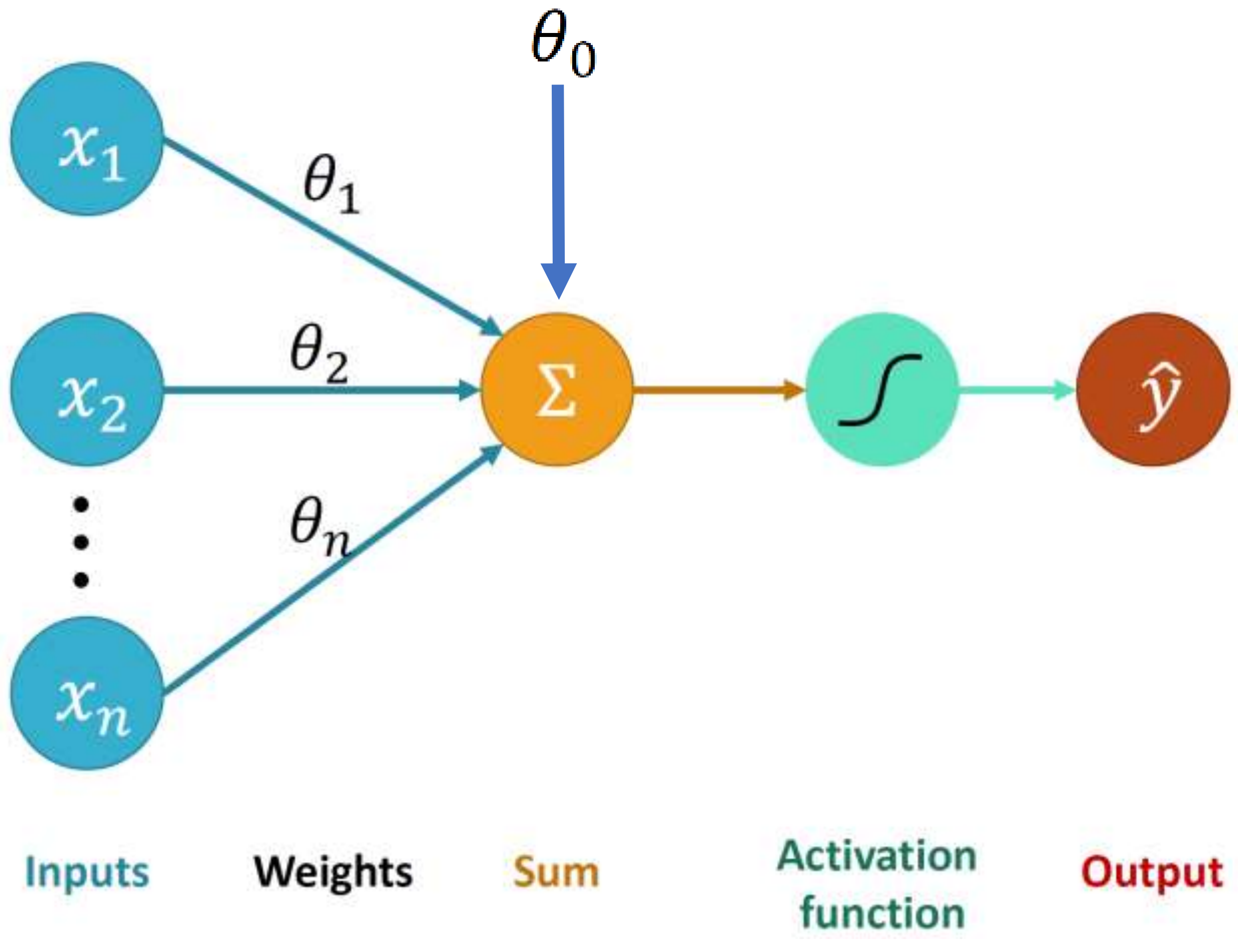

Output (prediction)

$$\hat{y} = g\left(\theta_0 + \overbrace{\sum_{i=1}^{n} x_i \ \theta_i}^{\text{Linear combination of inputs}}\right)$$

Bias

Non-linear activation function

Perceptron - example

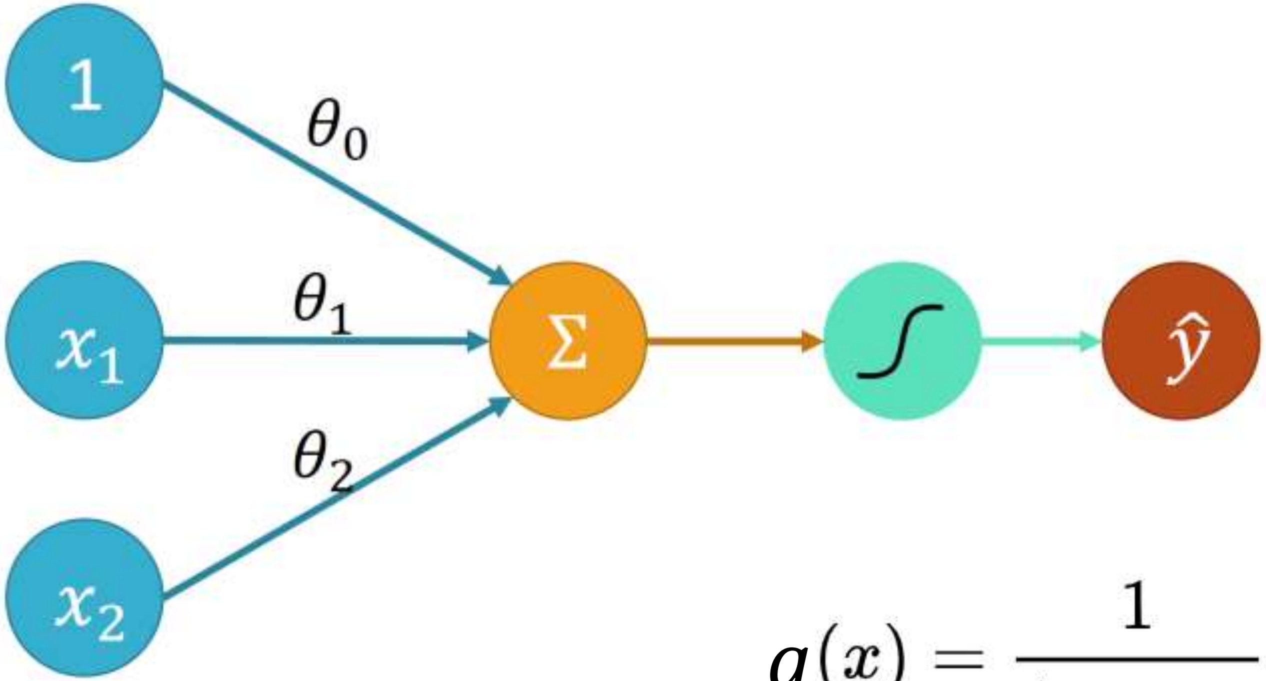

Let $\theta_0 = 1$ and $\boldsymbol{\theta} = \begin{bmatrix} 3 \\ -2 \end{bmatrix}$

$$g(x) = \frac{1}{1 + e^{-x}}$$

Sigmoid function

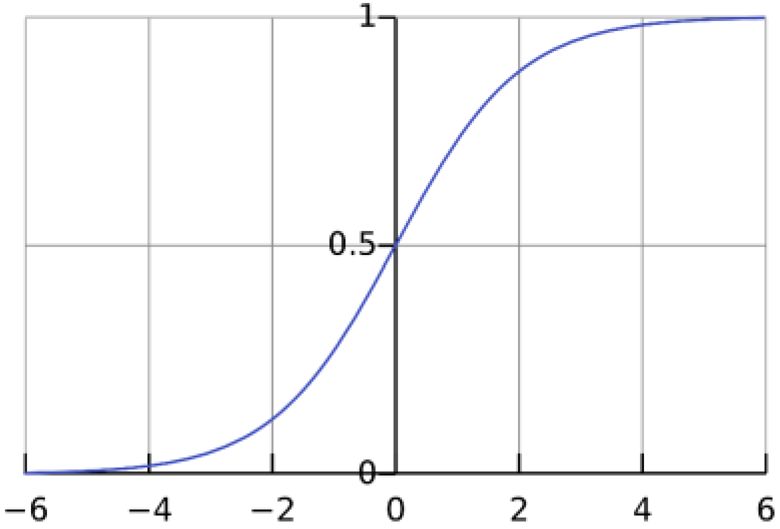

Perceptron - example

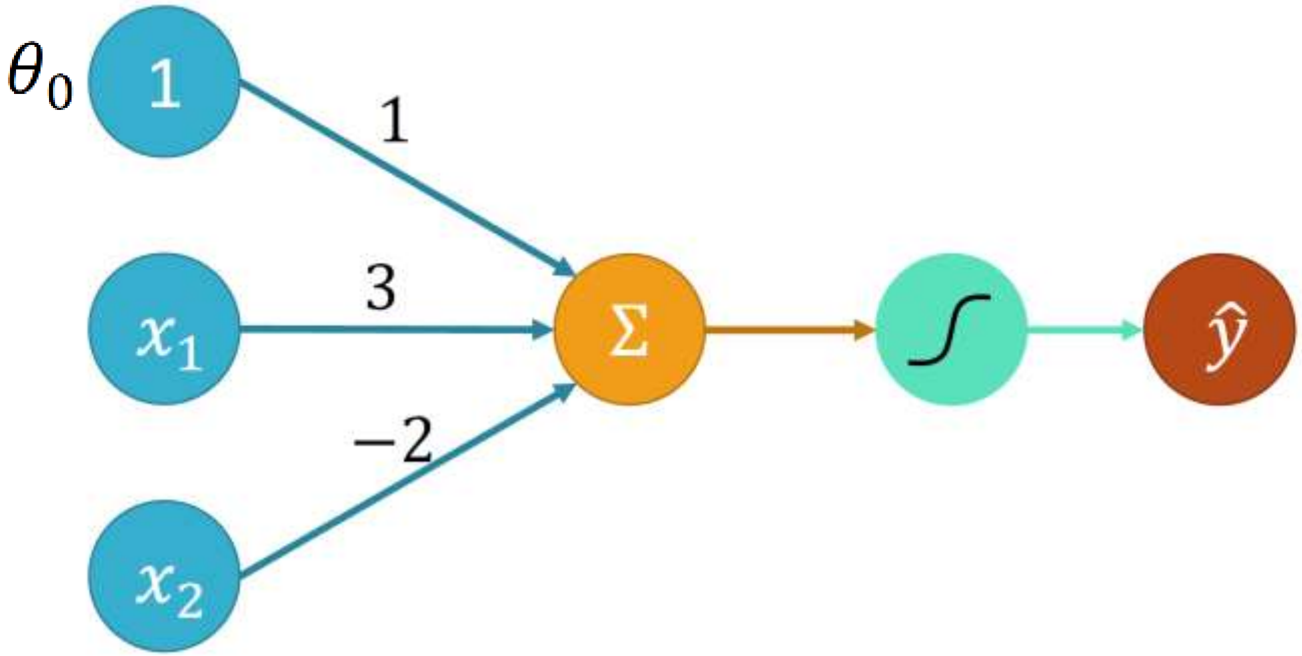

Let $\theta_0 = 1$ and $\boldsymbol{\theta} = \begin{bmatrix} 3 \\ -2 \end{bmatrix}$

Then:

$$\hat{y} = g\ (\theta_0 + \boldsymbol{X}^T\boldsymbol{\theta}\ )$$

$$= g\left(1 + \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}^T \begin{bmatrix} 3 \\ -2 \end{bmatrix}\right)$$

$$\hat{y} = g\ (\underbrace{1 + 3x_1 - 2x_2}_{\text{Just a 2D line!}})$$

Perceptron - example

$$\hat{y} = g\,(1 + 3x_1 - 2x_2)$$

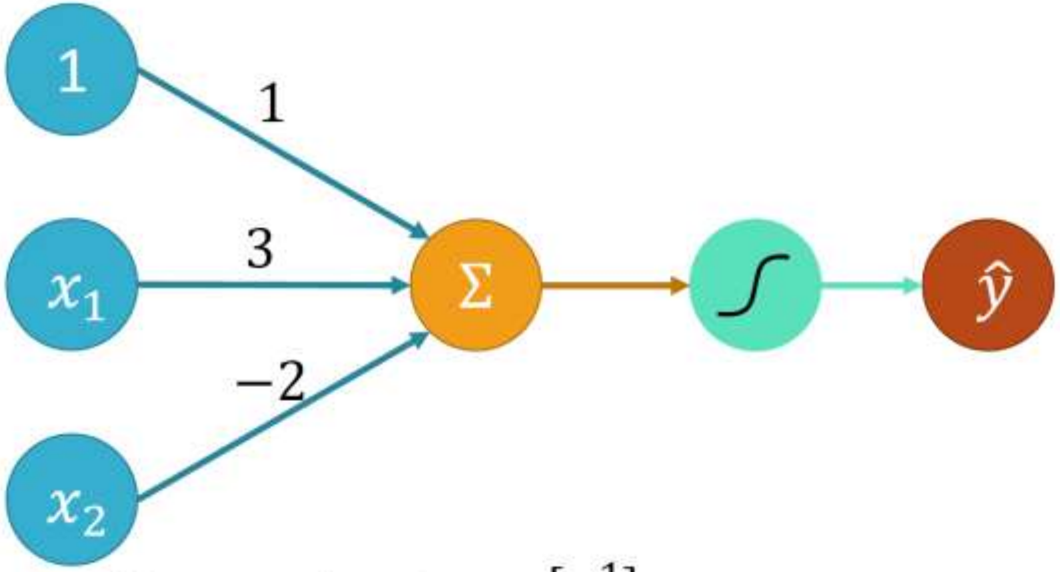

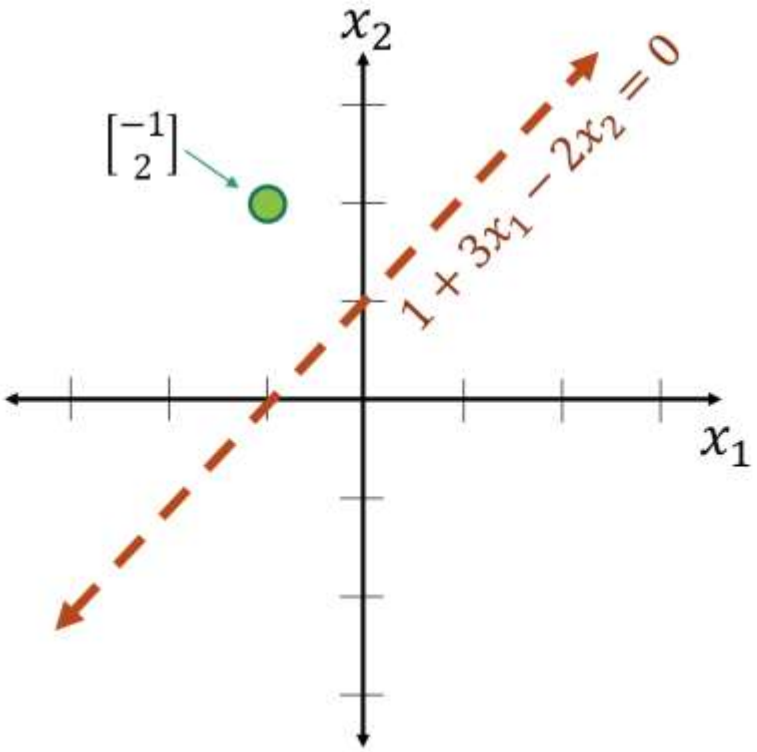

Given an input $\boldsymbol{X} = \begin{bmatrix}-1\\2\end{bmatrix}$

$$\hat{y} = g\,(1 + (3 * -1) - (2 * 2))$$
$$= g(-6) \approx 0.002$$

$$g(x) = \frac{1}{1 + e^{-x}}$$

Perceptron - example

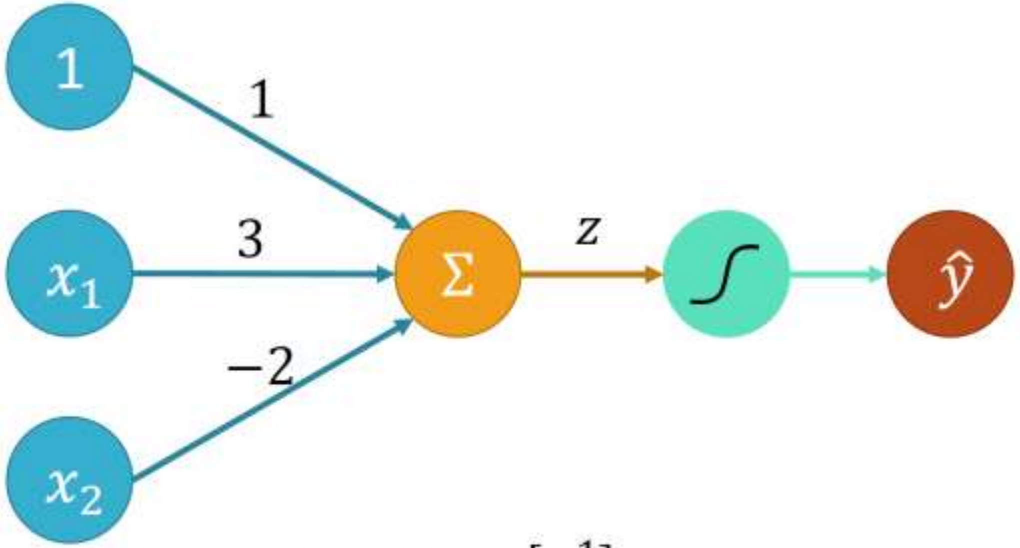

Given an input $\boldsymbol{X} = \begin{bmatrix} -1 \\ 2 \end{bmatrix}$

$$\hat{y} = g\,(1 + (3 * -1) - (2 * 2))$$
$$= g(-6) \approx 0.002$$

$$\hat{y} = g\,(1 + 3x_1 - 2x_2)$$

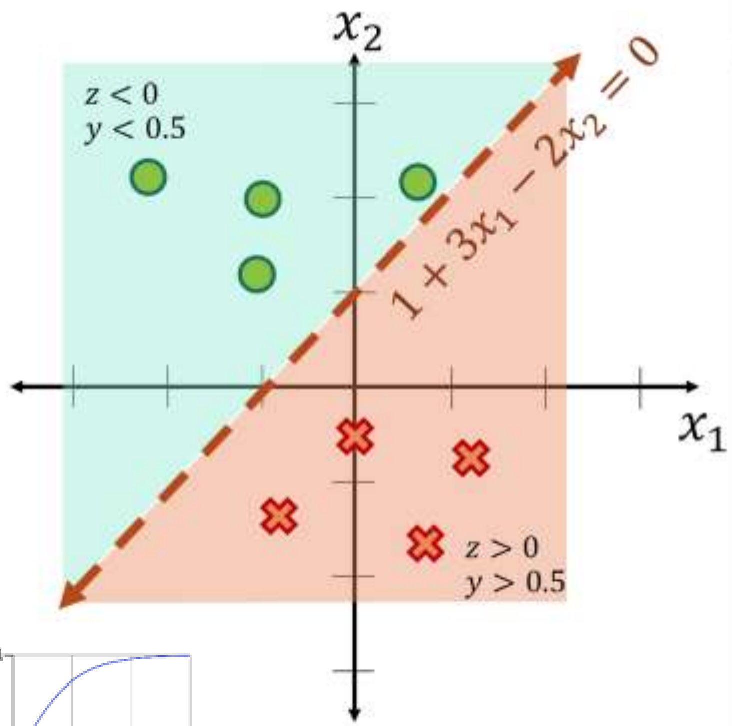

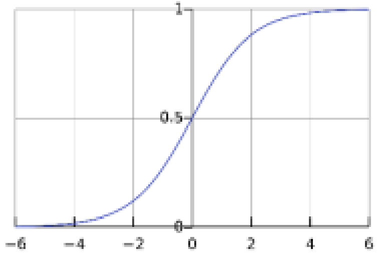

Single output network

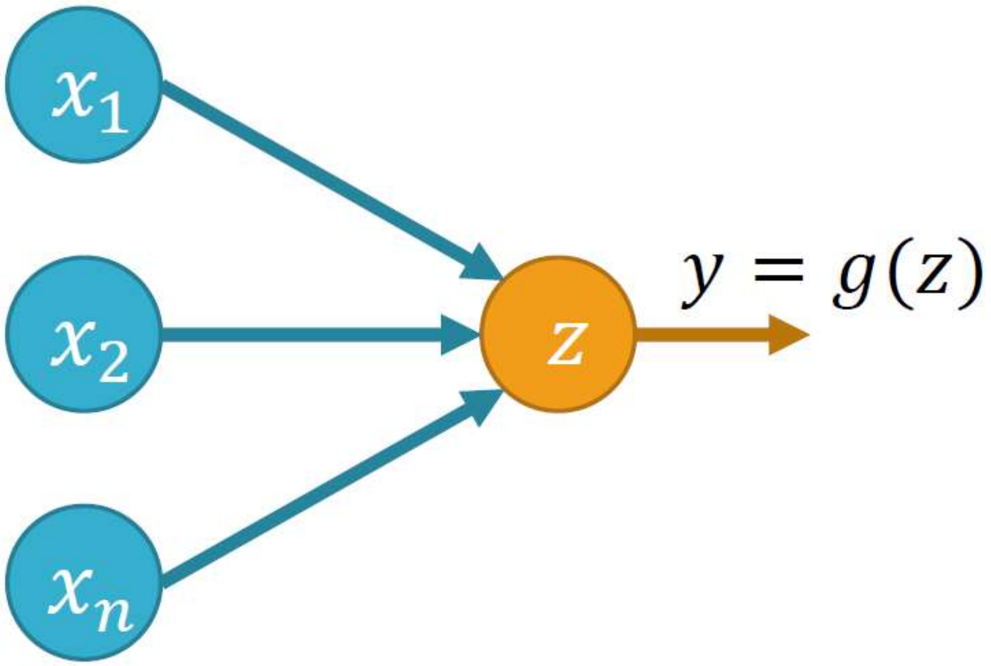

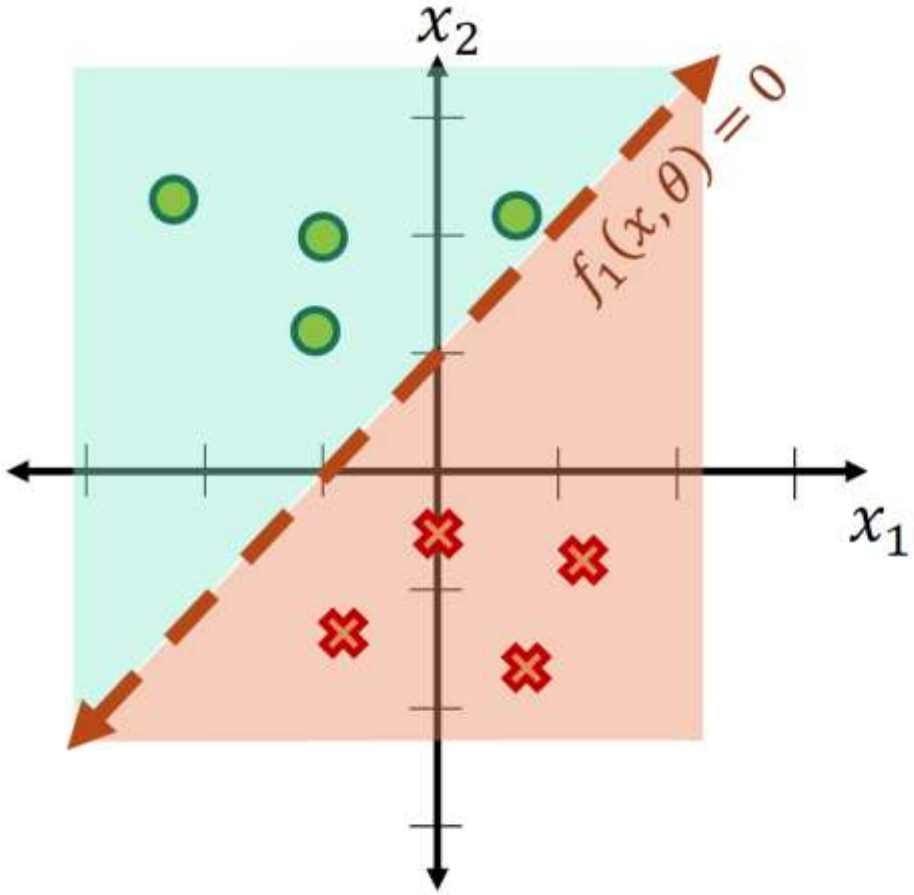

Frank Rosenblatt, 1957: Perceptron

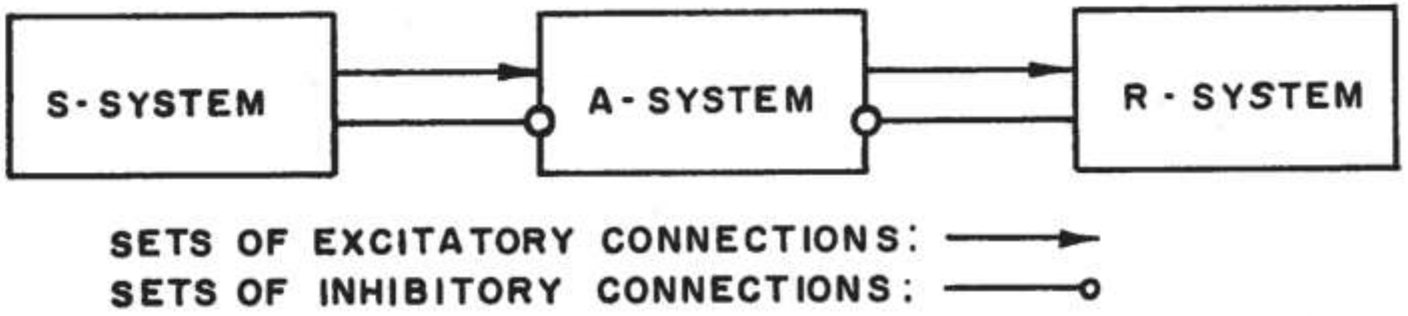

FIGURE I
GENERAL ORGANIZATION OF THE PERCEPTRON

NEW NAVY DEVICE LEARNS BY DOING; Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser

Give this article

July 8, 1958

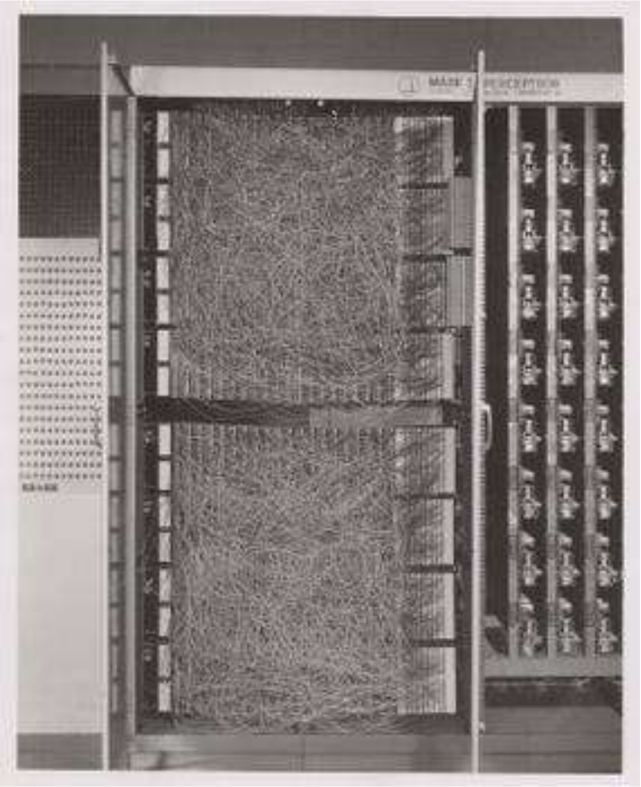

Mark I Perceptron machine, the first implementation of the perceptron algorithm.

AI Winter is coming

- XOR problem

x1	x2	y
0	0	0
0	1	1
1	0	1
1	1	0

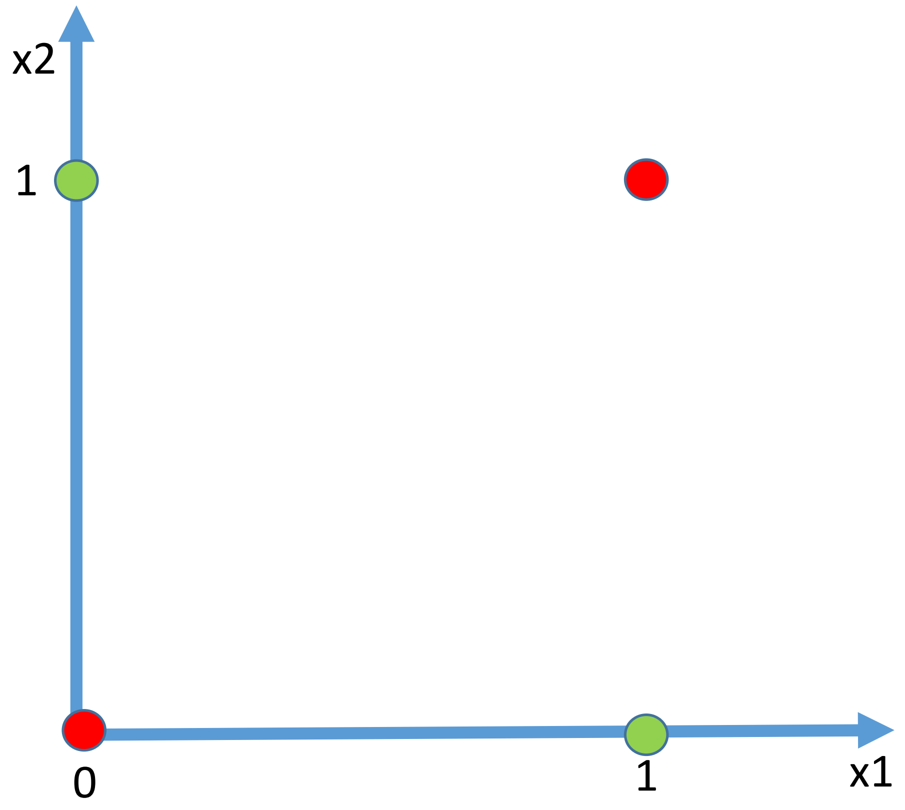

Multiple output network

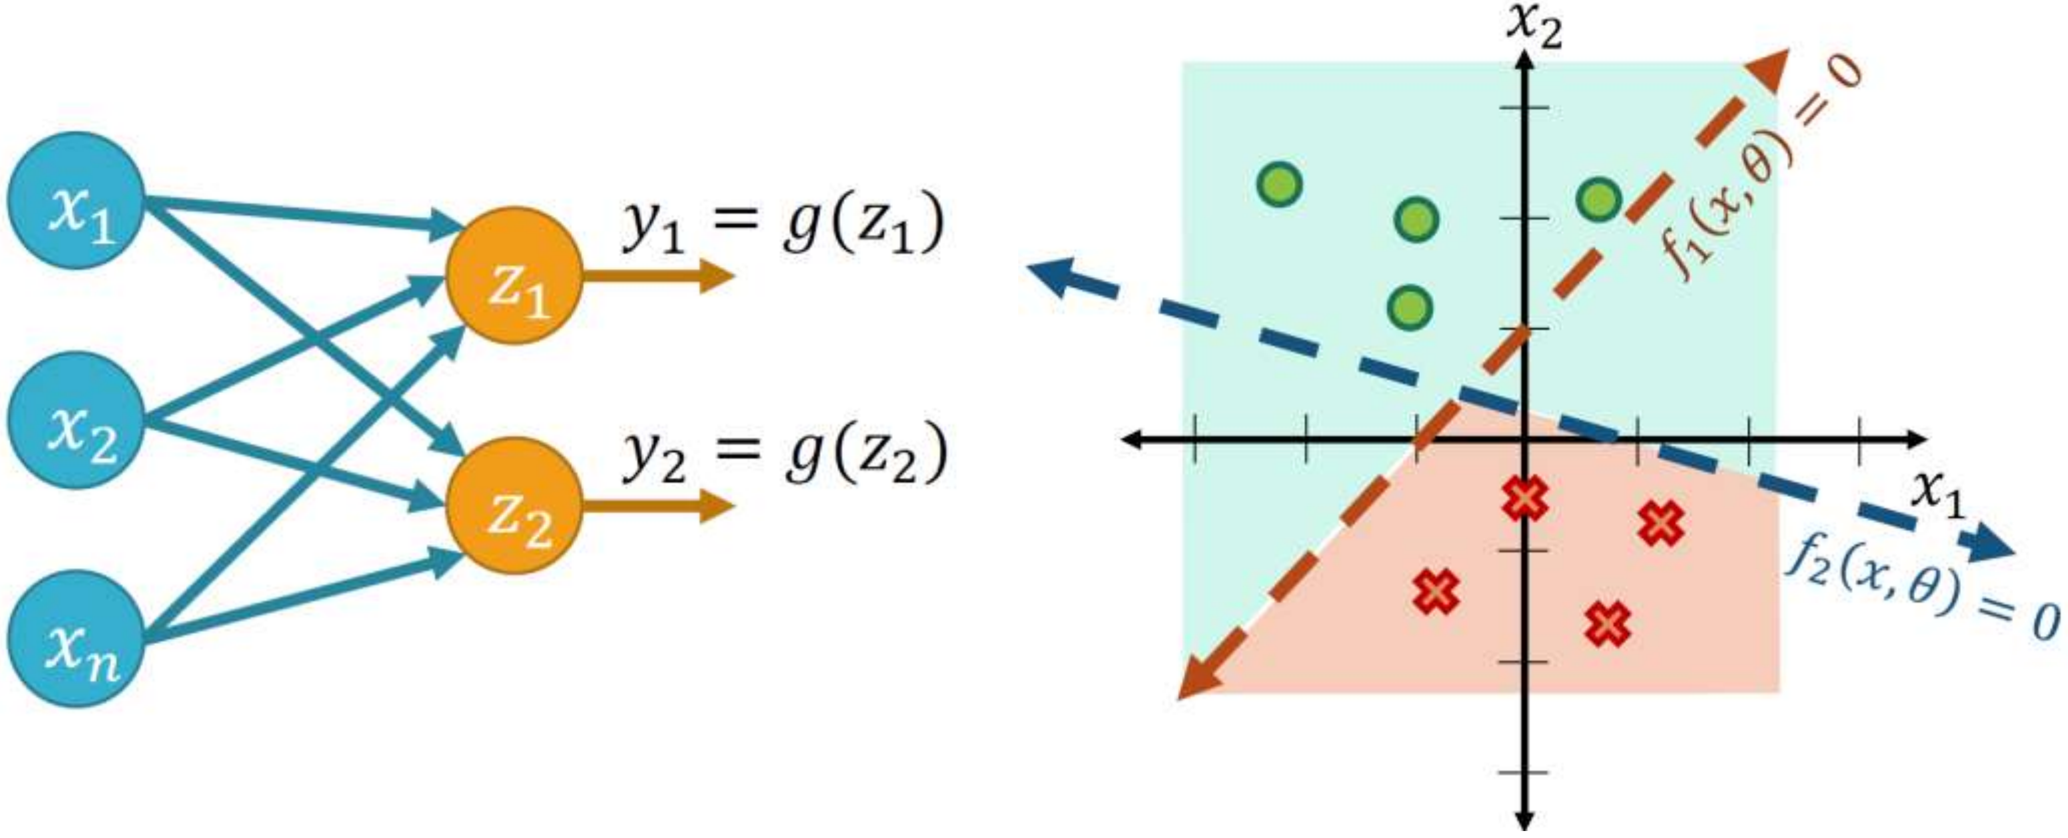

http://playground.tensorflow.org/

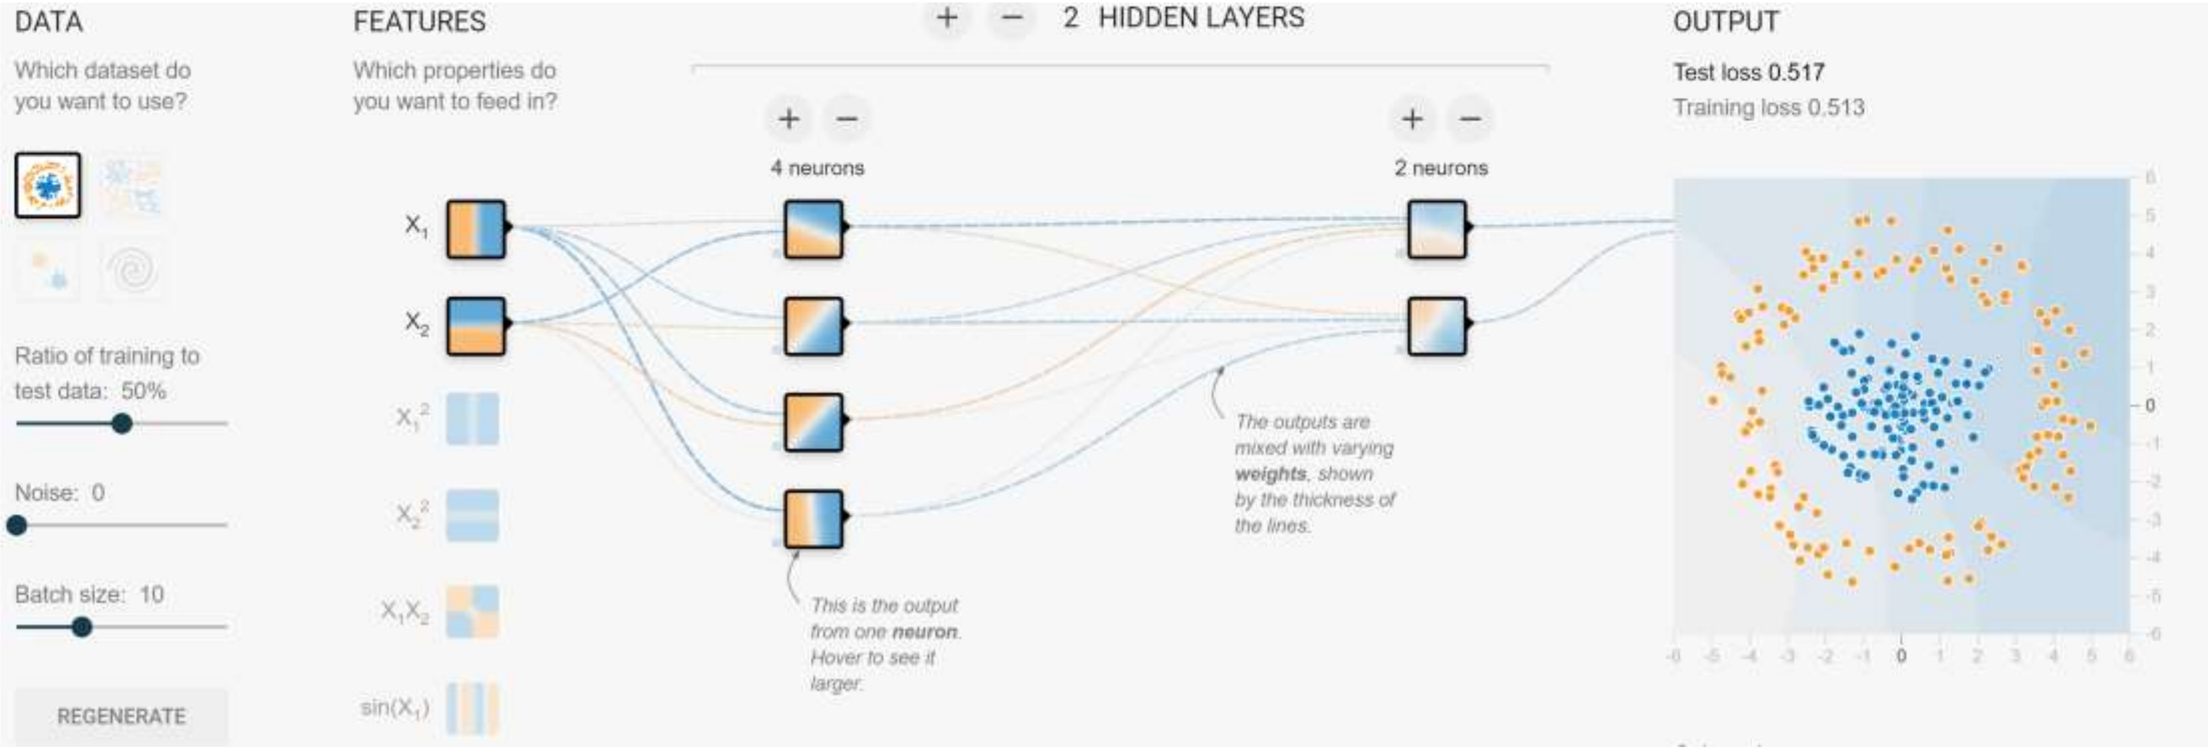

How to train network

in less than 10 minutes...

Anatomy of Perceptron

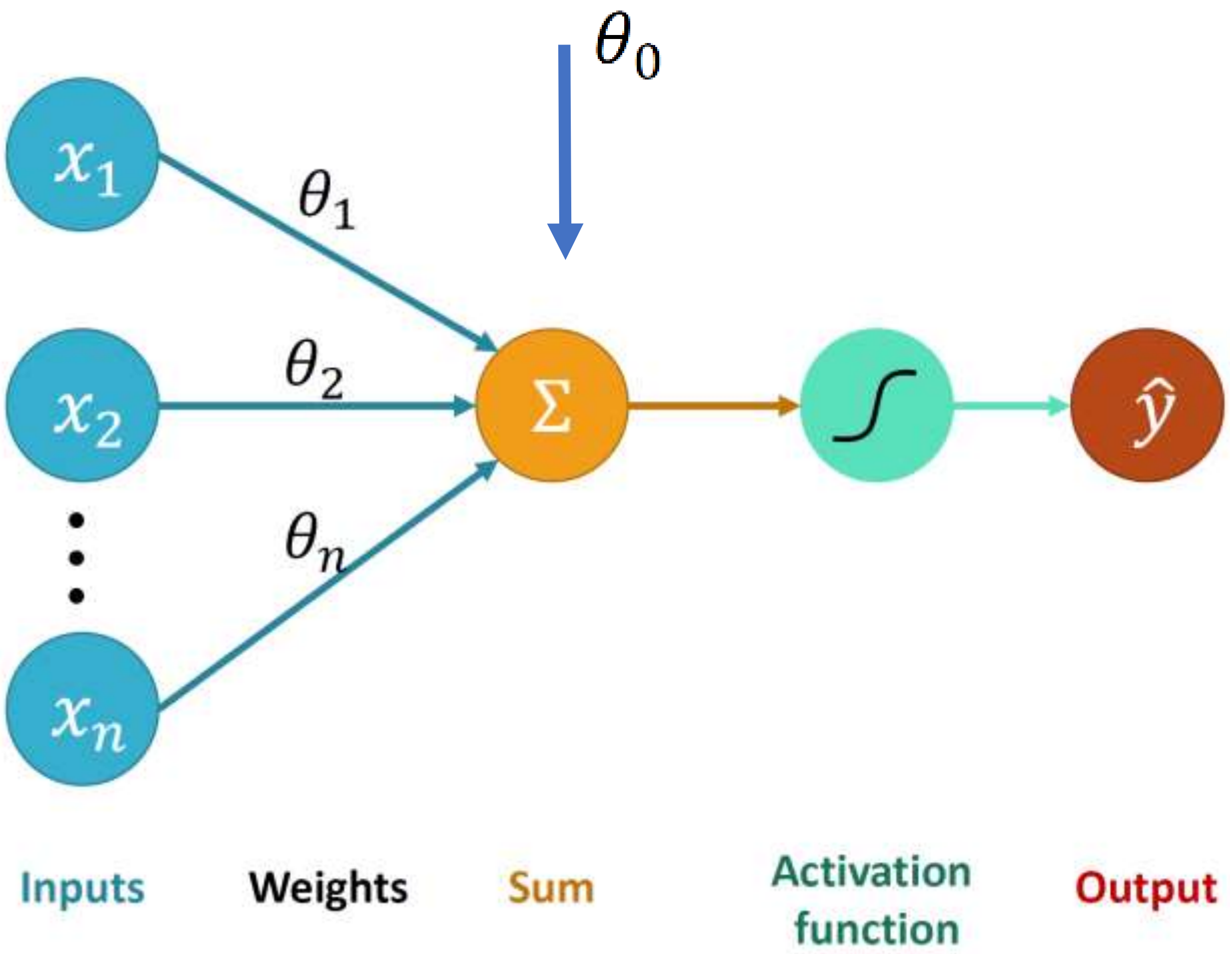

Output (prediction)

$$\hat{y} = g\left(\theta_0 + \overbrace{\sum_{i=1}^{n} x_i \; \theta_i}^{\text{Linear combination of inputs}}\right)$$

Bias

Non-linear activation function

Simple application of Neural Network

Does the image present goldfish?
[Carassius auratus]

Simple application of Neural Network

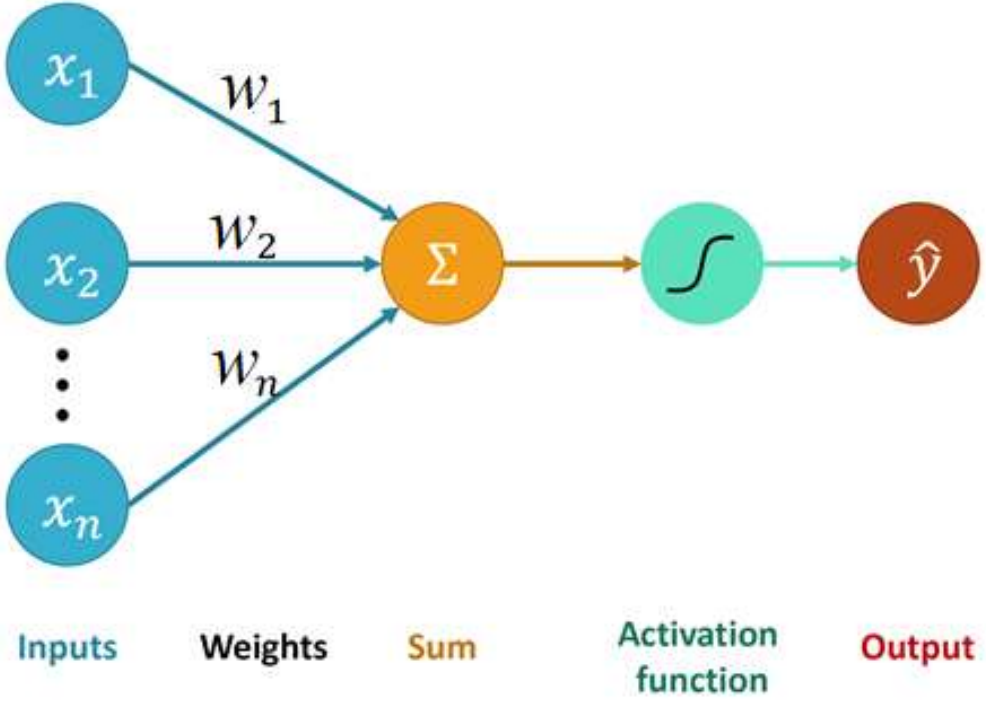

Is this object a goldfish?
(yes/no)

Simple application of Neural Network

Each pixel can be an input x

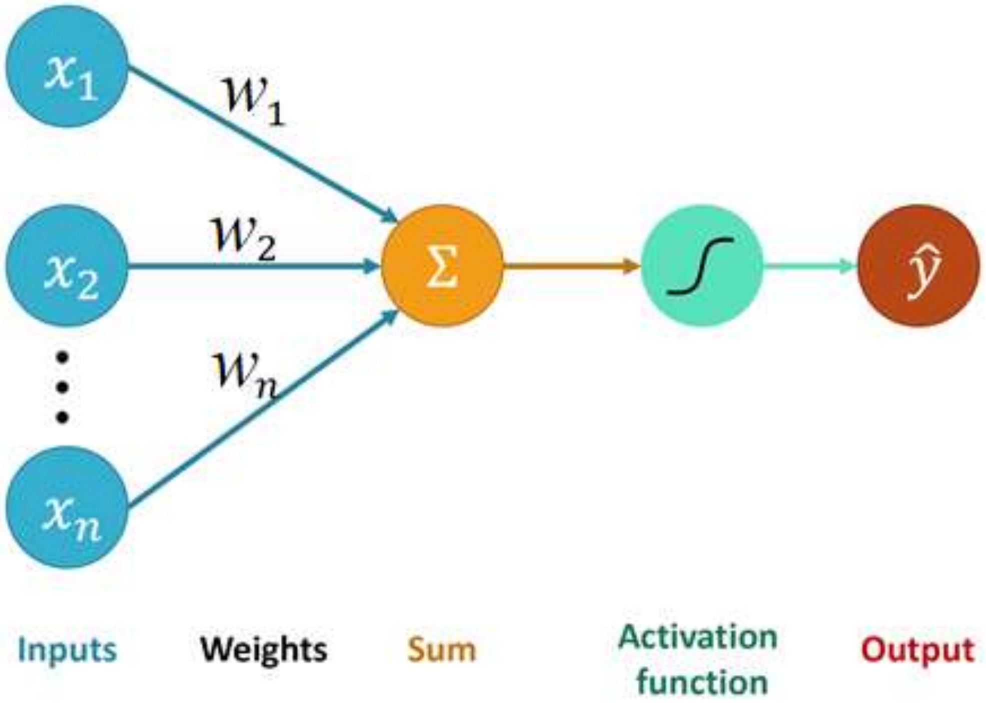

Is this object a goldfish?
(yes/no)

Simple application of Neural Network

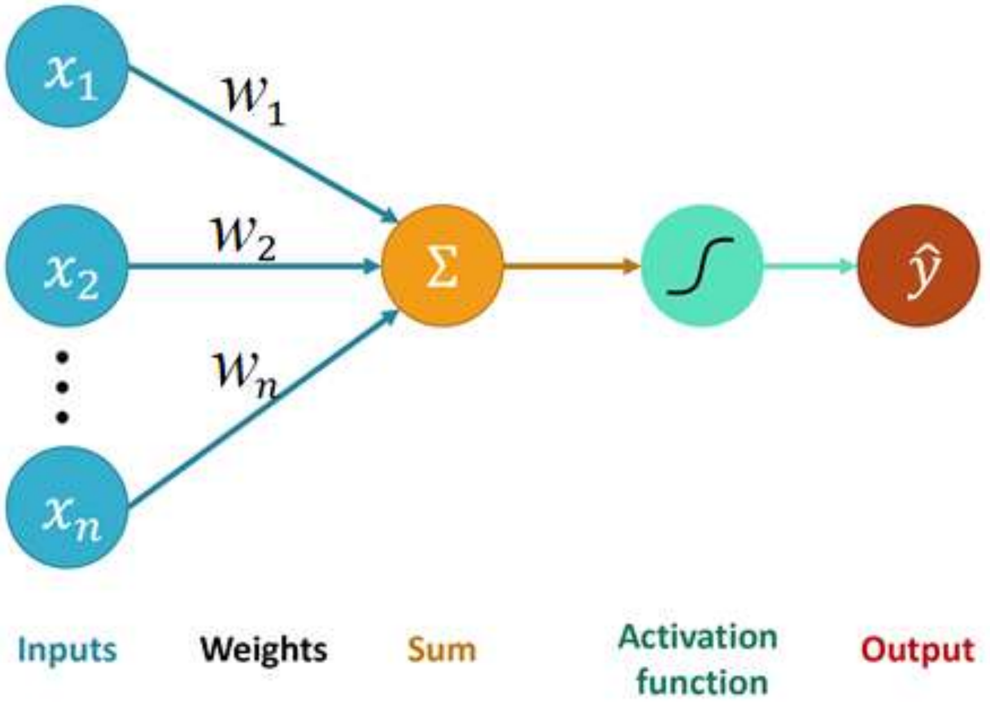

Is this object a goldfish?
(yes/no)

Hint:

$$f\left(\sum_i w_i x_i + b\right)$$

We do not know weights...

How to assign Weights?

- We don’t know weights, so let’s guess then!
- We can initialise our model with e.g. random weights
- Then we calculate error (loss function):

$$\text{error} = \left|\underbrace{\boldsymbol{d}}_{\text{known}} - \underbrace{x}_{\text{known}} \times \underbrace{W}_{\text{guess}}\right|$$

- If error is large we update weights

$w_i(t+1) = w_i(t) + r \cdot (d_j - y_j(t))x_{j,i}$, for all features $0 \leq i \leq n$, r is the learning rate

Collect some data

Calculate error and update weights

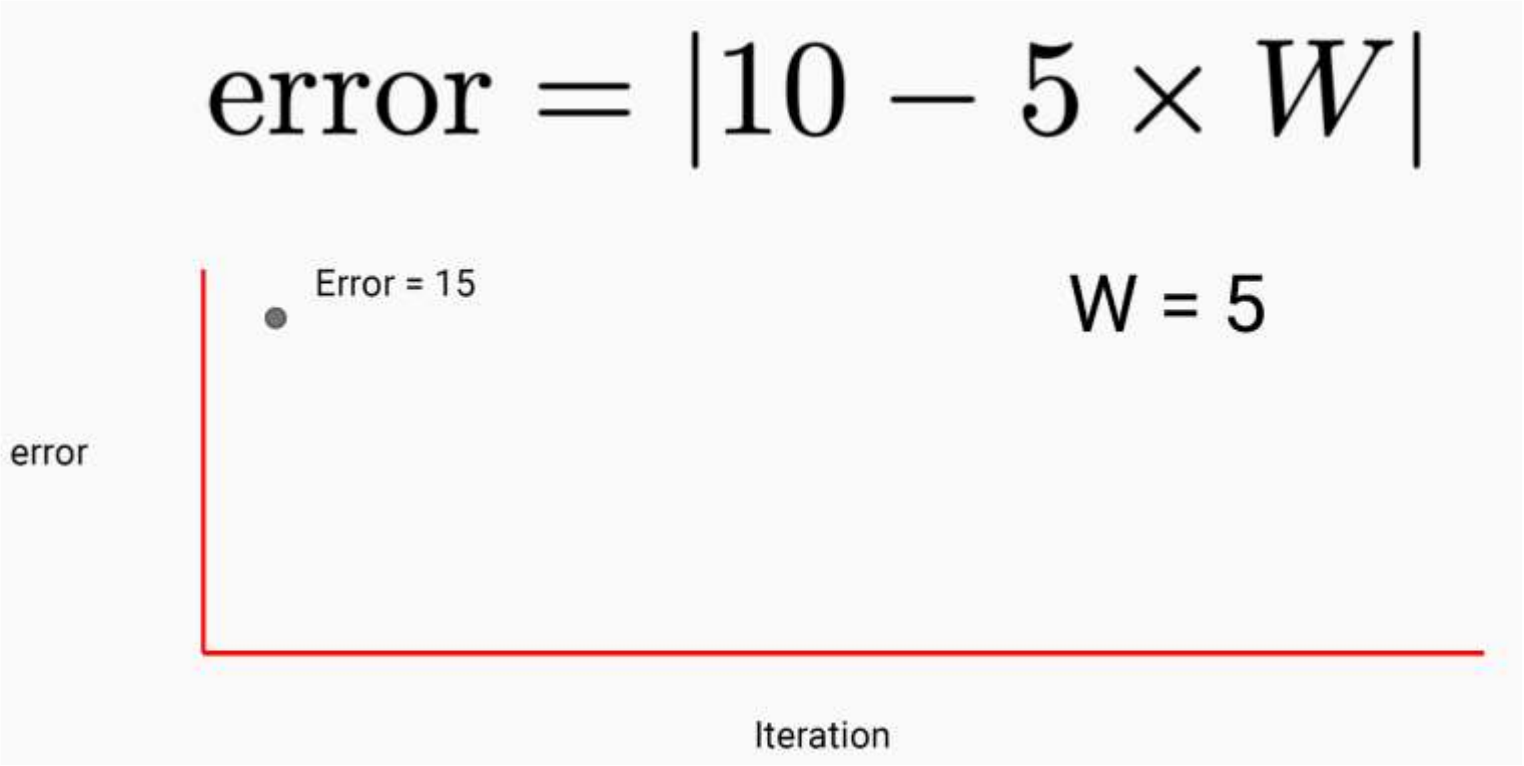

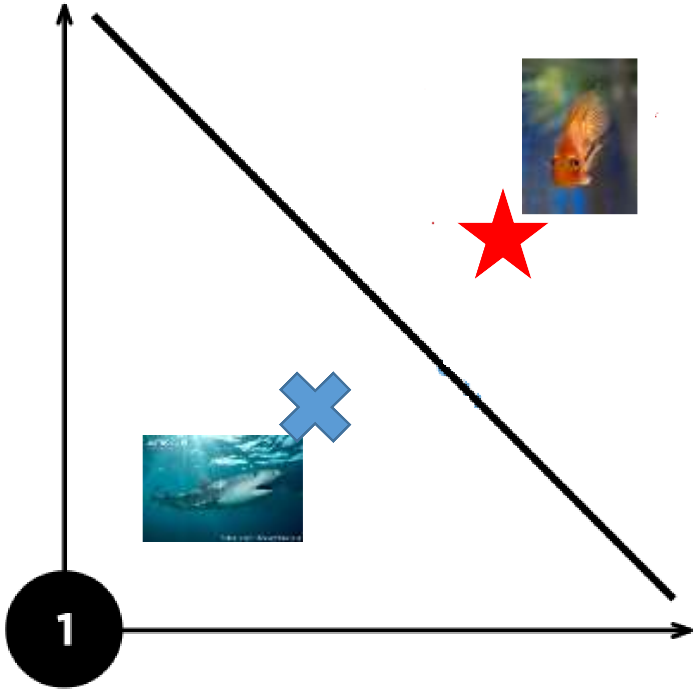

$$w_i(t+1) = w_i(t) + r \cdot (d_j - y_j(t))x_{j,i}$$

Collect More data

Calculate error and update weights...

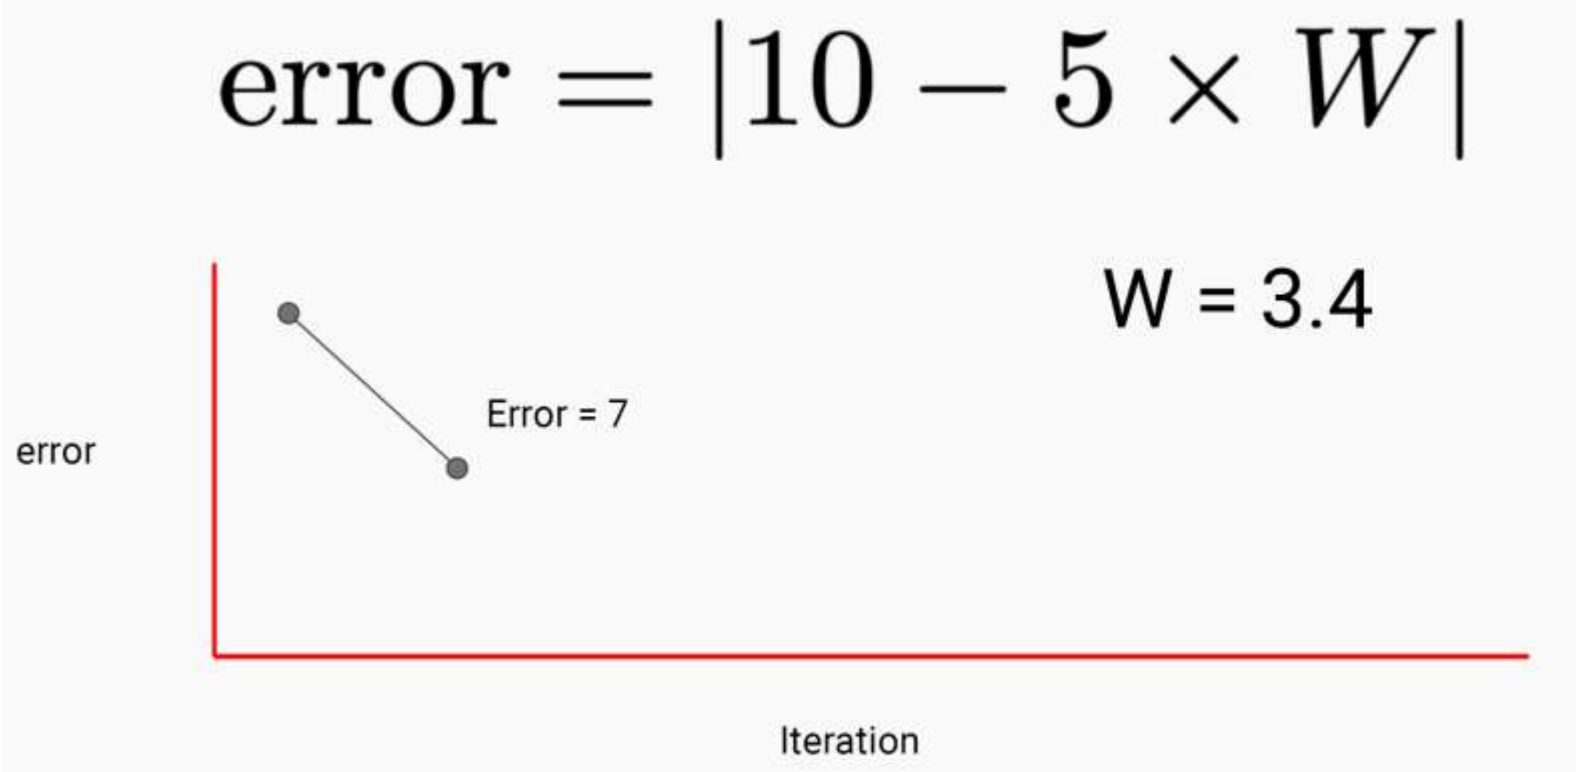

$$w_i(t+1) = w_i(t) + r \cdot (d_j - y_j(t))x_{j,i}$$

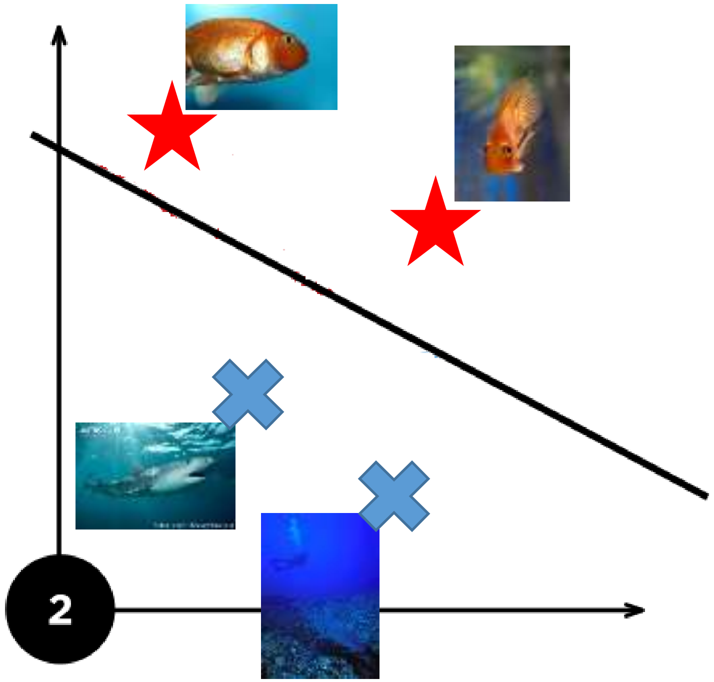

Collect even more data

Calculate error and update weights...

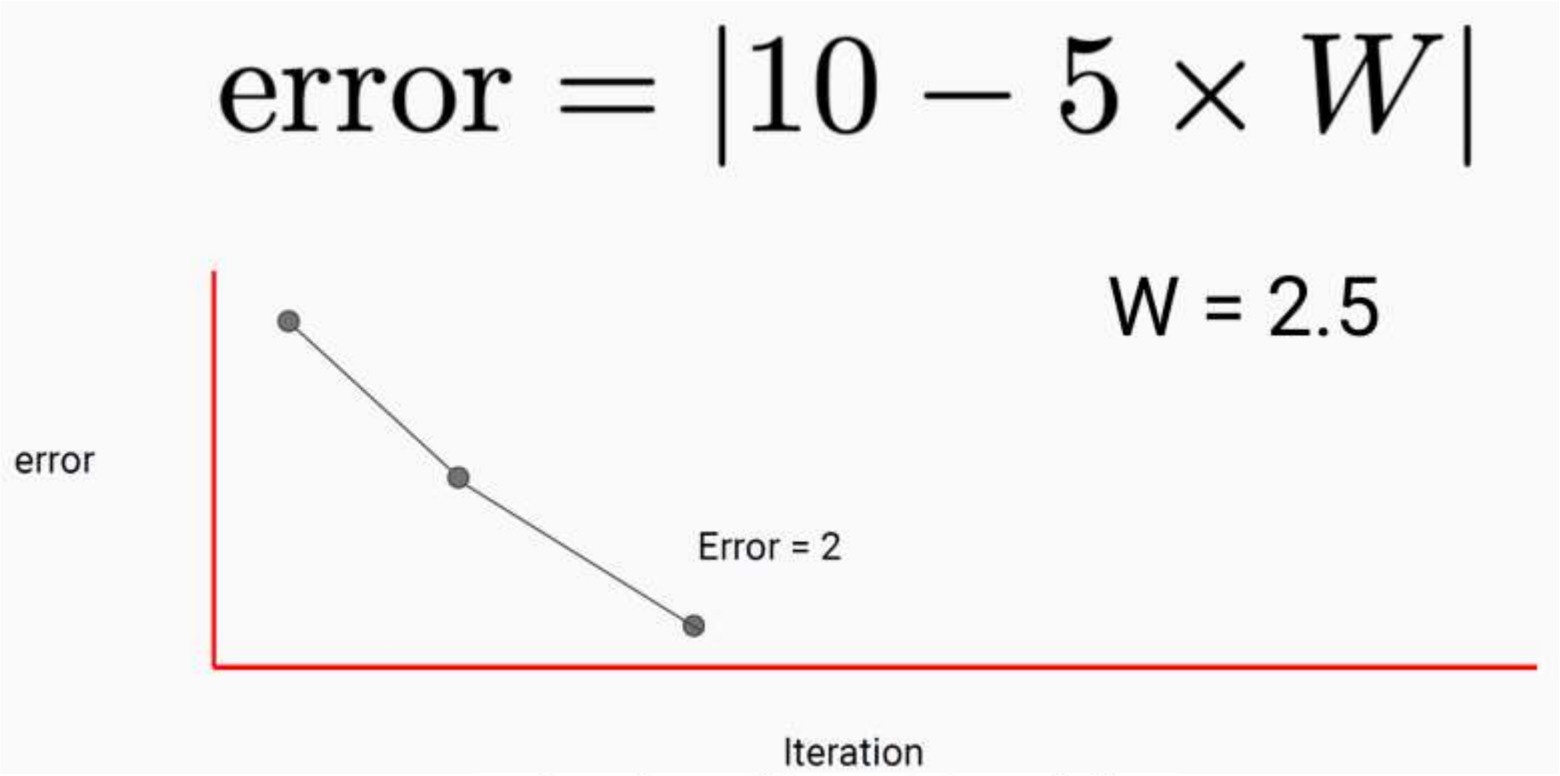

$$w_i(t+1) = w_i(t) + r \cdot (d_j - y_j(t))x_{j,i}$$

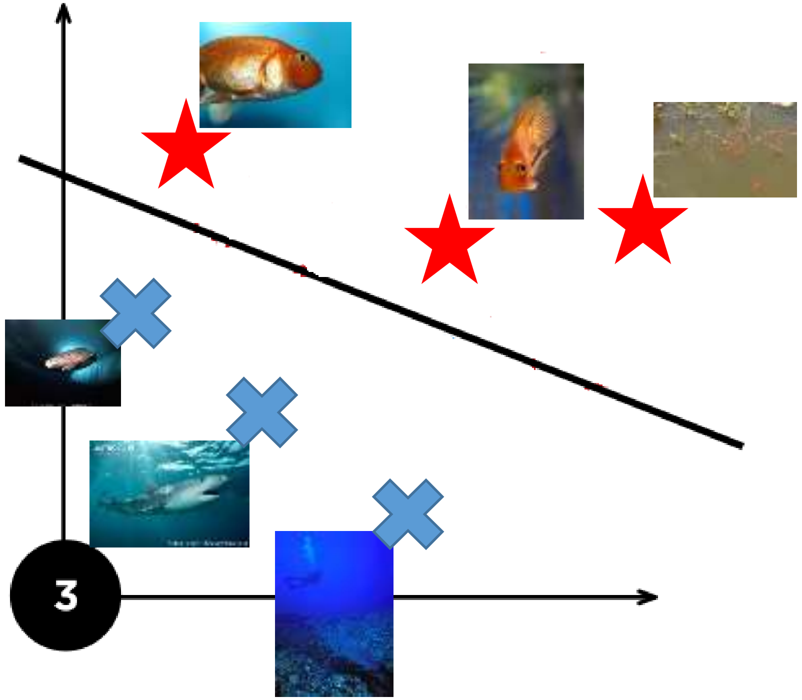

Collect a lot of data

Calculate error and update weights...

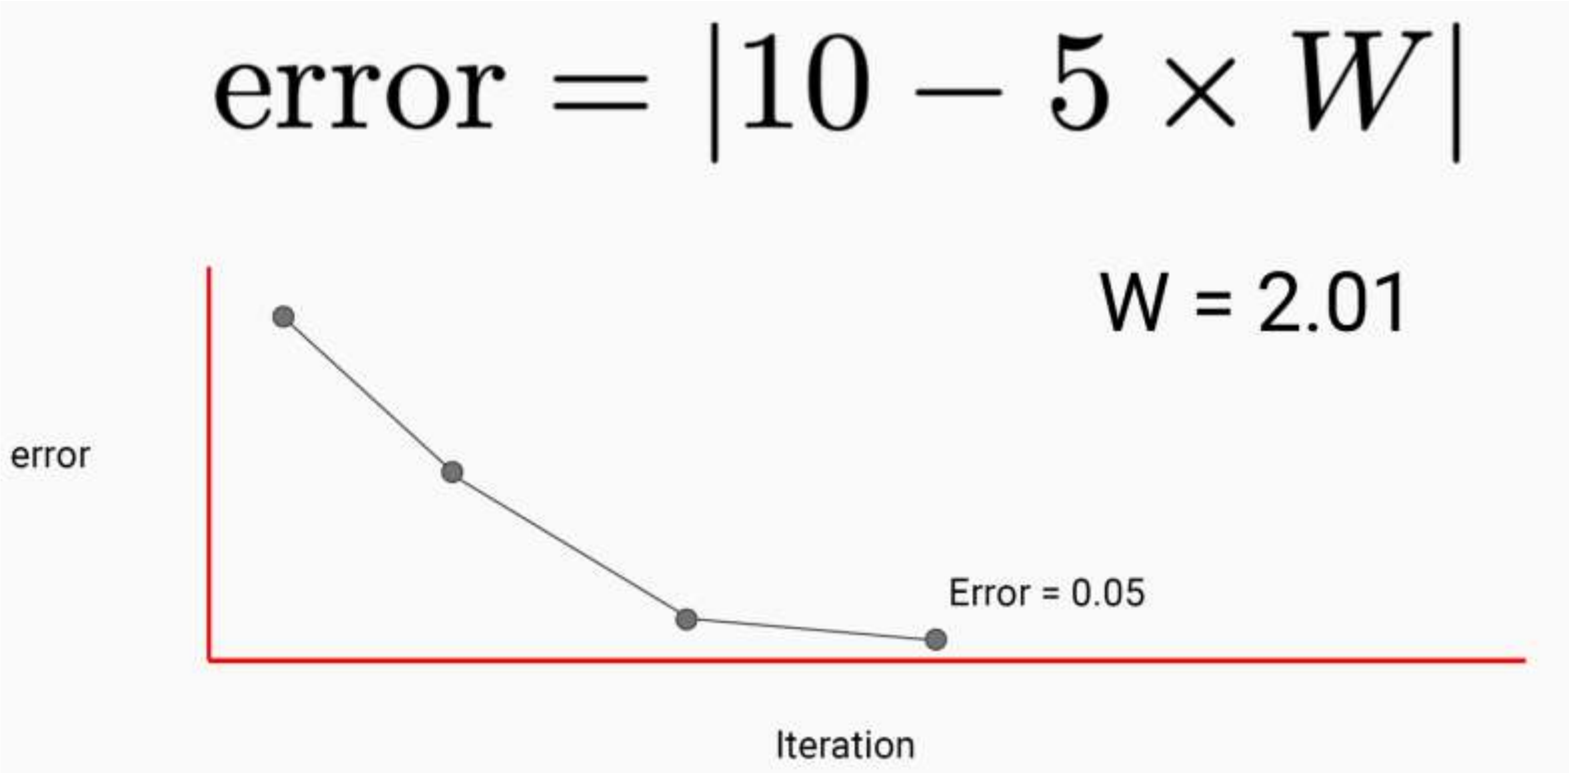

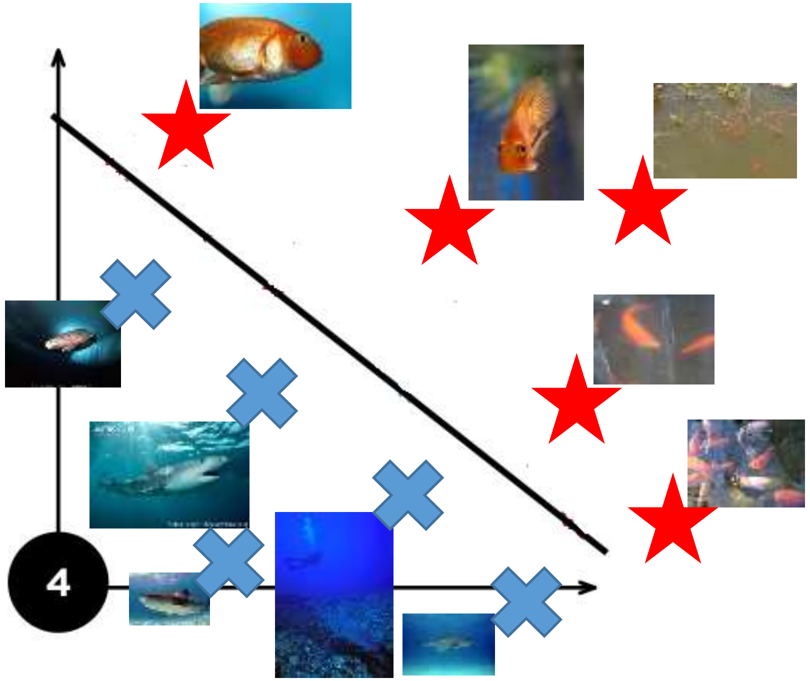

Single layer network

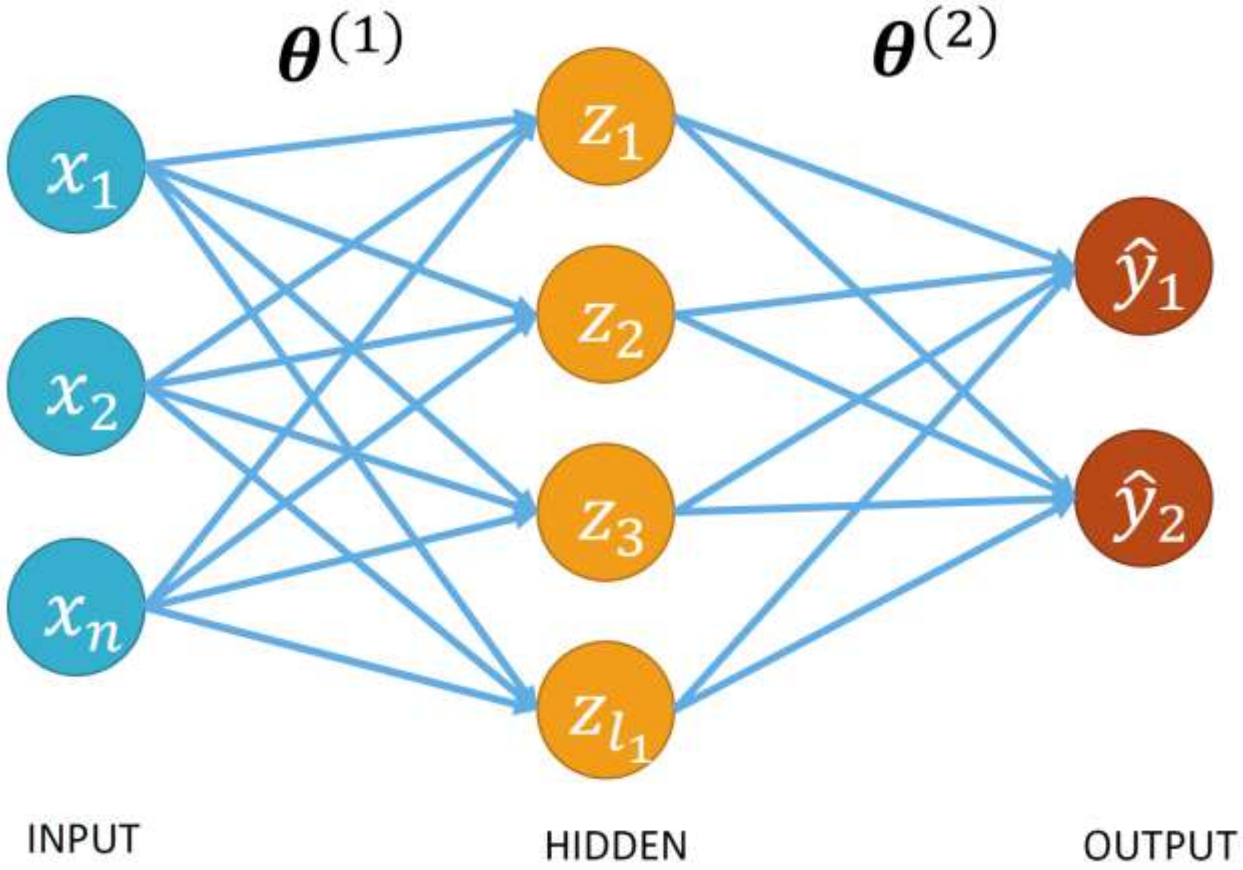

Shallow -> Deep Network

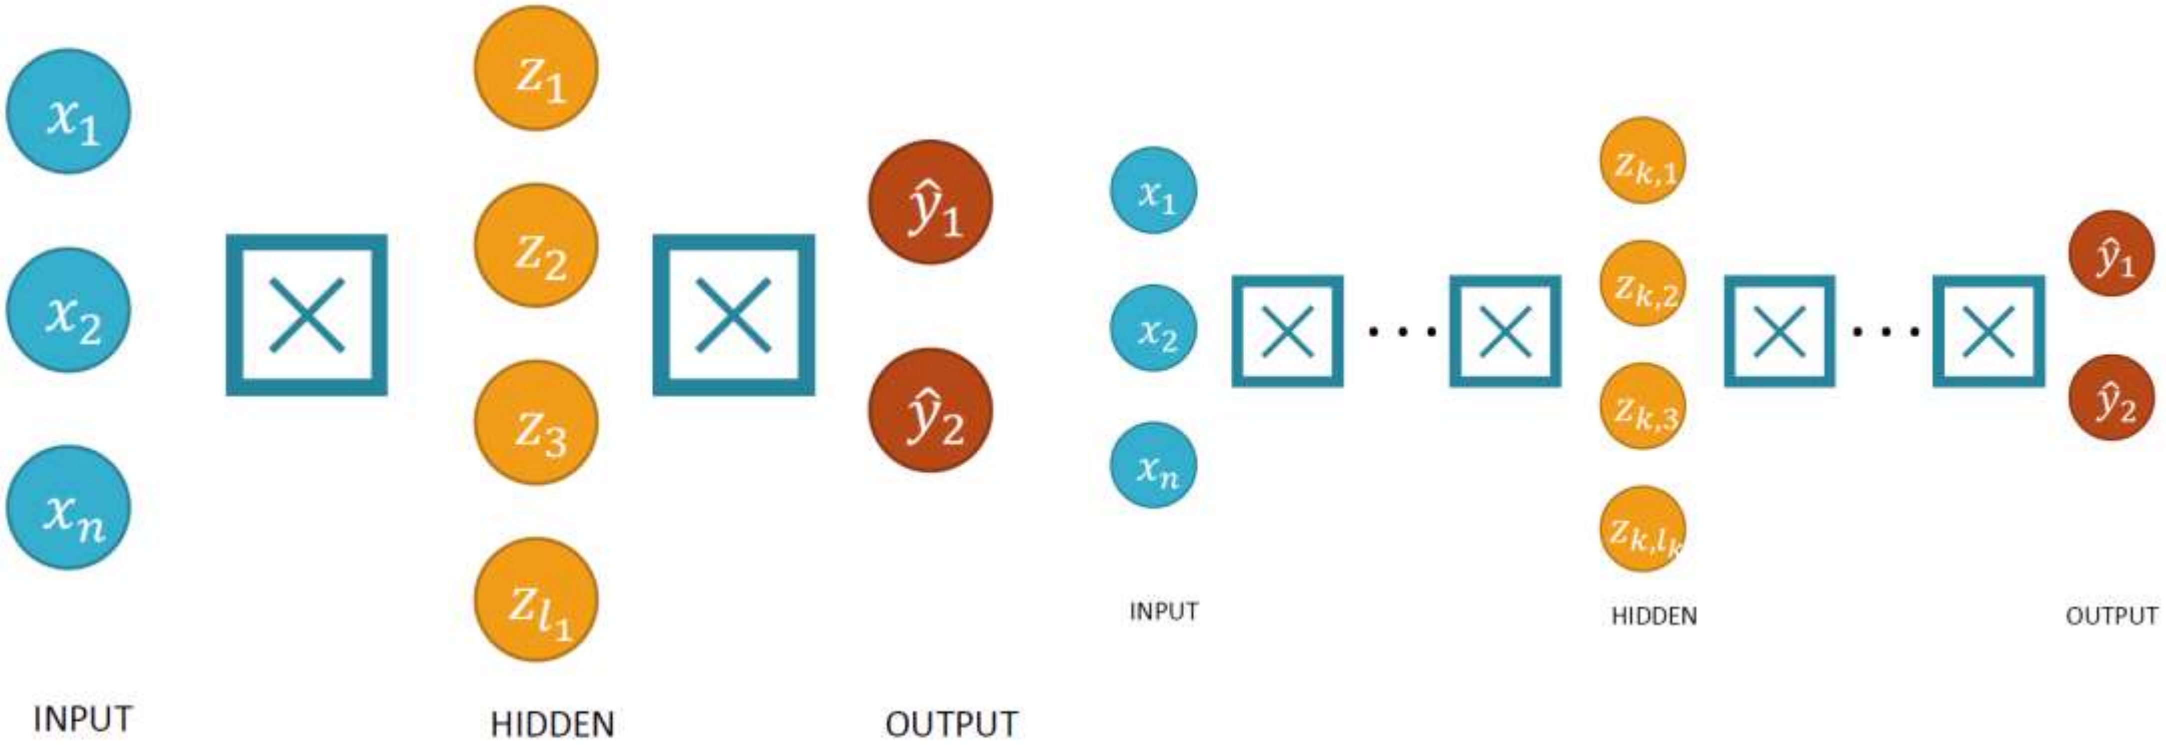

Optimisation

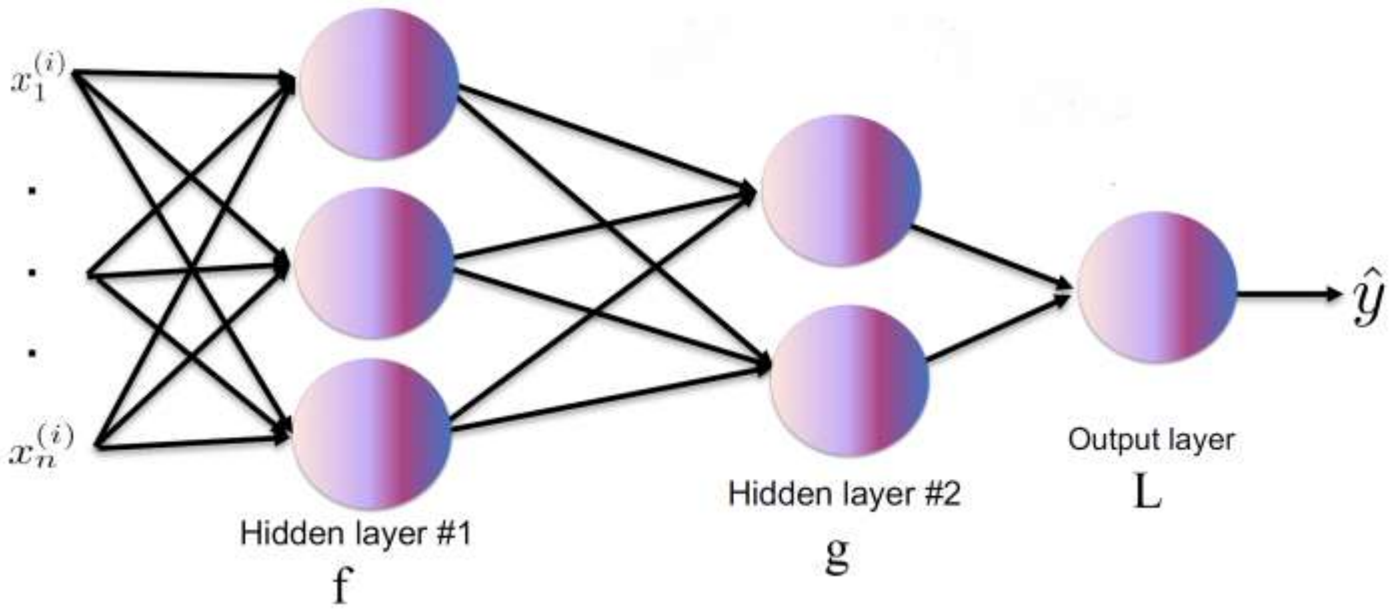

- Our (not so) deep network can be seen as composed functions

$$L(\boldsymbol{w}, \boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{L}(\boldsymbol{g}\big(\boldsymbol{f}\big(\boldsymbol{x}, \boldsymbol{w}_{\boldsymbol{f}}\big), \boldsymbol{w}_{\boldsymbol{g}}\big), \boldsymbol{y})$$

Optimisation - backpropagation

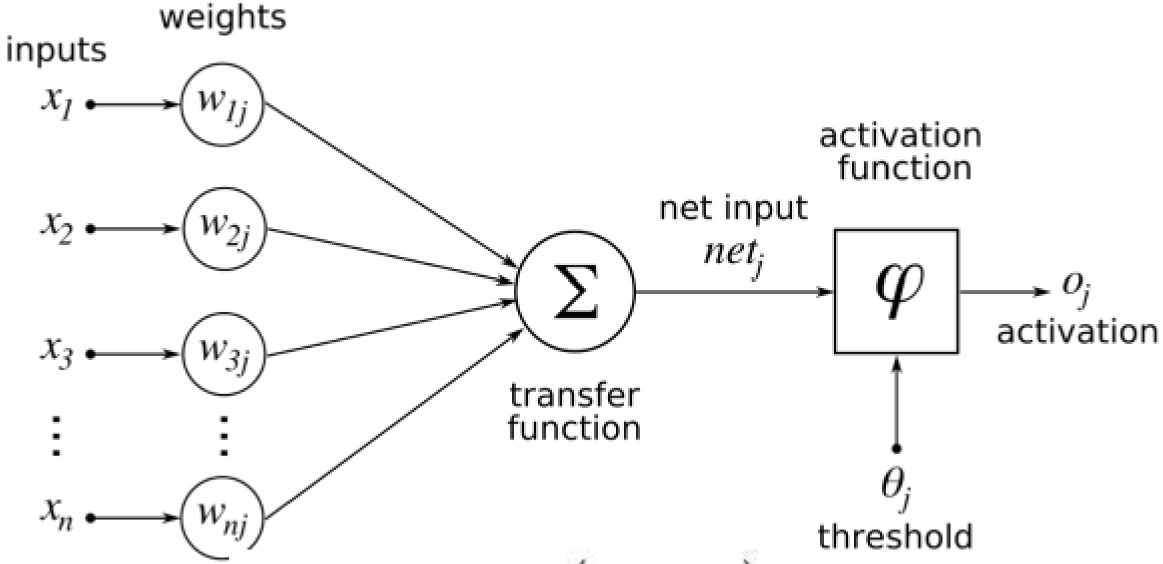

$$o_j = \varphi(\text{net}_j) = \varphi\left(\sum_{k=1}^{n} w_{kj} o_k\right)$$

$$\frac{\partial L}{\partial w_{ij}} = \frac{\partial L}{\partial o_j}\frac{\partial o_j}{\partial w_{ij}} = \frac{\partial L}{\partial o_j}\frac{\partial o_j}{\partial \text{net}_j}\frac{\partial \text{net}_j}{\partial w_{ij}}$$

$$\frac{\partial \text{net}_j}{\partial w_{ij}} = \frac{\partial}{\partial w_{ij}}\left(\sum_{k=1}^{n} w_{kj} o_k\right) = \frac{\partial}{\partial w_{ij}} w_{ij} o_i = o_i.$$

$$\frac{\partial o_j}{\partial \text{net}_j} = \frac{\partial}{\partial \text{net}_j}\varphi(\text{net}_j) = \varphi(\text{net}_j)(1 - \varphi(\text{net}_j)) = o_j(1 - o_j)$$

$$\frac{\partial L}{\partial o_j} = \sum\left(\frac{\partial L}{\partial \text{net}_\ell}\frac{\partial \text{net}_\ell}{\partial o_j}\right) = \sum\left(\frac{\partial L}{\partial o_\ell}\frac{\partial o_\ell}{\partial \text{net}_\ell}\frac{\partial \text{net}_\ell}{\partial o_j}\right) = \sum\left(\frac{\partial L}{\partial o_\ell}\frac{\partial o_\ell}{\partial \text{net}_\ell} w_{j\ell}\right)$$

Convolutional Neural Network

For Image Segmentation

Image segmentation using AI/ML

- Image segmentation is a process of estimating a plausible partitioning of image into two or more segments
- Most of the classic segmentation algorithms are based on one of two basic properties of image intensity values (similarities or edges)
- Currently segmentation is mostly done by convolutional neural networks

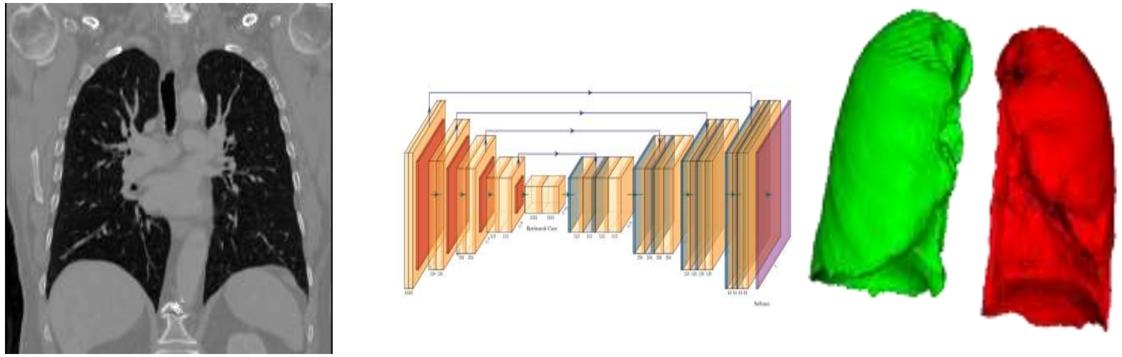

- Each pixel (voxel) gets a label
- Pixel level classification problem
  - (but pixel contains information about location)
- We (may) need some labels for training

Fully convolutional Network

- We have used Fully Convolutional Network (FCN) for classification

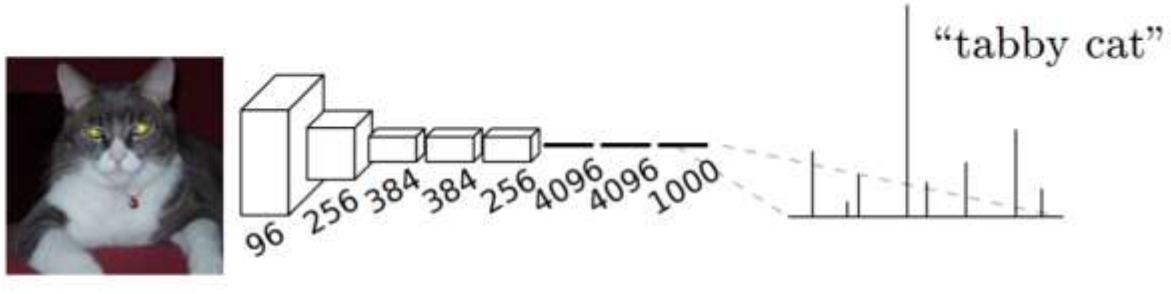

- The fully connected layers of this network have fixed dimensions and throw away spatial coordinates

Long, Jonathan, Evan Shelhamer, and Trevor Darrell. "Fully convolutional networks for semantic segmentation." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015.

NN vs. CNN

469 x 387

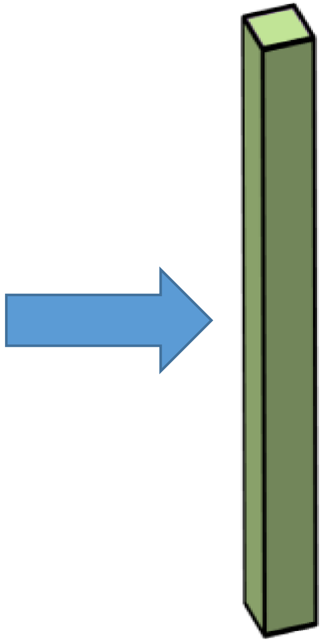

181503 x 1

- Image (set of pixels) as a vector?
- All (important) structural information is lost!
- All pixels are fully connected
- Size?

Fully convolutional Network

- Fully Convolutional Network (FCN) for segmentation

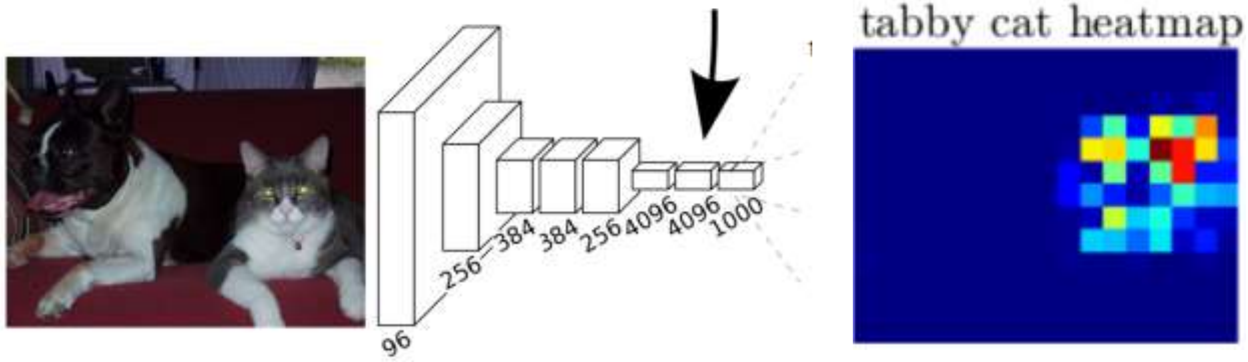

- We replace fully connected layers by convolutional layers
- Transforming fully connected layers into convolution layers enables a classification net to output a heatmap (very coarse segmentation)
- We need to add extra layers and spatial loss to make dense predictions for per-pixel tasks

Long, Jonathan, Evan Shelhamer, and Trevor Darrell. "Fully convolutional networks for semantic segmentation." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015.

Convolutional NN (CNN)

- Convolution with trainable filters:
  - Local connectivity is exploited by evaluating the pixels in close proximity (on the image grid)
  - Receptive field is the size of the filter (3x3, 5x5, etc)
  - We evaluate the same filter all across the image – translation invariance

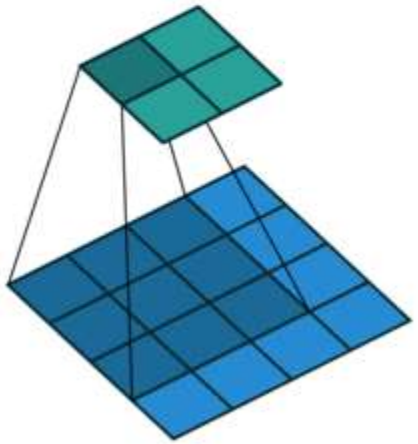

Vincent Dumoulin, Francesco Visin - A guide to convolution arithmetic for deep learning

Convolutional Layer - example

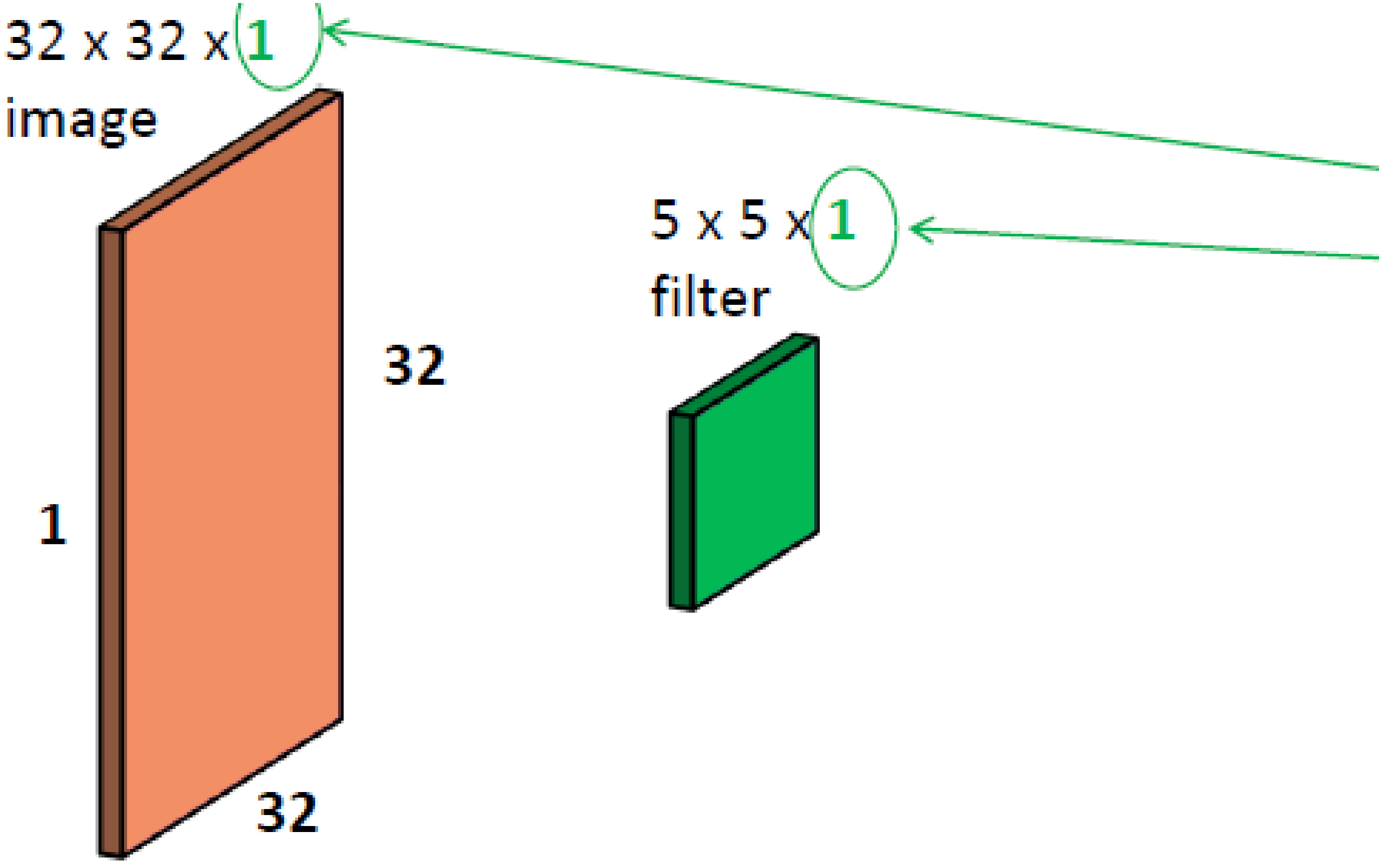

NB: the filter depth extends to the input image depth

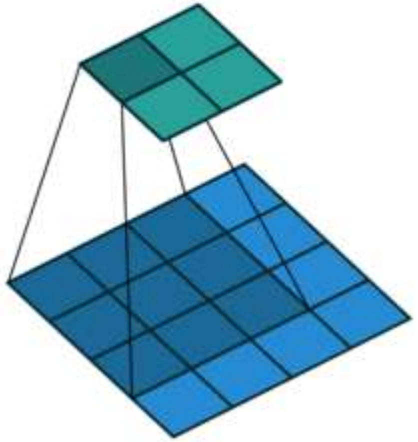

Convolutional Layer - example

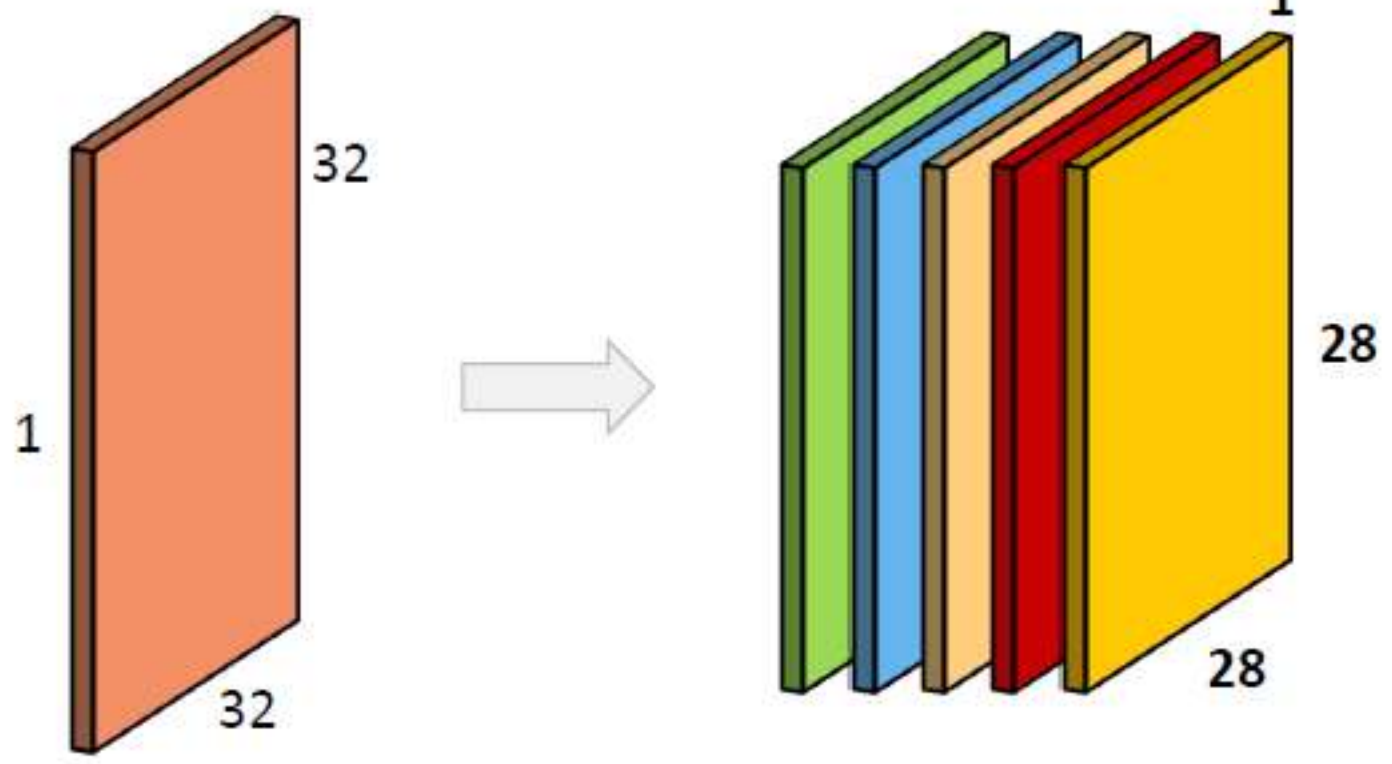

- Each new filter will create a new feature map

Learnt filters

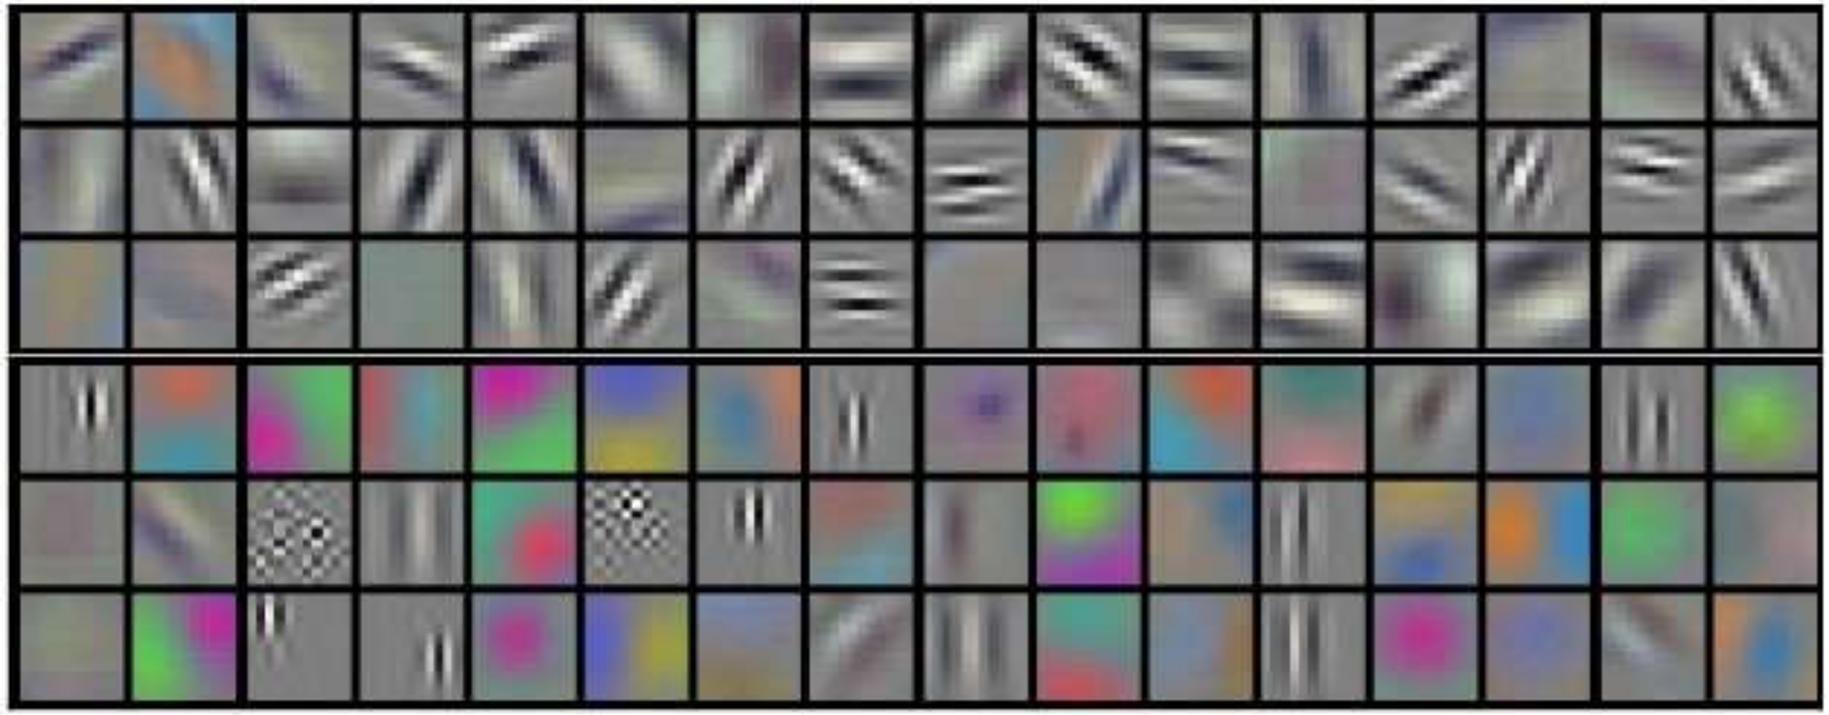

- Example filters learned by AlexNet. Each of the 96 filters shown here is of size [11x11x3], and each one is shared by the 55*55 neurons in one depth slice.

Feature map – AlexNet

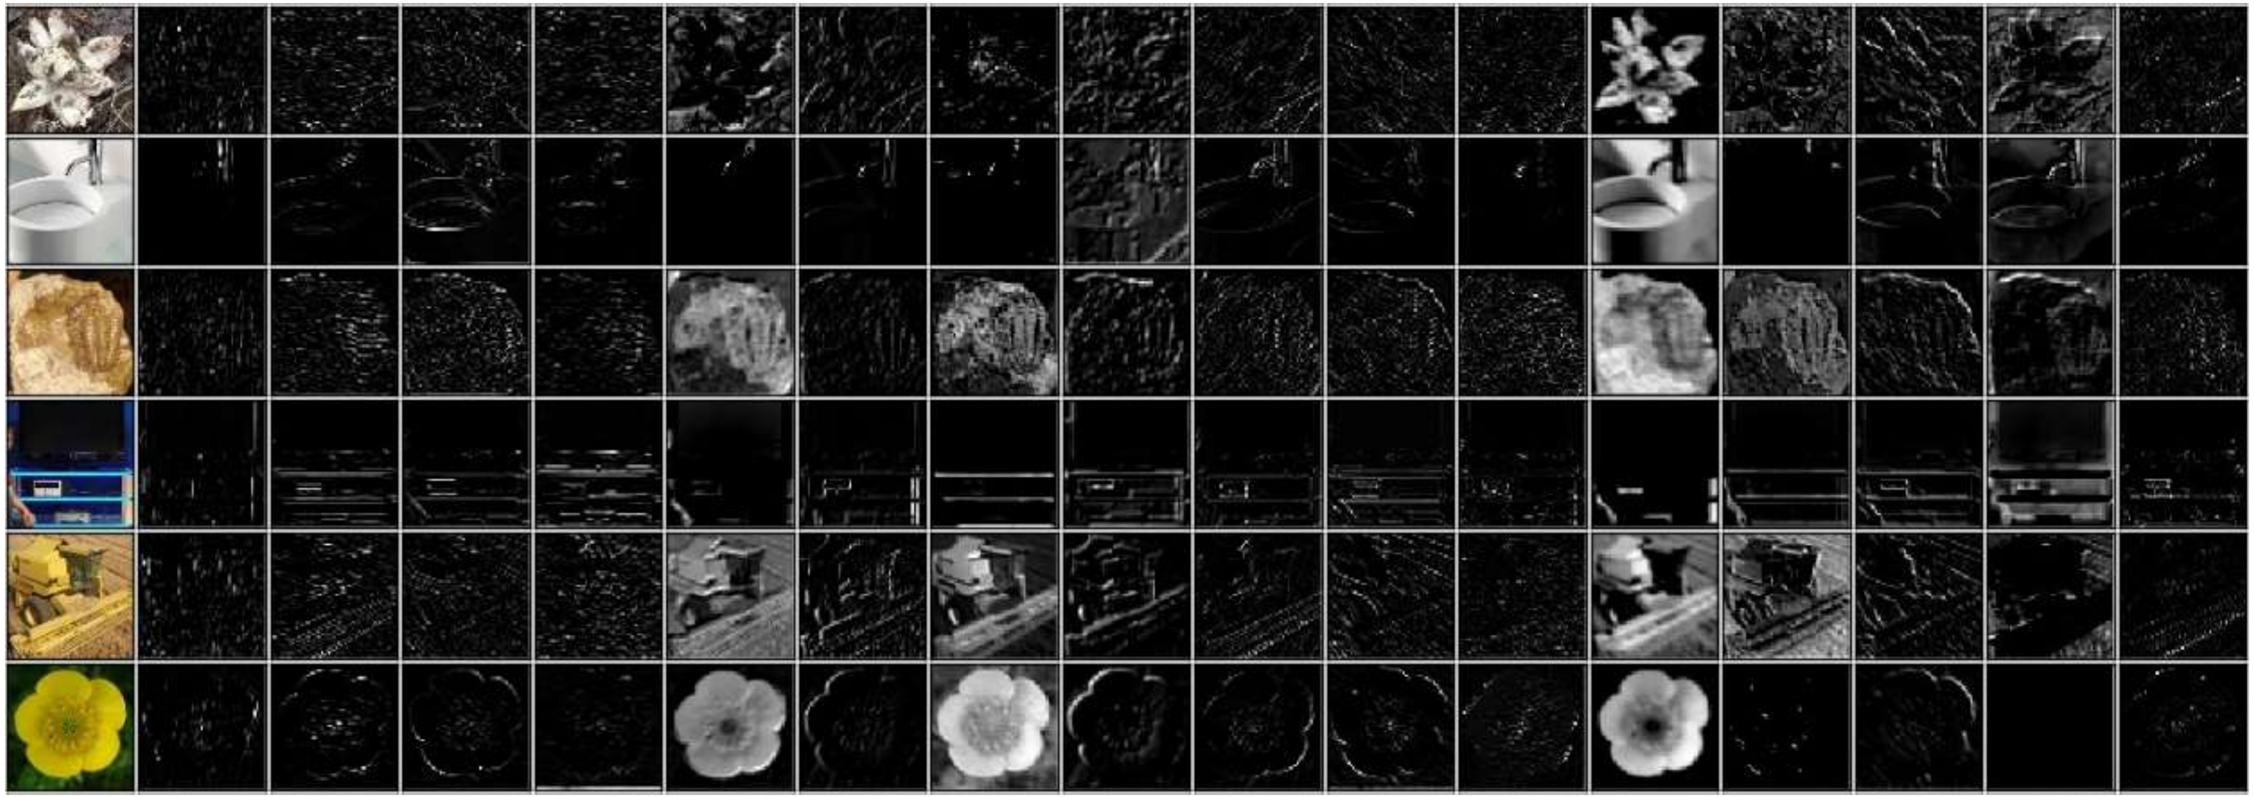

- The leftmost column contains random test images, while the remaining columns show a random subset of the feature maps in the first convolutional layer

Fully convolutional Network

- Alternatively, we can use a stack of same-sized convolutional layers to map the input image to the output one
- Quite good results, but it was extremely computationally expensive

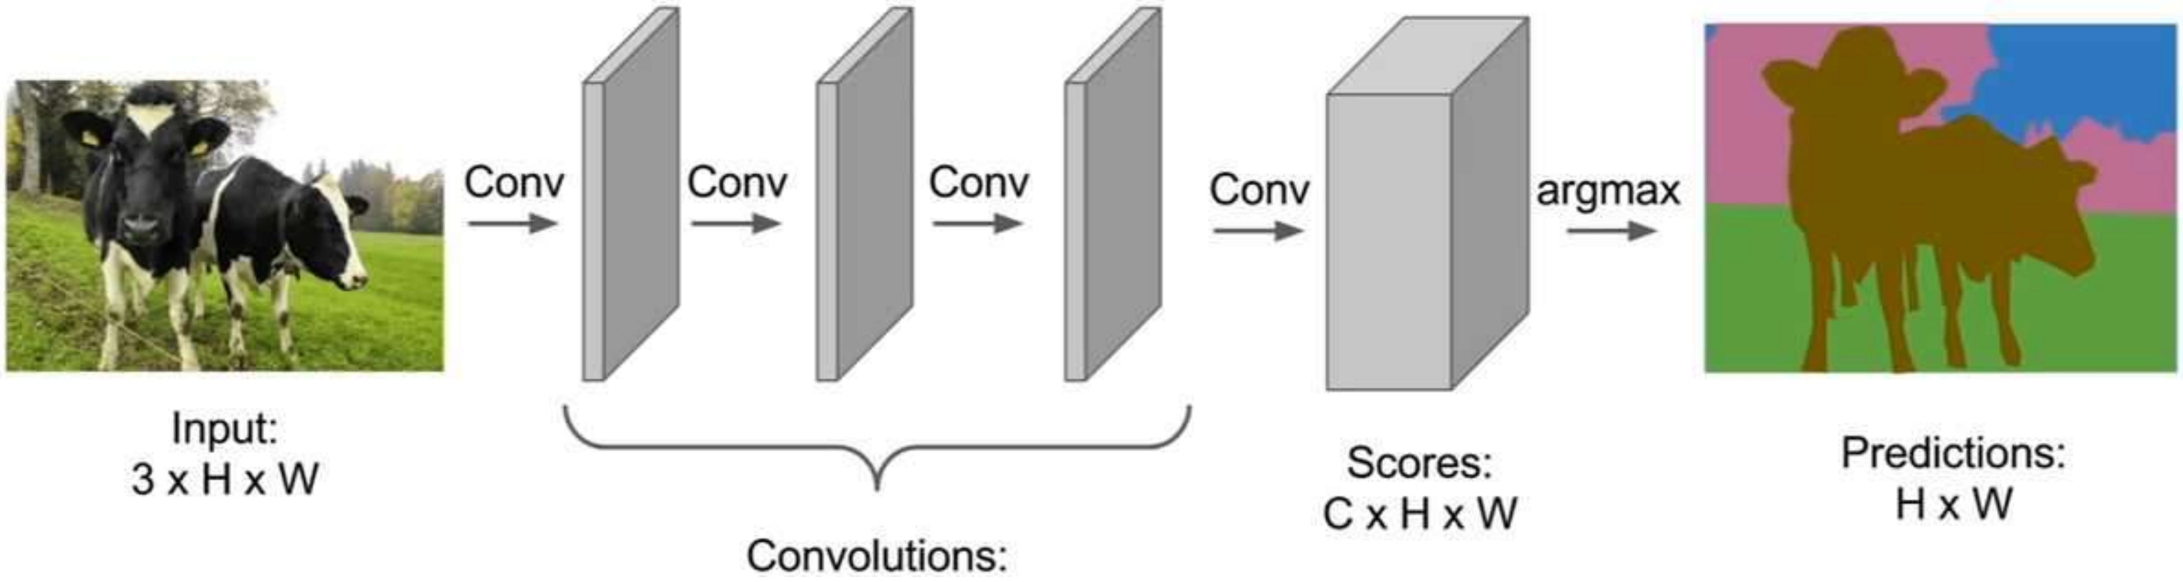

Encoder - decoder

- Typical architecture in image segmentation (Encoder-Decoder)

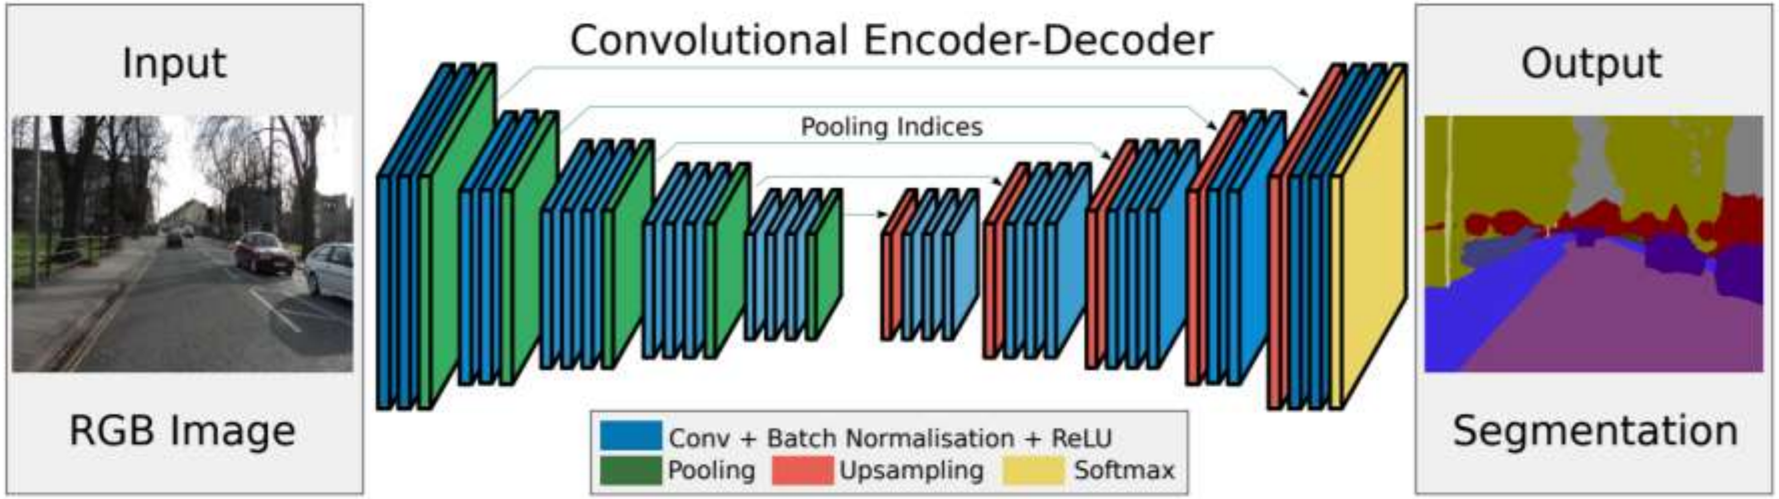

Badrinarayanan, Vijay, Alex Kendall, and Roberto Cipolla. "Segnet: A deep convolutional encoder-decoder architecture for image segmentation." *IEEE transactions on pattern analysis and machine intelligence* 39.12 (2017): 2481-2495.

U-Net

- U-Net consists of
  - Encoder to capture information about the object (the content)
  - Decoder to capture information about localisation of object
  - Skip connection to get more precise locations
- End-to-end solution

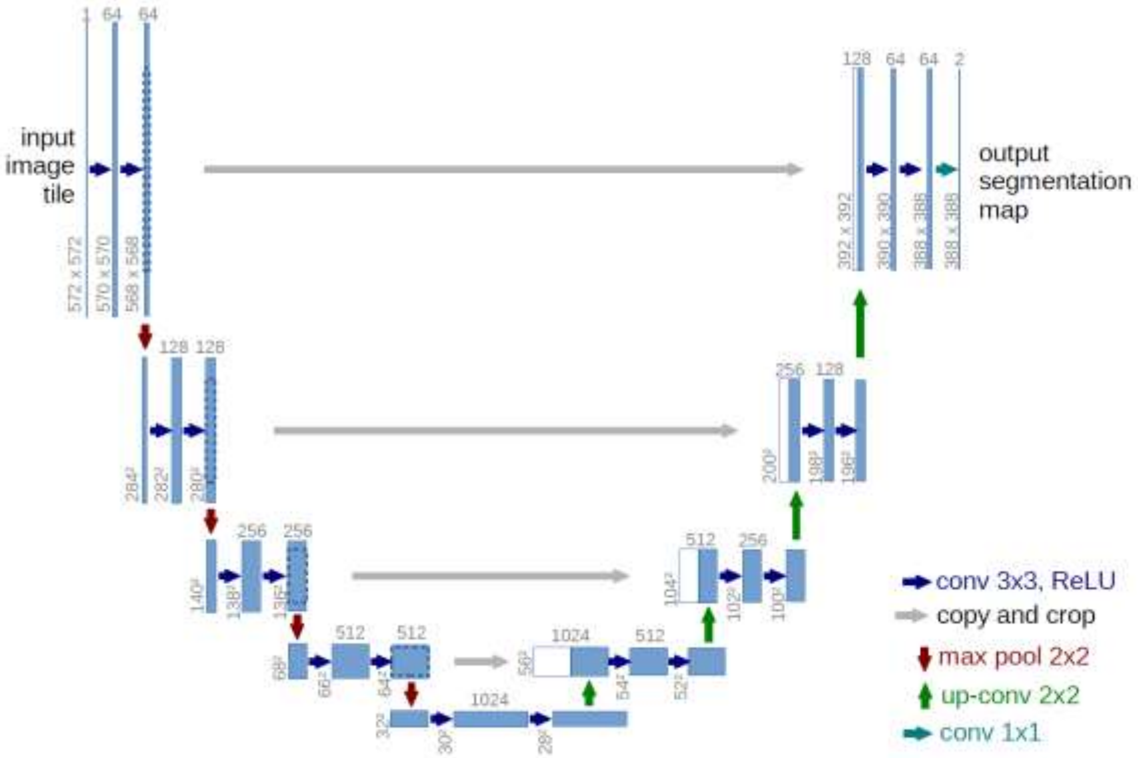

Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." *International Conference on Medical image computing and computer-assisted intervention*. Springer, Cham, 2015.

Mutlimodal Segmentation

Full-scale Unet - UNet3+

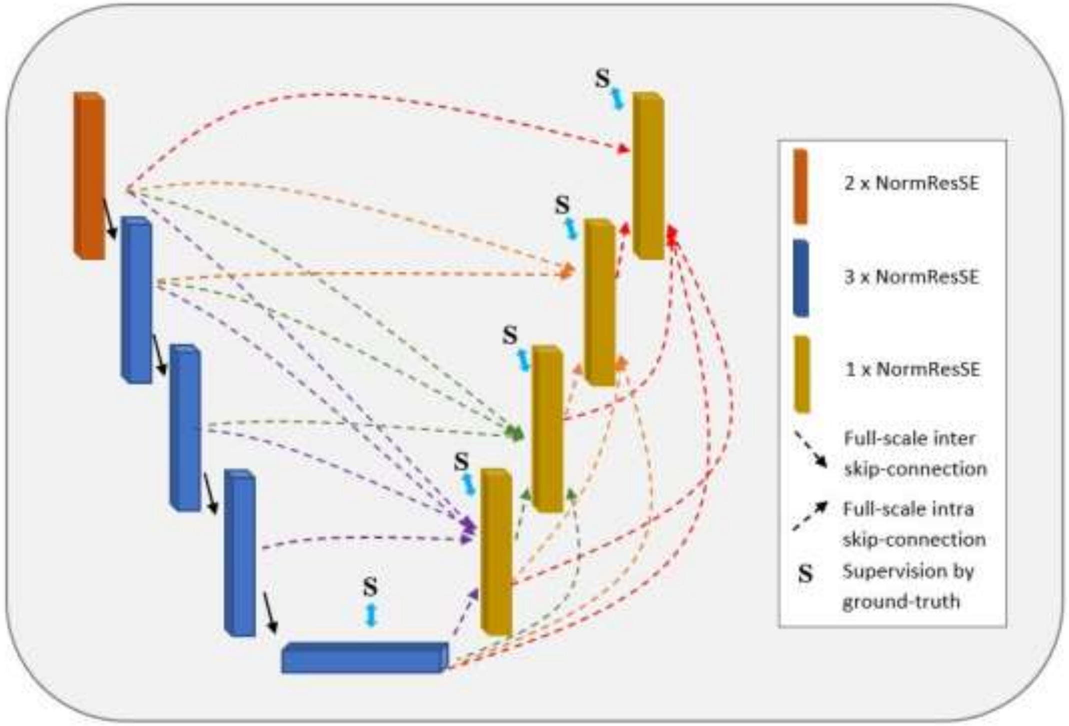

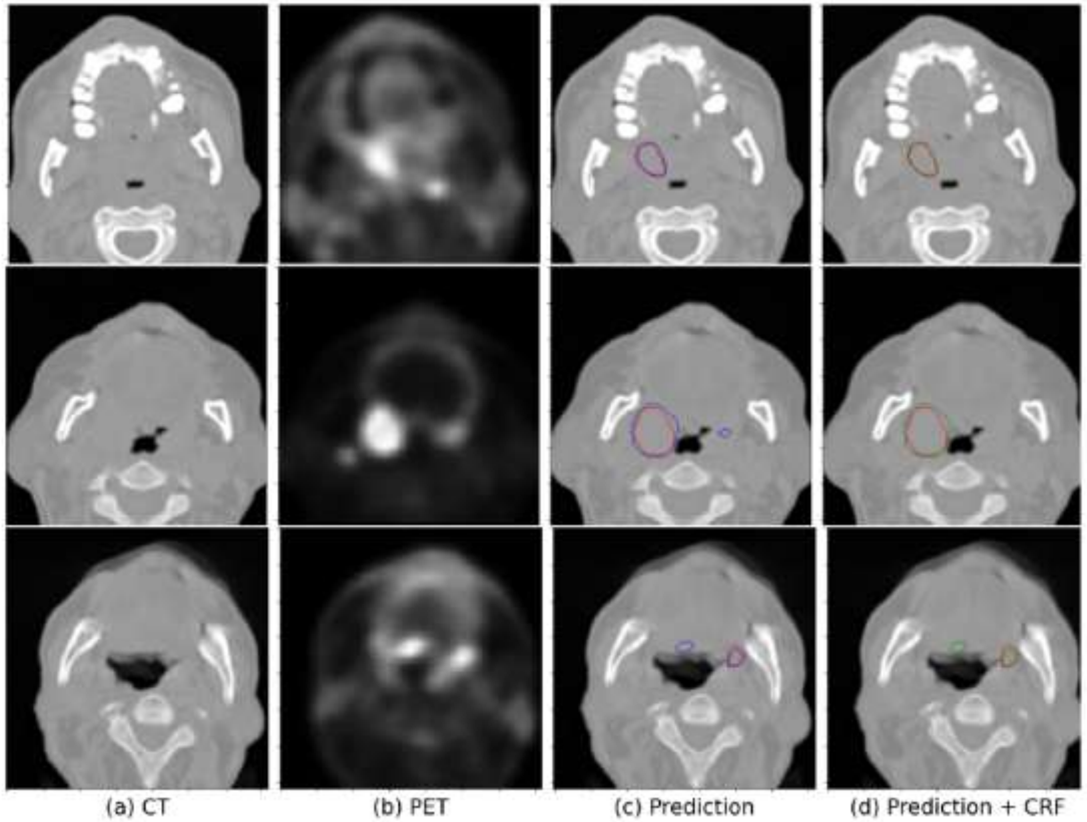

E. Bourigault et al. Multimodal PET/CT tumour segmentation and patient survival prediction using a full-scale UNet with attention (2021)

Radiomics

- With high-throughput computing, it is now possible to rapidly extract innumerable quantitative features from tomographic images (computed tomography [CT], magnetic resonance [MR], or positron emission tomography [PET] images).
- The conversion of digital medical images into mineable high-dimensional data, a process that is known as radiomics.
- Radiomics are motivated by the concept that biomedical images contain information that reflects underlying pathophysiology and that these relationships can be revealed via quantitative image analyses.

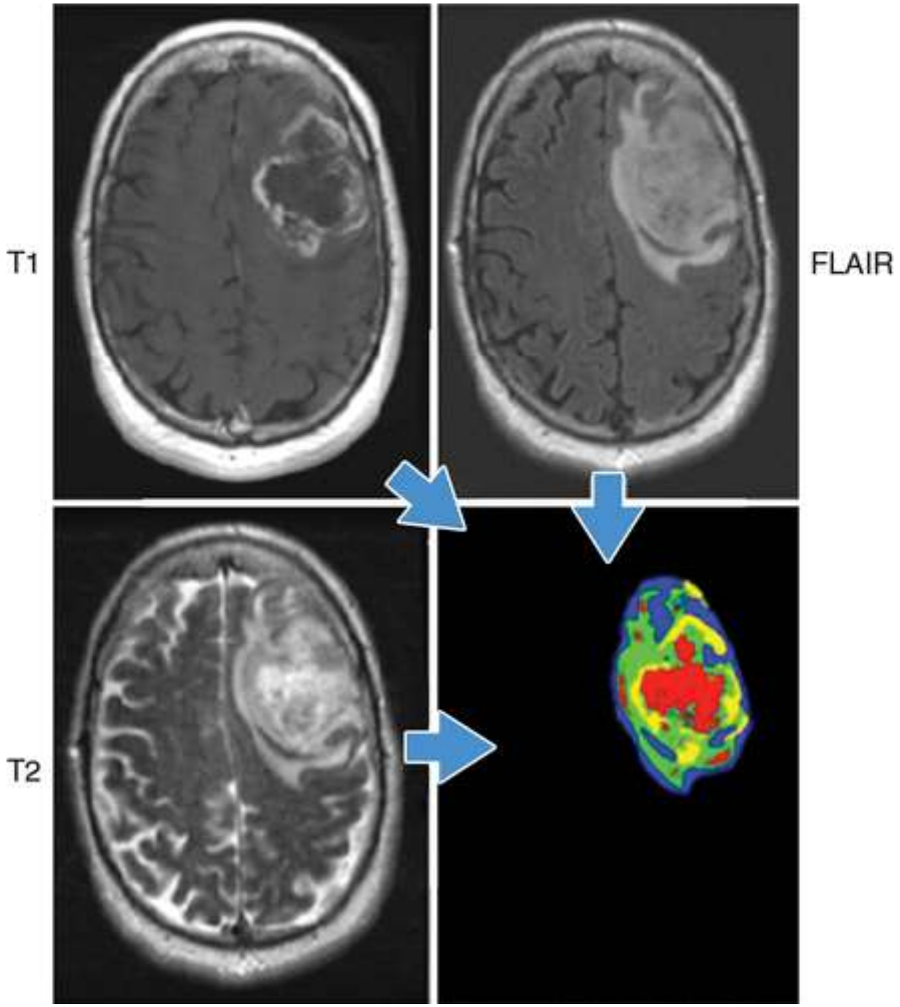

Habitats in a patient with glioblastoma multiforme.
Habitats were defined by combining unenhanced and contrast-enhanced T1-weighted, 120-msec echo time T2-weighted, and fluid-attenuated inversion recovery (FLAIR) images.

Gillies, Robert J., et al. "Radiomics: images are more than pictures, they are data." Radiology 278.2 (2016): 563-

Radiomics

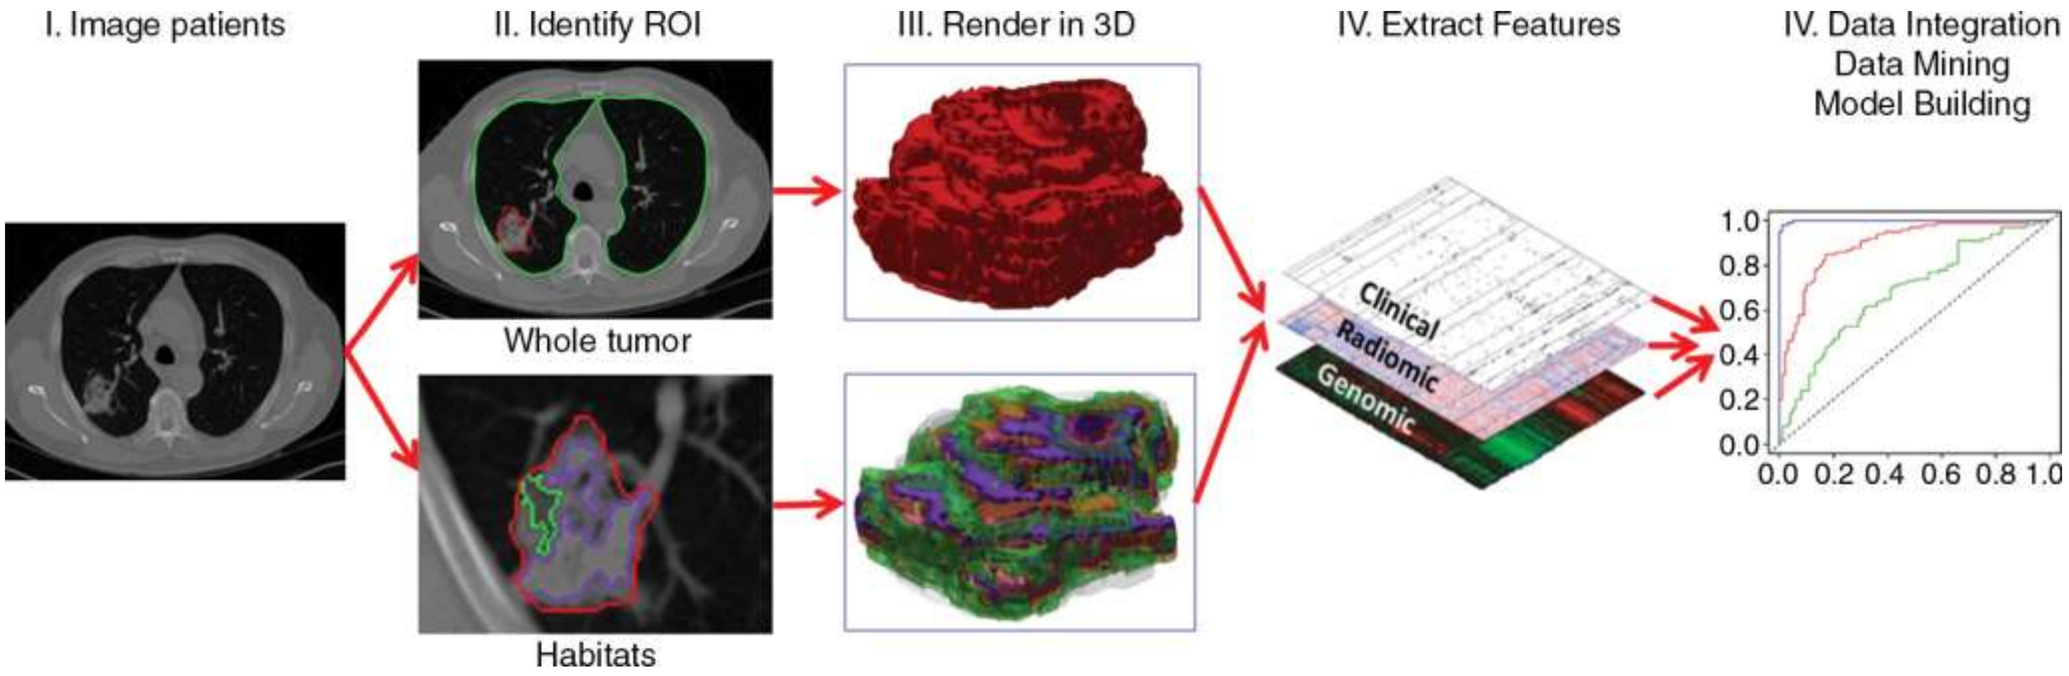

Flowchart shows the process of radiomics and the use of radiomics in decision support.

Gillies, Robert J., et al. "Radiomics: images are more than pictures, they are data." Radiology 278.2 (2016): 563-

Prediction of Progression-Free Survival

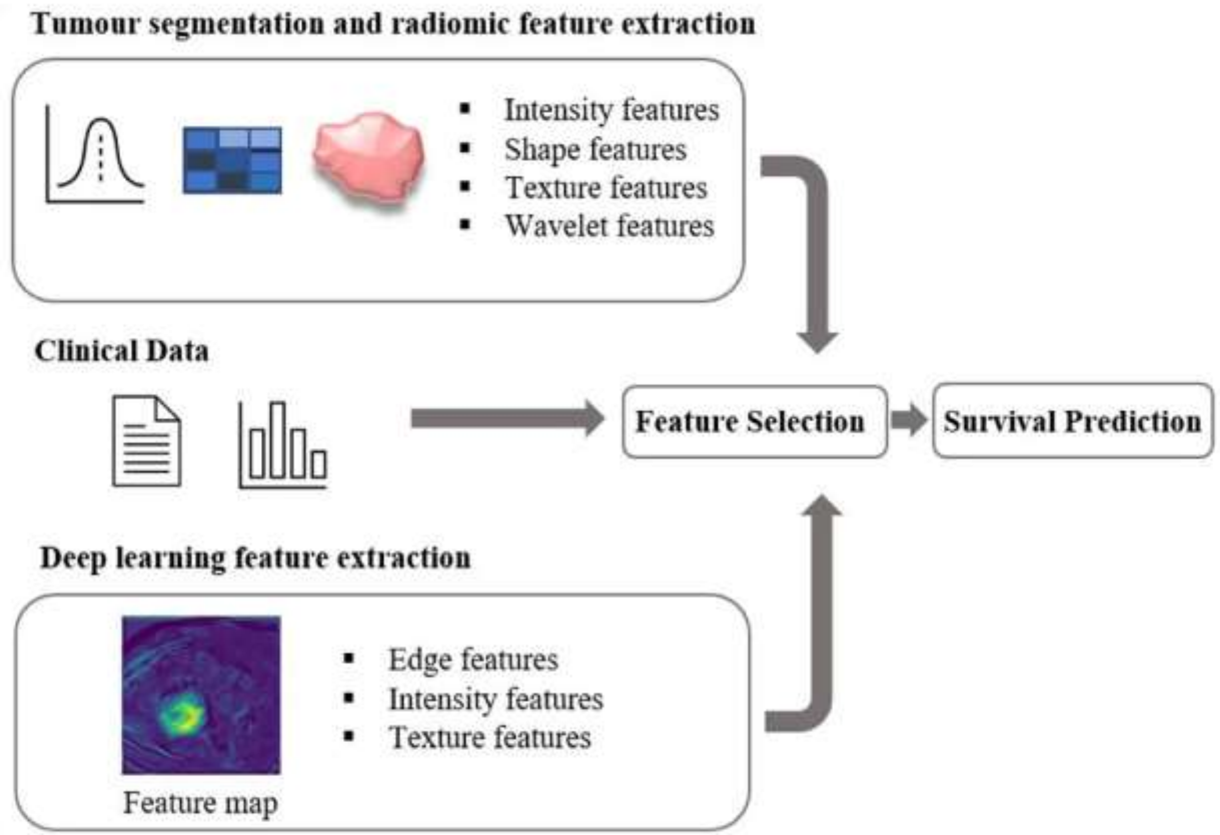

- Survival task
  - Perform tumor segmentation
  - Extract features (e.g. combination of clinical, CT radiomics, and deep learning features)
  - Train model for survival prediction (e.g. Cox proportional hazard regression)

E. Bourigault et al. Multimodal PET/CT tumour segmentation and patient survival prediction using a full-scale UNet with attention (2021)

Exemplar networks

1998 LeCun et al.

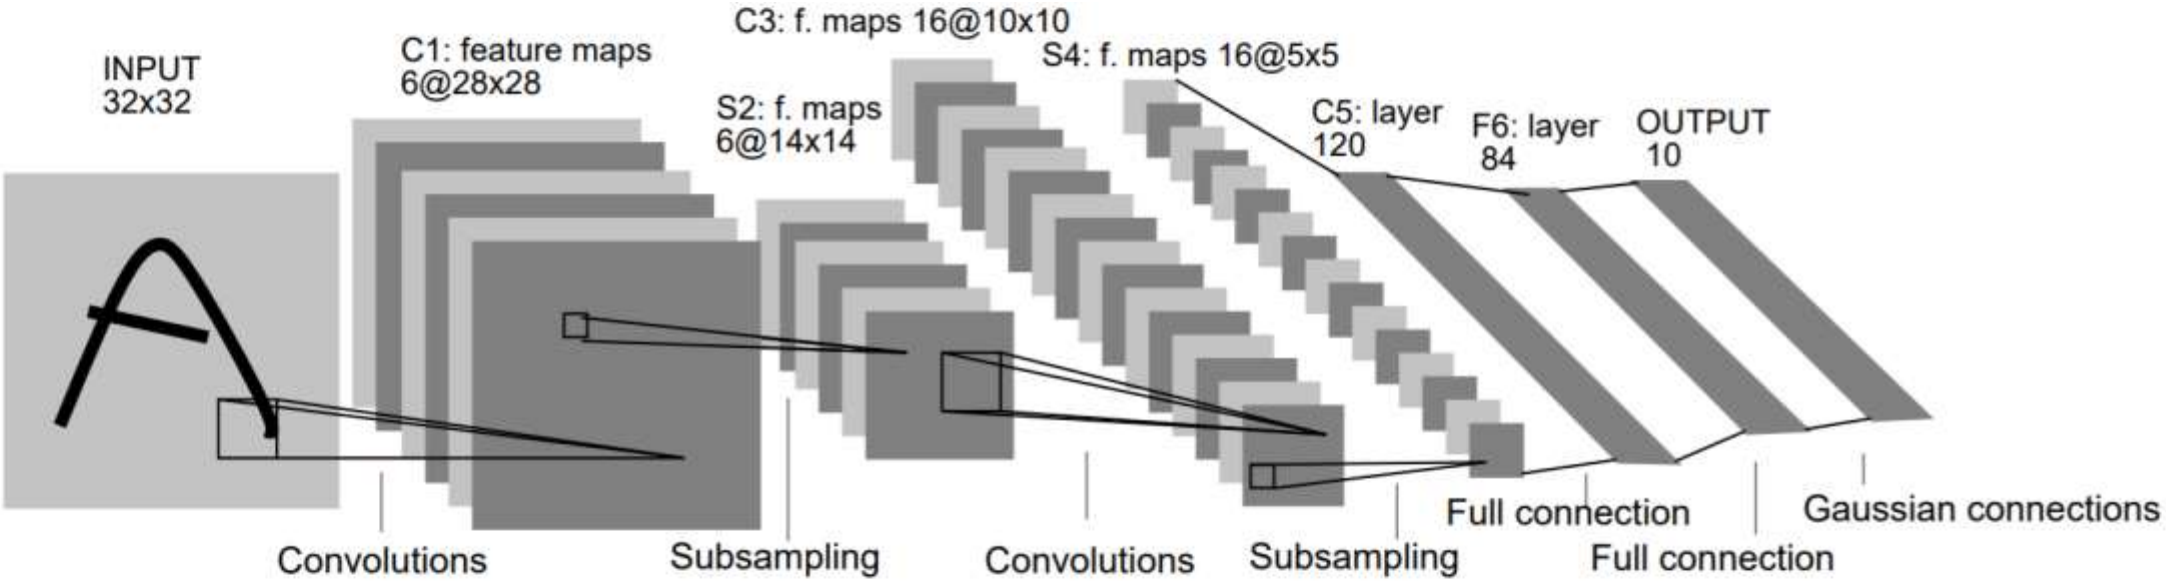

- LeNet 1 (1988): ~3,000 parameters
- LeNet 6 (1998): ~60,000 parameters


PROC. OF THE IEEE, NOVEMBER 1998

Gradient-Based Learning Applied to Document Recognition

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner

2012 Krizhevsky et al. (AlexNet)

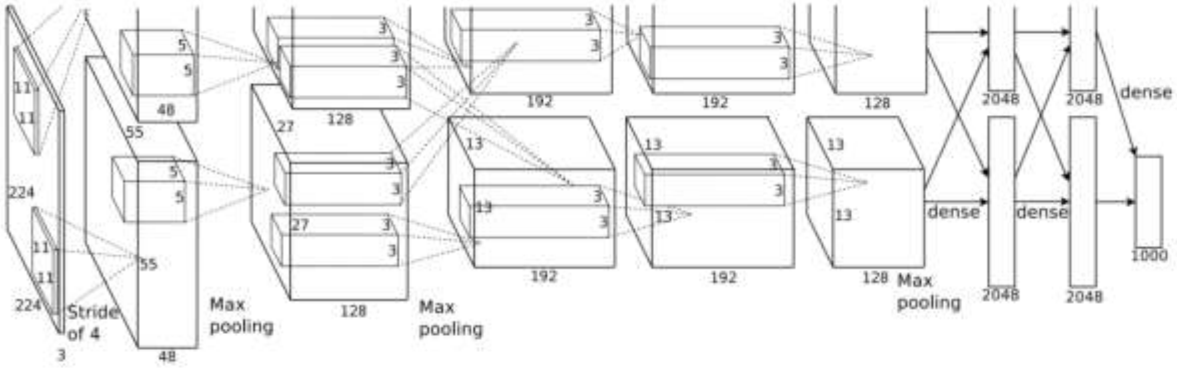

- 60 million parameters
- *‚a large, deep convolutional neural network is capable of achieving record breaking result'*
- *‚In the end, the network's size is limited mainly by the amount of memory available on current GPUs and by the amount of training time that we are willing to tolerate. Our network takes between five and six days to train on two GTX 580 3GB GPUs.'*

ImageNet

- ImageNet is an image database organized according to the WordNet hierarchy (currently only the nouns), in which each node of the hierarchy is depicted by hundreds and thousands of images. The project has been instrumental in advancing computer vision and deep learning research. The data is available for free to researchers for non-commercial use.
- WordNet® is a large lexical database of English. Nouns, verbs, adjectives and adverbs are grouped into sets of cognitive synonyms (synsets), each expressing a distinct concept.
- ImageNet consists of 14,197,122 images organized into 21,841 subcategories. These subcategories can be considered as sub-trees of 27 high-level categories

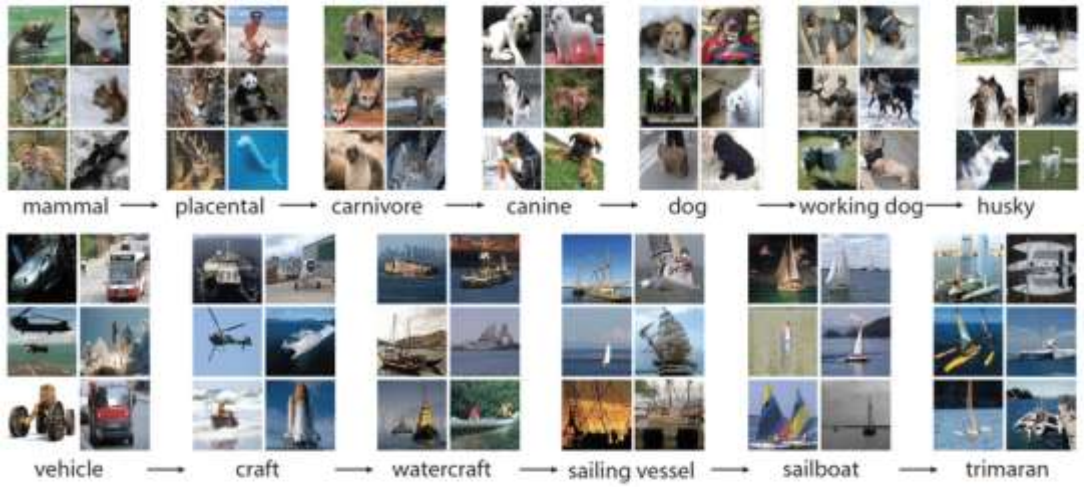

DL - AlexNet and ImageNet

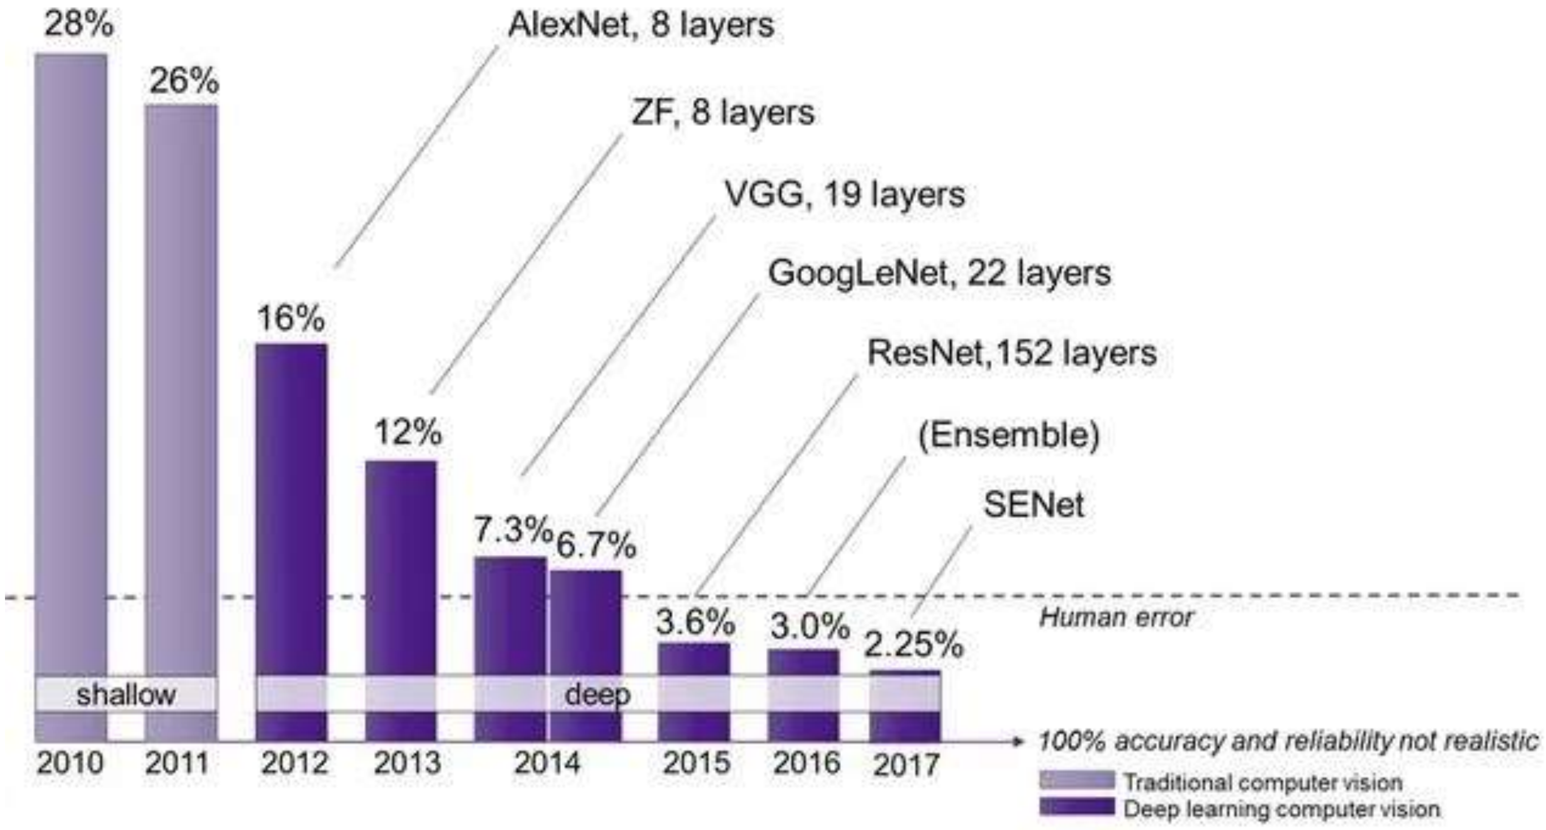

From LeNet to VGG

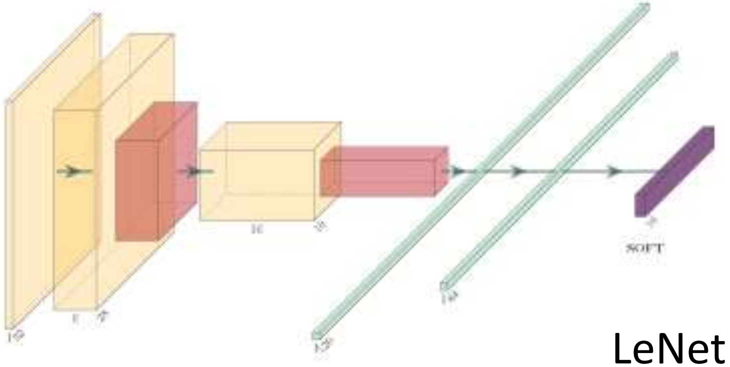

LeNet

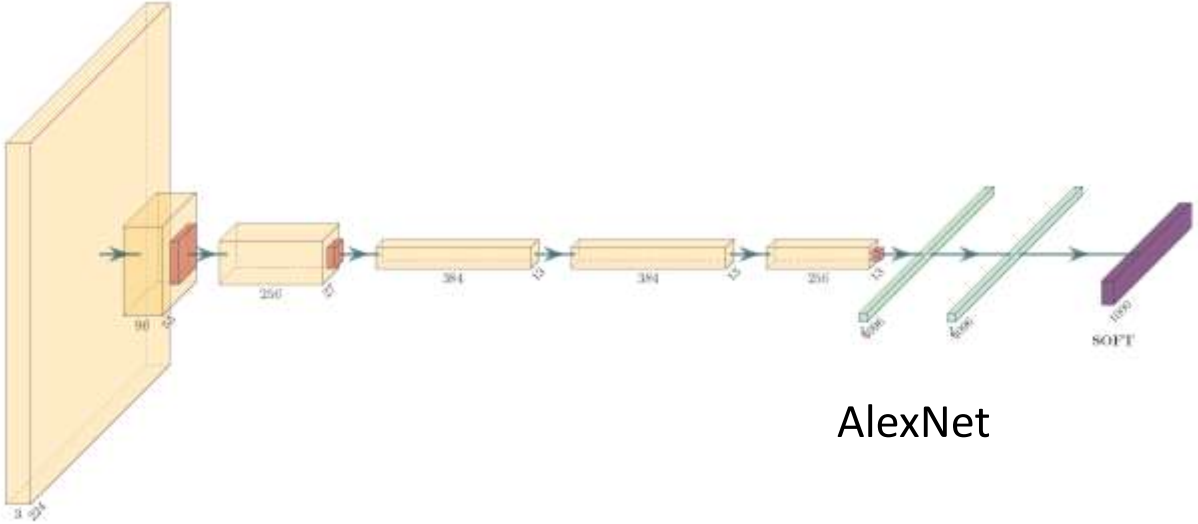

AlexNet

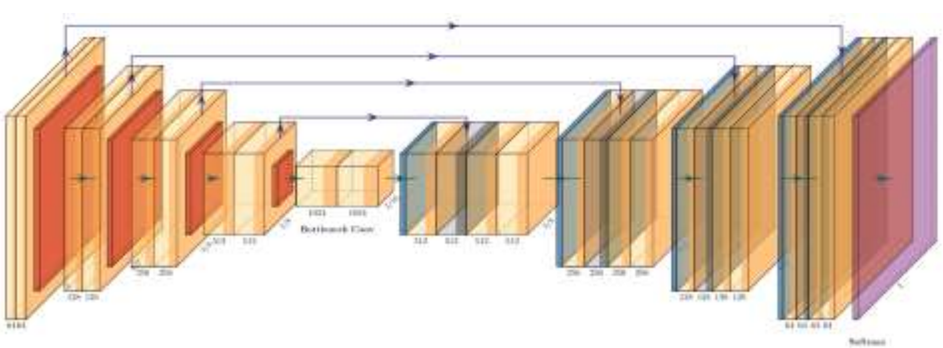

UNet

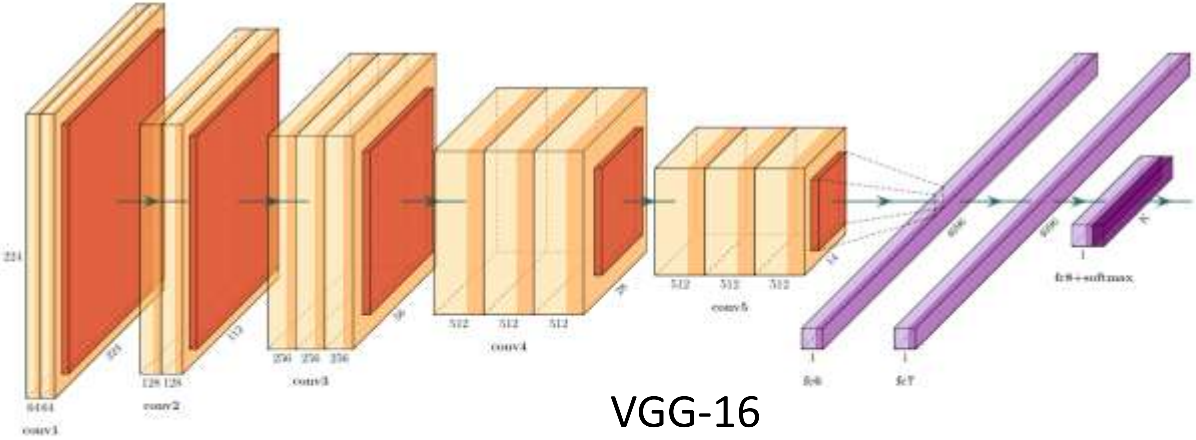

VGG-16

https://github.com/HarisIqbal88/PlotNeuralNet

Recent Developments

- **Generative Pre-trained Transformer 3** (**GPT-3**), a language model, that ~~holds~~ held the record for being the largest neural network ever created with 175 billion parameters
- **Megatron-Turing Natural Language Generation model (MT-NLG),** is the largest and the most powerful monolithic transformer English language model with 530 billion parameters.

Recent Developments

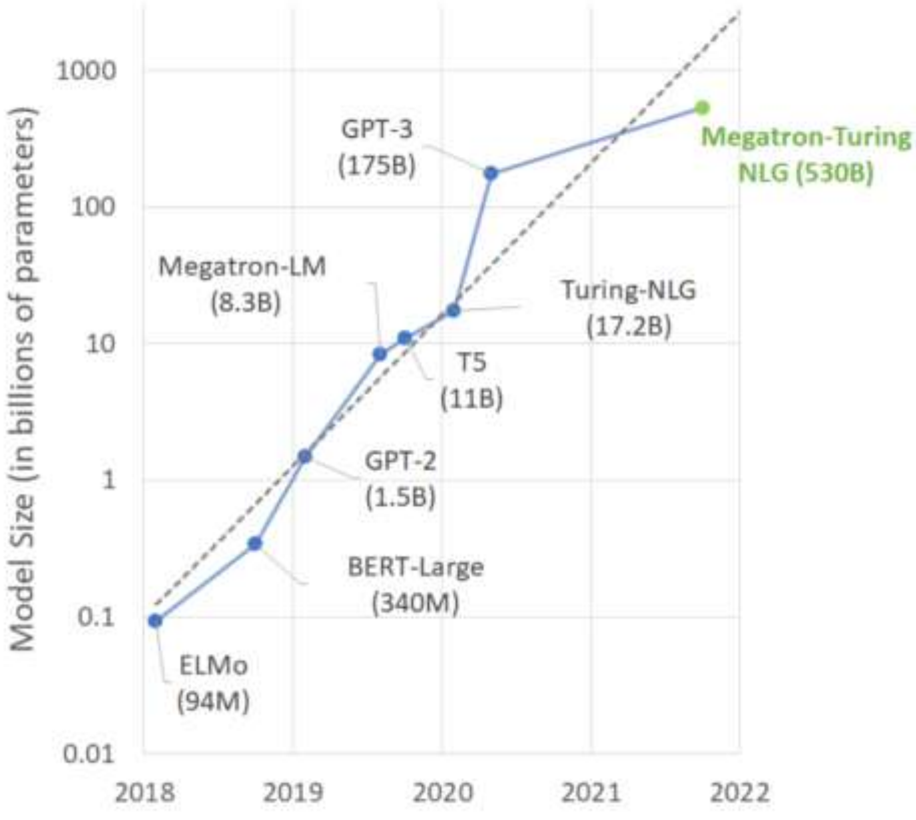

Figure 1: Trend of sizes of state-of-the-art NLP models with time.

- NVIDIA's Selene supercomputer with 560 DGX A100 nodes
- Each cluster node has 8 NVIDIA 80-GB A100 GPUs

- On average, the human brain contains about 100 billion neurons.
- Each neuron may be connected to up to 10,000 other neurons, passing signals to each other via as many as 1,000 trillion synapses

https://arxiv.org/pdf/2201.11990.pdf

DL and Medical Imaging

- Dermatologist-level classification of skin cancer with deep neural networks

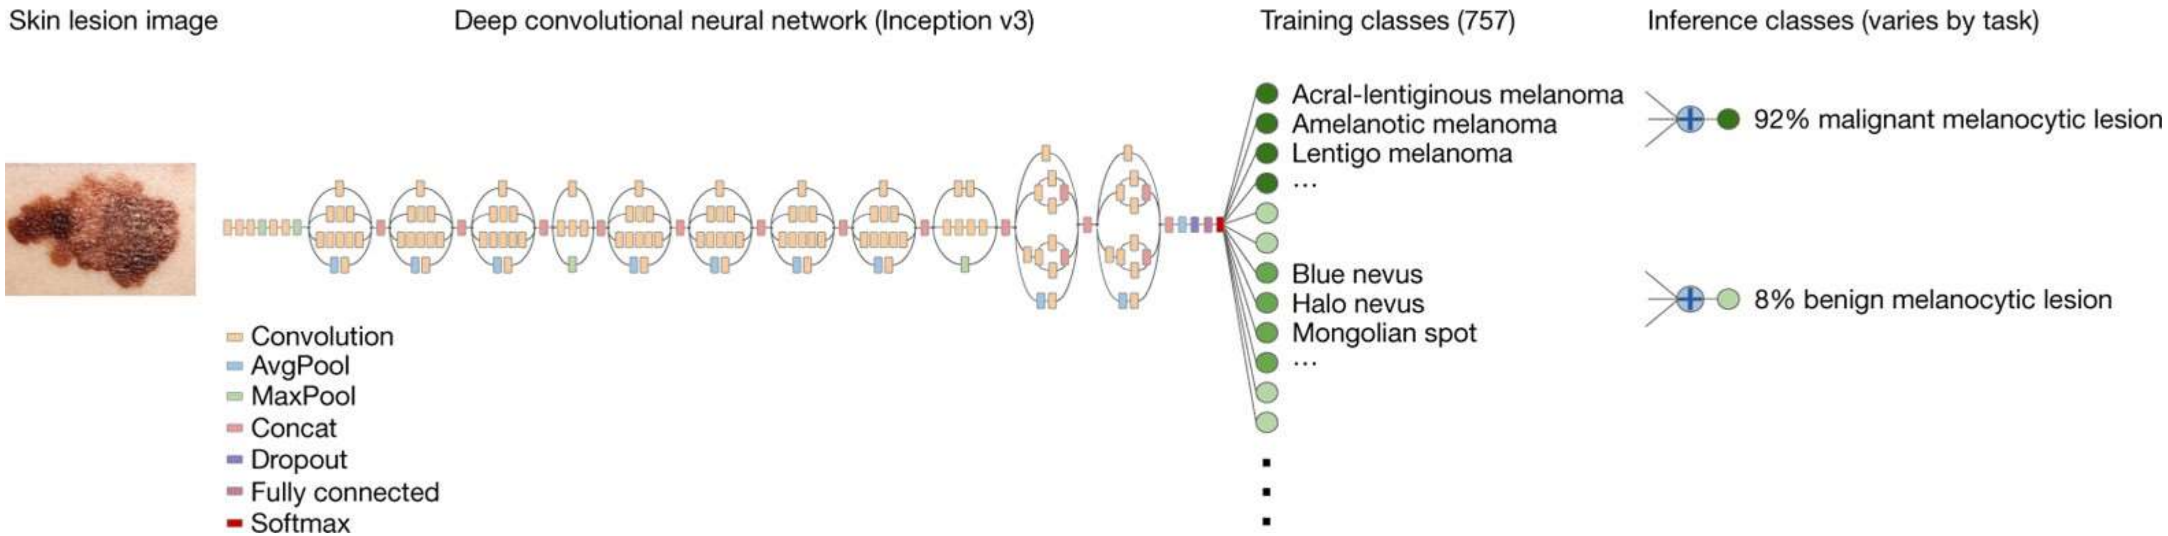

https://www.nature.com/articles/nature21056

DL and Medical Imaging

- Clinically applicable deep learning for diagnosis and referral in retinal disease

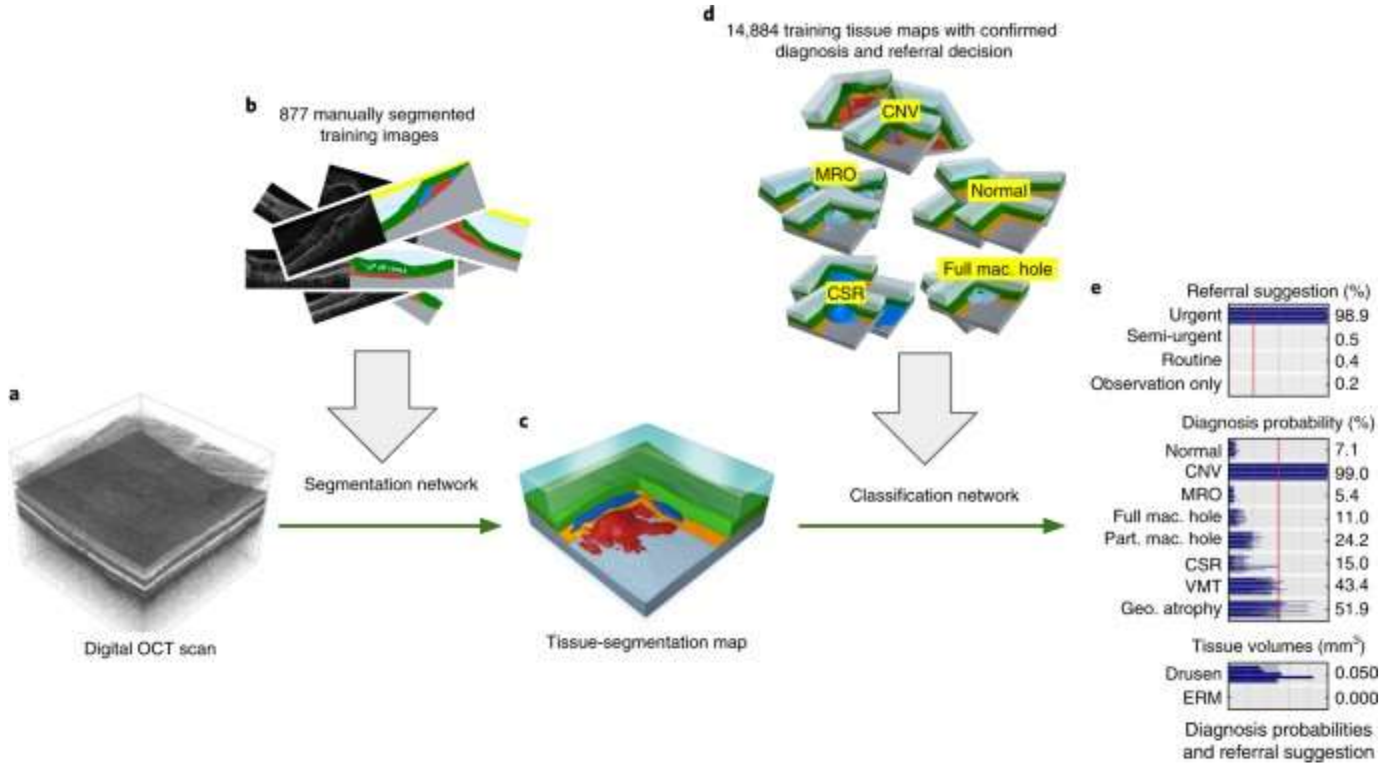

https://www.nature.com/articles/s41591-018-0107-6

DL and Medical Imaging

- CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning

https://arxiv.org/abs/1711.05225

Will AI take over?

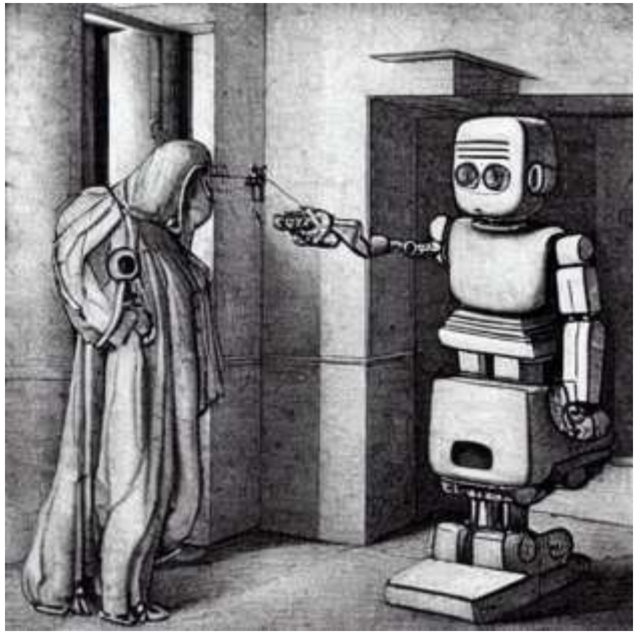

Robot with Artificial Intelligence working at hospital [by Leaonardo da Vinci]

Robot with Artificial Intelligence working at hospital [by van Gogh]

Robot with Artificial Intelligence working at hospital [by Beksiński]

https://huggingface.co/spaces/stabilityai/stable-diffusion

Checkpoint. COVID-19 pandemic...

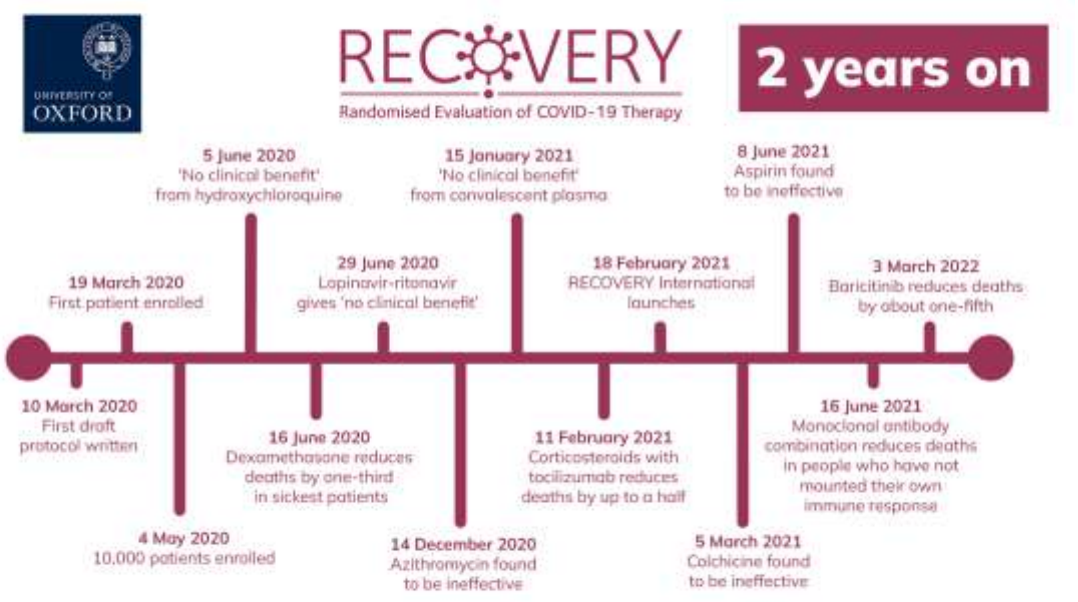

https://www.recoverytrial.net/

https://www.research.ox.ac.uk/area/coronavirus-research/vaccine

Has AI failed us?

- „Conventional data analysis has been at the heart of the COVID-19 response, not AI”
- „There were certainly innovative applications of AI during the pandemic, but the evidence suggests that more traditional methods of data collection and analysis were far more widespread ”
- „(...) in a healthcare setting, artificial intelligence did not play the outsized role many thought it would in relief efforts.”

Centre for Data Ethics and Innovation Blog

Organisations: Centre for Data Ethics and Innovation

Reflecting on the use of AI and data-driven technology in the pandemic

Has AI in medical imaging failed us too?

- „Despite the huge efforts of researchers to develop machine learning models for COVID-19 diagnosis and prognosis, we found methodological flaws and many biases throughout the literature, leading to highly optimistic reported performance.”

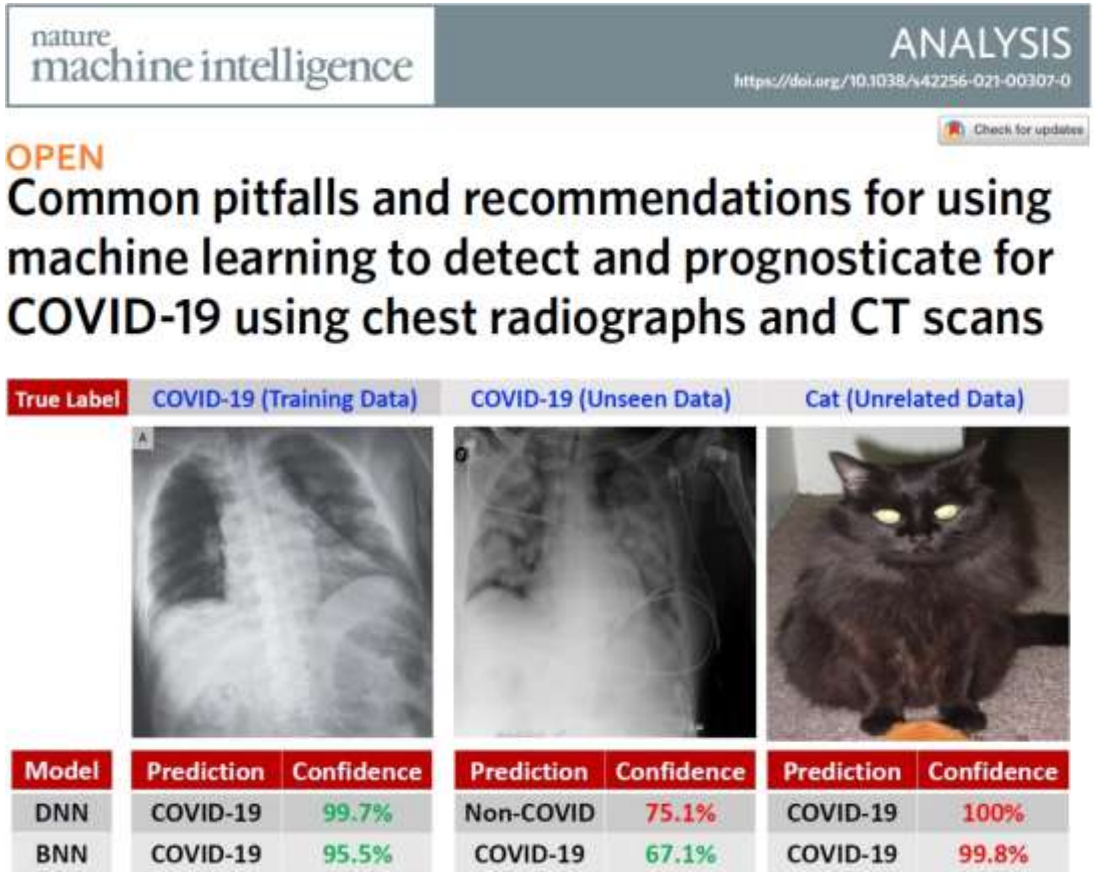

Mallick, Ankur, et al. "Probabilistic Neighbourhood Component Analysis: Sample Efficient Uncertainty Estimation in Deep Learning." *arXiv:2007.10800* (2020).

Challenges for Medical Imaging in Healthcare

- Data (access, standarisation, biases, inequality, exclusion, etc)
- Imaging repositories require enourmous (and costly) data curation before can be reliably used.

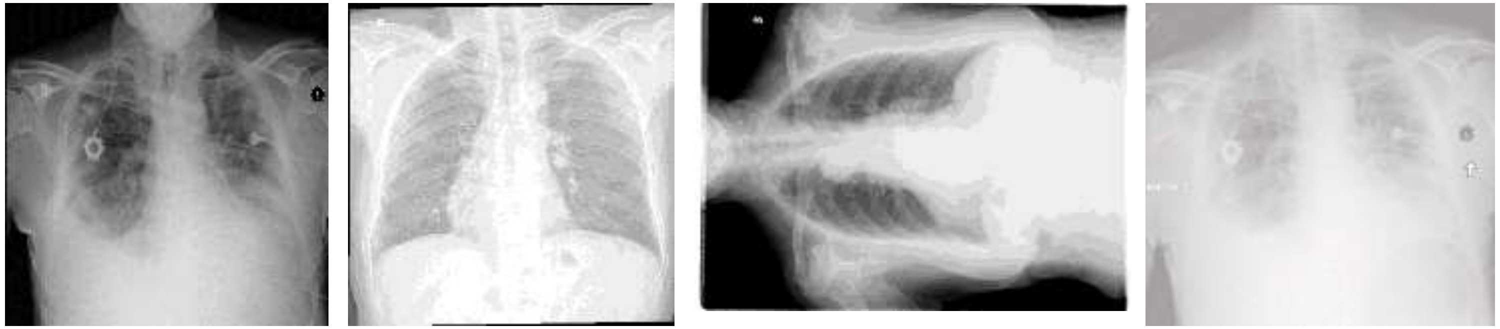

Randomly extracted chest radiographs (x-ray imaging) from hospital repository

How to annotate data?

Picture courtesy Prof. W. Sadłoń

- Let’s identify and annotate the Tatra chamois...
- Quite easy task even for non-experts
- But
  - We know context (this picture was really taken in the Tatra Mountains)

Challenges for Medical Imaging in Healthcare

- Clinical annotations... (intra-observer variability)

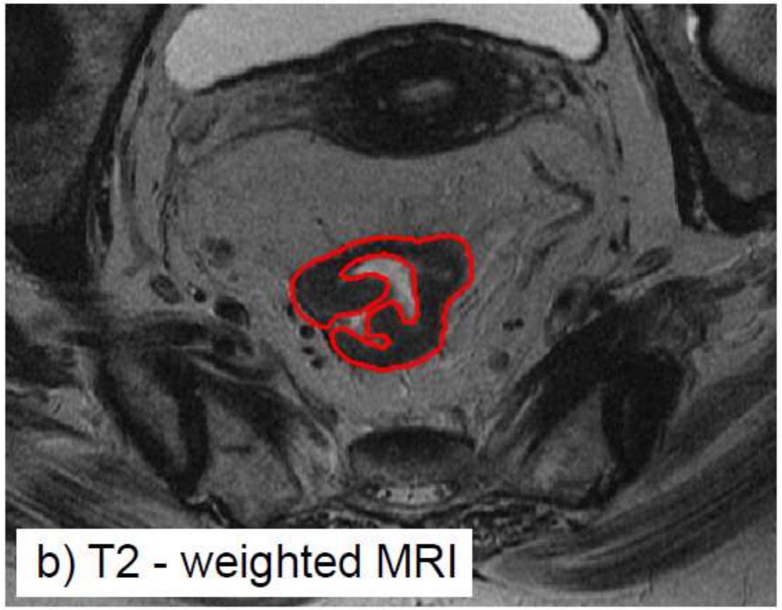

Annotation of colorectal tumour
(Expert 1)

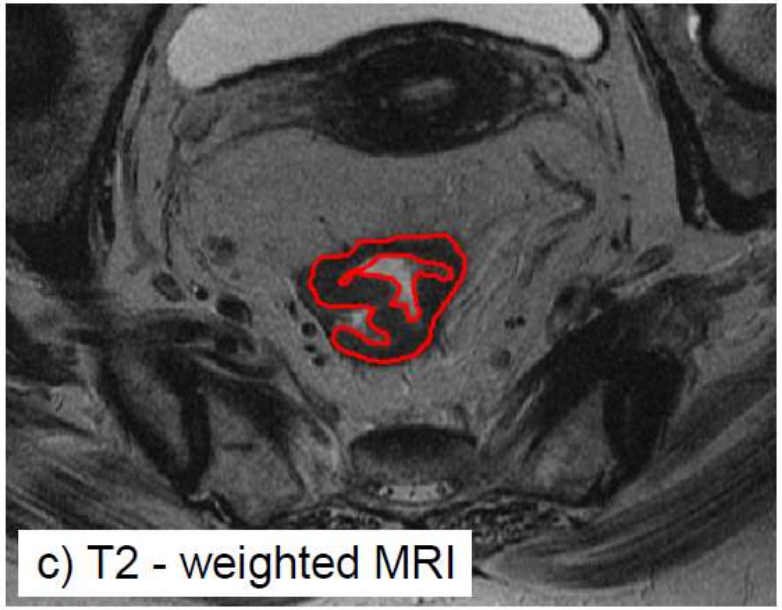

Annotation of colorectal tumour
(Expert 1 – week later)

Irving, Benjamin, et al. "Automated colorectal tumour segmentation in DCE-MRI using supervoxel neighbourhood contrast characteristics 2014.

Challenges for Medical Imaging in Healthcare

- More clinical annotations... (inter-observer variability)

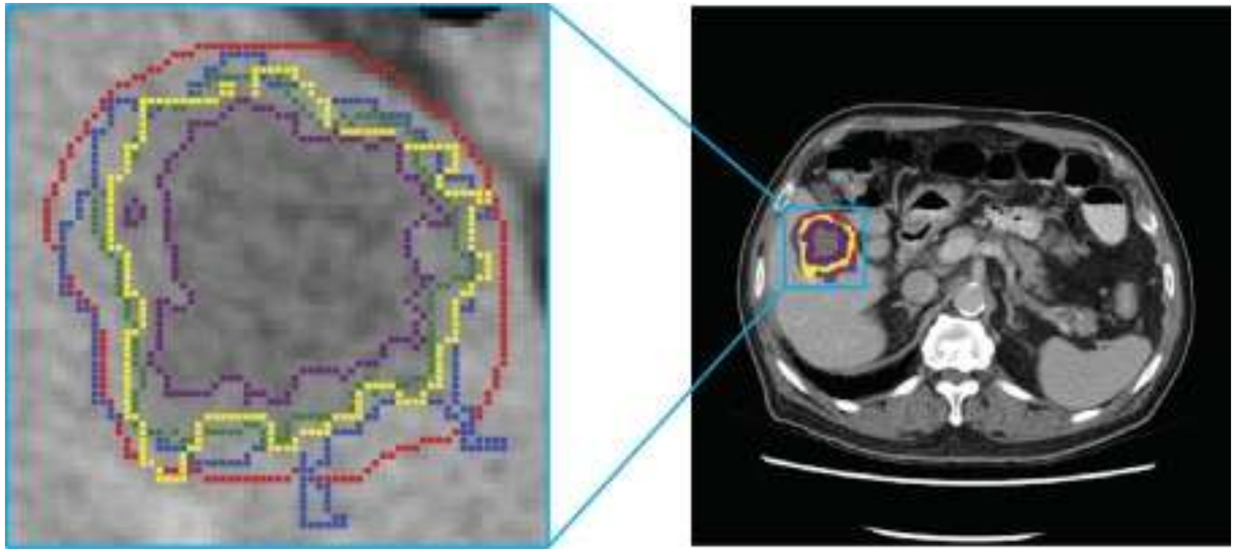

Manual segmentations of multiple observers of a colorectal liver metastasis on an axial slice of a CT scan

Martijn P.A. Starmans, Wiro J. Niessen, in Handbook of Medical Image Computing and Computer Assisted Intervention (2020)

Challenges for Medical Imaging in Healthcare

- Only a small fraction of the information contained in the imaging scan is used to assess medical condition (and health outcome)
- Information extracted from the imaging scan is often not quantifiable (+ human interpretation)

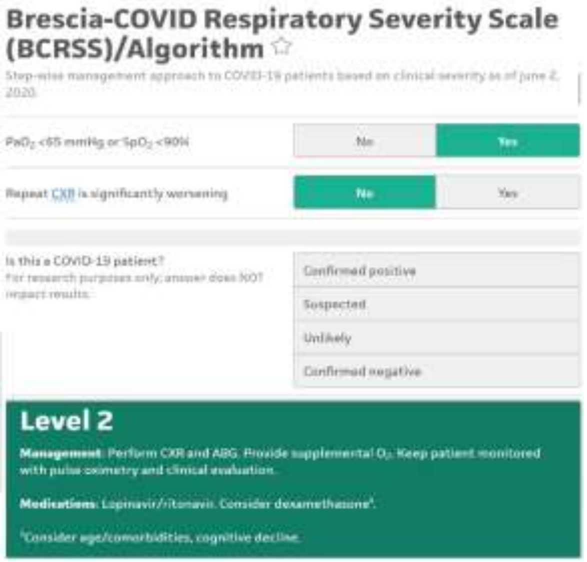

Brescia-COVID Respiratory Severity Scale (BCRSS)/Algorithm

Step-wise management approach to COVID-19 patients based on clinical severity as of June 2, 2020.

PaO_2 <65 mmHg or SpO_2 <90%	No	Yes
Repeat CXR is significantly worsening	No	Yes

Is this a COVID-19 patient?
For research purposes only; answer does NOT impact results.

- Confirmed positive
- Suspected
- Unlikely
- Confirmed negative

Level 2

Management: Perform CXR and ABG. Provide supplemental O_2. Keep patient monitored with pulse oximetry and clinical evaluation.

Medications: Lopinavir/ritonavir. Consider dexamethasone[1].

[1]Consider age/comorbidities, cognitive decline.

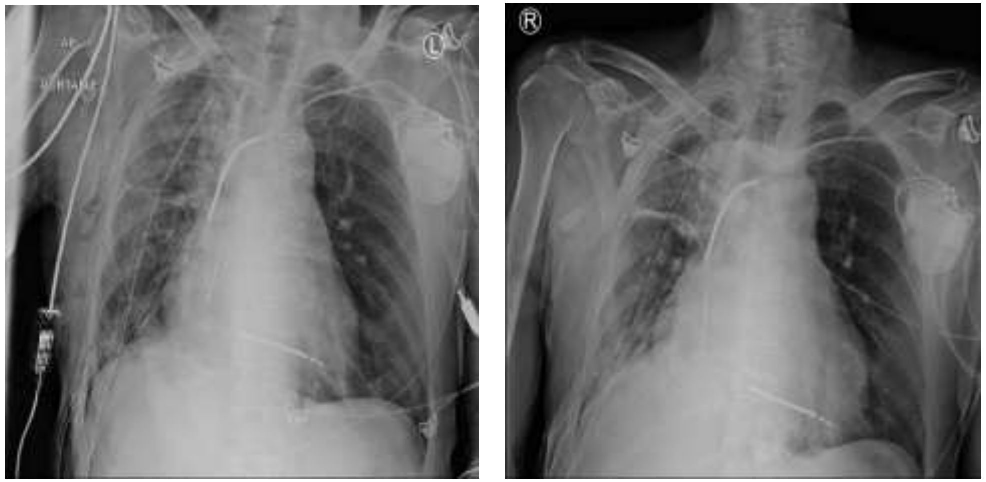

Chest radiograph (x-ray imaging) taken few days apart

Challenges of Medical Imaging in Healthcare

- AI algorithms ~~are~~ can be „black-box"

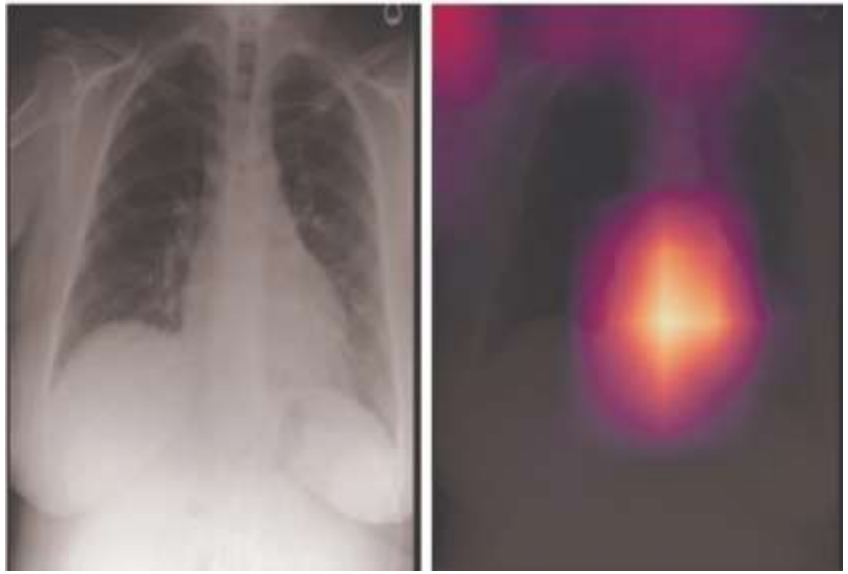

Example of saliency map from the image-based ResNet-50 model for true positive instance of cardiomegaly.

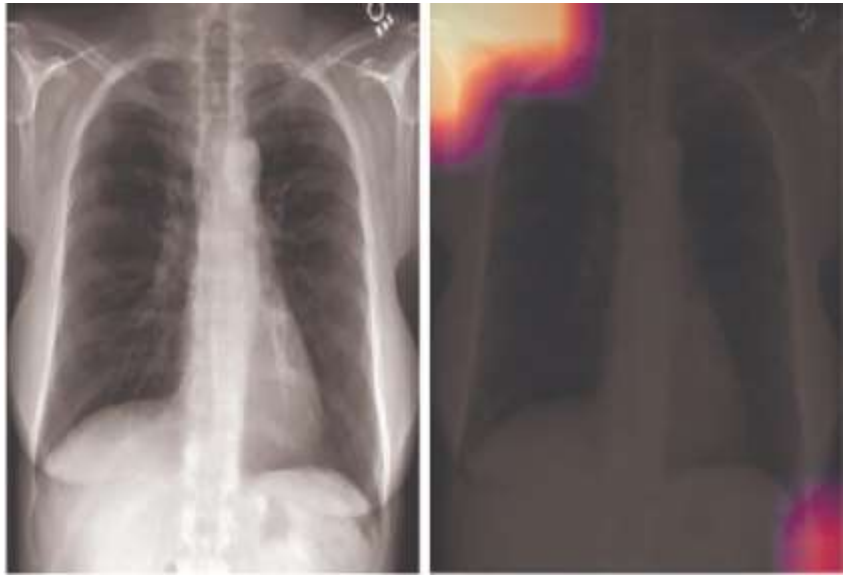

Example of saliency map from the image-based ResNet-50 model for true negative instance of cardiomegaly.

D. Grant, B.W. Papiez, et al. Deep learning classification of cardiomegaly using combined imaging and non-imaging ICU data, (2021)

Explainable AI for medical imaging

- Would be possible that AI could explain medical decision using Natural Language? Just do as radiologists do!

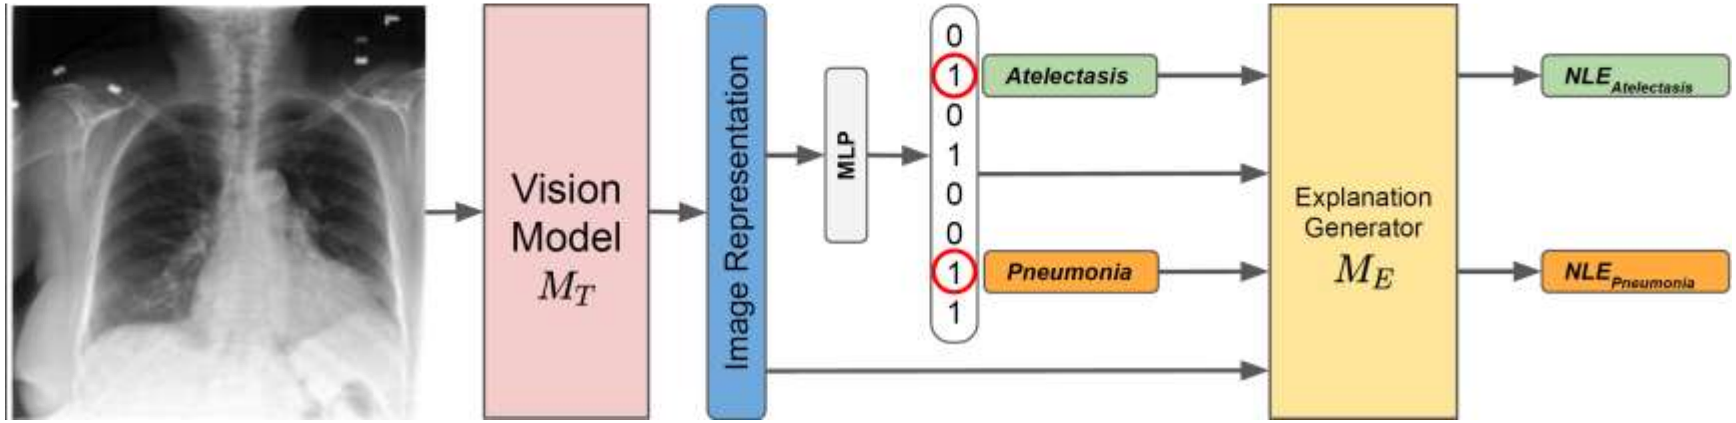

- The Vision Model MT provides multi-label classification across negative, uncertain, positive for 14 different radiological findings
- The Explanation Generator ME generates Natural Language Explanations (NLE)

Kayser, Maxime, et al. "Explaining Chest X-ray Pathologies in Natural Language." *MICCAI* (2022).

Explainable AI for medical imaging

- How to extract NLEs from raw radiology reports?

- We first extract the Findings and Impression sections, which contain the descriptive part of the report. Next, we identify the labels referred to in each sentence and the sentences that contain

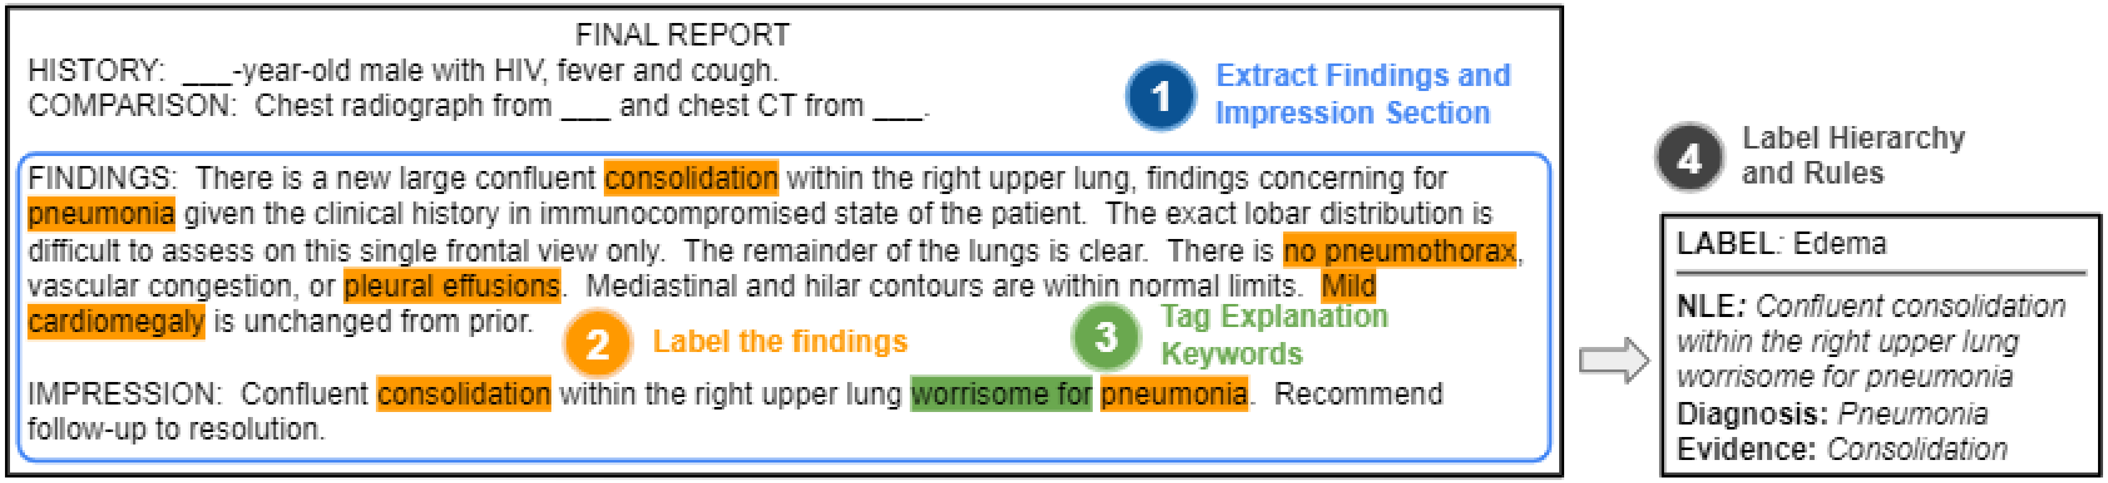

Kayser, Maxime, et al. "Explaining Chest X-ray Pathologies in Natural Language." *MICCAI* (2022).

Explainable AI for medical imaging

LABELS: Edema (Positive)

Natural Language Explanations for ***Edema:***	Clinical Evaluation:
Ground-Truth: *Indistinct appearance of the pulmonary vasculature is compatible with pulmonary edema.*	2
RATCHET: *Findings suggesting mild pulmonary edema.*	1
DPT: *Pulmonary edema and extensive bibasilar opacification appear slightly worse.*	3
TieNet: *Diffuse bilateral pulmonary opacities, likely edema.*	5

LABELS: Atelectasis (Uncertain) and Pneumonia (Uncertain)

Natural Language Explanations for ***Pneumonia:***	Clinical Evaluation:
Ground-Truth: *Interval appearance of patchy opacity at the left base could represent early pneumonia, although aspiration or patchy atelectasis would also be in the differential.*	3
RATCHET: *Patchy opacities in the lung bases may reflect atelectasis, but infection is not excluded in the correct clinical setting.*	2
DPT: *Peribronchial deformities in the right apex and elevation of the right hemidiaphragm are likely due to scarring No acute osseous disease or infection.*	1
TieNet: *Patchy bibasilar airspace opacities, likely atelectasis.*	2

The clinical evaluation is given on a Likert scale, where 5 is the highest, and 1 is the lowest score

Kayser, Maxime, et al. "Explaining Chest X-ray Pathologies in Natural Language." *MICCAI* (2022). Data set and code available here: https://github.com/maximek3/MIMIC-NLE

To AI, or not to AI?

https://huggingface.co/spaces/stabilityai/stable-diffusion

Resources

- Free books:

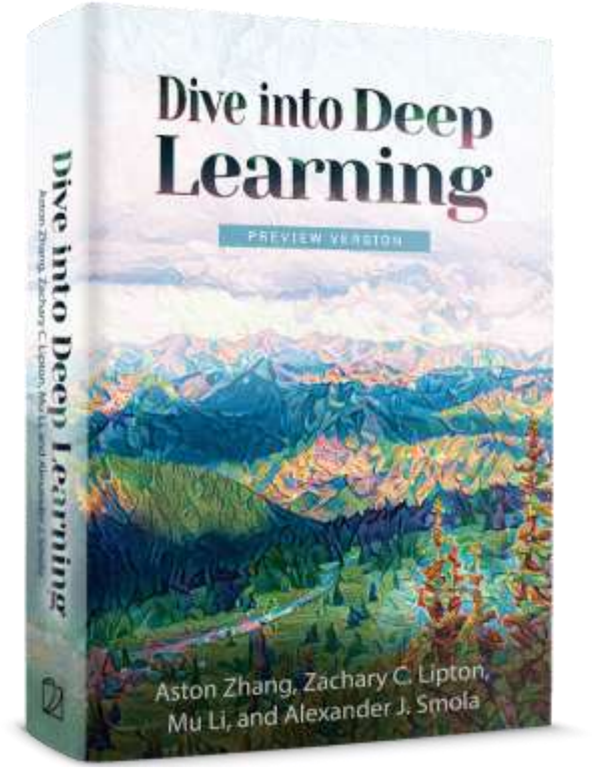

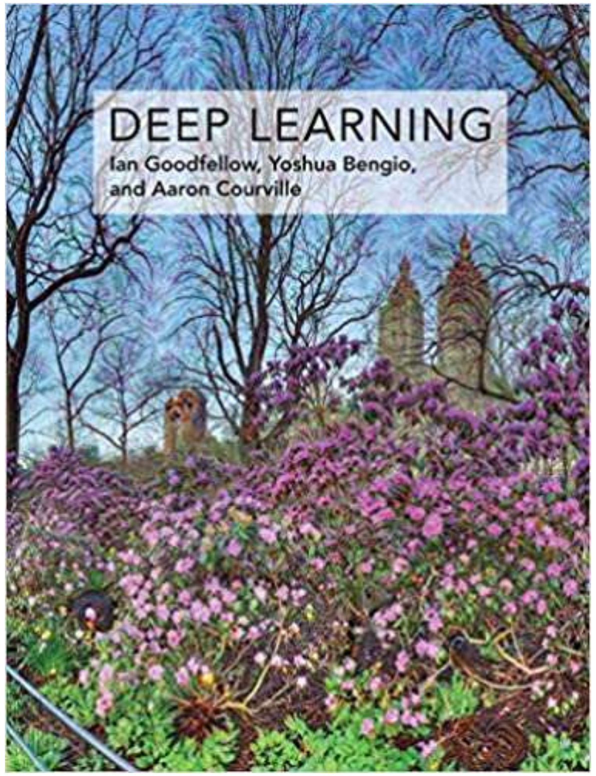

http://d2l.ai/index.html

https://www.deeplearningbook.org/

Questions?

Q&A