ENGINEERING & AUTOMOTIVE

MACHINE LEARNING INTRO

JEROEN VEEN

HAN_UNIVERSITY OF APPLIED SCIENCES

MINOR EMBEDDED VISION DESIGN

- Optimize and interpret camera images
- Segmentation, extraction, classification
- Embedded systems
- Machine learning and deep learning

https://www.minoren-han.nl/nl/222-embedded-vision-design-full-time
https://www.minoren-han.nl/nl/473-mobile-robotics-full-time
https://www.minoren-han.nl/nl/338-data-science

(bio)informatics

HAN_UNIVERSITY OF APPLIED SCIENCES

CONTENTS

- **Introduction**
- Machine learning applied to computer vision (CV)
- General approaches to machine learning
  - Pitfalls
- Example project
  - Data exploration
  - Training
  - Testing
  - Deployment
- Deep learning?
- Discussion

WHAT IS MACHINE LEARNING?

- Human vs machine learning?
- Machines can perform predictive analytics on large amounts of data far faster than humans
- Machines maximize performance on a certain task
  Typically function approximations
- Learning does not imply intelligence
  if a machine can learn it is not necessarily aware

DEFINING AI, DL & ML

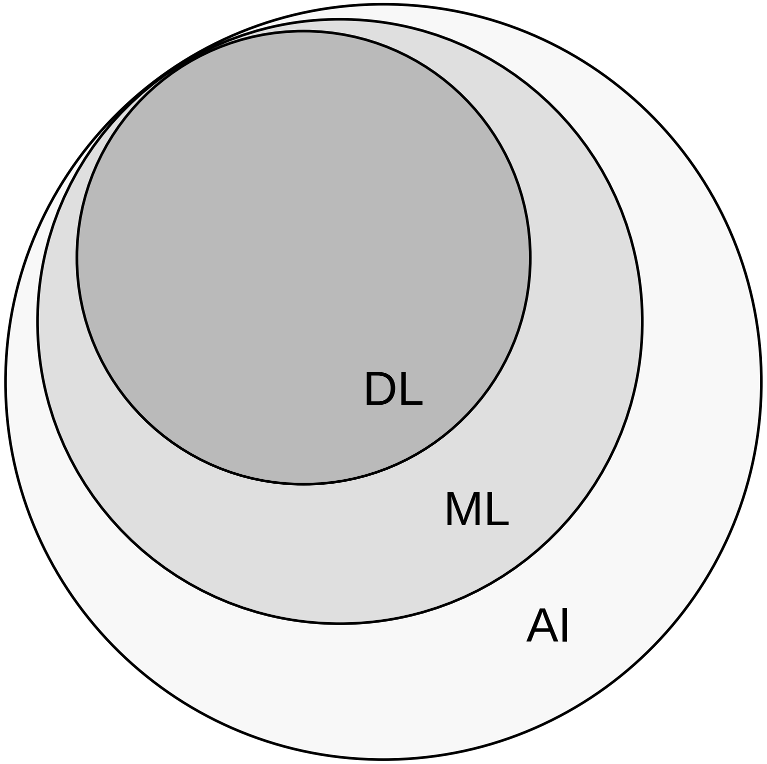

- Strong AI vs Applied AI
- Cognitive replication
- Rational process

Machine learning

- Performs predictive analysis
- Just fancy math & pattern matching

MACHINE LEARNING VS DEEP LEARNING

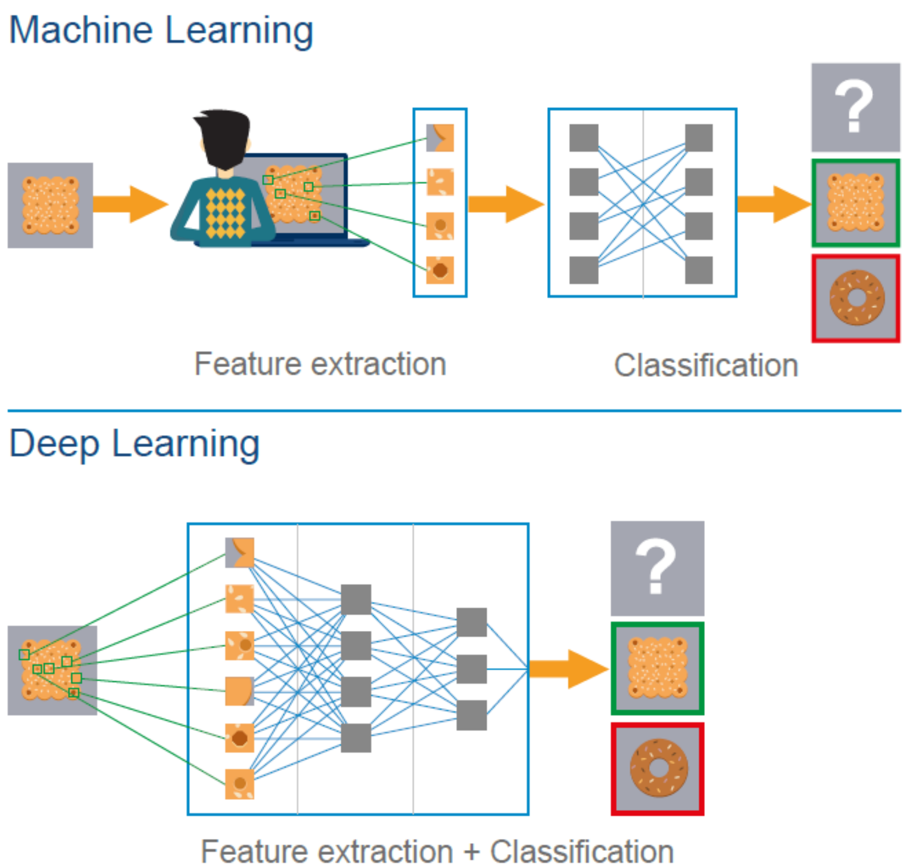

Source: Basler, Artificial Intelligence in Image Processing

WHY MACHINE LEARNING?

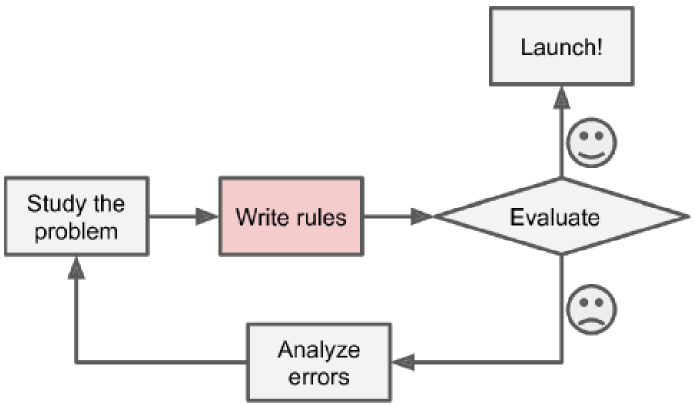

Traditional approach

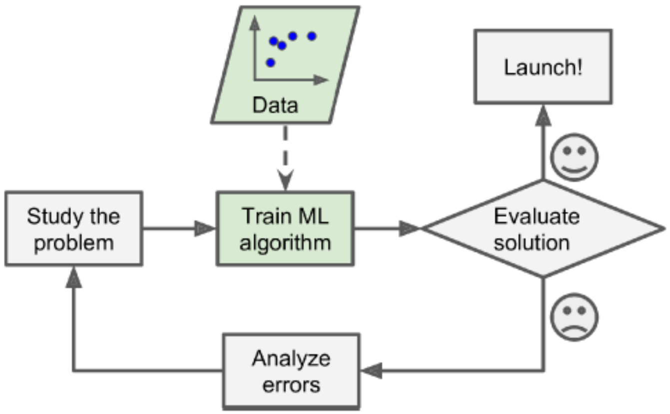

ML approach

Source: Géron, ISBN: 9781492032632

- Tackle problems for which existing solutions require a lot of fine-tuning or long lists of rules
- Deal with fluctuating environments by adapting to new data.
- Getting insights about complex problems and large amounts of data.

ETHICS

- Self-adjustment can go horribly wrong
- Think of ‘sampling bias’, ‘exclusion bias’ and ‘prejudice bias’
- Context matters
- Transparency is becoming important
  General Data Protection Regulation (GDPR)
  -> Explainable AI
- **It is vital that developers take responsibility!**

Uber drivers to launch legal bid to uncover app's algorithm

Union wants ride-sharing firm to increase transparency and disclose how data is used

CONTENTS

- Introduction
- **Machine learning applied to computer vision (CV)**
- General approaches to machine learning
- Pitfalls
- Example project
  - Data exploration
  - Training
  - Testing
  - Deployment
- Deep learning?
- Discussion

ML APPLIED IN COMPUTER VISION

- Classical image processing

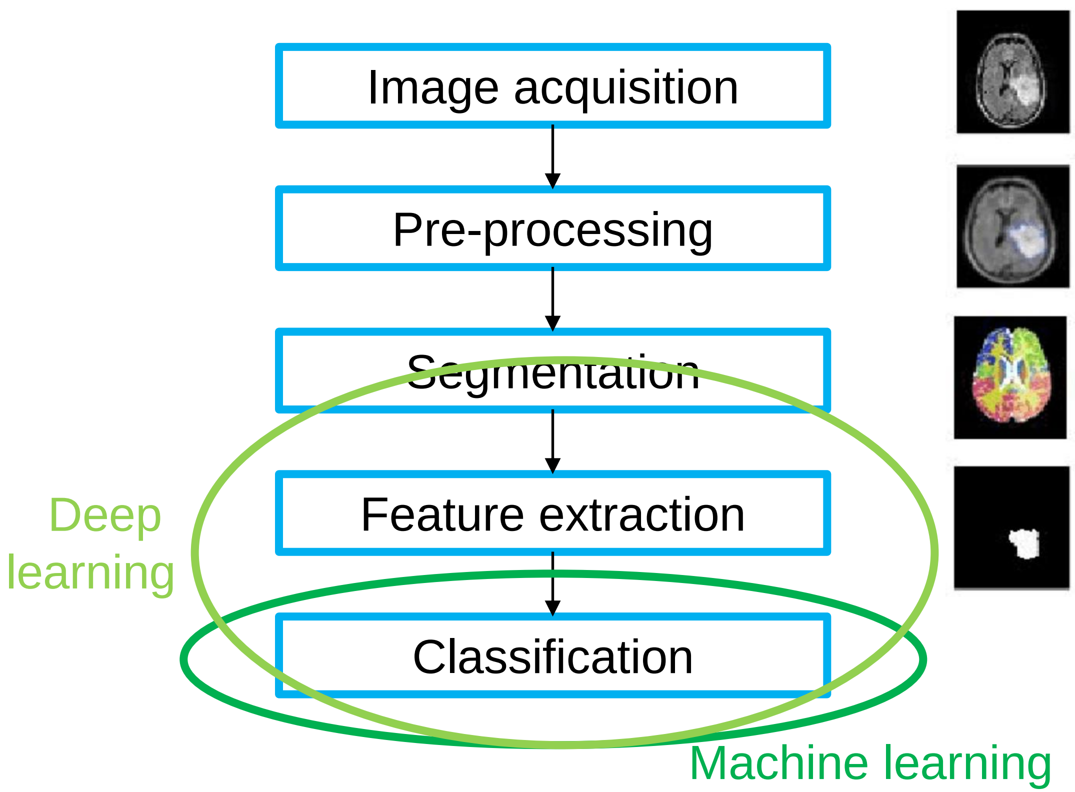

APPLICATIONS IN COMPUTER VISION

- Classification

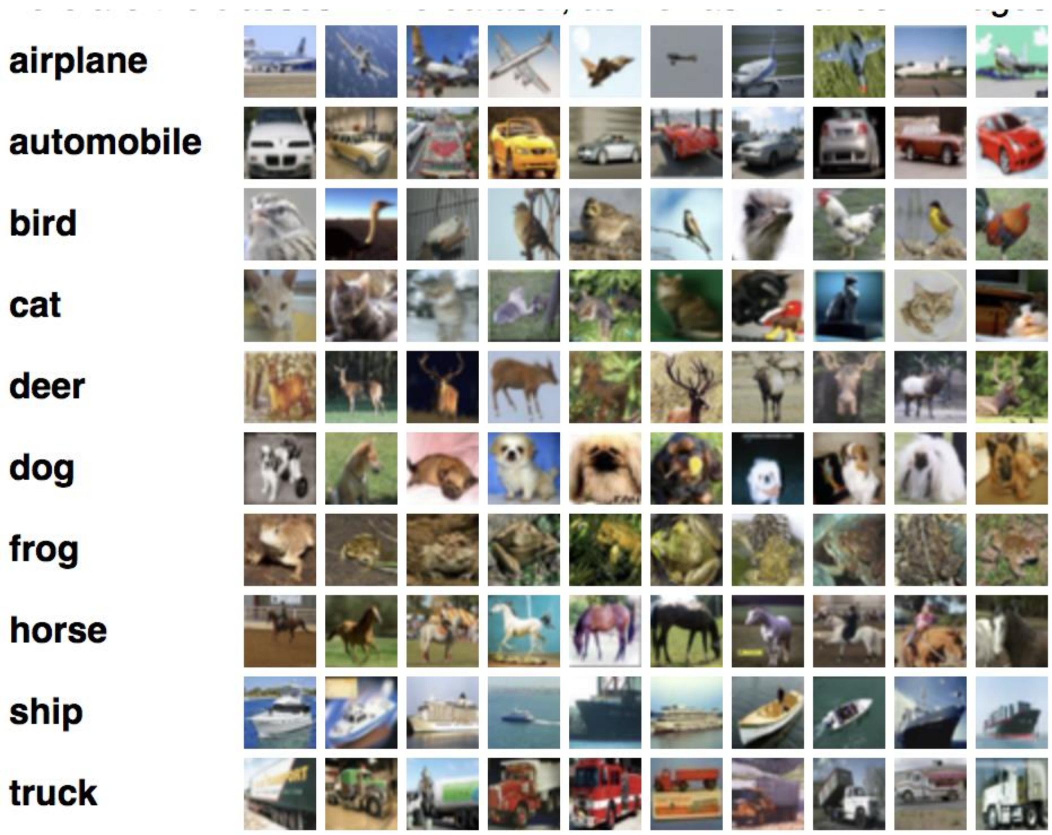

Source: https://machinelearningmastery.com/applications-of-deep-learning-for-computer-vision/

APPLICATIONS IN COMPUTER VISION

- Object detection

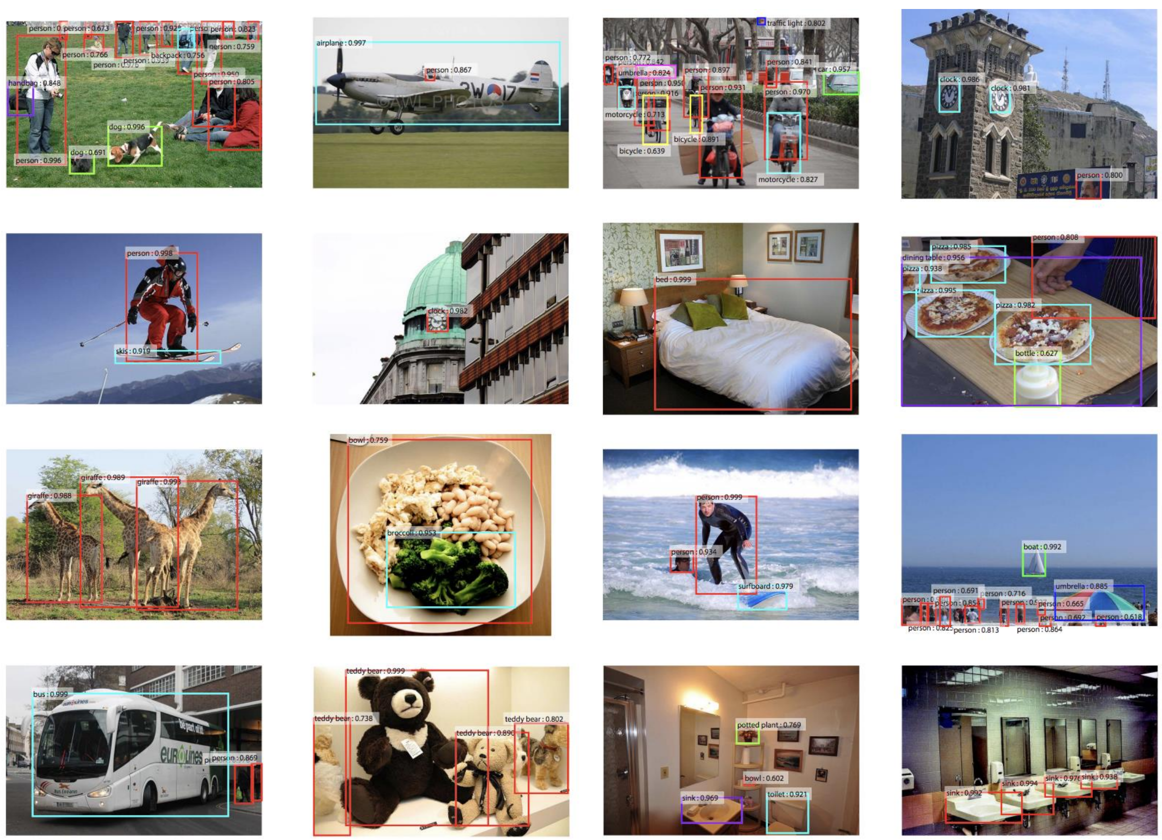

Source: https://machinelearningmastery.com/applications-of-deep-learning-for-computer-vision/

APPLICATIONS IN COMPUTER VISION

- segmentation

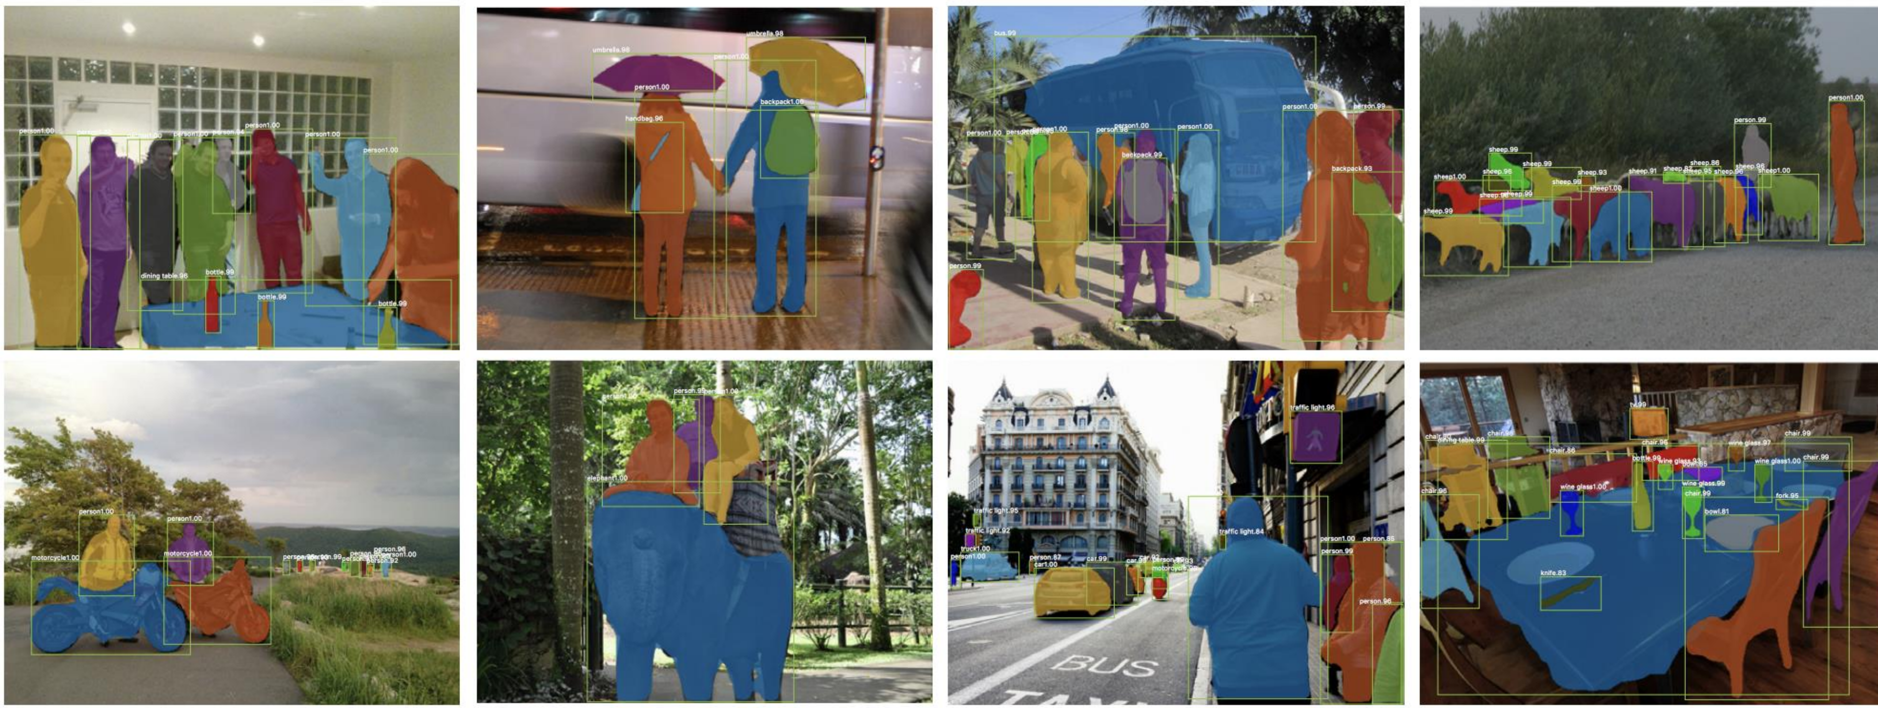

Source: https://machinelearningmastery.com/applications-of-deep-learning-for-computer-vision/

HAN_UNIVERSITY OF APPLIED SCIENCES

APPLICATIONS IN COMPUTER VISION

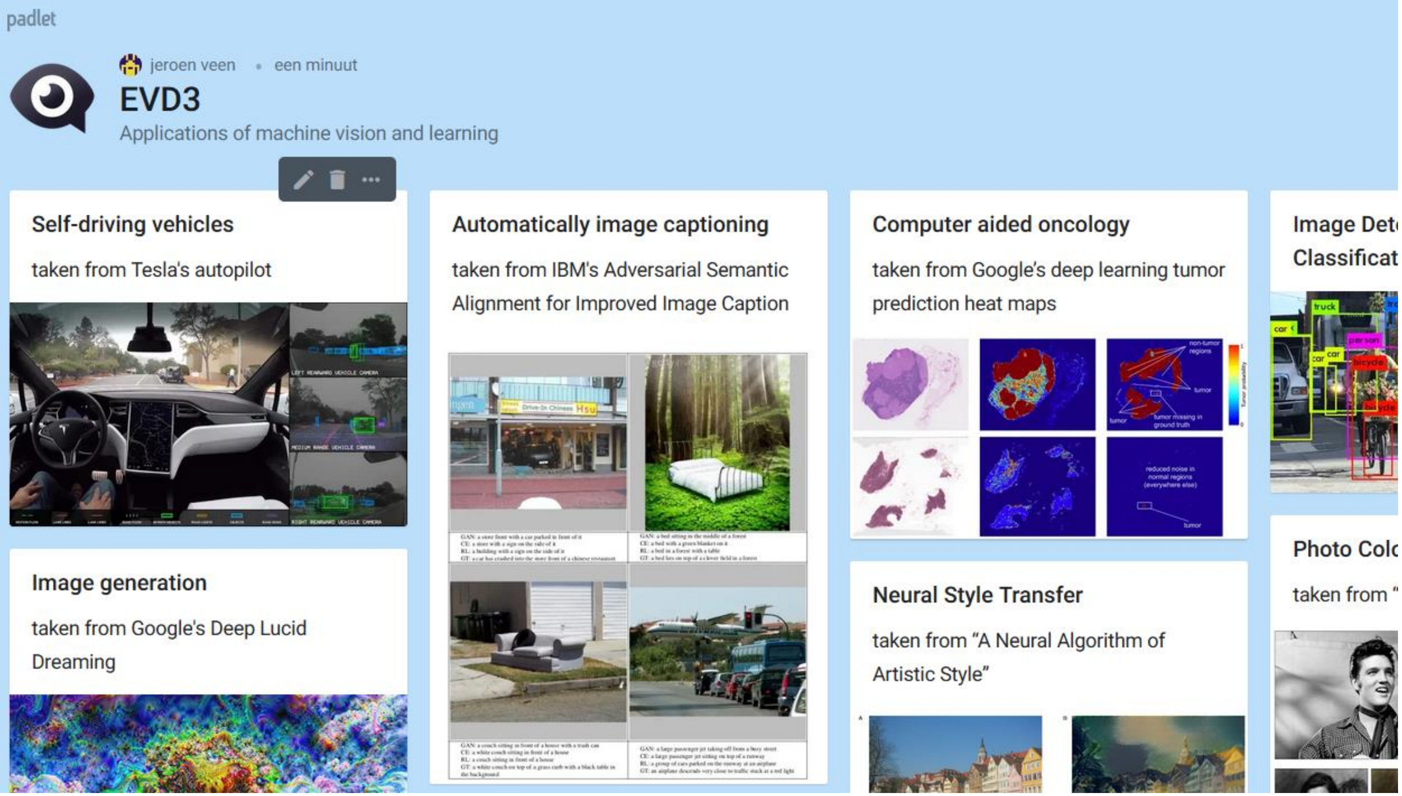

https://padlet.com/jeroen_veen/zul8z8tbvhqpvb8t

CONTENTS

- Introduction
- Machine learning applied to computer vision (CV)
- **General approaches to machine learning**
  - Pitfalls
- Example project
  - Data exploration
  - Training
  - Testing
  - Deployment
- Deep learning?
- Discussion

GENERAL MACHINE LEARNING APPROACHES

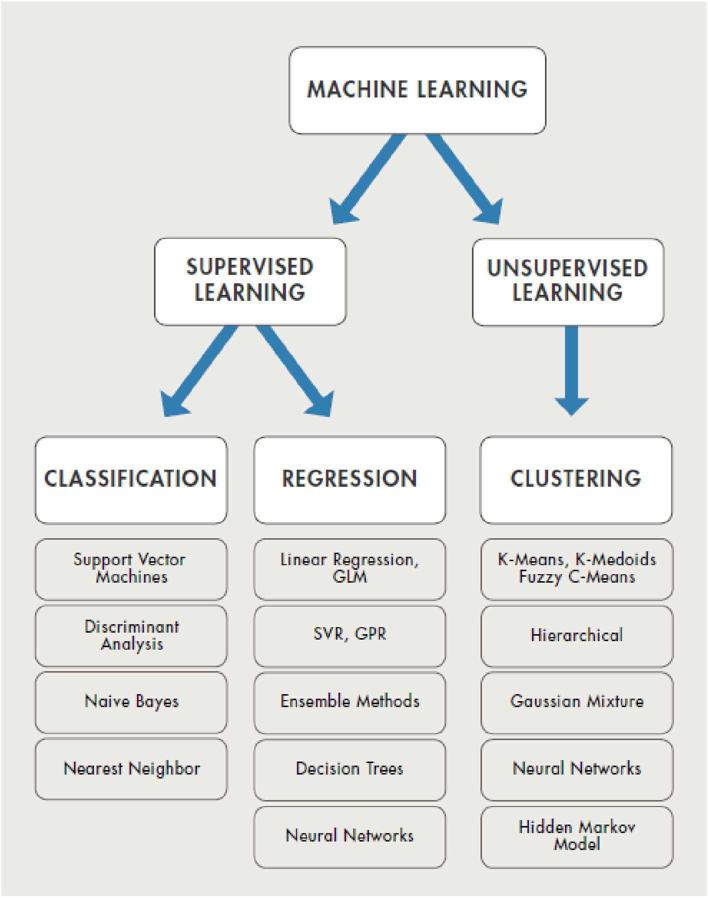

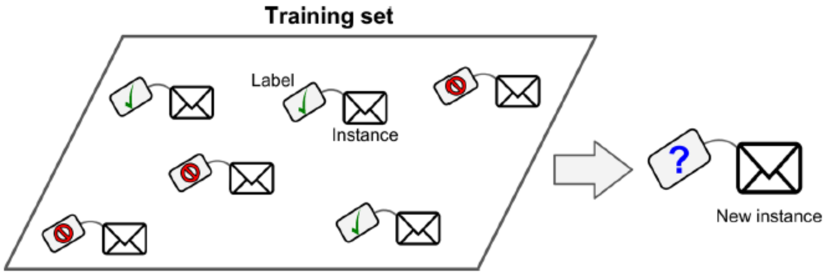

classification

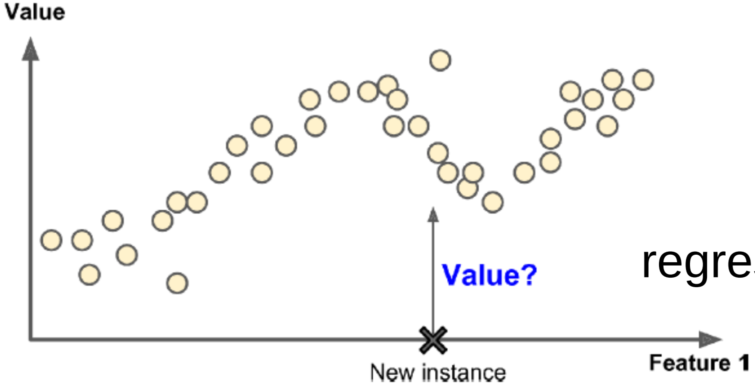

regression

Source: Géron, ISBN: 9781492032632

INSTANCE-BASED VERSUS MODEL-BASED LEARNING

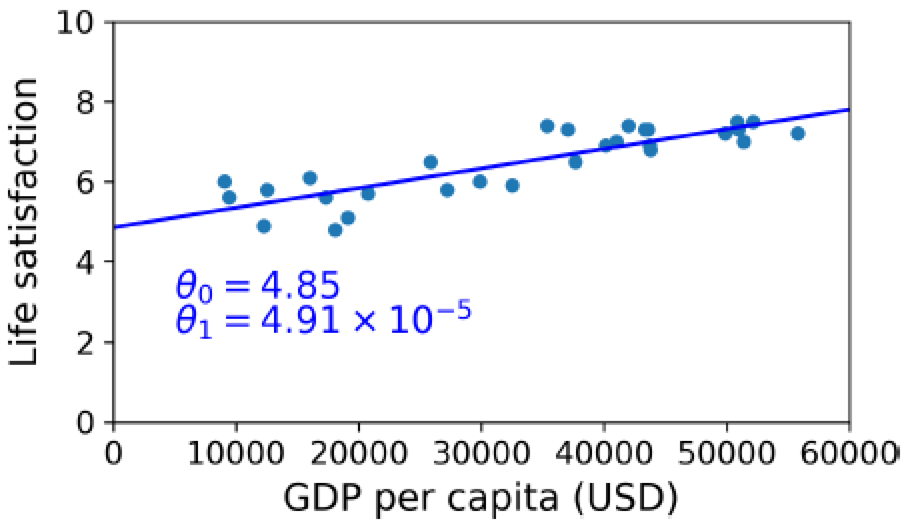

Source: Géron, ISBN: 9781492032632

ML PITFALLS

- Massive amounts of training data is needed
- Labelling is tedious and error prone
- No relationship exists between input and output
- Solution is not transparent
- Solution fails to generalize
- Bias

CONTENTS

- Introduction
- Machine learning applied to computer vision (CV)
- General approaches to machine learning
  - Pitfalls
- **Example project**
  - Data exploration
  - Training
  - Testing
  - Deployment
- Deep learning?
- Discussion

EXAMPLE: HAND GESTURE CLASSIFICATION

- E.g. sign language, rock-paper-scissors
- Min. 3 classes + unknown
- Pick silhouette gestures
- Alternatively, find a simple case within your main project, e.g. objects in autonomous robot

- Term1: solve with ML
- Term2: solve with DL

TYPICAL ML WORKFLOW

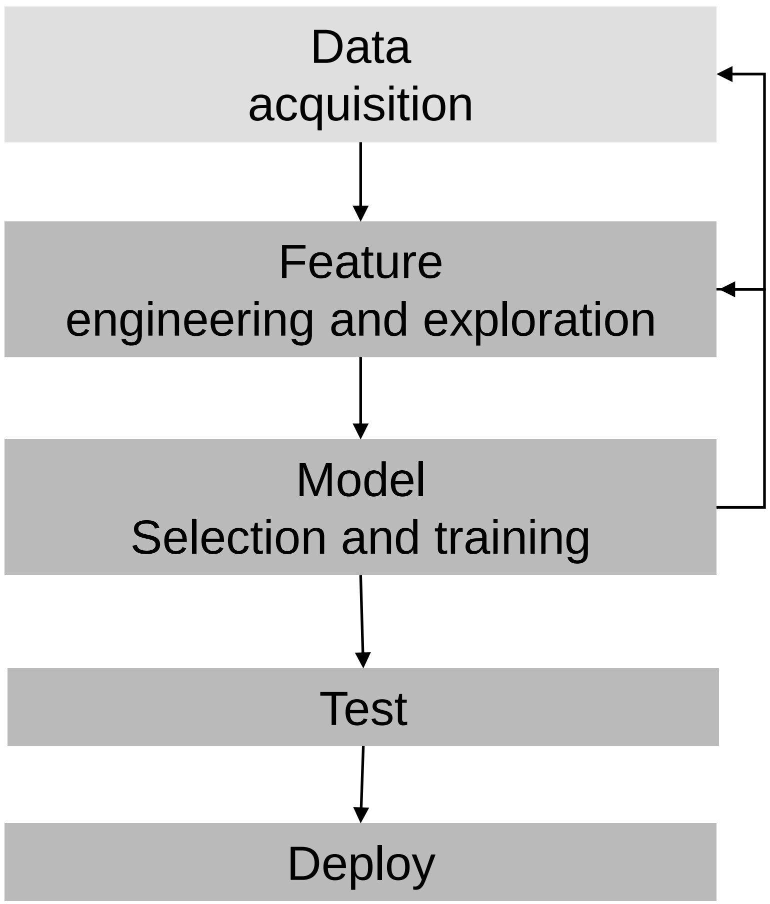

CONTENTS

- Introduction
- Machine learning applied to computer vision (CV)
- General approaches to machine learning
  - Pitfalls
- Example project
  - **Data exploration**
  - Training
  - Testing
  - Deployment
- Deep learning?
- Discussion

DATA ACQUISITION

ignore

paper

rock

scissors

hang loose

FEATURE ENGINEERING

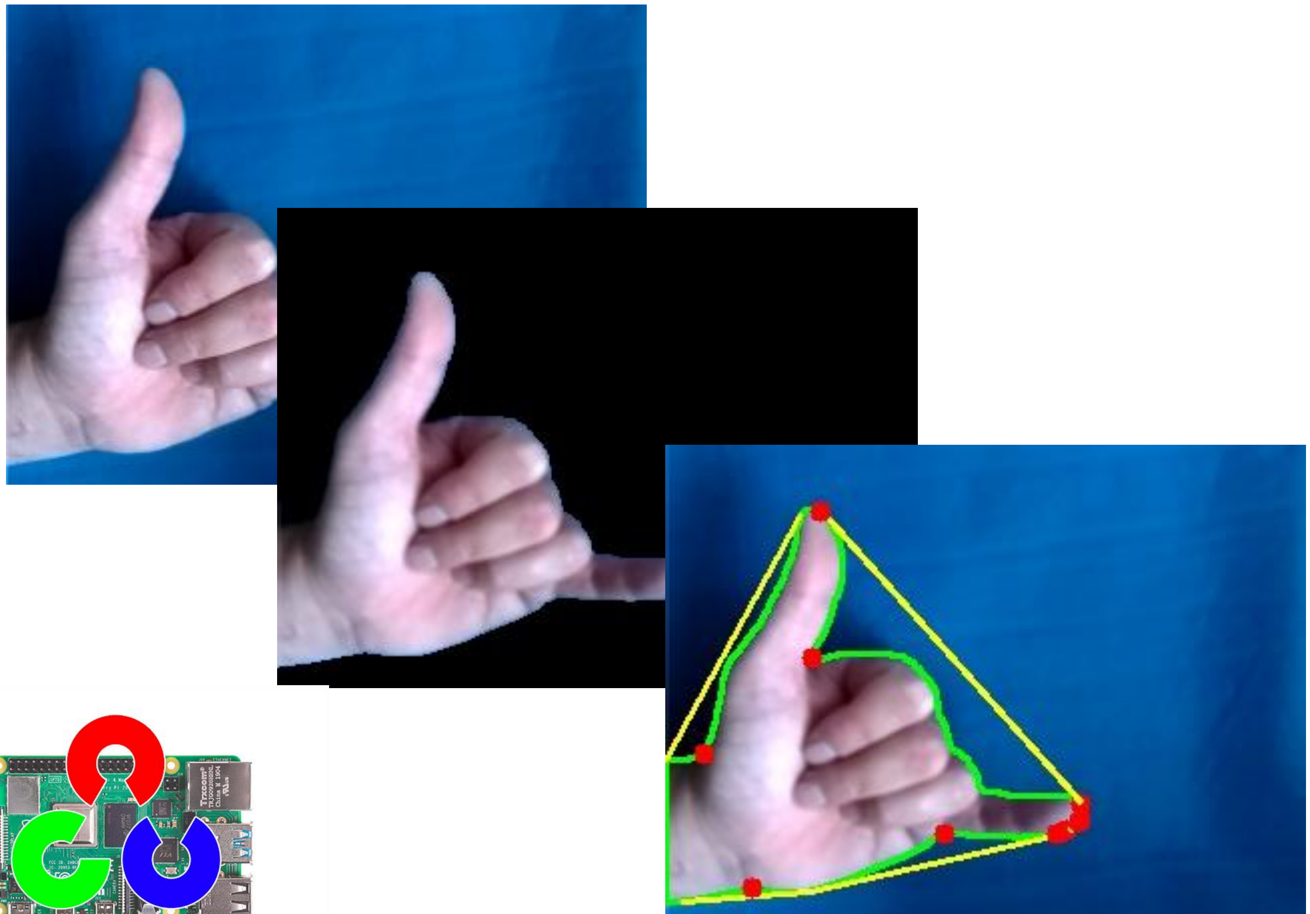

MEET THE DATA

- Examine with numpy or sklearn, or..

Example	Feature 1	Feature 2	Feature N	Label
1	,315	2,2	,51	Rock
2	,312	3,1	,56	Rock
3	,548	5,2	,23	Paper
4	,12	3,5	,1	Scissors
5	,8	6,5	1	Paper
6	,65	4,2	2,5	Paper

TRAINING AND TEST SETS: SPLITTING DATA

- **training set**—a subset to train a model.
- **test set**—a subset to test the trained model.
- You could imagine slicing the single data set as follows:

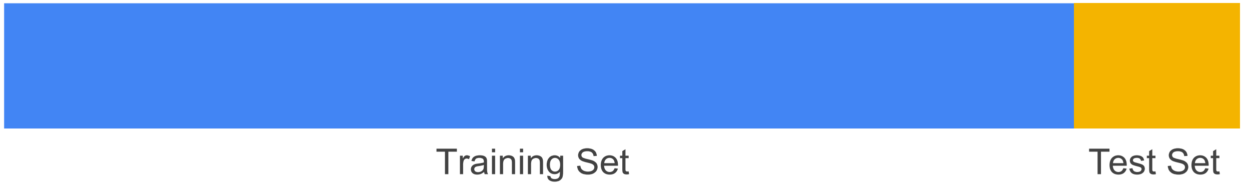

- Make sure that your test set meets the following two conditions:
  - Is large enough to yield statistically meaningful results.
  - Is representative of the data set as a whole. In other words, don't pick a test set with different characteristics than the training set.

https://developers.google.com/machine-learning/crash-course/training-and-test-sets/splitting-data

DATA EXPLORATION EXAMPLE

- explore.py

Total feature histograms

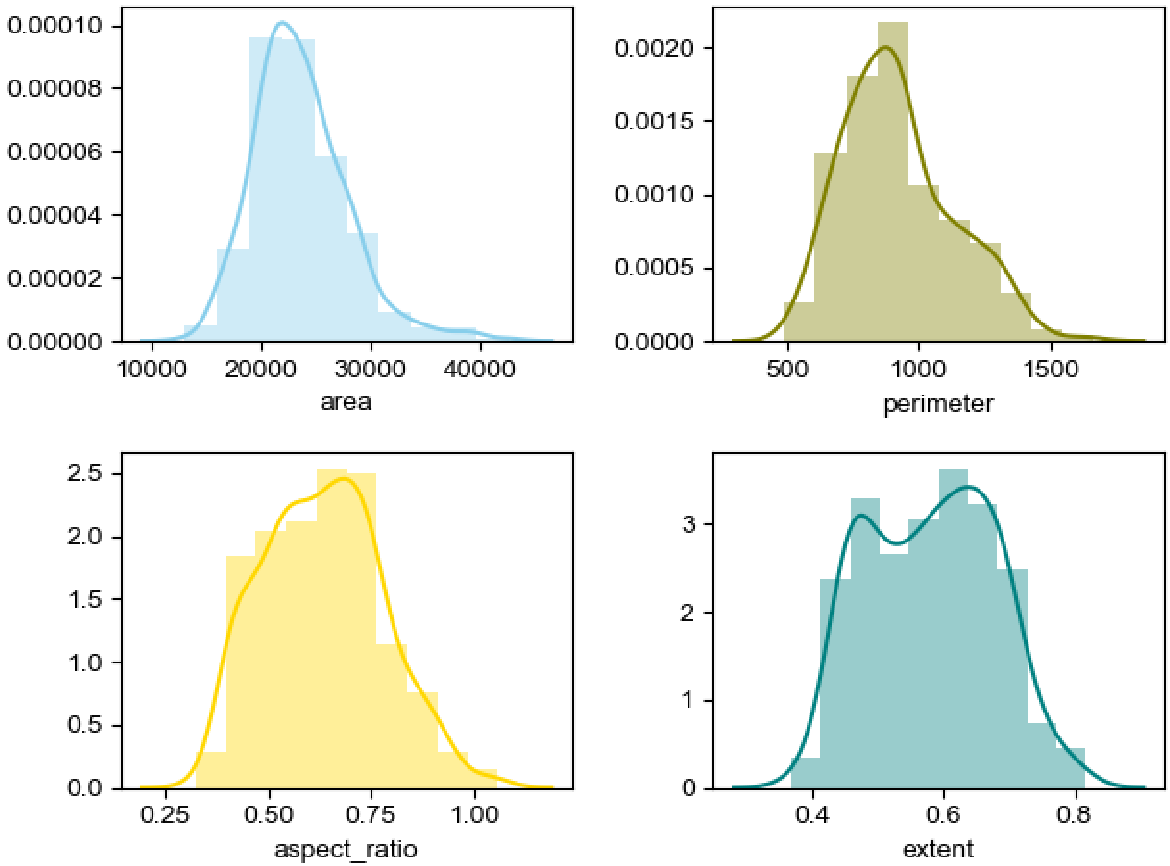

HAN_UNIVERSITY OF APPLIED SCIENCES

QUALITIES OF GOOD FEATURES

- Informative
- Discriminating
- Independent
- Nearly unique

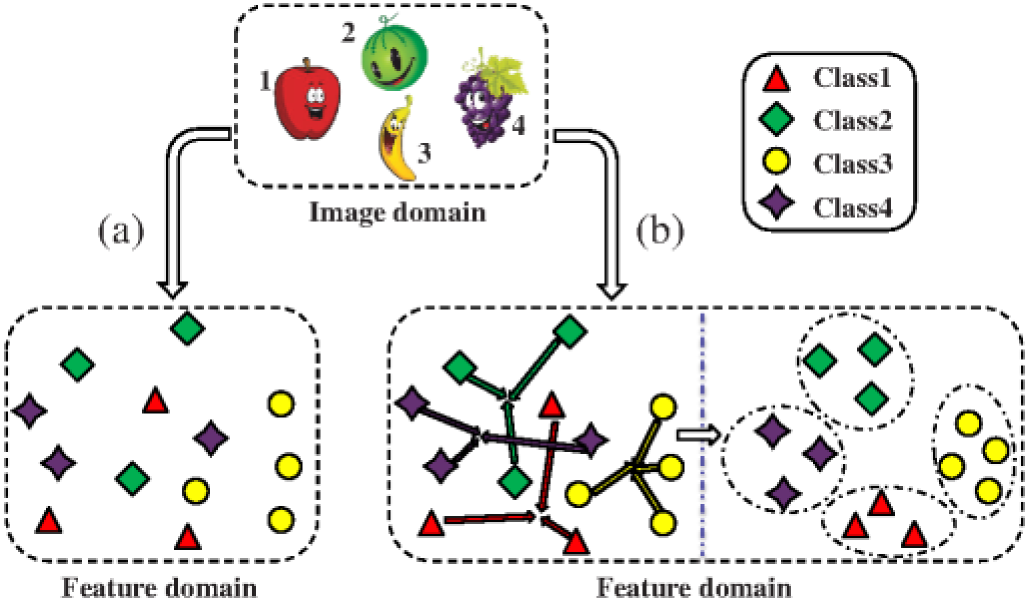

Source: https://www.spiedigitallibrary.org/ContentImages/Journals/JEIME5/26/1/013023

- NB feature scaling may be required

DATA EXPLORATION EXAMPLE

- explore.py

Boxplot of a single feature

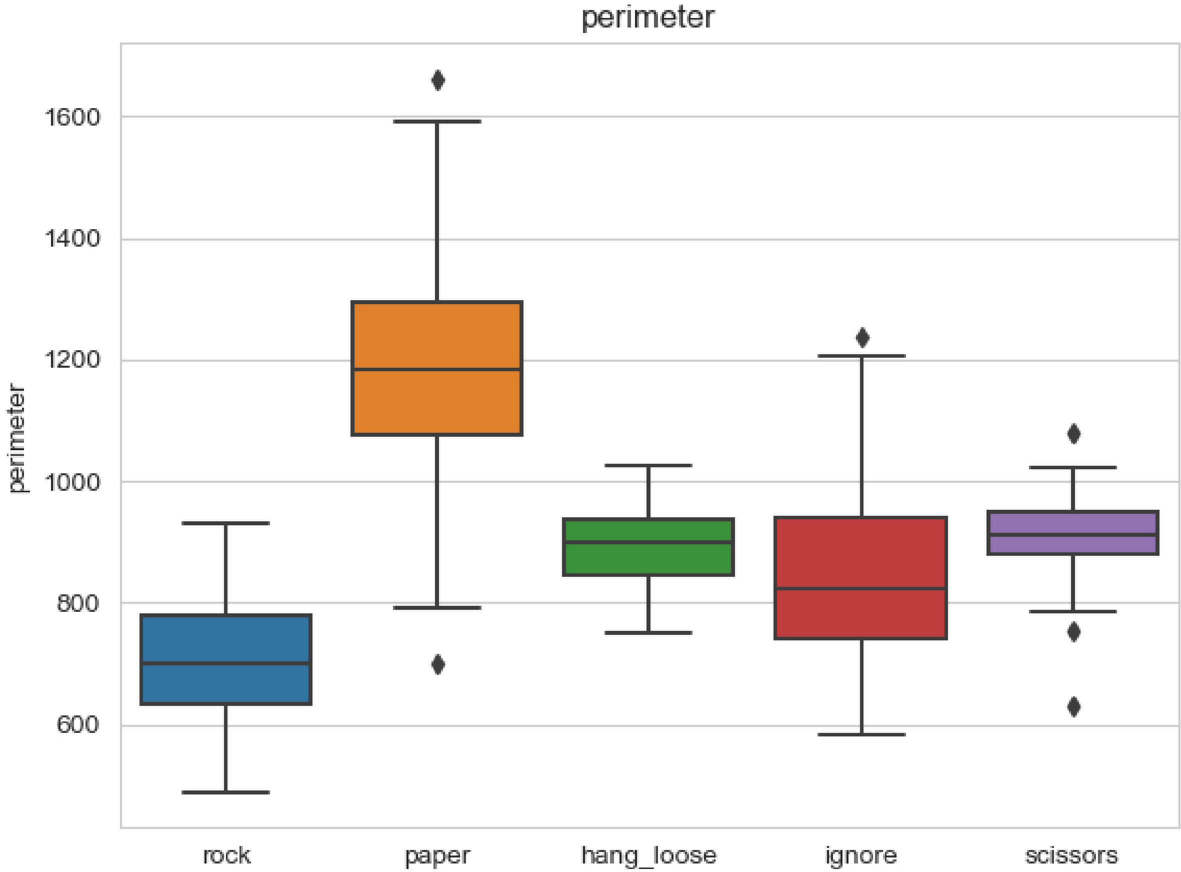

DATA EXPLORATION EXAMPLE

- explore.py

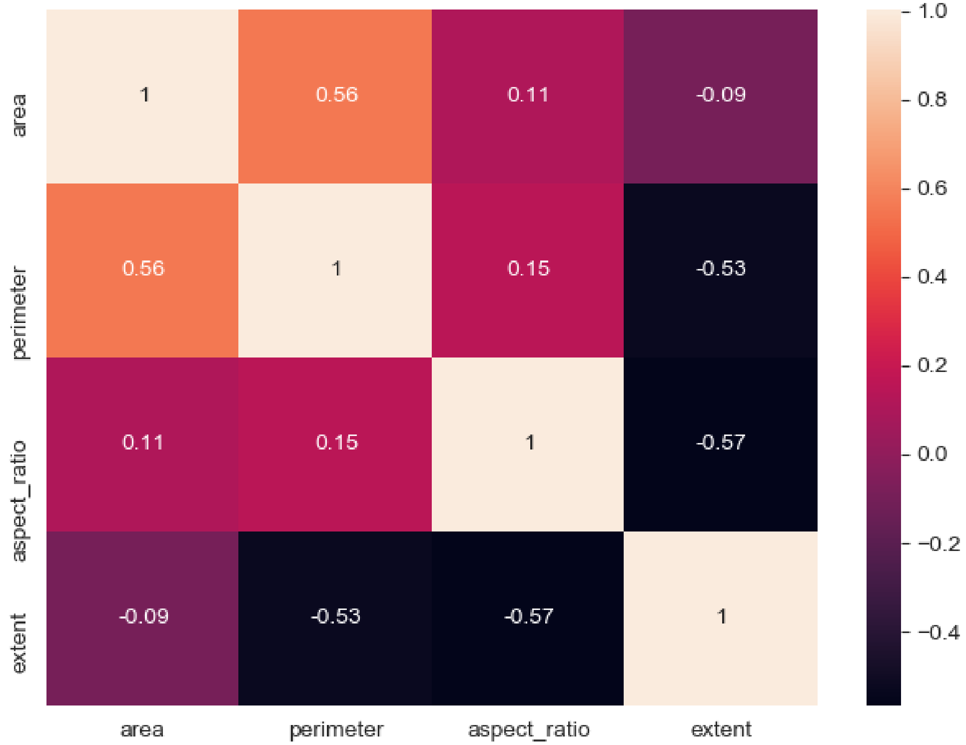

Feature correlation heatmap

See: https://www.statology.org/how-to-read-a-correlation-matrix/

CONTENTS

- Introduction
- Machine learning applied to computer vision (CV)
- General approaches to machine learning
  - Pitfalls
- Example project
  - Data exploration
  - **Training**
  - Testing
  - Deployment
- Deep learning?
- Discussion

ALGORITHMS, ALGORITHMS

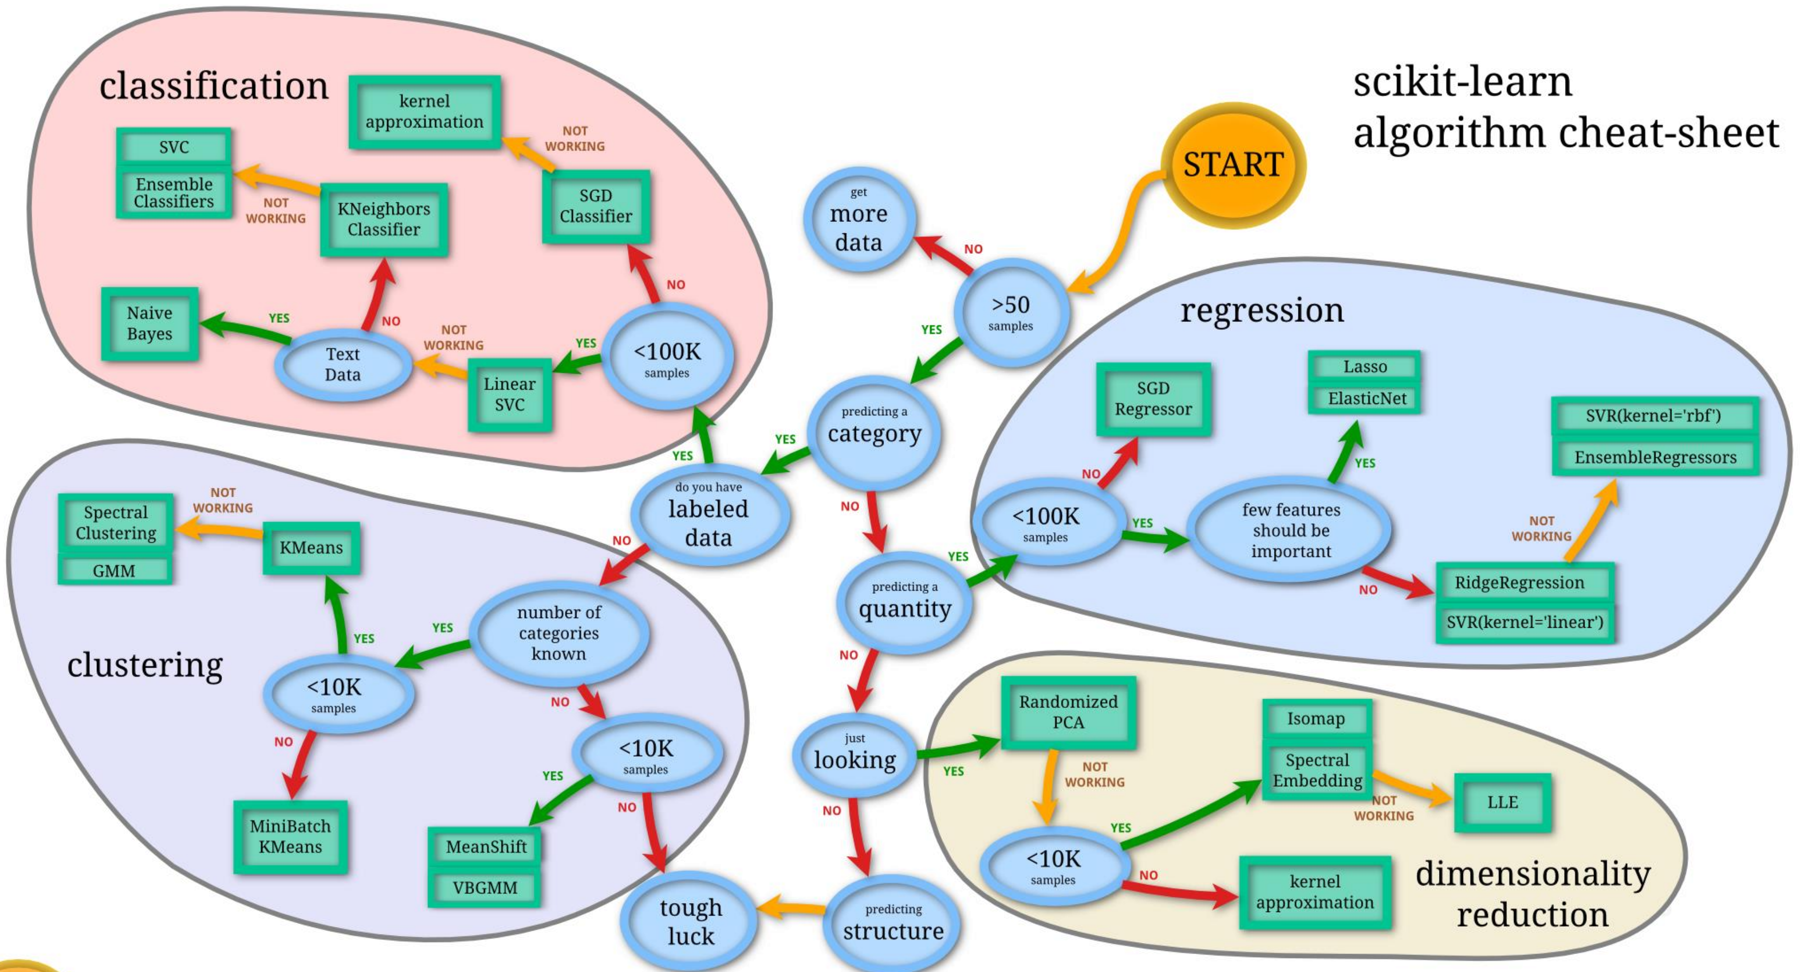

K NEAREST NEIGHBOR (KNN)

- The simplest classifier
- Assume feature vectors near each other are similar
- Categorizes objects based on the classes of their nearest neighbors
- No training required
- Intuitive
- Benchmark

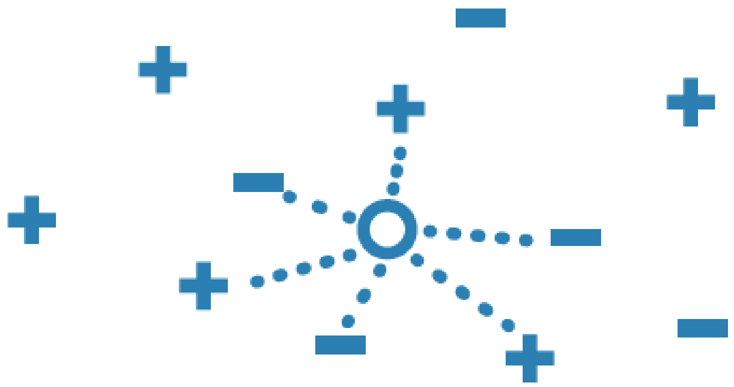

Source: Mathworks, Applying Supervised Learning

“Tell me who your neighbors are, and I’ll tell you who you are”

MAKING PREDICTIONS

- Comparing to every example is very slow
- More suited for low-dimensional feature spaces (which images are not)

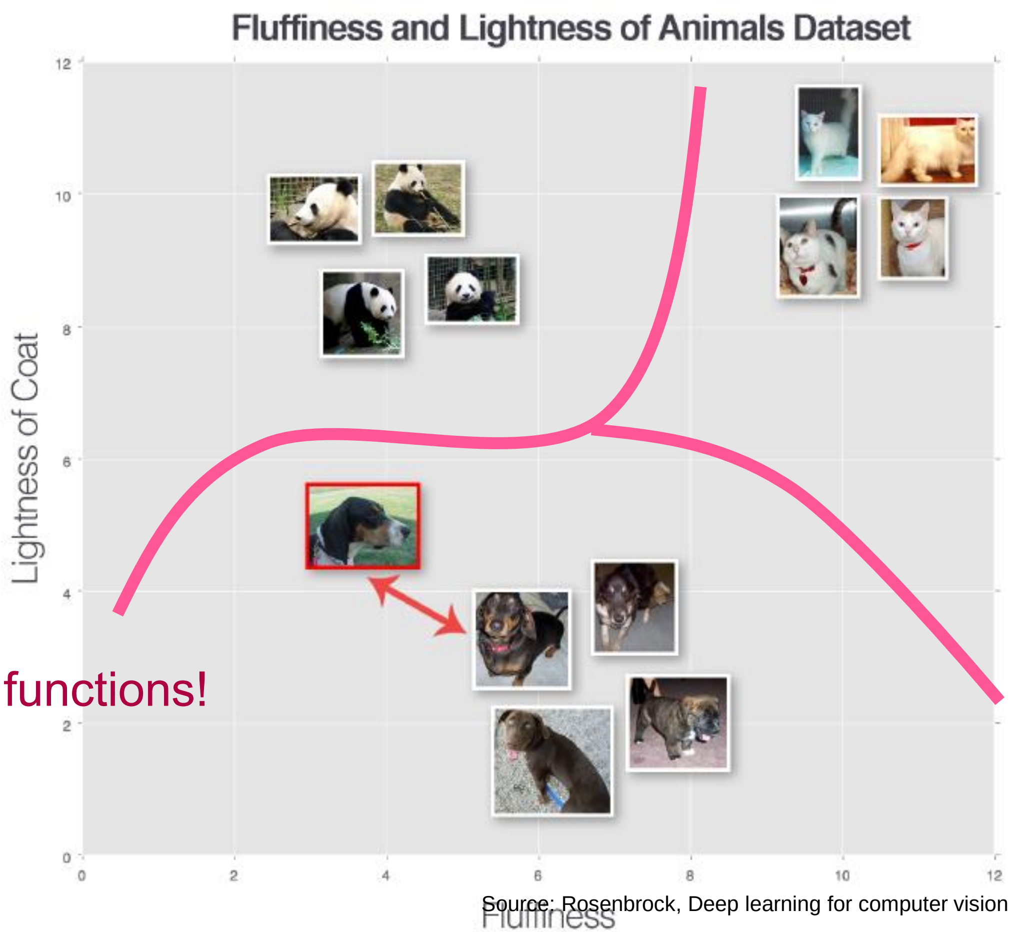

Let's just define some boundary functions!
-> we build a data model

Source: Rosenbrock, Deep learning for computer vision

SVM

- Powerful and versatile ML model
- Simple and easy to interpret.
- Intuitive
- Linear or nonlinear classification
- Binary classifier, so for multi-class data, reduction to several binary problems needed

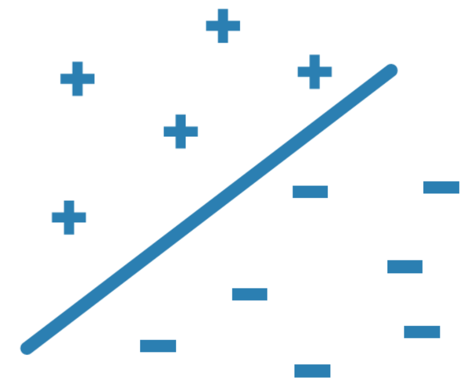

Source: Mathworks, Applying Supervised Learning

"Find the line that separates us"

NONLINEAR SVM CLASSIFICATION

- E.g. polynomial kernel

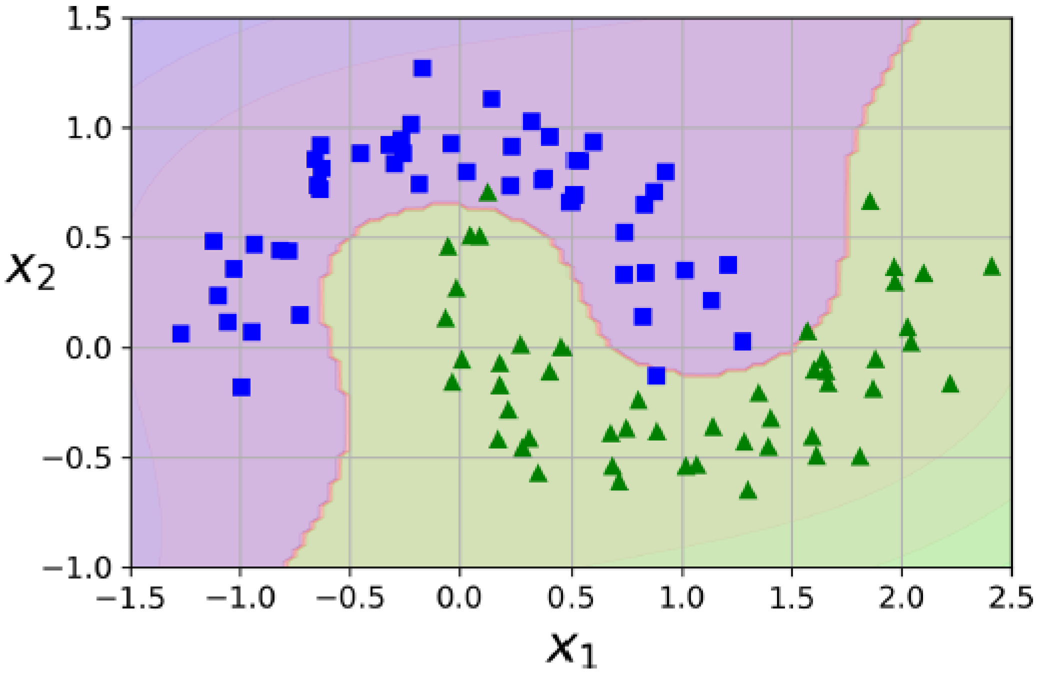

Source: Géron, ISBN: 9781492032632

DECISION TREES

- Predict responses to data by following the decisions in the tree from the down to a leaf node.
- Easy to interpret
- Fast to fit
- Minimize memory usage

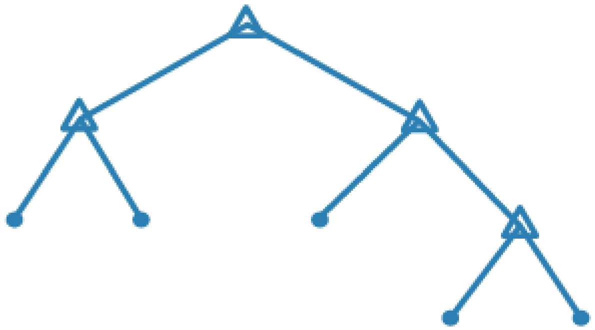

Source: Mathworks, Applying Supervised Learning

DECISION TREE BOUNDARIES

- white box models

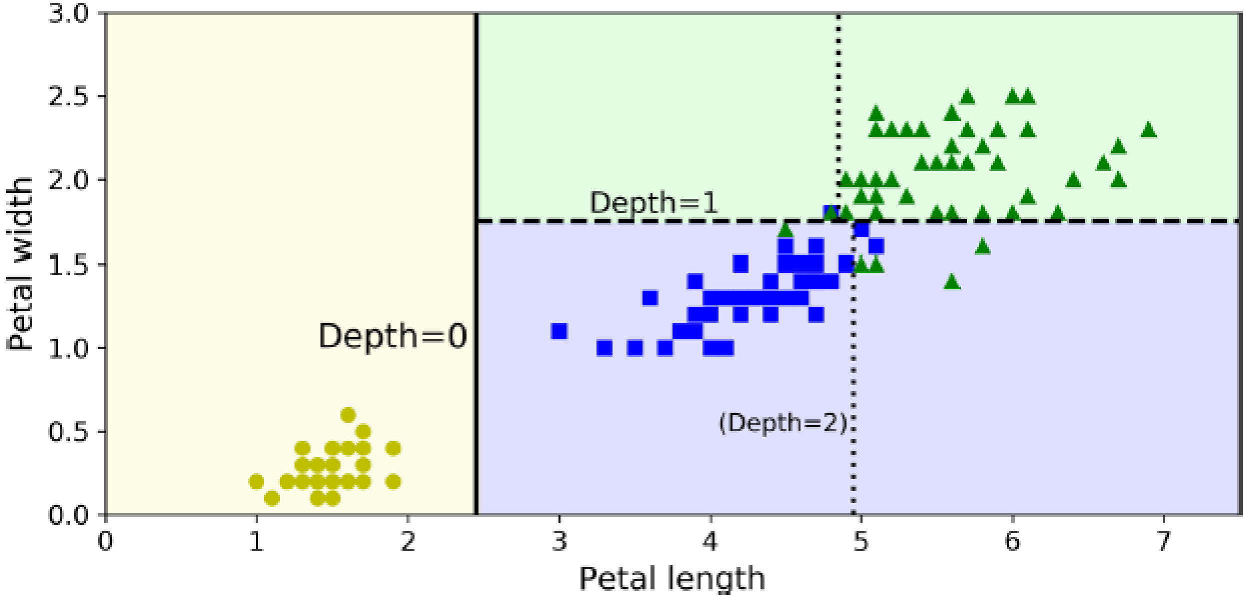

Source: Géron, ISBN: 9781492032632

HAN_UNIVERSITY OF APPLIED SCIENCES

ARTIFICIAL NEURAL NETWORKS

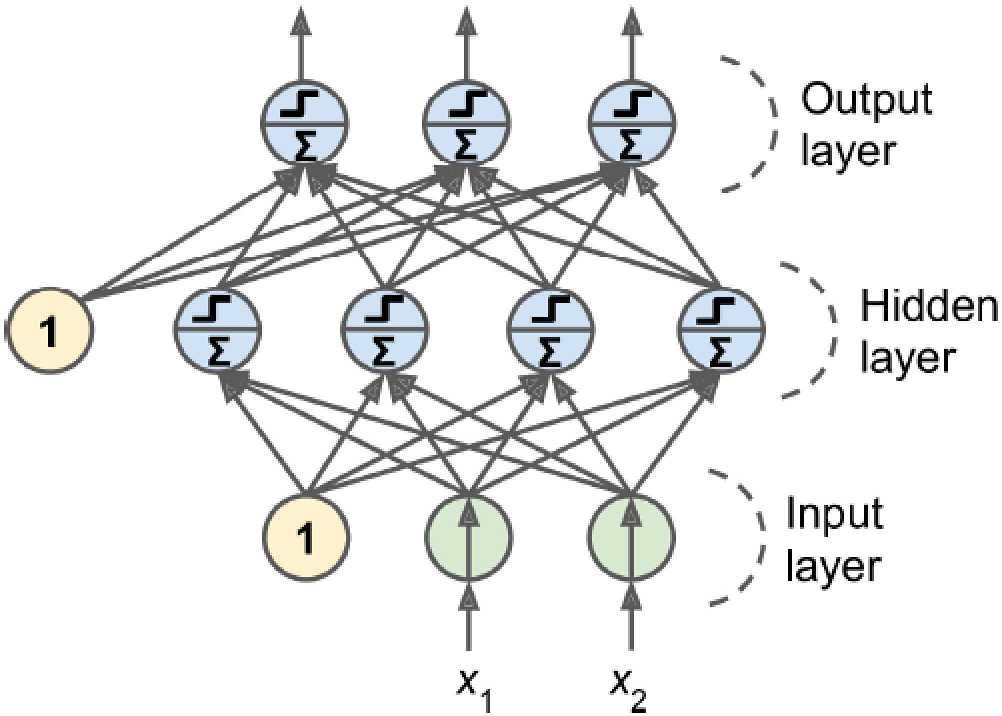

Source: Géron, ISBN: 9781492032632

CONTENTS

- Introduction
- Machine learning applied to computer vision (CV)
- General approaches to machine learning
  - Pitfalls
- Example project
  - Data exploration
  - Training
  - **Testing**
  - Deployment
- Deep learning?
- Discussion

THE BOY WHO CRIED WOLF

"Wolf" is a **positive class**.

"No wolf" is a **negative class**

An Aesop's Fable ~620 BCE

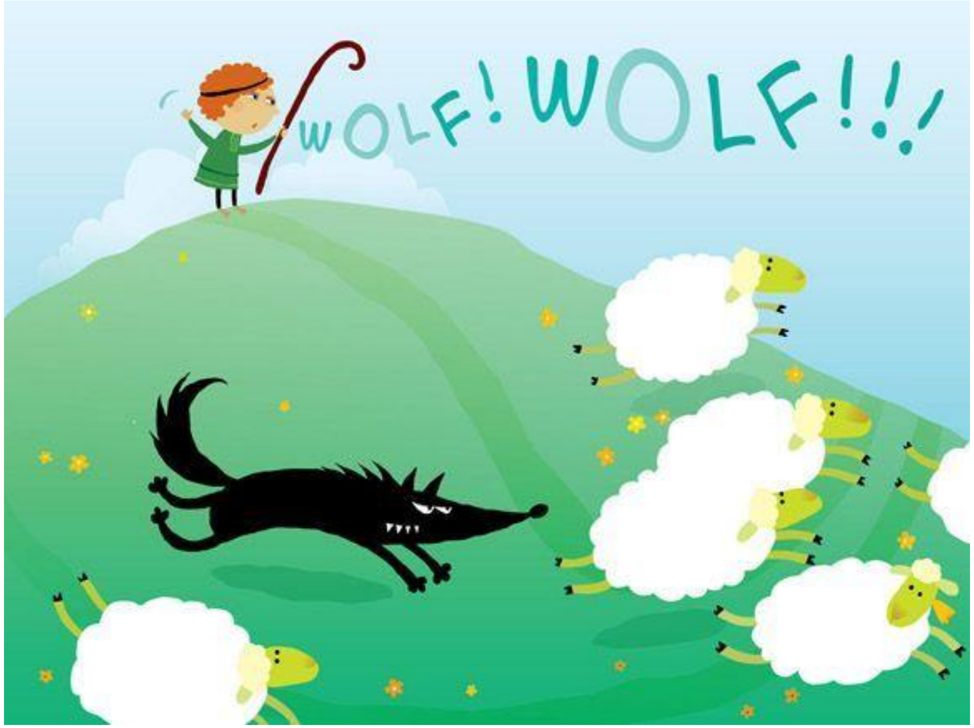

Source: Sam Taplin

CONFUSION MATRIX

ACTUAL

PREDICTED

(Type I error)

True Positive (TP) Reality: A wolf threatened. Shepherd said: "Wolf." Outcome: Shepherd is a hero.	**False Positive (FP)** Reality: No wolf threatened. Shepherd said: "Wolf." Outcome: Villagers are angry at shepherd for waking them up.
False Negative (FN) Reality: A wolf threatened. Shepherd said: "No wolf." Outcome: The wolf ate all the sheep.	**True Negative (TN)** Reality: No wolf threatened. Shepherd said: "No wolf." Outcome: Everyone is fine.

Type II error)

ACCURACY

- Fraction of predictions the model got right

$$\text{Accuracy} = \frac{\text{Number of correct predictions}}{\text{Total number of predictions}}$$

- For binary classification

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

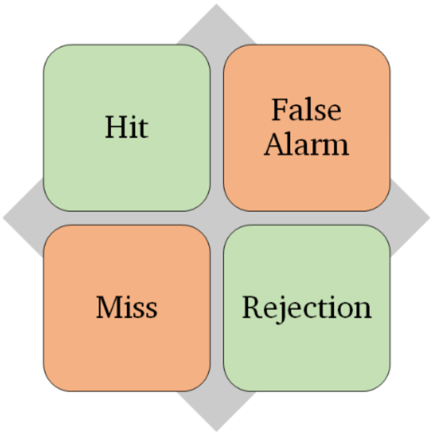

PREDICTION BIAS–VARIANCE TRADEOFF

- Central problem in supervised learning

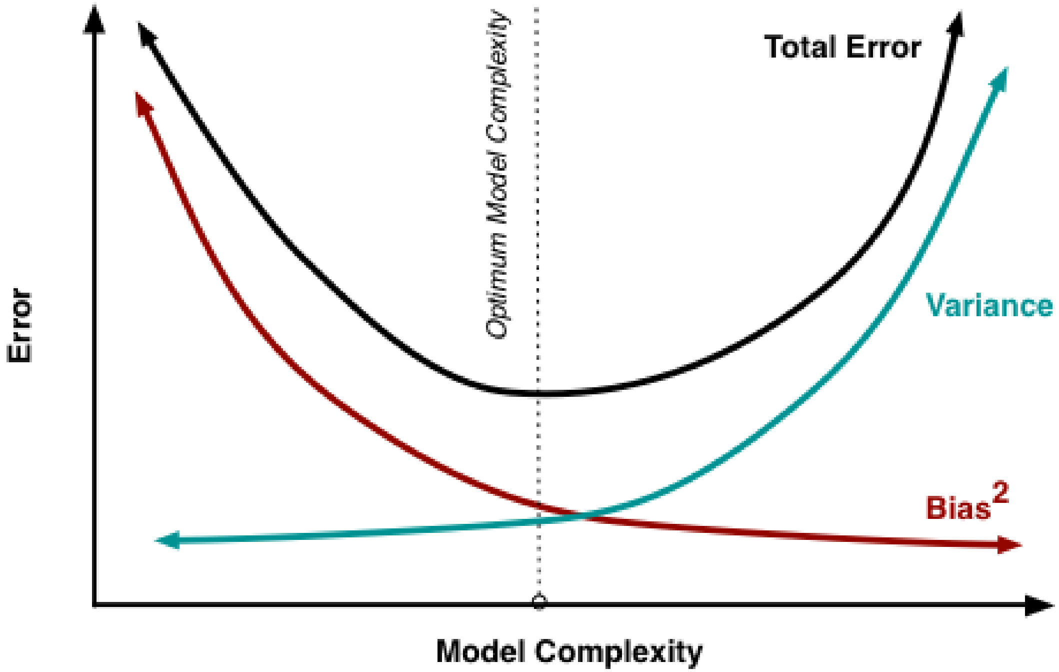

CONTENTS

- Introduction
- Machine learning applied to computer vision (CV)
- General approaches to machine learning
  - Pitfalls
- Example project
  - Data exploration
  - Training
  - Testing
  - **Deployment**
- Deep learning?
- Discussion

MODEL DEPLOYMENT

Target systems:

SBCs such as Rpi, Jetson Nano
MCUs such as STM32, ESP32, Kendryte K210,KD233 even Arduino

Models:
Custom models, proprietary models, or Tensorflow, tensorflowlite,

CONTENTS

- Introduction
- Machine learning applied to computer vision (CV)
- General approaches to machine learning
  - Pitfalls
- Example project
  - Data exploration
  - Training
  - Testing
  - Deployment
- **Deep learning?**
- Discussion

HAN_UNIVERSITY OF APPLIED SCIENCES