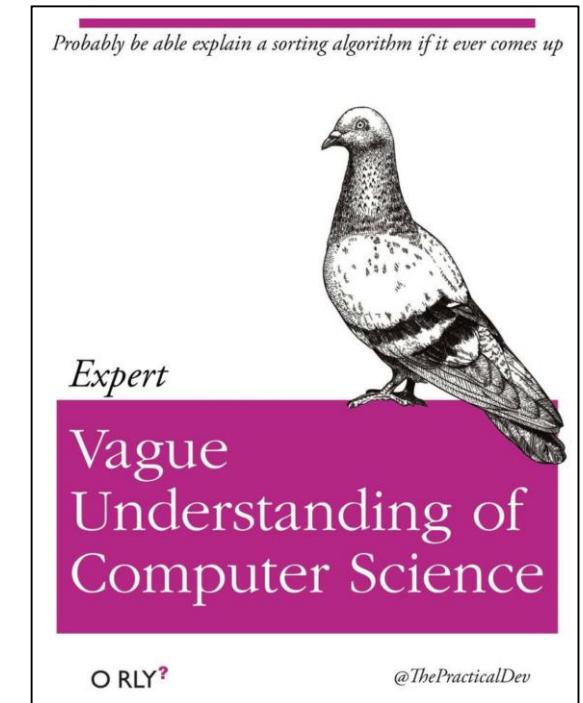
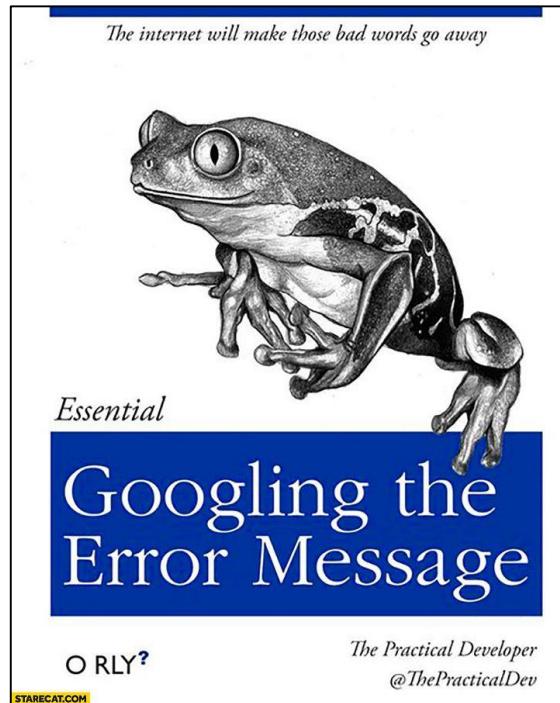
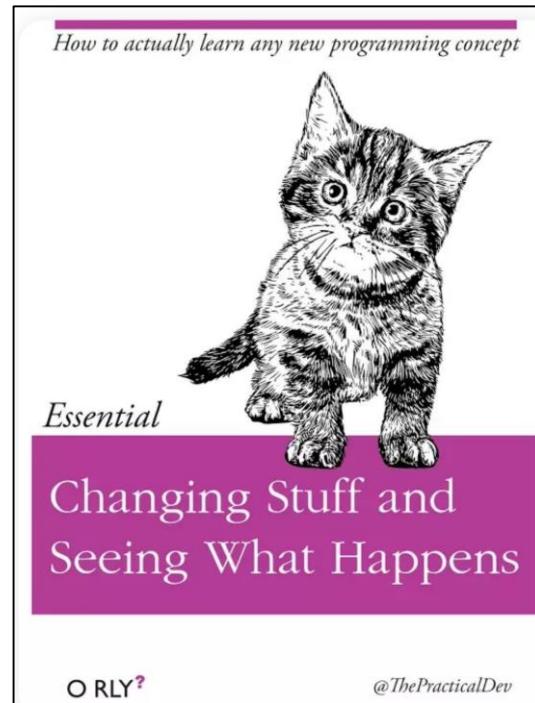
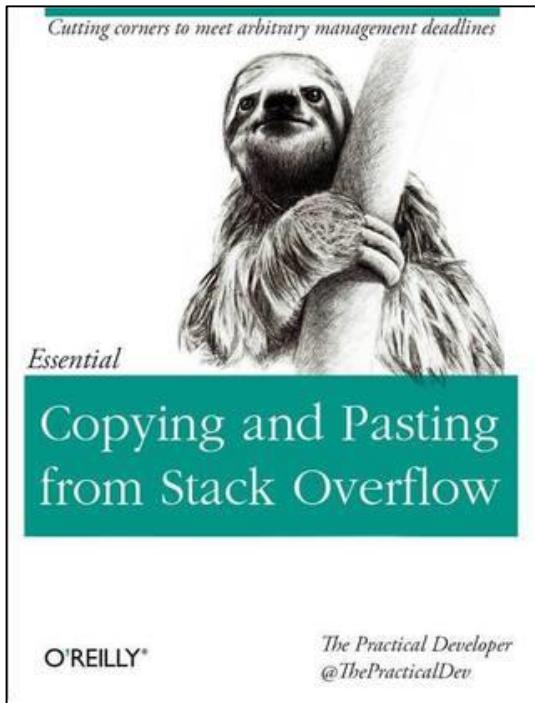


Fundamentals of AI

Introduction and the most basic concepts

How to learn fundamentals of AI?

Useful books***



***It is a joke

Any AI (ML) method in two lines of code in
any programming language

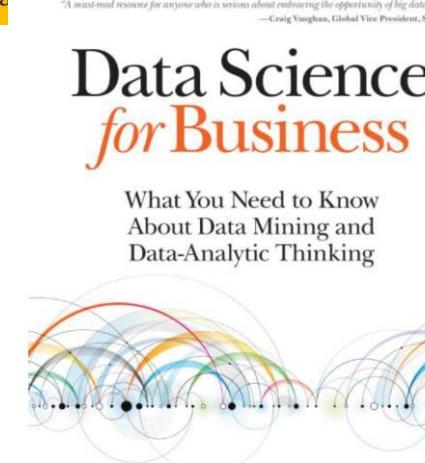
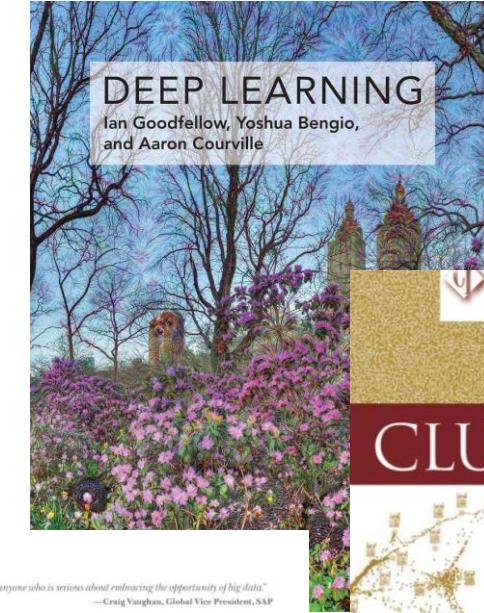
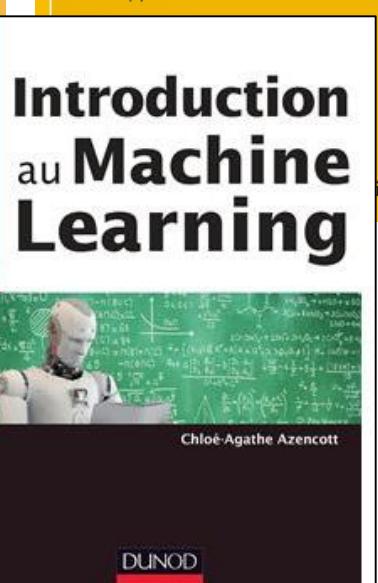
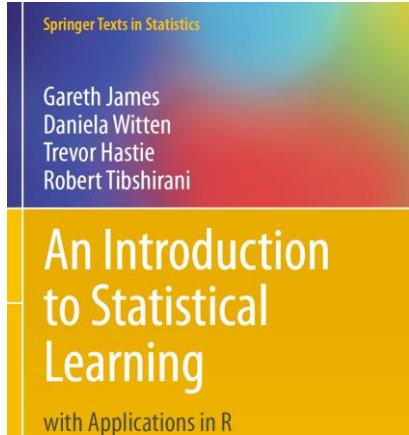
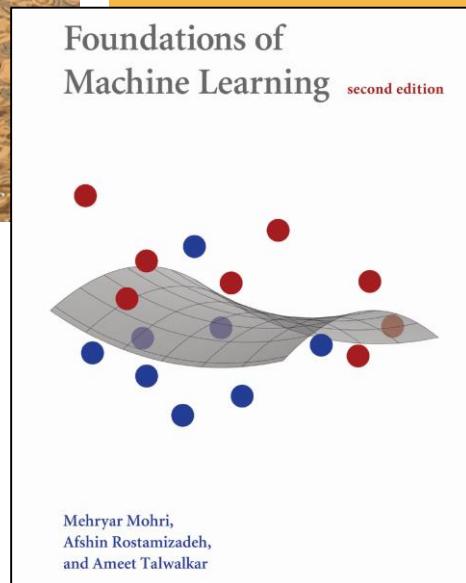
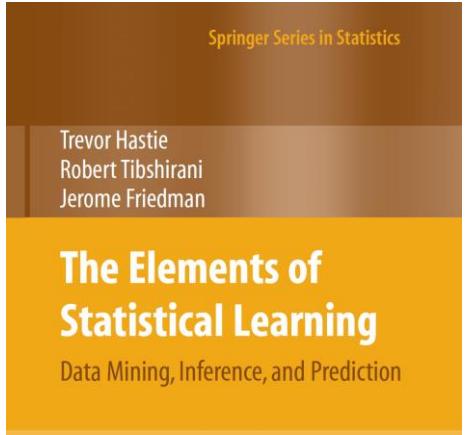
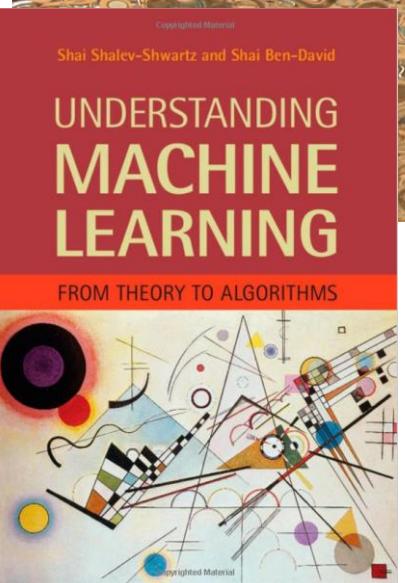
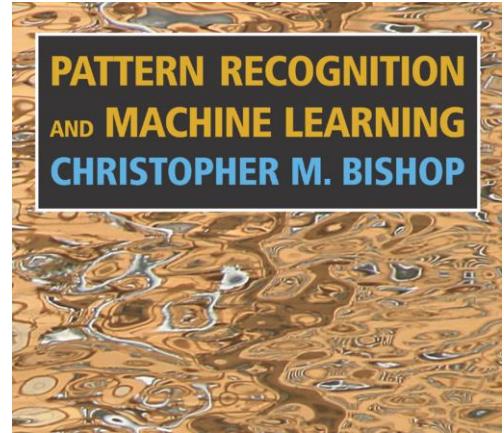
```
from libraryA import modelB as model
```

```
model.fit(X,Y)
```

```
model.predict(X)
```

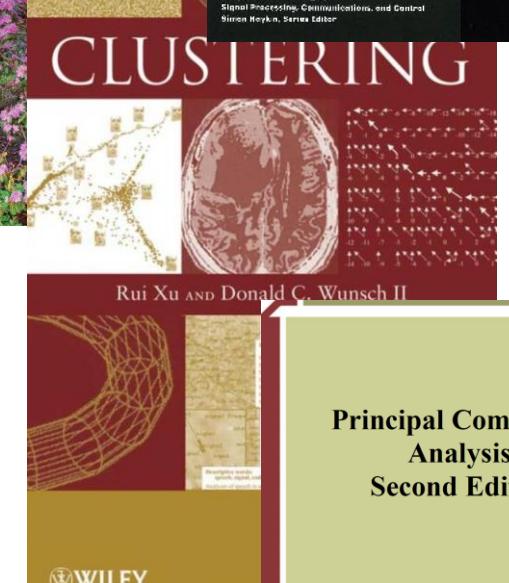
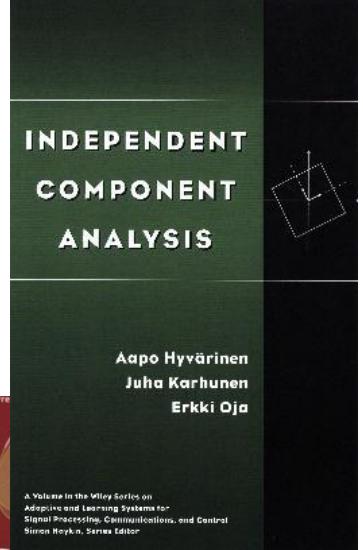
The rest is either data pre-processing or presenting the results...

Useful books***



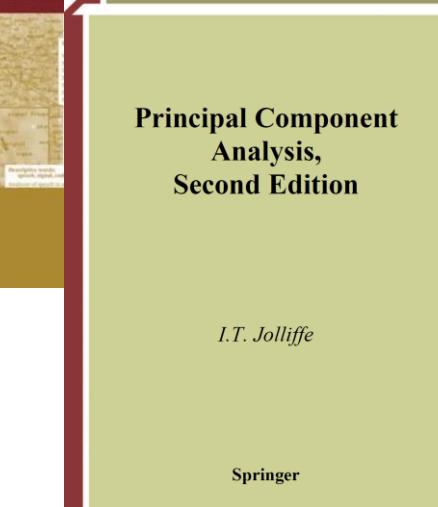
Foster Provost & Tom Fawcett

www.it-ebooks.info



Principal Component Analysis, Second Edition

I.T. Jolliffe



***now not a joke

Springer

Tons of YouTube channels and other stuff...

www.lebigdata.fr › Analytics › Data Analytics

Machine Learning et Big Data : définition et exlications de la ...



Le **Machine Learning** est une technologie d'intelligence artificielle permettant aux ordinateurs d'apprendre ...
6 juil. 2018 - Ajouté par Google Cloud Platform

openclassrooms.com › Accueil › Cours

Initiez-vous au Machine Learning - OpenClassrooms



Le **Machine Learning** est un ensemble de techniques utilisées par les Data Scientists qui a grandement fait ...
18 mars 2020

www.eskimoz.tr › machine-learning

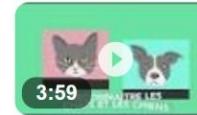
Machine learning et Deep learning : comment ça marche ?



Qu'est-ce que le **Machine Learning** ? Comment fonctionne le **Deep Learning** ? Quel est l'impact sur le ...
12 juin 2017 - Ajouté par Performance Web

www.youtube.com › watch

Le Machine learning - YouTube



Depuis toujours, les **machines** pensantes préoccupent les esprits. L'intelligence artificielle, comme on l ...
29 janv. 2020 - Ajouté par Centrale Nantes

developers.google.com › ... › Cours d'initiation

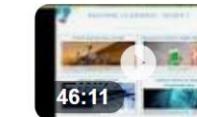
Cours d'initiation au machine learning | Google Developers



Une présentation intensive et pratique des principes fondamentaux du **machine learning** (durée : 20 heures ...
28 févr. 2018

www.youtube.com › watch

Le Machine Learning de A à Z - YouTube



Le **Machine Learning**, l'IA, le DeepLearning, les Statistiques, le Data Mining... bref, tous ces mots sont les ...
9 oct. 2017 - Ajouté par Salon Data

**What is valuable in learning with a teacher:
structuring, pinpointing, highlighting difficulties, expert's opinion:
ASK QUESTIONS!!!**

Andrew G. Moore's tutorials

- <http://www.cs.cmu.edu/~awm/>



- [Decision Trees](#)
- [Information Gain](#)
- [Probability for Data Miners](#)
- [Probability Density Functions](#)
- [Gaussians](#)
- [Maximum Likelihood Estimation](#)
- [Gaussian Bayes Classifiers](#)
- [Cross-Validation](#)
- [Neural Networks](#)
- [Instance-based learning \(aka Case-based or Memory-based or non-parametric\)](#)
- [Eight Regression Algorithms](#)
- [Predicting Real-valued Outputs: An introduction to regression](#)
- [Bayesian Networks](#)
- [Inference in Bayesian Networks \(by Scott Davies and Andrew Moore\)](#)
- [Learning Bayesian Networks](#)
- [A Short Intro to Naive Bayesian Classifiers](#)
- [Short Overview of Bayes Nets](#)
- [Gaussian Mixture Models](#)
- [K-means and Hierarchical Clustering](#)
- [Hidden Markov Models](#)

Andrew Ng's Coursera course on machine learning

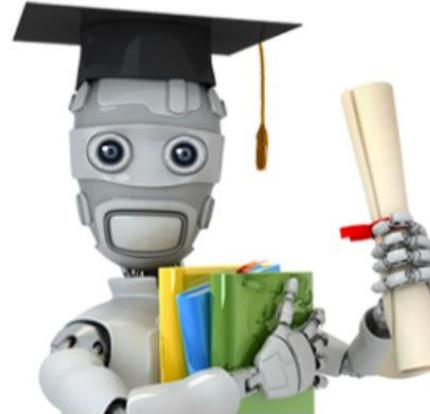
- <https://www.coursera.org/learn/machine-learning>



Andrew Ng

CEO/Founder Landing AI; Co-founder, Coursera; Adjunct Professor, Stanford University; formerly Chief Scientist, Baidu and founding lead of Google Brain

Stanford University
deeplearning.ai

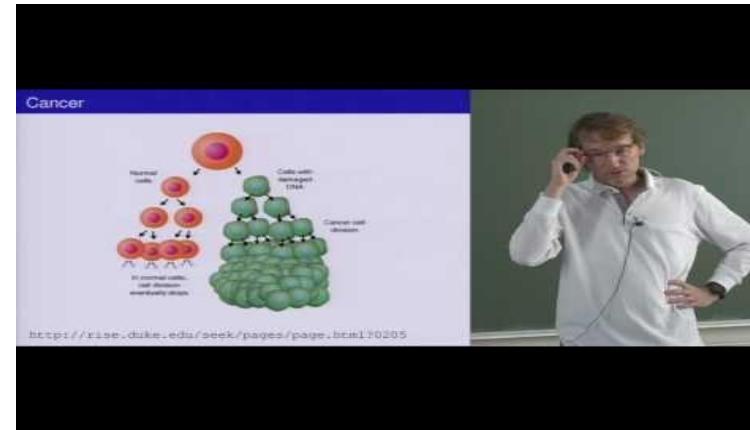


- 🌐 <https://www.andrewng.org/>
- FACEBOOK [andrew.ng.96](https://www.facebook.com/andrew.ng.96)
- TWITTER [AndrewYNg](https://twitter.com/AndrewYNg)
- LINKEDIN [andrewyng/](https://www.linkedin.com/in/andrewyng/)

bio

Andrew Ng is Co-founder of Coursera, and an Adjunct Professor of Computer Science at Stanford University. His machine learning course is the MOOC that had led to the founding of Coursera!

AI in France, French opinion-makers



... and many others!

Main difficulty (and – may be - mistake) of learning machine learning and AI

- Machine learning as a zoo of methods
(better to learn 'animal classification')
- No unifying theory so far
(probability theory, but...

Vapnik-Chervonenkis computational learning theory, but...)

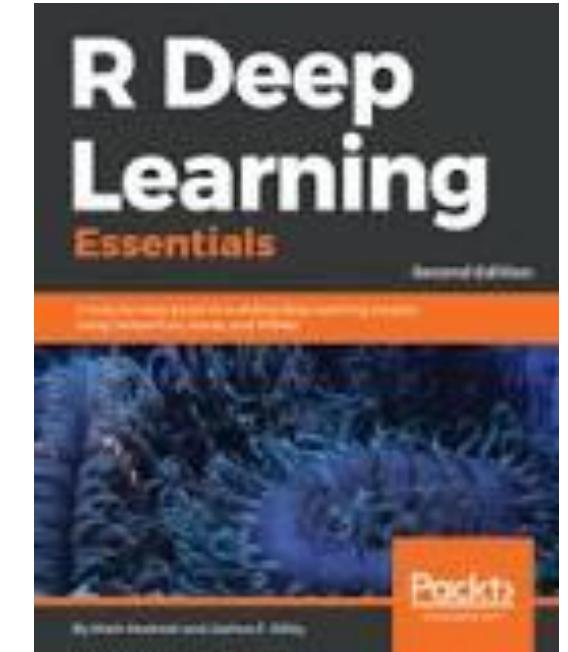
- '>model.fit(X,Y)
>model.predict(X)' trap
- What one should understand before and for studying ML and AI:

Linear algebra (vectors and matrices!), Methods of optimization, Probability theory, Functional analysis, Graph theory



Main current myth

- No need in all this zoo of machine learning methods
- No need to understand math behind
- One just need DEEP LEARNING
- However, despite the hype, deep learning probably accounts for less than 1% of the machine learning projects in production right now. Most of the recommendation engines and online adverts that you encounter when you browse the net are not powered by deep learning. Most models used internally by companies to manage their subscribers, for example churn analysis, are not deep learning models. The models used by credit institutions to decide who gets credit do not use deep learning.



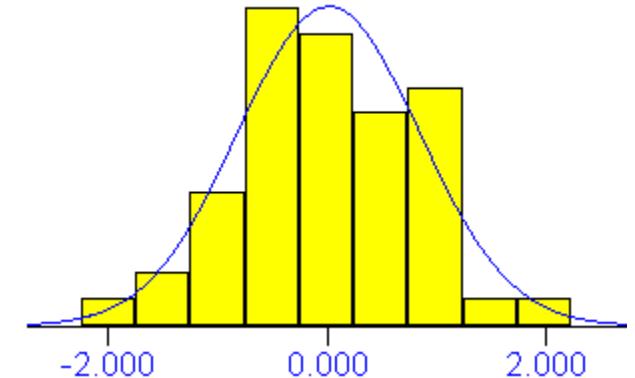
My objectives in this course:

- That you can ‘classify’ a new species (method) in ML-based AI
- That you can start reading WikiPedia, search StackOverflow and query Google in meaningfull and targeted way
- That you would see gears behind ‘model.predict(X,Y); model.fit(X)’

Two principal approaches to data mining (or statistics or AI or ML)

1) Hypothesis about underlying **probability distribution**

main notion is **probability density (PDF)**



2) **Geometrical approach** to the data analysis

main notion is metrics (**distance**)





High-dimensional post-classical world: Big Data, Bigger Dimension

D. Donoho, from Stanford University webpage

- **The number of attributes $p \gg$ The number of examples N**
- This *post-classical* world is different from the '*classical world*'.
- The classical methodology was developed for the '*classical world*' based on the assumption of $p < N$, and $N \rightarrow \infty$.
- These results all fail if $p > N$.
- **The $p > N$ case is not anomalous; it is the generic case.**

Donoho, D.L. High-Dimensional Data Analysis: The Curses and Blessings of Dimensionality.
Invited Lecture at Mathematical Challenges of the 21st Century, AMS.



High-dimensional post-classical world: Big Data, Bigger Dimension

D. Donoho, from Stanford University webpage

- **The number of attributes $p \gg$ The number of examples N**
- This *post-classical* world is different from the ‘classical’ world. Thus, the classical statistical learning theory become difficult to apply in the multidimensional post-classical world.
- **The $p > N$ case is not anomalous; it is the generic case.**

Donoho, D.L. High-Dimensional Data Analysis: The Curses and Blessings of Dimensionality. Invited Lecture at Mathematical Challenges of the 21st Century, AMS.



High-dimensional post-classical world: Big Data, Bigger Dimension

D. Donoho, from Stanford University webpage

Thus, the classical statistical learning theory become difficult to apply in the multidimensional post-classical world

- **Solution 1: Return to classics (dimensionality reduction)**
- **Solution 2: Exploit the properties of high-dimensions**

Donoho, D.L. High-Dimensional Data Analysis: The Curses and Blessings of Dimensionality. Invited Lecture at Mathematical Challenges of the 21st Century, AMS.