

Deep Learning and the Artificial Intelligence Revolution

Rita Fiorese, FaBiT, Unibo


Unibo, 2024

May 28, 2024

AI, Machine Learning and Deep Learning

ARTIFICIAL INTELLIGENCE

Early artificial intelligence stirs excitement.



MACHINE LEARNING

Machine learning begins to flourish.



DEEP LEARNING

Deep learning breakthroughs drive AI boom.



1950's

1960's

1970's

1980's


1990's

2000's

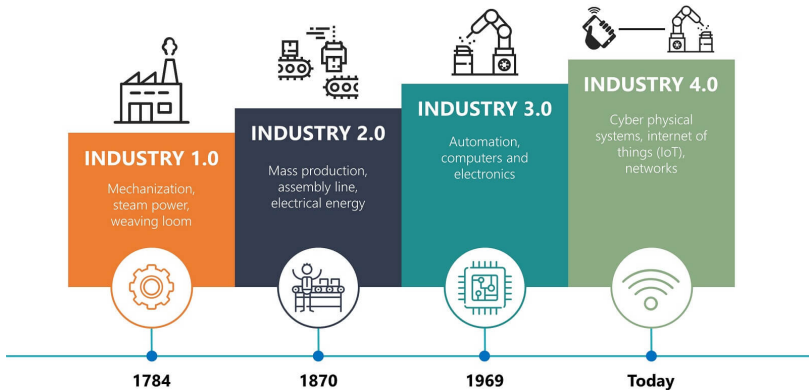
2010's



Deep Learning and Machine Learning




Industry 4.0



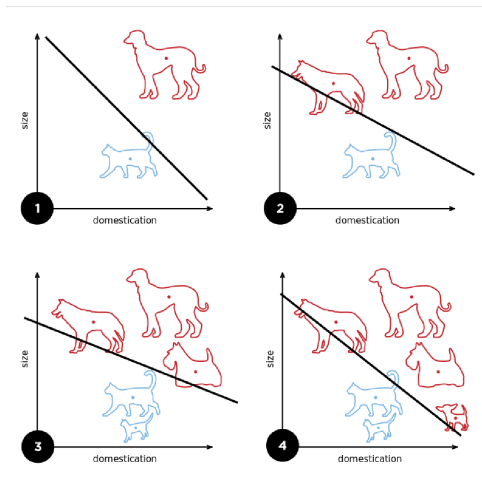
Some History: The Perceptron, a binary classifier

The first neural network: Mark I Perceptron (1957 Rosenblatt).
Navy Lab, Cornell University.



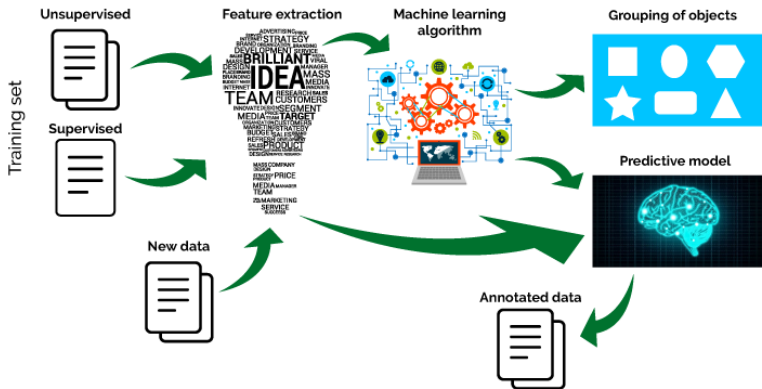
The Perceptron as Linear Classifier

The Perceptron works only if we can separate the two classes with a line.



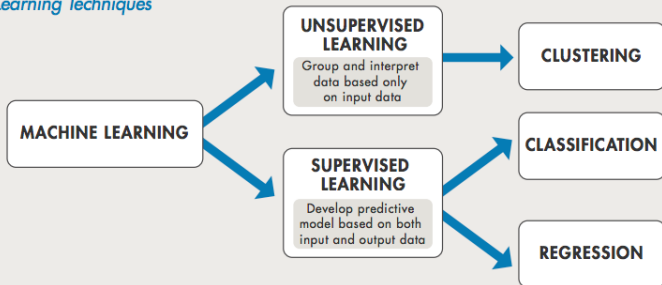
The Algorithms of Machine Learning

Machine Learning




Supervised/Unsupervised Learning


Machine Learning Techniques



Machine Learning Cycle



Imagenet Challenge ILSVRC: ImageNet Large Scale Visual Recognition Challenge



- **2010** 20000 images, 20 categories: 25% error.
- **2011** 1 million images, 1000 categories: 16% error.
- **2015** 1 million di images, 1000 categories: 4% error.

The Imagenet Challenge was declared won in 2017 by a Deep Learning algorithm.



Images in Imagenet



(a) ImageNet Synset: One sample image from each category



(b) Corel-1000 Dataset: Sample images from each category



(c) Caltech-256 Dataset: One sample image from each category



(d) Caltech-101 Dataset: One sample image from each category

Imagenet Challenge: the winners

- **2012 Alexnet:** Error 15% Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton (la prima CNN)
- **2013 ZF Net:** Error 11% Matthew Zeiler, Rob Fergus, NYU.
- **2014 VGG Net:** Error 7.3% Karen Simonyan, Andrew Zisserman, University of Oxford
- **2014 GoogLeNet:** Error 6.7%
- **2015 Microsoft ResNet:** Errore 3.6%



MNIST

0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
2 2 2 2 2 2 2 2 2 2 2 2 2 2 2
3 3 3 3 3 3 3 3 3 3 3 3 3 3 3
4 4 4 4 4 4 4 4 4 4 4 4 4 4 4
5 5 5 5 5 5 5 5 5 5 5 5 5 5 5
6 6 6 6 6 6 6 6 6 6 6 6 6 6 6
7 7 7 7 7 7 7 7 7 7 7 7 7 7 7
8 8 8 8 8 8 8 8 8 8 8 8 8 8 8
9 9 9 9 9 9 9 9 9 9 9 9 9 9 9



CIFAR10

airplane



automobile



bird



cat



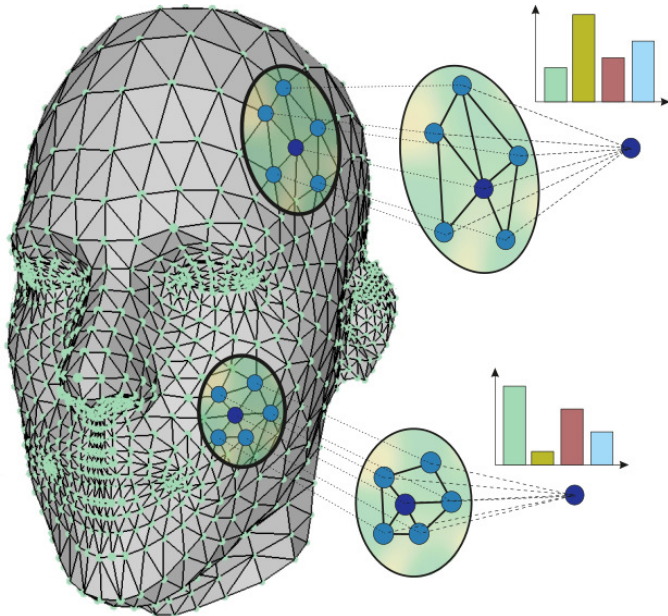
deer



dog



Geometric Deep Learning: Graph Convolutional Networks



Geometric Deep Learning

The eigenfunctions of the laplacian form the smoothest-possible basis function over a specific graph (they minimize the Dirichlet energy).

