

# ARTIFICIAL INTELLIGENCE

SARAH ALKASASBEH

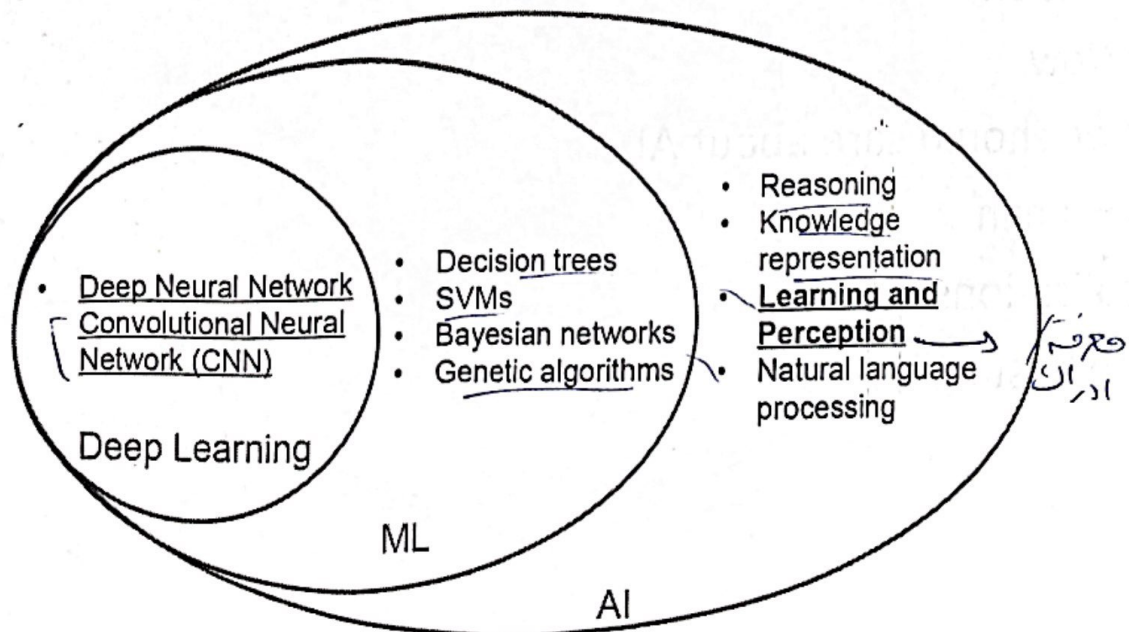
POWERUNIT

# Artificial Intelligence

- **Artificial Intelligence:** Build Machines which are capable of thinking like humans (mimic human behavior)
  - if-then statements programmed by experts
- **Machine Learning:** Give computers the ability to learn/make decisions without being explicitly programmed to do so
  - Adjust themselves in response to the data they're exposed to
- **Deep Learning:** Using Neural Networks to solve complex problems; Automatically discover patterns for feature detection

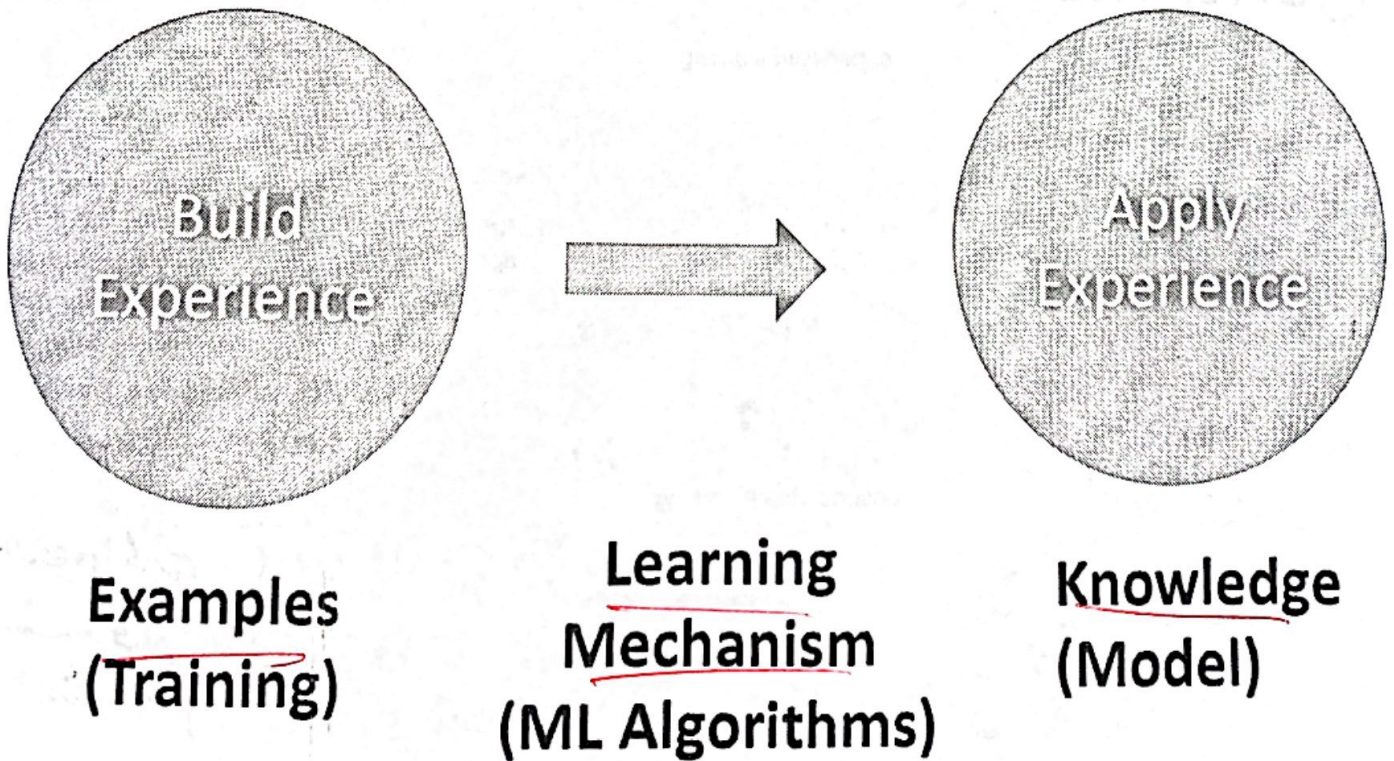
→ *Power of AI*

## AI Vs ML Vs Deep Learning

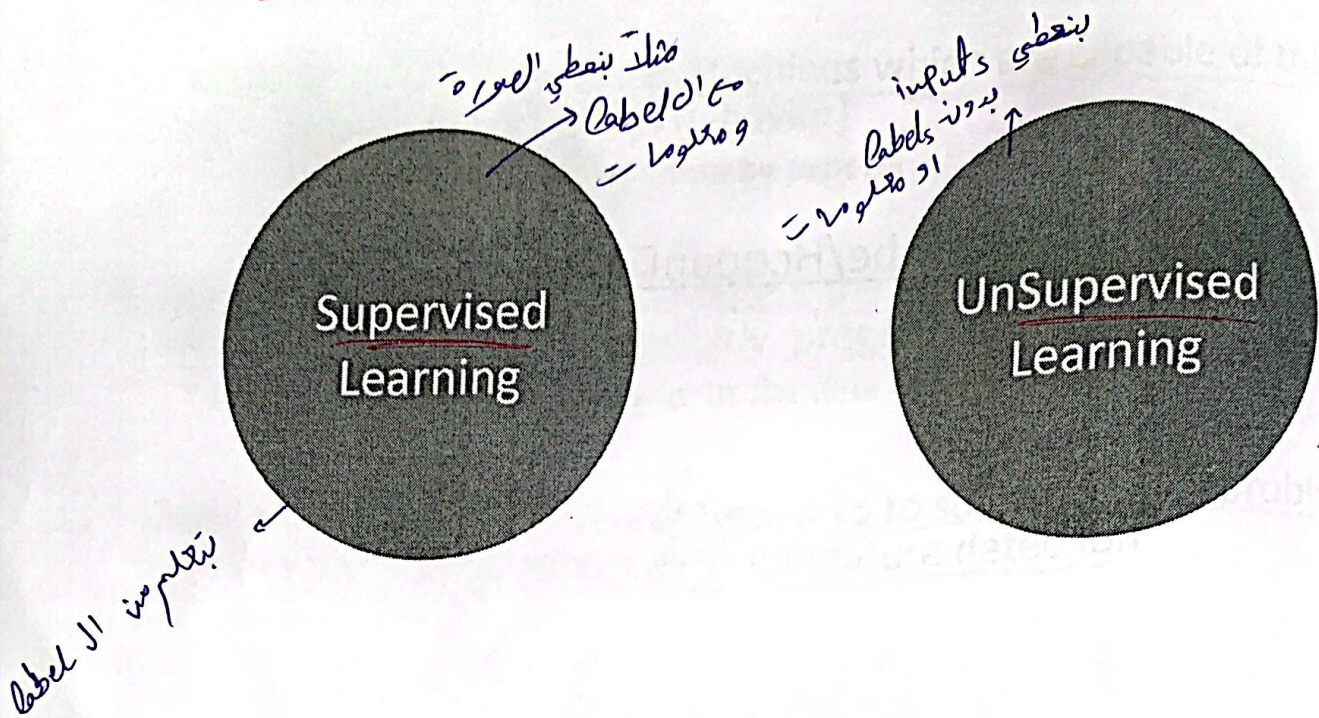


<https://youtu.be/HcqpanDadyQ>

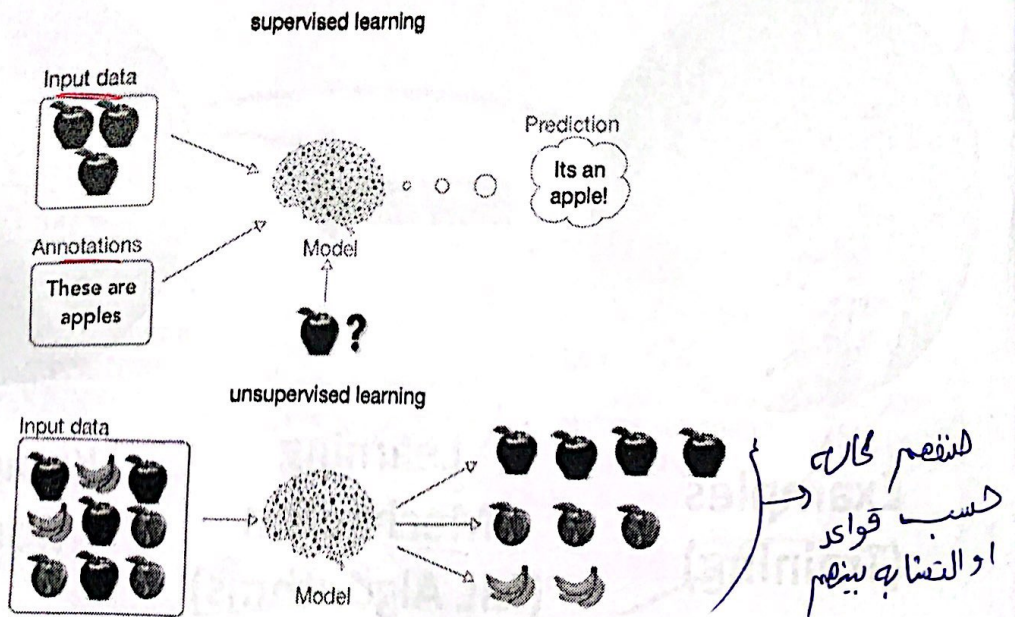
## Machine Learns Like the Human



# ML Algorithms

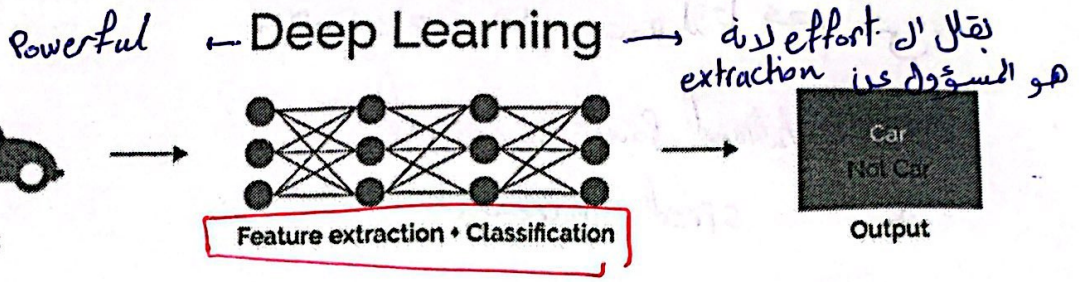
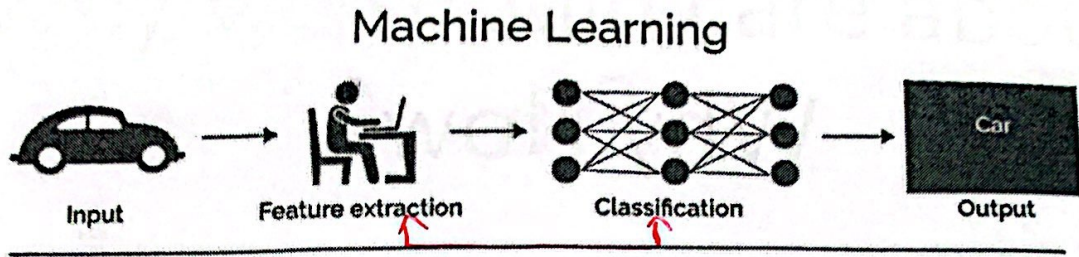


# ML Algorithms



Source: Background Augmentation Generative Adversarial Networks (BAGANs): Effective Data Generation Based on GAN-Augmented 3D Synthesizing


# Machine Learning vs Deep Learning



بأخذ قسم  
الـ pixels  
للصورة / الصورة

من الأتمتة الى  
Deep learning

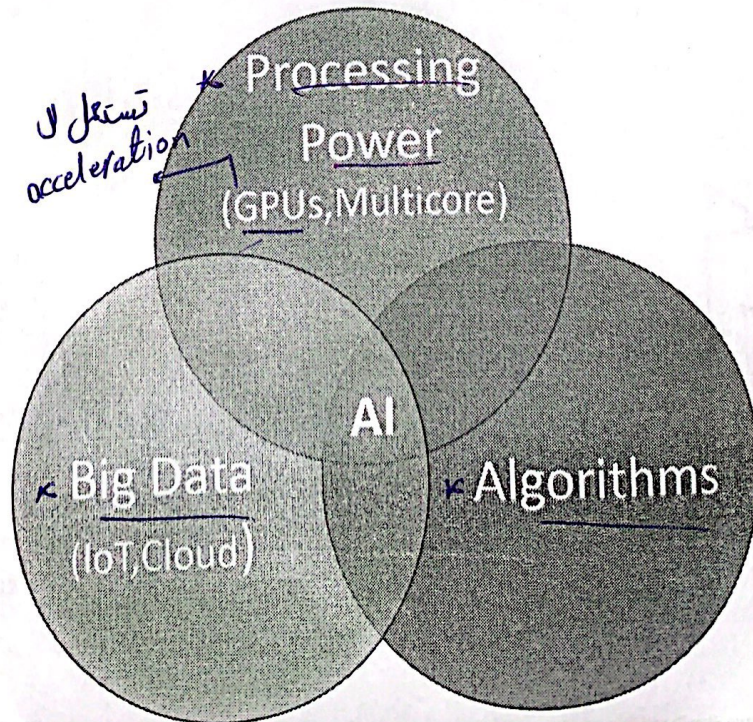
## Convolutional Neural Network      CNN



# Why Now?

- ضرورتاً الحاجة للـ processing لانا بحجم ايجين .
- Computational power increased
- & " speed increased

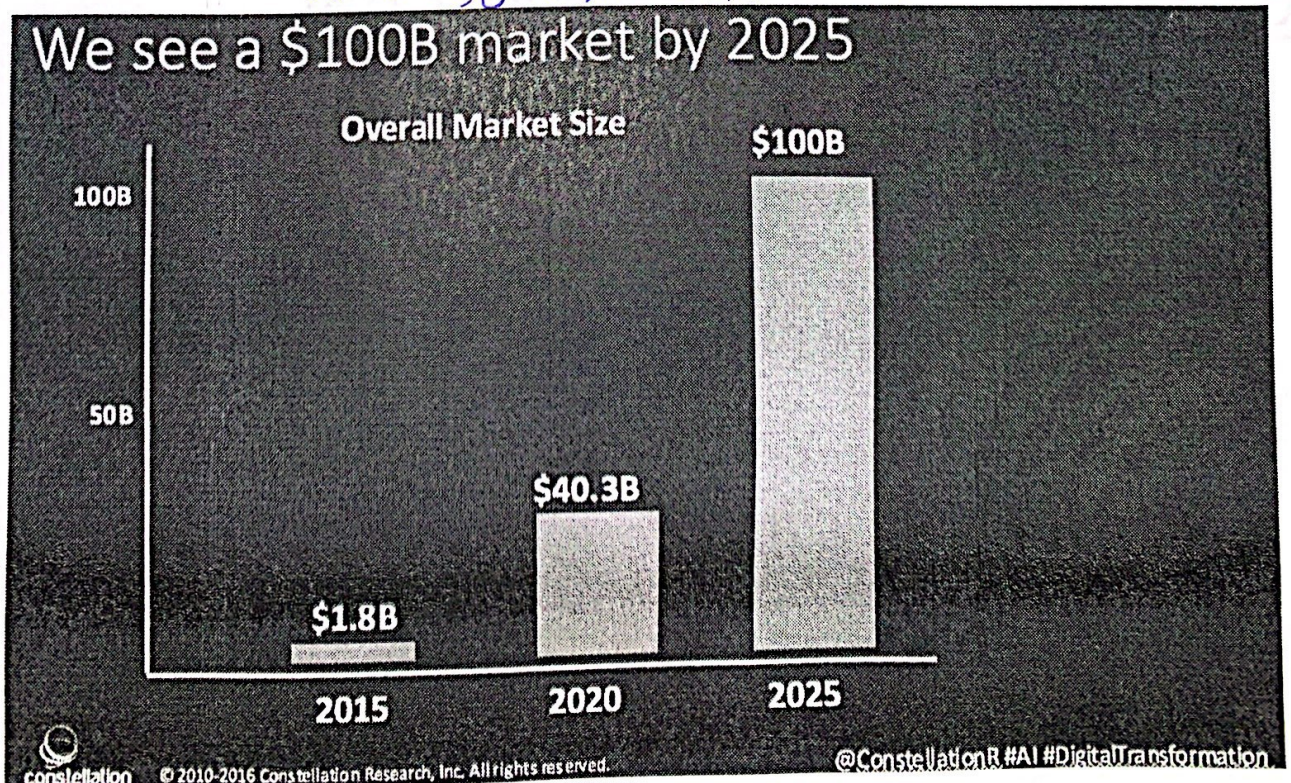
## AI emergence factors



# Why we should care about AI?

13

• عبء الأستثمار، باء الازح تزيد و التكلفة



# The Jobs Landscape in 2022



COMMITTED TO  
IMPROVING THE STATE  
OF THE WORLD

emerging  
roles,  
global  
change  
by 2022



## Top 10 Emerging

1. Data Analysts and Scientists
2. AI and Machine Learning Specialists
3. General and Operations Managers
4. Software and Applications Developers and Analysts
5. Sales and Marketing Professionals
6. Big Data Specialists
7. Digital Transformation Specialists
8. New Technology Specialists
9. Organisational Development Specialists
10. Information Technology Services

declining  
roles,  
global  
change  
by 2022



## Top 10 Declining

1. Data Entry Clerks
2. Accounting, Bookkeeping and Payroll Clerks
3. Administrative and Executive Secretaries
4. Assembly and Factory Workers
5. Client Information and Customer Service Workers
6. Business Services and Administration Managers
7. Accountants and Auditors
8. Material Recording and Stock Keeping Clerks
9. General and Operations Managers
10. Postal Service Clerks


Source: Future of Jobs Report 2018, World Economic Forum



[https://hai.stanford.edu/sites/g/files/sbiybj10986/f/ai\\_index\\_2019\\_report.pdf](https://hai.stanford.edu/sites/g/files/sbiybj10986/f/ai_index_2019_report.pdf)

Number of AI papers on arXiv, 2010-2019

Source: arXiv, 2019.



## Deep Learning Papers on arXiv

### Ranking Countries based on Total Number of Deep Learning Papers on arXiv, 2015-18

Source: arXiv, NESTA, 2019


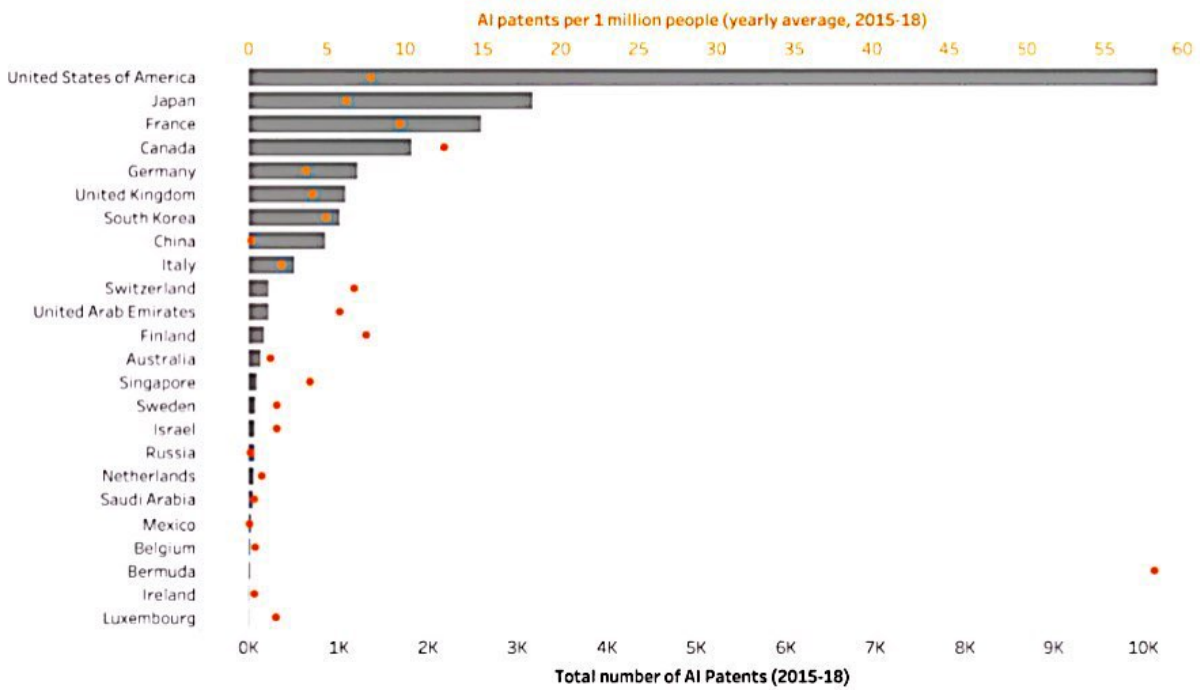


Fig. 1.7b

### Total Volume and average annual per capita AI Published Patents, 2015-2018

Source: MAG, 2019.



## Global AI startups that have received funding within the last year (July 2018-July 2019)

Source: CAPIQ, Crunchbase, Quid, 2019.




Fig. 4.2.6a

Network showing 4,403 global AI startups that received investment between July 2018 and July 2019. Colored by sector with top five highlighted.

Appendix: How to Read a Quid Network

 1980

### Othello

In the 1980s Kai-Fu Lee and Sanjoy Mahajan developed [BILL](#), a Bayesian learning-based system for playing the board game Othello. In 1989, the program won the US national tournament of computer players, and beat the highest ranked US player, Brian Rose, 56—8. In 1997, a program named Logistello won every game in a six game match against the reigning Othello world champion.

 1995

### Checkers

In 1952, Arthur Samuels built a series of programs that played the game of checkers and improved via self-play. However, it was not until 1995 that a checkers-playing program, [Chinook](#), beat the world champion.

 1997

### Chess

Some computer scientists in the 1950s predicted that a computer would defeat the human chess champion by 1967, but it was not until 1997 that IBM's [DeepBlue system](#) beat chess champion Gary Kasparov. Today, chess programs running on smartphones can play at the grandmaster level.

22

 2011

### Jeopardy!

In 2011, the IBM Watson computer system competed on the popular quiz show Jeopardy! against former winners Brad Rutter and Ken Jennings. Watson won the first place prize of \$1 million.

 2015

### Atari Games

In 2015, a team at Google DeepMind used a reinforcement learning system to learn how to play 49 Atari games. The system was able to achieve human-level performance in a majority of the games (e.g., Breakout), though some are still significantly out of reach (e.g., Montezuma's Revenge).

 2016

### Object Classification in ImageNet

In 2016, the error rate of automatic labeling of [ImageNet](#) declined from 28% in 2010 to less than 3%. Human performance is about 5%.

23

 2016

### Object Classification in ImageNet

In 2016, the error rate of automatic labeling of [ImageNet](#) declined from 28% in 2010 to less than 3%. Human performance is about 5%.

 2016

### Go

In March of 2016, the AlphaGo system developed by the Google DeepMind team beat [Lee Sedol](#), one of the world's greatest Go players, 4—1. DeepMind then released [AlphaGo Master](#), which defeated the top ranked player, Ke Jie, in March of 2017. In October 2017, a Nature paper detailed yet another new version, [AlphaGo Zero](#), which beat the original AlphaGo system 100—0.

 2017

### Skin Cancer Classification

In a 2017 [Nature article](#), Esteva et al. describe an AI system trained on a data set of 129,450 clinical images of 2,032 different diseases and compare its diagnostic performance against 21 board-certified dermatologists. They find the AI system capable of classifying skin cancer at a level of competence comparable to the dermatologists.

 2017

### Speech Recognition on Switchboard

In 2017, [Microsoft](#) and [IBM](#) both achieved performance within close range of "human-parity" speech recognition in the limited Switchboard domain

 2017

### Poker

In January 2017, a program from CMU called [Libratus](#) defeated four to human players in a tournament of 120,000 games of two-player, heads up, no-limit Texas Hold'em. In February 2017, a program from the University of Alberta called [DeepStack](#) played a group of 11 professional players more than 3,000 games each. [DeepStack](#) won enough poker games to prove the statistical significance of its skill over the professionals

 2017

### Ms. Pac-Man


[Maluuba](#), a deep learning team acquired by Microsoft, created an AI system that learned how to reach the game's maximum point value of 999,900 on Atari 2600

25

 2018

### Chinese - English Translation

A [Microsoft](#) machine translation system achieved human-level quality and accuracy when translating news stories from Chinese to English. The test was performed on newstest2017, a data set commonly used in machine translation competitions.

 2018

### Capture the Flag

A DeepMind agent reached human-level performance in a modified version of Quake III Arena [Capture the Flag](#) (a popular 3D multiplayer first-person video game). The agents showed human-like behaviours such as navigating, following, and defending. The trained agents exceeded the win-rate of strong human players both as teammates and opponents, beating several existing state-of-the-art systems

 2018

### DOTA 2

[OpenAI Five](#), OpenAI's team of five neural networks, defeats amateur human teams at [Dota 2](#) (with [restrictions](#)). OpenAI Five was trained by playing 180 years worth of games against itself every day, learning via self-play. ([OpenAI Five is not yet superhuman, as it failed to beat a professional human team](#))

 2018

### Prostate Cancer Grading

Google developed a [deep learning system](#) that can achieve an overall accuracy of 70% when grading prostate cancer in prostatectomy specimens. The average accuracy of achieved by US board-certified general pathologists in study was 61%. Additionally, of 10 high-performing individual general pathologists who graded every sample in the validation set, the deep learning system was more accurate than 8

26

 2018

### AlphaFold

DeepMind developed [AlphaFold](#) that uses vast amount of geometric sequence data to predict the 3D structure of protein at an unparalleled level of accuracy than before.

 2019

### Alphastar

DeepMind developed [Alphastar](#) to beat a top professional player in [Starcraft II](#).

 2019

### Detect diabetic retinopathy (DR) with specialist-level accuracy

Recent [study](#) shows one of the largest clinical validation of a deep learning algorithm with significantly higher accuracy than specialists. The tradeoff for reduced false negative rate is slightly higher false positive rates with the deep learning approach.

27

# AI Applications

28

## AI in Everyday Life

- Email Filters and smart replies in Gmail
- LinkedIn: match candidates
- Pinterest's LENS tool
- Chatbots
- Facebook : Relevant posts
- Product Recommendations
- Banking: Financial Institutions fraud prevention (not Common types of transactions)
- Ride-sharing Apps
- Unlock phone with face ID
- Voice assistants

29

**Artificial Intelligence**  
Contact  
[info@venturescanner.com](mailto:info@venturescanner.com)  
to see all 957 companies

**Machine Learning Gen (123 Companies)**  
big, lyric labs, SI, etc.

**Machine Learning App (100 Companies)**  
is, etc.

**Computer Vision Gen (100 Companies)**  
clarifai, etc.

**Computer Vision App (83 Companies)**  
flyby, percipio, etc.

**Smart Robots (66 Companies)**  
arabot, etc.

**Virtual Personal Assistants (92 Companies)**  
Vingo, sherpa, tempo, aivo, ejenta, etc.

**NLP-Speech Recog (70 Companies)**  
etc.

**NLP-General (154 Companies)**  
etc.

**Speech to Speech Trans. (16 Companies)**  
etc.

**Content Autore Comp (20 Companies)**  
grokr, etc.

**Gesture Control (100 Companies)**  
etc.

**Recommendation Eng (60 Companies)**  
etc.

**Video Content Recog (14 Companies)**  
etc.

Venture Scanner

30

# Healthcare

- Robot Assisted Surgery
- Administration and Workflow
- Cybersecurity
- Automated Image Diagnosis
- Fraud Detection
- Treatment Design
- Health Monitoring
- Drug Creation

## 10 AI Applications That Could Change Health Care



# Medical

## LETTER

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

We assess the most common forms of melanoma—a potentially fatal skin cancer—by using a deep convolutional neural network (CNN) trained on over 128,000 clinical images of skin lesions on various parts of the human body. This model is trained to identify the presence of skin cancer with a performance comparable to that of board-certified dermatologists. Our findings suggest that deep learning can be used to assist in the diagnosis of skin cancer, potentially reducing the burden on dermatologists. The model is trained on a large dataset of clinical images, and its performance is evaluated on a separate set of images. The model's performance is compared to that of board-certified dermatologists, and the results show that the model's performance is comparable to that of the experts. This suggests that deep learning can be used to assist in the diagnosis of skin cancer, potentially reducing the burden on dermatologists.

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

Pooneh Alipour-Far, Armin Arora, Rajeev Datta, Dheeraj Dingra, Pradyumn K. Mishra, Anshul K. Nayak, Ankur Rastogi, Corinna Saharia, Shashank Shekhar, Shuang Tang, Arun Tejasvi, Andrew Veit, David Wang, Michael Whalen, Yizhe Wang, and Anand Kulkarni

#### Abstract


We describe an algorithm that can detect pneumonia from chest X-rays at a level comparable to expert radiologists. The algorithm, CheXNet, is a deep convolutional neural network trained on 112,120 chest X-ray images publicly available from the Internet. The model is trained to detect pneumonia with a performance comparable to that of expert radiologists. The model's performance is evaluated on a separate set of images, and the results show that the model's performance is comparable to that of the experts. This suggests that deep learning can be used to assist in the diagnosis of pneumonia, potentially reducing the burden on radiologists.



Input Chest X-Ray image

CheXNet

Output Pneumonia: Positive (81%)




#### 1. Introduction

Medical image analysis is a critical component of clinical practice. Deep learning has emerged as a powerful tool for medical image analysis, enabling the development of algorithms that can perform tasks such as image classification, object detection, and segmentation. In this paper, we describe an algorithm that can detect pneumonia from chest X-rays at a level comparable to expert radiologists.




# Augmented Reality Microscope



Source: Augmented Reality Microscope for Real-time Automated Detection of Cancer

34

## Dental Pathologies




35

## Finance

- **Portfolio Management:**
  - Algorithms built to calibrate a financial portfolio to the goals and risk tolerance of the user (Betterment).
- **Algorithmic Trading**
  - Fast Trading Decisions
- **Loan Insurance underwriting**
  - Trained on millions of consumers examples
- **Fraud Detection**
  - Detect anomalies and flag them to the security team

36

## Books on AI & Finance



37

## Agriculture

- Monitoring
  - Agricultural crop conditions
  - Weather and climate
  - Ecosystems
- Planning and policy-making
- Intelligent environment control for plant production systems
- Intelligent robots in agriculture
- An expert geographical information system for land evaluation
- Artificial neural network for plant classification using image processing.
- Control of green house.

38

## Crop and Soil Health Monitoring

- PEAT: agriculture tech startup
  - Plantix Mobile App
    - Identifies possible defects through images captured by the user's smartphone camera.
    - Users are then provided with soil restoration techniques, tips and other possible solutions as explained in the short video below:



39

# Crop and Soil Health Monitoring

- **Trace Genomics: ML for diagnosing Soil Defects**



40

## Monitoring Crop Health and Sustainability

- **FarmShots: high-resolution satellite imagery that detects plant health by analyzing absorbed light from field images**



41

## Drones and Computer Vision for Crop Analysis

- **SkySquirrel Technologies: Data Analytics for drone-based imaging in agriculture**
- **aWhere: Deliver the most complete agricultural information and insight for real-time agriculture decisions, every day, global**

# Harvesting



Strawberry harvesting robot

43

# Energy Industry

## AI in the Power Grid

- Smart Grids
- Sector Coupling
- Monitoring of the Grid
- Coordination of Maintenance Work

## AI in the Virtual Power Plant

- Coordination of Decentralized Plants
- Forecasts

Artificial Intelligence

## AI for Power Consumption

- Smart Home & Smart Meter

## AI in Electricity Trading

- Forecasts
- Algorithmic Trading
- Monitoring Trade

Source: <https://www.next-kraftwerke.com/knowledge/artificial-intelligence>

44

# Sports

Technology

## NFL Incorporates Machine Learning, AI Technology to Prevent Player Injuries

Lindsey Gilber

July 30, 2022

Share on [LinkedIn](#) [Facebook](#) [Twitter](#) [Email](#)

Welcome to Thomas Insights — every day, we publish the latest news and analysis to keep our readers up to date on what's happening in industry. [Sign up today](#) to get the day's top stories delivered straight to your inbox.



**Thomas Industry Update**  
Stay up to date on industry news and trends, product announcements and the latest innovations.  
  
[Subscribe](#)

Top 5 Stories This Week

Source: <https://www.thomasnet.com/insights/nfl-incorporates-machine-learning-ai-technology-to-prevent-player-injuries/>

45

# Sports

**SportsPro** | [Home](#) | [Images](#) | [Sports Analytics](#) | [Patterns](#) | [Events](#) | [More](#)

**Liverpool partner with SkillCorner for AI-powered analysis**

Machine learning platform to measure player performance.

By [David T. King](#) | [Twitter](#) | [Facebook](#) | [LinkedIn](#)

46

# Companies and AI

Field	Organization	Applications
Energy	Arco and Tenneco Oil Company	Neural networks used to help pinpoint oil and gas deposits
Government	Internal Revenue Service	Software used to read tax returns and spot fraud
Human services	Merced County, California	Expert systems used to decide if applicants should receive welfare benefits
Marketing	Spiegel	Neural networks used to determine most likely buyers from a long list
Telecommunications	BT Group	Heuristic search used for a scheduling application that provides work schedules for more than 20,000 engineers
Transportation	American Airlines	Expert systems used to schedule the routine maintenance of airplanes
Inventory forecasting	Hyundai Motor	Neural networks and expert systems used to reduce delivery time by 20 percent and increase inventory turnover from 3 to 3.4
Inventory forecasting	SCI Systems	Neural networks and expert systems used to reduce on-hand inventory by 15 percent, resulting in \$180 million in annual savings
Inventory forecasting	Reynolds Aluminum	Neural networks and expert systems used to reduce forecasting errors by 2 percent, resulting in an inventory reduction of 1 million pounds
Inventory forecasting	Unilever	Neural networks and expert systems used to reduce forecasting errors from 40 percent to 25 percent, resulting in a multimillion-dollar savings

© Cengage Learning\*

47

# Required Skills

Top skills for the top 5 e

- MACHINE LEARNING ENGINEER**
  - Python
  - Algorithms
  - Data science
  - Artificial intelligence
  - Deep learning
  - Analytics
  - Data mining
  - Novel
  - Big data
  - Natural language processing
  - Hadoop
- APPLICATION DEVELOPMENT ANALYST**
  - Software development
  - Java
  - Programming
  - Salesforce.com
  - Business analysis
  - SAP products
  - Application support
  - Maintenance & repair
  - Accounts payable
  - TMI
  - Java
  - Research and development (R&D)
  - SAP ERP
  - ASAP
- MACHINE LEARNING ENGINEER**
  - Python
  - Algorithms
  - Data science
  - Artificial intelligence
  - Deep learning
  - Analytics
  - Data mining
  - Nosql
  - Big data
  - Natural language processing
  - Hadoop
- Full Stack Engineer**
  - JS
  - PHP
  - SQL
  - on rails
  - script
  - i
  - eDB
  - in
  - Y
  - at
  - Up
  - and web segment
- DATA SCIENTIST**
  - Data science
  - Machine learning
  - Analytics
  - Data mining
  - Big data
  - Hadoop
  - Python
  - R
  - Management
  - Statistical modeling
  - Matlab
  - Statistics
  - Predictive modeling

Source: <https://dazeinfo.com/2018/09/07/top-en>

48

## Online Resources



49

## Skills



Source: <https://www.sonayuti.com/trainings/artificial-intelligence-and-machine-learning-training-program.php>

50

## Projects and Datasets

- Kaggle
  - [www.kaggle.com](http://www.kaggle.com)

kaggle

- Github



51

# What Is Machine Learning?

- YouTube Video: What is Machine Learning? from Google Cloud Platform

<https://youtu.be/HcqpanDadyQ>

## What Is Machine Learning?

- The science (and art) of programming computers so they can learn from data.
- The field of study that gives computers the ability to learn without being explicitly programmed. Arthur Samuel, 1959
- A computer program is said to learn from experience E with respect to some task T and some performance measure P, if its performance on T, as measured by P, improves with experience E. Tom Mitchell, 1997  
→ Data for training only to improve experience
  - E: Training set made of training instances (samples)
  - T: Test set
  - P: Such as accuracy

- Task → what to do
- Experience → ability to do it
- Performance measure → ex: accuracy.

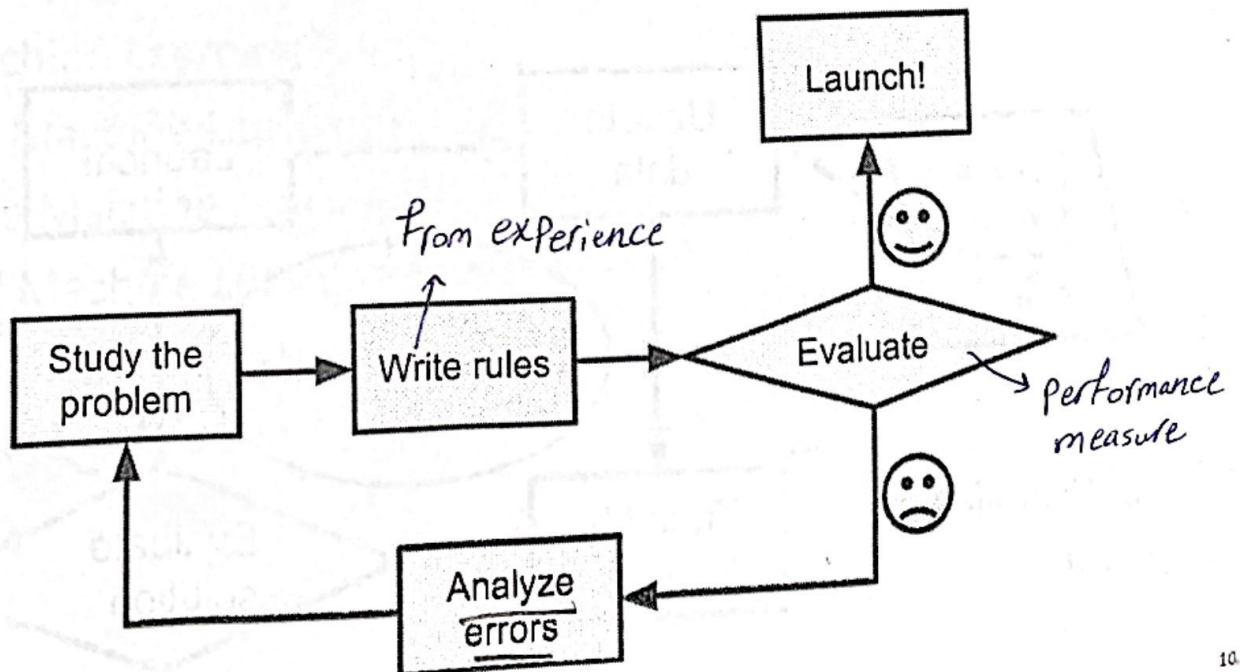
# Outline

- ✓ The Machine Learning Tsunami
- ✓ What Is Machine Learning?
  - Why Use Machine Learning?
  - Types of Machine Learning Systems
  - Main Challenges of Machine Learning
  - Testing and Validating
- Summary
- Exercises

9

## Why Use Machine Learning?

Spam filter using traditional programming techniques

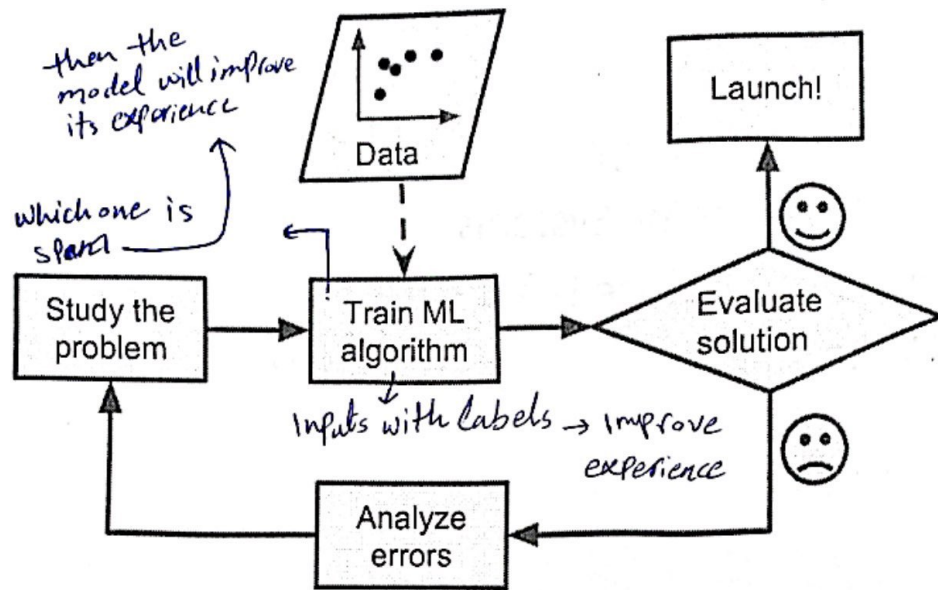


10



# Why Use Machine Learning?

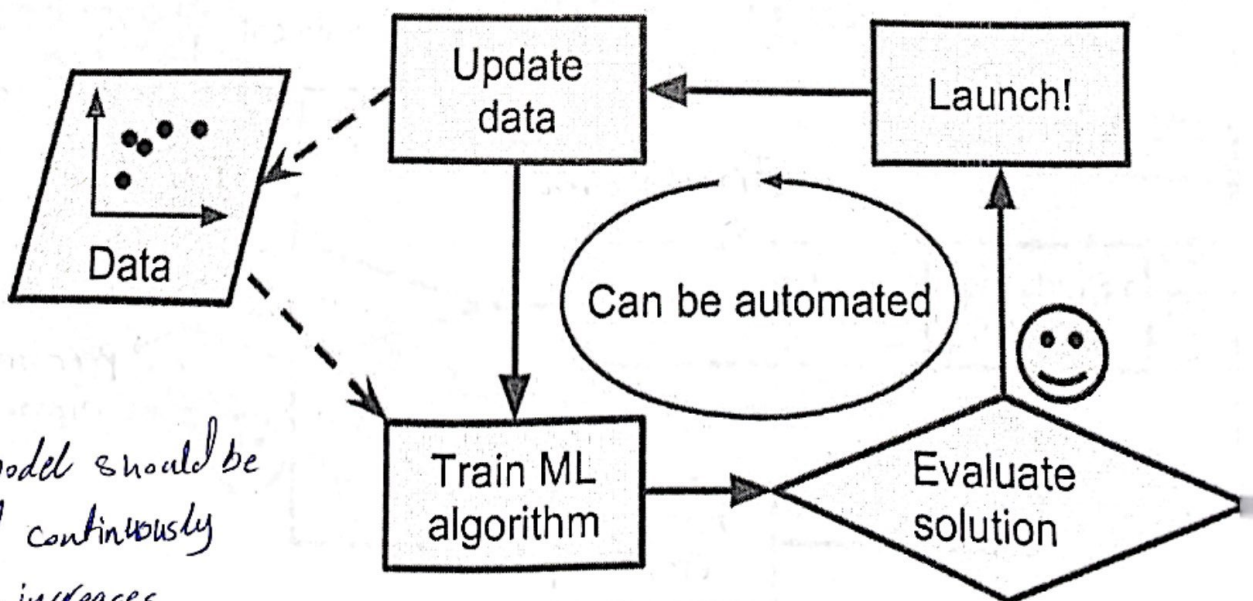
Spam filter using machine learning techniques 1/2



• Model يكتسب الخبرة من الـ model  
• Spam من الـ spam

# Why Use Machine Learning?

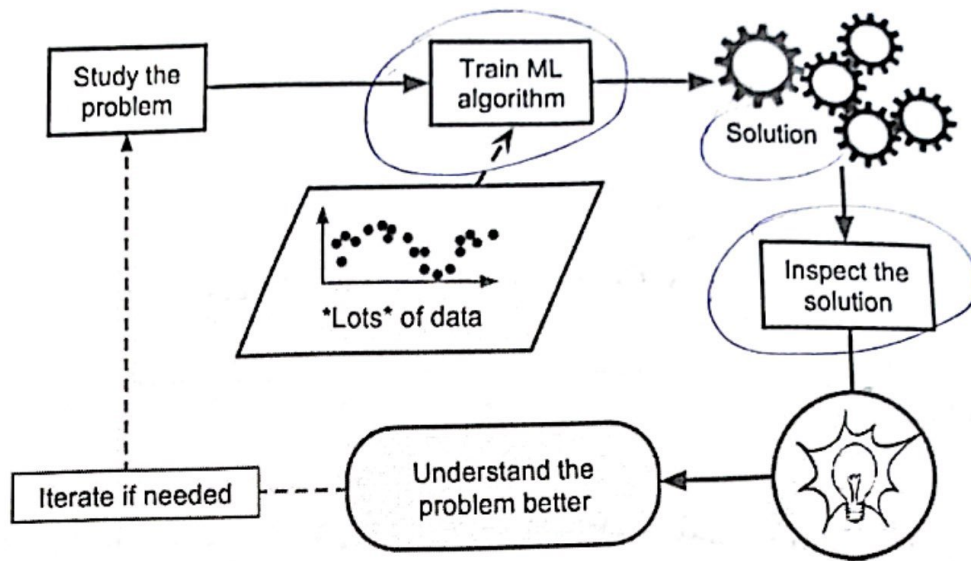
Automatically adapting to change 2/2



• The model should be updated continuously as data increases and old updated.

# Why Use Machine Learning?

ML can help humans learn (Data mining)



13

## Outline

- ✓ The Machine Learning Tsunami
- ✓ What Is Machine Learning?
- ✓ Why Use Machine Learning?
- Types of Machine Learning Systems
- Main Challenges of Machine Learning
- Testing and Validating
- Summary
- Exercises

14

# Types of Machine Learning Systems

## • Involves human supervision?

1. Supervised learning → label training data
2. Unsupervised learning
3. Semi-supervised learning
4. Reinforcement learning

## • Generalization approach

1. Instance-based learning
2. Model-based learning


## • Learns incrementally?

1. Batch learning
2. Online learning

5. Self-supervised learning

حسب المسئلة وكيفية تدبيرها model

## 1. Supervised Learning



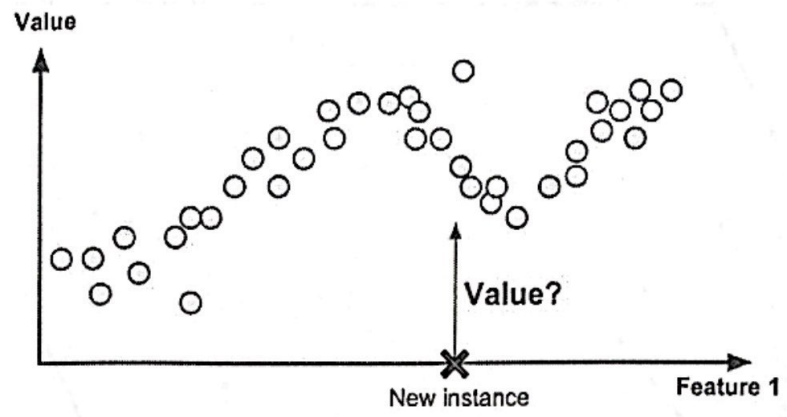
The training data you feed to the algorithm includes the desired solutions, called labels.

**Classification:** finds the class, e.g., email type (spam or ham)

(A, B, C...) ↓  
مثلا = كلمة الطالب

overhead (labeling data)

# 1. Supervised Learning



**Regression** finds the value, e.g., car price

مثال: توقع علامة الطالب (90.5 / 90.3) → Continuous range

# 1. Supervised learning algorithms

find linear relationship between inputs and outputs

Algorithm	Type
k-Nearest Neighbors	Both
Linear Regression	Regression
Logistic Regression	Classification
Support Vector Machines (SVMs)	Both
Decision Trees	Both
Random Forests	Both
Neural Networks	Both

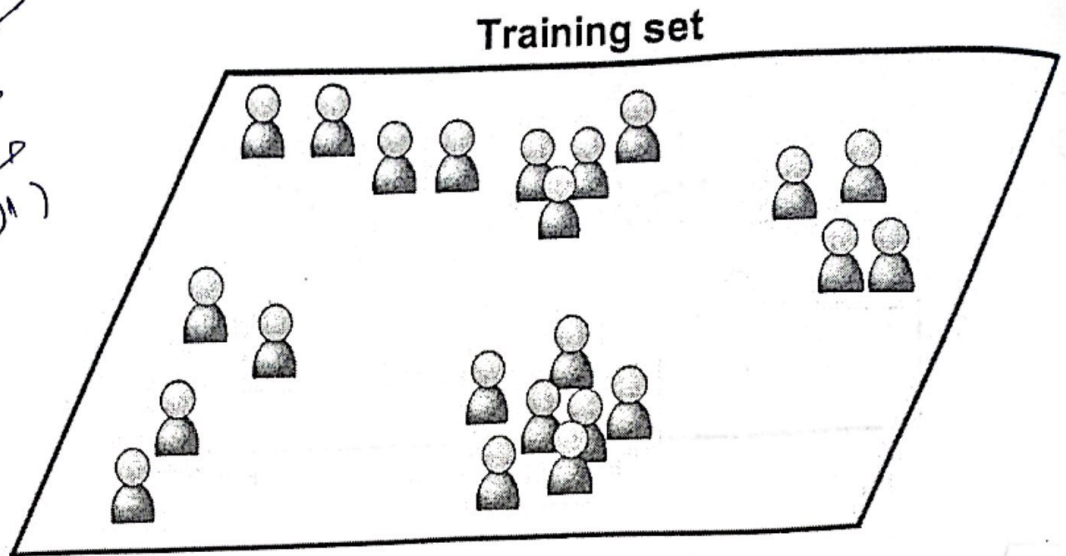
بعض classifier regression

## 2. Unsupervised Learning

Data without labels

بنحوده عدد ال clusters

يتم التصنيف حسب خصائصها (التشابه)



The training data is unlabeled.

## 2. Unsupervised learning algorithms

### Clustering

- k-Means
- Hierarchical Cluster Analysis (HCA)
- Expectation Maximization

needs distance measures → عثمان زيد  
العنوان اقرب  
كاي مجموعة


### Visualization and dimensionality reduction

- Principal Component Analysis (PCA)
- Kernel PCA
- Locally-Linear Embedding (LLE)
- t-distributed Stochastic Neighbor Embedding (t-SNE)

### Association rule learning

- Apriori
- Eclat


## 2.a Clustering



↓  
ex: Pixels

# 2.b Visualization

- + cat
- o automobile
- truck
- frog
- × ship
- airplane
- ◇ horse
- △ bird
- ▽ dog
- ▷ deer

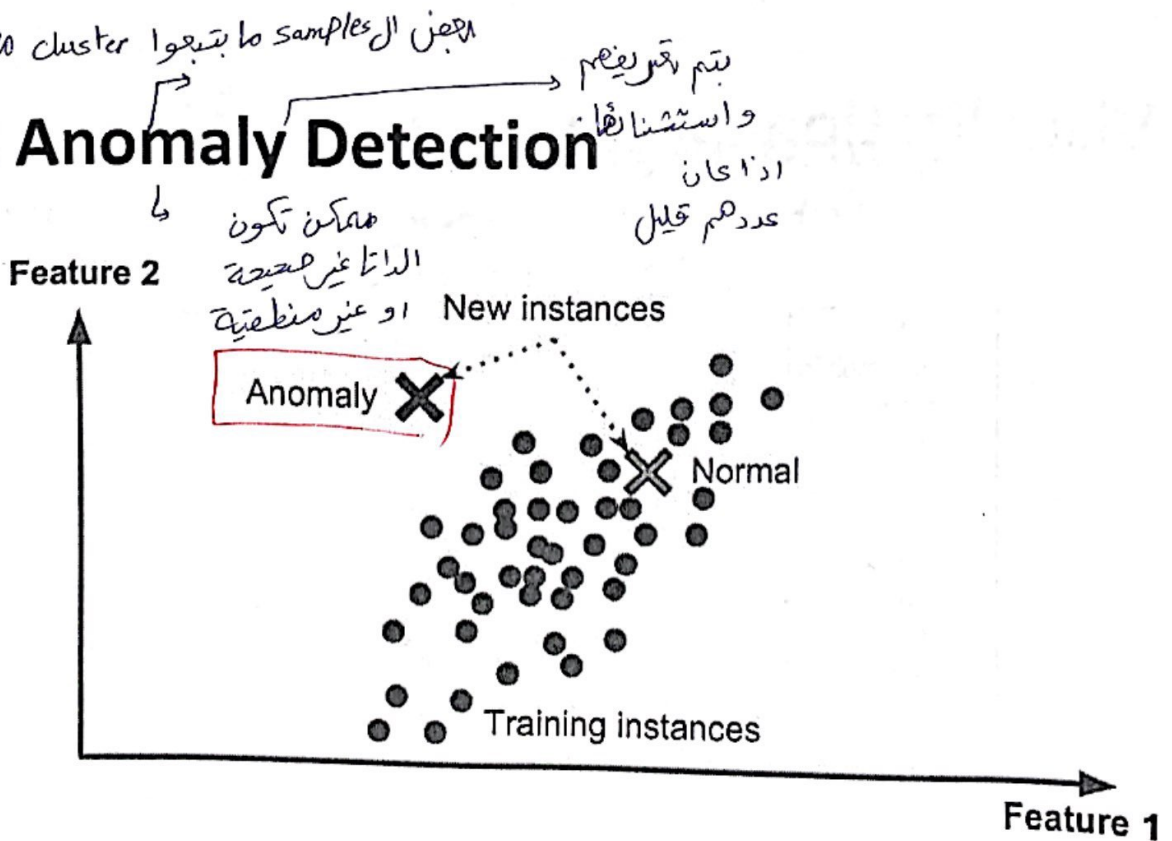


نقل الـ features (نقل منجز)  
او ندمجهم

## 2.c Dimensionality Reduction

- The goal is to simplify the data without losing too much information.
- One way to do this is to merge several correlated features into one. For example, a car's mileage may be very correlated with its age, so the dimensionality reduction algorithm will merge them into one feature that represents the car's wear and tear.
- Also called feature extraction.

## 2.d Anomaly Detection





خاصة  
لا يكون جميع  
البيانات

## 2.e Association Rule Learning

ترتيب البيانات

- The goal is to dig into large amounts of data and discover interesting relations between attributes.
- For example, suppose you own a supermarket. Running an association rule on your sales logs may reveal that people who purchase **barbecue sauce** and **potato chips** also tend to buy steak. Thus, you may want to place these items close to each other.

\* Medical diagnosis : using the multi-relational association rule, we can determine the probability of disease occurrence associated with various factors + symptoms in the data from past cases

\* Entertainment : Netflix + Spotify use association rules to fuel their content recommendation engines by analyzing user past behaviour

25

يعني مثلا الأشخاص حضور فيلم معين وتبعاً له حضور اسمه تاني


او اشخاص اشترى مثلا من امارو قطع وتبعاً الي اشترى اسمي تاني وهكذا GOT + hod

## 3. Semi-supervised Learning

لانه عليه ان labeling مكلف

(Partially labeled training data), usually a lot of unlabeled data and a little bit of labeled data. E.g., Google Photos.

Feature 2




ماكن ما يكون ال prediction صحيح ، لانه يتم الناس من ال prediction

Feature 1

26  
وبخيه بغير prediction

## 4. Self-supervised Learning

- Generating a fully labeled dataset from a fully unlabeled one



*Figure 1-12. Self-supervised learning example: input (left) and target (right)*

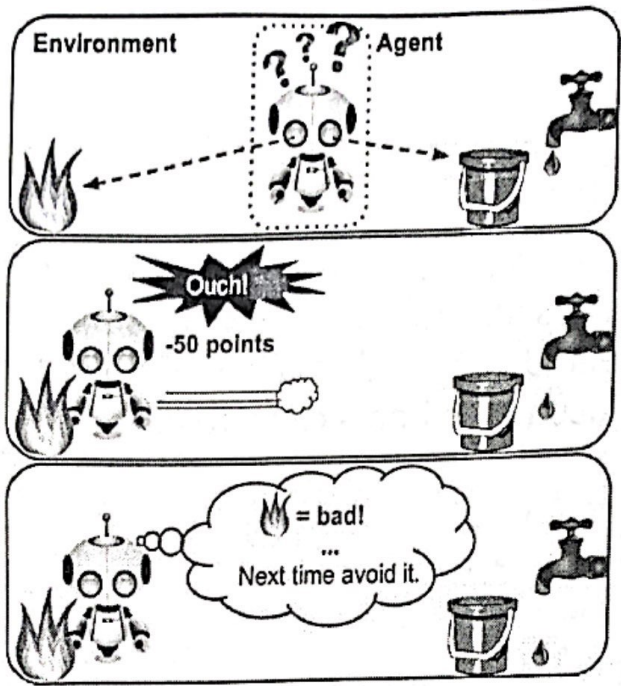
# 4. Reinforcement Learning

\* will not be covered

learning by experience

سائق السيارة  
والذي يتعلم  
القيادة  
من خلال التجربة  
هو reinforcement learning

بكون تجربة  
من كل محاولة  
او Past



- 1 Observe
- 2 Select action using policy
- 3 Action!
- 4 Get reward or penalty
- 5 Update policy (learning step)
- 6 Iterate until an optimal policy is found

## Types of Machine Learning Systems

7/13

### ✓ Involves human supervision?

1. Supervised learning
2. Unsupervised learning
3. Semi-supervised learning
4. Reinforcement learning

### • Generalization approach

1. Instance-based learning
2. Model-based learning

### • Learns incrementally?

1. Batch learning
2. Online learning

تدريب  
النموذج  
على كل  
البيانات

لتطوير  
النموذج  
واعادة  
تدريبه  
لأخذ المفضل  
بين الاختيار

# 1. Batch (offline) Learning

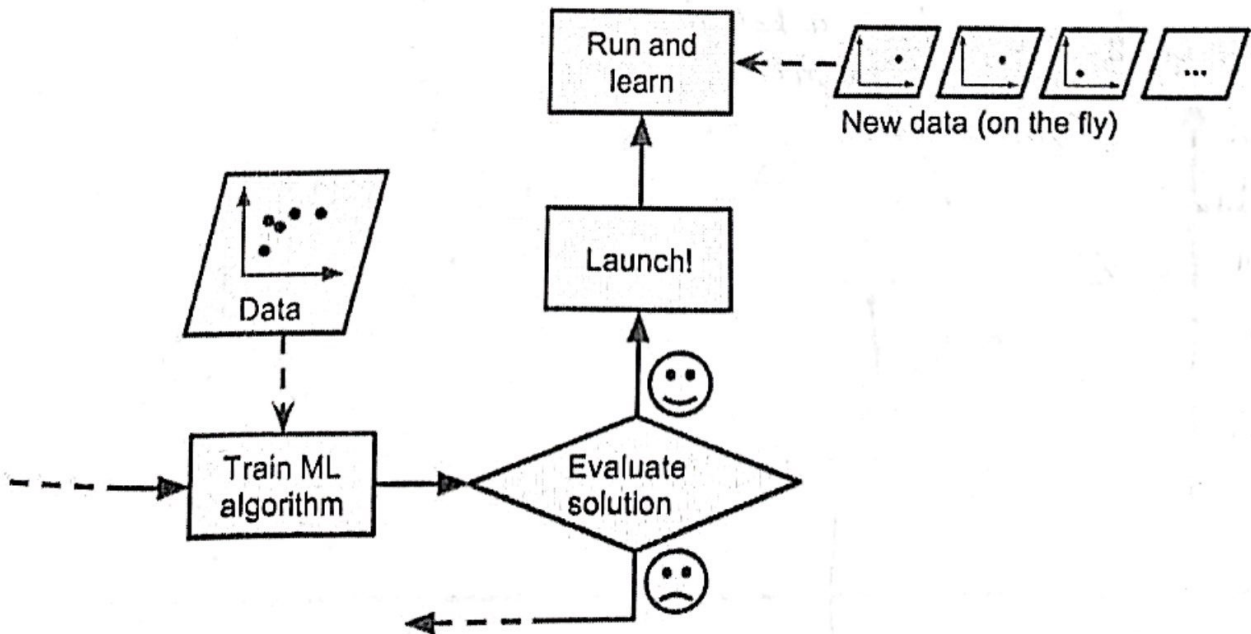
- Must be trained using (all the available data.)
- This will generally take a lot of time and computing resources, so it is typically done offline.
- First the system is trained, and then it is launched into production and runs without learning anymore; it just applies what it has learned.

أي تعديل على الـ data يتطلب إعادة تدريب الـ Test  
على الـ data القديمة والجديدة

29

## 2. Online Learning → Faster

Examples: Stock prices, huge data



30

# Types of Machine Learning Systems

## ✓ Involves human supervision?

1. Supervised learning
2. Unsupervised learning
3. Semi-supervised learning
4. Reinforcement learning

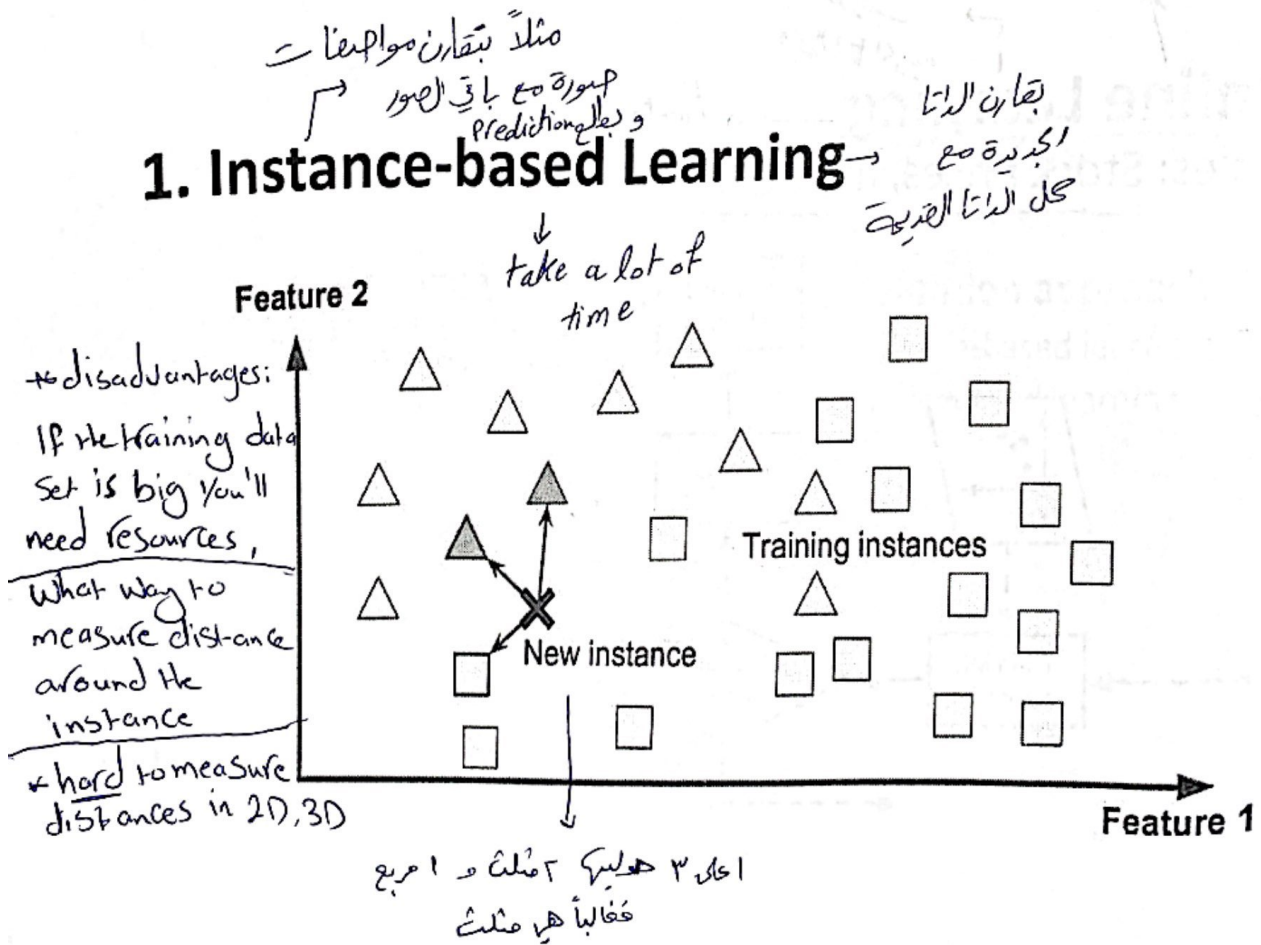
## • Generalization approach

1. Instance-based learning
2. Model-based learning

## ✓ Learns incrementally?

1. Batch learning
2. Online learning

## 1. Instance-based Learning




use the entire dataset as the model like (k-Nearest Neighbors).

use training data to create a model that has parameters ~~learned~~  
learned from the training data

## 2. Model-based Learning

\* Faster than instance based  
\* doesn't always need to bring out the training data while prediction  
\* less resources needed than instance based learning



• مثل هذا لا يتم تعويده  
• ال features فيها  
• ويكون الناتج  
• يتم مربوط بال  
decision boundary

• أكبر من ال decision boundary  
• اقل من ال // يكون من ربع

33

## Outline

- ✓ The Machine Learning Tsunami
- ✓ What Is Machine Learning?
- ✓ Why Use Machine Learning?
- ✓ Types of Machine Learning Systems
- Main Challenges of Machine Learning
- Testing and Validating
- Summary
- Exercises

34


# Main Challenges of Machine Learning (due to bad data)

1. Insufficient quantity of training data

• حجم البيانات

2. Non-representative training data

→ data shouldn't be biased, it should represent all classes



↑ data  
↑ accuracy

مثلاً في بلدان الدخل عندكم عالي وال life satisfaction عندكم 7 = وفي بلدان الدخل عندكم قليل وال life satisfaction عندكم 4 =  
البيانات ليست متوازنة classes وغير

# Main Challenges of Machine Learning (due to bad data)

3. Poor-quality data that contains:

- Errors
- Outliers
- Noise


features لا علاقة بالنتيجة

4. Irrelevant features: Need feature engineering:

- Feature selection: selecting the most useful features.
- Feature extraction: combining existing features to produce a more useful one.
- Creating new features by gathering new data.

# Main Challenges of Machine Learning (due to bad algorithm)

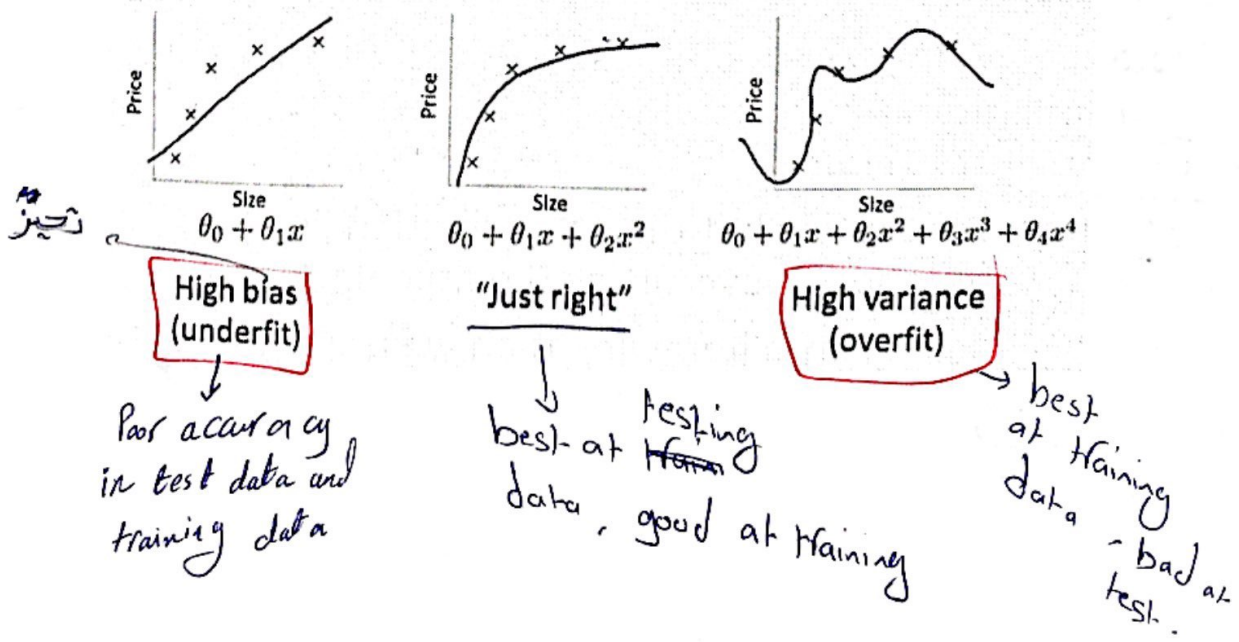
- Overfitting the training data (not general mode)
  - Regularization constrains the model's hyperparameters to make it simpler and (reduce the risk of overfitting.)



مش من النقاط بشكل حرفي، لو اعطيناه اي نقطة جديدة، ما راح يتعامل معها بشكل صحيح

# Main Challenges of Machine Learning (due to bad algorithm)

- Under-fitting the training data





# Outline

- ✓ The Machine Learning Tsunami
- ✓ What Is Machine Learning?
- ✓ Why Use Machine Learning?
- ✓ Types of Machine Learning Systems
- ✓ Main Challenges of Machine Learning
  - Testing and Validating
  - Summary
  - Exercises

## \* Testing and Validating

- ① **Split** your data into two sets (**cross validation**):
  - The training set (80%)
  - The test set (20%)
- ② **Evaluate**:
  - The training error
  - The generalization error
  - If the training error is low but the generalization error is high, it means that your model is overfitting the training data.
  - When the ML algorithm is iterative, often we use a third set:  
**validation set.**

test during the training

let's hint

↳ third set


- classes should be well represented in the training data and test data.

# Cross Validation

تقسيم البيانات لاجزاء واعمل  
 → train, test  
 اجزاء

- In **k-fold cross-validation**, the original sample is randomly partitioned into k equal size subsamples.

\* Cover all samples as validation set.



بكل مرة بنبدال model من الصفر

# Summary

- ML is about making machines get better at some task by learning from data, instead of having to explicitly code rules.
- Types of ML systems: supervised or not, batch or online, and instance-based or model-based.
- A model-based algorithm tunes some parameters to fit the model to the training set, and then hopefully it will be able to make good predictions on new cases.
- An instance-based algorithm learns the examples by heart and uses a similarity measure to generalize to new instances.
- The system will not perform well if your training set is too small, not representative, noisy, or polluted with irrelevant features.
- Your model needs to be neither too simple (under-fit) nor too complex (over-fit).

# Exercises


- ① • How would you define Machine Learning?
- ② • What is a labeled training set?
- ③ • Can you name four common unsupervised tasks?
- ④ • What type of Machine Learning algorithm would you use to allow a robot to walk in various unknown terrains?
- ⑤ • What type of algorithm would you use to segment your customers into multiple groups?
- ⑥ • What is an online learning system?
- ⑦ • What is the difference between a model parameter and a learning algorithm's hyperparameter?
- ⑧ • If your model performs great on the training data but generalizes poorly to new instances, what is happening? Can you name three possible solutions?
- ⑨ • What is the purpose of a validation set?  
↳ Finding and optimizing the best model to solve a given problem.

- ① giving computers the ability to learn without explicitly being programmed to.
- ② training set with its solutions.
- ③ clustering, visualization, dimensionality reduction, Association rule learning.
- ④ reinforcement learning.
- ⑤ supervised "if groups are labeled"  
↳ unsupervised "if no labels"

## End-to-End Machine Learning Project

- ⑥ online learning system: is a method of machine learning in which data becomes available in a sequential order and it's used to update the best predictor for future data at each step.  
Prof. Gheith Abandah
- ⑦ model parameters are estimated from data automatically, hyperparameters are set manually and used in process to help estimate model parameters.
- ⑧ overfitting, solutions: ① Cross validation  
② Regularization  
③ Simplify the model.

# Reference



- Chapter 2: End-to-End Machine Learning Project

- Aurélien Géron, **Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow**, O'Reilly, ~~2nd~~<sup>third</sup> Edition, 2019
  - Material: <https://github.com/ageron/handson-ml2>

2

## The 7 Steps of Machine Learning

- YouTube Video: The 7 Steps of Machine Learning from Google Cloud Platform

<https://youtu.be/nKW8Ndu7Miw>

*Caution: Alcohol is forbidden in the Islamic religion and causes addiction and has negative effects on health.*

3

# Outline

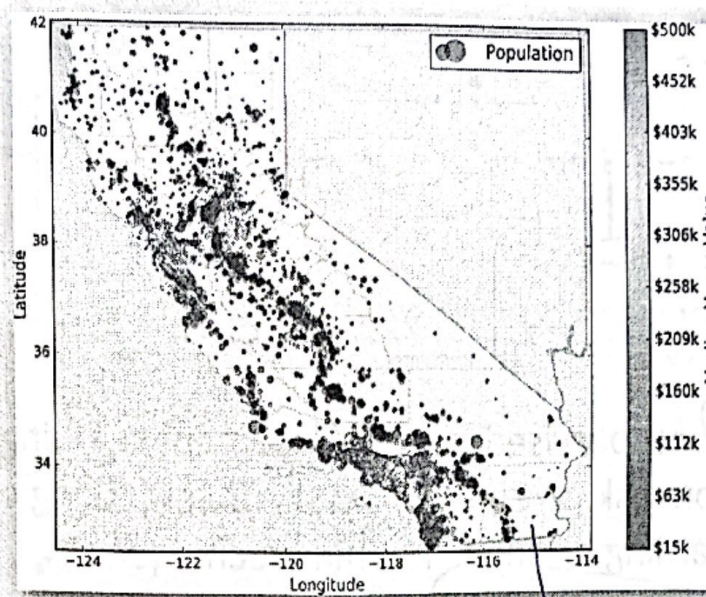
1. Look at the big picture
2. Get the data
3. Discover and visualize the data to gain insights
4. Prepare the data for Machine Learning algorithms
5. Select a model and train it
6. Fine-tune your model
7. Present your solution
8. Launch, monitor, and maintain your system
9. Exercises

## Working with Real Data

- Popular open data repositories:
  - [Tensorflow Datasets \(GitHub\)](#)
  - [UC Irvine Machine Learning Repository](#)
  - [Kaggle datasets](#)
  - [Amazon's AWS datasets](#)
  - [IEEE DataPort](#)
- Meta portals (they list open data repositories):
  - [Google Dataset Search](#)
  - <http://dataportals.org/>
  - <http://opendatamonitor.eu/>
  - <http://quandl.com/>
- Other pages listing many popular open data repositories:
  - [Wikipedia's list of Machine Learning datasets](#)
  - [Quora.com question](#)
  - [Datasets subreddit](#)

California  
↑

# 1. Look at the Big Picture: CA Housing Data

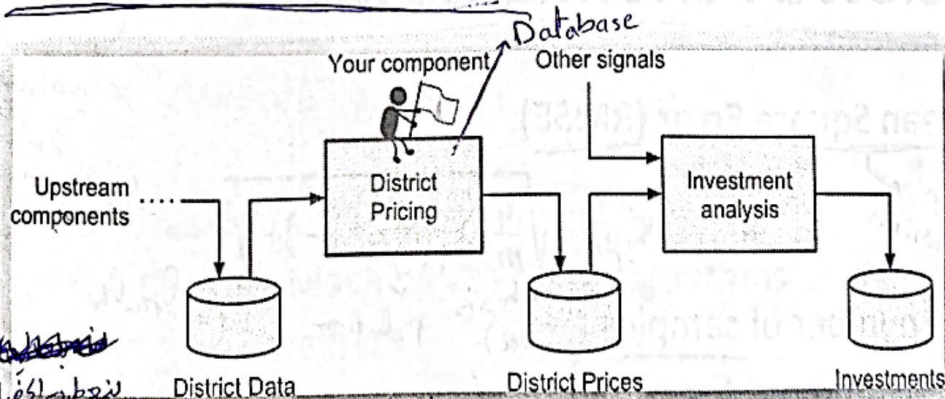


داتا ريزاليزيشن  
 visualization  
 عنوان زيمنه طيبه الراتا  
 قبل ما نستعمل علىها

data visualization

predict median house values in Californian districts, given a number of features from these districts.

## 1.1. Frame the Problem



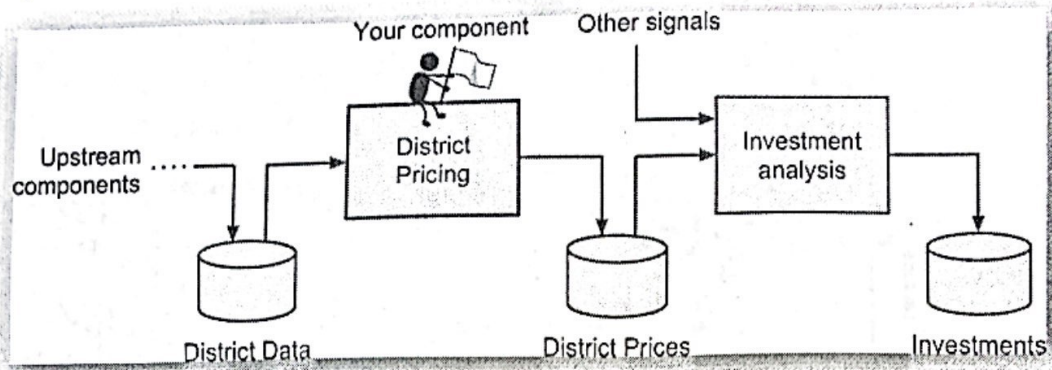
ننظر الكيفيه والاشياء  
 برنا سعر البيت بالارقام

Is it supervised, unsupervised, or Reinforcement Learning?  
 Is it a classification task, a regression task, or something else? Should you use batch learning or online learning techniques?  
Instance-based or Model-based learning?

لانه كل الراتا عندي وبنوريه  
 عن كل الراتا

better

# 1.1. Frame the Problem



Is it supervised, unsupervised, or Reinforcement Learning?  
 Is it a classification task, a regression task, or something else? Should you use batch learning or online learning techniques?  
 Instance-based or Model-based learning?

← بلزينا  
 distance measure  
 ← طرقة

← بنظير هو ايضا  
 ← عن طريق سعر  
 ← ال district مع  
 ← ال باقي و بشوف  
 ← ال ترب

# 1.2. Select a Performance Measure

• Root Mean Square Error (RMSE)

← يحتاجه اكثر  
 ← كما يكون  
 ← regression

← حالة ال  
 ← regression

$$RMSE(X, h) = \sqrt{\frac{1}{m} \sum_{i=1}^m (h(x^{(i)}) - y^{(i)})^2}$$

$\downarrow$  hypothesis (model)       $\downarrow$  Prediction       $\downarrow$  ground truth

- $m$  is the number of samples
- $x^{(i)}$  is the feature vector of Sample  $i$
- $y^{(i)}$  is the label or desired output
- $X$  is a matrix containing all the feature values

$$X = \begin{pmatrix} (x^{(1)})^T \\ (x^{(2)})^T \\ \vdots \\ (x^{(1999)})^T \\ (x^{(2000)})^T \end{pmatrix} = \begin{pmatrix} -118.29 & 33.91 & 1,416 \\ \vdots & \vdots & \vdots \end{pmatrix}$$

## 1.2. Select a Performance Measure

- Mean Absolute Error

$$\text{MAE}(\mathbf{X}, h) = \frac{1}{m} \sum_{i=1}^m |h(\mathbf{x}^{(i)}) - y^{(i)}|$$

- MAE is better than RMSE when there are outlier samples.

~~MAE is better than RMSE~~

Mean absolute error is better than root mean square error because if there is an outlier it'll have less impact but if ~~we~~ you don't have outliers use root mean square error.

10

## Outline

1. Look at the big picture
2. Get the data
3. Discover and visualize the data to gain insights
4. Prepare the data for Machine Learning algorithms
5. Select a model and train it
6. Fine-tune your model
7. Present your solution
8. Launch, monitor, and maintain your system
9. Exercises

11



## \*2. Get the Data 14/3

- If you didn't do it before, it is time now to **download** the **Jupyter notebooks** of the textbook from

<https://github.com/ageron/handson-ml2>

- Start Jupyter notebook and open Chapter 2 notebook.
- Hint: If you get kernel connection problem, try  
`C:\>jupyter notebook -port 8889`
- The following slides summarize the code used in this notebook.

## 2. Get the Data

1. Download the `housing.tgz` file from **Github** using `urllib.request.urlretrieve()` from the `urllib` package
2. Extract the data from this compressed tar file using `tarfile.open` and `extractall()`. The data will be in the CSV file `housing.csv`
3. Read the CSV file into a Pandas DataFrame called `housing` using `pandas.read_csv()`

## 2.1. Take a Quick Look at the Data Structure

- Display the top five rows using the DataFrame's `head()` method
- The `info()` method is useful to get a quick description of the data
- To find categories and repetitions of some column use `housing['key'].value_counts()`
- The `describe()` method shows a summary of the numerical attributes. *↳ Statistic*
- Show histogram using the `hist()` method and `matplotlib.pyplot.show()`

14

```
using.info()
```

```
class 'pandas.core.frame.DataFrame'>
RangeIndex: 20640 entries, 0 to 20639
Data columns (total 10 columns):
longitude      20640 non-null float64
latitude       20640 non-null float64
housing_median_age  20640 non-null float64
total_rooms    20640 non-null float64
total_bedrooms 20433 non-null float64
population     20640 non-null float64
households     20640 non-null float64
median_income  20640 non-null float64
median_house_value 20640 non-null float64
ocean_proximity 20640 non-null object
dtypes: float64(9), object(1)
memory usage: 1.6+ MB
```

207 missing features

```
>>> housing["ocean_proximity"].value_counts()
<1H OCEAN      9136
INLAND         6551
NEAR OCEAN     2658
NEAR BAY       2290
ISLAND          5
Name: ocean_proximity, dtype: int64
```

→ districts

15

## 2.2. Create a Test Set

- Split the available data randomly to:
  - Training set (80%)
  - Test set (20%) *must be representative*
- The example defines a function called `split_train_test()` for illustration.
- Scikit-Learn has `train_test_split()`.
- Scikit-Learn also has `StratifiedShuffleSplit()` that does stratified sampling.
- **Stratification** ensures that the test samples are representative of the target categories.

Preserve percentage of samples for each class

### 2.2.1. Create a Test Set: User-defined function

```
import numpy as np
def split_train_test(data, test_ratio):
    shuffled_indices = np.random.permutation(len(data))
    test_set_size = int(len(data) * test_ratio)
    test_indices = shuffled_indices[:test_set_size]
    train_indices = shuffled_indices[test_set_size:]
    return data.iloc[train_indices], data.iloc[test_indices]
```

You can then use this function like this:

```
>>> train_set, test_set = split_train_test(housing, 0.2)
>>> print(len(train_set), "train +", len(test_set), "test")
16512 train + 4128 test
```

`shuffled_indices[:test_set_size]` → من البداية لـ `test_set_size` و هو 4128  
`shuffled_indices[test_set_size:]` → من بعد `test_set_size` إلى النهاية

np.random.seed(42) → Save the state of randomness

generate same random numbers on multiple executions.

## 2.2.2. Create a Test Set: Using Scikit-Learn functions

```
from sklearn.model_selection import train_test_split
train_set, test_set = train_test_split(housing, test_size=0.2, random_state=42)
```

لانه موجود بال built-in libraries

Permutation 42

default value = 42 → magic number

Stratification is usually done on the target class.

```
from sklearn.model_selection import StratifiedShuffleSplit
split = StratifiedShuffleSplit(n_splits=1, test_size=0.2, random_state=42)
for train_index, test_index in split.split(housing, housing["income_cat"]):
    strat_train_set = housing.loc[train_index]
    strat_test_set = housing.loc[test_index]
```

to get the same train & test sets across different executions.

هون بنكليه ال data اللي بعلها split و بنكليه ال data اللي بعلها split column

### Outline

1. Look at the big picture
2. Get the data
3. Discover and visualize the data to gain insights
4. Prepare the data for Machine Learning algorithms
5. Select a model and train it
6. Fine-tune your model
7. Present your solution
8. Launch, monitor, and maintain your system
9. Exercises

80% class تدرج عليها  
 ال model ويفحص حاله فيها  
 مثلا لو كان عنا class 50%  
 نبيك بال data 50% و class تاني  
 نبيك 50%  
 80% تكون ال training set  
 و ال test set  
 مقسمة بين  
 ال 2 classes  
 بشكل متناسب  
 مع حجم ال data فيهم

# 3. Discover and Visualize the Data to Gain Insights

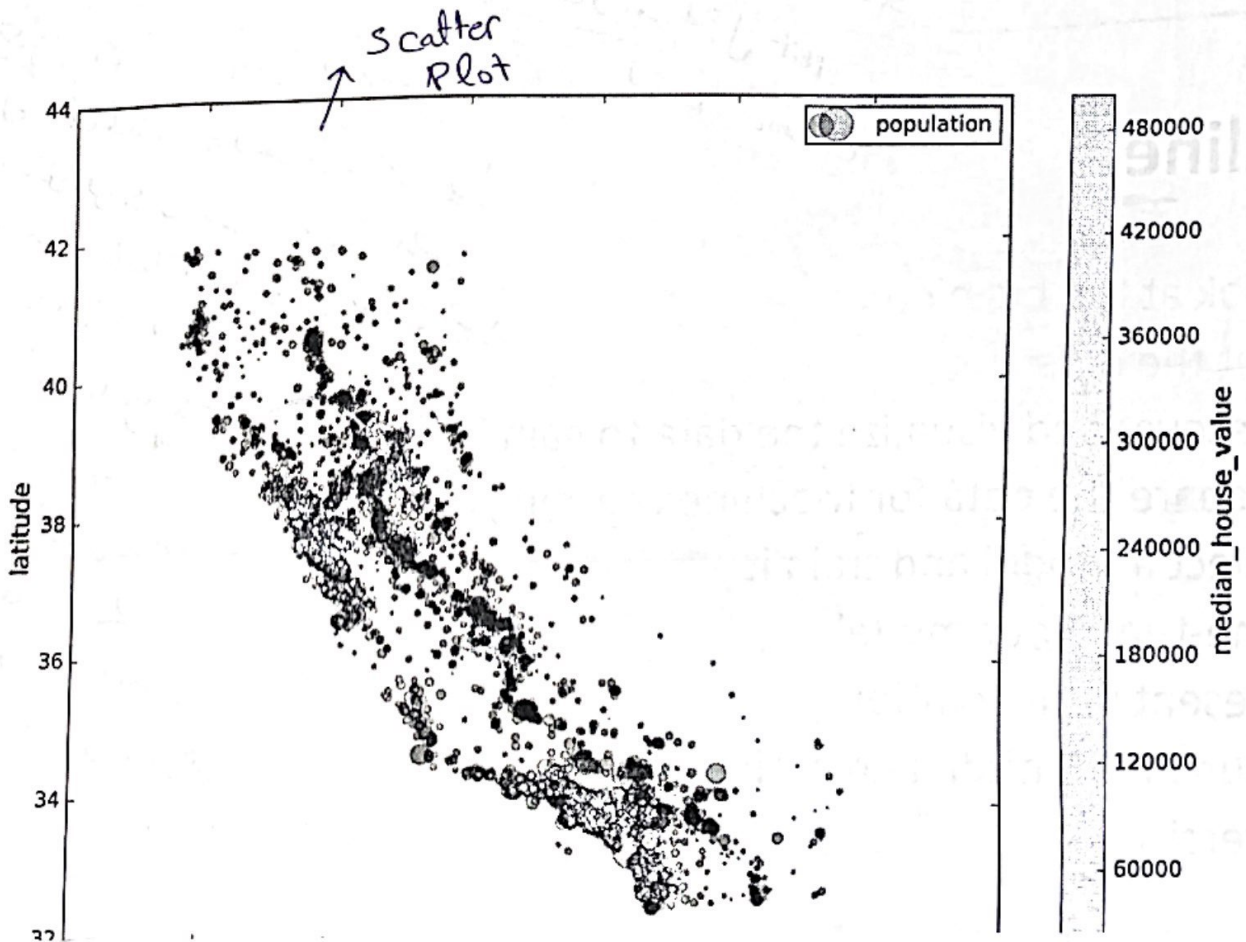
- Visualize geographical data using

```
housing.plot(kind="scatter", x="longitude", y="latitude", alpha=0.4,  
s=housing["population"]/100, label="population",  
c="median_house_value", cmap=plt.get_cmap("jet"), colorbar=True,  
)  
plt.legend()
```

البيانات الجغرافية

الشفافية  
alpha: Transparency, s: size, c: color, cmap: blue to red

size of the dot



Correlation can tell if two variables have a linear relationship



because I don't need all the features

### 3.1. Looking for Correlations

- Compute the standard correlation coefficient (also called Pearson's  $r$ ) between every pair of attributes using `corr_matrix = housing.corr()`

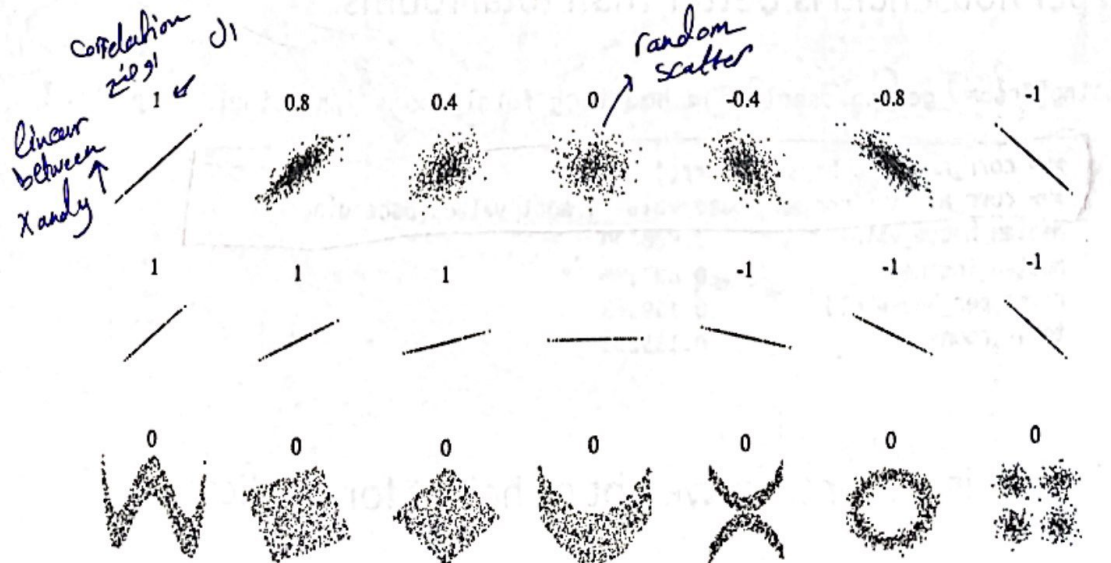
(to find correlations between each feature with others)  $r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}}$

```
>>> corr_matrix["median_house_value"].sort_values(ascending=False)
median_house_value    1.000000
median_income         0.687170
total_rooms           0.135231
housing_median_age   0.114220
households            0.064702
total_bedrooms        0.047865
population            -0.026699
longitude             -0.047279
latitude              -0.142826
```

- measured by
- ①. Corr Function
  - ②. scatter in pandas.

### 3.1. Looking for Correlations

- Zero linear correlation ( $r = 0$ ) does not guarantee independence.

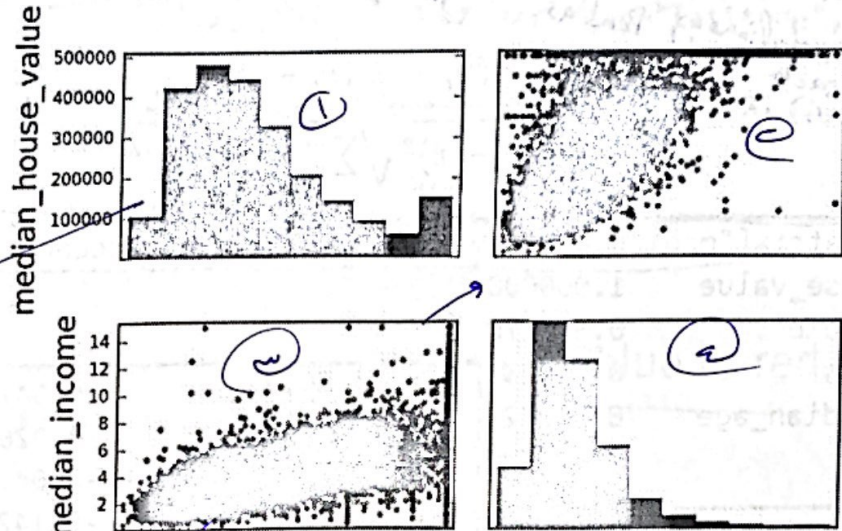


## 3.2. Pandas Scatter Matrix

*pandas.plotting*

```
from pandas.tools.plotting import scatter_matrix
attributes = ["median_house_value", "median_income"]
scatter_matrix(housing[attributes], figsize=(12, 8))
```

اللي بيها موجود  
ال correlations  
بيها



ال axis  
ال دي، ليا  
مع ال محور  
ال الثاني

histogram

correlation  
ما في داي بر  
ال feature  
ال دي = 1

Correlation  
Income House Value

## 3.3. Experimenting with Attribute Combinations

- Rooms per household is better than total rooms:

```
housing["rooms_per_household"] = housing["total_rooms"]/housing["households"]
```

```
>>> corr_matrix = housing.corr()
>>> corr_matrix["median_house_value"].sort_values(ascending=False)
median_house_value    1.000000
median_income         0.687170
rooms_per_household   0.199343
total_rooms           0.135231
```

- Similarly, BMI is better than weight or height for medical purposes.

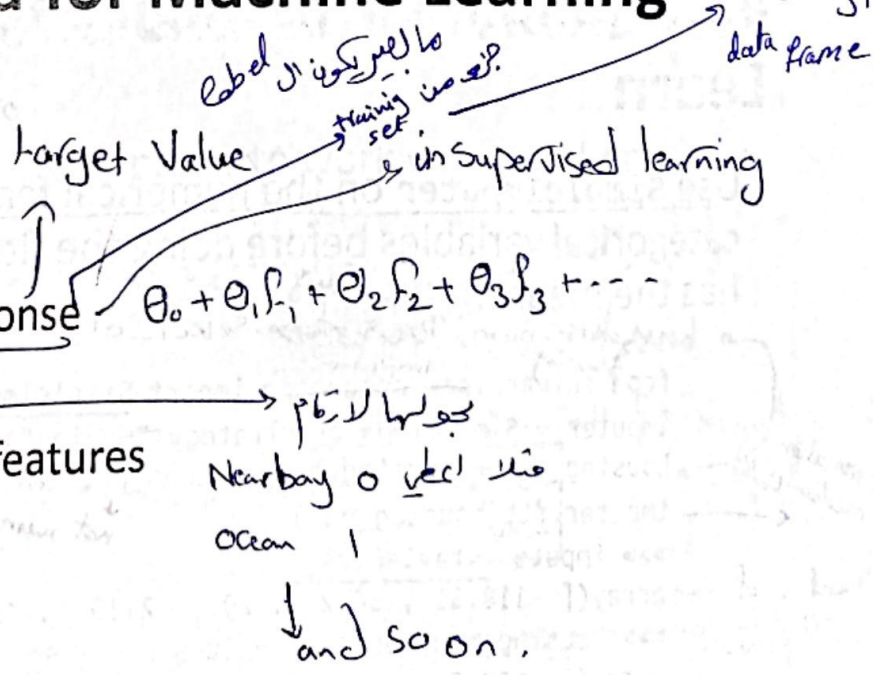
# Outline

1. Look at the big picture
2. Get the data
3. Discover and visualize the data to gain insights
4. Prepare the data for Machine Learning algorithms
5. Select a model and train it
6. Fine-tune your model
7. Present your solution
8. Launch, monitor, and maintain your system
9. Exercises

26

## 4. Prepare the Data for Machine Learning Algorithms

- Split to train and test (Done)
- Separate features from response
- Handle missing data
- Handle text and categorical features
- Scale (normalize) features
- Build preparation pipeline



27



# 4. Prepare the Data for Machine Learning Algorithms

- Separate the features from the response remove from training data set

```
housing = strat_train_set.drop("median_house_value", axis=1)
housing_labels = strat_train_set["median_house_value"].copy()
```

- Options of handling missing features:

1. Get rid of the corresponding districts
2. Get rid of the whole attribute
3. Set the values to some value (0, mean, median, etc.)

```
housing.dropna(subset=["total_bedrooms"]) # option 1
housing.drop("total_bedrooms", axis=1) # option 2
median = housing["total_bedrooms"].median() # option 3
housing["total_bedrooms"].fillna(median, inplace=True)
```

## 4.1. Handling Missing Features Using Scikit-Learn

- Use SimpleImputer on the numerical features. Need to remove categorical variables before doing the fit. The attribute `statistics` has the means.

```
housing_num = housing.select_dtypes(include=[np.number])
from sklearn.preprocessing import SimpleImputer
imputer = SimpleImputer(strategy="median")
housing_num = housing.drop("ocean_proximity", axis=1)
imputer.fit(housing_num)
>>> imputer.statistics_
array([-118.51, 34.26, 29., 2119., 433., 1164., 408., 3.5414])
>>> housing_num.median().values
array([-118.51, 34.26, 29., 2119., 433., 1164., 408., 3.5414])
X = imputer.transform(housing_num)
```

NumPy array

next slide

missing values feature كل feature مع missing values  
 Feature مع missing values مع كل feature

Convert the array X into dataframe

## 4.2. Handling Text and Categorical Attributes

- ocean\_proximity is categorical feature.

```
>>> housing_cat = housing[["ocean_proximity"]]
>>> housing_cat.head(10)
ocean_proximity
17606    <1H OCEAN
18632    <1H OCEAN
14650    NEAR OCEAN
3230     INLAND
3555     <1H OCEAN
19480    INLAND
8879     <1H OCEAN
13685    INLAND
4937     <1H OCEAN
4861     <1H OCEAN
```

indices why not 0, 1, 2--?

because of shuffling

Shuffling before splitting

## 4.2. Handling Text and Categorical Attributes

- Most machine learning algorithms prefer to work with numbers.

### Converting to numbers:

```
>>> from sklearn.preprocessing import OrdinalEncoder
>>> ordinal_encoder = OrdinalEncoder()
>>> housing_cat_encoded = ordinal_encoder.fit_transform(housing_cat)
>>> housing_cat_encoded[:10]
```

```
array([[0.],
       [0.],
       [4.],
       [1.],
       [0.],
       [1.],
       [0.],
       [1.],
       [0.],
       [0.]])
```

Numerical values imply distances

```
>>> ordinal_encoder.categories_
array(['<1H OCEAN', 'INLAND', 'ISLAND', 'NEAR BAY', 'NEAR OCEAN'],
      dtype=object)
```

رتبه هم تعدادی حسب ال code Asc

problem in ordinal encoding ?

الارقام التي بعدها حقا المعنى وما إلى علاقة بين

# 4.2. Handling Text and Categorical Attributes

categories عدال = digits عدال Common

• To ensure encoding neutrality, we can use the one-hot encoding.

```
>>> from sklearn.preprocessing import OneHotEncoder
>>> cat_encoder = OneHotEncoder()
>>> housing_cat_1hot = cat_encoder.fit_transform(housing_cat)
>>> housing_cat_1hot
<16512x5 sparse matrix of type '<class 'numpy.float64'>'
with 16512 stored elements in Compressed Sparse Row format>
>>> housing_cat_1hot.toarray()
array([[1., 0., 0., 0., 0.],
       [1., 0., 0., 0., 0.],
       [0., 0., 0., 0., 1.],
       ...,
       [0., 1., 0., 0., 0.],
       [1., 0., 0., 0., 0.],
       [0., 0., 0., 1., 0.]])
```

عدال rows = عدال columns  
 عدال rows = عدال columns  
 نوع من ال Optimization  
 يعني قدام يكون عازف  
 انواع matrix  
 طبانه اصغار عشان  
 اقدر اوفر بال memory

Converts sparse matrix to dense matrix.

يكون فيها عدال اصغار كثير  
 اذا كان عندي اصغار كثير قدام  
 بخرت بطريقة Compact  
 بخرت وبتوان ones موجودين (الاصغار)

• info →  
 attributes عدال  
 values ال قيم

• يحتاج ال اعداد تكون numerical  
 بالفاية يكون عننا  
 mathematical model

# 4.3. Custom Transformers

- Scikit-Learn allows you to create your own transformers.
- You can create a transformer to create derived features.
- Create a class and implement three methods: fit(), transform(), and fit\_transform(). Include base classes:
  - TransformerMixin to get fit\_transform()
  - BaseEstimator to get get\_params() and set\_params()

Fit then Transform

### 4.3. Custom Transformers

```

from sklearn.base import BaseEstimator, TransformerMixin
rooms_ix, household_ix = 3, 6

class CombinedAttributesAdder(BaseEstimator, TransformerMixin):
    def fit(self, X, y=None):
        return self # nothing else to do
    def transform(self, X, y=None):
        rooms_per_household = X[:, rooms_ix] / X[:, household_ix]
        return np.c_[X, rooms_per_household]

attr_adder = CombinedAttributesAdder()
housing_extra_attribs = attr_adder.transform(housing.values)
    
```

get and set parameters

fit\_transform

base classes to inherit some methods from.

Column

columns

number array

in Pandas

ما احتجنا ال fit لانه فينا column ال column ال average ال

describe -> show min / average / max / Percentaise for each column

### 4.4. Feature Scaling

features ال يكون في تفاوت بالقيم ال scaling نفضل نعمل

- ML algorithms generally don't perform well when the input numerical attributes have very different scales.

تزيد ال accuracy ويكون ال اوزان ادق ال features

Scaling techniques:

- Min-max scaling -> ~~more sensitive to outliers values~~
  - more sensitive to outliers values

$$x' = \frac{x - \min(x)}{\max(x) - \min(x)}$$

sample

- Standardization (standard scaler)

more popular  
better when we have outlier values


$$x' = \frac{x - \bar{x}}{\sigma}$$

sample  
average  
standard deviation

بالمقارنة بين كل ال feature

# 4.5. Transformation Pipelines

→ Fill missing values



السلسلة Pipeline

list of parameters "Stages" of the pipeline.

```
from sklearn.pipeline import Pipeline
from sklearn.preprocessing import StandardScaler
```

Work on numeric values.

```
num_pipeline = Pipeline([
    ('imputer', SimpleImputer(strategy="median")),
    ('attribs_adder', CombinedAttributesAdder()),
    ('std_scaler', StandardScaler()),
])
```

بعد الـ fit

```
housing_num_tr = num_pipeline.fit_transform(housing_num)
```

housing\_num - training

Data Frame of numeric values

without ocean proximity column.

Column of Strings

# 4.6. Full Pipeline

```
from sklearn.compose import ColumnTransformer
```

بخرن فيه اسما Columns

numeric values

Super Scaler

```
num_attribs = list(housing_num)
cat_attribs = ["ocean_proximity"]
full_pipeline = ColumnTransformer([
    ("num", num_pipeline, num_attribs),
    ("cat", OneHotEncoder(), cat_attribs),
])
```

```
housing_prepared = full_pipeline.fit_transform(housing)
```

Dense array

```
as : num_pipeline.fit_transform(housing[num_attribs])
```

go read ColumnTransformer sklearn → remainder function.

# Outline

1. Look at the big picture
2. Get the data
3. Discover and visualize the data to gain insights
4. Prepare the data for Machine Learning algorithms
5. Select a model and train it
6. Fine-tune your model
7. Present your solution
8. Launch, monitor, and maintain your system
9. Exercises


38

## 5. Select and Train a Model

- Linear regressor → linear relationship between feature and label
- Using RMSE for evaluation
- Decision tree regressor → for non-linear relationships
- k-fold cross validation
- Random forests regressor

Root Mean Square error

Decision trees مجموعة من ال



39

# 5. Select and Train a Model

- Let us start by training a simple linear regressor.

```
from sklearn.linear_model import LinearRegression  
lin_reg = LinearRegression()  
lin_reg.fit(housing_prepared, housing_labels)
```

Annotations:   
- "object" points to LinearRegression()  
- "البيانات التي نريد التوقعها" (the data we want to predict) points to housing\_labels  
- "ما فيها البيانات التي نريد ان نتعلمها" (the data we want to learn from) points to housing\_prepared  
- "Data" points to housing\_prepared

- Try it out on five instances from the training set.

```
>>> some_data = housing.iloc[:5]  
>>> some_labels = housing_labels.iloc[:5]  
>>> some_data_prepared = full_pipeline.transform(some_data)  
>>> print("Predictions:\t", lin_reg.predict(some_data_prepared))  
Predictions: [ 303104.  44800.  308928.  294208.  368704.]  
>>> print("Labels:\t\t", list(some_labels))  
Labels: [359400.0, 69700.0, 302100.0, 301300.0, 351900.0]
```

Annotations:   
- "data" points to some\_data  
- "labels" points to some\_labels  
- "5 samples" points to [:5]

50% off

Annotations:   
- "by pipeline" points to full\_pipeline.transform  
- "actual values" points to list(some\_labels)  
- "all data preparation on training data must be also done on Test data or you can just do all preparations on the whole data before you split it into train & test."

## 5.1. Evaluate the Model on the Entire Training Set

Annotations:   
- "LinearRegression" points to the class name  
- "RMSE" points to the metric

- Use RMSE manually

```
>>> from sklearn.metrics import mean_squared_error  
>>> housing_predictions = lin_reg.predict(housing_prepared)  
>>> lin_mse = mean_squared_error(housing_labels, housing_predictions)  
>>> lin_rmse = np.sqrt(lin_mse)  
>>> lin_rmse  
68628.413493824875
```

This is not a satisfactory result as the median\_housing\_values range between \$120,000 and \$265,000.

Since this is huge error, we will try another Model which is ~~Decision~~ Decision Tree Regressor.

## 5.2. Try the Decision Tree Regressor



```
from sklearn.tree import DecisionTreeRegressor
tree_reg = DecisionTreeRegressor()
tree_reg.fit(housing_prepared, housing_labels)

>>> housing_predictions = tree_reg.predict(housing_prepared)
>>> tree_mse = mean_squared_error(housing_labels, housing_predictions)
>>> tree_rmse = np.sqrt(tree_mse)
>>> tree_rmse
0.0
```

**Overfitting: It has memorized the entire training set!**

عمل Predict على نفس البيانات  
التي تم تدريب عليها فخطأه ان 0 = error

↳ If I want to test it, give the model test data that it ~~had~~ had never seen before and check its prediction



AiPages 139 - 150 From the book

• FunctionTransformer → estimator takes certain function we want to apply on a certain column as a parameter "we have to input it." → transform the data according to a certain function

optional

ex:-  $\text{log\_transformer} = \text{FunctionTransformer}(\text{np.log}, \text{inverse\_func}=\text{np.exp})$   
 $\text{log\_pop} = \text{log\_transformer}(\text{housing}["\text{Population}"])$

Takes the log of the elements in Population column.

• usually we use log when we have heavy tail. (then scaling).

• rbf\_kernel → similarity measure function.

• fit function → calculates the mean, median ... etc

• transform → apply operations ~~on the mean~~ that includes the mean, ... etc

• check\_array(x) → checks that x is an array with finite float values

• check\_is\_fitted(var) → checks that the parameters are already initialized.

↓ default\_num\_pipeline → imputation then scaling

• get\_feature\_names\_out() → returns columns names

• ratio\_pipeline() → calculates the ratio between 2 columns. returns the word "ratio"

## 5.1. Evaluate the Model on the Entire Training Set

```
>>> from sklearn.metrics import mean_squared_error
>>> lin_rmse = mean_squared_error(housing_labels, housing_predictions,
...                               squared=False)
...
>>> lin_rmse
68687.89176589991
```

If squared = false, RMSE  
If squared = true, MSE

42

## 5.2. Try the Decision Tree Regressor

```
from sklearn.tree import DecisionTreeRegressor

tree_reg = make_pipeline(preprocessing, DecisionTreeRegressor(random_state=42))
tree_reg.fit(housing, housing_labels)
```

Now that the model is trained, you evaluate it on the training set:

```
>>> housing_predictions = tree_reg.predict(housing)
>>> tree_rmse = mean_squared_error(housing_labels, housing_predictions,
...                               squared=False)
...
...
>>> tree_rmse
0.0
```

## 5.3. Better Evaluation Using Cross-Validation

Model to train




```
from sklearn.model_selection import cross_val_score  
  
tree_rmse = -cross_val_score(tree_reg, housing, housing_labels,  
                             scoring="neg_root_mean_squared_error", cv=10)
```

```
>>> pd.Series(tree_rmse).describe()  
count    10.000000  
mean     66868.027288  
std      2060.966425  
min      63649.536493  
25%      65338.078316  
50%      66801.953094  
75%      68229.934454  
max      70094.778246  
dtype: float64
```

Negative root mean squared  
error  
,Higher is better  
scoring هيك مبدأ ال

• Cross validation (CV)  $\rightarrow$  cross validation rounds

if  $CV = 3 \rightarrow$  3 rounds (it'll divide the data into 3 parts, 3 times)



# 5.4. Try the Random Forests Regressor

- Repeating training and evaluation:

```
from sklearn.ensemble import RandomForestRegressor
```

```
forest_reg = make_pipeline(preprocessing,  
                             RandomForestRegressor(random_state=42))  
forest_rmuses = -cross_val_score(forest_reg, housing, housing_labels,  
                                  scoring="neg_root_mean_squared_error", cv=10)
```

```
>>> pd.Series(forest_rmuses).describe()
```

```
count    10.000000  
mean     47019.561281  
std      1033.957120  
min      45458.112527  
25%     46464.031184  
50%     46967.596354
```

Best Accuracy  
Overfitting??

مجموعه من ال decision وكل واحد يعطي الجواب تبعه  
HFB

## 5.4. Try the Random Forests Regressor

- Repeating training and evaluation:

```
>>> from sklearn.ensemble import RandomForestRegressor
>>> forest_reg = RandomForestRegressor()
>>> forest_reg.fit(housing_prepared, housing_labels)
>>> [...]
>>> forest_rmse
18603.515021376355 ] → on training data
>>> display_scores(forest_rmse_scores)
Scores: [49519.80364233 47461.9115823 50029.02762854 52325.2806895
49308.39426421 53446.37892622 48634.8036574 47585.73832311
53490.10699751 50021.5852922 ]
Mean: 50182.303100336096
Standard deviation: 2097.0810550985693
→ on test data.
```

overfitting ← test  
الما يكون جواب ال training افضل من جواب ال test  
ال score

### Outline

الما يكونوا الجوابين سيئين - underfitting

1. Look at the big picture
2. Get the data
3. Discover and visualize the data to gain insights
4. Prepare the data for Machine Learning algorithms
5. Select a model and train it
6. Fine-tune your model
7. Present your solution
8. Launch, monitor, and maintain your system
9. Exercises

ال parameter الوحيه اليه ال  
Training data → default value  
and labels

لو بدنا نعدل ال hyperparameters بختناج عدد تجارب كبير

# 6. Fine-Tune Your Model

- Fine-tune your system by fiddling with:
  - The hyperparameters
  - Removing and adding features
  - Changing feature preprocessing techniques
- Can experiment manually. But it is best to automate this process using Scikit-Learn:

- GridSearchCV  $\rightarrow$  Cross Validation
  - or RandomizedSearchCV
- يعني ال parameters اللي بي اجرهم و لقيم  
اللي بي اعطوك لهاي ال Parameters و بهل تجارب ال Parameters

## 6.1. Grid Search

- Can automate exploring a search space of  $3 \times 4 + 2 \times 3 = 12 + 6 = 18$

```

from sklearn.model_selection import GridSearchCV
param_grid = [
    {'n_estimators': [3, 10, 30], 'max_features': [2, 4, 6, 8]},
    {'bootstrap': [False], 'n_estimators': [3, 10], 'max_features': [2, 3, 4]}
]
forest_reg = RandomForestRegressor()
grid_search = GridSearchCV(forest_reg, param_grid, cv=5,
                           scoring='neg_mean_squared_error',
                           return_train_score=True)
grid_search.fit(housing_prepared, housing_labels)
    
```

اول dictionary (12 تجارب)  
 عدد ال decision trees  
 list of dictionaries  
 ال ثاني dictionary  
 ال 11 تجارب  
 كل تجريب  
 5 مرات و 4 مرات

بهل كل التجارب  
بجرا بطبع النتائج



## 6.2 Examine the Results of Your Grid Search

- Can examine the best hyperparameters using:

```
>>> grid_search.best_params_
{'max_features': 8, 'n_estimators': 30}
```

*returns best result*

- Can examine all search results using:

```
>>> cvres = grid_search.cv_results_
>>> for mean_score, params in zip(cvres["mean_test_score"], cvres["params"]):
...     print(np.sqrt(-mean_score), params)
...
63669.05791727153 {'max_features': 2, 'n_estimators': 3}
55627.16171305252 {'max_features': 2, 'n_estimators': 10}
...
49682.25345942335 {'max_features': 8, 'n_estimators': 30}
```

*returns dictionary includes each set of parameters and their scores.*

Best Tuned Accuracy

## 6.2 Evaluate Your System on the Test Set

- The final model is the best estimator found by the grid search.
- To evaluate it on the test set, transform the test features, predict using transformed features, and evaluate accuracy.

```
→ final_model = grid_search.best_estimator_
X_test = strat_test_set.drop("median_house_value", axis=1)
y_test = strat_test_set["median_house_value"].copy()
X_test_prepared = full_pipeline.transform(X_test)
final_predictions = final_model.predict(X_test_prepared)
final_mse = mean_squared_error(y_test, final_predictions)
final_rmse = np.sqrt(final_mse) # => evaluates to 48,209.6
```

*returns best model*

Better than train set!

## 6.1. Grid Search (Updated)

- Can automate exploring a search space of  $3 \times 3 + 2 \times 3 = 9 + 6 = 15$

مكون من اكثر من pipeline  
وحسب طبيعة ال column كل  
pipeline ياخذ المناسب

```
from sklearn.model_selection import GridSearchCV
```

```
full_pipeline = Pipeline([
    ("preprocessing", preprocessing),
    ("random_forest", RandomForestRegressor(random_state=42)),
])
```

كل الداتا بتفوت على  
random forest regressor

الاسم الي بعد \_\_ هو  
Subclass  
من الاسم الي قبل \_\_

```
param_grid = [
    {'preprocessing__geo__n_clusters': [5, 8, 10],
     'random_forest__max_features': [4, 6, 8]},
    {'preprocessing__geo__n_clusters': [10, 15],
     'random_forest__max_features': [6, 8, 10]},
]
```

pipeline اسمه geo هو subclass من ال pipeline  
الي اسمه preprocessing,  
و n\_clusters هو parameter في geo.

```
grid_search = GridSearchCV(full_pipeline, param_grid, cv=3,
                           scoring='neg_root_mean_squared_error')
grid_search.fit(housing, housing_labels)
```

51

## 6.2 Examine the Results of Your Grid Search (Updated)

- Can examine the best hyperparameters using:

```
>>> grid_search.best_params_
{'preprocessing__geo__n_clusters': 15, 'random_forest__max_features': 6}
```

- Can examine all search results using:

```
>>> cv_res = pd.DataFrame(grid_search.cv_results_)
>>> cv_res.sort_values(by="mean_test_score", ascending=False, inplace=True)
>>> [...] # change column names to fit on this page, and show rmse = -score
>>> cv_res.head() # note: the 1st column is the row ID
```

	n_clusters	max_features	split0	split1	split2	mean_test_rmse
12	15	6	43460	43919	44748	44042
13	15	8	44132	44075	45010	44406
14	15	10	44374	44286	45316	44659
7	10	6	44683	44655	45657	44999
9	10	6	44683	44655	45657	44999

52

## 6.2 Evaluate Your System on the Test Set

- The final model is the best estimator found by the grid search.

```
>>>final_model = grid_search.best_estimator
```

- If GridSearchCV is initialized with **refit=True** (which is the default), then once it finds the best estimator using cross-validation, it retrains it on the whole training set.

When **refit=True**  
يعني باخذ ال best model بالنهاية  
وبدربه على كل الداتا  
لانه مثلا في حالة CV=3 تم تدريبيه  
على ثلثين الداتا فقط  
وفي حالة كان عنا ٨٠٪ من الداتا  
training و ٢٠٪ test في حال **refit**  
بتم تدريبيه بالنهاية على ١٠٠٪ من الداتا

53

## 6.2 Evaluate Your System on the Test Set

```
X_test = strat_test_set.drop("median_house_value", axis=1)  
y_test = strat_test_set["median_house_value"].copy()
```

```
final_predictions = final_model.predict(X_test)
```

```
final_rmse = mean_squared_error(y_test, final_predictions, squared=False)  
print(final_rmse) # prints 41424.40026462184
```

No overfitting



بس في حال كانت القيمة كثير بعيدة  
عن القيمة المتوقعة يعني الداتا غير  
كافية يعني في under-fitting

54

grid search

parameters لكل ال  
combinations الي بنعطيه اياهم  
وعلى عددهم بعمل التجارب  
في ال RandomizedSearch هاد  
بختار randomly عدد التجارب

## 6.2 Randomized Search

- Preferable, especially when the hyperparameter search space is large
- Run certain number of iterations
- Picks the hyperparameters values from the defined space.

```
from sklearn.model_selection import RandomizedSearchCV  
from scipy.stats import randint
```

```
param_distributions = {'preprocessing__geo__n_clusters': randint(low=3, high=50),  
                       'random_forest__max_features': randint(low=2, high=20)}
```

```
rnd_search = RandomizedSearchCV(  
    full_pipeline, param_distributions=param_distributions, n_iter=10, cv=3,  
    rnd_search.fit(housing, housing_labels) r, random_state=42)
```

بختار random number بين ال ٣  
وال ٥٠

55

## 6.3 Save Your Best Model for the Production System

- Save the model

```
import joblib
```

```
joblib.dump(final_model, "my_california_housing_model.pkl")
```

- Load the model      لما بدنا نعمله test على داتا جديدة

```
final_model_reloaded = joblib.load("my_california_housing_model.pkl")
```

```
new_data = [...] # some new districts to make predictions for  
predictions = final_model_reloaded.predict(new_data)
```

56

## 7. Present Your Solution

- Present your solution highlighting:
  - What you have learned
  - What worked and what did not
  - What assumptions were made
  - What your system's limitations are
- Document everything, and create nice presentations with:
  - Clear visualizations (*scaling / size of the data/encoding*)
  - Easy-to-remember statements, e.g., "the median income is the number one predictor of housing prices".

## 8. Launch, Monitor, and Maintain Your System

- Prepare your production program that uses your best trained model and launch it.
- Monitor the accuracy of your system. Also monitor the input data.
- Retrain your system periodically using fresh data.

# Summary

1. Look at the big picture
2. Get the data
3. Discover and visualize the data to gain insights
4. Prepare the data for Machine Learning algorithms
5. Select a model and train it
6. Fine-tune your model
7. Present your solution
8. Launch, monitor, and maintain your system
9. Exercises

54

## Exercise

- Try a Support Vector Machine regressor (`sklearn.svm.SVR`), with various hyperparameters such as `kernel="linear"` (with various values for the `C` hyperparameter) or `kernel="rbf"` (with various values for the `C` and `gamma` hyperparameters). Don't worry about what these hyperparameters mean for now. How does the best SVR predictor perform?

# Classification

Prof. Gheith Abandah

## Reference

- Chapter 3: **Classification**

- Aurélien Géron, **Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow**, O'Reilly, 2nd Edition, 2019
  - Material: <https://github.com/ageron/handson-ml2>

O'REILLY  
Hands-On  
Machine Learning  
with Scikit-Learn,  
Keras & TensorFlow  
Concepts, Tools, and Techniques  
to Build Intelligent Systems



# Introduction

- YouTube Video: Machine Learning - Supervised Learning Classification from Cognitive Class

<https://youtu.be/Lf2bCQlktTo>

• classification → discrete values called classes

• مثلاً لو رخطيه صورة، دكينا اختيارت شو ممكنة تكون فبعضنا، اكنينا، الى الى اي احتمال  
او ممكن نقارن الصورة مع اقرب حتمس صورة.

3

## Outline

1. MNIST dataset →
2. Training a binary classifier
3. Performance measures
4. Multiclass classification
5. Multilabel classification
6. Exercise

هي عبارة عن dataset مشهورة تستعمل عادة في بداية اي موضوع في ال ML وهي عبارة عن 70,000 صورة بكونوا hand written digits كل صورة عبارة عن 28X28 pixels

4



# 1. MNIST Dataset

- **MNIST** is a set of 70,000 small images of **handwritten digits**.
- Available from [mldata.org](http://mldata.org)
- **Scikit-Learn** provides **download** functions.



5

## 1.1. Get the Data

\*Fetching data needs Internet connection

```
from sklearn.datasets import fetch_openml
```

```
mnist = fetch_openml('mnist_784', as_frame=False)
```

Fetch the data from sklearn dataset

To get the data as Numpy array not Dataframe

Or we can use( make ) to generate data

(Load ) loads the data from your device

6

## 1.2. Extract Features and Labels

- There are 70,000 images, and each image has **784** features. This is because each image is **28×28** pixels, and each feature simply represents one pixel's intensity, from **0 (white)** to **255 (black)**.

مثلا اول row بمثل رقم 5 والثاني 0 وهكذا

كل row عبارة عن صورة

```
>>> X, y = mnist.data, mnist.target
>>> X
array([[0., 0., 0., ..., 0., 0., 0.],
       [0., 0., 0., ..., 0., 0., 0.],
       [0., 0., 0., ..., 0., 0., 0.],
       ...,
       [0., 0., 0., ..., 0., 0., 0.],
       [0., 0., 0., ..., 0., 0., 0.],
       [0., 0., 0., ..., 0., 0., 0.]])
>>> X.shape
(70000, 784)
>>> y
array(['5', '0', '4', ..., '4', '5', '6'], dtype=object)
>>> y.shape
(70000,)
```

## 1.3. Examine One Image

```
import matplotlib.pyplot as plt
```

```
def plot_digit(image_data):
    image = image_data.reshape(28, 28)
    plt.imshow(image, cmap="binary")
    plt.axis("off")
    some_digit = X[0]
    plot_digit(some_digit)
    plt.show()
```

```
>>> y[0]
'5'
```



## 1.4. Split the Data

- The MNIST dataset is actually already split into a **training set** (the first 60,000 images) and a **test set** (the last 10,000 images).

The training set is already shuffled.

↳ That's why we split in this way -

$X_{train}, X_{test}, y_{train}, y_{test} = X[:60000], X[60000:], y[:60000], y[60000:]$

↓  
rows

In general, in machine learning:

X → For data

Y → For labels



## Outline

1. MNIST dataset
2. Training a binary classifier
3. Performance measures
4. Multiclass classification
5. Multilabel classification
6. Exercise

Binary classifier

↓  
No or yes جواب

(two classes)

تقسيم على 2 classes y- labels  
true or false  
على الـ 2



## 2. Training a Binary Classifier

→ answer is True or False (Yes or No)

- A binary classifier can classify two classes.
- For example, classifier for the number 5, capable of distinguishing between two classes, 5 and not-5.

```
y_train_5 = (y_train == 5)
y_test_5 = (y_test == 5)

from sklearn.linear_model import SGDClassifier
sgd_clf = SGDClassifier(random_state=42)
sgd_clf.fit(X_train, y_train_5)

>>> sgd_clf.predict([some_digit])
array([ True])
```

Labels 5 or not 5 → For every element.

True for all 5s, False for all other digits.

model → SGDClassifier

data labels → sgd\_clf.fit(X\_train, y\_train\_5)

fit → boundaries, curve لحد  
جهد الـ 5 ان يجي  
(decision boundaries) جهد الـ 5 ان يجي

return → binary array with 60000 rows including values of True & False

astype(c) → to change the type of data.

### Outline

1. MNIST dataset
2. Training a binary classifier
3. Performance measures
4. Multiclass classification
5. Multilabel classification
6. Exercise

مقياس أداء  
Classifier dl

### 3. Performance Measures

correct answers / # of questions

عبر كذا فيه

- Accuracy: Ratio of correct predictions → over all predictions.
- Confusion matrix
- Precision and recall
- F1 Score
- Precision/recall tradeoff

### 3.1. Accuracy

Classifier

• %\_test → فيها الإجابات التي

• %\_predict → التي توقعها model

```

y_pred = clone_clf.predict(X_test_fold)
n_correct = sum(y_pred == y_test_fold) → returns binary array
print(n_correct / len(y_pred))

```

correct predictions

total number of predictions

Example how to find the accuracy. (manual)

another method

```

>>> from sklearn.model_selection import cross_val_score
>>> cross_val_score(sgd_clf, X_train, y_train_5, cv=3, scoring="accuracy")
array([0.96355, 0.93795, 0.95615])

```

fold accuracy

Using the cross\_val\_score() function to find the accuracy on three folds

sum of binary array

True

check in documentation what are other options for this parameter?

لما نطبق اي algorithm لازم يكون على baseline تكون الامتداد

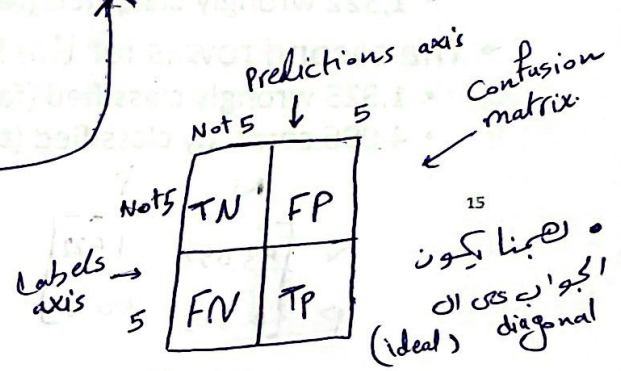

### 3.1. Accuracy

الادفضل نستعملها او cross\_val\_score  
 ما شيجي ندر البرات  
 عنان نقيصها

- Use `cross_val_predict()` to predict the targets of the entire training set.

```
from sklearn.model_selection import cross_val_predict
```

```
y_train_pred = cross_val_predict(sgd_clf, X_train, y_train_5, cv=3)
```



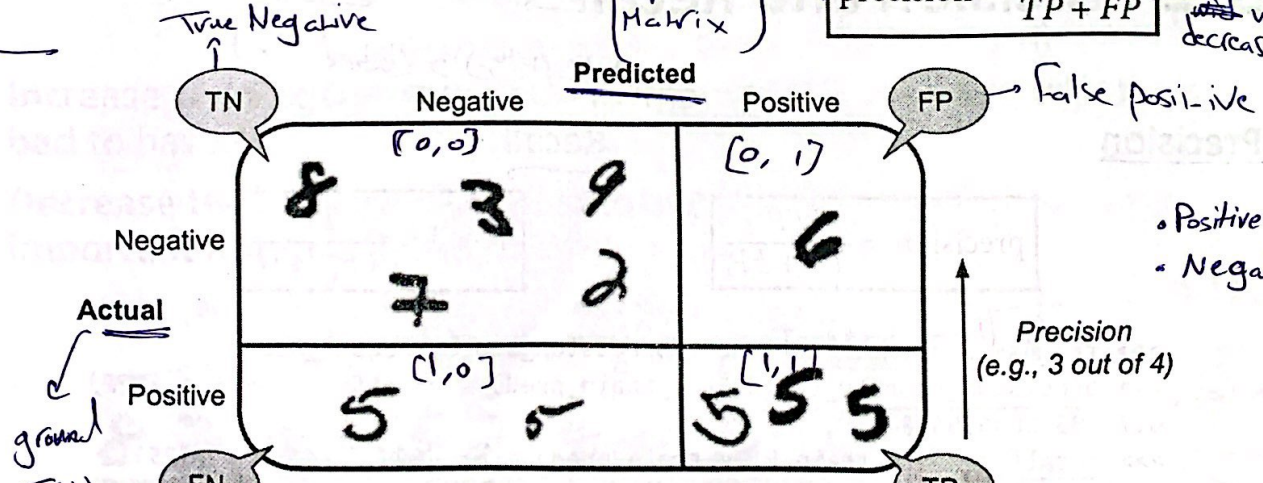
$$\text{accuracy} = \frac{TN + TP}{TN + TP + FN + FP}$$

### 3.2. Confusion Matrix

Square Matrix

$$\text{precision} = \frac{TP}{TP + FP}$$


↑ increase FP will decrease Precision



$$\text{recall} = \frac{TP}{TP + FN}$$

Recall (e.g., 3 out of 5) Ratio of TP for all positive (Actual)

if all predictions are right the diagonal of Matrix will only have values, others zero



اذا كان ال recall قليل في غير FN

من كل ال 5 كم وحدة ووقوعه 3

لما بدنا نعمل system الافضل من حيثى recall ونقل FN  
 (recall) من حيثى FN  
 فى حال ركزنا على Precision لازم نقل FP  
 وحسب ال system بنفنا Precision او recall او الاثنين

### 3.2. Confusion Matrix

- Scikit Learn has a function for finding the confusion matrix.

```
>>> from sklearn.metrics import confusion_matrix
>>> confusion_matrix(y_train_5, y_train_pred)
array([[53057, 1522],
       [1325, 4096]])
```

- The first row is for the non-5s (the negative class):
  - 53,057 correctly classified (true negatives)
  - 1,522 wrongly classified (false positives)
- The second row is for the 5s (the positive class):
  - 1,325 wrongly classified (false negatives)
  - 4,096 correctly classified (true positives)

	N	P
N	53,057	1,522
P	1,325	4,096

### 3.3. Precision and Recall

(Cross-val. Predict. الو Prediction ال) بسط  
 ال Label وال Prediction

Precision

Recall

$$\text{precision} = \frac{TP}{TP + FP}$$

$$\text{recall} = \frac{TP}{TP + FN}$$

open documentation

```
>>> from sklearn.metrics import precision_score, recall_score
>>> precision_score(y_train_5, y_train_pred) # == 4096 / (4096 + 1522)
0.7290850836596654
>>> recall_score(y_train_5, y_train_pred) # == 4096 / (4096 + 1325)
0.7555801512636044
```

The precision and recall are smaller than the accuracy.

Why?

FN ال و FP ال

accuracy-score (yt - - - ) Comparable

TP ال

فرق (y+1-est, y-pred)

### 3.4. F1 Score

بالهناك اهتمينا  
 بال recall وال precision  
 بهناك ال recall وال precision  
 بهناك ال recall وال precision  
 بهناك ال recall وال precision

- The **F1 Score** combines the precision and recall in one metric (harmonic mean). → both precision & recall must be high to get high harmonic mean.

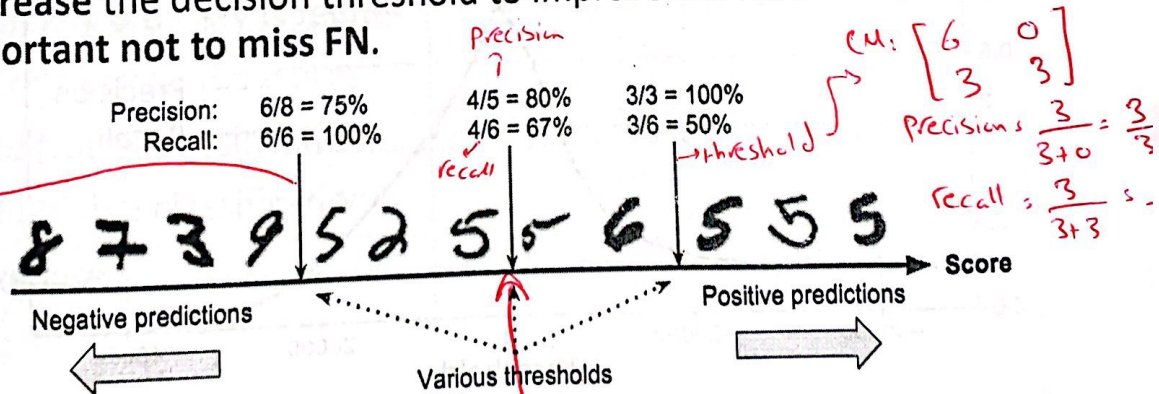
$$F_1 = \frac{2}{\frac{1}{\text{precision}} + \frac{1}{\text{recall}}} = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}} = \frac{2TP}{2TP + FN + FP}$$

```
>>> from sklearn.metrics import f1_score
>>> f1_score(y_train_5, y_train_pred)
0.7420962043663375
```

بناطبة ال labels وال predictions

### 3.5. Precision/Recall Tradeoff \*

- Increase the decision threshold to improve the precision when it is bad to have FP.
- Decrease the decision threshold to improve the recall when it is important not to miss FN.



Confusion Matrix. CM →

5	1
2	4

Accuracy =  $\frac{5+4}{5+1+2+4} = \frac{9}{12}$   
 Precision =  $\frac{4}{4+1} = \frac{4}{5}$   
 Recall =  $\frac{4}{4+2} = \frac{4}{6}$



• ↓ threshold ↑ FP ↓ Precision

• ↓ threshold ↓ FN ↑ Recall

يمكن تعديل default threshold ويمكن اننا نحدد.

### 3.5. Precision/Recall Tradeoff

• The function `cross_val_predict()` can return decision scores instead of predictions.

```
y_scores = cross_val_predict(sgd_clf, X_train, y_train_5, cv=3,
                             method="decision_function")
```

عشان يرجع score

• These scores can be used to compute precision and recall for all possible thresholds using the `precision_recall_curve()` function


```
from sklearn.metrics import precision_recall_curve
precisions, recalls, thresholds = precision_recall_curve(y_train_5, y_scores)
```

بحسب ال precisions & recalls For all thresholds زي سنة سلايه 20 (scores)

data	label	y_score
X	Y	
5	true	3000
1	False	-100

بتبدا ال threshold من اقل قيمة ال decision func

### 3.5. Precision/Recall Tradeoff



Threshold اذا بي ازيد ال precision بزيه ال Recall اقل ال

## 3.5. Precision/Recall Tradeoff

argmax() returns the index of the largest element in an array. If the array is boolean then it will return the index of the first occurrence of True

- For **larger precision**, **increase the threshold**, and **decrease it** for **larger recall**.

Precisions increase with the threshold

- **Example:** To get 90% precision.

```
>>> idx_for_90_precision = (precisions >= 0.90).argmax()
>>> threshold_for_90_precision = thresholds[idx_for_90_precision]
>>> threshold_for_90_precision
3370.0194991439557
y_train_pred_90 = (y_scores >= threshold_for_90_precision)
>>> precision_score(y_train_5, y_train_pred_90)
0.9000345901072293
>>> recall_at_90_precision = recall_score(y_train_5, y_train_pred_90)
>>> recall_at_90_precision
0.4799852425751706
```

# Outline

1. MNIST dataset
2. Training a binary classifier
3. Performance measures
4. Multiclass classification
5. Multilabel classification
6. Exercise

## 4. Multiclass Classification

our data is not only 5 or not 5 we have many classes 'numbers'

لدينا رقم detection و نقيس بينهم و نقطي اى منهم جوابه

- Multiclass classifiers can distinguish between more than two classes
- Some algorithms (such as Random Forest classifiers or Naive Bayes classifiers) are capable of handling multiple classes directly.
- Others (such as Support Vector Machine classifiers or Linear classifiers) are strictly binary classifiers.
- There are two main strategies to perform multiclass classification using multiple binary classifiers.

Probability ←

0-detector  $\Rightarrow$  0  $\rightarrow$  True 0.W False  
5-detector = 5  $\rightarrow$  True 0.W False

### 4.1. One-versus-All (OvA) Strategy

لدينا 10

- For example, classify the digit images into 10 classes (from 0 to 9) to train 10 binary classifiers, one for each digit (a 0-detector, a 1-detector, a 2-detector, and so on).
- Then to classify an image, get the decision score from each classifier for that image and select the class whose classifier outputs the highest score.

10 models يعني بنحتاج 10 classes ~~10~~

و بنحتاج ال model ال اعلى score

## 4.2. One-versus-One (OvO) Strategy

10 classifiers 1 or 5 classifera  
 True or False صحیح یا غلط  
 2 classes بنیاداً 2 classes  
 classifier بنیاداً classifier  
 - یز بیزم

- Train a binary classifier for every pair of digits.
- If there are N classes, need  $N \times (N - 1) / 2$  classifiers. For MNIST, need 45 classifiers.
- To classify an image, run the image through all 45 classifiers and see which class wins the most duels.
- The main advantage of OvO is that each classifier only needs to be trained on a subset of the training set.
- OvO is preferred for algorithms (such as Support Vector Machine) that scale poorly with the size of the training set.

45 @ model بنیاداً 2  
 10 classes 3 دلتیا

↓ SVM

which one is better? OVA or OvO?  
 حسب طبیعت ال model ال complexity

## 4.3. Scikit Learn Support of Multiclass Classification

- **Scikit-Learn** detects when you try to use a binary classification algorithm for a multiclass classification task, and it automatically runs **OvA** (except for **SVM** classifiers for which it uses **OvO**).

```
from sklearn.svm import SVC
```

```
svm_clf = SVC(random_state=42)
svm_clf.fit(X_train[:2000], y_train[:2000]) # y_train, not y_train_5
>>> svm_clf.predict([some_digit])
array(['5'], dtype=object)
>>> some_digit_scores = svm_clf.decision_function([some_digit])
>>> some_digit_scores.round(2)
array([[ 3.79,  0.73,  6.06,  8.3 , -0.29,  9.3 ,  1.75,  2.77,  7.21,
         4.82]])
>>> class_id = some_digit_scores.argmax()
>>> class_id
5
```

ال labels للداتا الاصلية الي القيم  
فيها ممكن تاخذ اي قيمة من صفر ل  
٩



## 4.3. Scikit Learn Support of Multiclass Classification

- Note that the multiclass task is harder than the binary task.
- Binary task

```
>>> from sklearn.model_selection import cross_val_score
>>> cross_val_score(sgd_clf, X_train, y_train_5, cv=3, scoring="accuracy")
array([0.96355, 0.93795, 0.95615])
```

↳  
بس بتكفي 5 اذ 5  
(Binary)

- Multiclass task

```
>>> cross_val_score(sgd_clf, X_train, y_train, cv=3, scoring="accuracy")
array([0.8489802 , 0.87129356, 0.86988048])
```


↳ 10 labels  
(more than 1 class)

الفرق بال accuracy

## 4.4. Error Analysis

Confusion matrix  
صارت اكبر  
بدل 2X2 بحالة binary classifier  
صارت ( 10labels ) 10X10

```
from sklearn.metrics import ConfusionMatrixDisplay
y_train_pred = cross_val_predict(sgd_clf, X_train_scaled, y_train, cv=3)
ConfusionMatrixDisplay.from_predictions(y_train, y_train_pred)
plt.show()
```



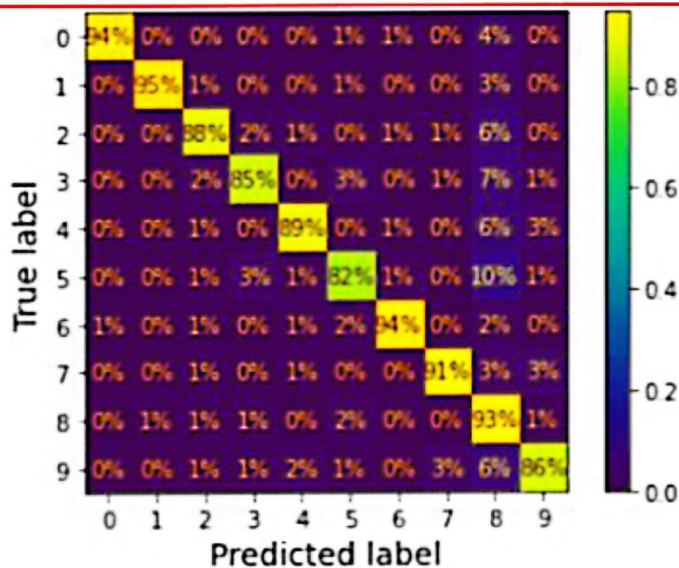
اي اشي برا ال diagonal يعني false prediction

31

مجموع عناصر ال row لازم يكون بساوي 1

## 4.4. Error Analysis

```
ConfusionMatrixDisplay.from_predictions(y_train, y_train_pred,
normalize="true", values_format=".0%")
plt.show()
```



32



# Outline

1. MNIST dataset
2. Training a binary classifier
3. Performance measures
4. Multiclass classification
5. Multilabel classification →
6. Exercise

ال sample له اكثر من Label  
ال Labels لا الهن ثلاثة - بربيع

31

## 5. Multilabel Classification

• Column for each label  
# of pred. = # of labels

- Classifiers that output multiple classes for each instance.

Concatenation  
is a simple  
affair

```
y_train_large = (y_train >= 7)
y_train_odd = (y_train % 2 == 1)
y_multilabel = np.c_[y_train_large, y_train_odd] → use like labels 1, 2, 3!
← knn_clf = KNeighborsClassifier() ← Popular algorithm 2 columns
knn_clf.fit(X_train, y_multilabel)
          samples      multilabel for each sample

>>> knn_clf.predict([some_digit])
array([[False,  True]], dtype=bool)
```

32

# **Training Models and Regression**

**Prof. Gheith Abandah**

# Reference

- Chapter 4: Training Models



- Aurélien Géron, **Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow**, O'Reilly, 2nd Edition, 2019
  - Material: <https://github.com/ageron/handson-ml2>

# Outline

1. Linear Regression
2. Gradient Descent → optimization techniques
3. Gradient Descent Variants
  1. Batch Gradient Descent
  2. Stochastic Gradient Descent
  3. Mini-batch Gradient Descent
4. Learning Curves
5. Early Stopping
6. Exercises

bias term → curve بال  
 كيف يطلع او ينزل  
 $\theta_0 \rightarrow$  slope


# Linear Regression

Relationship between  $x$  and  $y$   
 ("العلاقة بين المتغيرات  $x$  و  $y$ ")


$$\hat{y} = \theta_0 + \theta_1 x_1 + \theta_2 x_2 + \dots + \theta_n x_n$$

الجزء label

- $\hat{y}$  is the predicted value.
- $n$  is the number of features.
- $x_i$  is the  $i^{\text{th}}$  feature value.
- $\theta_j$  is the  $j^{\text{th}}$  model parameter (including the bias term  $\theta_0$  and the feature weights  $\theta_1, \theta_2, \dots, \theta_n$ ).



بضرب ال row بال  
 Column كل



$$\hat{y} = h_{\theta}(x) = \theta \cdot x$$

hypothesis (prediction)  
 فرضية

$n = 784 \rightarrow$  MNIST

## Analytical Solution

كم Prediction ال  
 بعد عن ال actual value

- The Root Mean Square Error (RMSE) is used as cost function.

في حال كان على بنفرد ال  
 $\begin{bmatrix} \theta_0 \\ \theta_1 \\ \theta_2 \\ \vdots \end{bmatrix}$

$$MSE(X, h_{\theta}) = \frac{1}{m} \sum_{i=1}^m (\theta^T x^{(i)} - y^{(i)})^2$$

higher complexity

- Minimizing this cost gives the following solution (normal function):

optimization

$$\hat{\theta} = (X^T X)^{-1} X^T y$$

Complexity  $O(mn^2)$

قيم ال parameters  
 بعد ما نعرفه  
 ارقام ونوعه  
 more complexity


- $\hat{\theta}$  is the value of  $\theta$  that minimizes the cost function.
- $y$  is the vector of target values containing  $y^{(1)}$  to  $y^{(m)}$ .

$m$ : no. of samples  
 $n$ : no. of features.




# Learning Rate < 1

Too Small

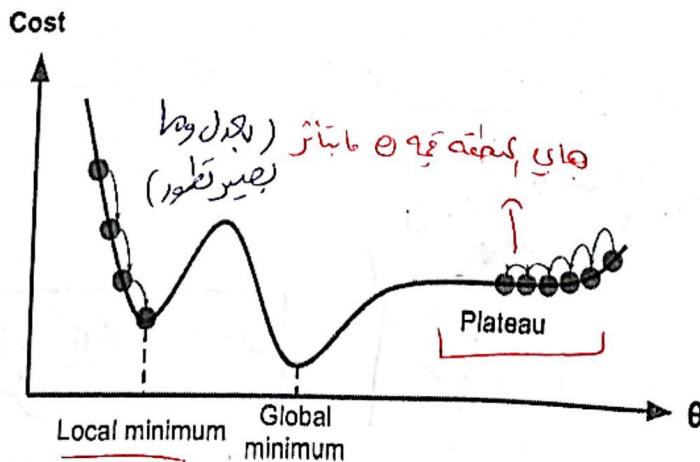


Too Large



learning rate = 1  
 optimal value skip  
 value

## \* Gradient Descent Pitfalls




Global min

# Feature Scaling

الهدف تكون قيم ال features قريبه من بعضه  
سببى ل convergence

- Ensure that all features have a similar scale (e.g., using Scikit-Learn's StandardScaler class).
- Gradient Descent with and without feature scaling.

used in Transformation Pipeline



10

مع ال Scaling ال Cost يتقل

## Outline

1. Linear Regression
2. Gradient Descent
3. Gradient Descent Variants
  1. Batch Gradient Descent
  2. Stochastic Gradient Descent
  3. Mini-batch Gradient Descent
4. Learning Curves
5. Early Stopping
6. Exercises

11

$$MSE(x, h_{\theta}) = \frac{1}{m} \sum_{i=1}^m (\theta^T x^{(i)} - y^{(i)})^2$$

# Batch Gradient Descent

Training set  
Average loss  
برای کل  
و حساب

- Partial derivatives of the cost function in  $\theta_j$

بعد اشتقاق از هر فاکتور از معادله MSE  
المره الأولى بالنسبة لـ  $\theta_0$  بقدرته

$$\frac{\partial}{\partial \theta_j} MSE(\theta) = \frac{2}{m} \sum_{i=1}^m (\theta^T x^{(i)} - y^{(i)}) x_j^{(i)}$$

$i \rightarrow$  رقم الساميل  
 $j \rightarrow$  feature

- Gradient vector of the cost function

بستق بانجه  $\theta$

$$\nabla_{\theta} MSE(\theta) = \begin{bmatrix} \frac{\partial}{\partial \theta_0} MSE(\theta) \\ \frac{\partial}{\partial \theta_1} MSE(\theta) \\ \vdots \\ \frac{\partial}{\partial \theta_n} MSE(\theta) \end{bmatrix} = \frac{2}{m} X^T (X\theta - y)$$

$X^T$  transpose  
The entire training Batch  
Features  
 $\theta^{(next\ step)} = \theta - \eta \nabla_{\theta} MSE(\theta)$

$\exists x = \theta_0 + \theta_1 F_1 + \theta_2 F_2$   
 $MSE = \frac{1}{m} \sum_{i=1}^m ((\theta_0 + \theta_1 F_1 + \theta_2 F_2) - y)^2$   
 $\frac{dMSE}{d\theta_0} = \frac{2}{m} \sum_{i=1}^m ((\theta_0 + \theta_1 F_1 + \theta_2 F_2) - y) \cdot 1$   
 $\frac{dMSE}{d\theta_1} = \frac{2}{m} \sum_{i=1}^m ((\theta_0 + \theta_1 F_1 + \theta_2 F_2) - y) \cdot F_1$

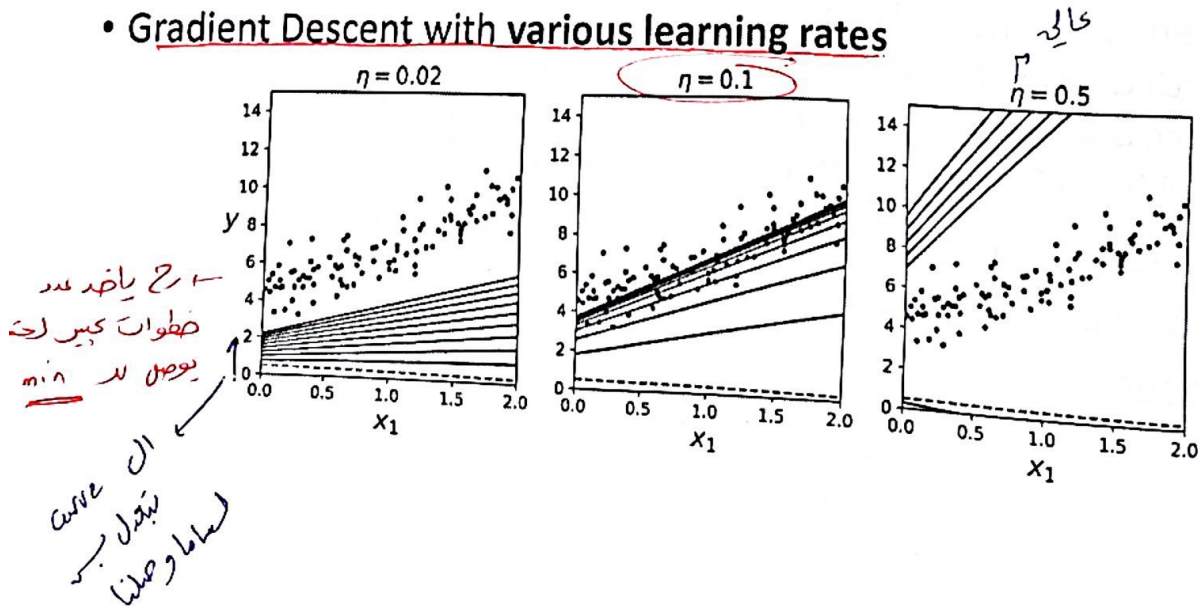
بعد الاشتقاق بطلع علينا ال slope لك Cost function  
بانجاه  $\theta$

# Batch Gradient Descent

- Gradient Descent step

$$\theta^{(next\ step)} = \theta - \eta \nabla_{\theta} MSE(\theta)$$

- Gradient Descent with various learning rates





Full gradient descent → لا تأخذ كل ال samples


Stochastic " " → تأخذ sample واحدة

# stochastic Gradient Descent (SGD)


$$(\theta_0 + \theta_1 F_1 + \theta_2 F_2) - y)^2 \rightarrow$$

- SGD picks a random instance in the training set at every step and computes the gradients.
- SGD is faster when the training set is large.
- Is bouncy
- Eventually gives good solution
- Can escape local minima

باص  
صين  
المعادلة  
التي  
بستقر  
تكون



× هاي الطريقة  
تعتمد على ال  
Sample



14


بار على ال weight التي يتطلع  
عنه بتعديل ال  
weights of  $\theta$

تأخذ مجموعة ال samples  
منه بس واحد

# Mini-batch Gradient Descent

ما يأخذ ال data كاملة ويقتطع  
الوقت ما يأخذ ال samples بشكل عشوائي.

- Computes the gradients on small random sets of instances called mini batches.
- Benefits from hardware accelerators (e.g., GPU).
- Less bouncy, better solution, escapes some local minima



15


# Outline

1. Linear Regression
2. Gradient Descent
3. Gradient Descent Variants
  1. Batch Gradient Descent
  2. Stochastic Gradient Descent
  3. Mini-batch Gradient Descent
4. Learning Curves
5. Early Stopping
6. Exercises

## Learning Curves

- The accuracy on the validation set generally increases as the training set size increases.

- Overfitting decreases with larger training set.



because the training set size is small

# Outline

بالعادة لما يكون ال data set مشيخه يكون عنده overfitting  
 لانه اذا اتاها يكون مشابهة لكل ال samples

1. Linear Regression
2. Gradient Descent
3. Gradient Descent Variants
  1. Batch Gradient Descent
  2. Stochastic Gradient Descent
  3. Mini-batch Gradient Descent
4. Learning Curves
5. Early Stopping
6. Exercises

18

اذا ما تبصنا ال RMSE "ما في تطور"  
 بوقف


## Early Stopping

→ to save time & prevent overfitting:

• ممكن بشرط فلا في ال Best model لما نشوف كل ال Epochs

- Stop training when the validation error reaches a minimum.
- Need to save the best model.

بالبدايه (bias)  
 Model isn't good  
 neither at training  
 Set nor validation



قل بتدريج ز اد

لجمله save

رجوع  
 يزيد

we can say here is  
 overfitting

19

لما كذا زديت يقل ال error  
 & training loss

يعني لما ادرب ال samples  
 واعدت ال weight - epoch

لانه ال training RMSE يقل  
 وال validation RMSE  
 دليل انه بعد ال bias training data

يعني مثلا بعطيه صورة وبدي يعطيني الجواب بالزبط  
 او الكلاس، بنطلع probability اذا كانت  
 اعلى من threshold معين يعني positive class  
 اذا اقل يعني negative class

ما بعطينا قيمة  
 بعطينا احتمال


# Logistic Regression

- Estimates the probability that an instance belongs to a particular class
  - **Positive Class: Probability greater than a given threshold**
- Instead of outputting the result directly like the linear regression model does, it outputs the *logistic* of this result

$$\hat{p} = h_{\theta}(\mathbf{x}) = \sigma(\theta^T \mathbf{x})$$

$$\sigma(t) = \frac{1}{1 + \exp(-t)}$$

$$\hat{y} = \begin{cases} 0 & \text{if } \hat{p} < 0.5 \\ 1 & \text{if } \hat{p} \geq 0.5 \end{cases}$$



t=hypotheses (h(x))

حولنا ال prediction لقيمة بين صفر  
 وواحد عشان نفهمه ك probability

20


القيم صارت بين صفر وواحد قال  
 RMSE اكيد بين الصفر وواحد  
 فبنستخدم ال log loss

# Logistic Regression-Training and Cost Function

- Log loss: Instances follow Gaussian distribution around the mean of their class
  - Log(p) is close to 0 when p is close to 1
  - Log(1-p) is close to 0 when p is close to 0

اذا ال probability قريبة من 1 وال  
 label (y) بساوي 1، ال cost قريبة  
 من صفر

$$c(\theta) = \begin{cases} -\log(\hat{p}) & \text{if } y = 1 \\ -\log(1 - \hat{p}) & \text{if } y = 0 \end{cases}$$



$$J(\theta) = -\frac{1}{m} \sum_{i=1}^m [y^{(i)} \log(\hat{p}^{(i)}) + (1 - y^{(i)}) \log(1 - \hat{p}^{(i)})]$$

$$\frac{\partial}{\partial \theta_j} J(\theta) = \frac{1}{m} \sum_{i=1}^m (\sigma(\theta^T \mathbf{x}^{(i)}) - y^{(i)}) x_j^{(i)}$$

21

$$\log(\hat{p}) \quad \text{label } y=1 \quad , \hat{p} \approx 1$$

$$y \log(\hat{p}) + \underbrace{(1-y) \log(1-\hat{p})}_{\substack{\text{هاد منحص } L \\ \text{لما يكون ال label صفر}}}$$

$$1 \cdot \log(\hat{p})$$

$$1 \cdot \log(1) = 1 \cdot 0 = 0$$

$$1 \cdot \log(0.2) \rightarrow \text{قيمة عالية}$$

زي كانه ال label صفر , كاي

$$\text{if } y=0 \quad \hat{p} \approx 0$$

$$(1-y) \log(1-\hat{p})$$

$$(1-0) \log(1-\hat{p}) \rightarrow 1 \cdot \log(1) = 0$$

التوقع يطابق  
ال ground truth

$$(1-0) \log(1-0.9) \rightarrow 1 \cdot \log(0.1) \rightarrow \text{قيمة عالية}$$

كل ما كانت القيمة داخل  $\log$  اصغر بتكون الجواب الاكبر.

لهمنا بدل minimization ال Cost فبنفق ال Cost اكبر

قيمة من قيم  $\ominus$

# Iris Dataset

- A famous dataset that contains the sepal and petal length and width of **150 iris flowers** of three different species: **Setosa**, **Versicolor**, and **Virginica**.



```
>>> from sklearn import datasets
>>> iris = datasets.load_iris()
>>> list(iris.keys())
['data', 'target', 'target_names', 'DESCR', 'feature_names', 'filename']
```

هي فعليا multiclass problem لانه  
عنا 3 classes  
بس احنا حليناها على انها binary  
classification problem

22

## Logistic Regression-Example (binary classifier)

- Predict\_proba(): returns the probability of the instance
- Predict(): return the predicted class for the instance

```
from sklearn.linear_model import LogisticRegression
from sklearn.model_selection import train_test_split


X = iris.data[["petal width (cm)"]].values
y = iris.target_names[iris.target] == 'virginica'
X_train, X_test, y_train, y_test = train_test_split(X, y,
random_state=42)

log_reg = LogisticRegression(random_state=42)
log_reg.fit(X_train, y_train)
```

23

# Logistic Regression-Example (binary classifier)

لما ناخذ ال petal width = 1.6  
بتكون ال probability = 0.5  
يعني اذا اقل من 1.6 not virginica



Predict -> returns if its virginica or not

Predict\_proba -> returns the probability

24

# Logistic Regression-Multiclass

بال predict باخذ ال max ال probability وبرجعلي اياها لما يكون عنا multiclass

- Softmax Regressor: Normalize the probability for each class.
- Cross entropy cost function

$$J(\Theta) = -\frac{1}{m} \sum_{i=1}^m \sum_{k=1}^K y_k^{(i)} \log(\hat{p}_k^{(i)})$$

Cost function here is the categorical cross entropy  
ينضرب كل label باحتمال ال class تبعته

```
>>> softmax_reg.predict([[5, 2]])
array([2])
>>> softmax_reg.predict_proba([[5, 2]]).round(2)
array([[0. , 0.04, 0.96]])
```

<- soft max regressor  
بعمل normalize لقيم ال  
prediction بحيث يصير مجموعهم  
بساوي 1

25

# **Classical Techniques**

**Prof. Gheith Abandah**




## Reference

- Chapter 5: Support Vector Machines
- Chapter 6: Decision Trees
- Chapter 7: Ensemble Learning and Random Forests

اجمع اكثر من Classifier مع بعض  
وباختار منهم


- Aurélien Géron, Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow, O'Reilly, 2nd Edition, 2019
  - Material: <https://github.com/ageron/handson-ml2>



## Outline

1. k-Nearest Neighbors
2. Support Vector Machines
3. Decision Trees
4. Ensemble Learning and Random Forests
5. Exercises  
Classifiers

بدی توقع های ار instance را ی Class  
 بنا علی الی حولی  
 حلا بین امانت بینک بین الی حولی  
 ۱۵ و ۲۵



# k-Nearest Neighbors

- Find a predefined number of training samples ( $k$ ) closest in distance to the new point and predict the label from them: regression or classification.
- The number of samples can be a user-defined constant (**k-nearest neighbor learning**), or vary based on the local density of points (**radius-based neighbor learning**).
- The distance can be any metric measure: standard **Euclidean distance** is the most common choice.
- Reference: <https://scikit-learn.org/stable/modules/neighbors.html>

ما عنی math made انو یعمل predict ، هون حسب امانت دهی نصف Class و reg  
 Classification ← بطلع علی ۳ الی حولی و یعطی ار Class الاقرب  
 regression ← ~ ~ ~ ~ ~ ال avg

لا او حلا شو ال instance الی بتبعد قطر ۱ او قطر ۲ ← radius based neighbour learning

• لهما توصل بسرعة لاقری بـ k-elements فینستغرا BallTree / KDTree الی یسجروا علیة ال indexing

## Nearest Neighbors Classification

# of neighbors by default = 5

```
class sklearn.neighbors.KNeighborsClassifier(n_neighbors=5, weights='uniform', ...)
```

# of neighbors  
 اقرب 5  
 یفضل لكون odd

don't forget to import

• weights can be: uniform: All points in each neighborhood are weighted equally, and distance: Weight points by the inverse of their distance.

له كلما زاد ابعد عن تهل ار decision تبعد

• Example:

```
from sklearn.neighbors import KNeighborsClassifier
knn_clf = KNeighborsClassifier() → object (default values)
knn_clf.fit(X_train, y_train)
```

على مرفها  $X_3$  بعيدة  
 $x_1 + x_2 + 0.8 x_3$   
 Weight ال

\* check the documentation of k-nearest neighbors.

م بدل ما احب انا في عني indexing nodes - في عني انا في عني  
وكل نقطة اذا بي استخدم مع large dataset احب باي data واهب انا في عني  
هنا كده بيستعمل memory

## Nearest Neighbors Regression

```
class sklearn.neighbors.KNeighborsRegressor(n_neighbors=5,  
weights='uniform', ... )
```

- The label assigned to a query point is computed based on the mean of the labels of its nearest neighbors.

- Example:

```
from sklearn.neighbors import KNeighborsRegressor  
model = KNeighborsRegressor(n_neighbors=3)  
model.fit(X, y)
```

## Outline

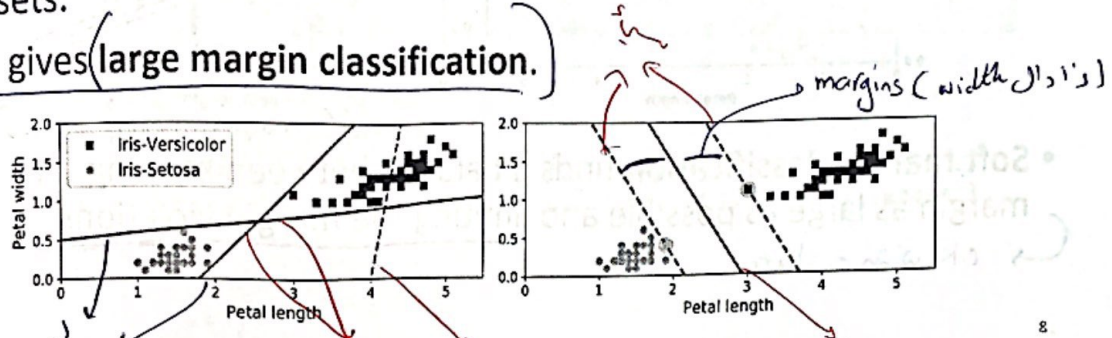
1. k-Nearest Neighbors
2. Support Vector Machines
3. Decision Trees
4. Ensemble Learning and Random Forests
5. Exercises

Complexity of SVM  $\rightarrow O(m^2)$   
 \* اذا جربت على data كبيرة يكون بطيء.

كمان يفضل ال classes في بعض  
 linear  $\rightarrow$  برسم خط مستقيم  
 non-linear  $\rightarrow$  مشرط مستقيم  
 binary  
 multi

# Support Vector Machine (SVM)

- Very powerful and versatile Machine Learning model, capable of performing linear or nonlinear classification, regression, and outlier detection.
- Well suited for classification of complex but small- or medium-sized datasets.
- SVM gives large margin classification.



ليس هون  
 المشكلة في نقاط  
 قريبة من ال line  
 صا في مسافة ليه  
 sample ال line  
 ال line Classifier  
 ليس هو

ممتاز  
 لانه  
 كافع نقاط  
 بنصف ال  
 class  
 مازة قابل بين ال  
 classes  
 بنظم يكون  
 ال decision boundary

## Linear SVM Classification $\rightarrow$


\* why to do Feature scale?

```

from sklearn.svm import LinearSVC
svc_linear = LinearSVC()
svc_linear.fit(x_train, y_train)
y_pred = svc_linear.predict(x_test)
    
```


- The decision boundary is fully determined by the instances located on the edge. These instances are called the **support vectors**.
- SVMs are **sensitive to the feature scales**.

على لما اكون بقدر اديسكي على اساس مسافات ال distance  
 لازم اعمل scaling (الانظروا)



الوزن  
 تبع  $X_1$   
 رح يكون  
 تأثيره على  
 المسافة اكب

دفع يقلل من تأثير  $X_0$   
 في العادة.



منوع يكون  
 عندي  
 instances  
 داخل ال  
 margin

Surface ال  
 ال افضل

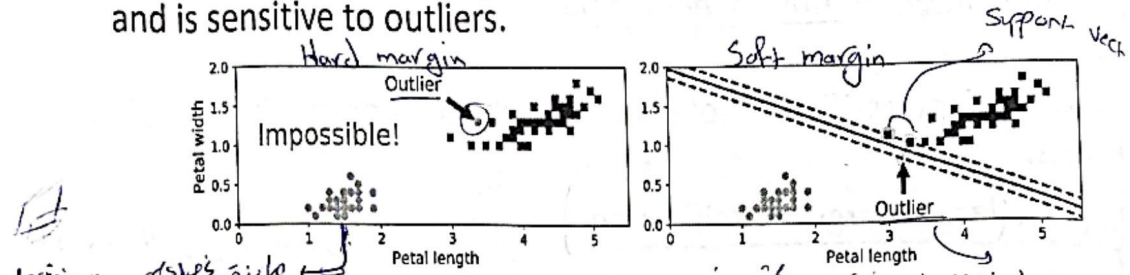
Hard margin classification

margin هو المسافة بين الخط الفاصل واقرن تقام عليه

(support vector)

## Soft Margin Classification

- All instances must be off the street and on the correct side
- Hard margin classification cannot handle linearly inseparable class and is sensitive to outliers.




Soft margin classification finds a balance between keeping the margin as large as possible and limiting the margin violations.

OK with outliers.

instances ليعبر لحدود  
wrong side يكونوا على

## Soft Margin Classification (by default) soft

- You can control the number of violations using the  $c$  hyperparameter



- If your SVM model is overfitting, you can try regularizing it by reducing  $C$ .


$C \uparrow$  Violations  $\downarrow$  margins

this model accuracy is less than this model (because more violations).

$C = \infty \rightarrow$  hard margin

# Iris Dataset

- A famous dataset that contains the sepal and petal length and width of **150 iris flowers** of three different species: **Setosa**, **Versicolor**, and **Virginica**.



```
>>> from sklearn import datasets
>>> iris = datasets.load_iris()
>>> list(iris.keys())
['data', 'target', 'target_names', 'DESCR', 'feature_names', 'filename']
```

If we want to use SVM then we must scale the data

## SVM Classification Example

```
from sklearn.datasets import load_iris
from sklearn.pipeline import make_pipeline
from sklearn.preprocessing import StandardScaler
from sklearn.svm import LinearSVC

iris = load_iris(as_frame=True)
X = iris.data[["petal length (cm)", "petal width (cm)"].values
y = (iris.target == 2) # Iris virginica

svm_clf = make_pipeline(StandardScaler(),
                        LinearSVC(C=1, random_state=42))
svm_clf.fit(X, y)

>>> X_new = [[5.5, 1.7], [5.0, 1.5]]
>>> svm_clf.predict(X_new)
array([ True, False])

>>> svm_clf.decision_function(X_new)
array([ 0.66163411, -0.22036063])
```

اسماء ال columns

We have 3 labels, (0,1,2)

C=1 -> Soft



```
poly_kernel_svm_clf.fit(X, y)
```

# Nonlinear SVM Classification

مرات الداتا ما بتكون linearly separable يعني ما بنقدر نعمل خط مستقيم يفصل الداتا عن بعض

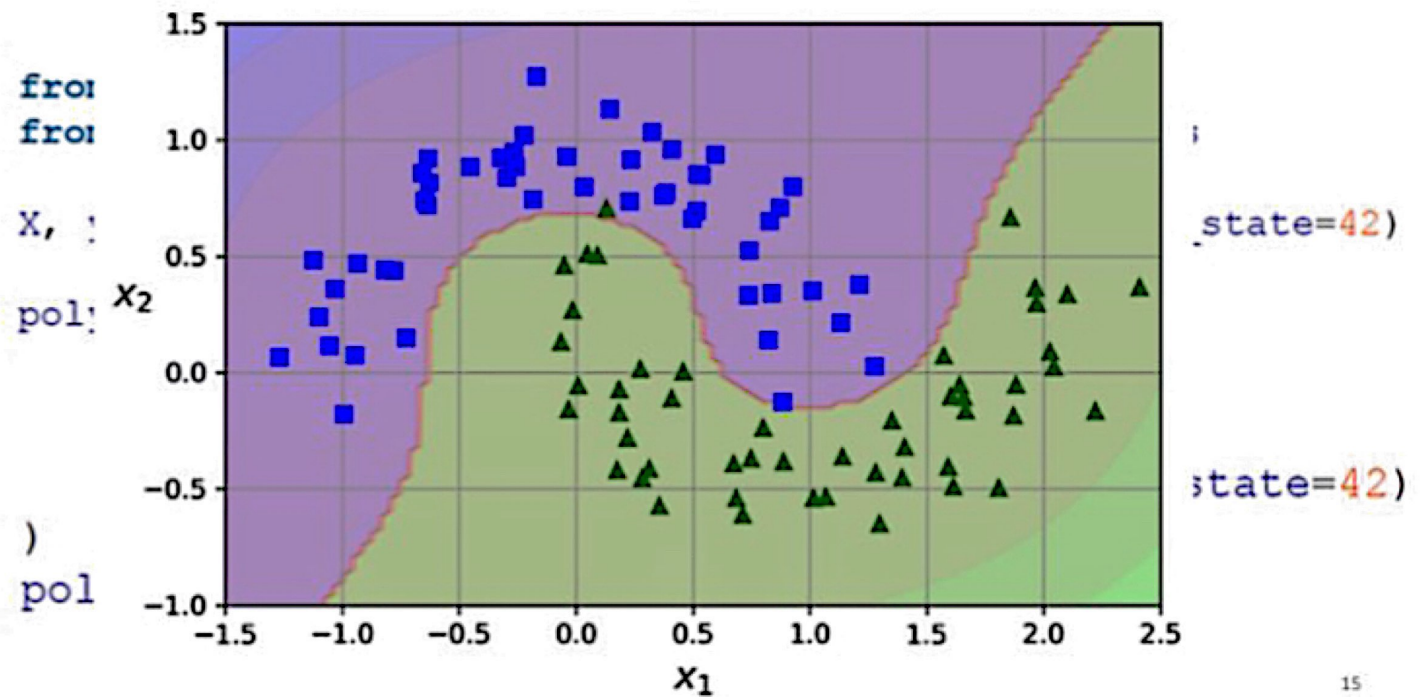
- Some dataset are not linearly separable
- Adding more features such as polynomial features can making the function linearly separable

كانت بس  $X_1$  فعملناها  $X_1^2$



او ممكن نستعمل classes تانية غير linear SVC ال وجوبها بكون curve

# Nonlinear SVM Classification



## Nonlinear SVM Classification

كيف نحسب  $X_1^2$  باستخدام ال  
polynomialFeatures و بنحدد ال  
degree واستخدمنا linearSVC لانه  
بدنا خط مستقيم

```
from sklearn.datasets import make_moons
from sklearn.preprocessing import PolynomialFeatures

X, y = make_moons(n_samples=100, noise=0.15, random_state=42)

polynomial_svm_clf = make_pipeline(
    PolynomialFeatures(degree=3),
    StandardScaler(),
    LinearSVC(C=10, max_iter=10_000, random_state=42)
)
polynomial_svm_clf.fit(X, y)
```



from sklearn.svm import SVC  
 SVC\_nonlinear = SVC()  
 SVC\_nonlinear.fit(X\_train, Y\_train)  
 Y\_pred = SVC\_nonlinear.predict(X\_test)

لازم  
العمل  
Scaling

## Nonlinear SVM Classification


- The SVM class supports nonlinear classification using the kernel option

```

from sklearn.svm import SVC
poly_kernel_svm_clf = Pipeline([
    ("scaler", StandardScaler()),
    ("svm_clf", SVC(kernel="poly", degree=3, coef0=1, C=5))
])
poly_kernel_svm_clf.fit(X, y)
  
```

Controls how much the model is influenced by high-degree polynomials versus low-degree

← فضل بالعامة  
 parameter واحد بالمرة  
 الخاصة




كزيادة ال degree يخلي ال curve يصير مع ال data بشكل احسن  
 بس الزيادة الكبيرة ممكنة تكون overfitting

## Gaussian Radial Basis Function → RBF Function

نضيف Features جديدة

$$\phi_\gamma(x, \ell) = \exp(-\gamma \|x - \ell\|^2)$$

- The Gaussian RBF can be used to find similarity features ( $x_2$  and  $x_3$ ) of the one-dimensional dataset with two landmarks to it at  $x_1 = -2$  and  $x_1 = 1$



reference point  
 اوجب ال  
 فكل نقطة ال  
 ال 2 Features  
 ال نقطة ال  
 ال

ال  
 ال  
 ال  
 ال

# Gaussian RBF Kernel

- Is popular with SVM to solve nonlinear problems.

```
rbf_kernel_svm_clf = Pipeline([
    ("scaler", StandardScaler()),
    ("svm_clf", SVC(kernel="rbf", gamma=5, C=0.001))
])
rbf_kernel_svm_clf.fit(X, y)
```





- Transforms a training set with  $m$  instances and  $n$  features to  $m$  instances and  $m$  features.
- gamma and C are used for regularization with smaller values.

↳ like learning rate.

remember :-  $C \uparrow$  violation  $\downarrow$

# Gaussian RBF Kernel

violations دى كبر



decision boundaries

non-linear kernel

جلى violations دى كبر

كل الارقام الخطر اذ العلاقة التي يتصل اليها النقاط ( line يوصف , العلاقة بين X & Y )

## Linear SVM Regression → output (cont. values).

→ can be used as classifier & as regressor.

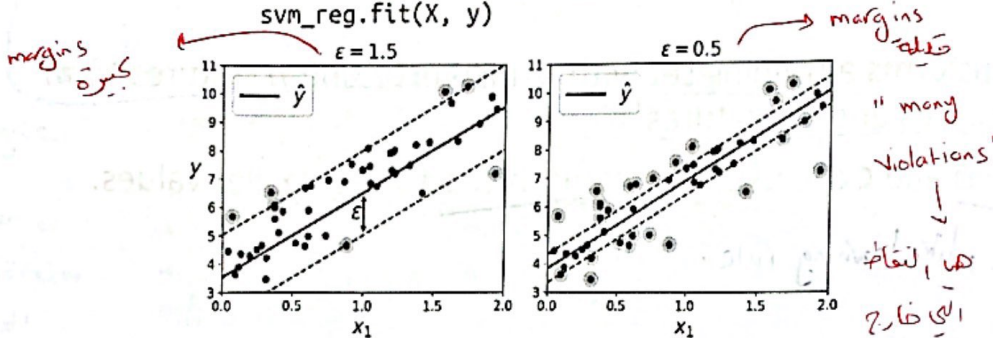
- Fits as many instances as possible on the margin while limiting margin violations. The width of the street is controlled by a hyperparameter  $\epsilon$ .

$\epsilon$ .

```
from sklearn.svm import LinearSVR
```

```
svm_reg = LinearSVR(epsilon=1.5)
```

```
svm_reg.fit(X, y)
```



SVC → classifier  
 SVR → regressor

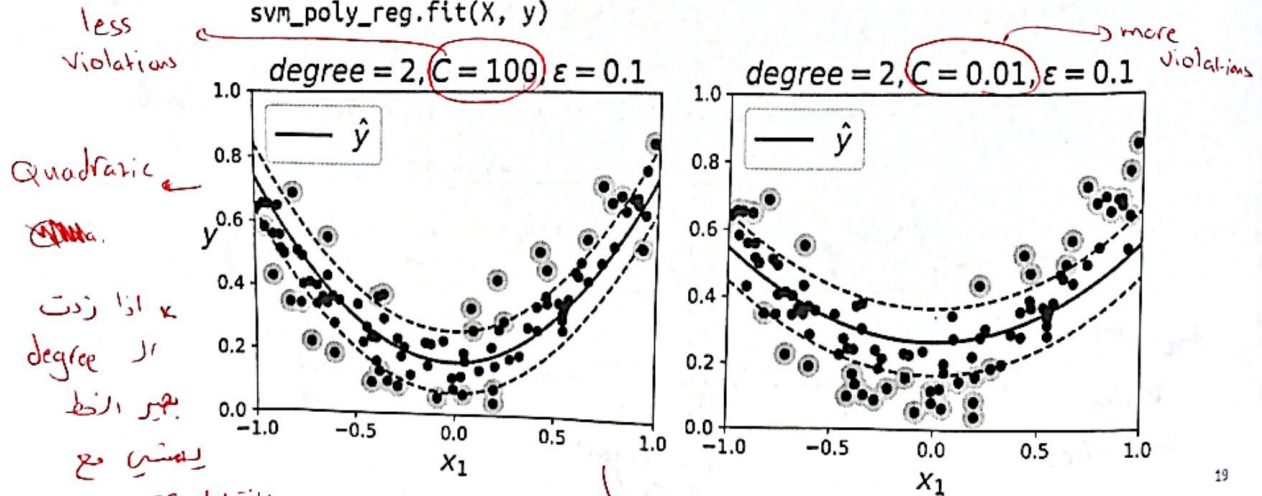
المargins عكس ال classification  
 بينا نوجه ال margin الي بتعمل معظم النقاط

## Nonlinear SVM Regression

```
from sklearn.svm import SVR
```

```
svm_poly_reg = SVR(kernel="poly", degree=2, C=100, epsilon=0.1)
```

```
svm_poly_reg.fit(X, y)
```



better results

# SVM Conclusion

can be multi-classifier

- The Linear SVC has complexity of  $O(m \times n)$ .
- The SVC time complexity is usually between  $O(m^2 \times n)$  and  $O(m^3 \times n)$ .  
# of samples (pointing to  $n$ )  
# of samples (pointing to  $m$ )
- This algorithm is perfect for complex but small or medium training sets. However, it scales well with the number of features.

SVC →  
بجهد كبير  
= بيهد  
↑ m ↑ complexity

## Outline

1. k-Nearest Neighbors
2. Support Vector Machines
3. Decision Trees
4. Ensemble Learning and Random Forests
5. Exercises



# Decision Trees < /span> classification regression

- Decision Trees are versatile Machine Learning algorithms that can perform both classification and regression tasks, and even multioutput tasks.
- They are very powerful algorithms, capable of fitting complex datasets.

```

from sklearn.datasets import load_iris
from sklearn.tree import DecisionTreeClassifier
iris = load_iris() → iris dataset
X = iris.data[:, 2:] # petal length and width
y = iris.target → labels
tree_clf = DecisionTreeClassifier(max_depth=2)
tree_clf.fit(X, y)
    
```

Y-pred = tree\_clf.predict(X-test)

we have predict-proba

no. of samples class 1    class 2

$$gini = 1 - \left(\frac{50}{150}\right)^2 - \left(\frac{50}{150}\right)^2 - \left(\frac{50}{150}\right)^2$$


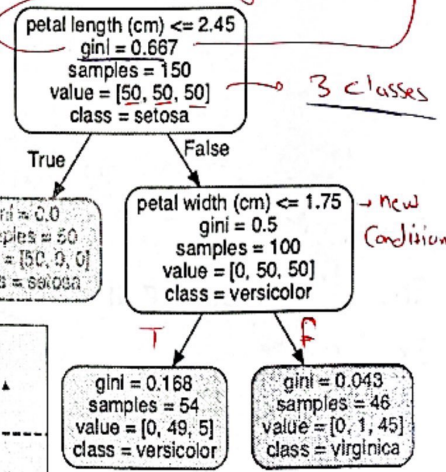
total no. of samples

ل يكون اقرب للصفر كما يكونوا ال elements  
كلهم في class واحد

## Visualizing a Decision Tree

```

from sklearn.tree import export_graphviz
export_graphviz(
    tree_clf,
    out_file=image_path("iris_tree.dot"),
    feature_names=iris.feature_names[2:],
    class_names=iris.target_names,
    rounded=True,
    filled=True
)
    
```



1.75

2.45

ليس يفرق ال classes من بعضه  
هو وقف صون



هو وقف عند قيمة ال Petal ال width (impurity) max depth = 2  
عنا بهن وقف هون

• Decision tree

# Regularization Hyperparameters

- Increase min\_\* or decrease max\_\*: `max_depth=None`, `min_samples_split=2`, `min_samples_leaf=1`, `min_weight_fraction_leaf=0.0`, `max_leaf_nodes=None`, `max_features=None`

min number of samples in each node





overfitting في جزيء من data او outliers  
Probability Class

## Decision Trees Regression

```
from sklearn.tree import DecisionTreeRegressor
tree_reg = DecisionTreeRegressor(max_depth=2)
tree_reg.fit(X, y)
```

① predict  
tree-reg-mse = mean-squared-error (y-test, y-pred)  
rmse = np.sqrt( )  
rmse

Linear يكون range الواسع



صينية ودرجة "بطين انا" تابعة الى samples

مفترضا الفترزة عشان يكون الحل احسن عن طريق زيادة max depth فترزا  
عدد ال splits  
صالح SVM  
سماح smooth اكثر  
وبسبب العلاقة التربيعية

# Outline

1. k-Nearest Neighbors
2. Support Vector Machines
3. Decision Trees
4. Ensemble Learning and Random Forests
5. Exercises

classifiers contains decision trees.

هذه tree بدرجتها على subset من البيانات ، لاني لو درتهم كلهم على كل البيانات او نفس البيانات بطوري نفس الجواب

احد الاشياء على Ensemble learning

بدرجات عدد هائل من Decision trees

مع بعضها .

أدرب أكثر من model مع بعضها وأخذ برأيهم كلهم بصحة .

مثلاً استعمل أكثر من classifier وناخذ برأيهم مثلاً نأخذ ال majority

• ممكن كل واحد منهم يتدرب على البيانات كلها وكل واحد منهم ال طريقة

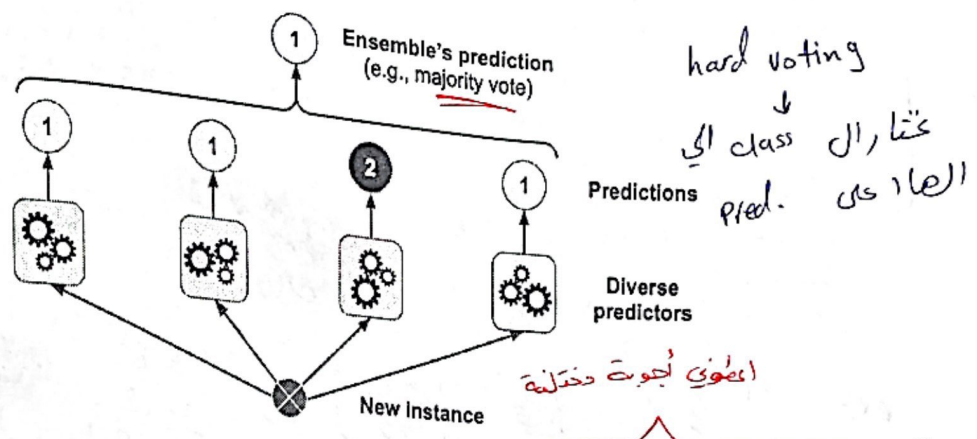
## Ensemble Learning and Random Forests

- A group of predictors is called an ensemble.
- You can train a group of Decision Tree classifiers, each on a different random subset of the training set.
- To make predictions, obtain the predictions of all individual trees, then predict the class that gets the **most votes.**
- Such an ensemble of Decision Trees is called a Random Forest.

# Voting Classifiers

strong classifier يعطوي بعضه بعضه weak classifiers

If each classifier is a **weak learner** (meaning it does only slightly better than random guessing), the **ensemble can be a strong learner** (achieving high accuracy).



example

Data

- C1 [0.51, 0.49]
- C2 [0.2, 0.8]
- C3 [0.51, 0.49]

Hard Voting

0
1
0
0

Soft voting

$P(0) = \frac{0.51 + 0.2 + 0.51}{3} \approx 0.4$

$P(1) = \frac{0.49 + 0.49 + 0.8}{3} \approx 0.6$

P(1) is higher → 1

## Scikit-Learn Voting Classifier 1/2

Hard Voting → لما يكون عدد ال classifiers كبير

```
from sklearn.ensemble import RandomForestClassifier
from sklearn.ensemble import VotingClassifier
from sklearn.linear_model import LogisticRegression
from sklearn.svm import SVC
```

```
log_clf = LogisticRegression()
rnd_clf = RandomForestClassifier()
svm_clf = SVC()
```

يمكن ادرهم كمام على نفس ال data  
 كذا classifiers مختلفين

```
voting_clf = VotingClassifier(
    estimators=[('lr', log_clf), ('rf', rnd_clf), ('svc', svm_clf)],
    voting='hard')
voting_clf.fit(X_train, y_train)
```

voting='soft' predict the class with the highest class probability

اذا binary  
 القيمة الأكثر تكرر  
 يعني كل classifier مع يطين 1 أو 0  
 بتوفتوا أكثر نسبة طلع


والمع  
 بمستين ال احتمال 0 واحتمال 1  
 متا قيم 0 و 1



# Scikit-Learn Voting Classifier 2/2

Soft Voting → Prob. د باض  
ال اضع و بناسا اى اضع بقى

```
>>> from sklearn.metrics import accuracy_score
>>> for clf in (log_clf, rnd_clf, svm_clf, voting_clf):
...     clf.fit(X_train, y_train)
...     y_pred = clf.predict(X_test)
...     print(clf.__class__.__name__, accuracy_score(y_test, y_pred))
...
LogisticRegression 0.864
RandomForestClassifier 0.896
SVC 0.888
VotingClassifier 0.904
```



Sample اى تنكر ال

With replacement → Samples can be given more than once to the same classifier  
كيف حننا، ال انا ال classifiers

## Bagging and Pasting

اما لما يكونوا classifiers مختلفين  
تسرا اعطهم كلهم نفس ال data

تمكن  
استخدم  
ال  
data  
تبعي  
صغيرة

- Use the same training algorithm for every predictor, but train them on different random subsets of the training set.
- When sampling is performed with replacement, this method is called bagging (short for bootstrap aggregating).
- When sampling is performed without replacement, it is called pasting.
- The aggregation function is the most frequent prediction (hard voting) for classification, or the average for regression.

bagging → with replacement

(ال اى متلا لوعنا 10,000 row و randomly او classifier)

Pasting →

ال اى متلا لوعنا 2 rows ، التاني لما شتا ، يكونوا من هتصنق ال انا



# Bagging and Pasting

```
from sklearn.ensemble import BaggingClassifier
from sklearn.tree import DecisionTreeClassifier
```

```
bag_clf = BaggingClassifier(
    DecisionTreeClassifier(), n_estimators=500,
    max_samples=100, bootstrap=True, n_jobs=-1)
bag_clf.fit(X_train, y_train)
y_pred = bag_clf.predict(X_test)
```

no. of decision tree classifiers

with replacement and use all available cores



# Random Forests

هون ٥٨ م  
لستخدا  
tree

الافضل بالواحد  
bagging او voting

- An ensemble of Decision Trees trained via the bagging with `max_samples` set to the size of the training set, and choosing the best random splits.

```
from sklearn.ensemble import RandomForestClassifier

rnd_clf = RandomForestClassifier(n_estimators=500, max_leaf_nodes=16, n_jobs=-1)
rnd_clf.fit(X_train, y_train)

y_pred_rf = rnd_clf.predict(X_test)
```

- Equivalent to:

```
bag_clf = BaggingClassifier(
    DecisionTreeClassifier(splitter="random", max_leaf_nodes=16),
    n_estimators=500, max_samples=1.0, bootstrap=True, n_jobs=-1)
```

split randomly

هل هو  
نفس الشيء

# Outline

1. k-Nearest Neighbors
2. Support Vector Machines
3. Decision Trees
4. Ensemble Learning and Random Forests
5. Exercises

34

## Exercises

1. Train an **SVM classifier** on the **MNIST** dataset. Since SVM classifiers are binary classifiers, you will need to use one-versus-all to classify all 10 digits. You may want to tune the hyperparameters using small validation sets to speed up the process. What accuracy can you reach?

35

## Exercises

2. Train and fine-tune a **Decision Tree** for the **moons dataset**.
  - a) Generate a moons dataset using `make_moons(n_samples=10000, noise=0.4)`.
  - b) Split it into a training set and a test set using `train_test_split()`.
  - c) Use grid search with cross-validation (with the help of the `GridSearchCV` class) to find good hyperparameter values for a `DecisionTreeClassifier`. Hint: try various values for `max_leaf_nodes`.
  - d) Train it on the full training set using these hyperparameters, and measure your model's performance on the test set. You should get roughly 85% to 87% accuracy.

36

## Exercises

3. Load the **MNIST** data and split it into a training set, a validation set, and a test set (e.g., use 50,000 instances for training, 10,000 for validation, and 10,000 for testing). Then train various classifiers, such as a **Random Forest classifier**, an **Extra-Trees classifier**, and an **SVM**. Next, try to combine them into an **ensemble** that outperforms them all on the validation set, using a **soft** or **hard** voting classifier. Once you have found one, try it on the test set. How much better does it perform compared to the individual classifiers?

37

# Summary


1. k-Nearest Neighbors
2. Support Vector Machines
3. Decision Trees
4. Ensemble Learning and Random Forests
5. Exercises

## Unsupervised Learning and Clustering

Prof. Gheith Abandah

# Reference

- Chapter 8: Dimensionality Reduction
- Chapter 9: Unsupervised Learning Techniques



- Aurélien Géron, **Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow**, O'Reilly, 2nd Edition, 2019
  - Material: <https://github.com/ageron/handson-ml2>

2

## Outline

- Dimensionality Reduction
  - Projection and Manifold
  - Principal Component Analysis (PCA)

## • Unsupervised Learning

- Clustering
  - K-Means
  - DBSCAN
- Gaussian Mixtures and Anomaly Detection
- Exercises

ما بتغير و ما البر تأثير على اللى ممكن استير  
+ ما يكون ال Features حافى بينه variance  
( اخلى ال Factor تبعه قليل )  
لما بقل ال execution time

ما بتقل عدد ال samples

used when data has alot of features  
بيستخ PCA لما انزل عدد ال Features  
دا بقل الحافى على ال variance  
ممكن يصير loss صغير

3

# Dimensionality Reduction

- Many Machine Learning problems involve thousands or even millions of features for each training instance.
- All these features make training extremely **slow** and make it much harder to find a good solution.
- This problem is often referred to as the **curse of dimensionality**.



## Dimensionality reduction approaches

1. Drop not useful features
2. Merge correlated features
3. Projection and manifold
4. Transform features

مثلاً بال MNIST لو لخصنا حواف الصور بنقل عدد ال features بالتالي بنقل 784 على افتراض انه كل الصور الرقم فيها بالنص.

## Projection and Manifold

حساب كل Feature  
عبارة عن dimension




Projection  
2 new features  
Plane  
فقلنا عدد ال features من 3 ل 2

دیتا کے انداز

# Projection and Manifold

Swissroll data



داتا  
ال  
Features.



- Simply projecting onto a plane may not give better solution.
- Projecting to a proper manifold is better.

فی تاخیر !!

Manifold




axis-2  
بہار، شکل

جہازتہ ال داتا  
Projection




Surface  
z1  
new feature

# Projection and Manifold

- The decision boundary may not always be simpler with lower dimensions.



Projection



decision boundary

manifold

Classification

decision tree  
with depth 3

decision boundaries



Surface الجهد الی بیسی اصل علی projection یکنون بانظ علی ال  
 variance ل data - عتبان الی غیر new Features

انظر بانظوا  
 Variance علی ال


one of dimensionality reduction techniques.

# Principal Component Analysis (PCA)

• و عدد من الطرق الی نقل منصف projection

- Is the most popular dimensionality reduction algorithm.
- First it identifies the hyperplane that lies closest to the data, and then it projects the data onto it.
- PCA identifies the axis that accounts for the largest amount of variance in the training set. Then it finds the next orthogonal axes that accounts for the largest amount of remaining variance.

$C_1, C_2$   
 مقاطع



یعنی هاد الخط صتا کثیر بعبر  
 عن ال variation فی ال data  
 و صافی دایم الی غیره new Feature

• ال variance عالی (التغیر بال feature عالی)  
 ال variance عالی (التغیر بال feature عالی)


اول سنی طبعت  $C_1$  بعین  $C_2$   
 کل next بكون عمودي علی  $C_1$  بیک  
 Principal Component

# Principal Component Analysis (PCA)

- Use PCA to reduce the dimensionality of the dataset down to two dimensions.
- Instead of specifying the number of principal components you want to preserve, you can set n\_components to be a float between 0.0 and 1.0, indicating the ratio of variance you wish to preserve.

```
from sklearn.decomposition import PCA
pca = PCA(n_components = 2)
X2D = pca.fit_transform(X)
```

```
pca = PCA(n_components=0.95)
X_reduced = pca.fit_transform(X_train)
```



بتریزد Features  
 ما یبقی variance

بدل  $X_1$  و  $X_2$  و  $X_3$   
 بقیه داتا جدیدة  $Z_1$  و  $Z_2$


# Outline

- Dimensionality Reduction
  - Projection and Manifold
  - Principal Component Analysis (PCA)
- Unsupervised Learning
- Clustering
  - K-Means
  - DBSCAN
- Gaussian Mixtures and Anomaly Detection
- Exercises

بافتوا اليراع الابلاب  
 cluster center  
 د بر جوال  
 مع ال  
 cluster

# Clustering

- The task of identifying similar instances and assigning them to clusters, i.e., groups of similar instances.
- Classification (left) versus clustering (right)



ما في اي طريقة تميز اليراع  
 لازم نظيرة algorithms

# Clustering Applications

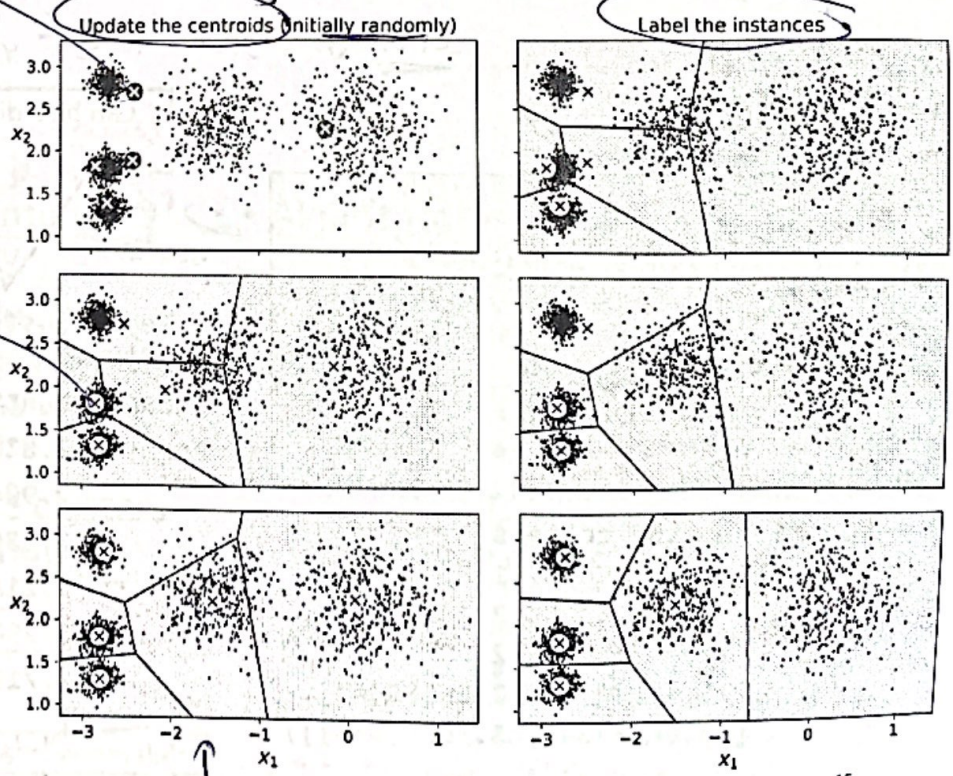
- Customer segmentation: useful for recommender systems. *مثلا حساب ال shopping behavior*
- Data analysis: discover clusters of similar instances as it is often easier to analyze clusters separately.
- Dimensionality reduction: find affinity features to the found clusters
- Anomaly detection: any instance that has a low affinity to all the clusters is likely to be an anomaly.
- Semi-supervised learning: perform clustering and propagate the labels to all the instances in the same cluster.
- Search engines for images *كل صورة من ال pixels / شوية تمشي*
- Image segmentation *كل نقطة في ال data*

$MSE$  ← كل نقطة في ال data  
 اذا كانت ال center عالية منا كويس يعني المسافة كبيرة  
 Score → avg distance between data & center point  
 ال data تتبع اقرب center اليها

## K-Means

- Quick and efficient algorithm
- Scale before clustering
- Need to specify the number of clusters

\* لانه بيتكلم  
 وبتغير المسافة  
 بحرب ارقام



The best

كل نقطة مع تاني label  
 وهو ال cluster الي  
 قريب


Score = inertia

# K-Means

(كل instance  $x$  في اي cluster هو موجود)

برجع array مع كل نقطة في اى cluster  
تم اى cluster

- Cluster to 5 clusters




```

from sklearn.cluster import KMeans
k = 5
kmeans = KMeans(n_clusters=k)
y_pred = kmeans.fit_predict(X)
y_pred
array([4, 0, 1, ..., 2, 1, 0],
      dtype=int32)
# Hard clustering: cluster 1 cluster 2
X_new = np.array([[0, 2], [-3, 3]])
kmeans.predict(X_new)
array([1, 2], dtype=int32)

```

## Ch.9 K-Means in the basic.



```

kmeans.cluster_centers_
array([[ -2.80389616,  1.80117999],
       [  0.20876306,  2.25551336],
       [-2.79290307,  2.79641063],
       [-1.46679593,  2.28585348],
       [-2.80037642,  1.30082566]])

```

Can be a dimensionality reduction technique.

مركز كل نقطة

Sample Size 5 cluster center

# Soft clustering, a score per cluster:

```

kmeans.transform(X_new)
array([[2.81093633, 0.32995317,
       2.9042344 , 1.49439034,
       2.88633901],
       [1.21475352, 3.29399768,
       0.29040966, 1.69136631,
       1.71086031]])

```

كل نقطة تابعة Cluster 2

ياخذ اقل قيمة

# K-Means- Centroid Initialization

5 cluster centers "user defined"

لما يبدأ عملية ال clustering ابتدا

منهم

- User Defined Initial values

```
good_init = np.array([[ -3, 3], [ -3, 2], [ -3, 1], [ -1, 2], [ 0, 2]])
kmeans = KMeans(n_clusters=5, init=good_init, n_init=1)
```

- Random Initialization

- Randomly initialize centroids
- Repeat experiment **n\_init** times
- Select the model with lower **inertia** (Minimum mean diastance between the instances and the centroids)

بعمل التجربة مرة وحدة، لما تكون ٢  
مثلا بعمل التجربة مرتين ب  
initialization مختلف



بحكي كيف شغلي وكم المسافة الي  
وصلناها عشان نقدر نقارن  
(score) المسافة بين كل نقطة وال  
cluster center الي ربطتها فيه

```
>>> kmeans.inertia_
211.59853725816856
```

كل ما كانت قيمته اقل بكون افضل

# K-Means

- It is important to specify the **right** number of clusters **k**.
- Inertia is not a good performance measure because it is getting lower as we increase **k**.



من 4 ل 5 ومن 5 ل 6 وبعد صارت  
التعديلات قليلة وال improvement  
قل كثير فمممكن نحكي 5 او 4 منيح

# K-Means


امتي ممكن يطلع معنا 1 ؟ لما تكون a قليلة جدا مقارنة بال b او تساوي صفر يعني بنكون بالمكان الصح يعني كل ال samples بنفس ال cluster قريبة وببقي ال clusters بعيدة مسافة

- Find  $k$  that gives highest mean silhouette coefficient.


$$\text{Silhouette coefficient} = \frac{b - a}{\max(a, b)}$$

- a: the mean distance to the instances in the same cluster
- b: the mean distance to the instances in the next closest cluster
- The score is between -1 and 1

```
from sklearn.metrics import
    silhouette_score
silhouette_score(X, kmeans.labels_)
0.655517642572828
```



# K-Means




The best

Figure 9-10. Analyzing the silhouette diagrams for various values of k

# DBSCAN

كل ما يتكون في cluster هو كانه يعمل clustering لهذا

- Defines clusters as continuous regions of high density.
- Works well if all the clusters are dense enough, and they are well separated by low-density regions.
- Behaves well when the clusters have varying sizes or non-spherical shapes.



Can detect anomalies

البيانات التي لها كثافة منخفضة بقدر اقل وتساوي ع يكونوا بقية ال cluster

DBSCAN is faster than K-means

## Algorithm


- For each instance, counts how many instances are located within a small distance  $\epsilon$ -neighborhood.
- If an instance has at least min\_samples instances in its  $\epsilon$ -neighborhood, then it is considered a core instance.
- All instances in the neighborhood of a core instance belong to the same cluster. This may include other core instances; therefore, a long sequence of neighboring core instances forms a single cluster.
- Any instance that is not a core instance and does not have one in its neighborhood is considered an anomaly (-1).

ع اذا ما لقيت حولي نقطة اي core قريب فتر او وة نقطة لها كثافة بينه وبينها جملد ع

بغني يكون عذي نقطه وهو ليس متعلق  
المساحة فتر وبينهم ع او اقل min-samples  
core بقية ال instance

كل النقاط البعيدة يكون ال behavior واد  
 يكون ال  
 1- = cluster  
 cluster number

# DBSCAN



```
from sklearn.cluster import DBSCAN
from sklearn.datasets import make_moons
X, y = make_moons(n_samples=1000, noise=0.05)
dbscan = DBSCAN(eps=0.2, min_samples=5)
dbscan.fit(X)
```

7 clusters  
 Red X → anomaly  
 كل هاي يتسكه عند طريق  
 اني أكبر E

2 clusters  
 ما في شئ  
 anomaly

# DBSCAN

→ have fit\_predict()


• DBSCAN class does not have a predict() method.

• Can use other classifiers.

\* check DBSCAN doc.

```
from sklearn.neighbors import KNeighborsClassifier
knn = KNeighborsClassifier()
knn.fit(dbscan.components_, dbscan.labels_[dbscan.core_sample_indices_])
X_new = np.array([[ -0.5, 0], [0, 0.5], [1, -0.1], [2, 1]])
knn.predict(X_new)
array([1, 0, 1, 0])
```

ش ريم ار  
 cluster



2 clusters

non linear decision boundary



بتقادل تفل mimic  
للجهاز العصبي جسم  
الإنسان


# Neural Networks

ميزات استخراج  
Features extraction

Prof. Gheith Abandah

## Reference

- Chapter 10: Introduction to Artificial Neural Networks with Keras



- Aurélien Géron, **Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow**, O'Reilly, 2nd Edition, 2019
  - Material: <https://github.com/ageron/handson-ml2>

# Introduction

YouTube Video: *But what \*is\* a Neural Network?* from 3Blue1Brown

<https://youtu.be/aircAruvnKk>

3

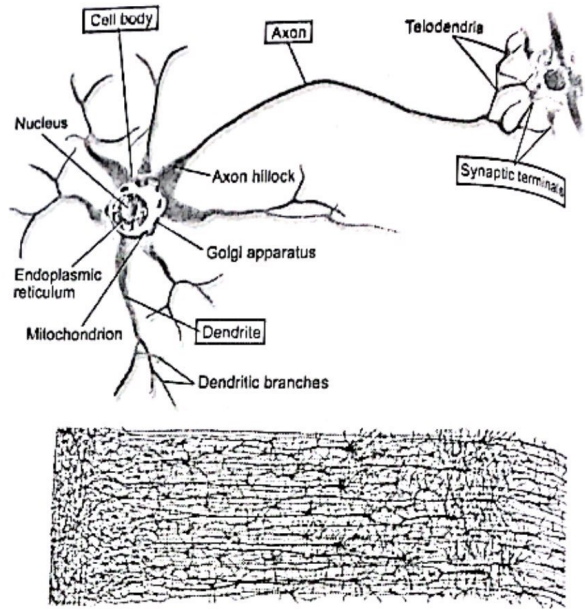
## Outline

1. Introduction
2. The perceptron
3. Multi-layer perceptron (MLP)
4. Regression MLPs
5. Classification MLPs

4

# 1. Introduction

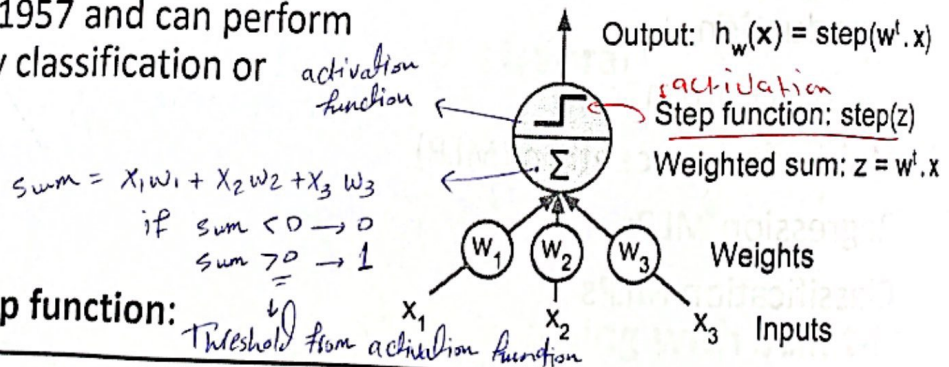
- Artificial neural networks (ANNs) are inspired by the brain's architecture.
- First suggested in 1943. Is now flourishing due to the availability of:
  - Data
  - Computing power
  - Better algorithms



# 2. The Perceptron *Feature weights and activation function*

- The Perceptron is a simple ANN, invented in 1957 and can perform linear binary classification or regression.

## Linear threshold unit (LTU)




$$\text{sum} = x_1 w_1 + x_2 w_2 + x_3 w_3$$

if  $\text{sum} < 0 \rightarrow 0$   
 $\text{sum} > 0 \rightarrow 1$

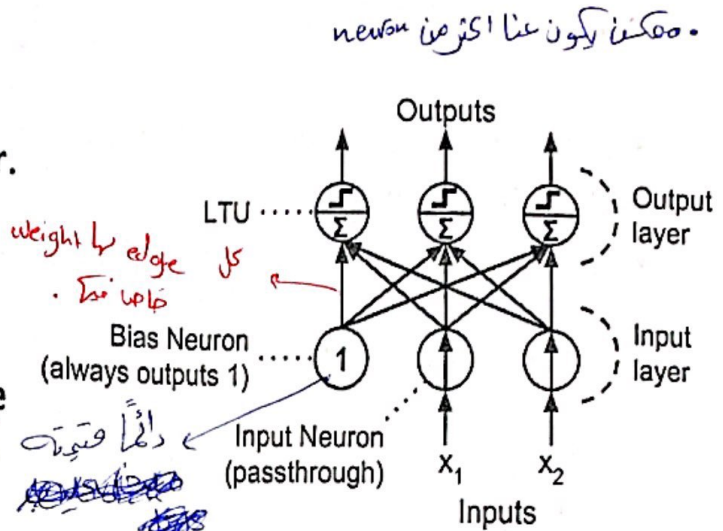
- Common step function:

$$\text{heaviside}(z) = \begin{cases} 0 & \text{if } z < 0 \\ 1 & \text{if } z \geq 0 \end{cases} \quad \text{sgn}(z) = \begin{cases} -1 & \text{if } z < 0 \\ 0 & \text{if } z = 0 \\ +1 & \text{if } z > 0 \end{cases}$$



## 2. The Perceptron

- The Perceptron has an input layer with bias and output layer.
- With multiple output nodes, it can perform multiclass classification.
- Hebbian learning "Cells that fire together, wire together."



$$w_{i,j} \text{ (next step)} = w_{i,j} + \eta(y_j - \hat{y}_j)x_i$$

كيف نعدل ال weights  
بالتقريب من gradient descent

→ زي فكرة ال learning rate

• ممكن نعدل ال weights  
بكل neuron مشان نعدل  
ال نتيجة منيرة

• ممكن نعدل ال weight of  $x_2$  ب neuron 1 و neuron 2  
ال neuron بدي ال combination ال features ال دخلتي ال weight او باجي ال weight

## 2. The Perceptron

- Scikit-Learn provides a Perceptron class.

```
import numpy as np
from sklearn.datasets import load_iris
→ from sklearn.linear_model import Perceptron

iris = load_iris()
X = iris.data[:, (2, 3)] # petal length, petal width
y = (iris.target == 0).astype(np.int) # Iris Setosa?

per_clf = Perceptron(random_state=42)
per_clf.fit(X, y)

y_pred = per_clf.predict([[2, 0.5]])
```

sklearn

• ال Perceptron ال

neural ال network

ال ال

data preprocessing

• data split

\* يسفل اكثر من Mathematical model مع بعض



هدفنا نعرف مشوار  
 عشان نطلع accurate answer  
 على ما يتغيرال network deeper  
 بتغير العلاقات = ادق

## 2. The Perceptron

- The perceptron cannot solve non-linear problems such as the XOR problem.
- The Multi-Layer Perceptron (MLP) can.

$x_1$	$x_2$	$F$
0	0	0
0	1	1
1	0	1
1	1	0

Hy for all



2 layers


bias

$$x_2 w + x_1 w + 1 * -0.5 = -0.5 \rightarrow \text{step function} = 0$$

Cont.

$$F = \theta_0 + \theta_1 x_1 + \theta_2 x_2$$

### Outline




1. Introduction
2. The perceptron
3. Multi-layer perceptron (MLP)
4. Regression MLPs
5. Classification MLPs

### 3. Multi-Layer Perceptron (MLP)

needs computational power

- An MLP is composed of a (pass-through) **input layer**, one or more layers of LTUs, called **hidden layers**, and a final layer of LTUs called the **output layer**.



- When an ANN has two or more hidden layers, it is called a deep neural network (DNN).

output ال hidden ال 15 edge.

ليا Input layer وحدة و output layer وحدة.

deep network hidden layers  
neuron feature بتدخل على ال neuron  
12 weights 12 edges

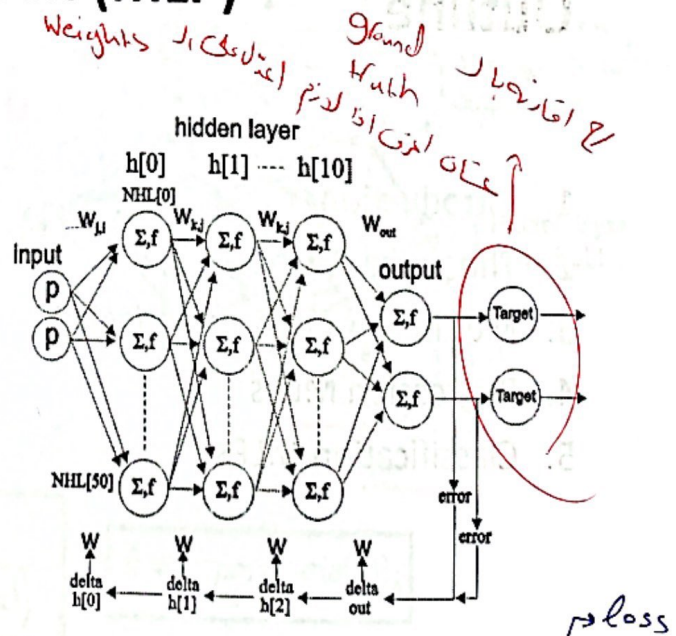
neuron ما بتكون  
حين عبارة عن features

عدد ال weights وال edges  
هو نفس عدد ال parameters

### 3. Multi-Layer Perceptron (MLP)

- Trained using the **backpropagation training algorithm**.

- For each training instance the algorithm first makes a prediction (**forward pass**), measures the **error**,
- then goes through each layer in reverse to measure the error contribution from each connection (**reverse pass**),
- and finally slightly **tweaks the connection weights to reduce the error (Gradient Descent step)**.



- Forward pass "forward propagation" → from input to output. (to find pred.)  
بدينا حسب قيمة ال prediction weights بترجع عدد ال weights.
- backpropagation ← weights ال عدد ال

backpropagation ↘

بمنشغال  $cost$  بالنسبة لكل  $weight$  ونشوف كيف كل  $weight$  تأثيره  
 وبعده اذا ازديده او نقله وهكذا بنعرف قيمه الجديدة

### \* 3. Multi-Layer Perceptron (MLP)

- Common activation functions: logistic, hyperbolic tangent, and rectified linear unit.

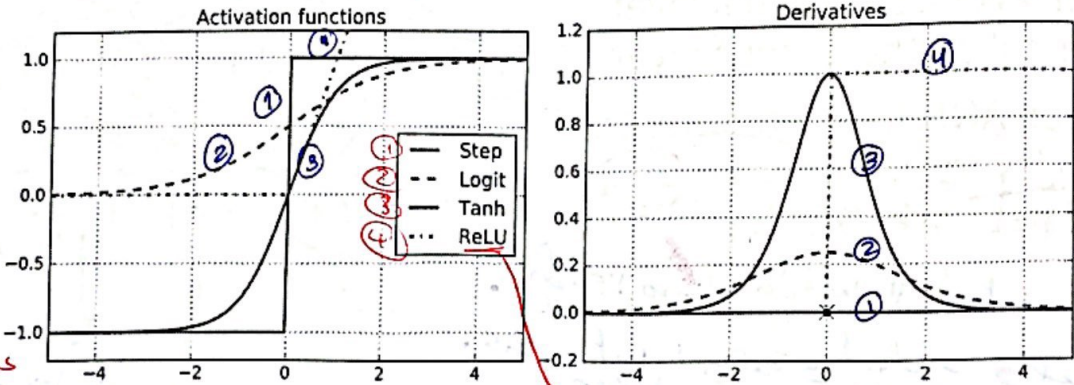
$$\sigma(z) = 1 / (1 + \exp(-z))$$

$$\tanh(z) = 2\sigma(2z) - 1$$

$$\text{ReLU}(z) = \max(0, z)$$

4 step function  
 خيرا قابل للاشتغال عند ال zero  
 فدما بنستعمله

بهذا ال  
 activation  
 function  
 يكون قابل  
 للاشتغال  
 عتبان بحتاج  
 عملية للاشتغال



لما تكونا بيدي اثره  
 شولانم ابدال على ال  
 edges weight.

بناء على قيمة ال step  
 يعرف قيمة ال step  
 التي لازم احدثها.

is the best activation function.

### Outline

1. Introduction
2. The perceptron
3. Multi-layer perceptron (MLP)
4. Regression MLPs
5. Classification MLPs

" Backpropagation "

$$\sigma = \frac{1}{1 + e^{-x}}$$


input layer → 2 inputs (features)

1 hidden layer

one output layer

$z_1 / z_2 / z_3$  → summation

$f_1 / f_2 / f_3$  → activation functions



$$z_1 = X_1 w_1 + X_2 w_3$$

$$f_1 = \sigma(z_1)$$

$$z_2 = X_1 w_2 + X_2 w_4$$

$$f_2 = \sigma(z_2)$$

$$z_3 = f_1 w_5 + f_2 w_6$$

$$f_3 = \sigma(z_3)$$

$$f = \sigma(z)$$

cost funct. → MSE

• assume "Regression" and our target is 500, and the output is 700  
↳ prediction

• error = 200

• if binary classifier → cost function =  $-\log(\hat{y}) - (1 - \hat{y}) \log(1 - \hat{y})$

learning rate ← → كيف الـ cost بتغير

$$w_{5 \text{ new}} = w_{5 \text{ old}} - \eta \left( \frac{\partial \text{Cost}}{\partial w_5} \right)$$

بتغير  $w_5$

remember: chain rule →  $\frac{dy}{dz} = \frac{dy}{dx} \cdot \frac{dx}{dz}$  لاسه العلوية كذا، لا تسيء matrix multiplication

$$\text{so } \rightarrow \frac{dl}{dw_5} = \frac{dl}{f_3} \cdot \frac{df_3}{dz_3} \cdot \frac{dz_3}{dw_5} \text{ where } f_3 = \hat{y}$$

بتغير الـ cost بالنسبة لـ  $f_3$  وبتغير  $f_3$  بالنسبة لـ  $z_3$

$$= f_3(1 - f_3) \quad \hookrightarrow = f_1$$



# 4. Regression MLPs

- Typical MLP architecture for regression:

neuron activation function

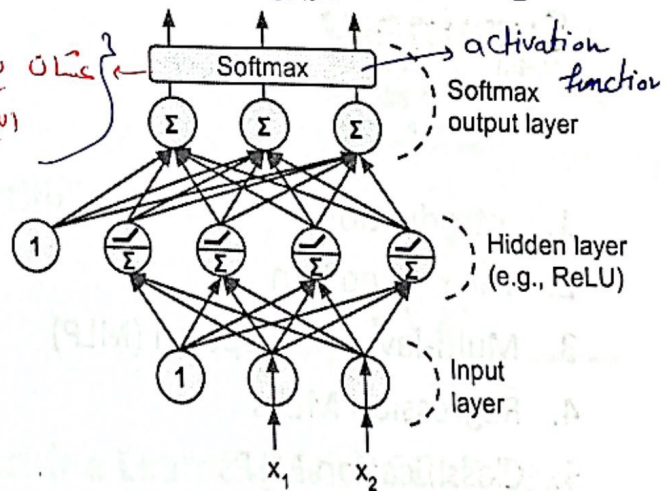
Hyperparameter	Typical Value
# input neurons	One per input feature (e.g., 28 x 28 = 784 for MNIST)
# hidden layers	Depends on the problem. Typically 1 to 5.
# neurons per hidden layer	Depends on the problem. Typically 10 to 100.
# output neurons	1 per prediction dimension
Hidden activation	ReLU (or SELU, see Chapter 11)
Output activation	None or ReLU/Softplus (if positive outputs) or Logistic/Tanh (if bounded outputs)
Loss function	MSE or MAE/Huber (if outliers)

If regression  $\sigma(x)$  not activation or you can use ReLU.   
 Mean Absolute Error   
 graph of MAE function

# 5. Classification MLPs

probabilities normalization  $\sum = 1$

- For classification, the output layer uses the softmax function.
- The output of each neuron corresponds to the estimated probability of the corresponding class.



$$\hat{p}_k = \sigma(s(x))_k = \frac{\exp(s_k(x))}{\sum_{j=1}^K \exp(s_j(x))}$$

$$\hat{y} = \operatorname{argmax}_k \sigma(s(x))_k$$

output its 10 neurons  $\rightarrow$  10 classes   
 softmax multi class activation function

# 5. Classification MLPs

- Typical MLP architecture for classification:

Hyperparameter	Binary classification	Multilabel binary classification	Multiclass classification
Input and hidden layers	Same as regression	Same as regression	Same as regression
# output neurons	1	1 per label	1 per class
Output layer activation	Logistic	Logistic	Softmax
Loss function	Cross-Entropy	Cross-Entropy	Cross-Entropy

كل label يعطى جواب معين

$$\sum y \log \hat{y} + (1-y) \log (1-\hat{y})$$

$y$ : Actual

$\hat{y}$ : Predicted

## Summary

1. Introduction
2. The perceptron
3. Multi-layer perceptron (MLP)
4. Regression MLPs
5. Classification MLPs

# \* Artificial Neural Networks with Keras

Prof. Gheith Abandah


1

## Reference

- Chapter 10: Introduction to Artificial Neural Networks with Keras

O'REILLY  
Hands-On  
Machine Learning  
with Scikit-Learn,  
Keras & TensorFlow  
Concepts, Tools, and Techniques  
to Build Intelligent Systems

powered by  
MISTRA




Aurélien Géron

- Aurélien Géron, **Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow**, O'Reilly, 2nd Edition, 2019
  - Material: <https://github.com/ageron/handson-ml2>

2

## Reference



- **Deep Learning with Python**, by François Chollet, Manning Pub, 2018
- **Introduction to Keras** by Francois Chollet, March 9th, 2018 (slides)

3

## Outline

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

\* Keras is higher level than  
TensorFlow

4

# Introduction

- YouTube Video: *Keras Explained* from Siraj Raval

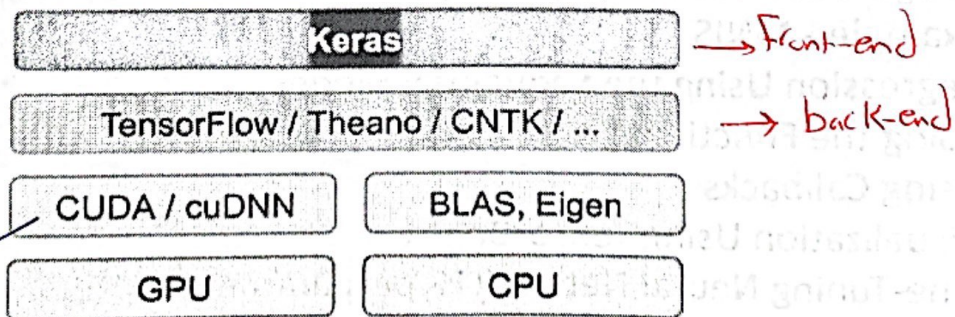
[https://youtu.be/j\\_pJmXJwMLA](https://youtu.be/j_pJmXJwMLA)

5

## 1. Introduction

- Keras is a high-level API to build and train deep learning models.

by default  
Keras dl  
TensorFlow dl  
deep neural dl  
CPU dl  
2. training acceleration  
GPU dl



6

# 1. Introduction – Advantages

- **User friendly**: Keras has a simple, consistent interface optimized for common use cases. It provides clear and actionable feedback for user errors.
- **Modular and composable**: Keras models are made by connecting configurable building blocks together, with few restrictions.
- **Easy to extend**: Write custom building blocks to express new ideas for research. Create new layers, loss functions, and develop state-of-the-art models.

## Outline

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

\*In Neural Network, you don't need to extract features

## 2. Keras API Styles

كيف نكتب ال model ؟

### 1. The Sequential Model

- Dead simple
- Only for single-input, single-output, sequential layer stacks
- Good for 70+% of use cases

### 2. The functional API → ال layer ممكن تاخذ من اي layer

- Like playing with Lego bricks
- Multi-input, multi-output, arbitrary static graph topologies
- Good for 95% of use cases

### 3. Model subclassing

- Maximum flexibility
- Larger potential error surface

→ more advanced.

Flexibility ↑, Error ↑

كل layer بتاخذ من ال layer ابي قبلها

activation

Relu

Relu

layer-1

## Outline


1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

### 3. TensorFlow Keras

- Keras is the official high-level API of TensorFlow
- tensorflow.keras (tf.keras) module
- Part of core TensorFlow since v1.4
- Full Keras API
- With useful extra features such as tf.data

↳  
 lines لكتابة  
 data preparation

هو Keras الواجهة  
 tensorflow من



↓  
 For acceleration

### 3. TensorFlow Keras

- To install TensorFlow

```
[ $ pip install --upgrade tensorflow ]
```

بكون من كلاس  
 Keras ال

- To import Keras from TensorFlow

```
>>> import tensorflow as tf
>>> from tensorflow.keras import Layers
>>> from tensorflow import keras
>>> tf.__version__
'2.1.0'
>>> keras.__version__
'2.2.4-tf'
```

عشان نعلم  
 الجواب

- Dense
- Activations
- Dropout
- Conv1D, 2D, 3D
- Polling
- RNN, LSTM, GRU
- ...



# Outline

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

13

## \* 4. Image Classifier Using the Sequential Model

بداية نبنى neural network  
تخبر نوع الملابس بالصورة

يتكون من 10 classes

input

- Fashion MNIST is similar to MNIST (70,000 grayscale images of 28x28 pixels each, with 10 classes).



14

# 4. Fashion MNIST

1. Get and prepare the dataset.
2. Build sequential model of layers that maps your inputs to your targets.   
 *عشان خذ شوية loss*
3. Compile the model and configure the learning process by choosing a loss function, an optimizer, and some metrics to monitor.
4. Train the model by calling the fit() method of your model.
5. Evaluate and use the model.

بنبدأ على  
بكون عن افكرة  
neuron  
و عن عدد  
hidden layer  
و عن neuron  
output

ال validation ثابتة كل لحظة ال training  
بسرنا ببطي hint

## 4.1 Get and Prepare the Dataset

```
import tensorflow as tf
from tensorflow import keras
```

ال بنحول كل صورة ل 1D array وهاي  
ال 1D هي ال "feature"  
كل Pixel هي "feature"


```
# Get the Fashion MNIST
fashion_mnist = keras.datasets.fashion_mnist
(X_train_full, y_train_full), (X_test, y_test) =
    fashion_mnist.load_data()
```

ال deep neural  
ما بافدوا 2D  
Arrays

```
# Prepare the data train (55000), val (5000), test (10000)
```

```
X_valid = X_train_full[:5000] / 255.
X_train = X_train_full[5000:] / 255.
y_valid, y_train = y_train_full[:5000], y_train_full[5000:]
X_test = X_test / 255.
```

Validation  
normalization  
عشان كل  
Pixels  
ديمر قيمه بين 0 و 1  
بدلنا 0 و 255



# 4.2 Build the Model

جول الصورة من 28x28 الى vector

```

model = keras.models.Sequential()
model.add(keras.layers.Flatten(input_shape=[28, 28]))
model.add(keras.layers.Dense(300, activation="relu"))
model.add(keras.layers.Dense(100, activation="relu"))
model.add(keras.layers.Dense(10, activation="softmax"))
    
```

The default is no activation function, i.e., linear layer.

hidden layer  
output layer

```

>>> model.summary()
Layer (type)                 Output Shape         Param #
-----
flatten_1 (Flatten)         (None, 784)         0
dense_3 (Dense)              (None, 300)         235500
dense_4 (Dense)              (None, 100)         30100
dense_5 (Dense)              (None, 10)          1010
-----
Total params: 266,610
Trainable params: 266,610
Non-trainable params: 0
    
```

For more details: [keras.io](https://keras.io)

$$300 \times 784 = 235200$$

$$+ 300$$

$$= 235500$$

بي ال probabilities الى  
بتطلع يكون مجموعها = 1

$$(300 \times 100) + 100$$

From bias

طالع فيه 300 edge

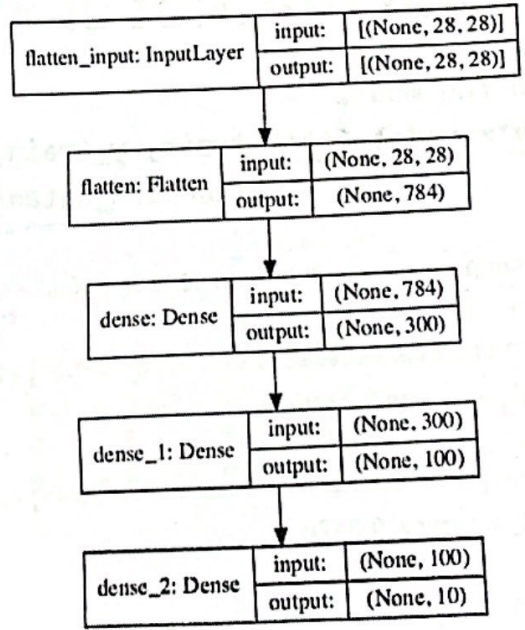
عدد ال (samples) فووس  
نتر عدد  
بتحدد ال  
الطب Samples  
(ال ك اليا بس)  
بنيت model

```
model.summary()
```

# 4.2 Build the Model

```

# Plot the model
keras.utils.plot_model(
    model,
    "my_model.png",
    show_shapes=True)
    
```



\* For binary classification  
=> Single neuron output  
activation -> could be **Sigmoid**

## 4.3 Compile the Model

multi classification


```
model.compile(loss="sparse_categorical_crossentropy",
              optimizer="sgd",
              metrics=["accuracy"])
```

كيف نعدل ال weight

Stochastic Gradient Descent

```
# For sparse labels (0-9):
loss = "sparse_categorical_crossentropy"
# For one-hot labels:
loss = "categorical_crossentropy"
# For binary labels:
loss = "binary_crossentropy"
# For regression:
loss = "mean_squared_error"
```

$$-\sum_{c=1}^M y_{o,c} \log(p_{o,c})$$



optimization for loss

```
validation data
X_valid = X_train.iloc[10000:]
Y_valid = Y_train.iloc[10000:]
```

categorical\_crossentropy  
Hot encoding

## 4.4 Train the Model

one epoch:

```
# Train the model
history = model.fit(X_train, y_train, epochs=30,
                   validation_data=(X_valid, y_valid))
```

round around all data

بجرون

Train on 55000 samples, validate on 5000 samples

Epoch 1/30

55000/55000 [=====] - 2s 44us/sample - loss: 0.7226 - accuracy: 0.7641 - val\_loss: 0.5073 - val\_accuracy: 0.8320

Epoch 2/30

55000/55000 [=====] - 2s 39us/sample - loss: 0.4844 - accuracy: 0.8321 - val\_loss: 0.4541 - val\_accuracy: 0.8478

...

Epoch 30/30


55000/55000 [=====] - 2s 39us/sample - loss: 0.2256 - accuracy: 0.9195 - val\_loss: 0.3049 - val\_accuracy: 0.8882

epoch كل 20 اشياء accuracy و loss

epoch ليس احسن

## 4.4 Train the Model

```
import pandas as pd
pd.DataFrame(history.history).plot(figsize=(8, 5))
plt.grid(True)
plt.gca().set_ylim(0, 1)
save_fig("keras_learning_curves_plot")
plt.show()
```



ووصلنا المرحلة  
التي ان  
Validation loss  
لا يتغير  
في حين ان  
ال loss training  
تقل

## 4.5 Evaluate and Use the Model

لذا كبرية same as predict() in scikit learn.

```
model.evaluate(X_test, y_test)
10000/10000 [=====] - 0s 21us/sample - loss: 0.3378 -
accuracy: 0.8781
[0.33780701770782473, 0.8781]
```

```
X_new = X_test[:3]
y_proba = model.predict(X_new)
y_proba.round(2)
array([[0. , 0. , 0. , 0. , 0. , 0. , 0. , 0.01, 0. , 0.99], → sample 1
       [0. , 0. , 0.99, 0. , 0.01, 0. , 0. , 0. , 0. , 0. ], → sample 2
       [0. , 1. , 0. , 0. , 0. , 0. , 0. , 0. , 0. , 0. ]], → " 3
dtype=float32)
```

ما يتعلق ال loss / ما يتعلق ال accuracy  
لذا كبرية / ما يتعلق ال loss / ما يتعلق ال accuracy  
بشوف كل sample شو جوبه

```
model.predict_classes(X_new)
array([9, 2, 1])
```

Classes

كانه افنا  
نتائج ال Predict  
وكانه ال max

# Outline

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

23

## 5. Example - MNIST

1. Define your training data: input tensors and target tensors.
2. Define a network of layers (or model) that maps your inputs to your targets.
3. Configure the learning process by choosing a loss function, an optimizer, and some metrics to monitor.
4. Iterate on your training data by calling the `fit()` method of your model.

24

## 5. Example – Prepare the data

```
from keras.datasets import mnist
(train_images, train_labels), (test_images, test_labels) =
    mnist.load_data()
```

```
train_images = train_images.reshape((60000, 28 * 28))
train_images = train_images.astype('float32') / 255
test_images = test_images.reshape((10000, 28 * 28))
test_images = test_images.astype('float32') / 255
```

reshape to make them as 1 vector "1D"

```
from keras.utils import to_categorical #one hot
```

```
train_labels = to_categorical(train_labels)
```

```
test_labels = to_categorical(test_labels)
```

عنانه قيم ال  
pixels  
بالصوت بلا ما تكون  
من 0 - 255  
بتصير من 0 - 1

25

one hot encoding  
كل Label عبارة عن array  
من 10 ارقام واحد منهم يساوي 1  
اي هو رقم الكلاس

## 5. Example – Define and configure the network

```
from keras import models
from keras import layers
```

first hidden

```
network = models.Sequential()
network.add(layers.Dense(512, activation='relu', input_shape=(28 * 28,)))
```

```
network.add(layers.Dense(10, activation='softmax'))
```

output layer

```
network.compile(optimizer='rmsprop',
                loss='categorical_crossentropy',
                metrics=['accuracy'])
```

لا يكون ال  
Hot encoding

26

epochs کی تعداد سے ← 60000/128

## 5. Example – Training and evaluation

```
network.fit(train_images, train_labels, epochs=5, batch_size=128)
```

Validation data  
↓  
Validation accuracy

```
Epoch 1/5  
60000/60000 [=====] - 2s - loss: 0.2577 - acc: 0.9245  
Epoch 2/5  
60000/60000 [=====] - 1s - loss: 0.1042 - acc: 0.9690  
Epoch 3/5  
60000/60000 [=====] - 1s - loss: 0.0687 - acc: 0.9793  
Epoch 4/5  
60000/60000 [=====] - 1s - loss: 0.0508 - acc: 0.9848  
Epoch 5/5  
60000/60000 [=====] - 1s - loss: 0.0382 - acc: 0.9890
```

```
test_loss, test_acc = network.evaluate(test_images, test_labels)
```

```
9536/10000 [=====>..] - ETA: 0s
```

test\_acc: 0.9777

27

## Outline

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

28



neuron output

بال sequential layers  
كيفية layers  
وذا يعني

dense  
لا neuron  
في layer  
من input  
output

## 6. Regression Using the Sequential Model

- Solve the California housing problem using a regression neural network.
- Scikit-Learn has fetch\_california\_housing() function to load the data
- This dataset contains only numerical features and there are no missing values.

use encoding

Regression →  
Probability →  
Soft max

### 6.1 Get and Prepare the Dataset

```
from sklearn.datasets import fetch_california_housing
from sklearn.model_selection import train_test_split
from sklearn.preprocessing import StandardScaler
```

```
housing = fetch_california_housing()
```


The default is 75% : 25%

```
X_train_full, X_test, y_train_full, y_test =
    train_test_split(housing.data, housing.target, random_state=42)
X_train, X_valid, y_train, y_valid = train_test_split(X_train_full,
    y_train_full, random_state=42)
```

```
scaler = StandardScaler()
X_train = scaler.fit_transform(X_train)
X_valid = scaler.transform(X_valid)
X_test = scaler.transform(X_test)
```

check IP ocean proximity

80% : 20%  
75% : 25%



## 6.2 Build and Compile the Model

# Building by passing a list of layers when creating the Sequential model. *طريقة ثانية*  
 # the Sequential model *عشان نعمل المودل*

```

model = keras.models.Sequential([
    keras.layers.Dense(30, activation="relu",
        input_shape=X_train.shape[1:]),
    keras.layers.Dense(1)
])
model.compile(loss="mean_squared_error",
    optimizer=keras.optimizers.SGD(lr=1e-3))
    
```


*ادخل layer* (pointing to the list of layers)  
*layer* (pointing to each layer)  
*برجع ال rows و ال columns للبيانات* (pointing to input\_shape)  
*يعني لو كانت ال data 15 feature* (pointing to [1:])  
*برجع ال 15* (pointing to [1:])  
*The default is 0.01* (pointing to lr=1e-3)  
*learning rate* (pointing to lr=1e-3)  
*لا بد regression* (pointing to the compile line)

## 6.3 Train and Evaluate the Model

```

history = model.fit(X_train, y_train, epochs=20,
    validation_data=(X_valid, y_valid))
    
```

*بيكون optional*  
*بعض ال options احسن*  
*عشان نعمل monitoring*



*ال error عم يقل*

```

mse_test = model.evaluate(X_test, y_test)
5160/5160 [=====] - 0s
15us/sample - loss: 0.421
    
```

*لا بد اني عاين*  
 check → *Scaling*  
*Labels*

## 6.4 Save and Restore the Model

- After training a model save it to a file.

```
model.save("my_keras_model.h5")
```

بمحتاج نعمل save للمودل  
الذي تم تدريبه  
لنستخدمه في داتا  
تانية بدون ما نرجه  
نعدو؟ كمان مرة.

- In the production program, load the trained model.

```
model = keras.models.load_model("my_keras_model.h5")
```

لما اعمله load وارجع اعمله  
بمحتاج اني اوقف عنده  
ما بعد من اول.

33

### \* Outline

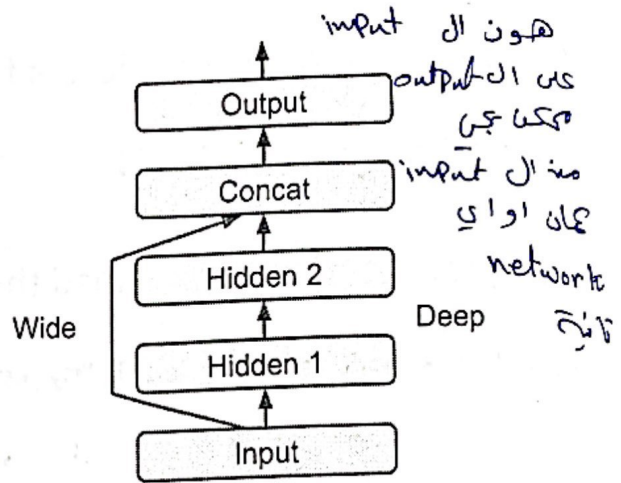
1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

34

في ال  
sequential  
API

## 7. Using the Functional API

- Keras functional API can be used to build arbitrary static graph topologies.
- Create a layer and as soon as it is created, call it like a function, passing it the input.
- Example 1: the wide and deep network that learns both deep patterns (using the deep path) and simple rules (through the short path).

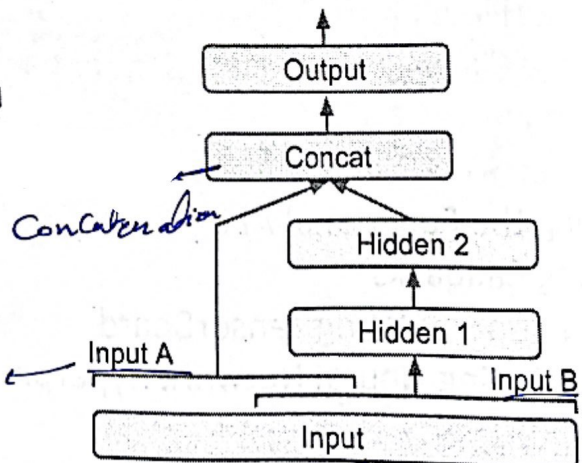


35

functional API يمكن نبي sequential model باستخدام ال

## 7. Using the Functional API

2. Multi-input: You can send a subset of the features through the wide path, and a different subset (possibly overlapping) through the deep path.



(feature) output ال output ال

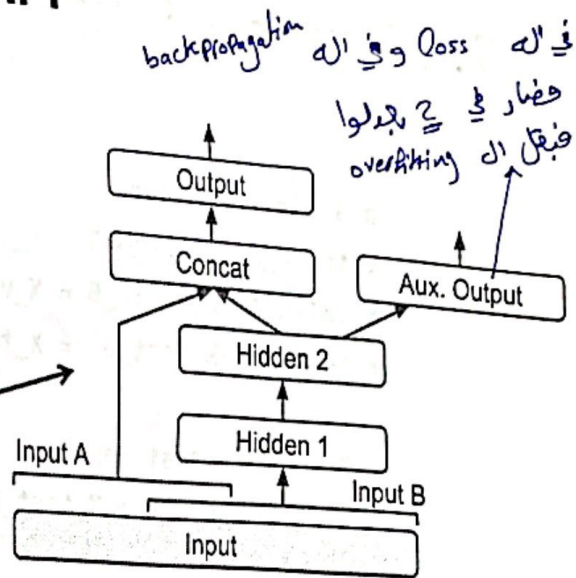
3

# 7. Using the Functional API

## 3. Multiple Outputs

- To locate and classify the main object in a picture.
- Multiple independent tasks to perform based on the same data.
- Regularization technique (to ensure that the deep network learns something useful on its own).

overfitting (by randomness)   
 overfitting



output weight

## 7.1 Auxiliary Output for Regularization

# Build the model

input\_A = keras.layers.Input(shape=[5], name="wide\_input")

input\_B = keras.layers.Input(shape=[6], name="deep\_input")

hidden1 = keras.layers.Dense(30, activation="relu")(input\_B)

hidden2 = keras.layers.Dense(30, activation="relu")(hidden1)

concat = keras.layers.concatenate([input\_A, hidden2]) → Concat the 2 outputs

output = keras.layers.Dense(1, name="main\_output")(concat) → input

aux\_output = keras.layers.Dense(1, name="aux\_output")(hidden2)

```
model = keras.models.Model(inputs=[input_A, input_B],
                           outputs=[output, aux_output])
```

## 7.1 Auxiliary Output for Regularization

```
# Split the input
X_train_A, X_train_B = X_train[:, :5], X_train[:, 2:]
X_valid_A, X_valid_B = X_valid[:, :5], X_valid[:, 2:]
X_test_A, X_test_B = X_test[:, :5], X_test[:, 2:]

# Take some test samples
X_new_A, X_new_B = X_test_A[:3], X_test_B[:3]
```

\* overlap between A & B "2,3,4"

39

## 7.1 Auxiliary Output for Regularization

```
# Compile, train, evaluate, and predict
model.compile(loss=["mse", "mse"], loss_weights=[0.9, 0.1],
              optimizer=keras.optimizers.SGD(lr=1e-3))

history = model.fit([X_train_A, X_train_B], [y_train, y_train], epochs=20,
                  validation_data=([X_valid_A, X_valid_B], [y_valid, y_valid]))

total_loss, main_loss, aux_loss = model.evaluate([X_test_A, X_test_B],
                                                [y_test, y_test])

y_pred_main, y_pred_aux = model.predict([X_new_A, X_new_B])
```

Loss = L1  
Loss = L2  
الحدود  
التقسيم

for output

for aux-output

\* الترتيب

Regularization

output 40%  
aux. " " 10%

40

# Outline

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

41

• شرط يكون احسن نموذج هو النموذج عند ان epoch الاخير.  
• ومن شرط كل ما زدنا عدد ال models يصير النموذج احسن.  
• مرات يكون اسوأ ولما نكمل ان epochs .  
• فبنستخدم ان callbacks

## 8. Using Callbacks

لما نيجي نعمل fit

- The `fit()` method accepts a `callbacks` argument that lets you specify a list of objects that Keras will call during training
  - at the start and end of **training**
  - at the start and end of each **epoch**
  - before and after processing each **batch**
- There are many callbacks available in the `keras.callbacks` package.  
See

<https://keras.io/callbacks/>

42

## 8.1 Saving Best Model

- Save your best model when its performance on the validation set is the best so far.

```

→ checkpoint_cb = keras.callbacks.ModelCheckpoint(
    "my_keras_model.h5", save_best_only=True)
history = model.fit(X_train, y_train, epochs=10,
    validation_data=(X_valid, y_valid),
    callbacks=[checkpoint_cb])
    
```

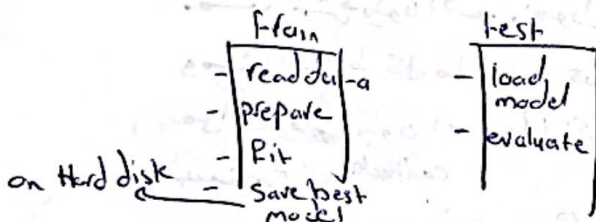
to save the best model "epoch with best performance"

```

# rollback to best model
model = keras.models.load_model("my_keras_model.h5")
mse_test = model.evaluate(X_test, y_test)
    
```

↓  
→ min loss or  
→ highest accuracy

Weights of model load for



→ 2 files (notebooks)

## 8.2 Early Stopping

- Interrupt training when there is no progress on the validation set for a number of epochs (defined by the patience argument)
- Optionally roll back to the best model.

```

early_stopping_cb = keras.callbacks.EarlyStopping(
    patience=10, restore_best_weights=True)
    
```

```

history = model.fit(X_train, y_train, epochs=100,
    validation_data=(X_valid, y_valid),
    callbacks=[checkpoint_cb, early_stopping_cb])
    
```

10 epochs

اذا ال 10 epochs  
best model  
بوقف ال train



# Outline

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

ان Callbacks ای سوکین نستولجا .  
read about them - { 1) Learning Rate Scheduler  
2) TensorBoard  
↓  
بجج Statistics  
عن ان epochs  
ویرججا → GUI


## 9. Visualization Using TensorBoard

- TensorBoard is a great interactive visualization tool that comes with TensorFlow.
- Use it using its callback

```
tensorboard_cb =  
    keras.callbacks.TensorBoard(run_logdir)  
history = model.fit(X_train, y_train, epochs=30,  
    validation_data=(X_valid, y_valid),  
    callbacks=[tensorboard_cb])
```

• Start TensorBoard server → کتان فریج  
\$ tensorboard --logdir=./my\_logs --port=6006  
terminal ای سو فرج

## 9. Open <http://localhost:6006>



47

## Outline

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

48

# 10. Fine-Tuning Neural Network Hyperparameters

لو استخدمنا early-stopping كادى  
لو استخدمنا عدد layers او neurons  
اكثر من اللى بنحتاجه .

- **Number of Hidden Layers**
  - One hidden layer can theoretically model even the most complex functions, provided it has enough neurons.
  - But for complex problems, deep networks have a much higher parameter efficiency than shallow ones.
- **Number of Neurons per Hidden Layer**
  - Pyramid across layers or same size
  - Stretch pants: pick a model with more layers and neurons than you actually need, then use early stopping and other regularization techniques to prevent it from overfitting.
- Better to increase the number of layers instead of the number of neurons per layer.

بنقول بنقل  
عدد ال neurons  
لكل layer  
اللي بنحتاجه  
اللي بنحتاجه  
اللي بنحتاجه

from keras.wrappers.scikit\_learn import //  
Keras classifier  
Keras regressor } → شئ مكافئ ل grid-search  
نقل ال model جوا wrapper بيتر compatible للمودل  
اللي تقابلنا معاه .

او بنقول  
keras tuner

# 10. Fine-Tuning Neural Network Hyperparameters

- **Learning Rate**: the optimal LR is about half of the maximum LR.
- **Optimizer**: There are other than the Mini-batch Gradient Descent optimizer.
- **Batch Size**
  - Larger gives better speed up with hardware accelerators.
  - Smaller makes the models more general.

## • Activation Functions

ال RELU من اكثر  
من funct. ال  
بتعطين نتائج حيدة

بالتة بتكون  
10<sup>-3</sup> / 10<sup>-2</sup>

بالتة بتكون  
convergence

بالتة بتكون  
128 او اقل

عشان ما تستهلك  
memory او لستقل  
ال resources  
عدد ال data وما تقدر تكمل

## 11. Tutorials

- <https://keras.io/>
- <https://www.tensorflow.org/guide/keras>
- Keras Tutorial: Deep Learning in Python from DataCamp, <https://www.datacamp.com/community/tutorials/deep-learning-python>
- Keras Tutorial: The Ultimate Beginner's Guide to Deep Learning in Python, from EliteDataScience, <https://elitedatascience.com/keras-tutorial-deep-learning-in-python>

51

## 12. Exercise

From Chapter 10, solve exercise:

- 10. Train a deep MLP on the MNIST dataset (you can load it using `keras.datasets.mnist.load_data()`). See if you can get over 98% precision. Try searching for the optimal learning rate by using the approach presented in this chapter (i.e., by growing the learning rate exponentially, plotting the error, and finding the point where the error shoots up). Try adding all the bells and whistles—save checkpoints, use **early stopping**, and plot learning curves using **TensorBoard**.

52

# Summary

1. Introduction
2. Keras API Styles
3. TensorFlow Keras
4. Image Classifier Using the Sequential Model
5. Example - MNIST
6. Regression Using the Sequential Model
7. Using the Functional API
8. Using Callbacks
9. Visualization Using TensorBoard
10. Fine-Tuning Neural Network Hyperparameters
11. Tutorials
12. Exercise

*Hands-on github ml2*


53

# Deep Neural Networks

**Prof. Gheith Abandah**

# Reference

- Chapter 11: Training Deep Neural Networks



- Aurélien Géron, **Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow**, O'Reilly, 2nd Edition, 2019
  - Material: <https://github.com/ageron/handson-ml2>

# Outline

1. Introduction
2. Vanishing/Exploding Gradients Problems
  - Glorot and He Initialization
  - Nonsaturating Activation Functions
  - Batch Normalization
  - Gradient Clipping
3. Reusing Pretrained Layers
4. Faster Optimizers → we have used SGD
5. Avoiding Overfitting
  - $\ell_1$  and  $\ell_2$  Regularization
  - Dropout
6. Summary
7. Exercise

تعديلات تحسني

على ال weights  
تعديلات كبيرة

# 1. Introduction

- Deep neural networks can solve complex problems and provide end-to-end solutions.
- When you train a deep network, you may face the following problems:

- ① • Vanishing or exploding gradients: The gradients grow smaller and smaller, or larger and larger.
- ② • Not enough data
- ③ • Long training time
- Overfitting

التقديرات الي  
weight  
بمنحله كل weight  
ممكن قديما توصل للصفر (بين تقديرات اي تقديرات)  
وممكن يصير exploding ويصير في تقديرات كبيرة

ممكن خلاص  
optimizers  
تجعل convergence  
بشكل أسرع

## Outline

1. Introduction
2. Vanishing/Exploding Gradients Problems
  - Glorot and He Initialization
  - Nonsaturating Activation Functions
  - Batch Normalization
  - Gradient Clipping
3. Reusing Pretrained Layers
4. Faster Optimizers
5. Avoiding Overfitting
  - $\ell_1$  and  $\ell_2$  Regularization
  - Dropout
6. Summary
7. Exercise






## 2.1 Glorot and He Initialization

- Recommended initialization parameters for each type of activation function.

Initialization	Activation functions	$\sigma^2$ (Normal)
Glorot	None, Tanh, Logistic, Softmax	$1 / fan_{avg}$
He	ReLU & variants	$2 / fan_{in}$
LeCun	SELU	$1 / fan_{in}$

- For the uniform distribution, use  $r = \sqrt{3\sigma^2}$
- Keras uses **Glorot initialization** with a uniform distribution.



*↳ default*

8

## 2.1 Glorot and He Initialization

- To change it to He initialization:

```
keras.layers.Dense(10, activation="relu",
    kernel_initializer="he_normal") # Or "he_uniform"
```

- He initialization with a uniform distribution but based on  $fan_{avg}$ :

```
he_avg_init = keras.initializers.VarianceScaling(scale=2.,
    mode='fan_avg', distribution='uniform')
```

```
keras.layers.Dense(10, activation="sigmoid",
    kernel_initializer=he_avg_init)
```

9

\* In Keras, no default for activation.

→ check Keras.io

saturation activation function کا نسیبہ کا تعلق ہے۔


## 2.2 Nonsaturating Activation Functions

ReLU Nonsaturating کی زیادہ ان input اور output مابین

• Step does not work with the back propagation algorithm.   
 یہ saturation ہے input کا جب output = 0 ہے

• ReLU is better than sigmoid because it does not saturate for positive values and is fast.

→ • Dying ReLUs: A neuron dies when its input is negative for all training instances.



• saturation on negative values

↓  
• دying ReLU کی وجہ سے neuron کی activation = 0 ہے۔




یہ دying ReLU کی وجہ سے neuron کی دying ہے

## 2.2 Nonsaturating Activation Functions

• Leaky ReLU performs better than ReLU.

$$\text{LeakyReLU}_\alpha(z) = \max(\alpha z, z)$$

•  $\alpha$  between 0.01 and 0.3



small slope

```
model = keras.models.Sequential([
    ...
    keras.layers.Dense(10, kernel_initializer="he_normal"),
    keras.layers.LeakyReLU(alpha=0.2), # added as a layer
    ...
])
```

zero parameter layer


slope

## 2.2 Nonsaturating Activation Functions

\* increase simulation time & it's slower

- Exponential linear unit (ELU) also performs better than ReLU but is slower. → complex derivative
- Scaled ELU (SELU) performs best with dense and CNN, (but must scale inputs and use `lecun_normal`). → training time ↓

$$ELU_{\alpha}(z) = \begin{cases} \alpha(\exp(z) - 1) & \text{if } z < 0 \\ z & \text{if } z \geq 0 \end{cases}$$



```
layer = keras.layers.Dense(10, activation="selu",
kernel_initializer="lecun_normal")
```

## 2.2 Nonsaturating Activation Functions

- Summary:
  - SELU > ELU > leaky ReLU > ReLU > tanh > logistic
- If you cannot use SELU, use ELU.
- For fast response, use leaky ReLU or ReLU.

## 2.3 Batch Normalization

بلقي الاصل  
 Initialization  
 normalization في كل  
 activation layer  
 layer في كل  
 ↓ scale في  
 activation  
 convergence  
 ↓

to solve Vanishing / Exploding

weights  
 ↓  
 parameters  
 ← بعد كل بنية ال

The techniques in §2.1 and §2.2 can significantly reduce the vanishing/exploding gradients problems at the beginning of training, but don't guarantee that they won't come back during training.

- Batch Normalization (BN) zero-centers and normalizes each layer input using statistics from the mini batch (> 30).
- Other benefits: Works even without §2.1 and §2.2, allows using larger LR, and have regularization effect.

std, mean في  
 across all batches (training samples)  
 batch في كل layer  
 layer في كل

## \*2.3 Batch Normalization

7  
 10  
 20  
 $\frac{7+10+20}{3} = 12.5$

• Implementing batch normalization with Keras is easy.

```

model = keras.models.Sequential([
    keras.layers.Flatten(input_shape=[28, 28]),
    → keras.layers.BatchNormalization(),
    keras.layers.Dense(300, activation="elu",
    kernel_initializer="he_normal"),
    → keras.layers.BatchNormalization(),
    keras.layers.Dense(100, activation="elu",
    kernel_initializer="he_normal"),
    → keras.layers.BatchNormalization(),
    keras.layers.Dense(10, activation="softmax")
])
    
```

7 - 12.5 = □ ?  
 10 - 12.5 = □  
 20 - 12.5 = □  
 ← std  
 next layer

1)

ممكن يكون ال weight كاي  
ظاهرة بال exploding

## 2.4 Gradient Clipping

اذا قاطر ال  
max value →

بمسجل ال  
clip value

- Mitigates the exploding gradients problem by **clipping the gradients** during backpropagation so that they never exceed some threshold.
- Use it when you observe that the gradients are exploding during training. You can **track the size of the gradients** using TensorBoard.

```
optimizer = keras.optimizers.SGD(clipvalue=1.0)
```

```
model.compile(loss="mse", optimizer=optimizer)
```

كيف ال  
Gradient  
بغير at run time

16

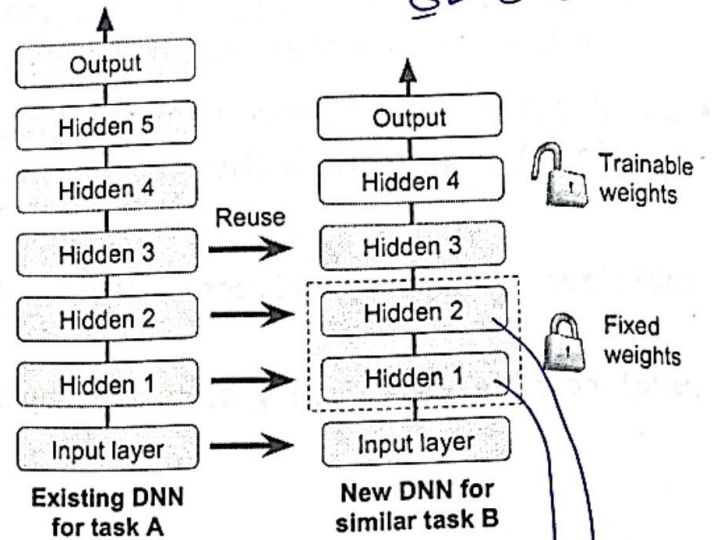
## Outline

1. Introduction
2. Vanishing/Exploding Gradients Problems
  - Glorot and He Initialization
  - Nonsaturating Activation Functions
  - Batch Normalization
  - Gradient Clipping
3. Reusing Pretrained Layers
4. Faster Optimizers
5. Avoiding Overfitting
  - $\ell_1$  and  $\ell_2$  Regularization
  - Dropout
6. Summary
7. Exercise

17

### 3. Reusing Pretrained Layers

- **Transfer Learning:** Using one NN developed for a certain task to solve another task.
- Useful to shorten training time or with small datasets.



كانه النموذج تعلم سابق ونقل معرفته لنموذج ثاني

لا يتعدى ال weights بعدد ال layers

### Transfer Learning with Keras

```
# Load the ready model
model_A = keras.models.load_model("my_model_A.h5")
# Create a new model using all but the last layer
model_B_on_A = keras.models.Sequential(model_A.layers[:-1])
model_B_on_A.add(keras.layers.Dense(1, activation="sigmoid"))
# Freeze loaded layers then compile
for layer in model_B_on_A.layers[:-1]:
    layer.trainable = False
model_B_on_A.compile(loss="binary_crossentropy",
                    optimizer="sgd", metrics=["accuracy"])
```

ال layers ال model A  
output ال layer

new output layer

locked) weights ال على ال layer

# Transfer Learning with Keras

```
# Train the model for a few epochs
history = model_B_on_A.fit(X_train_B, y_train_B, epochs=4,
                           validation_data=(X_valid_B, y_valid_B))
# Unfreeze loaded layers
for layer in model_B_on_A.layers[:-1]:
    layer.trainable = True
# Compile with small learning rate (default = 1e-2)
optimizer = keras.optimizers.SGD(lr=1e-4)
model_B_on_A.compile(loss="binary_crossentropy",
                    optimizer=optimizer, metrics=["accuracy"])
```

في حال ما عينا  
النتيجة  
بفتح الـ unlock  
(بدون طلب  
تدريب)  
عشان يعطينا  
نتائج افضل

# Transfer Learning with Keras

```
# Train the model for more epochs
history = model_B_on_A.fit(X_train_B, y_train_B, epochs=16,
                           validation_data=(X_valid_B, y_valid_B))
```


# Outline

1. Introduction
2. Vanishing/Exploding Gradients Problems
  - Glorot and He Initialization
  - Nonsaturating Activation Functions
  - Batch Normalization
  - Gradient Clipping
3. Reusing Pretrained Layers
4. Faster Optimizers
5. Avoiding Overfitting
  - $\ell_1$  and  $\ell_2$  Regularization
  - Dropout
6. Summary
7. Exercise

22

## 4. Faster Optimizers

- The SGD optimizer can be made faster using momentum optimization



$$\theta \leftarrow \theta - \eta \nabla_{\theta} J(\theta)$$

$$1. \quad m \leftarrow \beta m - \eta \nabla_{\theta} J(\theta)$$

$$2. \quad \theta \leftarrow \theta + m$$

$\beta$

```
optimizer = keras.optimizers.SGD(lr=0.001, momentum=0.9)
```

weight کی بجائے history

23




## 4. Faster Optimizers

حسب ان gradient  
ببدا يتحرك للادناه.

- Nesterov momentum optimization measures the gradient of the cost function not at the local position  $\theta$  but slightly ahead in the direction of the momentum, at  $\theta + \beta m$

momentum coeff.  $\rightarrow$   $\beta m$   $\rightarrow$  average history

1.  $m \leftarrow \beta m - \eta \nabla_{\theta} J(\theta + \beta m)$
2.  $\theta \leftarrow \theta + m$




$\rightarrow$  optimizer = keras.optimizers.SGD(lr=0.001, momentum=0.9, nesterov=True)

## 4. Faster Optimizers

بتنقص على ان history حمان

- The adaptive optimizers such as AdaGrad, RMSProp, Adam, and Nadam scale down the gradient vector along the steepest dimensions.



$\rightarrow$  optimizer = keras.optimizers.RMSprop()  
 $\rightarrow$  optimizer = keras.optimizers.Adam()

## 4. Faster Optimizers

- RMSProp, Adam and Nadam often **converge fast**. But they can give poor **generalization**.
- Solution: Use Nesterov accelerated gradient.

Class	Speed	Quality
SGD	*	***
→ SGD with momentum, Nesterov	**	***
Adagrad	***	*
RMSProp, Adam, Nadam, AdaMax	***	** or ***

26

## Outline

1. Introduction
2. Vanishing/Exploding Gradients Problems
  - Glorot and He Initialization
  - Nonsaturating Activation Functions
  - Batch Normalization
  - Gradient Clipping
3. Reusing Pretrained Layers
4. Faster Optimizers
5. Avoiding Overfitting
  - $\ell_1$  and  $\ell_2$  Regularization
  - Dropout
6. Summary
7. Exercise

27

## 5. Avoiding Overfitting

ليس لانه ينحط  
 عدد epochs  
 عكس وما ينحرف  
 كم عدد لهم الخنا

• Deep neural networks typically have many parameters, giving them ability to fit a huge variety of complex datasets.

Useful regularization techniques: → to avoid overfitting →

- Early stopping → call back
- Batch normalization
- $\ell_1$  and  $\ell_2$  regularization
- Dropout

المودل ما يكون دقيق  
 لعدد samples معينة  
 ليدنا يكون general

تفريغ كالتالي loss

توقف stop بعد  
 epochs معينة وتقبل  
 best model لا restore

28

### 5.1 $\ell_1$ and $\ell_2$ Regularization

• Constrain a neural network's connection weights. → regularization factor

- $\ell_1$ :
- $\ell_2$ :

$$\text{Cost function} = \text{Loss} + \frac{\lambda}{2m} \cdot \sum \|w\|$$

$$\text{Cost function} = \text{Loss} + \frac{\lambda}{2m} \cdot \sum \|w\|^2$$

التعديل هون يكون  
 اقل لانه يتعامل  
 مع ال weights تقريبا  
 بالتالي ال features المهمة  
 ال weight بتعطي يكون الجين

```
layer = keras.layers.Dense(100, activation="elu",
    kernel_initializer="he_normal",
```

```
kernel_regularizer=keras.regularizers.l1(0.01))
```

# The other regularization functions:

- ② `keras.regularizers.l2(0.01)`
- ③ `keras.regularizers.l1_l2(l1=0.01, l2=0.01)`

0.01 < λ < 0.05

29

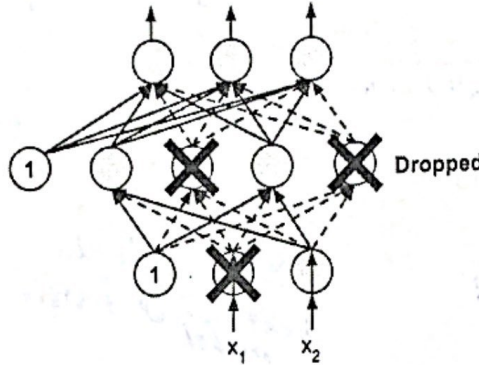
التدريب  
 نغني بعض neurons وما يدخلوا  
 Forward Pass  
 ما يتعدى عدد neurons  
 ما يكون في neuron واحد مستخدم عند التدريب  
 ما يتعدى عدد neurons  
 ما يكون في neuron واحد مستخدم عند التدريب

## 5.2 Dropout

\* better generalization

- Popular technique to improve accuracy.
- At every training step, every neuron (excluding the output neurons) has a probability  $p$  of being temporarily dropped out.

المعدل الشافعي ما  
 يفعل فيه Dropout  
 النسبة العامة 0.2 يعني  
 iteration 20% بكل iteration



في حال صار عندي مشكلة ب neuron معين  
 مثل dying neuron ما بتأش التود لي كامل وهاي  
 برهني طريقة لتجنب ال overfitting

يفضل ان سكر layer منفصلة

## 5.2 Dropout

استخدم بعد ال activation  
 وقبل ال Batch Normalization

```
model = keras.models.Sequential([
    keras.layers.Flatten(input_shape=[28, 28]),
    → keras.layers.Dropout(rate=0.2),
    keras.layers.Dense(300, activation="elu",
        kernel_initializer="he_normal"),
    → keras.layers.Dropout(rate=0.2),
    keras.layers.Dense(100, activation="elu",
        kernel_initializer="he_normal"),
    → keras.layers.Dropout(rate=0.2),
    keras.layers.Dense(10, activation="softmax")
])
```

20% من  
300

بقدر ازيد نسبة  
 بس يفضل ما تعلق

←  
in train

1)

لا عمل test يفضل اضرب ال weights كل  
 ب (1-rate) ← keep ratio  
 ربه ما عندي drop لا عمل test

# Outline

1. Introduction
2. Vanishing/Exploding Gradients Problems
  - Glorot and He Initialization
  - Nonsaturating Activation Functions
  - Batch Normalization
  - Gradient Clipping
3. Reusing Pretrained Layers
4. Faster Optimizers
5. Avoiding Overfitting
  - $\ell_1$  and  $\ell_2$  Regularization
  - Dropout
6. Summary
7. Exercise

32

## 6. Summary

- Recommended default DNN configuration

*Epochs can be divided into steps*

Hyperparameter	Default value
Kernel initializer	He initialization
Activation function	ELU <i>or leaky</i>
Normalization	None if shallow; Batch Norm if deep
Regularization	Early stopping (+ $\ell_2$ reg. if needed)
Optimizer	<u>Momentum optimization</u> (or RMSProp or Nadam)
Learning rate schedule	1 cycle

*↳ in Callbacks*

33

## 6. Summary

- For a simple stack of dense or CNN layers.

Hyperparameter	Default value
Kernel initializer	LeCun initialization
Activation function	SELU
Normalization	None (self-normalization)
Regularization	Alpha dropout if needed
Optimizer	Momentum optimization (or RMSProp or Nadam)
Learning rate schedule	1 cycle

34

## 7. Exercise

- 11.8. Practice training a deep neural network on the CIFAR10 image dataset:
- Build a DNN with 20 hidden layers of 100 neurons each (that's too many, but it's the point of this exercise). Use He initialization and the ELU activation function.
  - Using Nadam optimization and early stopping, train the network on the CIFAR10 dataset. You can load it with `keras.datasets.cifar10.load_data()`. The dataset is composed of 60,000 32 × 32-pixel color images (50,000 for training, 10,000 for testing) with 10 classes, so you'll need a softmax output layer with 10 neurons. Remember to search for the right learning rate each time you change the model's architecture or hyperparameters.
  - Now try adding Batch Normalization and compare the learning curves: Is it converging faster than before? Does it produce a better model? How does it affect training speed?
  - Try replacing Batch Normalization with SELU, and make the necessary adjustments to ensure the network self-normalizes (i.e., standardize the input features, use LeCun normal initialization, make sure the DNN contains only a sequence of dense layers, etc.).
  - Try regularizing the model with alpha dropout. Then, without retraining your model, see if you can achieve better accuracy using MC Dropout.
  - Retrain your model using 1cycle scheduling and see if it improves training speed and model accuracy.

35

# Deep Computer Vision Using Convolutional Neural Networks

CNN

الهدف منها نقل extract features

Prof. Gheith Abandah

تستعمل  
بشكل  
تتم تحويل الصور الى  
vector


matching  
بينه وبين الصور

بنقل عددا features الكيس بعد اقل  
neural network  
مهمه بناخذها الى  
representing للصور الكيس

## Reference

- Chapter 14: Deep Computer Vision Using Convolutional Neural Networks

O'REILLY  
Hands-On  
Machine Learning  
with Scikit-Learn,  
Keras & TensorFlow  
Concepts, Tools, and Techniques  
to Build Intelligent Systems




Aurélien G

- Aurélien Géron, Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow, O'Reilly, 2nd Edition, 2019

- Material: <https://github.com/ageron/handson-ml2>

## Reference



- **Deep Learning with Python**, by François Chollet, Manning Pub 2018

3

## Outline

1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

4



# Introduction


- YouTube Video: Convolutional Neural Networks (CNNs) explained from Deeplizard

[https://youtu.be/YRhxdVk\\_sls](https://youtu.be/YRhxdVk_sls)

## 1. Introduction

كيف عرفنا انه الي بالصورة عبارة عن عذبة ؟ هاد الاشي مسؤول عنه ال CNN  
لانه في مثلك ومرجع ومستطيل ضمن الشكل

- Convolutional neural networks (CNNs) emerged from the study of the brain's visual cortex.
- Many neurons in the visual cortex have a small local receptive field.



ال CNN بتفسي على الصورة بأول من اجل بيتوف  
ال edges عن طريق الفروقات و بأفتر حلة

كولعال Vector

هو يمين بس اذا في بيت او لا بس ما بقدر عدد  
وين او موقته بالصورة

# Outline


1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

## 2. Convolutional Layer


ان neurons ما يكونوا بالتركيبة نفسها  
 dense network

Kernel  
 Filters

- **Neurons** in one layer are not connected to every single pixel/neuron in the previous layer, but only to pixels/neurons in their **receptive fields**.
- This architecture allows the network to concentrate on **low-level features** in one layer, then assemble them into **higher-level features** in the next layer.
- Each layer is represented in 2D.



لما تنتقل من layer ل layer الثانية  
 بنظرن ال features




صورة واول مربع مشينا  
 ال filter هي اول  
 Pixels 4  
 بنظرن على نقطة بالفلتر  
 بالقيمة اى يتقابلها وينجمه

النتائج برزت  
 على ال next layer

بنمستى  
 الظلن على باقى  
 Pixels ال

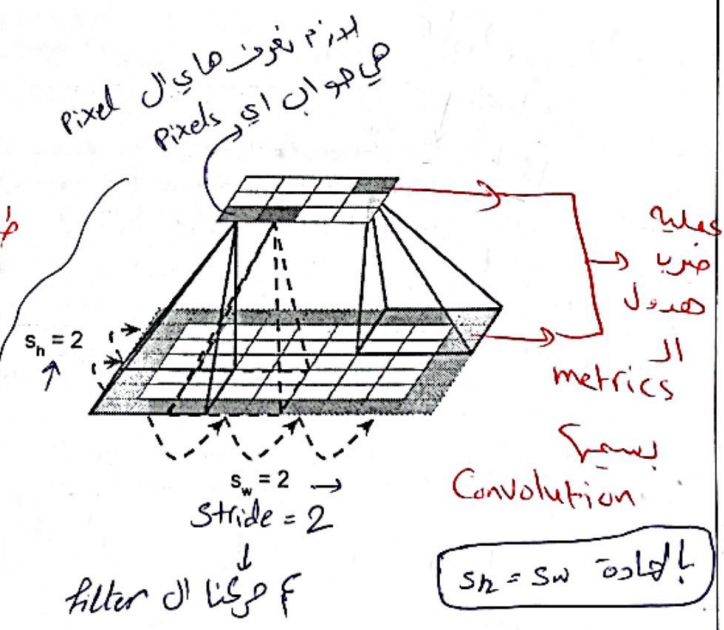
## 2. Convolutional Layer

- $f_h$  and  $f_w$  are the height and width of the receptive field.
- **Zero padding:** In order for a layer to have the same height and width as the previous layer, it is common to add zeros around the inputs.



## 2. Convolutional Layer

- It is also possible to connect a large input layer to a smaller layer by spacing out the receptive fields.
- The distance between two consecutive receptive fields is called the **stride**.
- A neuron located in row  $i$ , column  $j$  is connected to the neurons in the previous layer located in:
  - Rows:  $i \times s_h$  to  $i \times s_h + f_h - 1$
  - Cols:  $j \times s_w$  to  $j \times s_w + f_w - 1$



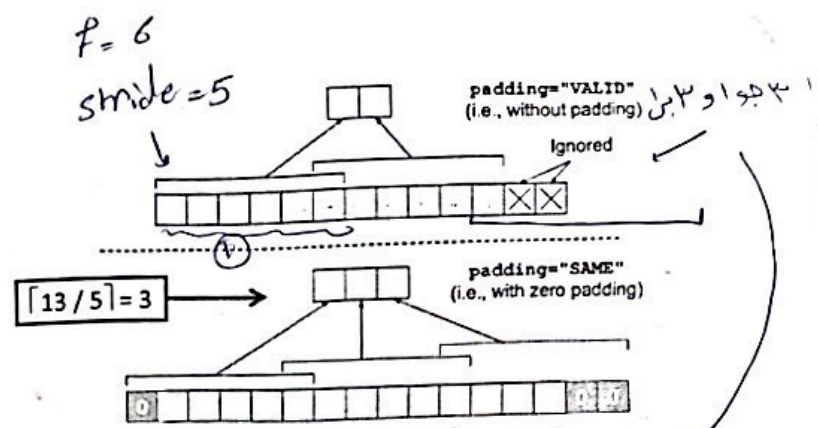
horizontal & vertical.

لأنه إذا كان stride يساوي 2 يعني في كل مرة ننتقل بحجم 2 في كل اتجاه (أفقي وعمودي).  
 إذا كان stride يساوي 1 يعني ننتقل بحجم 1 في كل اتجاه (أفقي وعمودي).  
 low level details

لا يتبدل أبعاد الخرج من Filter بل هي نفس الصورة

## 2. Convolutional Layer

- Keras supports
  - No padding (default) padding="VALID"
  - Zero padding padding="SAME"
- Example:
  - Input width: 13
  - Filter width: 6
  - Stride: 5



لو كان ال filter ما يطبق على pixels كل من الصورة فلهذا ما يكمل (ما يجتمع الخطوة) وعمل بالاتجاه الثاني بسر ريس في معلومات ما غطيناها.   
 if no padding → لو كان ال filter ما يطبق على pixels كل من الصورة فلهذا ما يكمل (ما يجتمع الخطوة) وعمل بالاتجاه الثاني بسر ريس في معلومات ما غطيناها.

$$\text{Padding} = \lceil \frac{\text{Pixels per row}}{\text{stride}} \rceil$$

والأفضل نوزعها ال padding عاشرمان و عالصيرة

if padding = same → ال pixels الزيادة اعتبرهم اصغار، بالتالي بنشمل كل الصورة.

## 2. Convolutional Layer

هاي الصورة الاصلية 5x5  
 بينا غنشي عليها filter 3x3  
 stride=1  
 f=3x3  
 رطردول بنشغل خطوة لخصه  
 لو ال padding مسوي 2 نعمل كمان خطوة تانية  
 Vertically


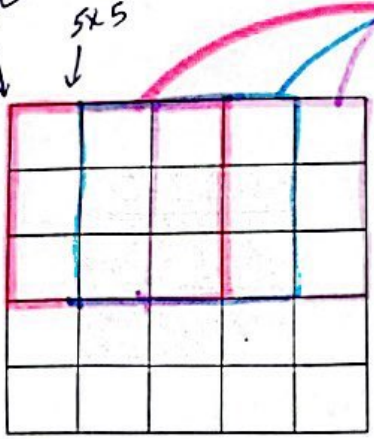


Figure 5.5 Valid locations of 3 x 3 patches in a 5 x 5 input feature map

لو ال padding = same ال 5x5 = (Feature map) Size of feature map before convolution = 11 after stride=1

ال Filter يكون ال weights ال initialization تبع ال Filter  
 وكما بي ارجع امدوح امدوح ال initialization تبع ال Filter


## 2.1 Filters

ان filter هو نفس ال kernel

- A neuron's weights can be represented as a small image the size of the receptive field, called **filters**.

- When all neurons in a layer use the same line filters, we get the **feature maps** on the top.

every filter gives a Feature map.



0	1	0
0	1	0
0	1	0

0	0	0
1	1	1
0	0	0

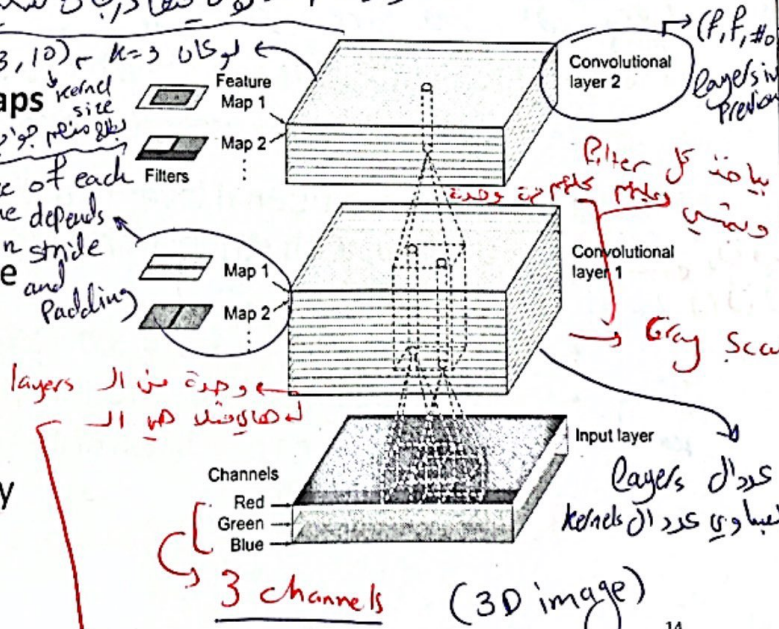
بس 0, 1 لانه عننا ابيضه و اسود الالوان

مكن تكون 1, 0, 1 - ال ابيض  
 ابيض اسود

## 2.2 Stacking Feature Maps

الصور ممكن تكون gray scale - بتكون 1 channel  
 ان Pixels فيها من 0 ل 255 / الالوان فيها درجات السكيني والابيض والاسود

- In reality, each layer is 3D composed of several feature maps of equal sizes.
- Within one feature map, all neurons share the same parameters, but different feature maps may have different parameters.
- Once the CNN has learned to recognize a pattern in one location, it can recognize it in any other location.



OpenCV -> library to apply image processing techniques.

Processing techniques.

مينا: تحويل الصورة من color coding  
 ل color coding تاني

Feature map = From Filter =



# Outline

1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

اذا استعملت strides و padding حذرت من Size الصورة

بعد convolution نضيق

فبنتقل ان pixels الي يكون قريب من بعضه



## 3. Pooling Layer ما ال Parameters

Conv. ل تطبيقها بعد ان

- Its goal is to subsample (i.e., shrink) the input image in order to reduce the computational load, the memory usage, and the number of parameters.
- It aggregates the inputs using max or mean.

تقليل ال size ال Feature maps

Pixel with most brightness



نتيجتي على 4 بتاخذ ال max او min

بتروح على ال 4 الي بعد

Pooling mean

Overlap ما في

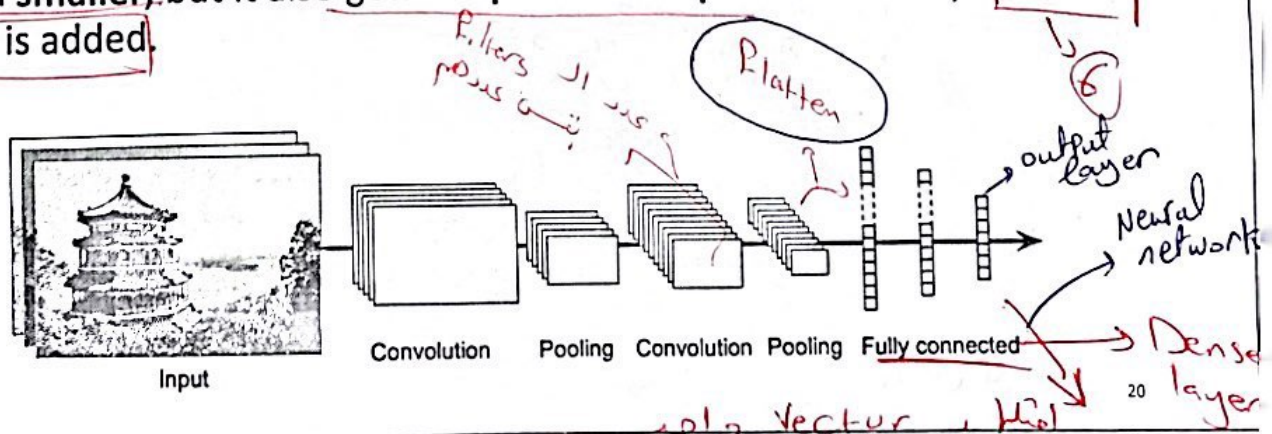
# Outline

1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

## 4. CNN Architectures

اختراع ال features غي كاعني

• Stack few convolutional layers (each one generally followed by a ReLU layer), then a pooling layer, then another few convolutional layers, then another pooling layer, and so on. The image gets smaller and smaller, but it also gets deeper and deeper. At the end, a dense NN is added.



CNN not stand alone, it's almost feature extraction



# 4.1 Example – Fashion MNIST

```

model = keras.models.Sequential([
    keras.layers.Conv2D(64, 7, activation="relu", padding="same",
        input_shape=[28, 28, 1]),
    keras.layers.MaxPooling2D(2),
    keras.layers.Conv2D(128, 3, activation="relu", padding="same"),
    keras.layers.Conv2D(128, 3, activation="relu", padding="same"),
    keras.layers.MaxPooling2D(2),
    keras.layers.Conv2D(256, 3, activation="relu", padding="same"),
    keras.layers.Conv2D(256, 3, activation="relu", padding="same"),
    keras.layers.MaxPooling2D(2),
    keras.layers.Flatten(),
    keras.layers.Dense(128, activation="relu"),
    keras.layers.Dropout(0.5),
    keras.layers.Dense(64, activation="relu"),
    keras.layers.Dropout(0.5),
    keras.layers.Dense(10, activation="softmax")
])
    
```

*Handwritten notes:*

- Filters عدد ال** (Filters count)
- kernel size (7,7,1)** → **Filter size = 7 x 7 x 1** (for an image)
- output size = 28 x 28 x 64**
- Feature maps** (2x2)
- 3 x 3 x 64** (Previous Layer)
- 14 x 14 x 28**
- 3 x 3 x 128**
- 2x2 window and stride 2**
- عدد ال Filters و بنظر ال** (Filters count and perspective)
- Size للصورة** (Image size)
- 2D** (2D)
- mnist 2 channel**
- Out size = 14 x 14 x 64**
- 10 classes** (10 classes)
- 10 classes** (10 classes)

# 4.1 Example – Fashion MNIST

```

model.compile(loss="sparse_categorical_crossentropy",
              optimizer="nadam", metrics=["accuracy"])
    
```

```

history = model.fit(X_train, y_train, epochs=10,
                   validation_data=(X_valid, y_valid))
    
```

Train on 55000 samples, validate on 5000 samples

Epoch 1/10 55000/55000 [=====] - 51s 923us/sample - loss: 0.7183 - accuracy: 0.7529 - val\_loss: 0.4029 - val\_accuracy: 0.8510

... Epoch 10/10

55000/55000 [=====] - 50s 911us/sample - loss: 0.2561 - accuracy: 0.9145 - val\_loss: 0.2891 - val\_accuracy: 0.9036

*لتعدد ال Labels ال*

## 4.1 Example – Fashion MNIST

```
→ score = model.evaluate(X_test, y_test)
X_new = X_test[:10] # pretend we have new images
y_pred = model.predict(X_new)
```

```
10000/10000 [=====] - 2s 239us/sample - loss:
0.2972 - accuracy: 0.8983
```

Can reach 92% with more epochs

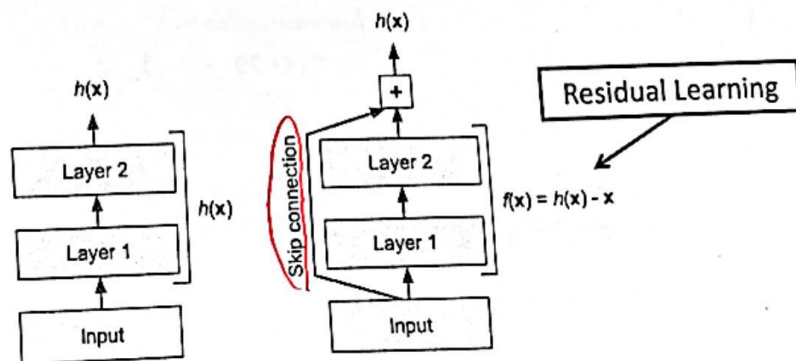
23

## 4.2 ResNet

شبكة network

- Residual Network (or ResNet) won the ILSVRC 2015 challenge.
- Top-5 error rate under 3.6%, using an extremely deep CNN composed of **152 layers**.
- To train such a deep network, it uses **skip connections**.

اسى 5 Pred.  
و ستون  
من الـ 100  
اسى 5 Pred.

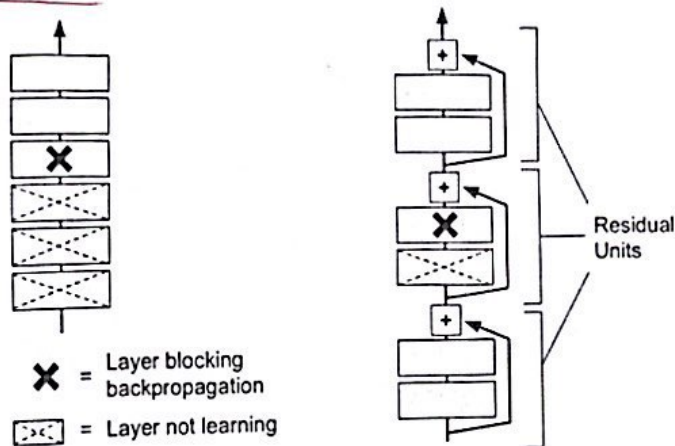


24

## 4.2 ResNet

ما استعملنا ان sequential  
 عشان لو بتا تعمل layer update  
 لا ان من كذا ال قبلها  
 فاستعملنا skip connections  
 backward propagation

- The network can start making progress even if several layers have not started learning yet.



25

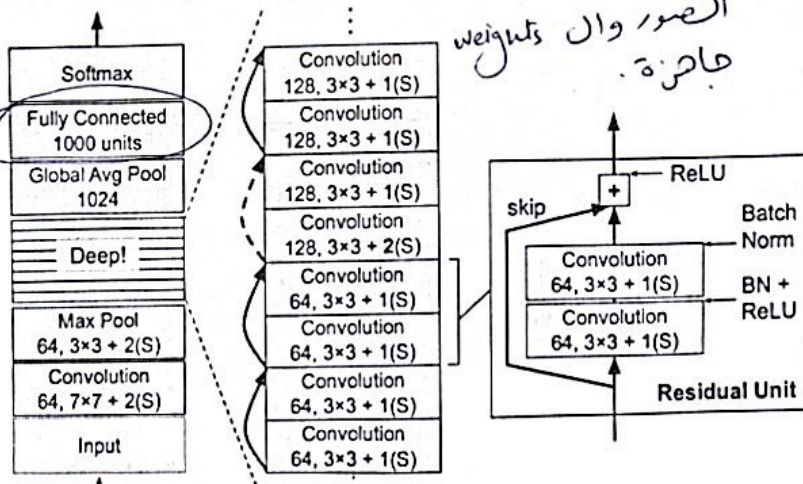
## 4.2 ResNet

موجود  
 kelas model?

- ResNet is a stack of residual units.

model جاهز  
 مدرسه دي اعرف  
 الصور وان weights جاهزة.

ال output layer  
 لو كان به تاكتيك  
 بعد كذا  
 اقل منه ال 1000  
 او اشي ما تدرب  
 عليه 8, 16  
 نفس ال output layer



26

او ممكن تدرب  
 ال hidden layers  
 من ال اول

## Outline

1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

27

## 5. Using Pretrained Models

- Pretrained networks are readily available from the `keras.applications` package.
- Check <https://github.com/keras-team/keras-applications>
- You can load the **ResNet-50** model, pretrained on **ImageNet**, with the following line of code:

```
model = keras.applications.resnet50.ResNet50(weights="imagenet")
```

حمار  
عن model.  
جائز

اعل load مع ال weights  
اي وعلها كما تدرج على ال صور

28

## 5. Using Pretrained Models

```
# Input: 224 x 224-pixel images  
images_resized = tf.image.resize(images, [224, 224])
```

بعل مع resize للصورة  
← حساب ابعاد model  
تسأل Input shape  
التي بقبله

```
# Preprocess images, should be scaled 0-255  
→ inputs = keras.applications.resnet50.preprocess_input(  
    images_resized * 255)
```

```
→ Y_proba = model.predict(inputs)
```

```
# Get top predictions out of the 1000-class probs.
```


```
→ top_K = keras.applications.resnet50.decode_predictions(Y_proba, top=3)
```

## 5. Using Pretrained Models

```
# Print results  
for image_index in range(len(images)):  
    print("Image #{}".format(image_index))  
    for class_id, name, y_proba in top_K[image_index]:  
        print(" {} - {:12s} {:.2f}%".format(class_id, name, y_proba * 100))  
    print()
```

```
Image #0  
n03877845 - palace 42.87%  
n02825657 - bell_cote 40.57%  
n03781244 - monastery 14.56%
```

```
Image #1  
n04522168 - vase 46.83%  
n07930864 - cup 7.78%  
n11939491 - daisy 4.87%
```



# Outline

1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

## 6. Pretrained Models for Transfer Learning

- Training a pretrained network (Xception) for a dataset from TFDS (<https://www.tensorflow.org/datasets>).

- `tf_flowers`: 3670 images, 5 classes

# Load the dataset

```
import tensorflow_datasets as tfds
```

```
dataset, info = tfds.load("tf_flowers",  
                          as_supervised=True, with_info=True)
```

```
dataset_size = info.splits["train"].num_examples # 3670
```

```
n_classes = info.features["label"].num_classes # 5
```

```
class_names = info.features["label"].names
```

Class: roses



## 6. Pretrained Models for Transfer Learning

---

```
# Reload the dataset with three splits tf.data.Dataset
test_set_raw, valid_set_raw, train_set_raw = tfds.load(
    "tf_flowers", split=["train[:10%]",
        "train[10%:25%]", "train[25%:]"],
    as_supervised=True)

# Define the preprocessing function
def preprocess(image, label):
    resized_image = tf.image.resize(image, [224, 224])
    final_image = keras.applications.xception.preprocess_input(
        resized_image)
    return final_image, label
```

33

---

## 6. Pretrained Models for Transfer Learning

```
# Apply this preprocessing function to the 3 datasets
# Shuffle the training set
# Add batching and prefetching to all the datasets
batch_size = 32 → 32000
train_set = train_set_raw.shuffle(3000).repeat()
train_set = train_set.map(preprocess).batch(
    batch_size).prefetch(1)
valid_set = valid_set_raw.map(preprocess).batch(
    batch_size).prefetch(1)
test_set = test_set_raw.map(preprocess).batch(
    batch_size).prefetch(1)
```

## 6. Pretrained Models for Transfer Learning

```
# Load an Xception model, pretrained on ImageNet  
# excluding the global avg pool. and dense o/p layers
```

```
base_model = keras.applications.xception.Xception(  
    weights="imagenet", include_top=False)
```

*model.summary()*

```
# Add global avg pool. layer based on model output
```

```
avg = keras.layers.GlobalAveragePooling2D()(base_model.output)
```

```
output = keras.layers.Dense(n_classes, # Add dense o/p  
    activation="softmax")(avg)
```

```
model = keras.models.Model(inputs=base_model.input,  
    outputs=output) # Create the Keras Model
```

35

## 6. Pretrained Models for Transfer Learning

```
# Freeze the weights of the pretrained layers
```

```
for layer in base_model.layers:
```

```
    layer.trainable = False
```

*ماتریک  
تدریس و عملیاتی نفس ال  
weights*

```
# Compile the model and start training
```

```
→ optimizer = keras.optimizers.SGD(lr=0.2, momentum=0.9,  
    decay=0.01) # LR=0.2 with schedule, k=1/0.01
```

```
model.compile(loss="sparse_categorical_crossentropy",  
    optimizer=optimizer, metrics=["accuracy"])
```

```
history = model.fit(train_set, epochs=5,
```

```
    validation_data=valid_set) # Tops at 75-80% acc.
```

$$\eta(t) = \eta_0 / (1 + t/k)$$

36



## 6. Pretrained Models for Transfer Learning

# Unfreeze the weights of the pretrained layers  
for layer in base\_model.layers:

```
layer.trainable = True → unfreeze
```

• در بنام اول و بجزیره

# Recompile with lower LR and decay

```
optimizer = keras.optimizers.SGD(lr=0.01, momentum=0.9,  
nesterov=True, decay=0.001)
```

```
model.compile(loss="sparse_categorical_crossentropy",  
optimizer=optimizer, metrics=["accuracy"])
```

```
history = model.fit(train_set, epochs=40,  
validation_data=valid_set) # Result: 95% acc.
```

37

## Outline

1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps *CNN*
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

38

اذا كان عندي مثلا ورقة واحدة في الصورة


## 7. Classification and Localization

يقارن ان  
ال  
Pre. Box  
Ground  
truth


Classification

تحدد مكان ال  
object

- Localizing an object in a picture can be expressed as a regression task.
- Predict the horizontal and vertical coordinates of the object's center and its height and width.



bounding  
box.



metric to determine if the bounding box is true

Common metric: the Intersection over Union (IoU)

→ If overlapped

↑ IoU better

IoU =  $\frac{\text{Intersection}}{\text{Union}}$

39

2D bounding  
box  
bounding  
box

3 metrics.  
3 losses

## 7. Classification and Localization

network with softmax classification & regression → 2 out-puts

```
base_model = keras.applications.xception.Xception(
    weights="imagenet", include_top=False)
avg = keras.layers.GlobalAveragePooling2D()(base_model.output)
class_output = keras.layers.Dense(n_classes, activation="softmax")(avg)
loc_output = keras.layers.Dense(4)(avg)
model = keras.Model(inputs=base_model.input,
    outputs=[class_output, loc_output])

model.compile(loss=["sparse_categorical_crossentropy", "mse"],
    loss_weights=[0.8, 0.2],
    optimizer=optimizer, metrics=["accuracy"])
```

→ 4 neurons  
بجسوا احداثيات ال  
bounding box

40

# outline

1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

41


## 8. Object detection

كل object شو هو وورين موجود

لا يكون عندي اكثر من object بنفس الصورة

- The task of classifying and localizing multiple objects in an image.
- A **slow** approach is use a CNN trained to classify and locate a single object, then **slide** it across the image.

Confidence ← هاد كذا عريف

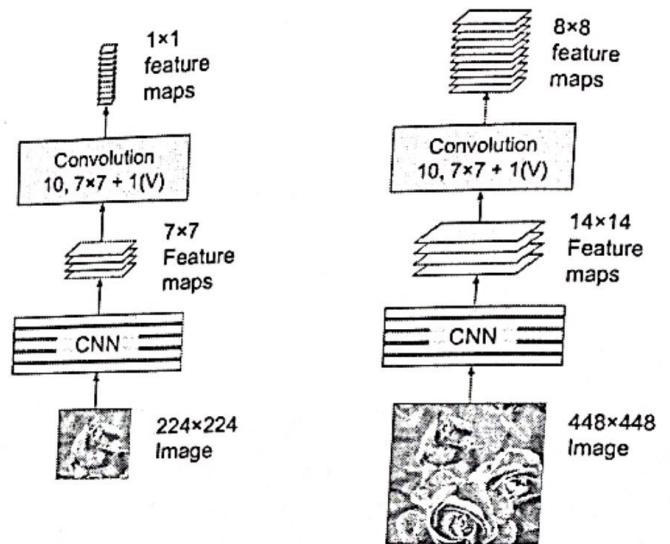


\* بطايعي اكثر من  
banding box  
بضاي اللي اله  
Confidence اعلى  
اللي يكونا نظمه اعرض

42

# 8.1 Fully Convolutional Networks

- FCN has also a convolution layer at the output with valid padding. *map ال صوليس ال zeros*
- FCN can process images of any size.
- Example:
  - Train the CNN for classification and localization on small images, 10 outputs.
  - For larger image, it output  $8 \times 8$  grid where each cell contains 10 numbers.




*هذه ال models ال object detection*

*there are eight versions from Yolo  
↳ you only look once.*


# 8.2 You Only Look Once (YOLO)

- YOLO is an extremely fast and accurate object detection architecture.
  1. Resizes the input image to  $448 \times 448$
  2. Runs a single convolutional network on the image
  3. Thresholds the resulting detections by the model's confidence.



# 8.2 You Only Look Once (YOLO)

- Models detection as a regression problem. It divides the image into an  $S \times S$  grid.
- For each grid cell predicts  $B$  bounding boxes, confidence for those boxes, and  $C$  class probabilities.



sliding window  
بیشتر کاره  
تقسیم تصویر به جابجاء  
و شیوه تقاطع هر کس به

## Outline


1. Introduction
2. Convolutional layer
  1. Filters
  2. Stacking feature maps
  3. Mathematical summary
  4. Memory requirements
3. Pooling layer
4. CNN architectures
  1. Example – Fashion MNIST
  2. ResNet
5. Using pretrained models
6. Pretrained models for transfer learning
7. Classification and localization
8. Object detection
9. Semantic segmentation
10. Exercises

## 9. Semantic Segmentation

كل الصورة لشيء  
بجمع

قابلية تصنيف الصورة حسب فئة

- Each pixel is classified according to the class of the object it belongs to.
- Can use FCN followed by up sampling layers.



supervised learning

- instant segmentation → object

more complex

## Exercises

- 14.9. Build your own CNN from scratch and try to achieve the highest possible accuracy on MNIST.
- 14.10. Use transfer learning for large image classification, going through these steps:
  - a) Create a training set containing at least 100 images per class. For example, you could classify your own pictures based on the location (beach, mountain, city, etc.), or alternatively you can use an existing dataset (e.g., from TensorFlow Datasets).
  - b) Split it into a training set, a validation set, and a test set.
  - c) Build the input pipeline, including the appropriate preprocessing operations, and optionally add data augmentation.
  - d) Fine-tune a pretrained model on this dataset.

Good  
Luck!

