Brief History of Recent Deep Learning

AI for ecologists

Paul Tresson

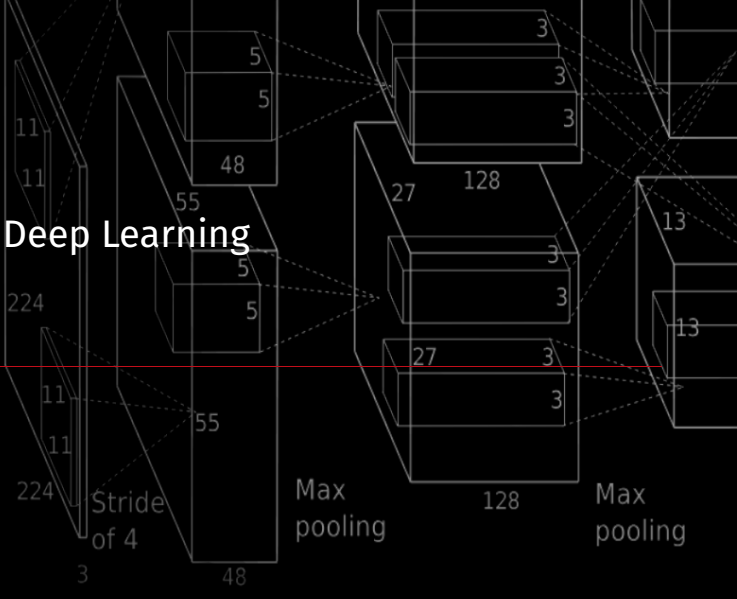

Institut de Recherche pour le Développement FRANCE
AMAPlab
LIRMM
FRB
CESAB

Reminder

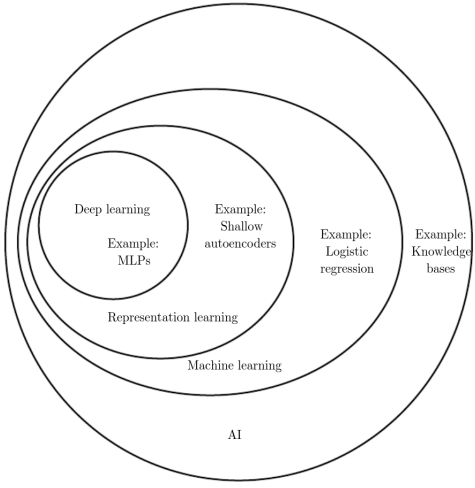

Different types of Artificial Intelligence. From Goodfellow et al., 2016

First Neural Network Model : The Perceptron

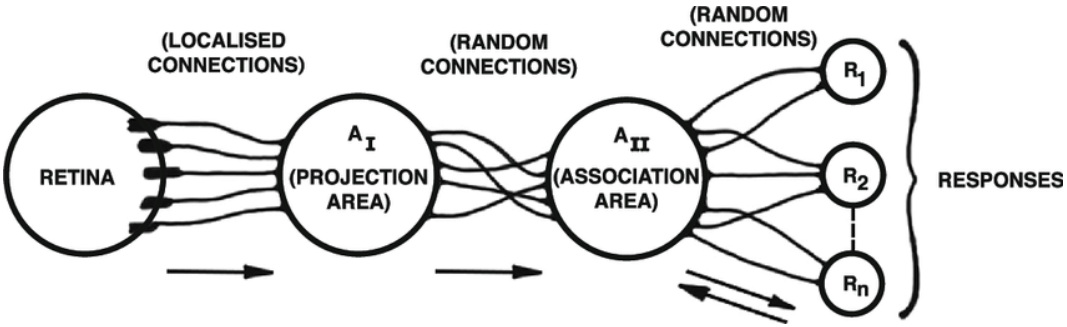

First probabilistic model of a neuron proposed by Rosenblatt, 1958

First Neural Network Model : The Perceptron

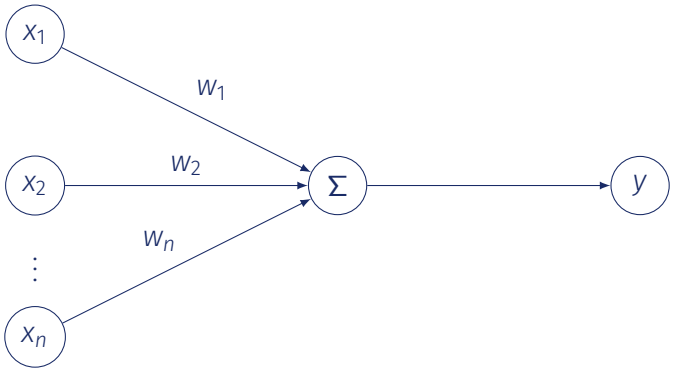

First Neural Network Model : The Perceptron

$$y = \begin{bmatrix} w_1 & w_2 & \dots & w_n \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$$

First Neural Network Model : The Perceptron

$$y = \begin{bmatrix} w_1 & w_2 & \dots & w_n \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}$$

First Roadblock : XOR problem

x_1	x_2	y
0	0	0
1	0	1
0	1	1
1	1	0

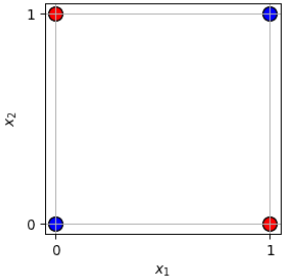

First Roadblock : XOR problem

x_1	x_2	y
0	0	0
1	0	1
0	1	1
1	1	0

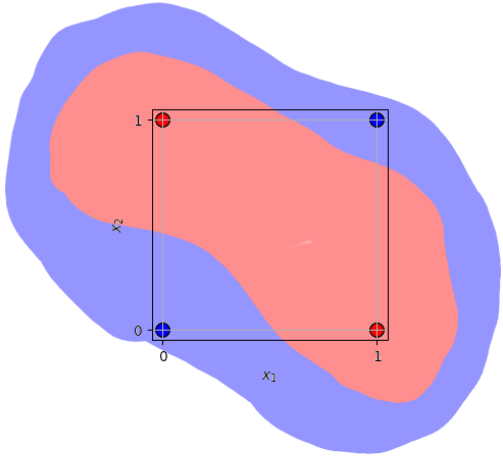

1980's : First Deep Neural Networks

- Non-linearity (*e.g.* ReLu)

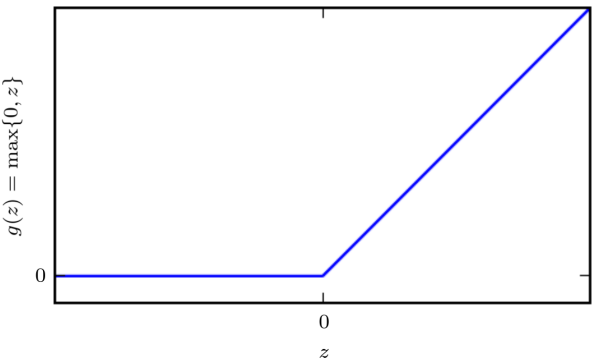

Rectified Linear Unit. Figure from Goodfellow et al., 2016

1980's : First Deep Neural Networks

- Non-linearity (*e.g.* ReLu)
- Multi Layer Perceptron

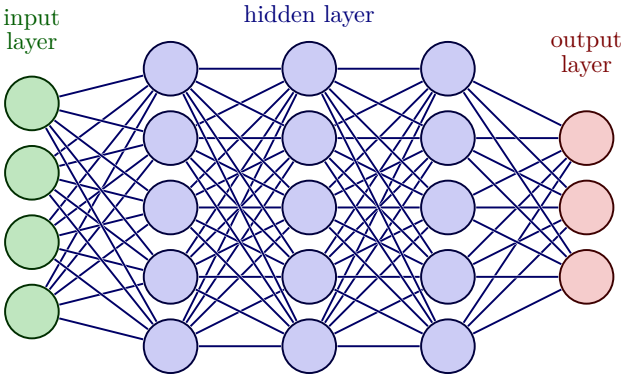

Multi Layer Perceptron

1980's : First Deep Neural Networks

- Non-linearity (*e.g.* ReLu)
- Multi Layer Perceptron
- Gradient Descent and Backpropagation
  Rumelhart, 1986

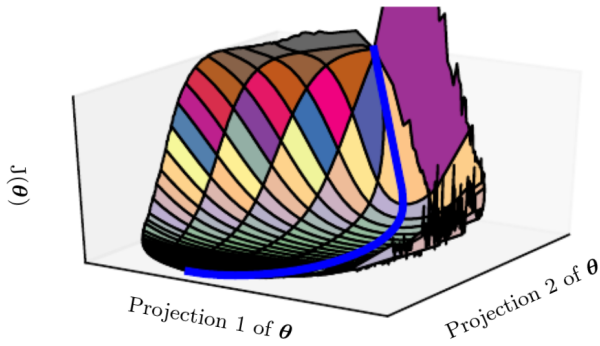

Gradient descent example. Figure from Goodfellow et al., 2016

As matrices

$$y = ReLU\left(\begin{bmatrix} w_{1,1} & w_{1,2} & \dots & w_{1,n} \\ w_{2,1} & w_{2,2} & \dots & w_{2,n} \\ \vdots & & & \vdots \\ w_{k,1} & w_{k,2} & \dots & w_{k,n} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}\right)$$

1990's : First Successes

- MNIST
  LeCun, 1998

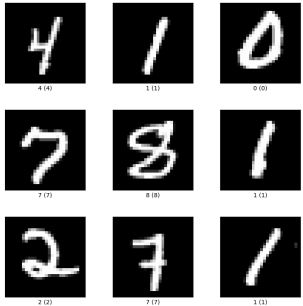

Images from MNIST

Modified National Institute of Standards and Technology

1990's : First Successes

- MNIST
  LeCun, 1998
- LeNet
  LeCun et al., 1989b

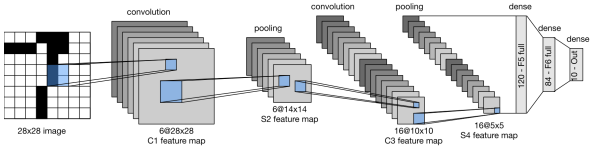

LeNet5.
Figure from Zhang et al. - `https://github.com/d2l-ai/d2l-en`

Modified National Institute of Standards and Technology

1990's : First Successes

- MNIST
  LeCun, 1998
- LeNet
  LeCun et al., 1989b
- Convolutional Neural Network (CNN)
  LeCun et al., 1989a

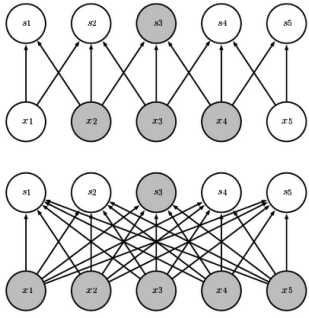

sparse connectivity with CNNs.
Figure from Goodfellow et al., 2016

Good proof of concept but too costly in computing power and datasets

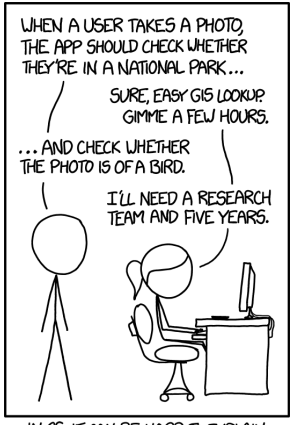

XKCD comic from 2014

Good proof of concept but too costly in computing power and datasets ...until

Good proof of concept but too costly in computing power and datasets ...until

Good proof of concept but too costly in computing power and datasets ...until

2010's : CNN revolution

- ImageNet
  Deng et al., 2009

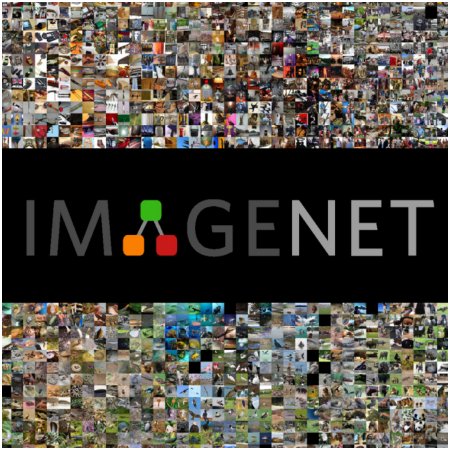

2010's : CNN revolution

- ImageNet
  Deng et al., 2009
- AlexNet
  Krizhevsky et al., 2012

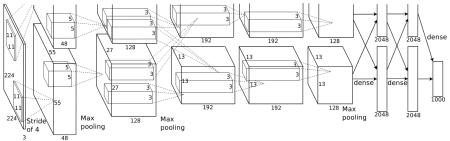

AlexNet. Figure from Krizhevsky et al., 2012

2010's : CNN revolution

- ImageNet
  Deng et al., 2009
- AlexNet
  Krizhevsky et al., 2012
- ResNet
  He et al., 2016

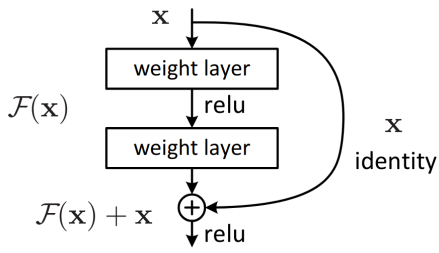

A residual connection. Figure from He et al., 2016

2010's : CNN revolution

- ImageNet
  Deng et al., 2009
- AlexNet
  Krizhevsky et al., 2012
- ResNet
  He et al., 2016

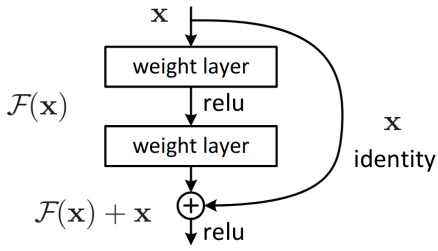

A residual connection. Figure from He et al., 2016

His 2016 paper *Deep Residual Learning for Image Recognition* is the most cited research paper in 5 years according to Google Scholar's reports in 2020 and 2021.[7][8]

2010's : CNN revolution

- ImageNet
  Deng et al., 2009
- AlexNet
  Krizhevsky et al., 2012
- ResNet
  He et al., 2016
- YOLO, mask-RCNN...
  He et al., 2017; Redmon et al., 2016

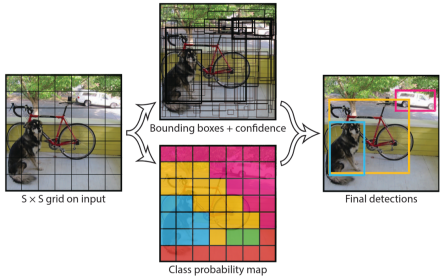

YOLO. Figure from Redmon et al., 2016

2020's : Transformers ...

- Attention is all you need
  Vaswani et al., 2017

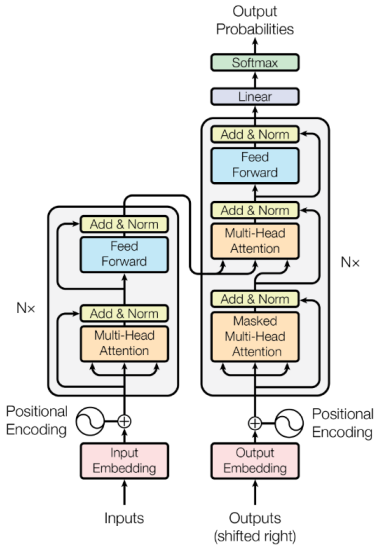

2020's : Transformers ... and Self Supervised Learning

- Attention is all you need
  Vaswani et al., 2017
- LLMs (GPT, BERT...)
  Devlin et al., 2019

2020's : Transformers ... and Self Supervised Learning

- Vision Transformer
  Dosovitskiy et al., 2020

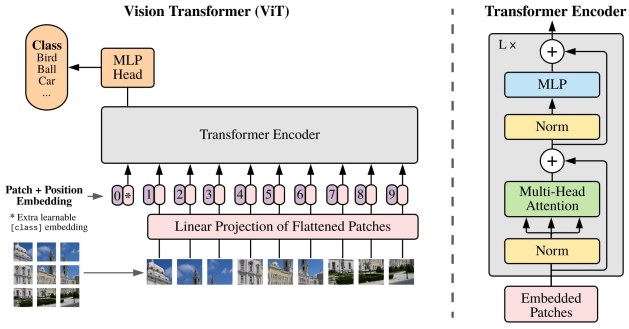

Figure from Dosovitskiy et al., 2020

2020's : Transformers ... and Self Supervised Learning

- Vision Transformer
  Dosovitskiy et al., 2020
- Masked Auto Encoder
  He et al., 2022

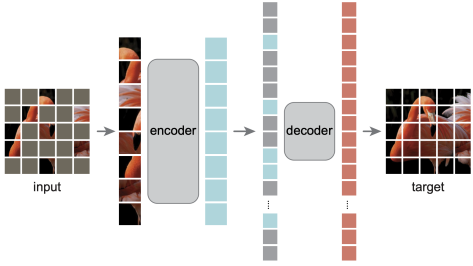

Figure from He et al., 2022

2020's : Transformers ... and Self Supervised Learning

- Vision Transformer
  Dosovitskiy et al., 2020
- Masked Auto Encoder
  He et al., 2022
- DINO
  Caron et al., 2021

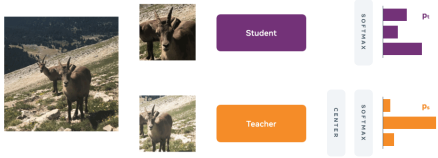

Figure from Caron et al., 2021

A lot left out !

Finetuning

A lot left out !

Finetuning

Games

A lot left out !

Reinforcement learning

Finetuning

Games

A lot left out !

Reinforcement learning

Finetuning

Games

Time series

A lot left out !

Reinforcement learning

Finetuning

Games

Time series

Speech recognition

A lot left out !

Reinforcement learning

Finetuning

Games

Image generation

Time series

Speech recognition

A lot left out !

Reinforcement learning

Finetuning

Games

Image generation

Diffusion models

Time series

Speech recognition

A lot left out !

Reinforcement learning

Finetuning

Games

Image generation

Diffusion models

Time series

Speech recognition

...

Useful ressources

State of the art

- Huggingface
- PapersWithCode

Getting started

- Pytorch
- Keras

Understanding papers

- Yannic Kilcher
- AI coffe break

Understanding visually

- 3blue1brown
- deepia

Thanks for you attention !

Let’s practice !

References i

Caron, Mathilde et al. (2021). **“Emerging properties in self-supervised vision transformers”.** In: *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9650–9660.

Deng, Jia et al. (2009). **“Imagenet: A large-scale hierarchical image database”.** In: *2009 IEEE conference on computer vision and pattern recognition*. Ieee, pp. 248–255.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2019). **“Bert: Pre-training of deep bidirectional transformers for language understanding”.** In: *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pp. 4171–4186.

References ii

Dosovitskiy, Alexey et al. (2020). **“An image is worth 16x16 words: Transformers for image recognition at scale”.** In: *arXiv preprint arXiv:2010.11929*.

Goodfellow, Ian, Yoshua Bengio, Aaron Courville, and Yoshua Bengio (2016). ***Deep learning***. Vol. 1. 2. MIT press Cambridge.

He, Kaiming, Georgia Gkioxari, Piotr Dollár, and Ross Girshick (2017). **“Mask r-cnn”.** In: *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun (2016). **“Deep residual learning for image recognition”.** In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.

He, Kaiming et al. (2022). **“Masked autoencoders are scalable vision learners”.** In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009.

References iii

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E Hinton (2012). **"Imagenet classification with deep convolutional neural networks".** In: *Advances in neural information processing systems* 25.

LeCun, Yann et al. (1989a). **"Generalization and network design strategies".** In: *Connectionism in perspective* 19.143-155, p. 18.

LeCun, Yann (1998). **"The MNIST database of handwritten digits".** In: *http://yann.lecun.com/exdb/mnist/*.

LeCun, Yann et al. (1989b). **"Backpropagation applied to handwritten zip code recognition".** In: *Neural computation* 1.4, pp. 541–551.

Redmon, Joseph, Santosh Divvala, Ross Girshick, and Ali Farhadi (2016). **"You only look once: Unified, real-time object detection".** In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788.

References iv

Rosenblatt, Frank (1958). **“The perceptron: a probabilistic model for information storage and organization in the brain.”.** In: *Psychological review* 65.6, p. 386.

Rumelhart, DE (1986). **“Learning representations by error propagation”.** In: *Parallel distributed processing* 1, pp. 318–362.

Vaswani, Ashish et al. (2017). **“Attention is all you need”.** In: *Advances in neural information processing systems* 30.

Sidestep : Transformers

All
data
in
deep
learning
must
be
represented
as
vectors

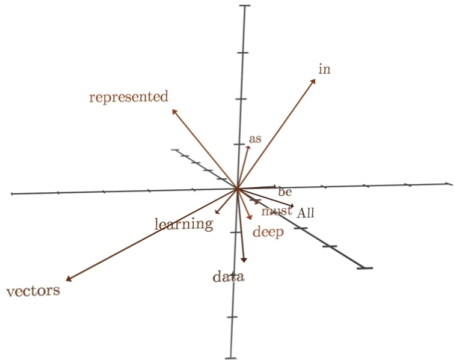

Adapted from 3blue1brown

Sidestep : Transformers

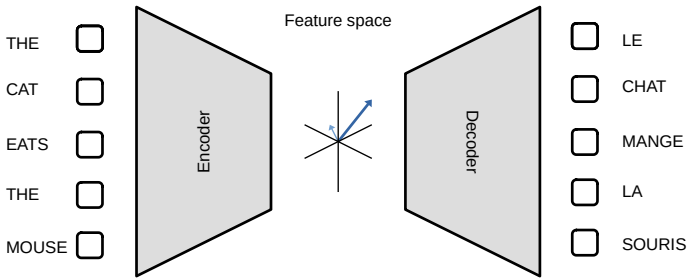

Sidestep : Transformers

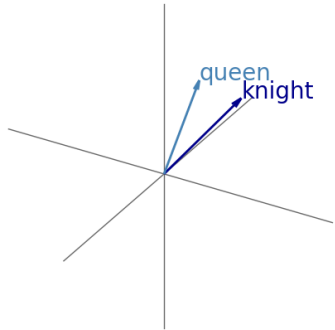

The knight saves the queen

Sidestep : Transformers

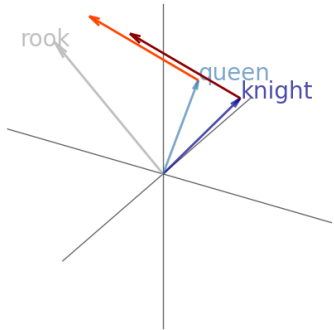

Queen takes knight , checkmate !

Sidestep : Transformers

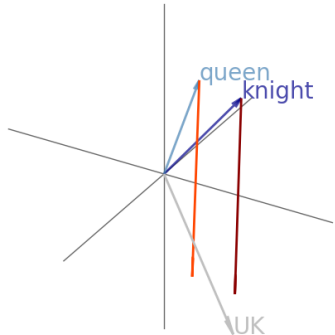

He was knighted by Queen Elizabeth II

Sidestep : Transformers