POLITECNICO
MILANO 1863

Artificial Neural Networks and Deep Learning

- Introduction to the course -

https://boracchi.faculty.polimi.it/

Prof. Giacomo Boracchi – *giacomo.boracchi@polimi.it*
Loris Giulivi – *loris.giulivi@polimi.it*
but also …
Prof. Matteo Matteucci – *matteo.matteucci@polimi.it*
Eng. Eugenio Lomurno – *eugenio.lomurno@polimi.it*
Eng. Francesco Lattari – *francesco.lattari@polimi.it*

Who I am

Giacomo Boracchi (https://boracchi.faculty.polimi.it/)

- Mathematician (Università Statale degli Studi di Milano 2004),
- PhD in Information Technology (DEIB, Politecnico di Milano 2008)
- Associate Professor since 2019 at DEIB, Polimi (Computer Science)

My Research Interests are mathematical and statistical methods for:

- Machine Learning and in particular unsupervised learning, change and anomaly detection
- Image analysis and processing

… and the two combined

Teaching

Advanced courses taught:

- Artificial Neural Networks and Deep Learning (MSc)
- Mathematical Models and Methods for Image Processing (MSc, spring 2023)
- Advanced Deep Learning Models And Methods (PhD, Winter 2022 with Prof. Matteucci)
- Online Learning and Monitoring (PhD, Spring 2022 with Prof Trovò)

- Computer Vision and Pattern Recognition (MSc in USI, Spring 2020)
- Learning Sparse Representations for image and signal modeling (PhD)

- Informatica A (Mathematical Engineering!!!)

Course Objectives

*"The course major goal is to provide students with the theoretical background and the practical skills to **understand and use Neural Networks,** and, at the same time, become familiar and with **Deep Learning for solving complex engineering problems** … especially **in vision tasks**"*

A Course with Code Sharing

This course is offered to Bioengineering and Mathematical Engineering

- 056869 - ARTIFICIAL NEURAL NETWORKS AND DEEP LEARNING - 5 CFU
- Prof. Giacomo Boracchi, Eng. Eugenio Lomurno

… **equivalent** course for Computer Science and Engineering students

- 054307 - ARTIFICIAL NEURAL NETWORKS AND DEEP LEARNING - 5 CFU
- Prof. Matteo Matteucci, Eng. Francesco Lattari

The **same teachers will teach the same topics to both classes**, but you need to be enrolled in the right course and attend the right lectures …

The Teachers

Official teacher, please refer to me for bureaucratic stuff!

Prof. Matteo Matteucci

- Neural Networks
- Deep Learning
- Sequence Learning

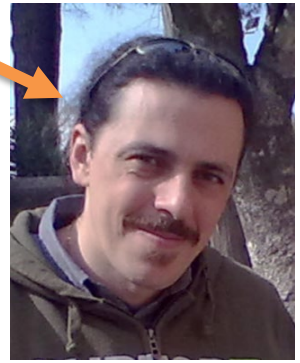



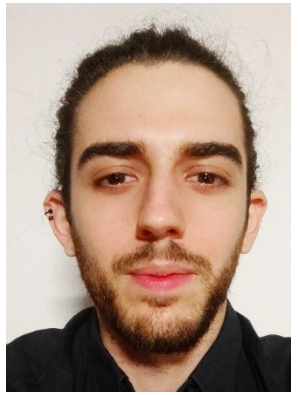

Prof. Giacomo Boracchi

- Deep Learning for visual recognition (Classification, Segmentation, Detection..)

Loris Giulivi, Francesco Lattari and Eugenio Lomurno

- Programming DL in Py
- Online Challenges

https://boracchi.faculty.polimi.it/teaching/AN2DLCalendar.htm

A detailed schedule on Google Calendar

Each event includes

- Teacher
- Possibly last-minute slides
- Links to video recordings

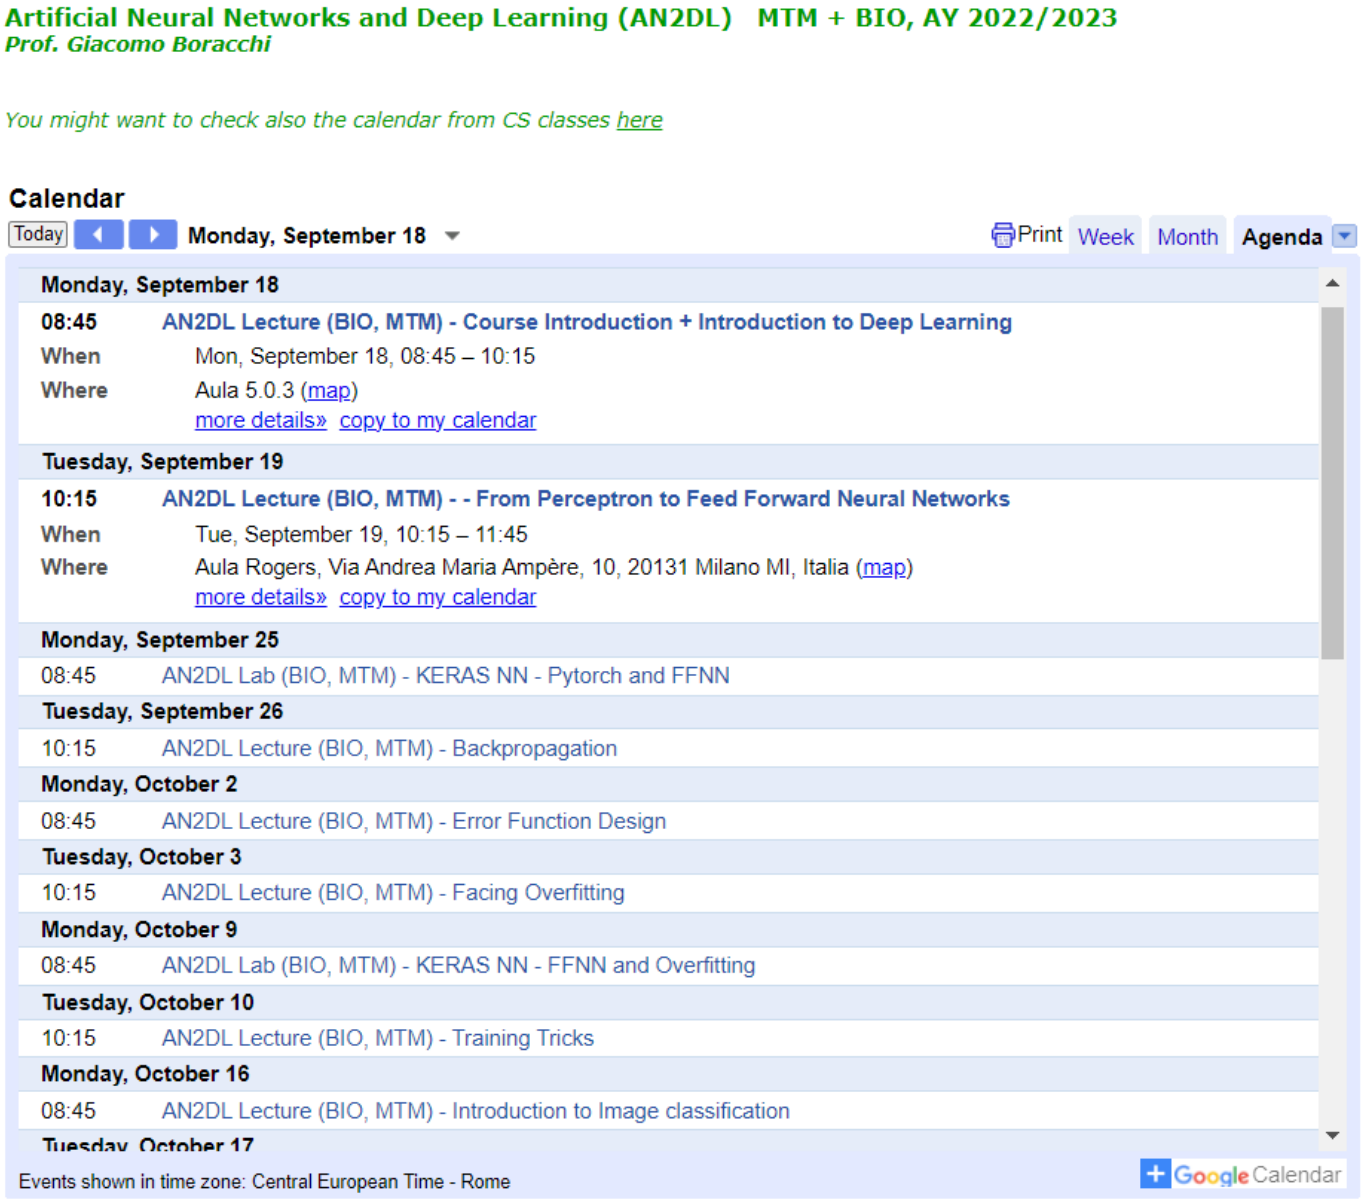

Artificial Neural Networks and Deep Learning (AN2DL) MTM + BIO, AY 2022/2023
Prof. Giacomo Boracchi

You might want to check also the calendar from CS classes here

Calendar

Today | Monday, September 18 | Print | Week | Month | Agenda

Monday, September 18

08:45 **AN2DL Lecture (BIO, MTM) - Course Introduction + Introduction to Deep Learning**
When Mon, September 18, 08:45 – 10:15
Where Aula 5.0.3 (map)
more details» copy to my calendar

Tuesday, September 19

10:15 **AN2DL Lecture (BIO, MTM) - - From Perceptron to Feed Forward Neural Networks**
When Tue, September 19, 10:15 – 11:45
Where Aula Rogers, Via Andrea Maria Ampère, 10, 20131 Milano MI, Italia (map)
more details» copy to my calendar

Monday, September 25

08:45 AN2DL Lab (BIO, MTM) - KERAS NN - Pytorch and FFNN

Tuesday, September 26

10:15 AN2DL Lecture (BIO, MTM) - Backpropagation

Monday, October 2

08:45 AN2DL Lecture (BIO, MTM) - Error Function Design

Tuesday, October 3

10:15 AN2DL Lecture (BIO, MTM) - Facing Overfitting

Monday, October 9

08:45 AN2DL Lab (BIO, MTM) - KERAS NN - FFNN and Overfitting

Tuesday, October 10

10:15 AN2DL Lecture (BIO, MTM) - Training Tricks

Monday, October 16

08:45 AN2DL Lecture (BIO, MTM) - Introduction to Image classification

Tuesday, October 17

Events shown in time zone: Central European Time - Rome

Google Calendar

The Students

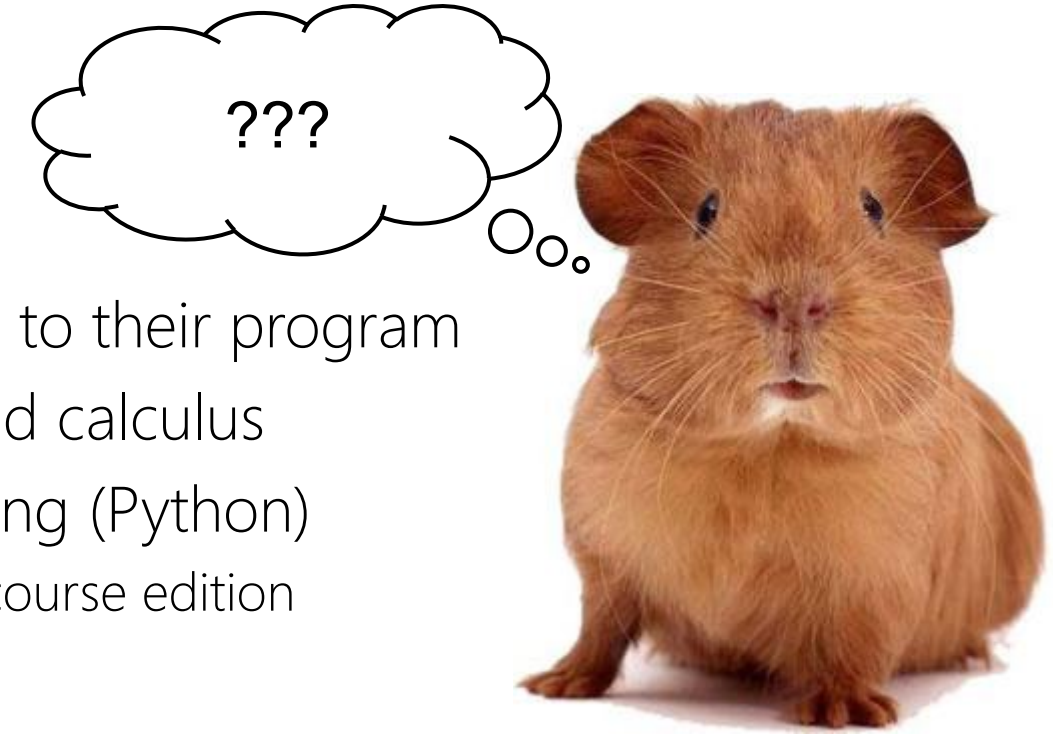

Students are expected to:

- To attend the proper classes according to their program
- Feel confortable with basic statistics and calculus
- Feel confortable with basic programming (Python)
  - Be ready to act as «guinea pigs» for this course edition
- Be curious and willing to learn …

Students are not expected to:

- Know more than what is usually tought in basic engineerig courses
- Know already about machine learning (althought it doesn't hurt)
- Be hyper-skilled python hackers (you'll not need it)
- …

Course syllabus

Introduction to Neural Network and Deep Learning — 2h lectures

Neural Networks and Deep Learning — 16h lectures

- From the Perceptron to neural networks
- Backpropagation and neural networks training
- Best practices in neural network training
- Recurrent architectures
- Autoencoders and long short-term memories

Visual Recognition with Deep Neural Networks — 16h lectures

- Image Classification and Convolutional Neural Networks
- CNN Training Tricks and Best Practices
- CNN for Advanced Vision Tasks (Segmentation, Detection,…)

ANN and Deep Learning Coding (with Keras) — 16h practicals

Course Website and Detailed Schedule

All details and info are on the course website

https://boracchi.faculty.polimi.it/teaching/AN2DL.htm

https://chrome.deib.polimi.it/index.php?title=Artificial_Neural_Networks_and_Deep_Learning

How to get there?

- From our websites https://boracchi.faculty.polimi.it/
- Select "Teaching And Available Thesis", then "Artificial Neural Network and Deep Learning"

What do you find there:

- Detailed schedule
- Lecture slides / links

Lectures Schedule and Timings MTM/BIO

Classes (there is no distinction between lecture and exercises):

- Monday in Room B.4.2. from **8.45** till 10.15 (we postpone a bit the lecture start)
- Tuesday in Room 7.1.3 from 10.30 till 12.00

Check the teacher who will be in class on the detailed schedule

- Lectures will be recorded and made available afterwards
- Lectures won't be streamed

Lectures Schedule and Timings CSE

You might consider attending these to avoid overlaps

- Wednesday, 16:15 – 18:15, in T2.2 (starts at 16:30 ends at 18:10)
- Thursday, 14:15 – 16:15, in T2.1 (starts at 14:30 end at 16:10)

Please drop us an email if you plan to attend the other course, as this needs to be authorized

Check the detailed schedule (including Lecture Topics) as due to calendar issues, the two courses are not perfectly aligned!

CSE: https://boracchi.faculty.polimi.it/teaching/AN2DLCalendar_CS.htm

BIO+MTM: https://boracchi.faculty.polimi.it/teaching/AN2DLCalendar.htm

Course Evaluation AN2DL

Grading comprises a theoretical part and a practical part:

- Written examination covering the whole program up to 20/30 ✚
- Home project in the form of 2 coding challenges up to 10/30 =
- Final score will be the sum of the grades of the two 30/30

Challenges are graded based on what you do, not based on the position in the rank!

Written Examination

- Digital exam on moodle: bring your own laptop
- We will use the platform: https://remoteexam.polimi.it/
- Safe Exam Browser (SEB) will be required
  - It does not run on Linux…. Sorry for that… make sure you can borrow a Windows or Mac laptop

Please, make sure you can run the test quiz well ahead the exam.

Go to https://remoteexam.polimi.it/

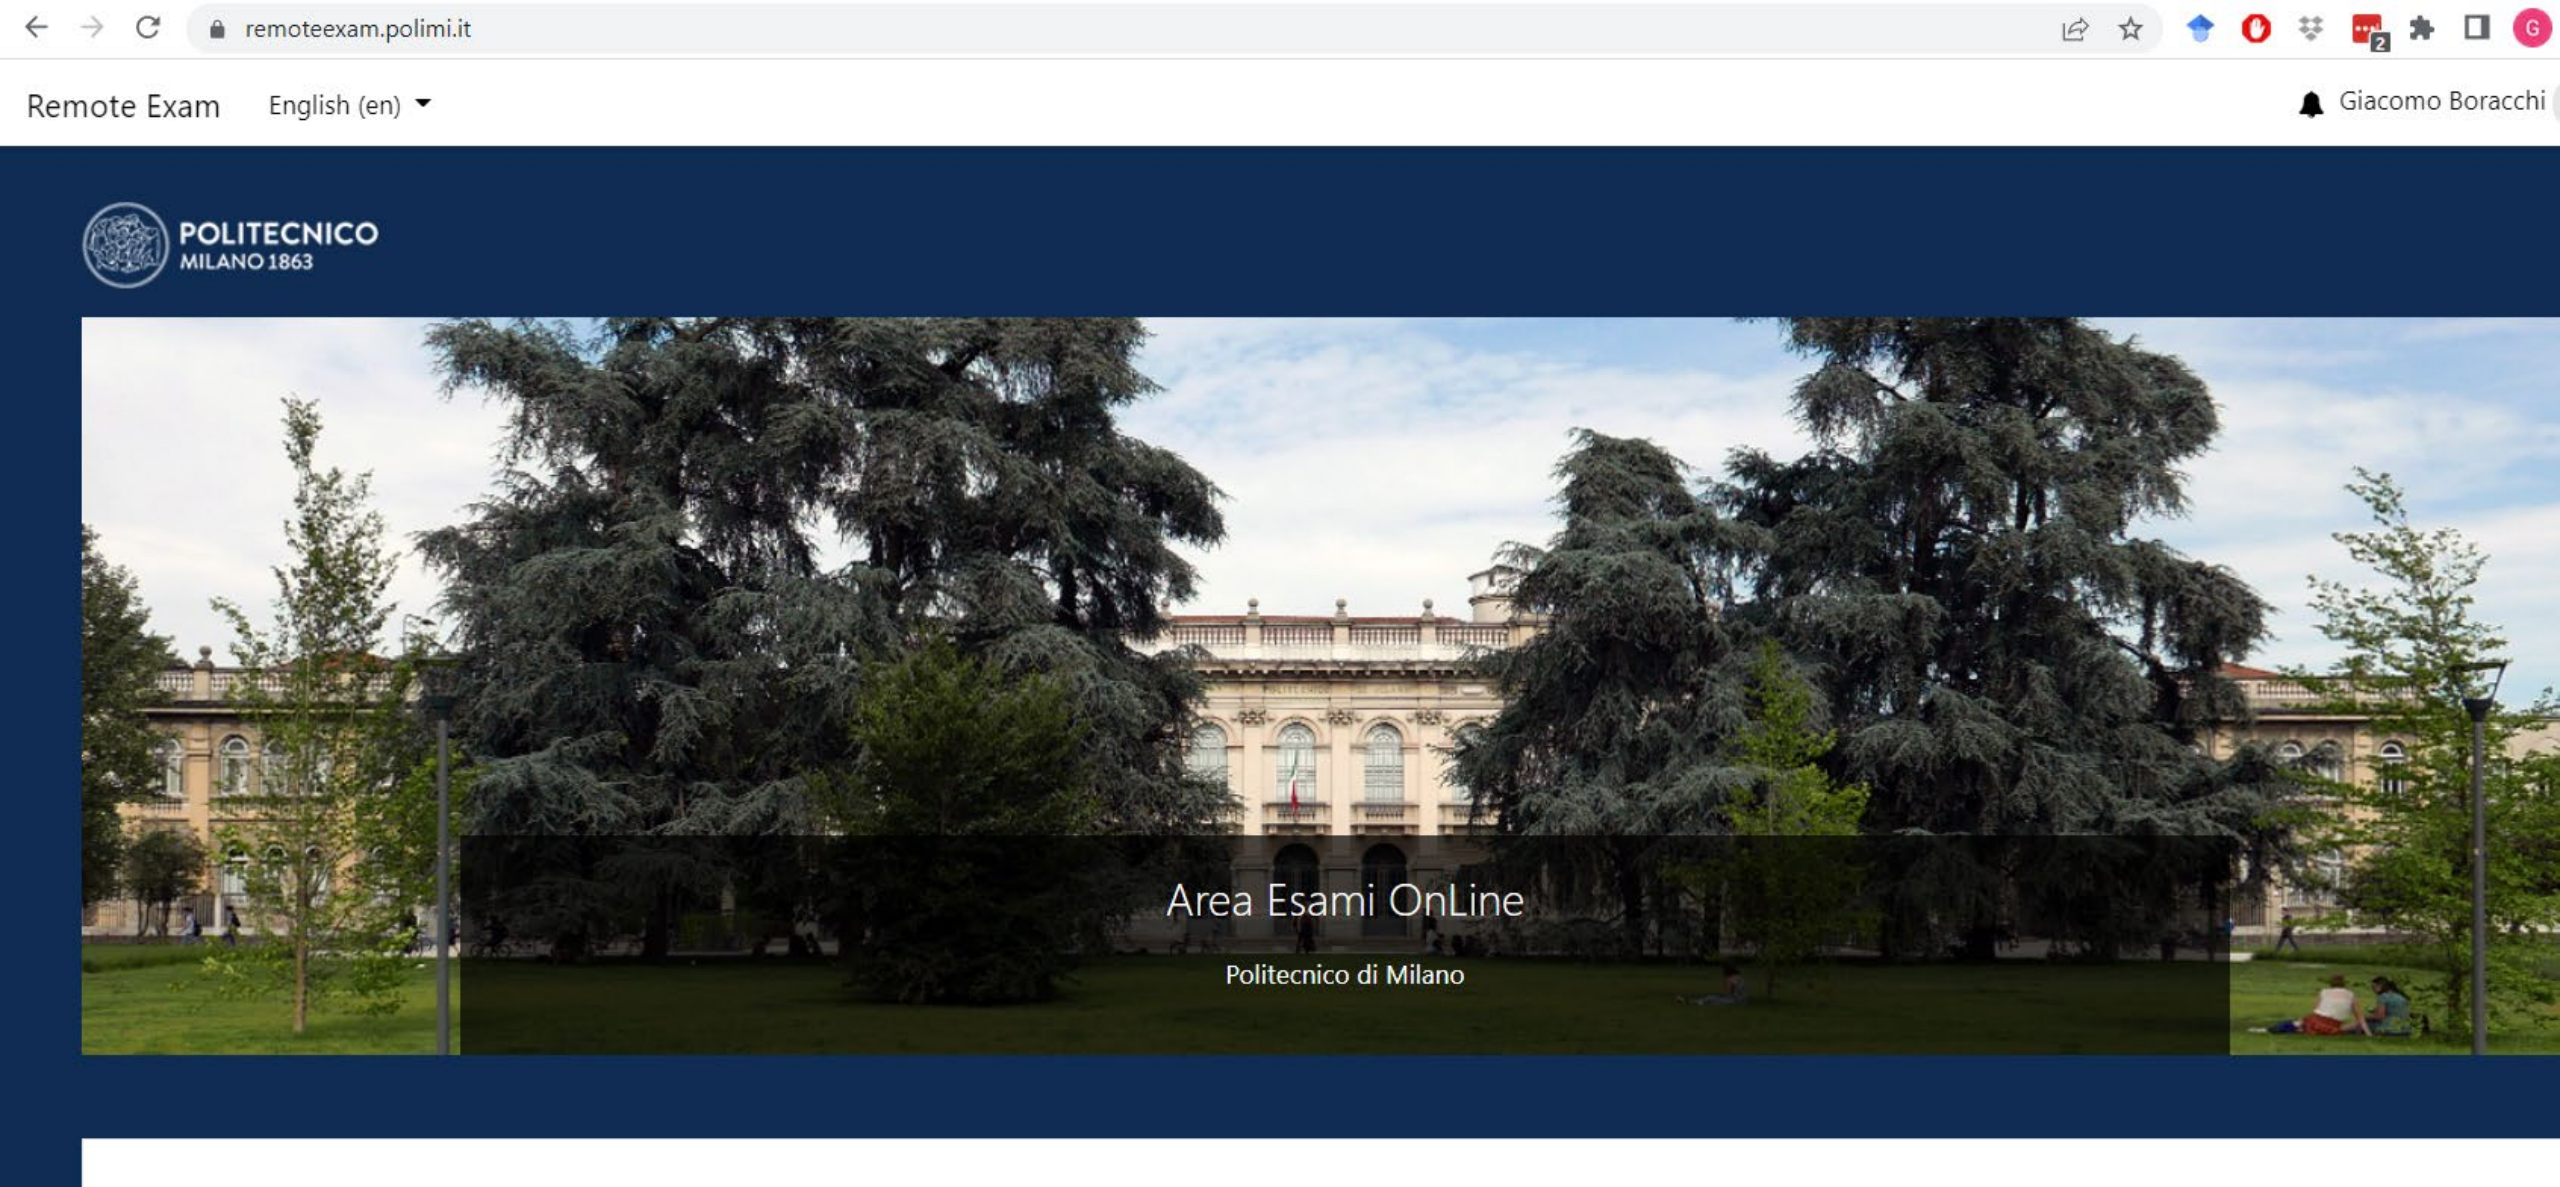

Search for our course (that will be updated…)

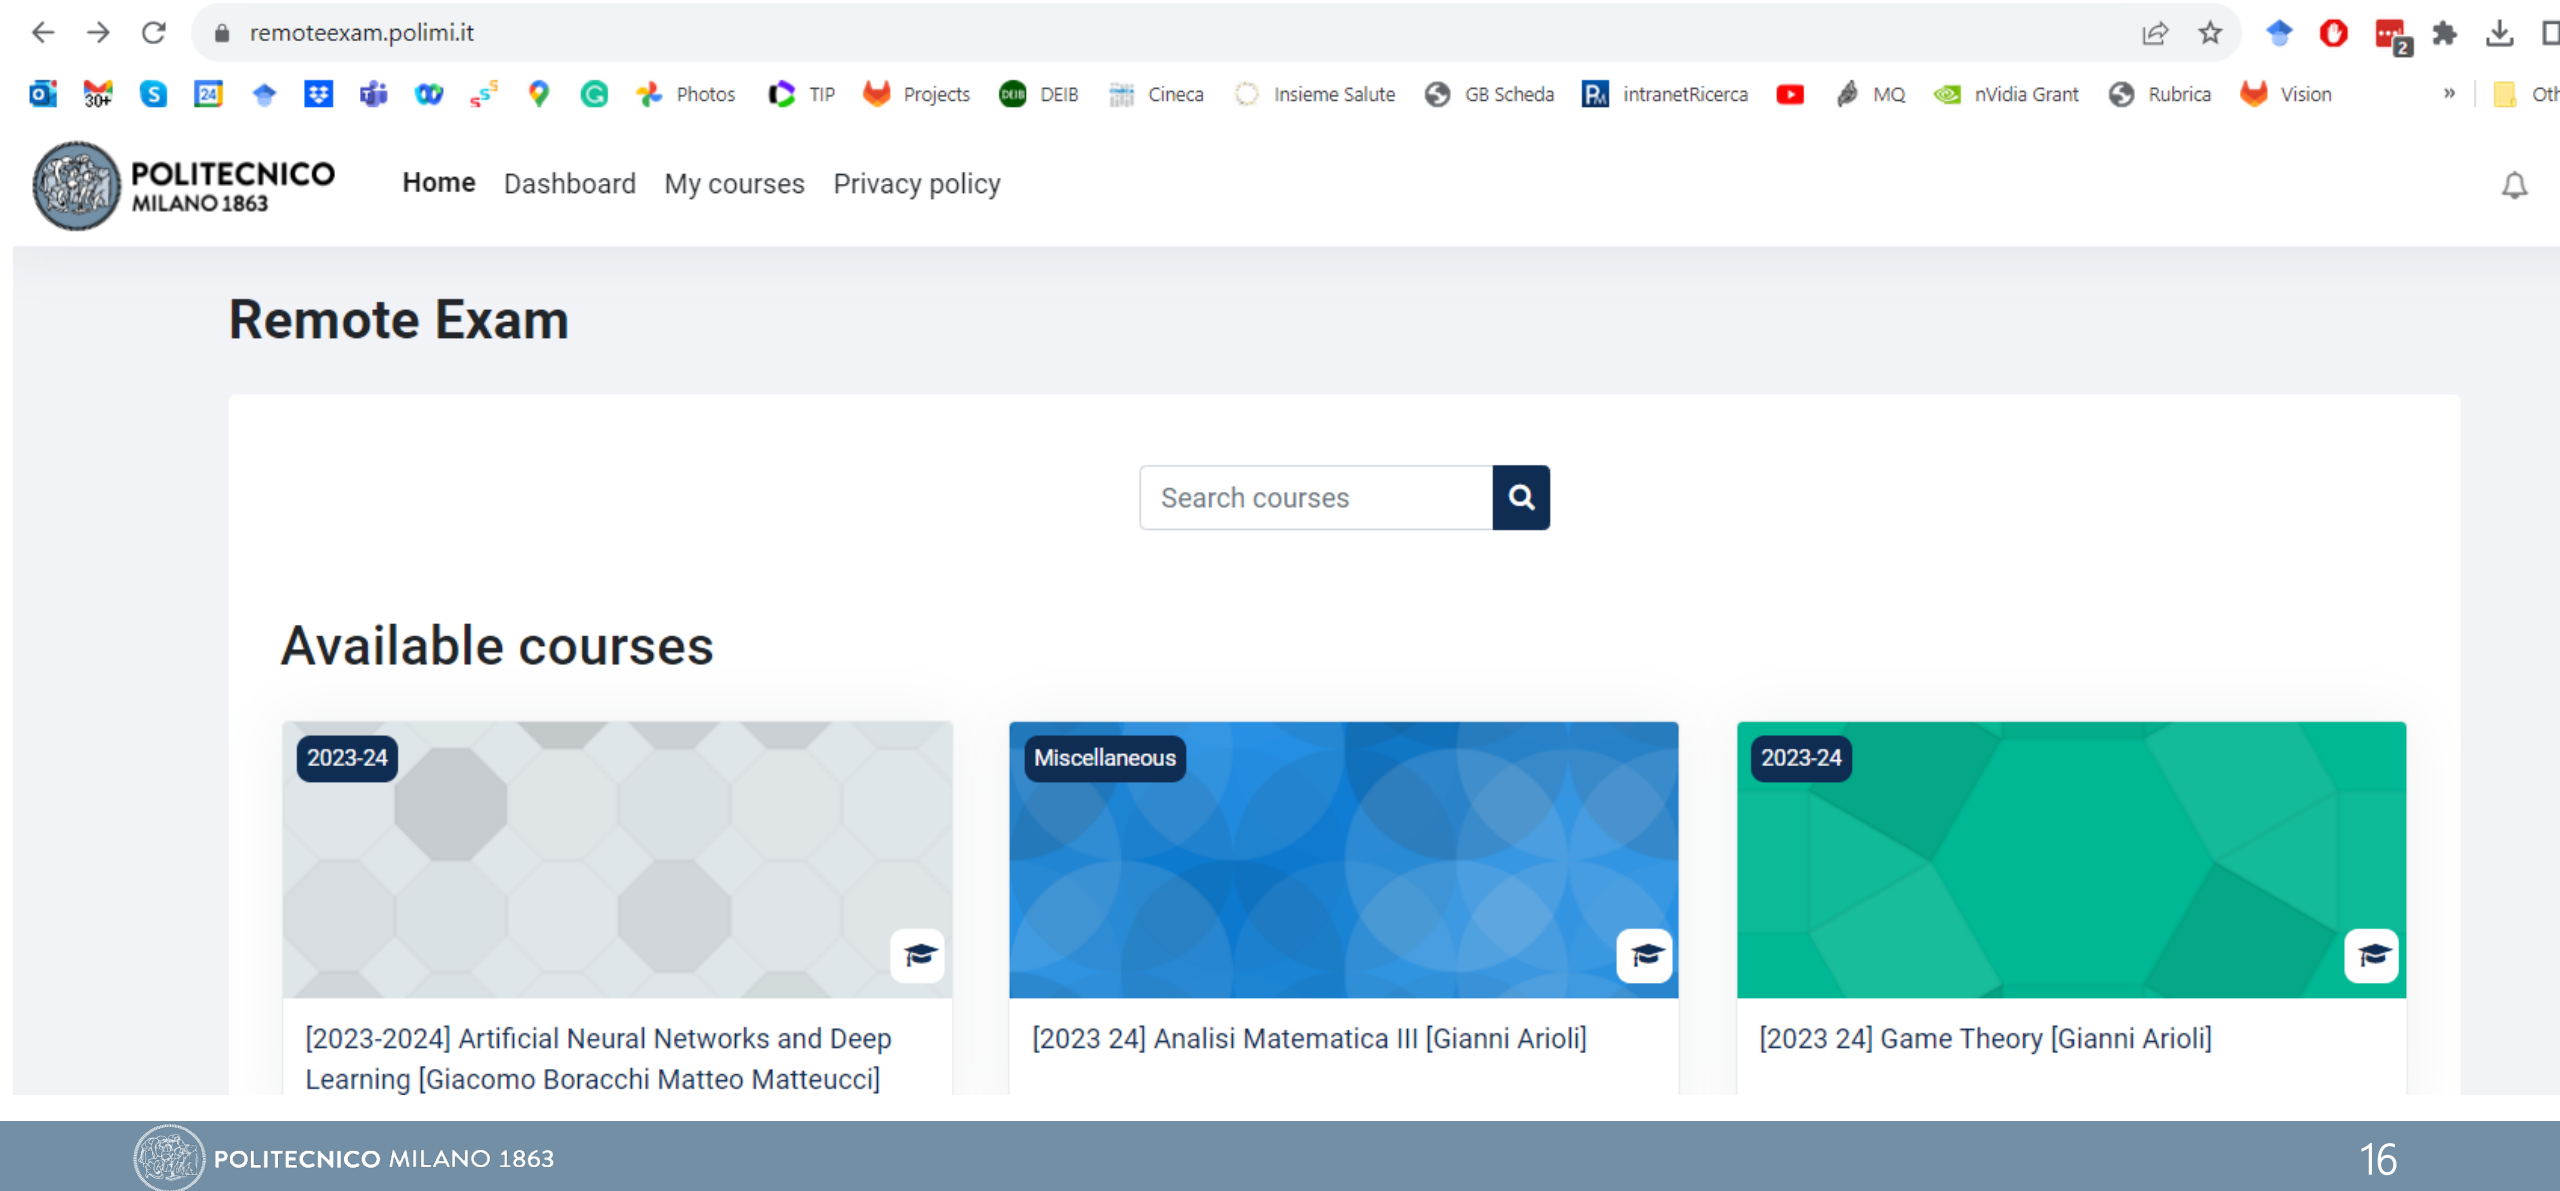

Run the test (it's already there)

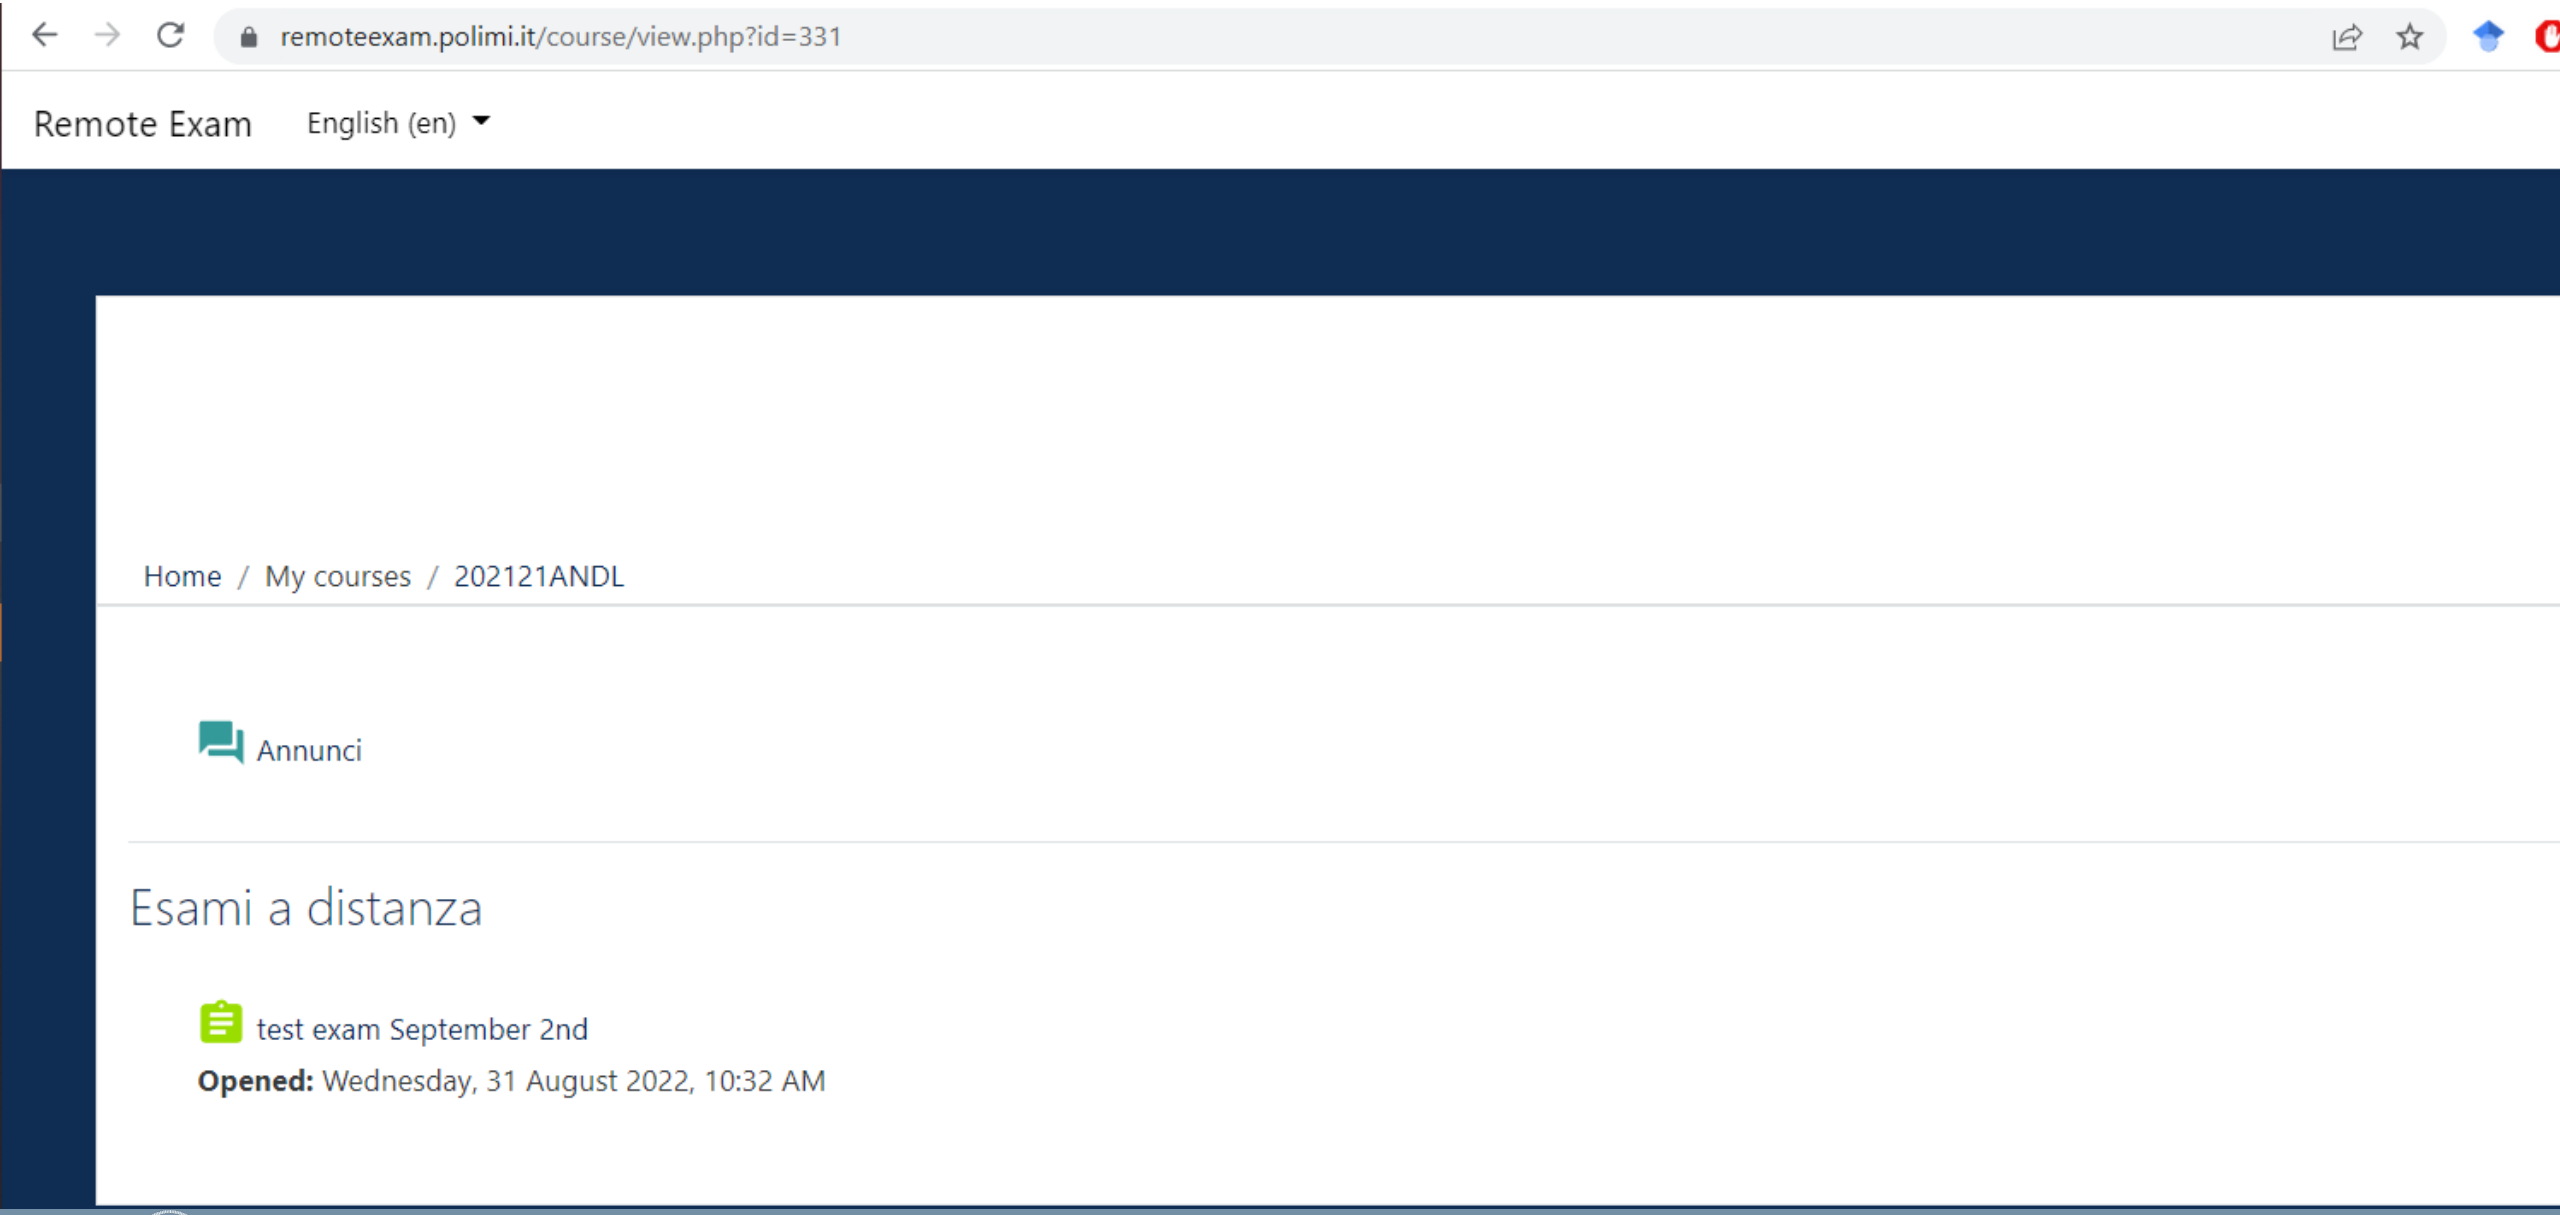

Run the test (but read instructions)!

The test is configured exactly as the exam form. Please, make sure you can successfully accomplish the following tasks with your laptop at your earliest convenience:

- Go to the https://remoteexam.polimi.it/ website and select the "[2023-2024] Artificial Neural Networks and Deep Learning [Giacomo Boracchi Matteo Matteucci]" course
- **Log in with your student credentials** and select the "test exam for AY 23-24 sessions".
- In case you do not have SEB installed, you will have the option to **install SEB**. Make sure SEB is installed before the exam.
- Make sure everything goes smoothly and you can fill in the quiz.
- **After submitting the answers**, you can quit the SEB session. You will be prompted a quit password, use **"IAmDone".**

If you cannot access the form through the SEB, you won't have the chance to give the exam.

Course Evaluation!

Grading comprises a theoretical part and a practical part:

- Written examination covering the whole program up to 20/30 ✚
- Home project in the form of 2 coding challenges up to 10/30 =
- Final score will be the sum of the grades of the two 30/30

Comments and notes about the grading

- 10 points of the theoretical part will be given by Prof. Matteucci
- 10 points of the theoretical part will be given by Prof Boracchi
- 5 points for each homework challenge are given by Francesco Lattari
- Homework challenges are not repeated, they are just run once a year
- Challenge 1 around 2nd November, Challenge 2 around 6th December

Challenges are graded based on what you do, not based on the position in the rank!

Course Evaluation!

Gradi

-
-
-

Comm

-
-
-
-
-

Chall *nk!*

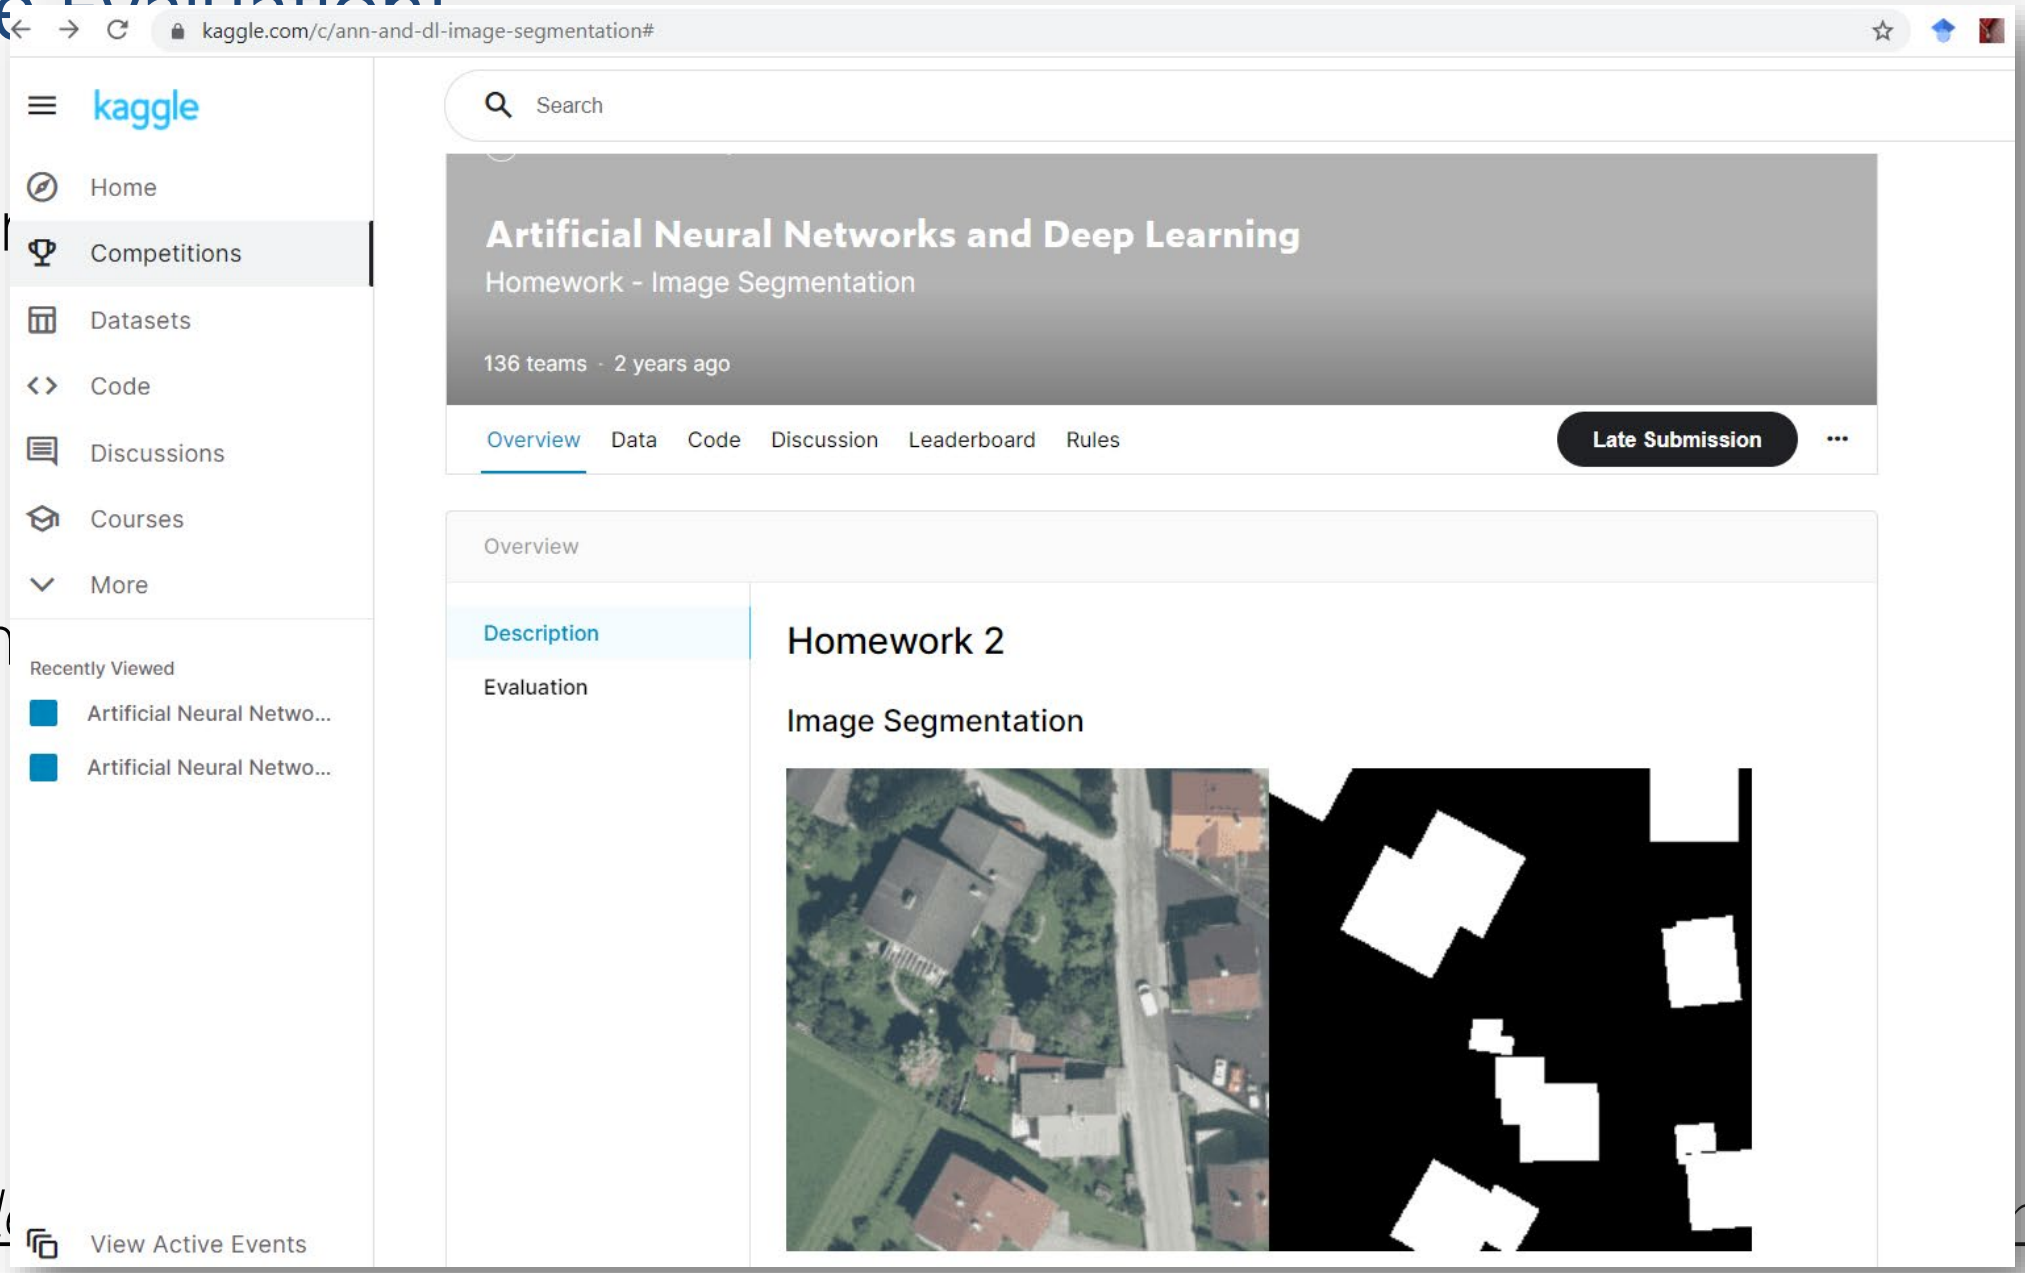

Laude

Laude is meant to reward brilliant students that:

- Actively participate to lectures
- Provide outstanding homework solutions
- Solve the written exam very timely

IC Course Evaluation

METHODS & APPLICATIONS OF AI IN BIOMEDICINE [I.C. 10 CFUs]

The final grade will be the average of

- Applied AI in Biomedicine (5 CFUs from Prof. Valentina Corino)
- AN2DL (5 CFUs).

Synergies with Other Courses

AN2DL is a course on machine learnin… courses on the same topic, but it has b… …ith:

Even taking them all the overlap ends up to be at most 10h (<20%)

- *Machine Learnig*: there you see classical machine learning tools, some concepts such as generalization, overfitting, and crossvalidation might be similar …

Up to 4-5h out of 50h (< 10%)

- *Uncertainty in Artificial Intelligence*: neural networks have been removed from this course and they have been replaced by Bayesian Networks and Graphical Models …

0h out of 50h (0%)

- *Image Analysis and Computer Vision*: Image classification part has been removed… there is a shared background on image filtering…

0h out of 50h (0%)

- *Data Mining and text Mining*: does not cover neural networks and it is mostly based on unsupervised methods

up to 4h out of 50h (< 10%)

Ironing out the kinks …

Some details have not been sorted out yet today, working on those ..

- WeBeep Use
  - No we use the calendar and enrolled students emails
- Projects/Competitions:
  - How many people per group (2-3 people)
  - Competitions out 2nd November & 6th December
- Practical evaluation of challenges:
  - Not doing it scores up to 0 points
  - Doing it with basic tools present in class up to 1-4 points (?)
  - Doing it with passion and in a propositive manner up to 5 points (?)
  - Automated scoring / code plagiarism check (?)

Frequently Asked Question (up to now)

I cannot attend all classes, do you follow a book?

You can find all covered topics on the Deep Learning book, but we are going to present the course in a personalized manner. We suggest you to attend and follow our material then check the book to complete your preparation. Slides will be made available as well as lecture recordings.

We are not computer scientist, will we be able to do the competition?

We are going to use simple libraries, we expect with basic competencies in programming you should be able to do it autonomously at least to a minimum level.

Are you going to stream/record lectures?

We are going to record and share links on the Google Calendar. No lecture streaming, though.

Frequently Asked Question (up to now)

I have overlaps can I attend AN2DL with CS?

Sure, that's fine by us. However, please inform us so that we can keep track of how many students are going to attend

- Wednesday from 16.30 till 18.00
- Thursday from 14.30 till 16.00

Other questions?

https://boracchi.faculty.polimi.it/teaching/AN2DLCalendar_CS.htm

Similar Calendar for AN2DL for CSE students…

You might want to check thi out in case of lecture overlaps

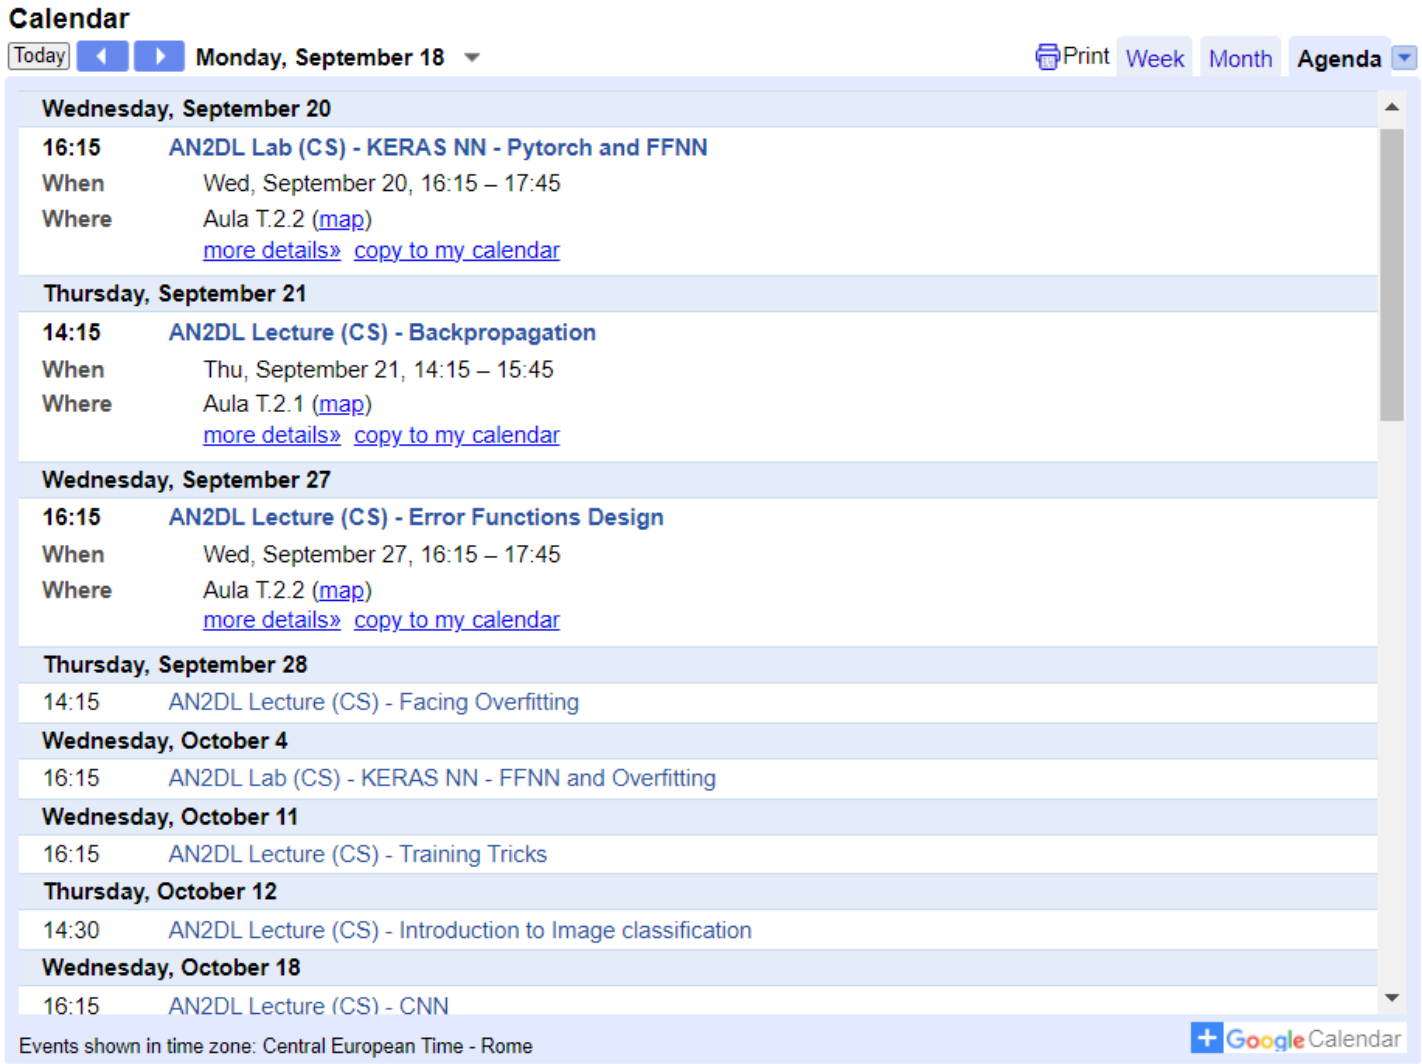

POLITECNICO
MILANO 1863

Artificial Neural Networks and Deep Learning

- Machine Learning vs Deep Learning-

Giacomo Boracchi, PhD
https://boracchi.faculty.polimi.it/
Politecnico di Milano

AIRLAB
ARTIFICIAL INTELLIGENCE AND ROBOTICS LAB

Standard Programming

```c
/* What is this program about?*/
# include<stdio.h>
int main()
{
   int a, sum;
   sum = 0;
   printf("\nInsert a:");
   scanf("%d", &a);
   while (a > 0)
   {
      sum += a;
      printf("\nInsert a:");
      scanf("%d", &a);
   }
   printf("\nSum = %d", sum);
}
```

Standard Programming

```
/* What is this program about?*/
# include<stdio.h>
int main()
{
   int a, sum;
   sum = 0;
   printf("\nInsert a:");
   scanf("%d", &a);
   while (a > 0)
   {
      sum += a;
      printf("\nInsert a:");
      scanf("%d", &a);
   }
   printf("\nSum = %d", sum);
}
```

Can you write a program that takes as input an image and tells whether it contains a car or a motorbike?

Machine Learning Paradigms

ML is the solution, as the program **becomes a very big parameteric function** f_θ, whos paramters θ are learned from data!

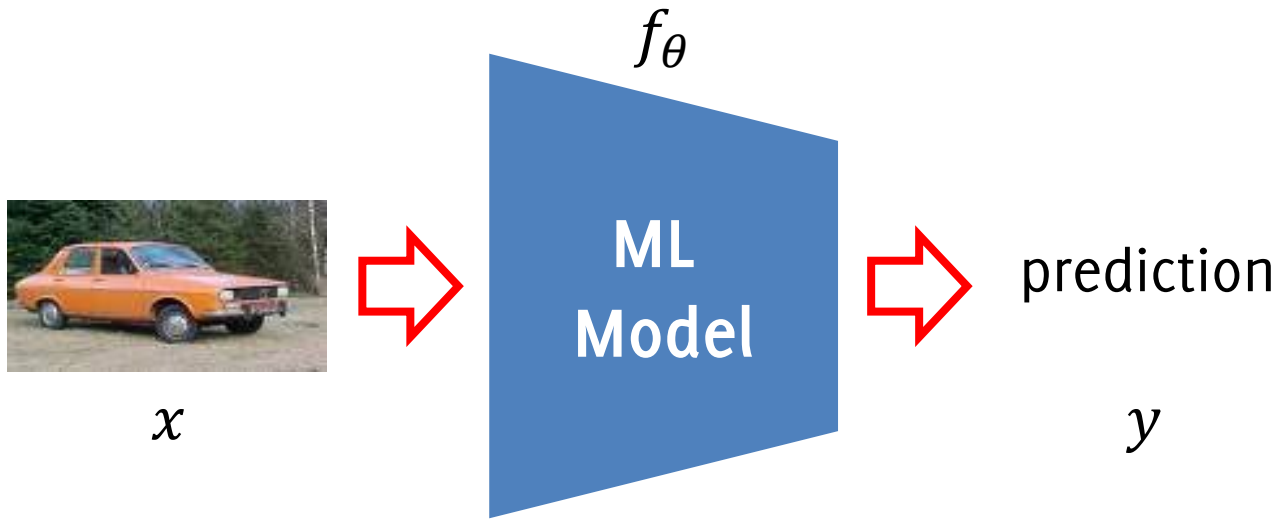

x ⇨ $f_\theta(x) = y$

Machine Learning Paradigms

ML is the solution, as the program becomes a very big parameteric function f_θ, whos paramters θ are learned from data!

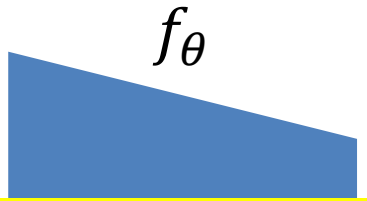

Learning consists is (automatically) defining the parameters θ of the model f.
Different settings applies, which give rise to the supervised and unsupervised settings

prediction

y

x ⇨ $f_\theta(x) = y$

Machine Learning Paradigms

ML is the solution, as **the program becomes a very big parameteric function** f_θ, whos paramters θ are learned from data!

Supervised Learning

- Classification
- Regression

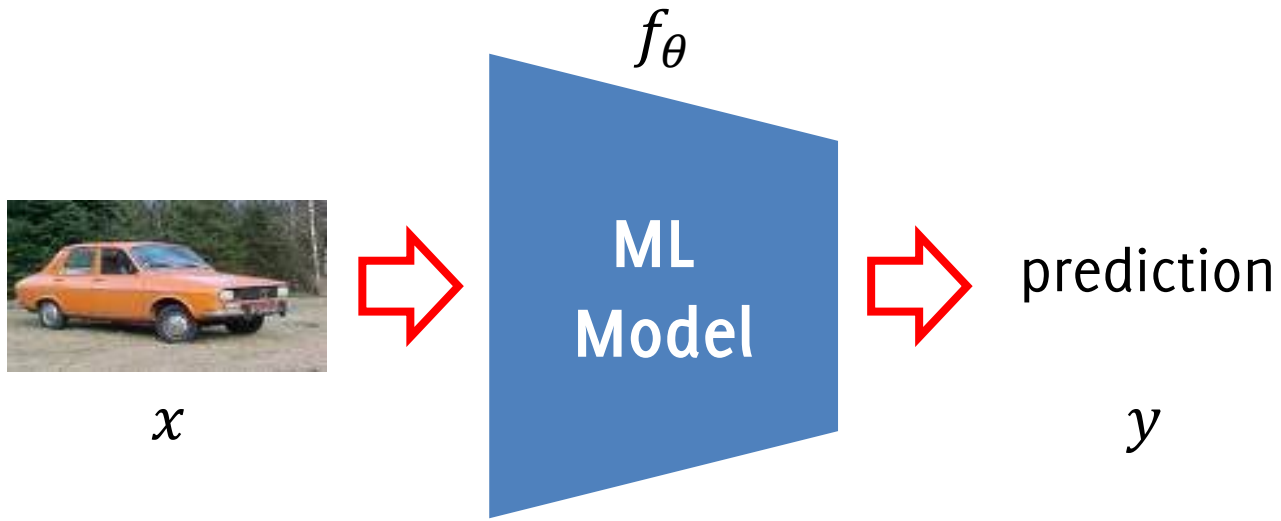

$x \Rightarrow f_\theta(x) = y$

Supervised Learning

In **Supervised Learning** we are given a training in the form:

$$TR = \{(x_1, y_1), \dots, (x_n, y_n)\}$$

where

- $x_i \in \mathbb{R}^d$ is the input
- $y_i \in \Lambda$ is the target, the expected output of the model to x_i

The set Λ can be

- A discrete set, as in classification $\Lambda = \{$"brown", "green", "blue"$\}$ (e.g., possible eye colors)
- An ordinal set (often continuous set, $\mathbb{R}$) in case of regression.

Λ can be also multivariate (e.g., regressing weight and height of an individual or estimating they eye colors and heirs color)

Training Set for (binary) Image Classification

Cars

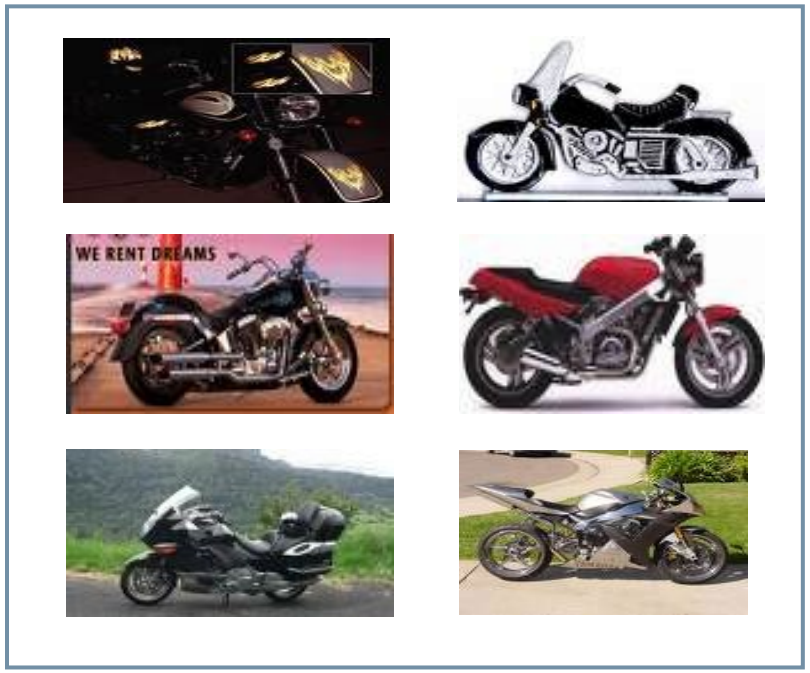

Motorcycles

$$TR = \{(x_1, y_1), \ldots, (x_n, y_n)\}$$

- $x_i \in \mathbb{R}^{R \times C \times 3}$ is the input image
- $y_i \in \{\text{"car"}, \text{"motorcycle"}\}$

Inference Using the Trained Classifier

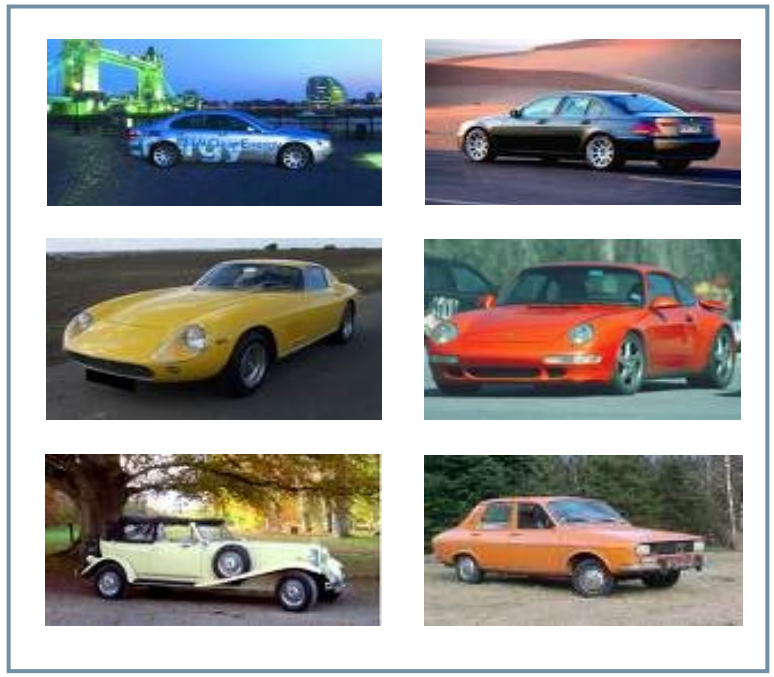

Cars

Motorcycles

Classifier

Motorcycle

Supervised learning: Regression

12000 $

15000 $

6000 $

2000 $

8000 $

22000 $

4000 $

28000 $

6000 $

35000 $

Regressor

3800 $

Training Set for Regression


12000 $


15000 $

6000 $


2000 $


8000 $


22000 $


4000 $


28000 $


6000 $


35000 $

$$TR = \{(x_1, y_1), \ldots, (x_n, y_n)\}$$

- $x_i \in \mathbb{R}^{R \times C \times 3}$ is the input image
- $y_i \in \mathbb{R}$

Supervised learning: Regression

12000 $

15000 $

6000 $

2000 $

8000 $

22000 $

4000 $

28000 $

6000 $

35000 $

Regressor

3800 $

Remarks

- Number of classes can be larger than two (multiclass classification, e.g., {"car", "motorcycle","truck"})
- The input size in general needs to be fixed
- The number of outputs for regression can be larger (multivariate regression, e.g., estimating cost and weight of the vehicle)
- Training a Classifier or a Regressor requires different losses
- Difference between classification or regression is not only on the fact that Λ discrete, but whether it is ordinal
  - Λ categorical (no ordinal) -> classification
  - Λ ordinal (either discrete or continuous) -> regression

Give a few examples of

Regression problems on images

-
-
-
-
-

Classification problems on images

-
-
-
-
-

Machine Learning Paradigms

ML is the solution, as the program becomes a very big parameteric function f_θ, whos paramters θ are learned from data!

Supervised Learning

- Classification
- Regression

Unsupervised Learning

- Clustering
- Anomaly Detection
- …

$x \Rightarrow f_\theta(x) = y$

Unupervised Learning

In **Unsupervised Learning**, the training set contains only inputs,

$$TR = \{x_1, \dots, x_n\}$$

and the goal is to find structure in the data, like

- grouping or clustering of data points
- estimating probability density distribution
- detecting outliers
- …

Unsupervised learning: Clustering



















Unsupervised learning: Clustering









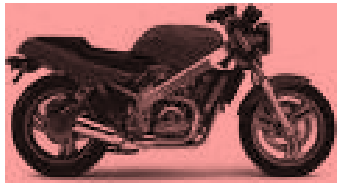









Unsupervised learning: Clustering

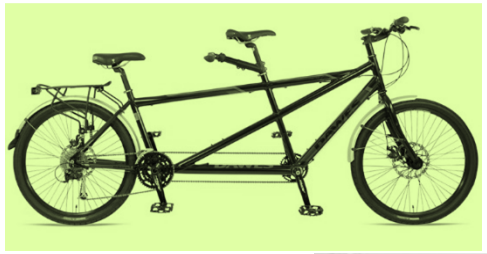


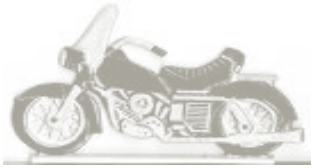








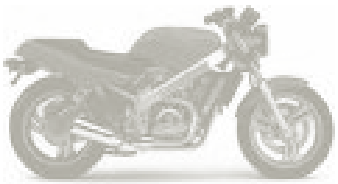






Unsupervised learning: Clustering




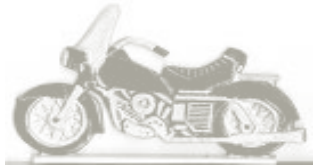




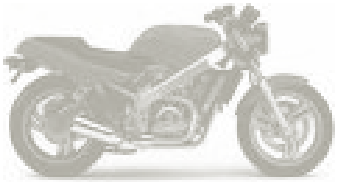









Unsupervised learning: Clustering









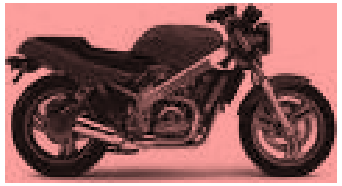









Unsupervised learning: Anomaly Detection

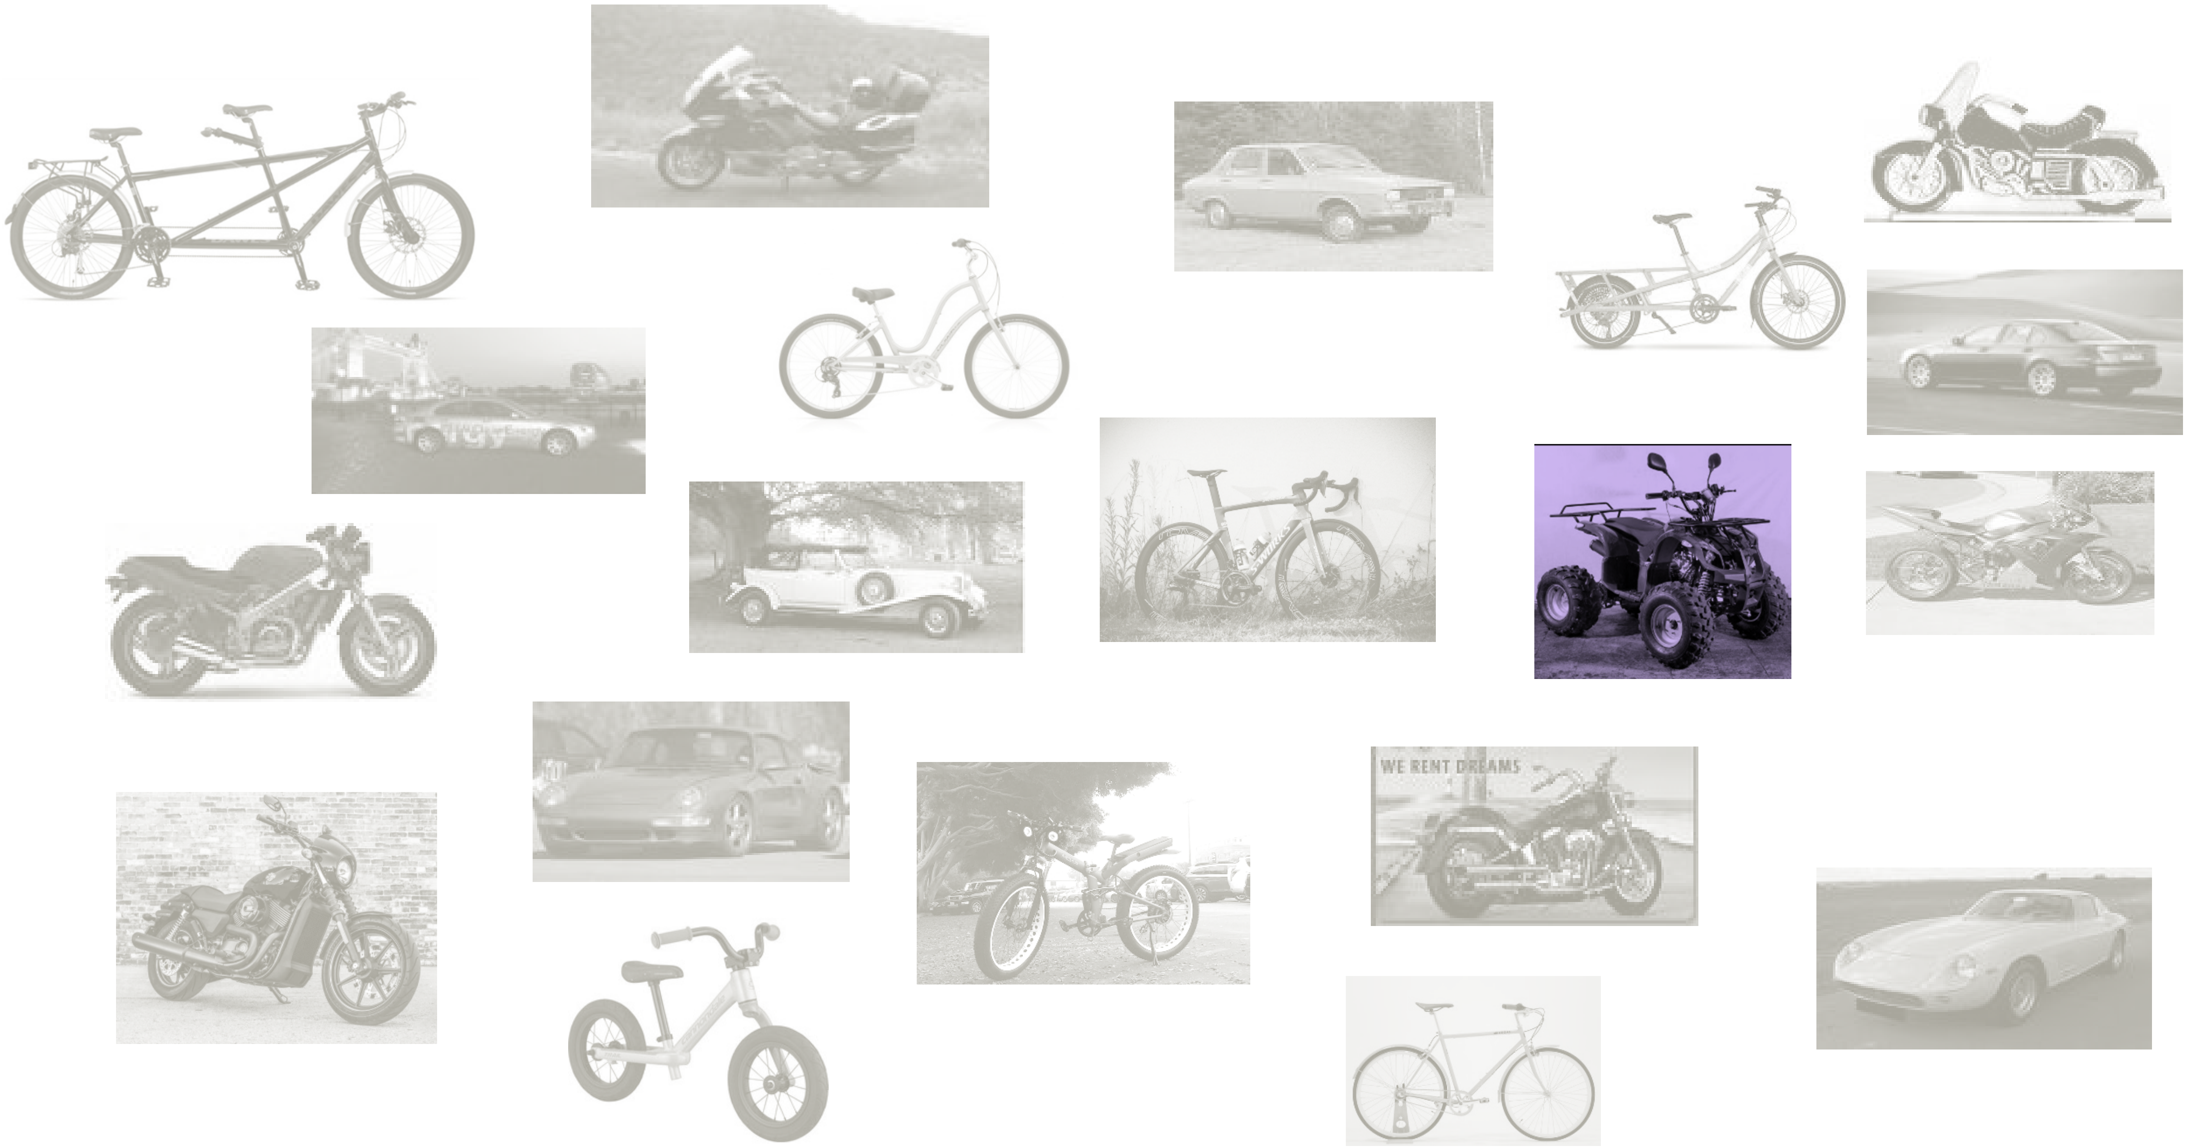

To Summarize: Machine Learning Paradigms

Immagine you have a certain experience E, i.e., data, and let's name it

$$D = x_1, x_2, x_3, \dots, x_N$$

- ***Supervised learning***: given a training set of pairs (input, desired output) $\{(x_1, y_1), \dots, (x_N, y_N)\}$, learn to produce the correct output of new inputs
- ***Unsupervised learning***: exploit regularities in D to build a meaningful/compact representation of these, which can help regression/prediction
- ***Reinforcement learning***: producing actions $a_1, a_2, a_3, \dots, a_N$ which affect the environment, and receiving rewards $r_1, r_2, r_3, \dots, r_N$ learn to act in order to maximize rewards in the long term

To Summarize: Machine Learning Paradigms

Immagine you have a certain experience E, i.e., data, and let's name it

$$D = x_1, x_2, x_3, \dots, x_N$$

- ***Supervised learning***: given a training set of pairs (input, desired output) $\{(x_1, y_1), \dots, (x_N, y_N)\}$, learn to produce the correct output of new inputs
- ***Unsupervised learning***: exploit regularities in $\boldsymbol{D}$ to build a meaningful/compact representation of these, which can help regressio
- ***Reinforcement learning***: producing actions the environment, and receiving rewards r_1, r_2 to maximize rewards in the long term

This course focuses most on Supervised Learning (with some unsupervised spots)

Machine Learning

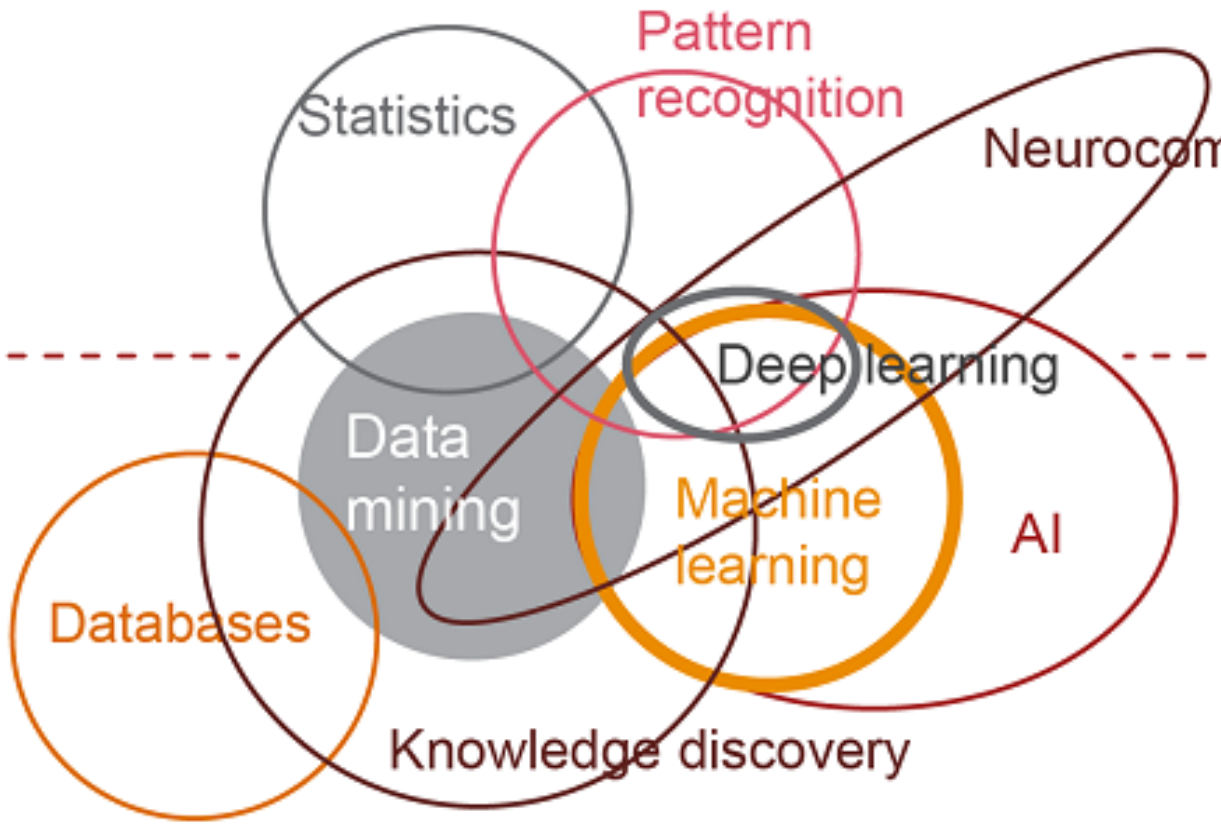

Machine learning is a category of research and algorithms focused on finding patterns in data and using those patterns to make predictions. Machine learning falls within the artificial intelligence (AI) umbrella, which in turn intersects with the broader field of knowledge discovery and data mining.

Source: SAS, 2014 and PwC, 2016

Machine Learning

MIT Technology Review

10 BREAKTHROUGH TECHNOLOGIES 2013

Deep Learning

With massive amounts of computational power, machines can now recognize objects and translate speech in real time. Artificial intelligence is finally getting smart.

Temporary Social Media

Messages that quickly self-destruct could enhance the privacy of online communications and make people freer to be spontaneous.

Prenatal DNA Sequencing

Reading the DNA of fetuses will be the next frontier of the genomic revolution. But do you really want to know about the genetic problems or musical aptitude of your unborn child?

Additive Manufacturing

Skeptical about 3-D printing? GE, the world's largest manufacturer, is on the verge of using the technology to make jet parts.

Baxter: The Blue-Collar Robot

Rodney Brooks's newest creation is easy to interact with, but the complex innovations behind the robot show just how hard it is to get along with people.

Memory Implants

Smart Watches

Ultra-Efficient Solar Power

Big Data from Cheap Phones

Supergrids

MIT Technology Review

10 BREAKTHROUGH TECHNOLOGIES 2013

Deep Learning

With massive amounts of computational power, machines can now recognize objects and translate speech in real time. Artificial intelligence is finally getting smart.

Temporary Social Media

Messages that quickly self-destruct could enhance the privacy of online communications and make people freer to be spontaneous.

Prenatal DNA Sequencing

Reading the DNA of fetuses will be the next frontier of the genomic revolution. But do you really want to know about the genetic problems or musical aptitude of your unborn child?

Additive Manufacturing

Skeptical about 3-D printing? GE, the world's largest manufacturer, is on the verge of using the technology to make jet parts.

Baxter: The Blue-Collar Robot

Rodney Brooks's newest creation is easy to interact with, but the complex innovations behind the robot show just how hard it is to get along with people.

Memory Implants

A maverick

Smart Watches

Ultra-Efficient Solar Power

Big Data from Cheap Phones

Supergrids

Hand-Crafted Features

How images / signals were classified before deep learning

Assume you need to automatize this process

Assume you need to automatize this process

Assume you need to automatize this process

An Illustrative Example: Parcel Classification

Images acquired from a RGB-D sensor:

- No color information provided
- A few pixels report depth measurements
- Images of 3 classes
  - ENVELOPE
  - PARCEL
  - DOUBLE

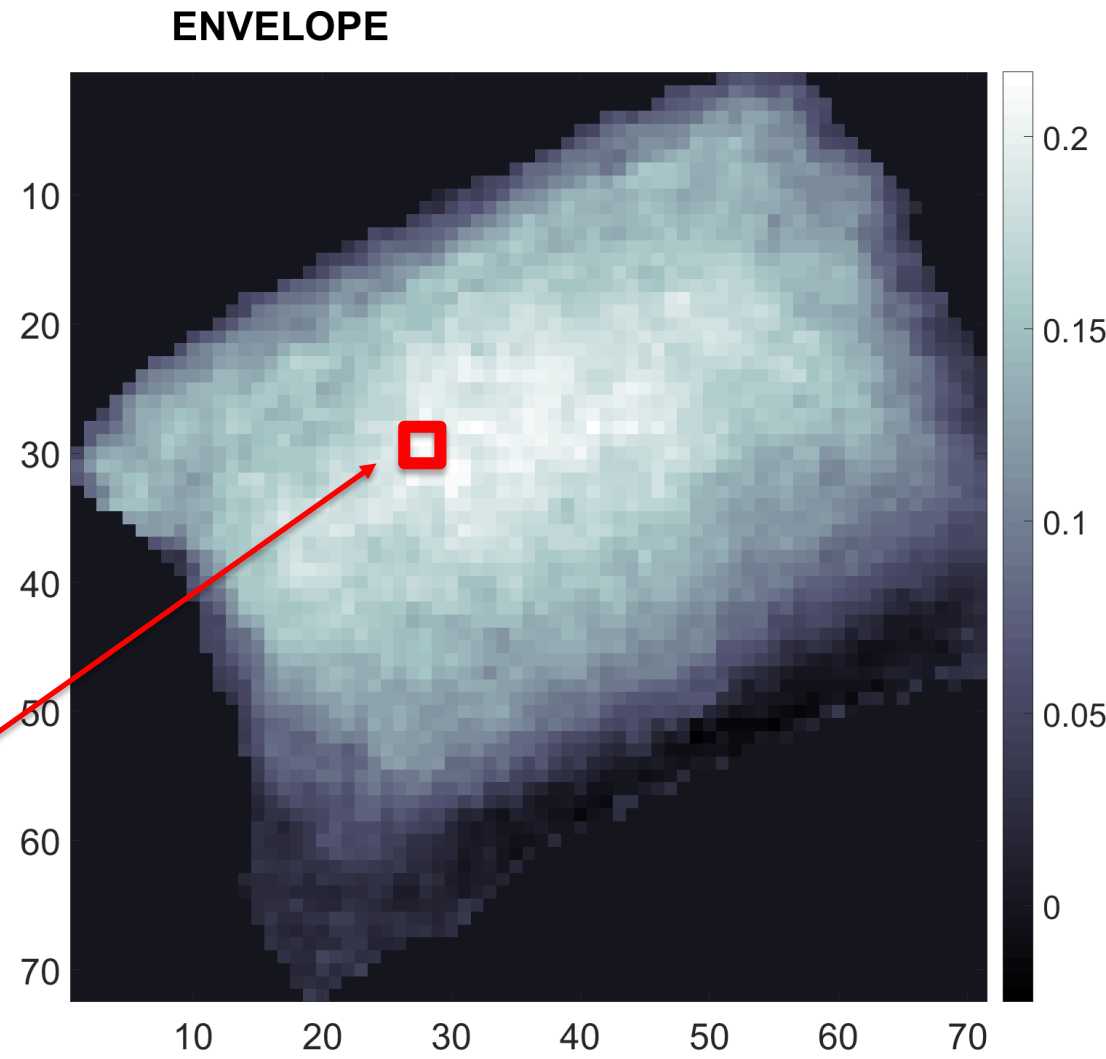

Envelop height at that pixel

An Illustrative Example: Parcel Classification

Images acquired from a RGB-D sensor:

- No color information provided
- A few pixels report depth measurements
- Images of 3 classes
  - ENVELOPE
  - PARCEL
  - DOUBLE

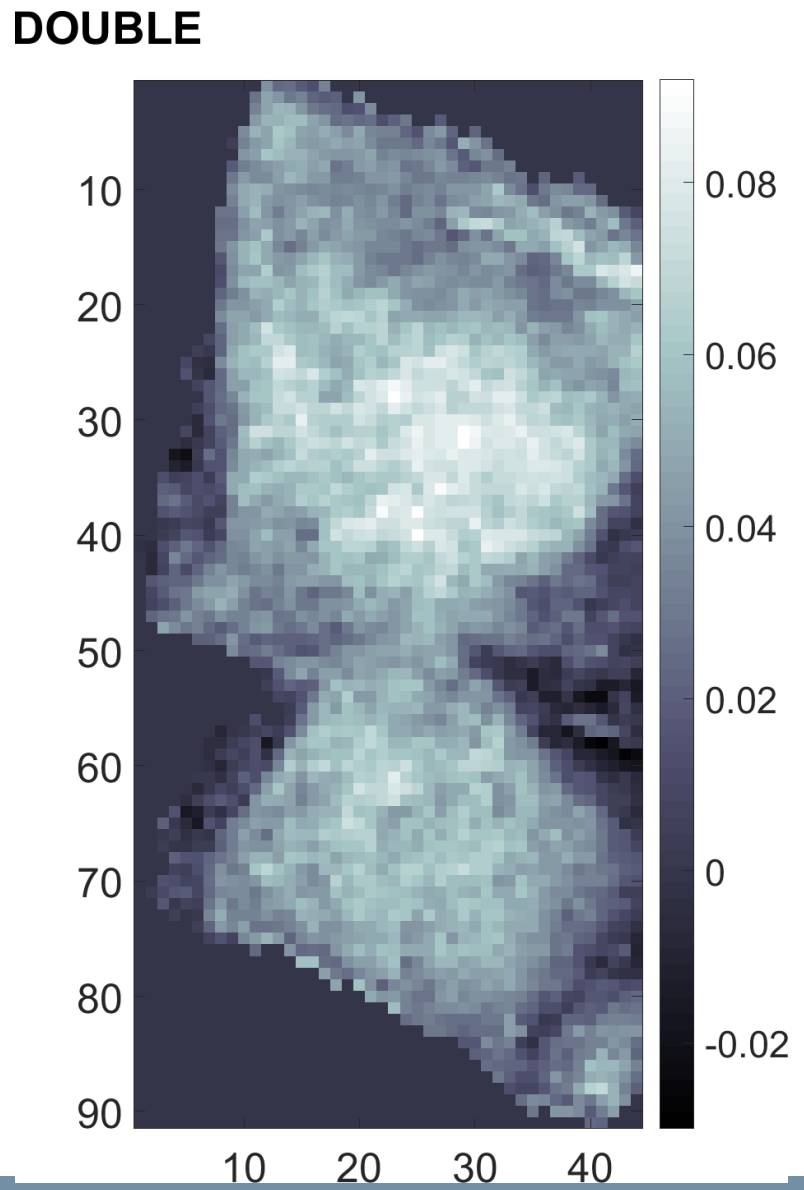

An Illustrative Example: Parcel Classification

Images acquired from a RGB-D sensor:

- No color information provided
- A few pixels report depth measurements
- Images of 3 classes
  - ENVELOPE
  - PARCEL
  - DOUBLE

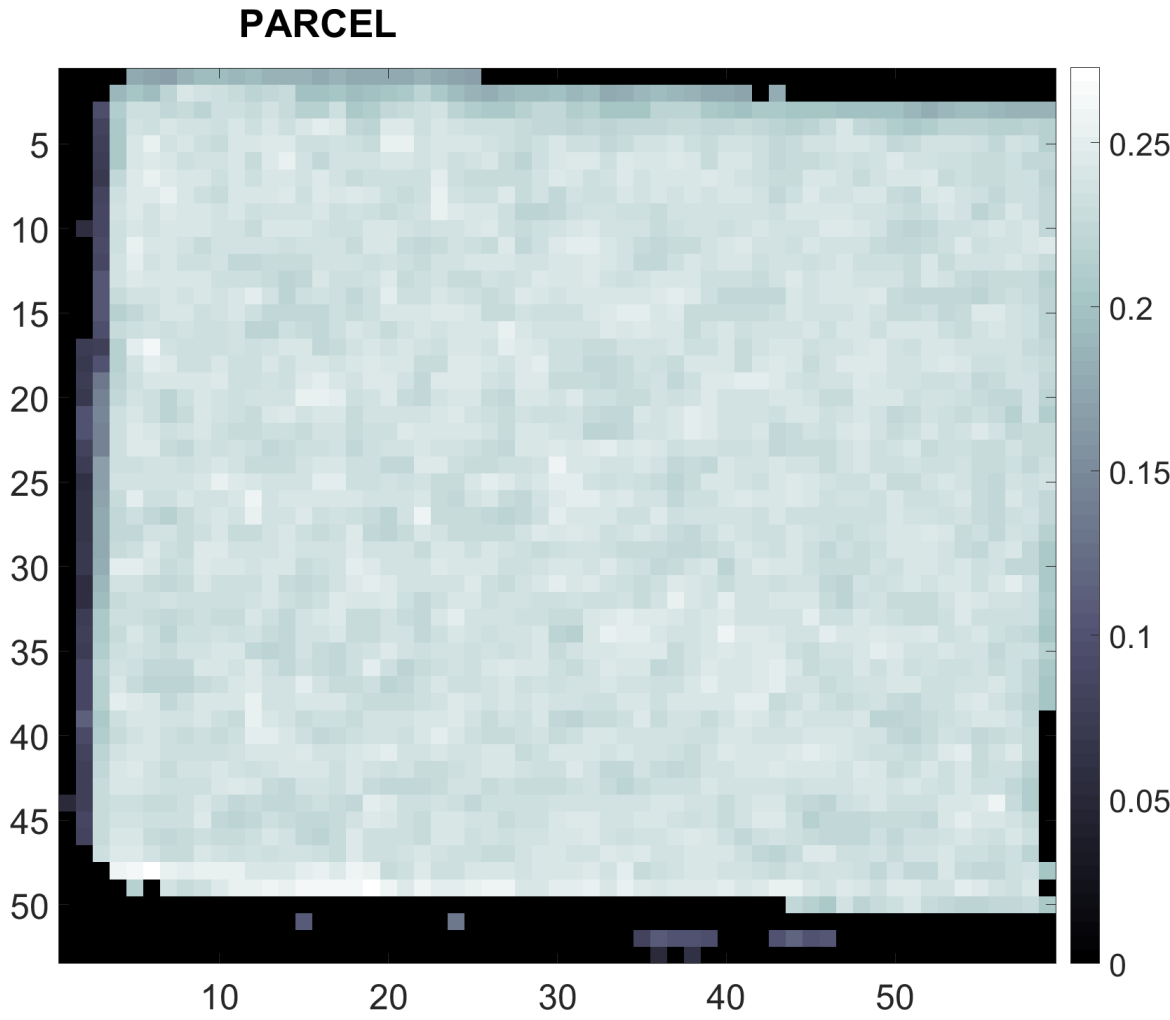

Hand Crafted Featues

Engineers:

- know what's meaningful in an image (e.g. a specific color/shape, the area, the size)
- can implement algorithms to map this information in a set of measurements, a **feature vector**

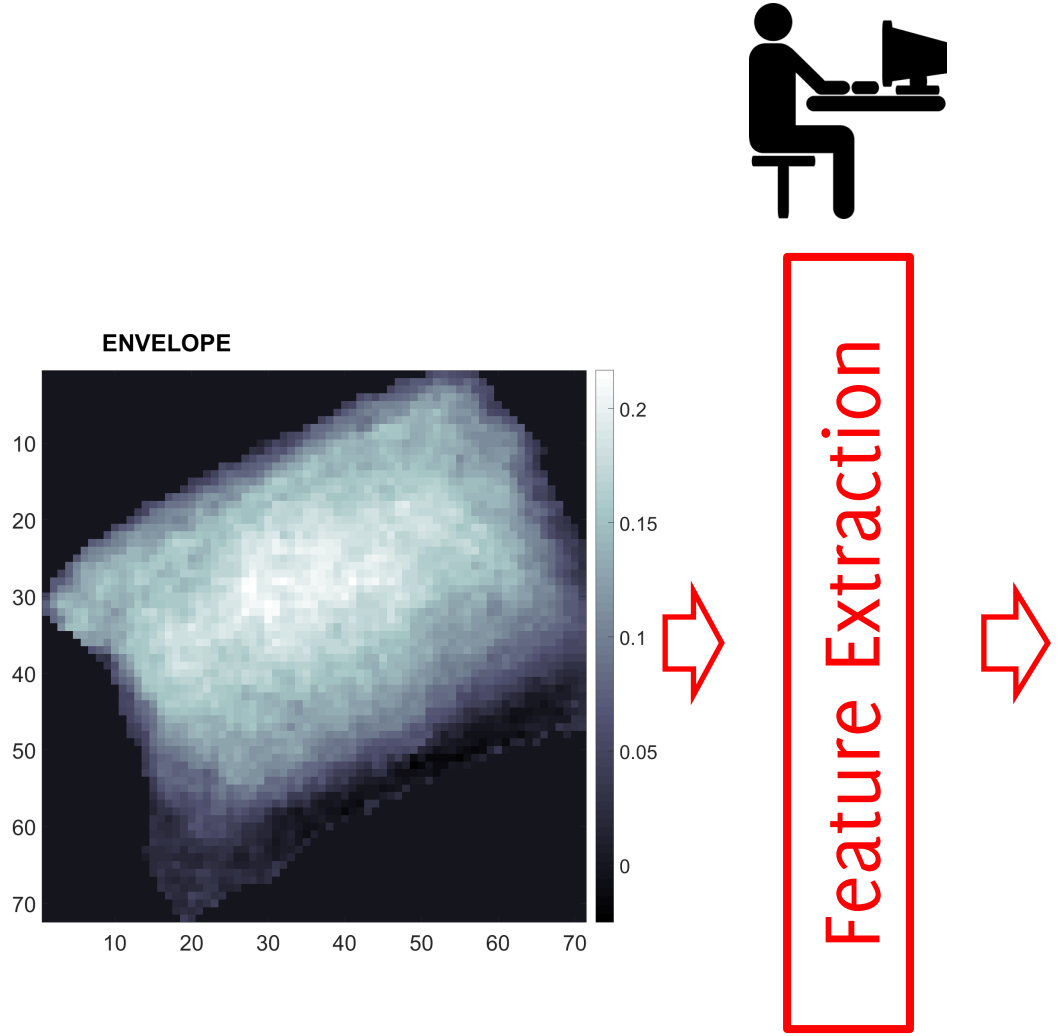

Hand Crafted Featues

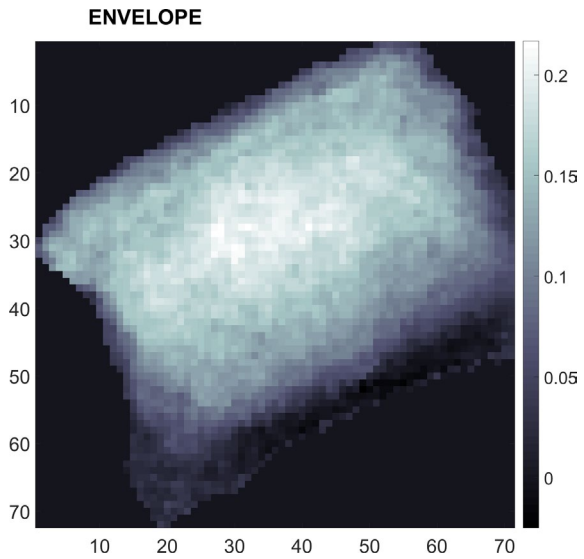

Feature Extraction

h average

area

h max

h min

perimeter

ratio

$\mathbf{x} \in \mathbb{R}^d$

This is exactly what a doctor would to to classify ECG tracings

Heartbeats morphology has been widely investigated

Doctors know which patterns are meaningful for classifying each beat

Features are extracted from landmarks indicated by doctors:

e.g. QT distance, RR distance…

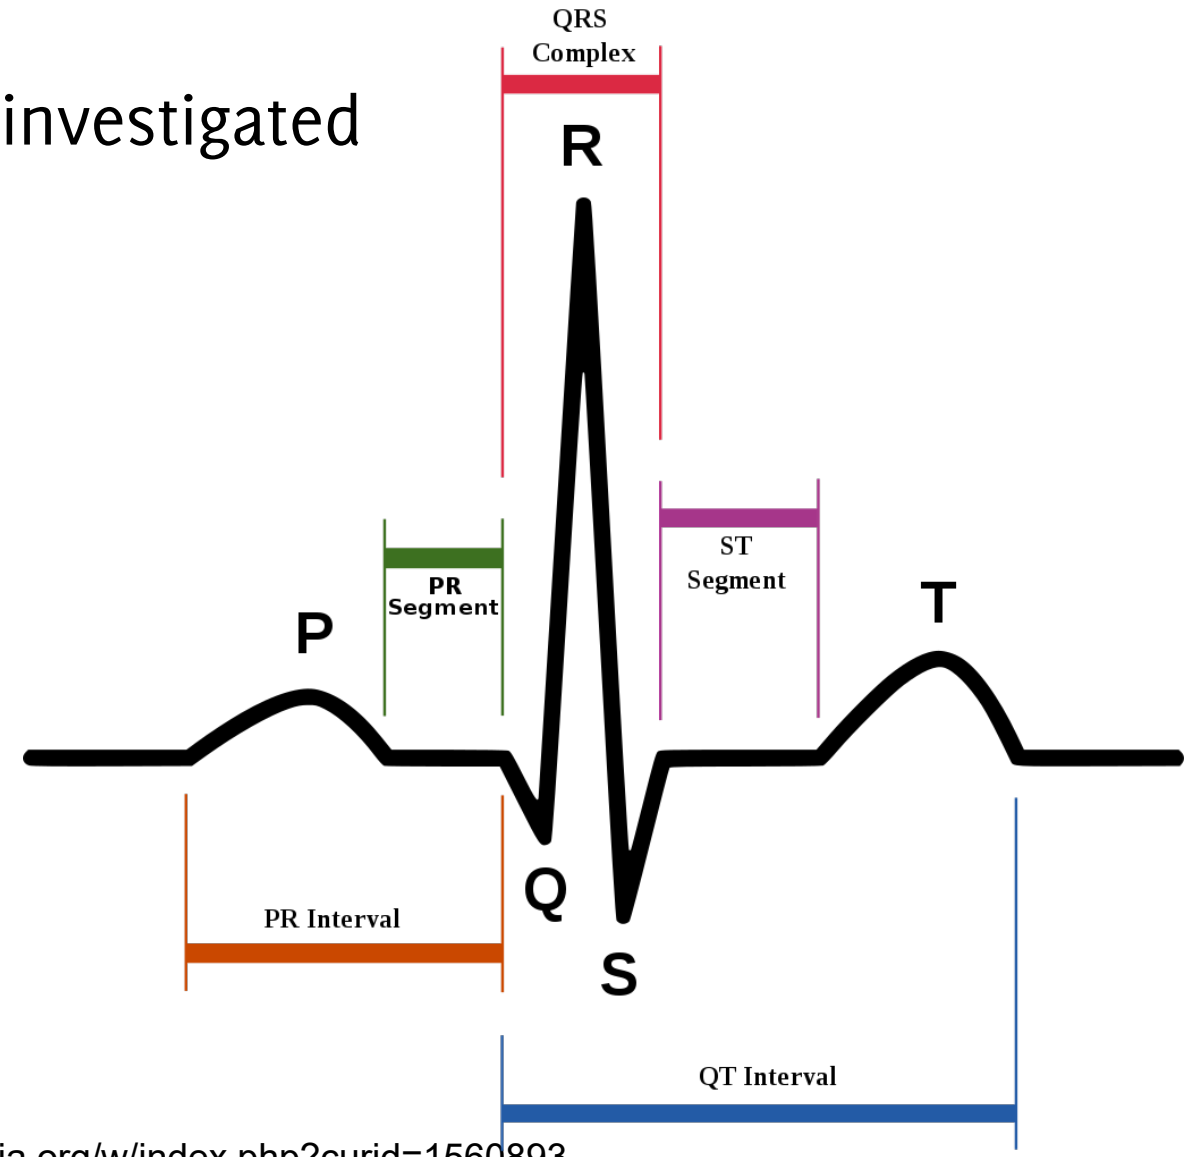

Created by Agateller (Anthony Atkielski), Public Domain, https://commons.wikimedia.org/w/index.php?curid=1560893

The Training Set

The training set is a set of annotated examples

$$TR = \{(\boldsymbol{x}, \boldsymbol{y})_i, i = 1, \dots, N\}$$

Each couple $(\boldsymbol{x}, \boldsymbol{y})_i$ corresponds to:

- an image $\boldsymbol{x}_i$
- the corresponding label $\boldsymbol{y}_i$

This is meant for a **Supervised** Learning Problem!

The Training Set: images + labels

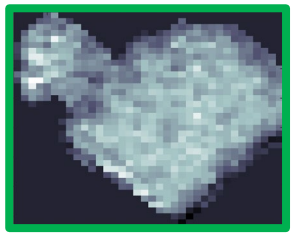

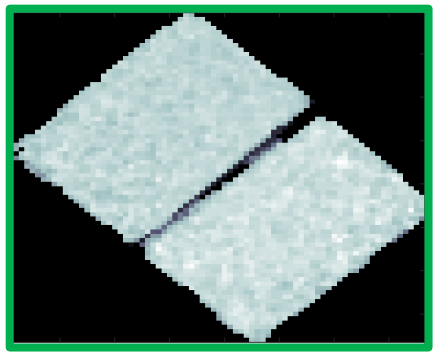

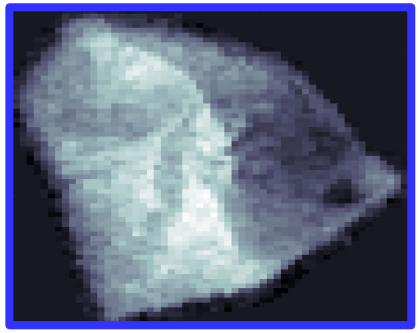

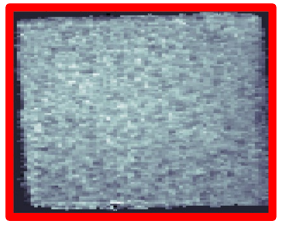

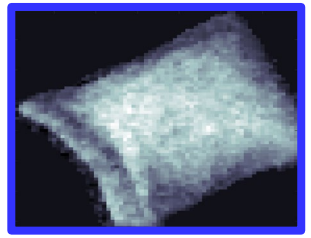

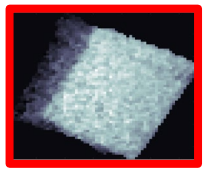

The Training Set: images + labels

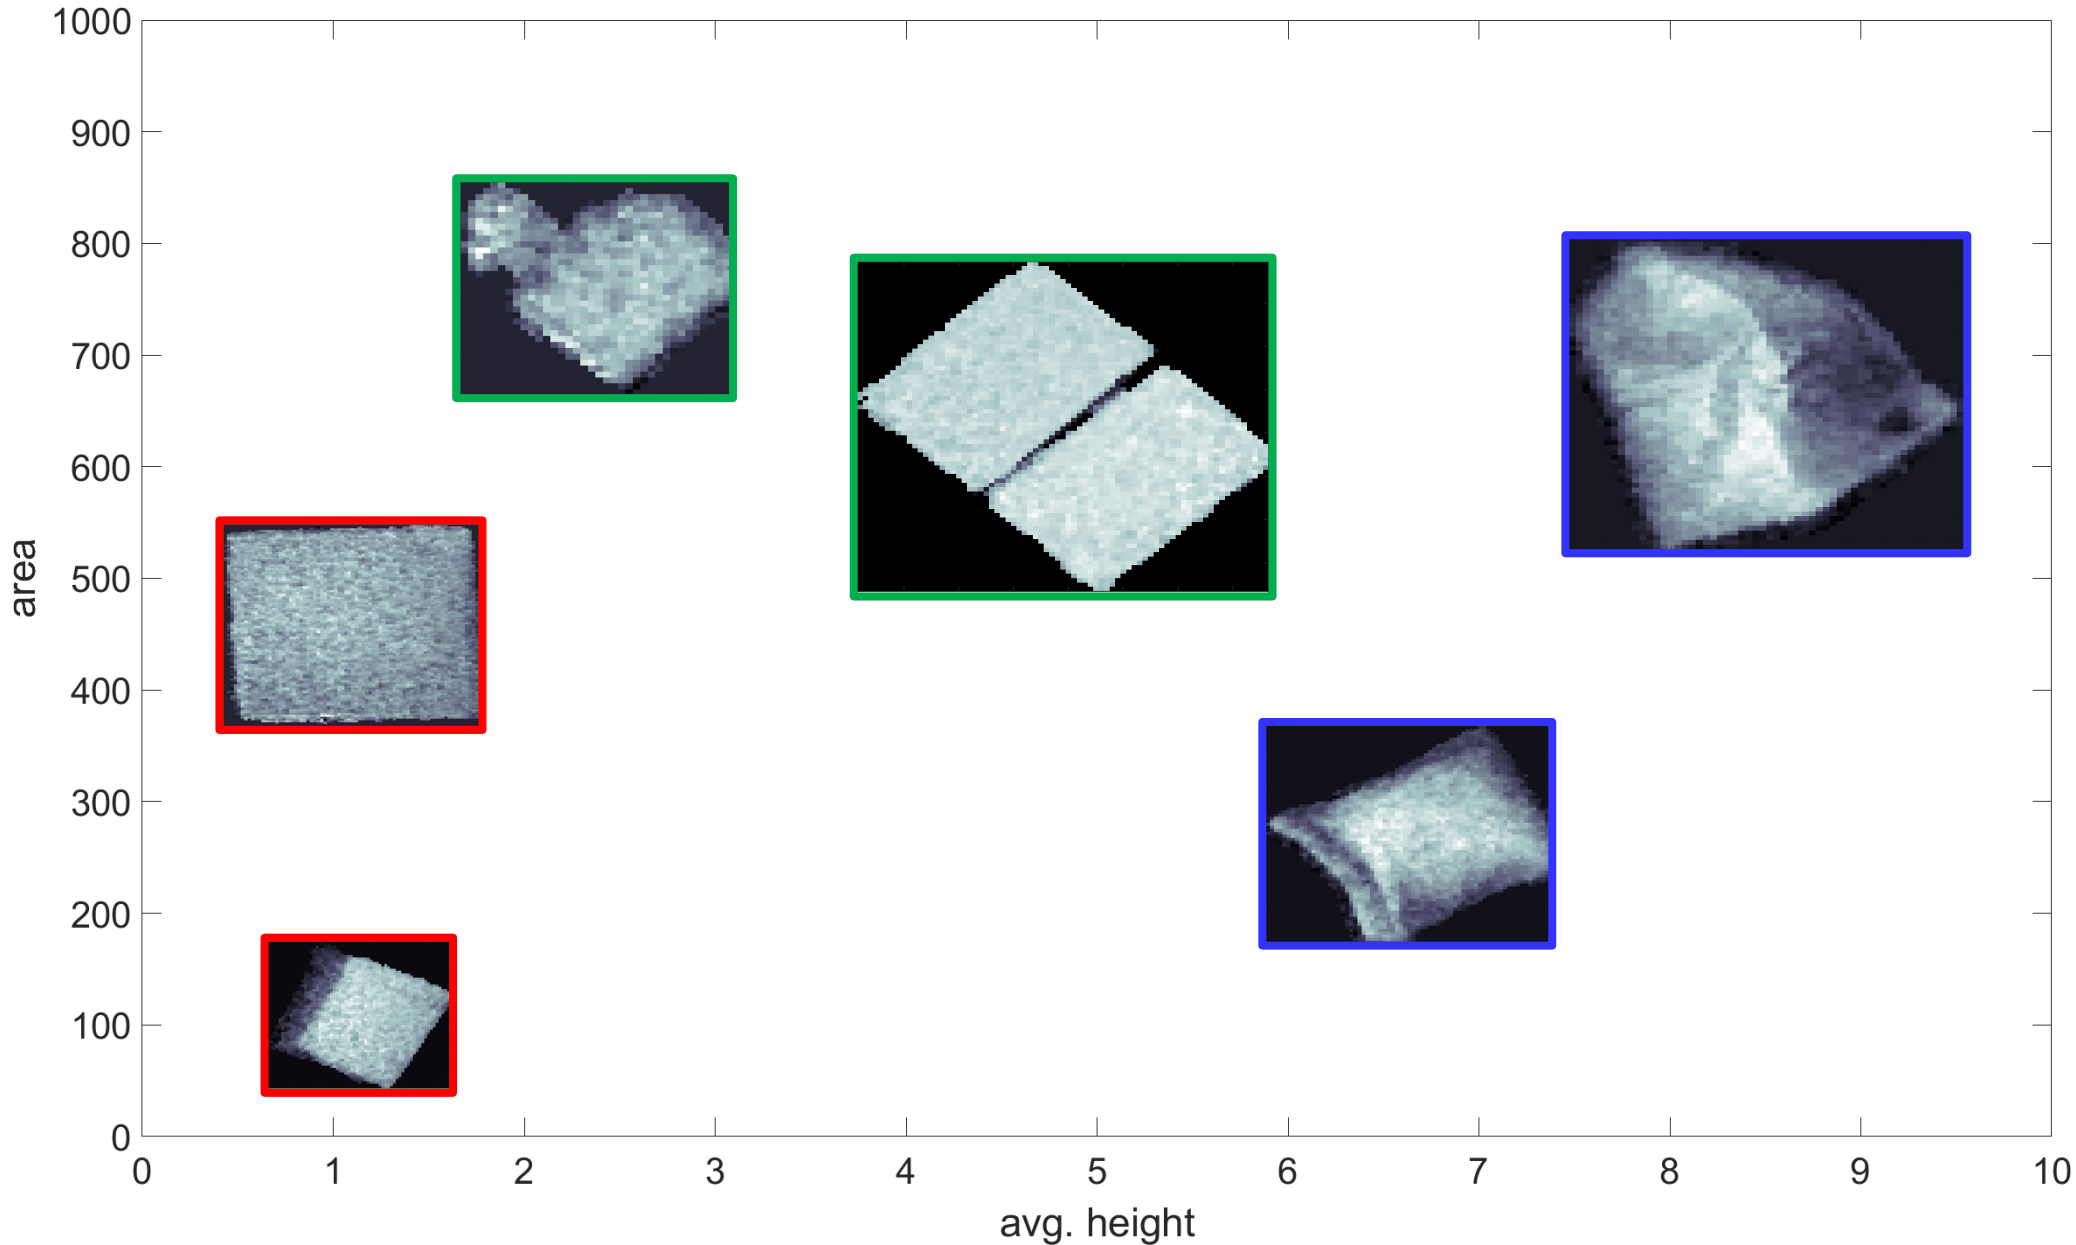

The Training Set: features + labels

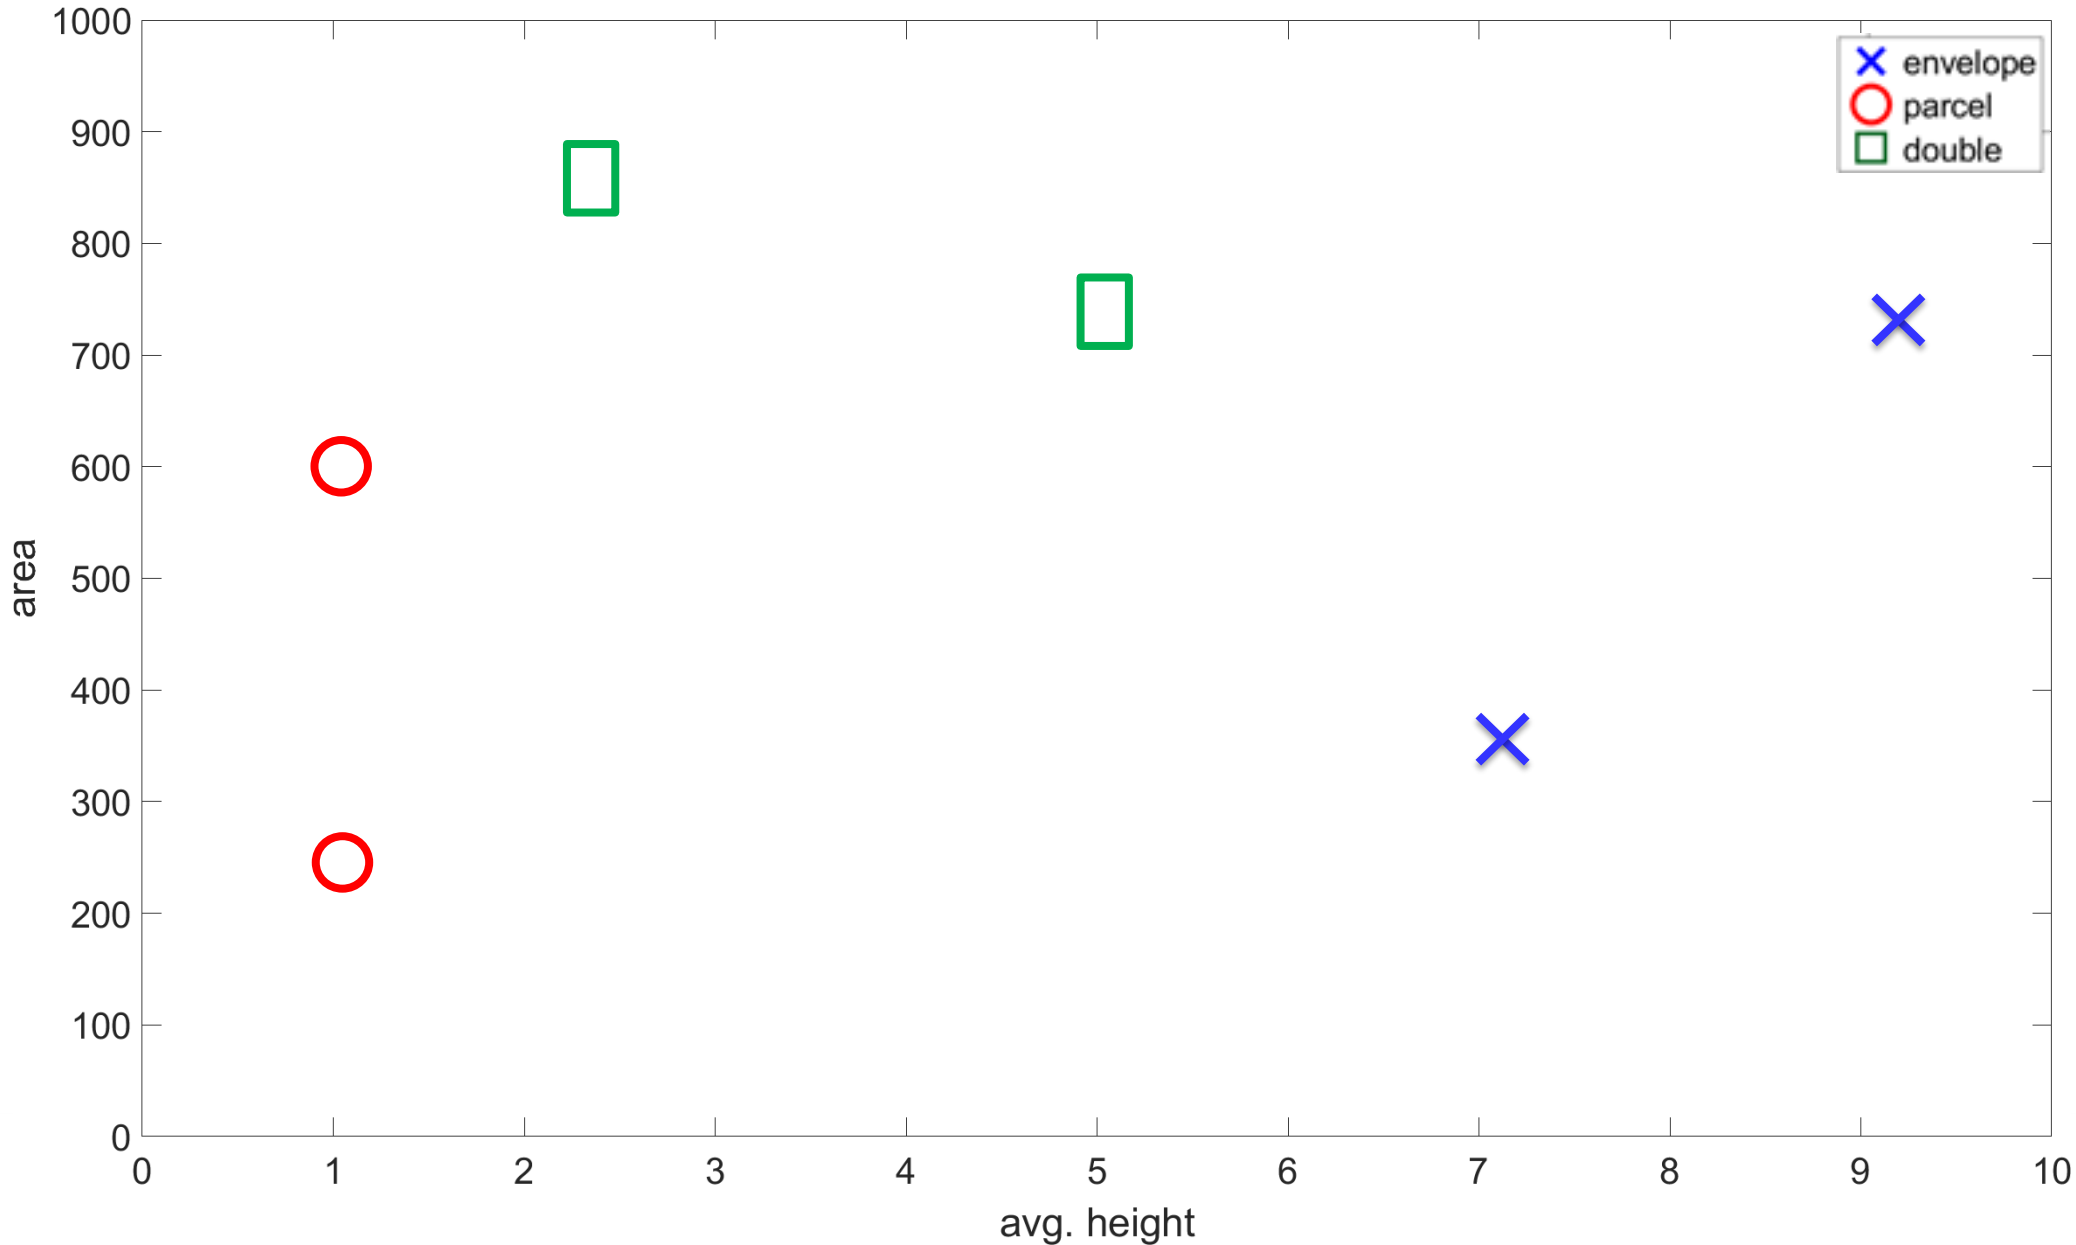

The Training Set

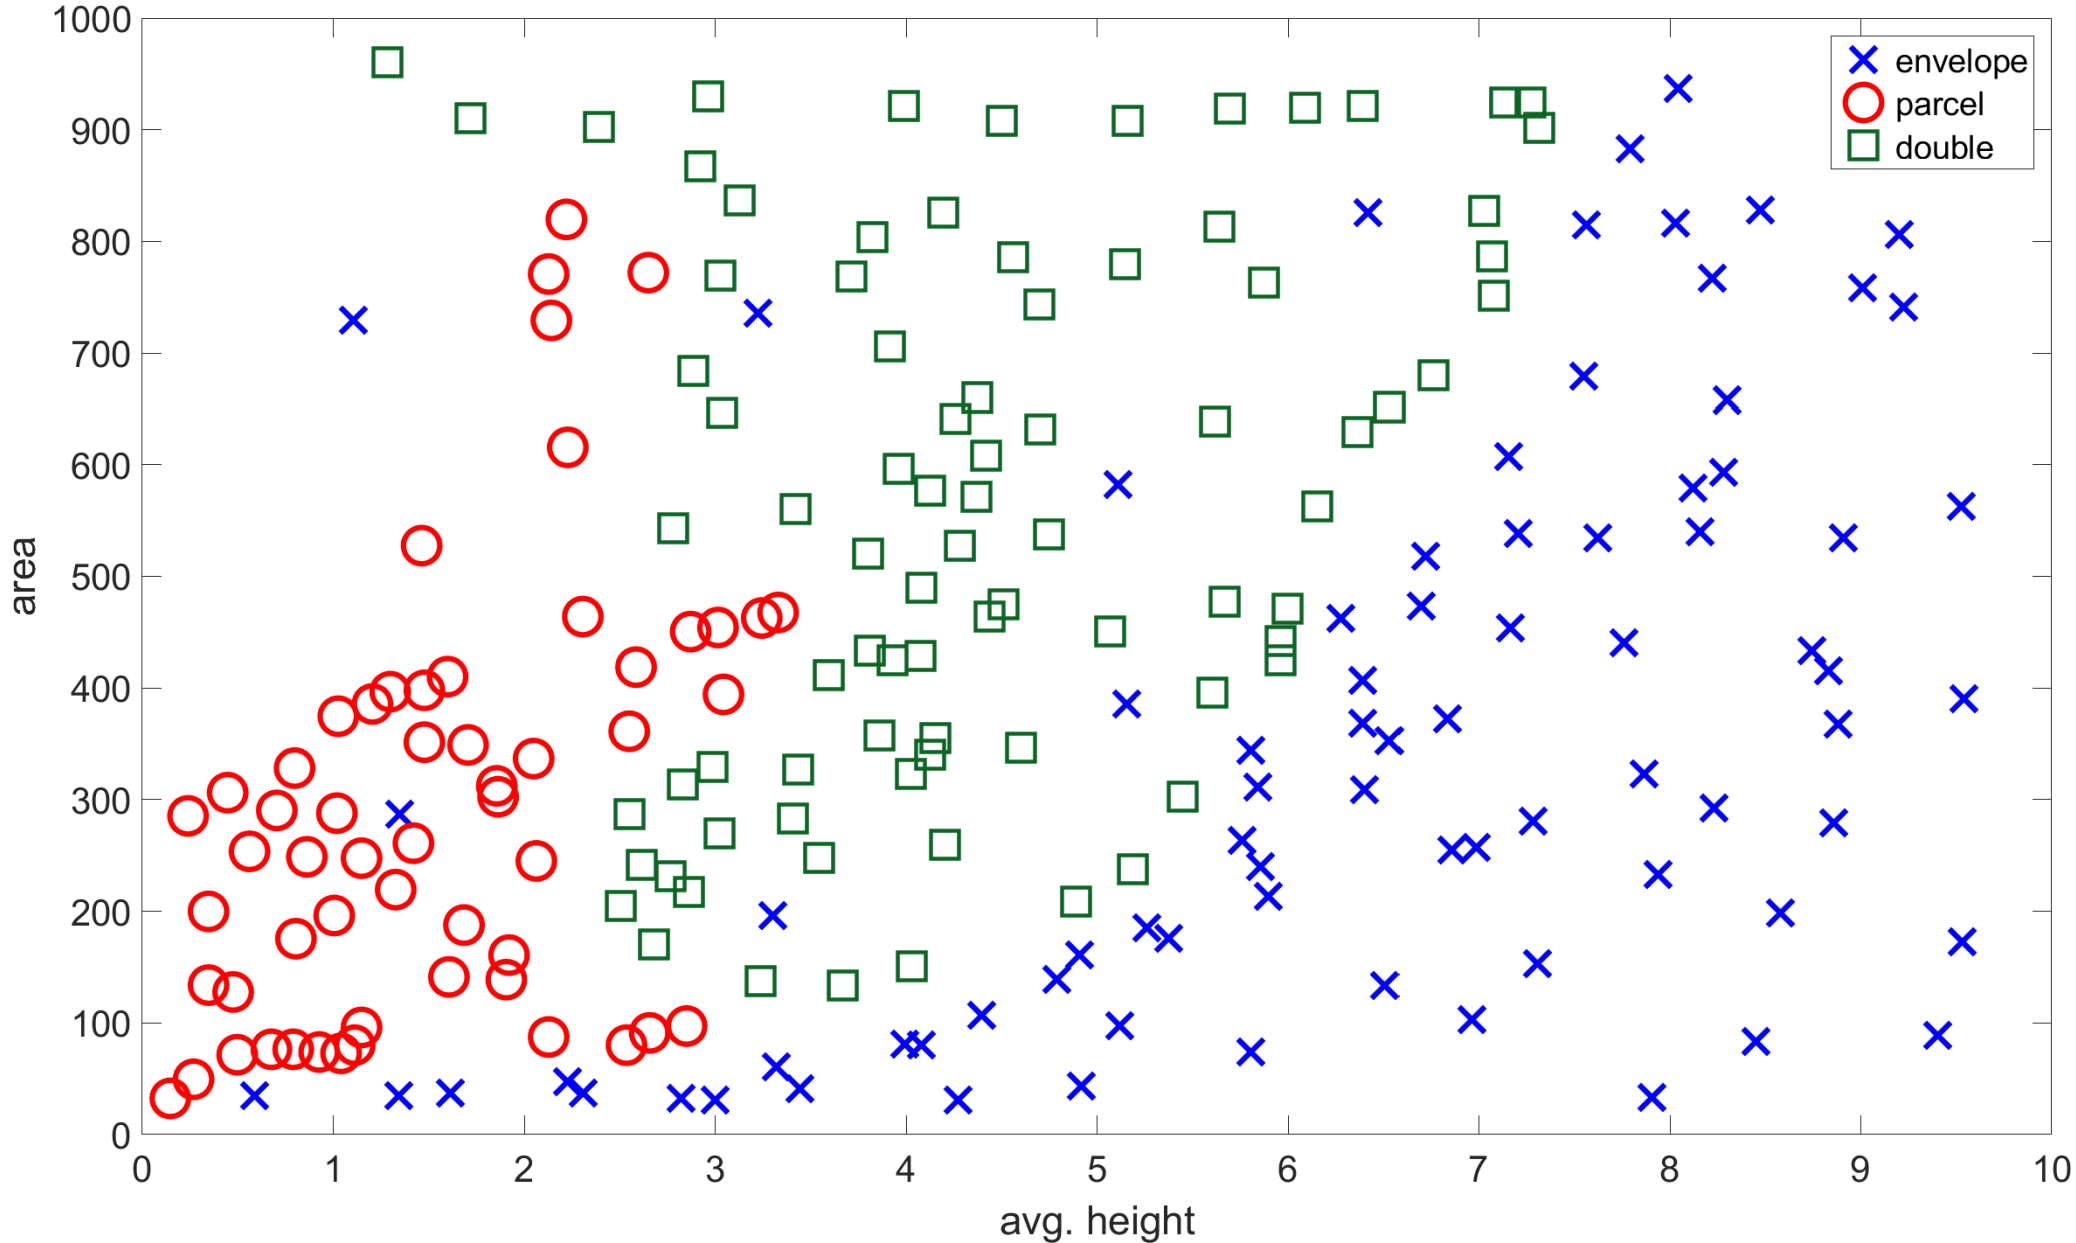

Training Set

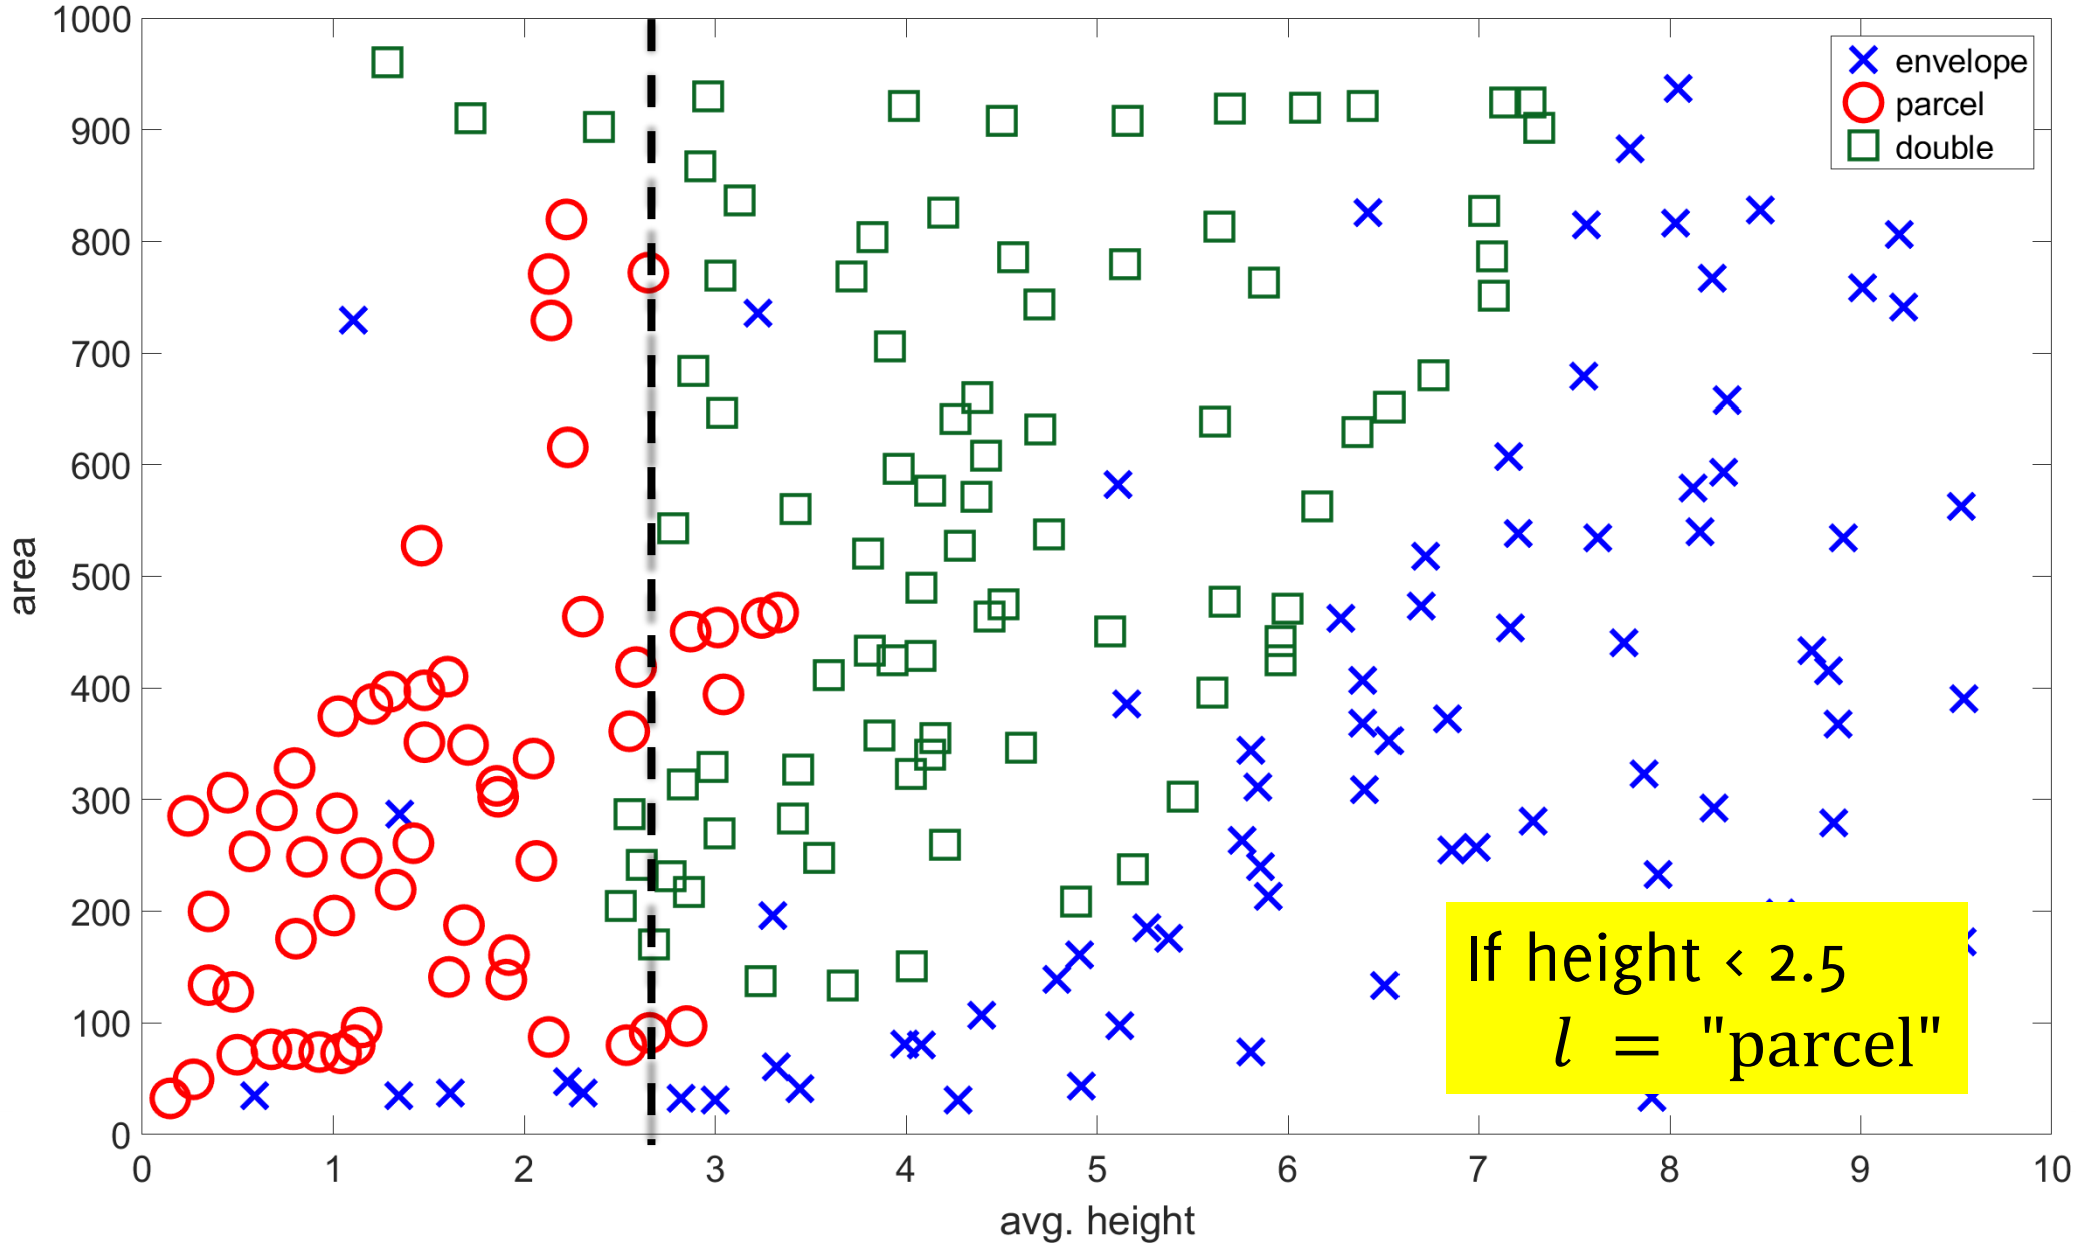

Training Set

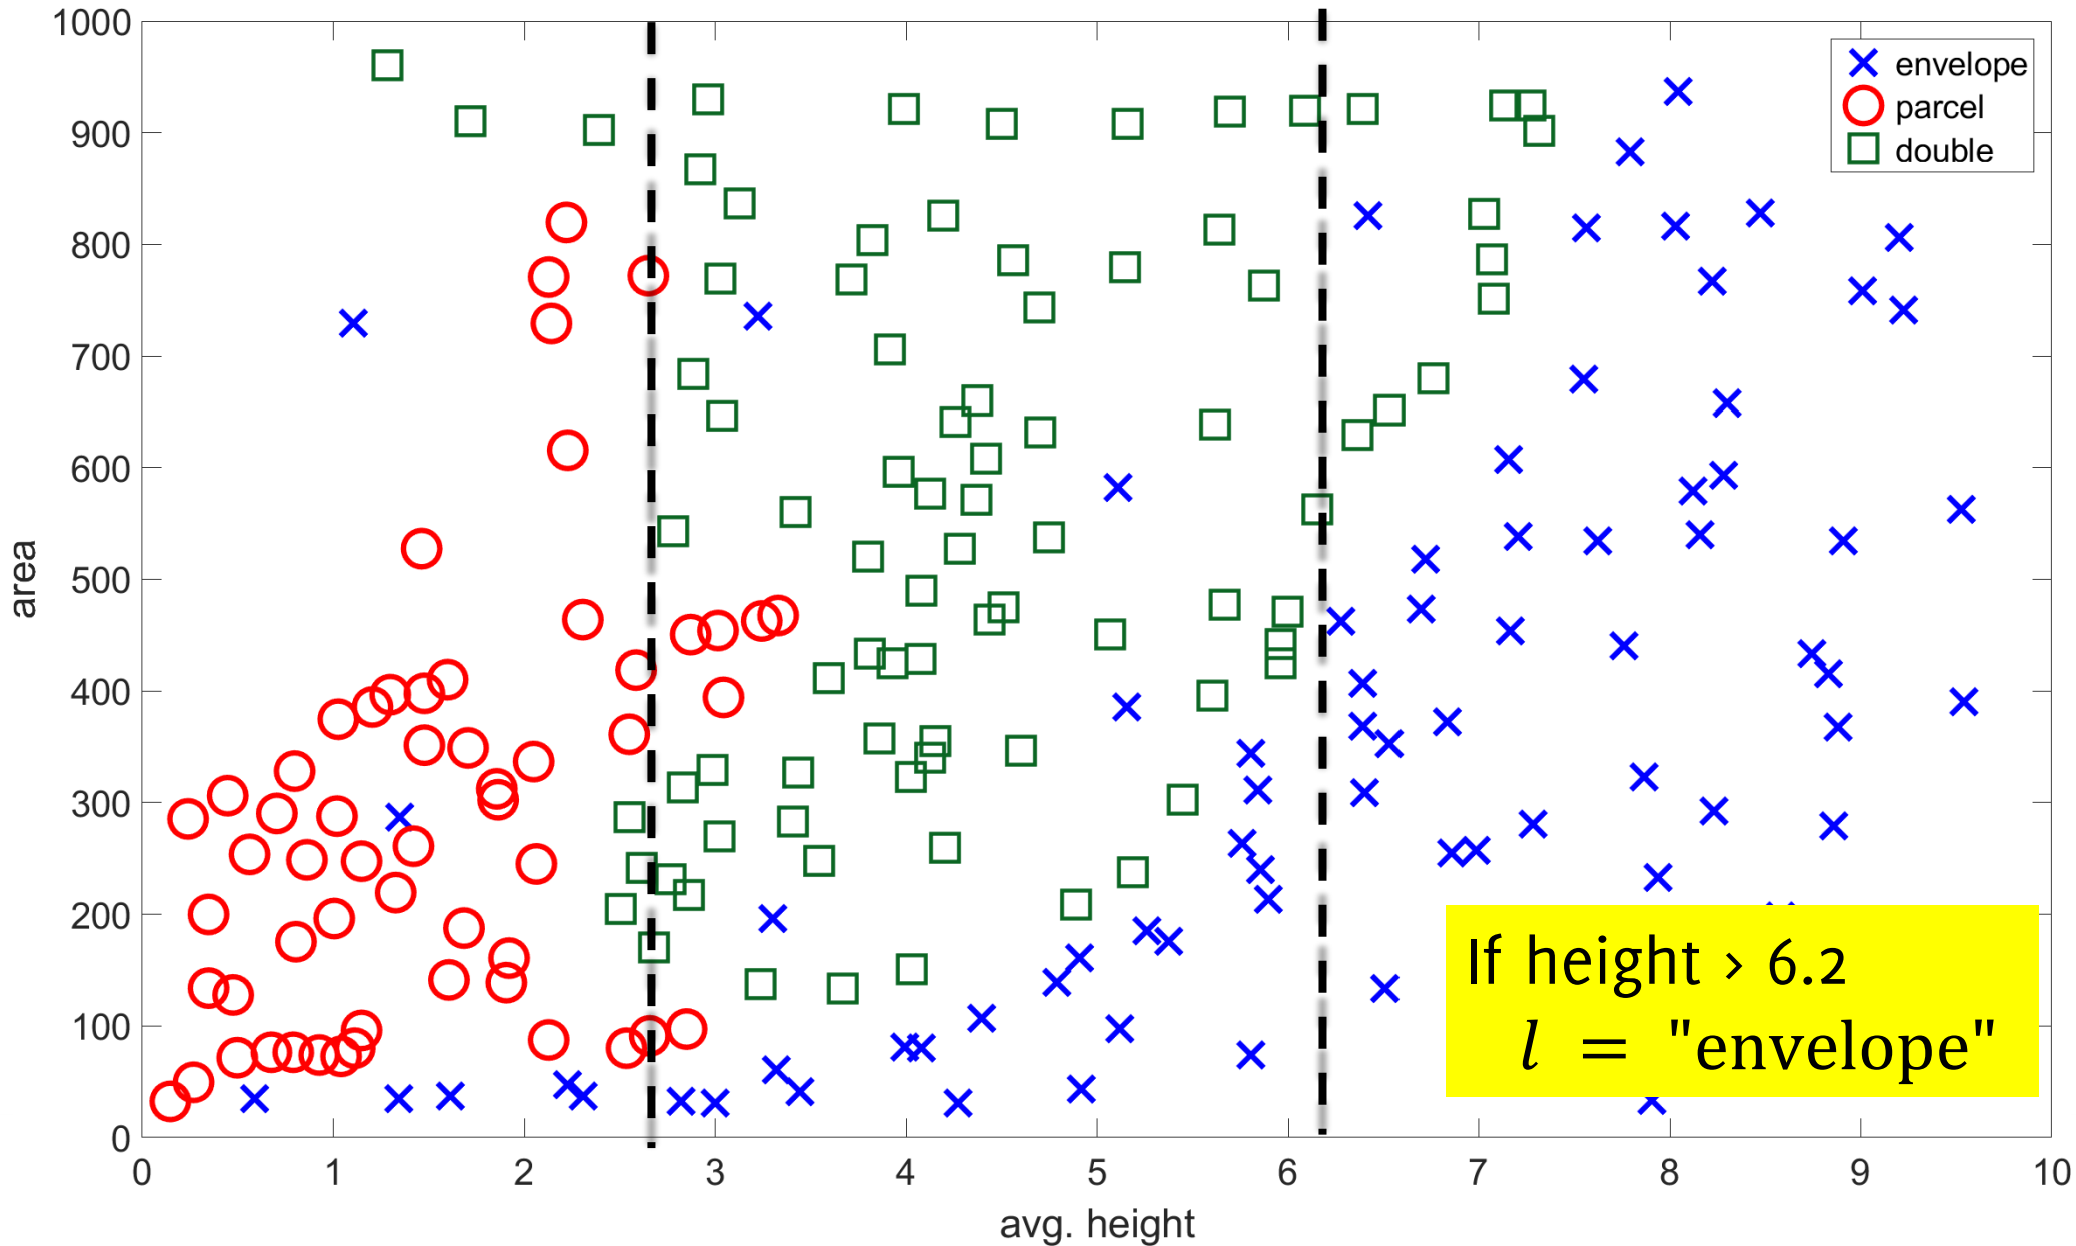

Training Set

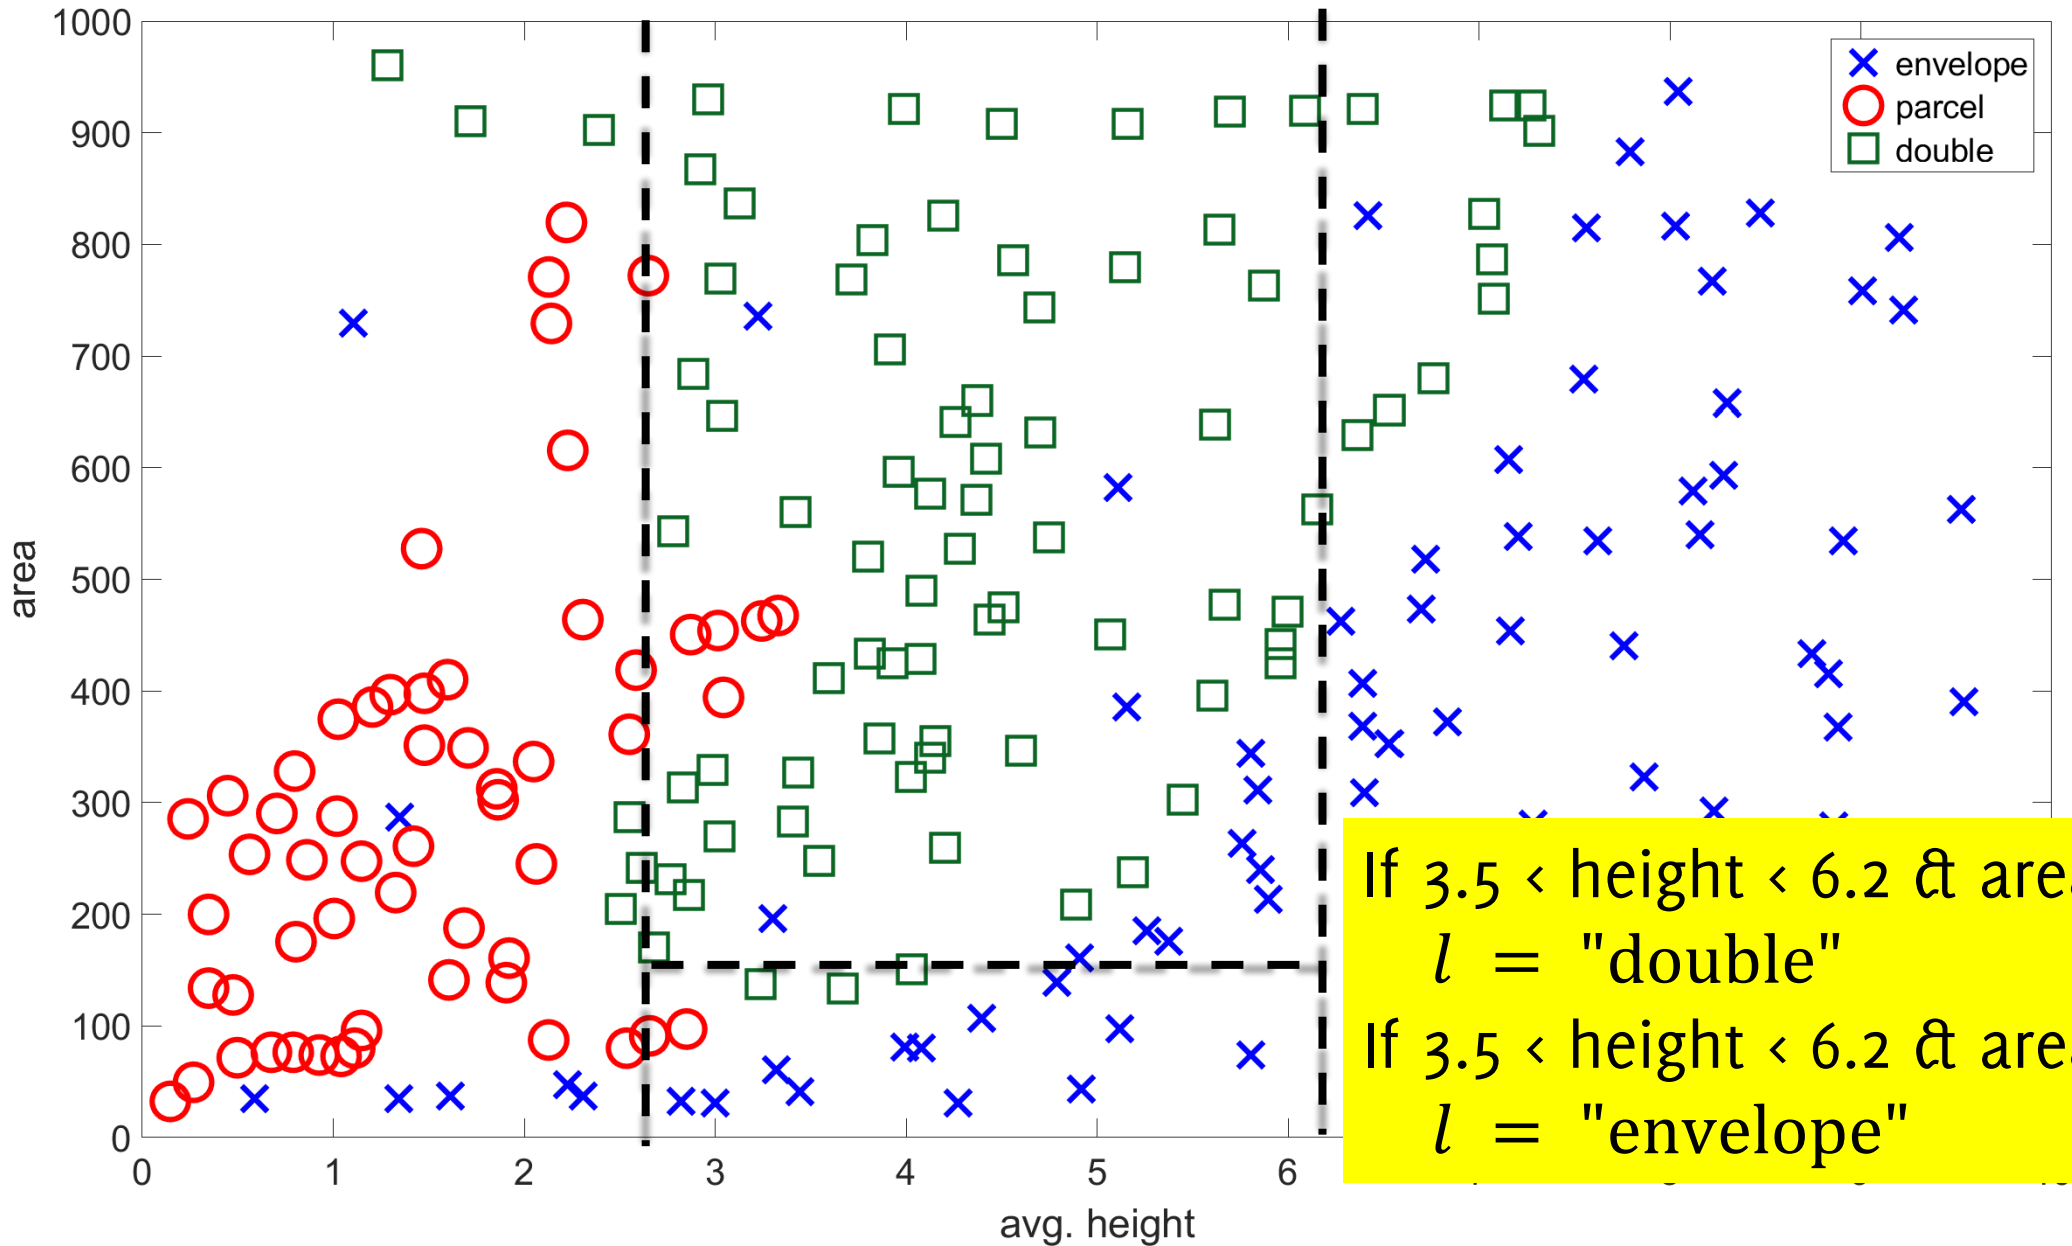

If 3.5 < height < 6.2 & area > 200
l = "double"

If 3.5 < height < 6.2 & area < 200
l = "envelope"

Classifier output

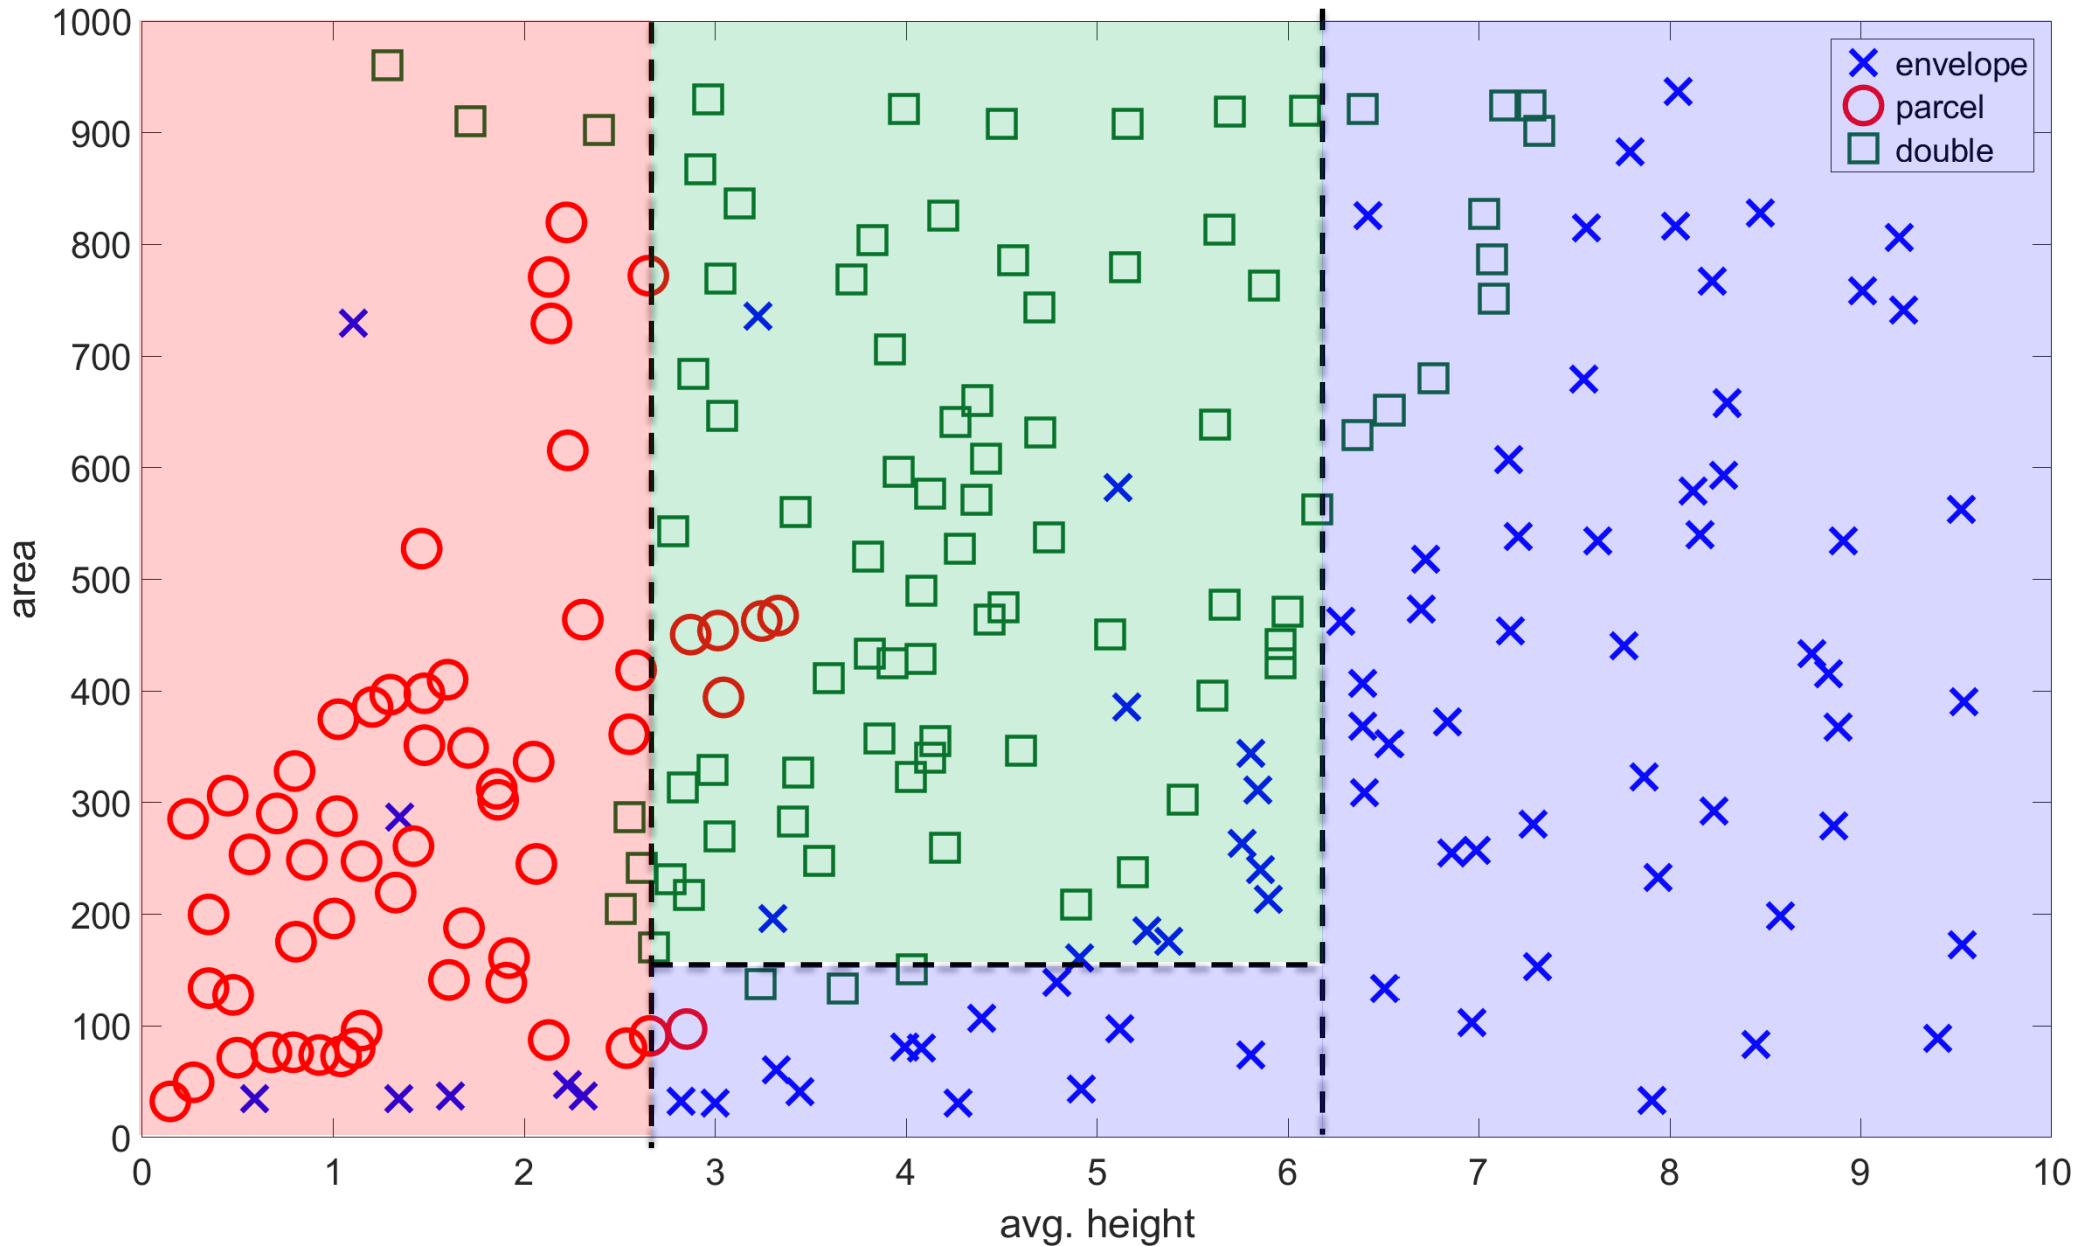

A tree classifying image features

Input image

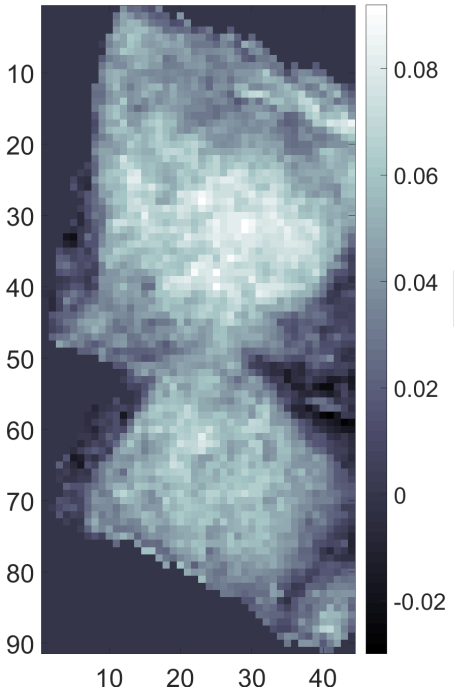

$I_1 \in \mathbb{R}^{r_1 \times c_1}$

Feature Extraction Algorithm

h
a

$\mathbf{x} \in \mathbb{R}^2$

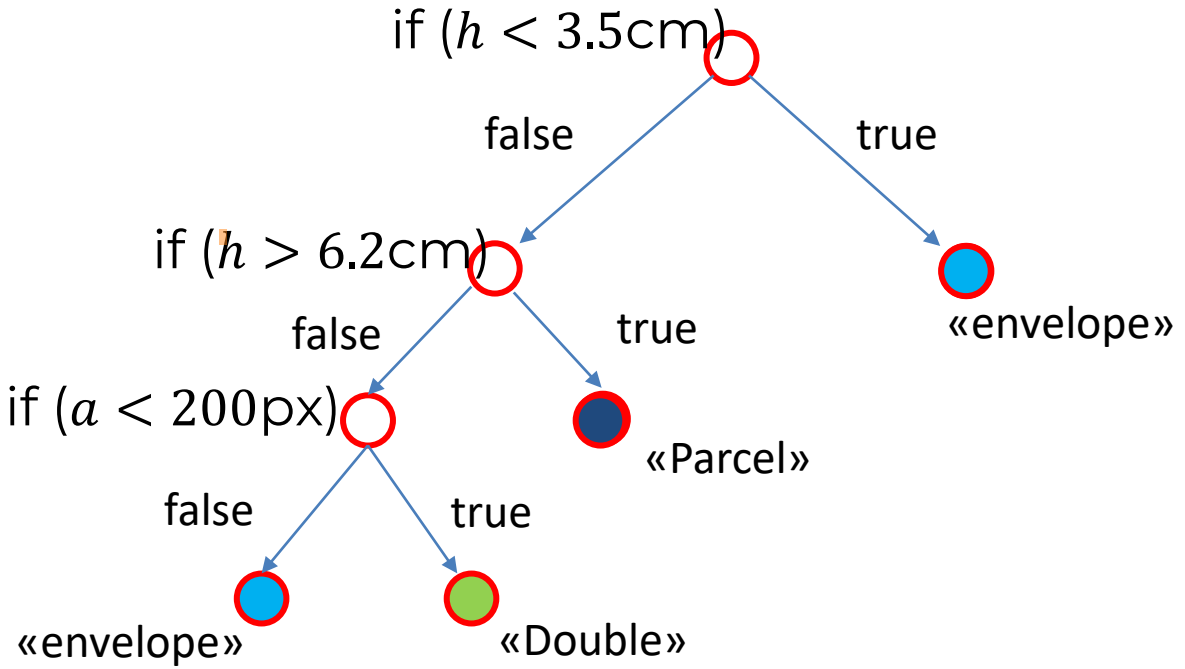

"double" "envelope" "parcel"

Limitations of Rule Based Classifier

It is difficult to grasp what are meaningful dependencies over multiple variables (it is also impossible to visualize these)

Let's resort to a **data-driven model** for the only task of separating feature vectors in different classes.

How can a classifier achieve better performance?

A tree classifying image features

The classifier has a few patameters:

- The **splitting criteria**
- The **splitting thresholds T_i**

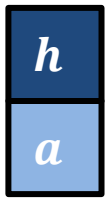

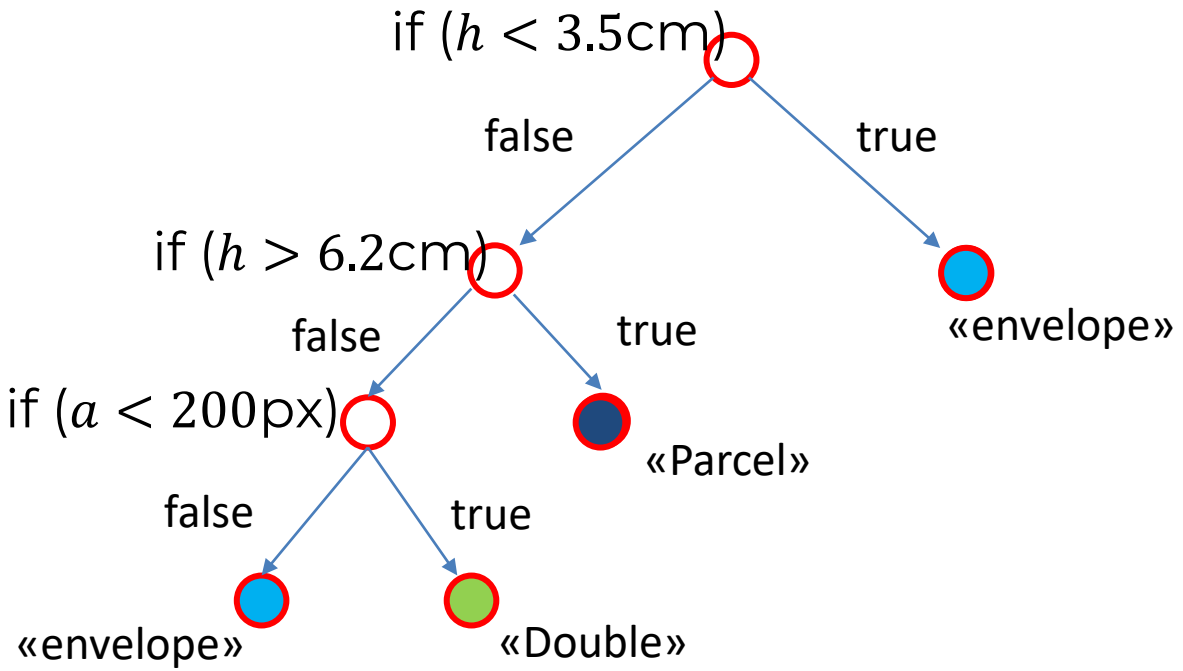

$\mathbf{x} \in \mathbb{R}^2$

"double" "envelope" "parcel"

This is our first solution

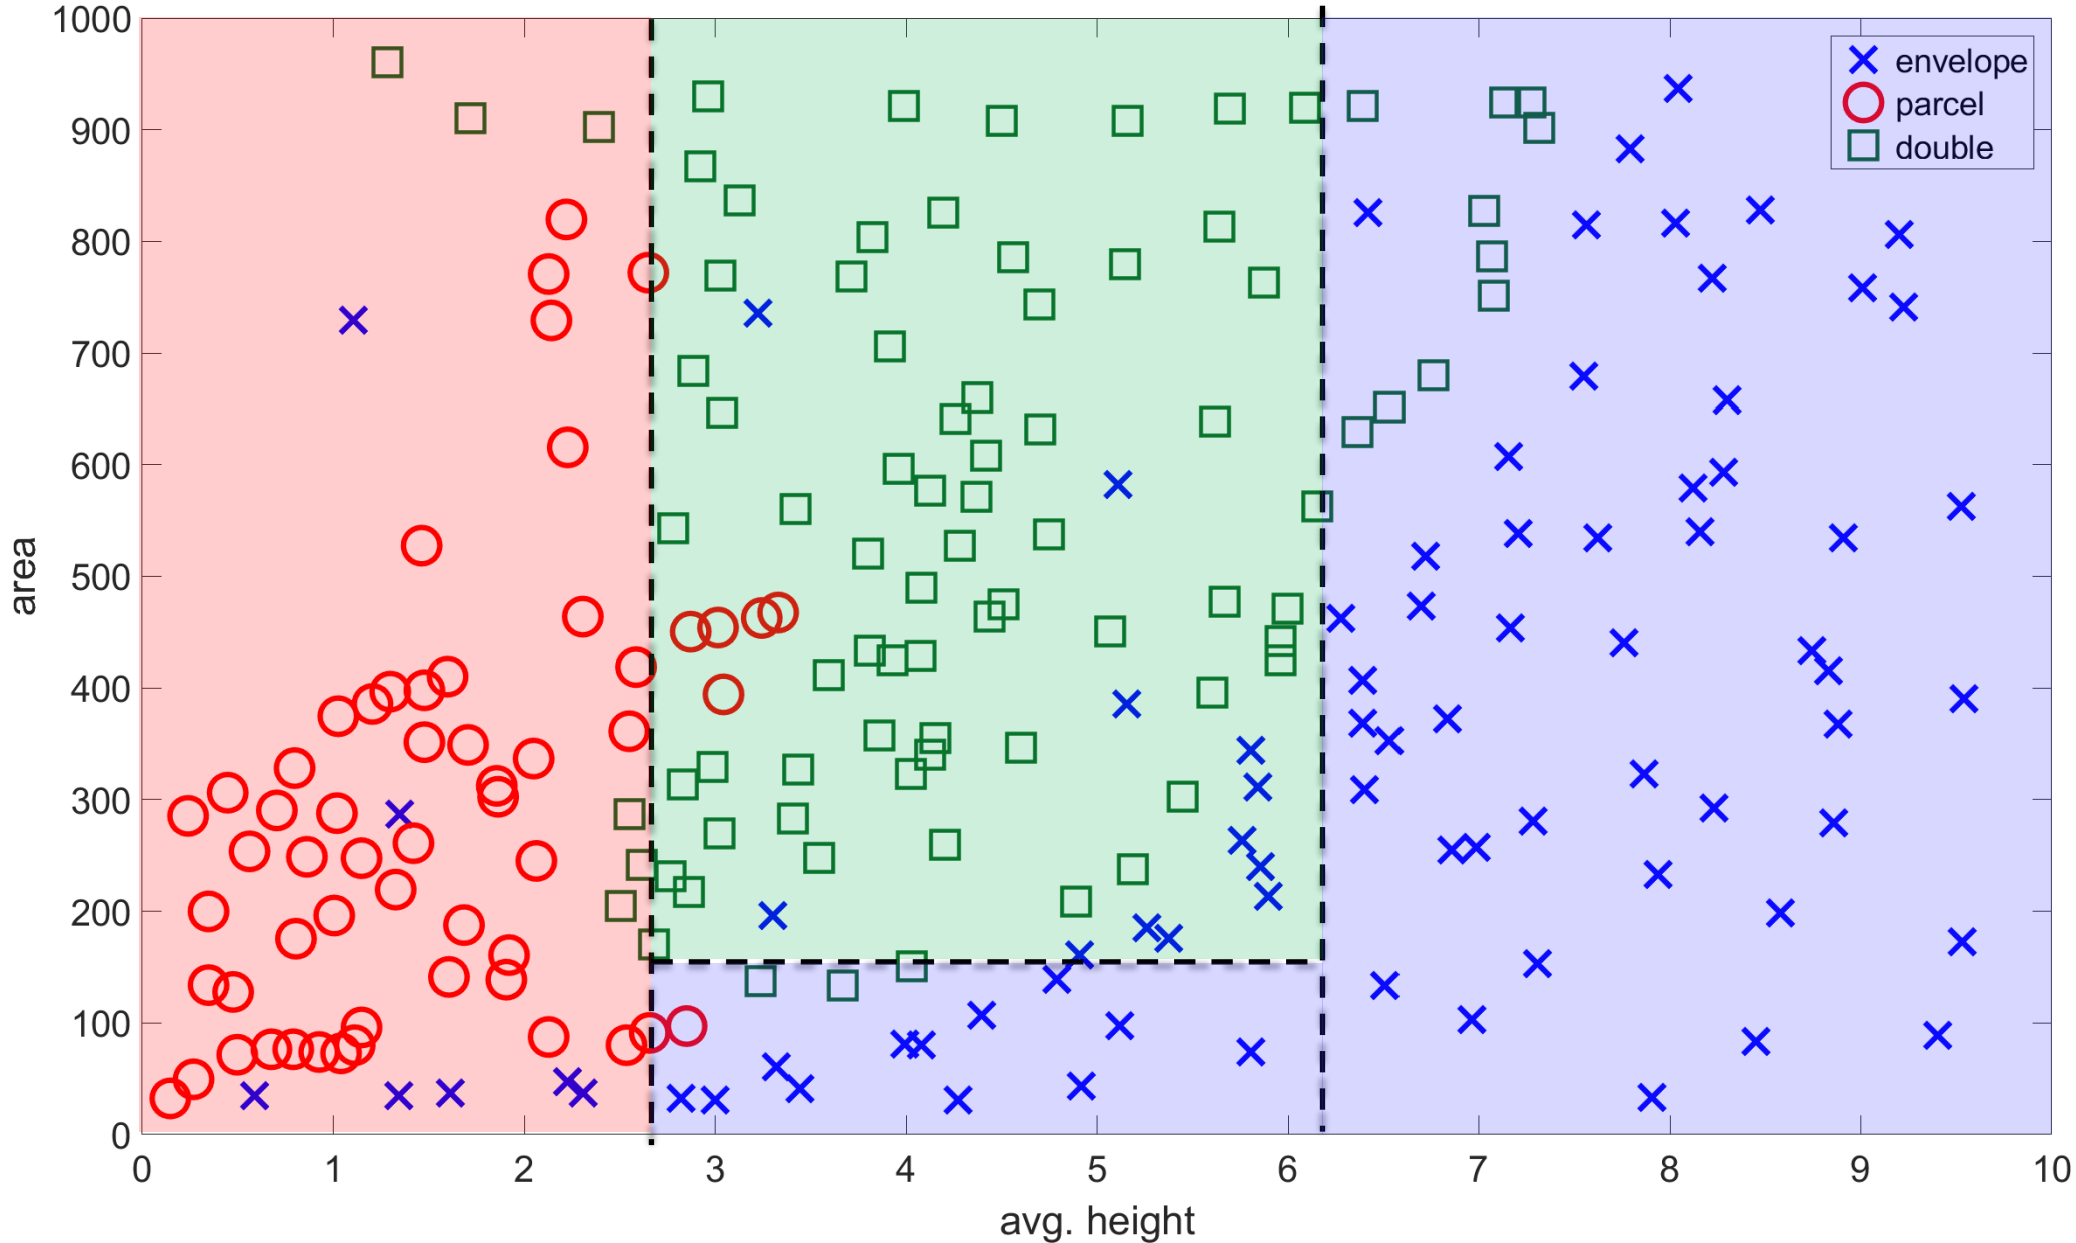

There are a few errors

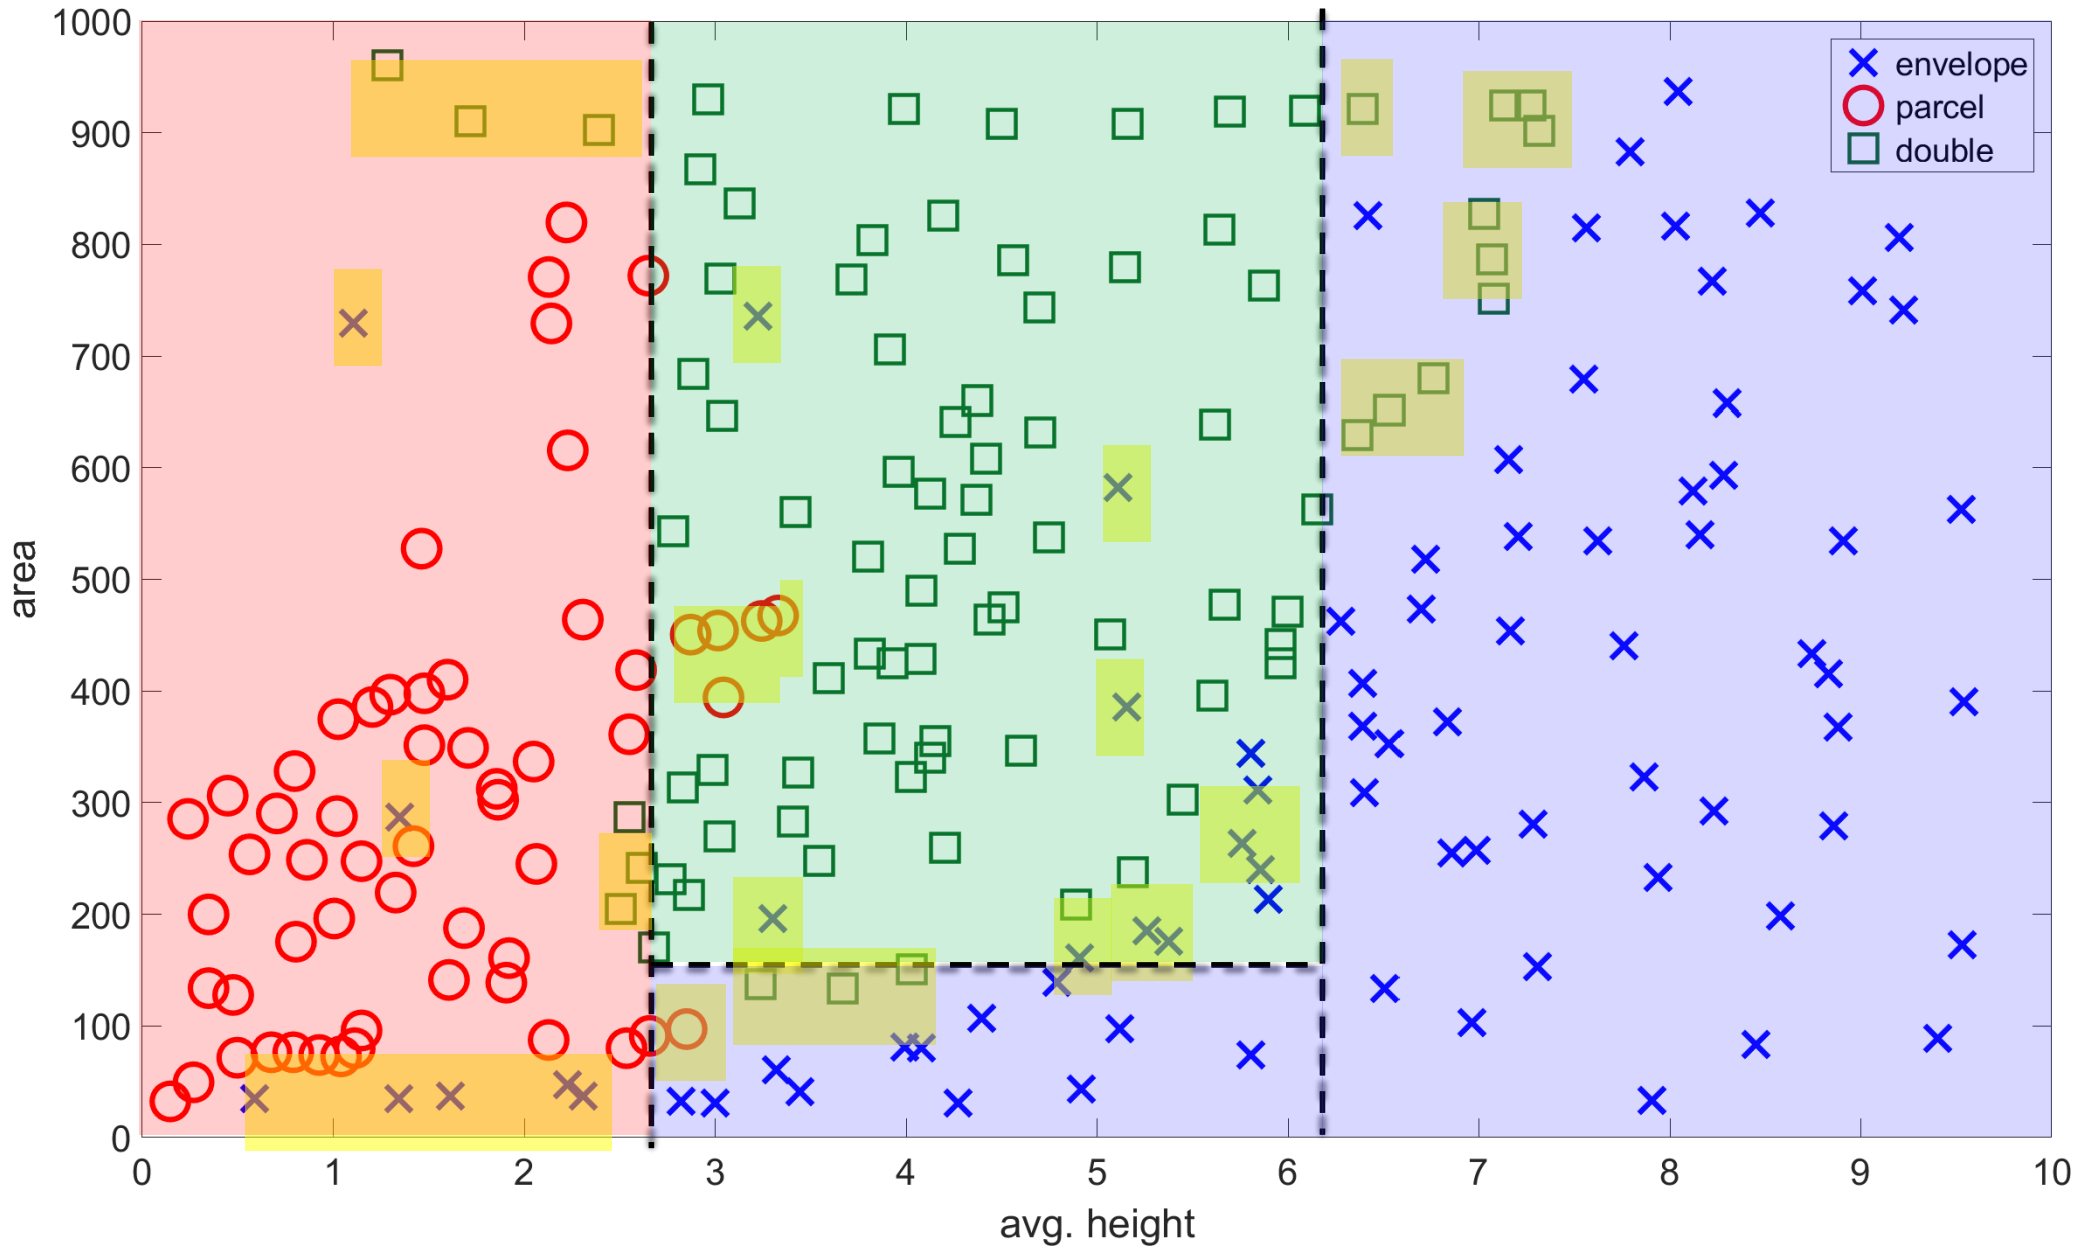

Can I do better?

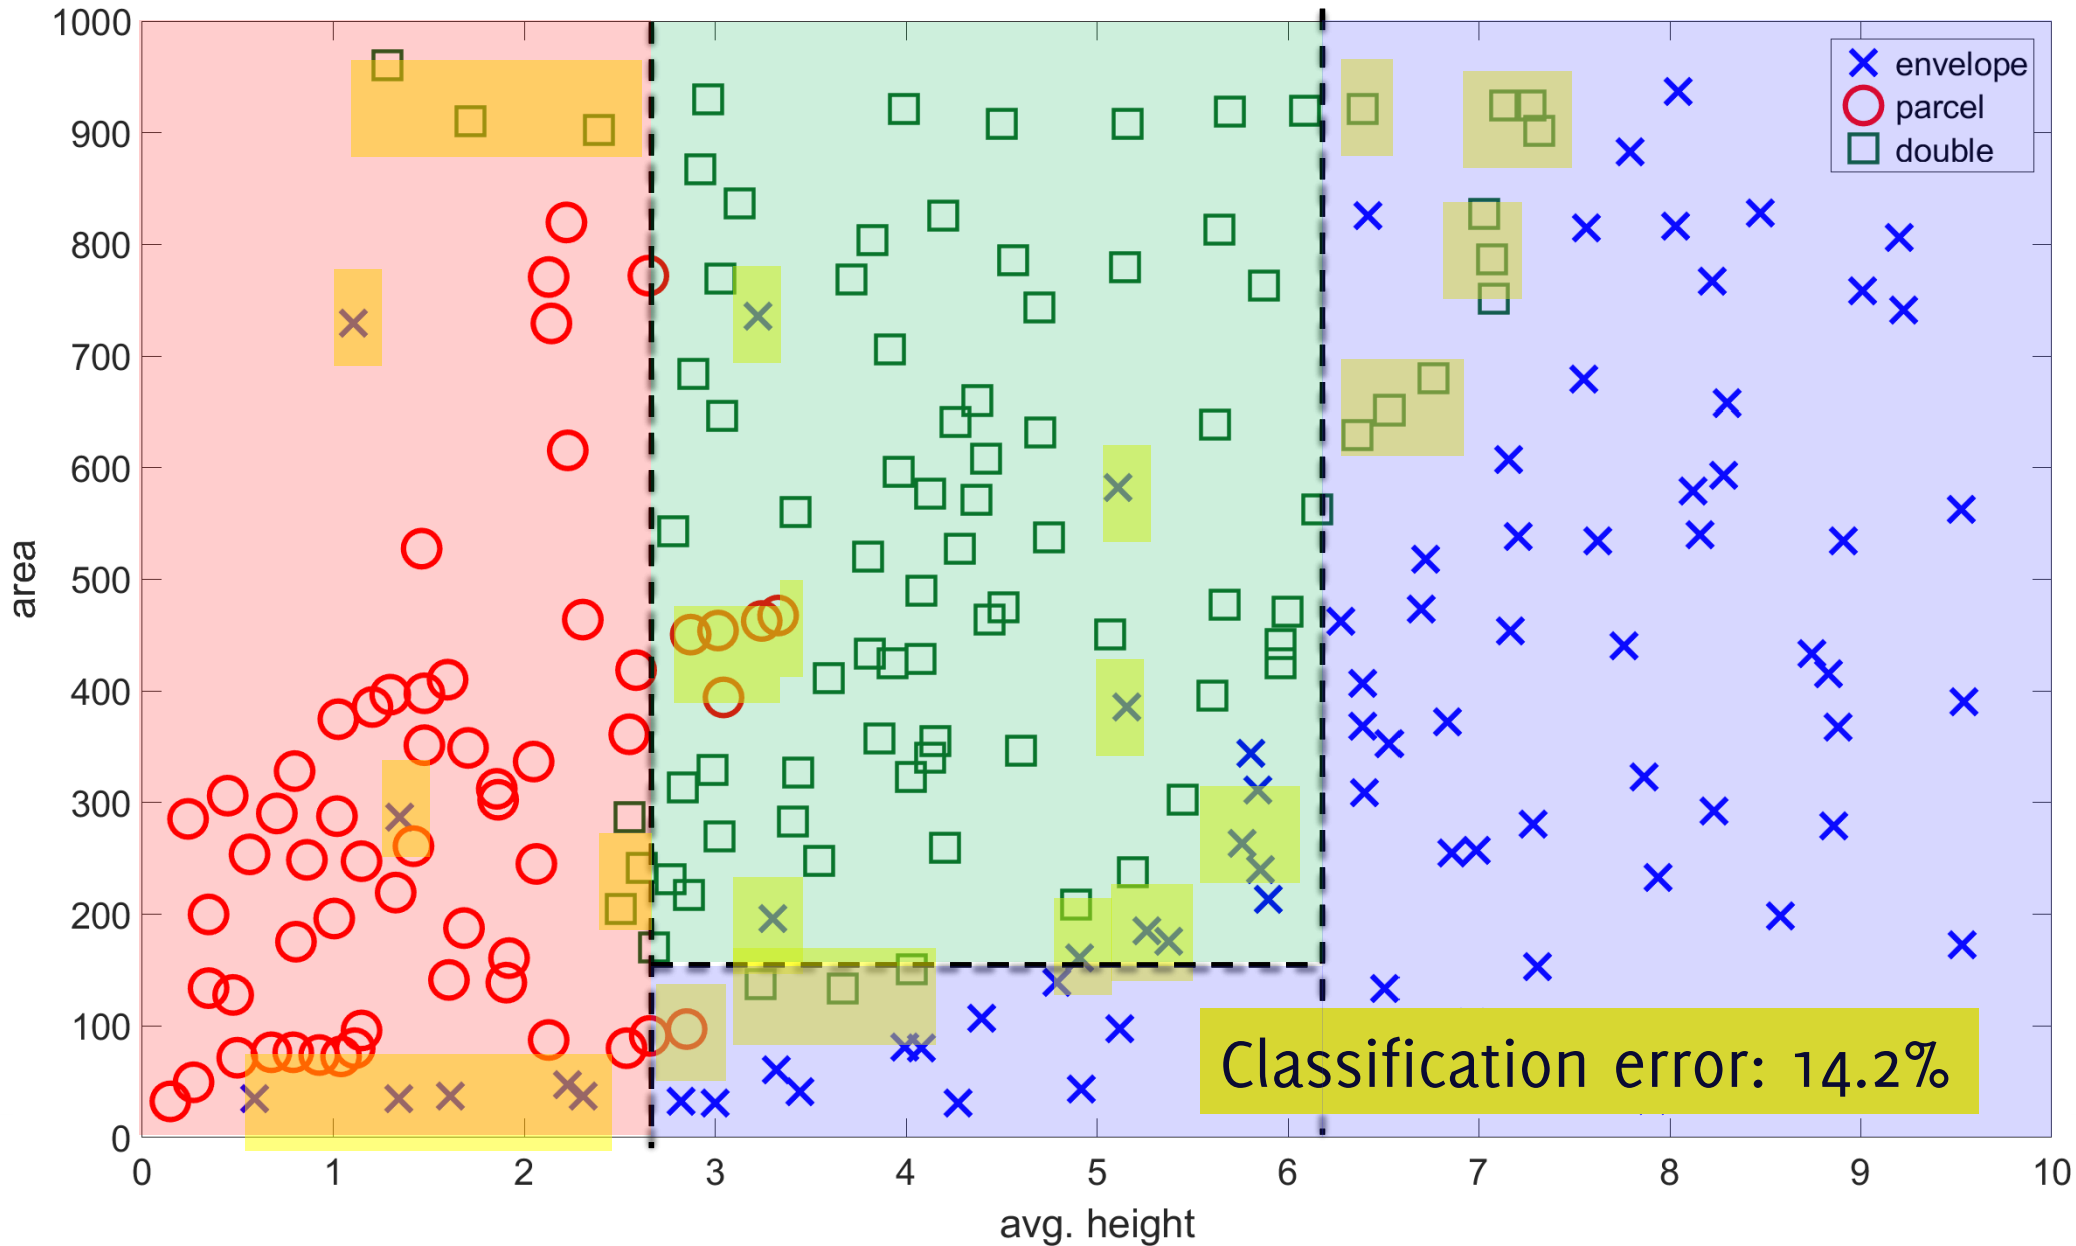

Let's try different parameters

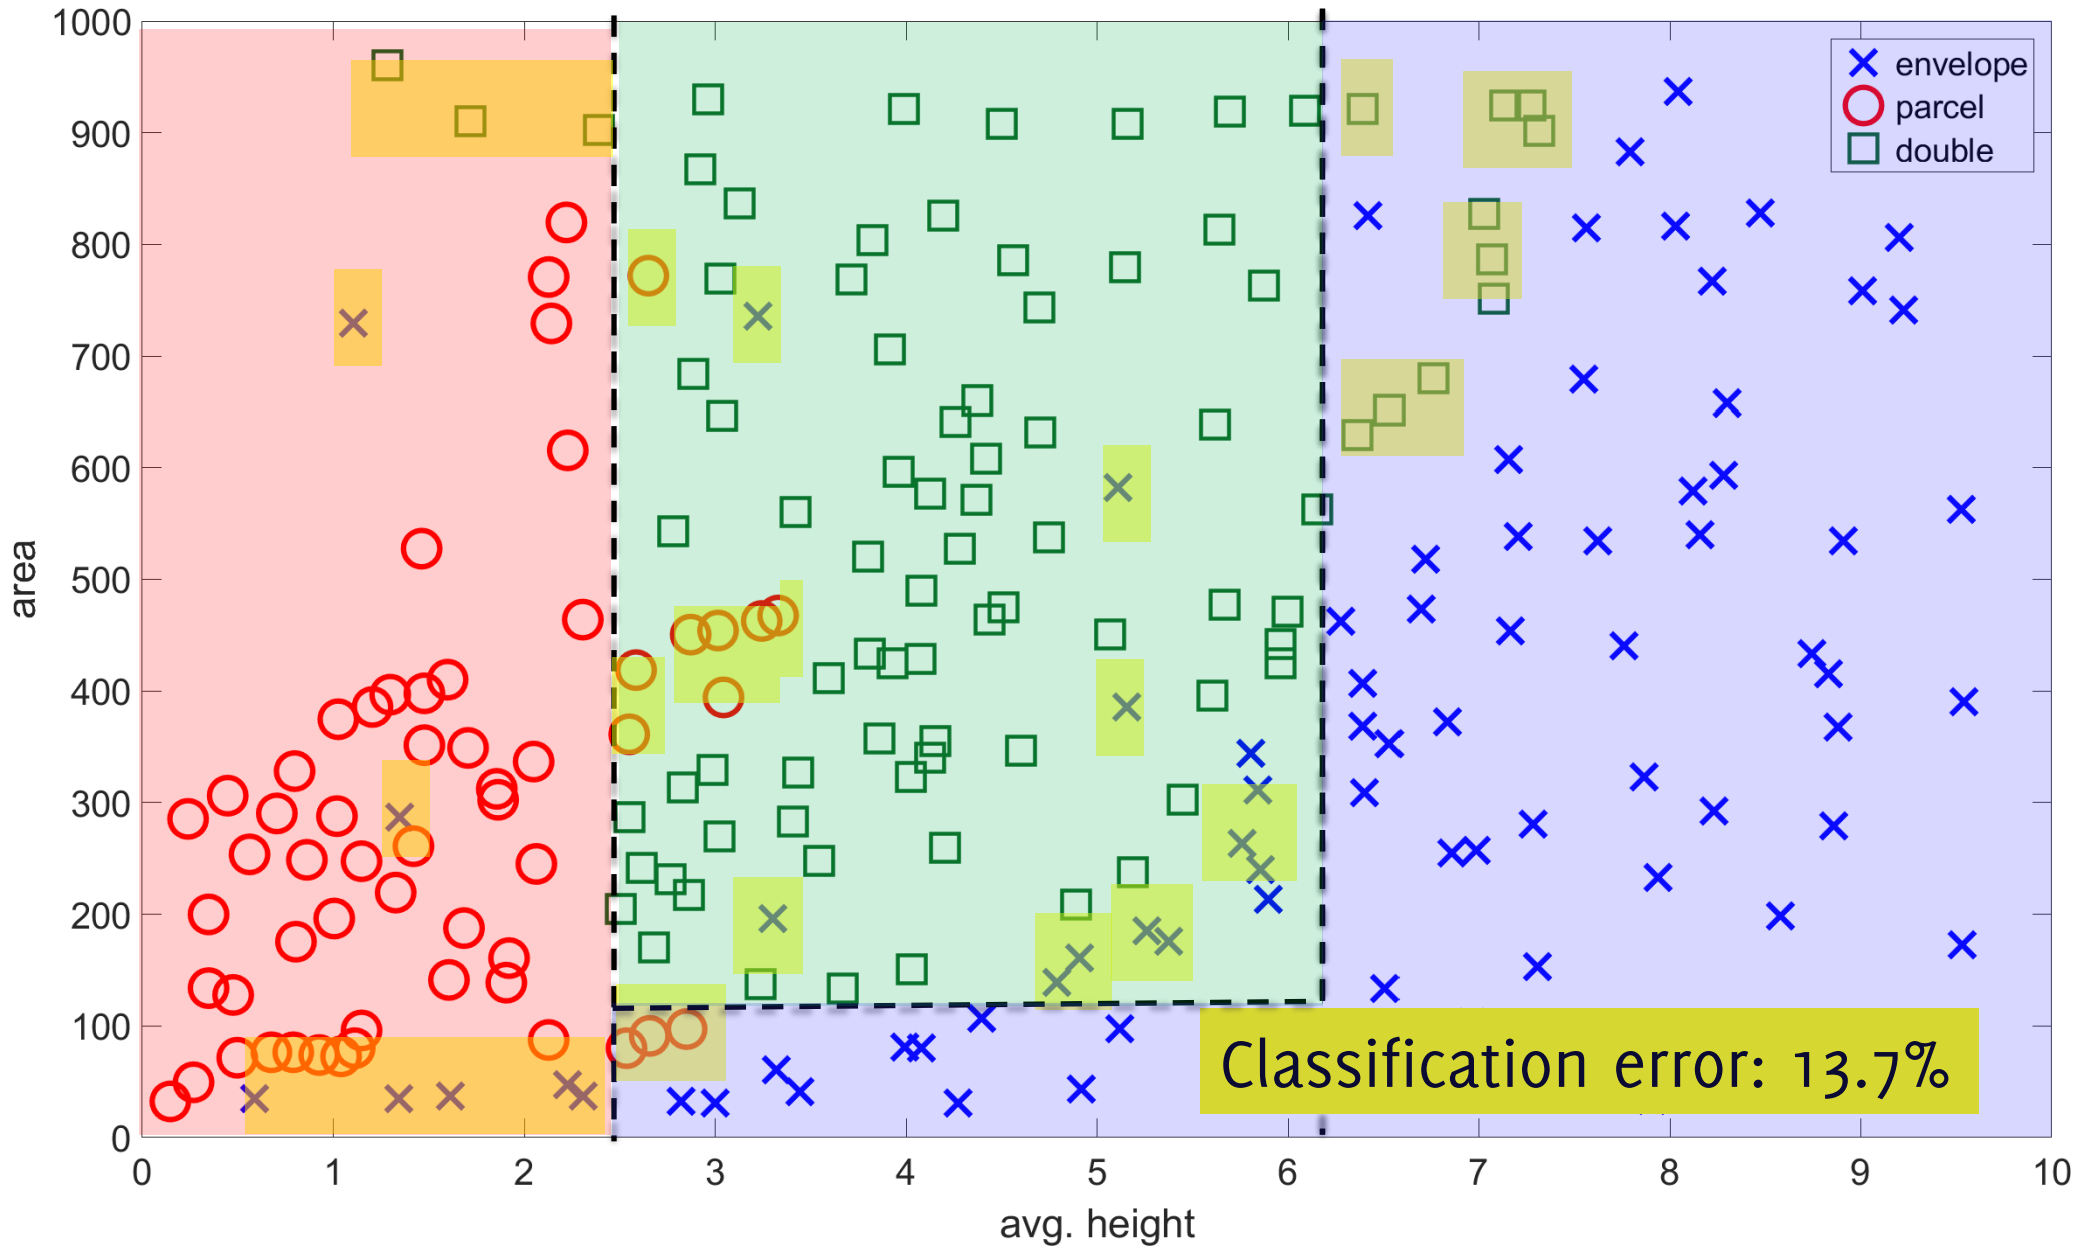

Data Driven Models

Data Driven Models are defined from a training set of (supervised) pairs

$$TR = \{(\boldsymbol{x}, \boldsymbol{y})_i, i = 1, \dots, N\}$$

The model parameters θ (e.g. Neural Network weights) are set to minimize a **loss function** (e.g., the classification error in case of discrete output or the reconstruction error in case of continuous output)

$$\theta^* = \underset{\theta}{\operatorname{argmin}}\, \mathcal{L}(\theta, TR)$$

Network training is an optimization process to find params minimizing the loss function.

Can definitvely boost the image classification performance

- Annotated training set is always needed
- Classification performance depends on the training set
- Generalization is not guaranteed

Hand Crafted Feature Extraction, data-driven Classification

Input image

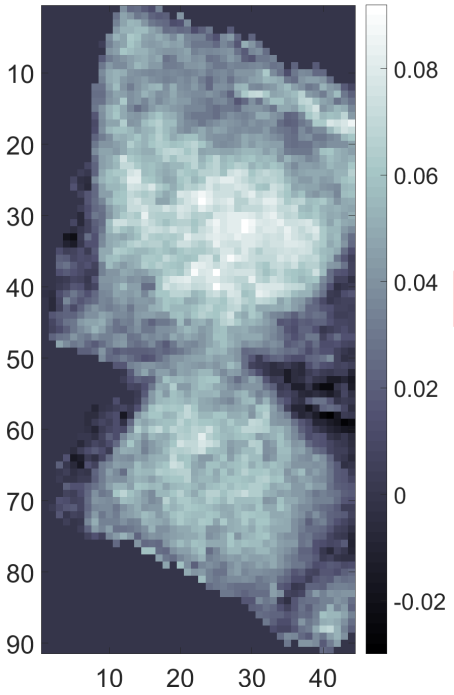

$I_1 \in \mathbb{R}^{r_1 \times c_1}$

Feature Extraction Algorithm

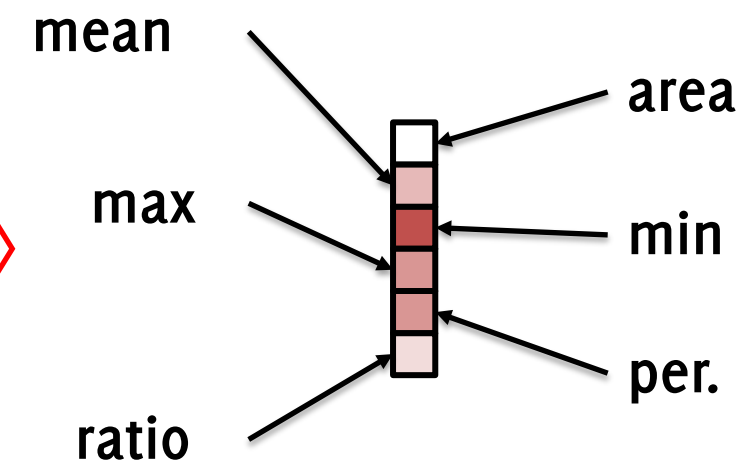

$\mathbf{x} \in \mathbb{R}^d$

$(d \ll r \times c)$

Classifier

"double"

$t \in \Lambda$

Are there better classifiers?

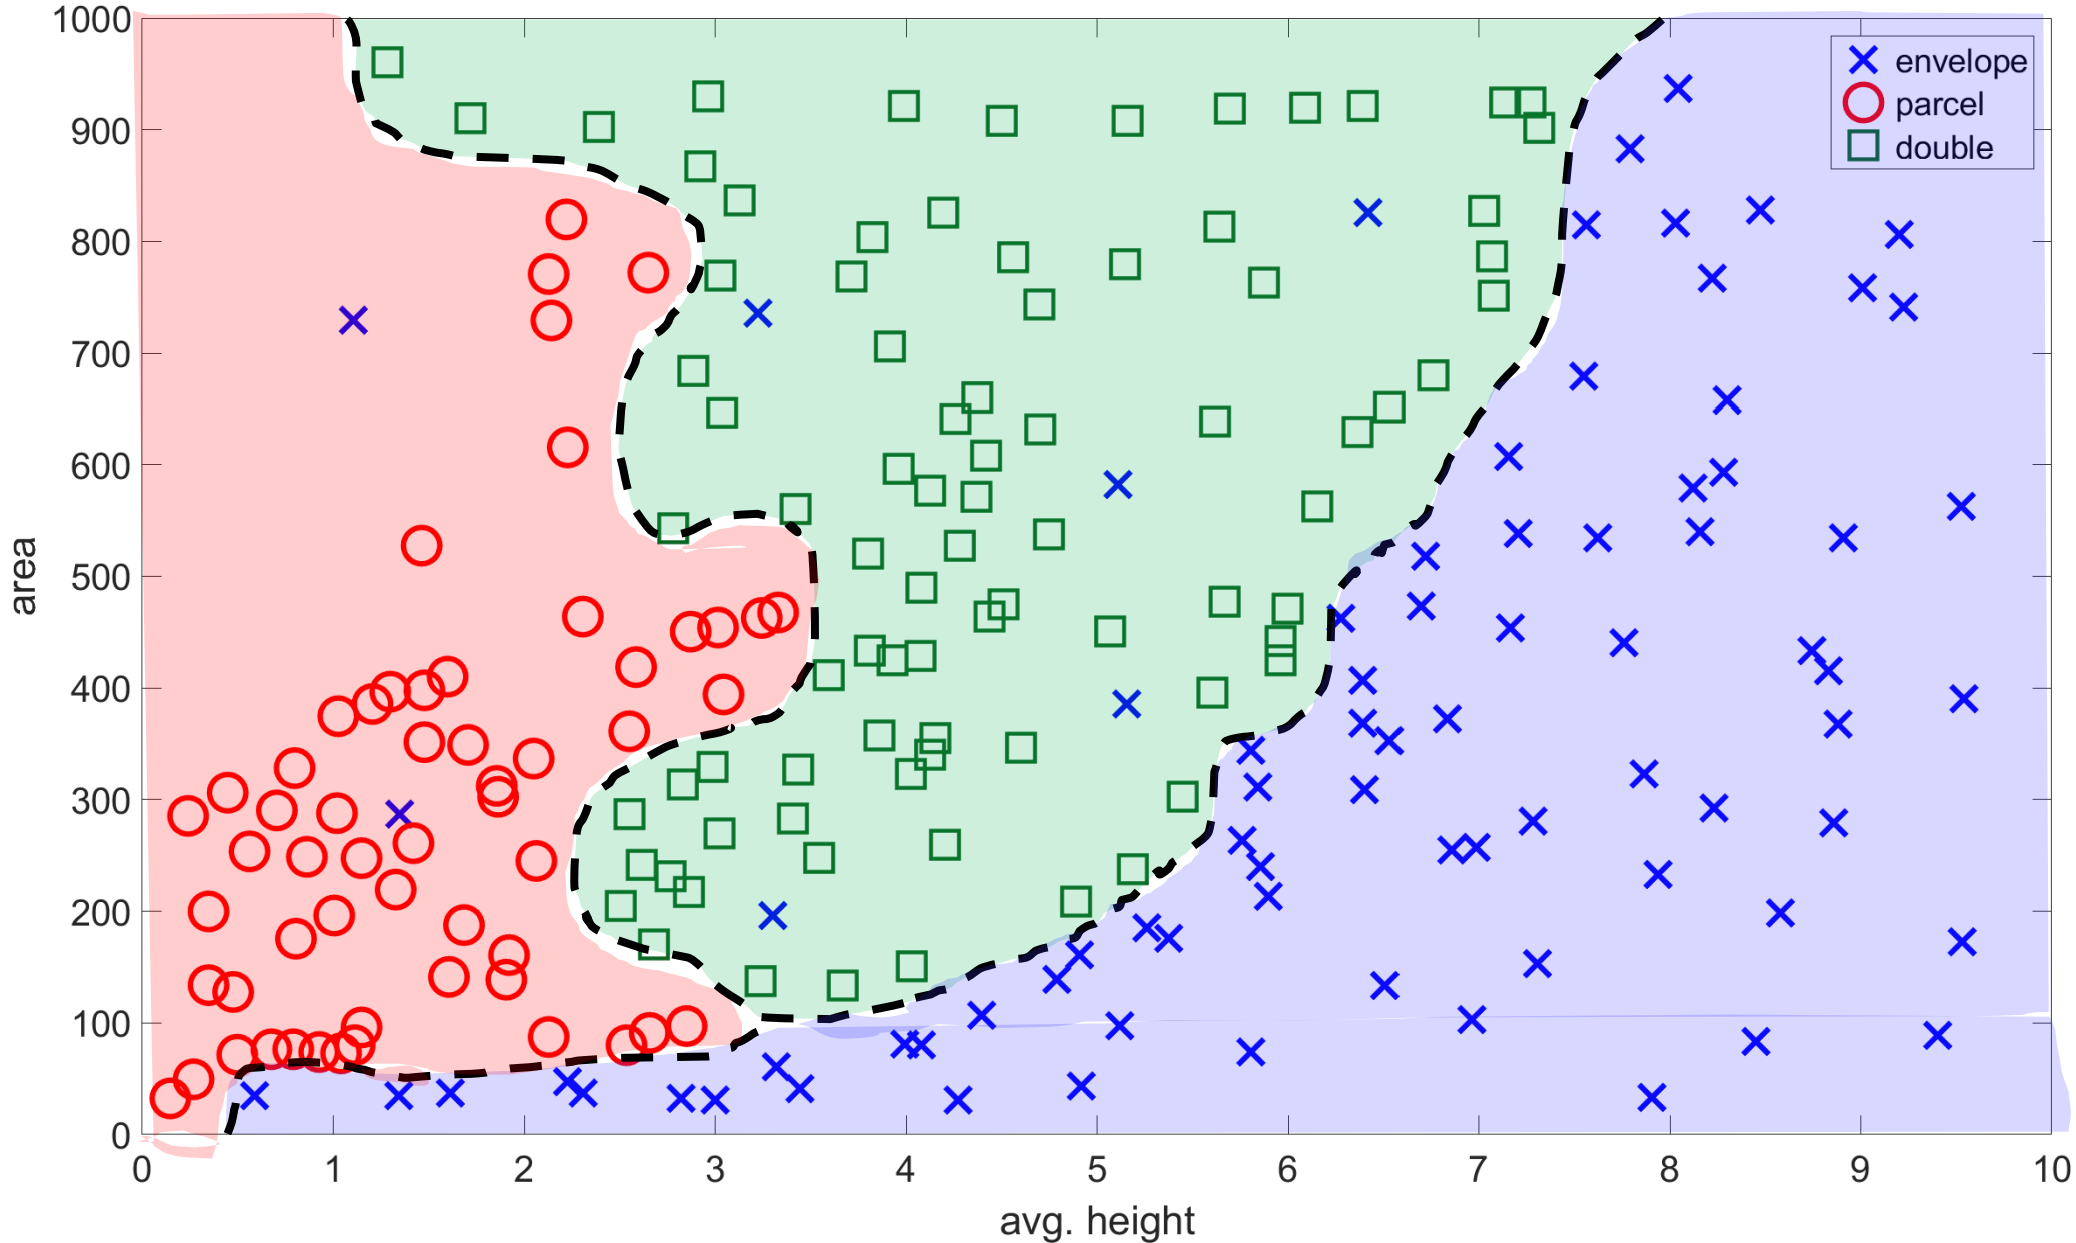

Are there better classifiers?

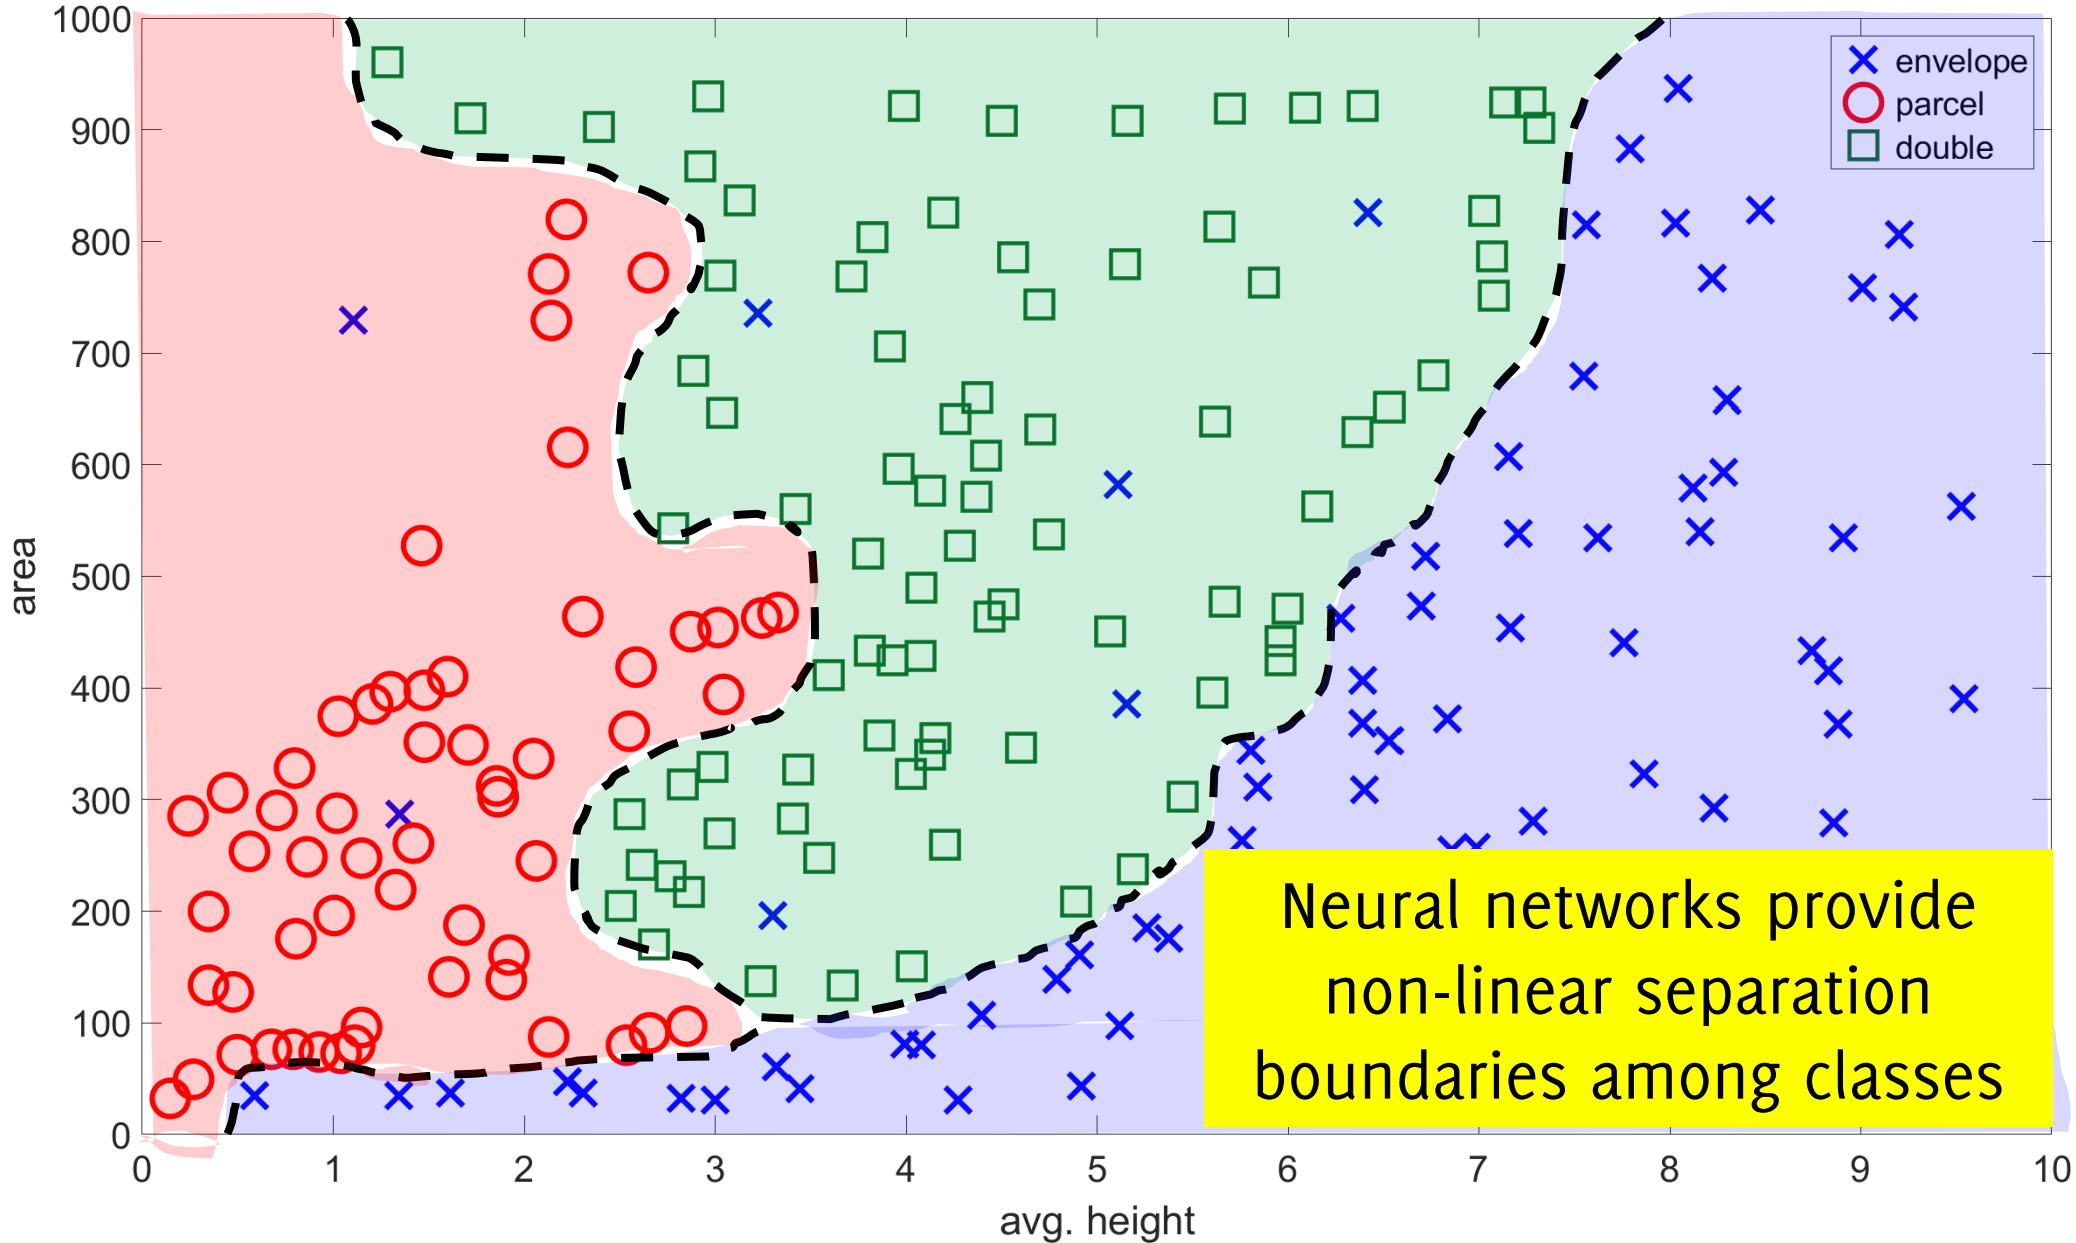

And Neural Networks are not the only..

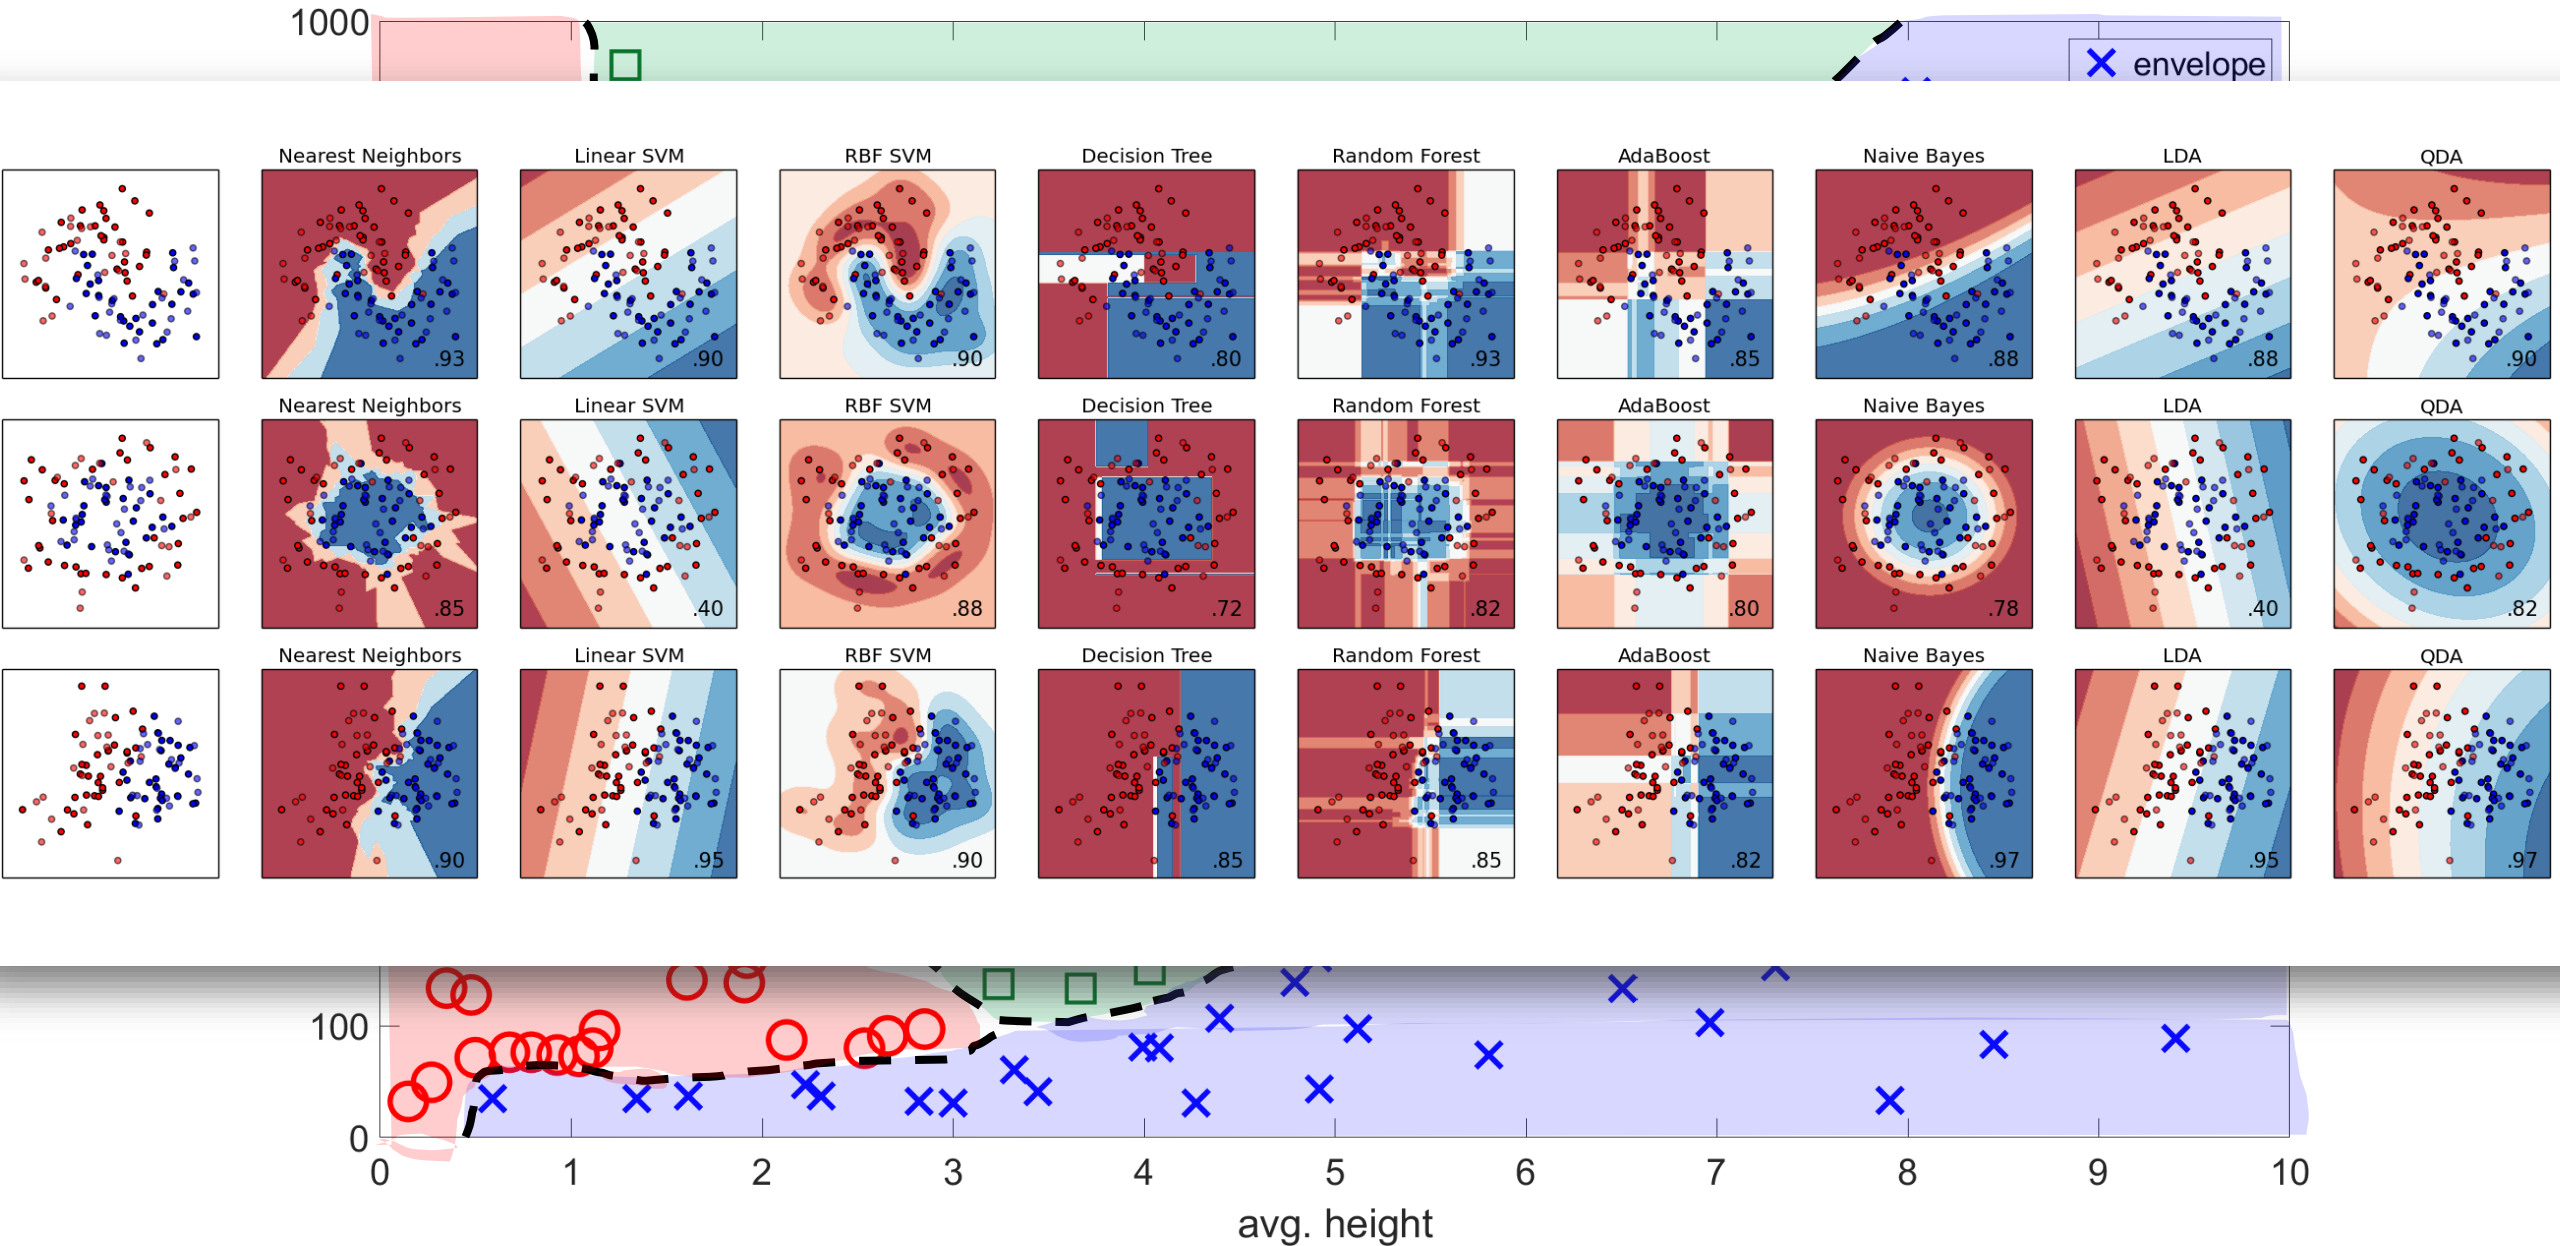

Neural Networks

Input image

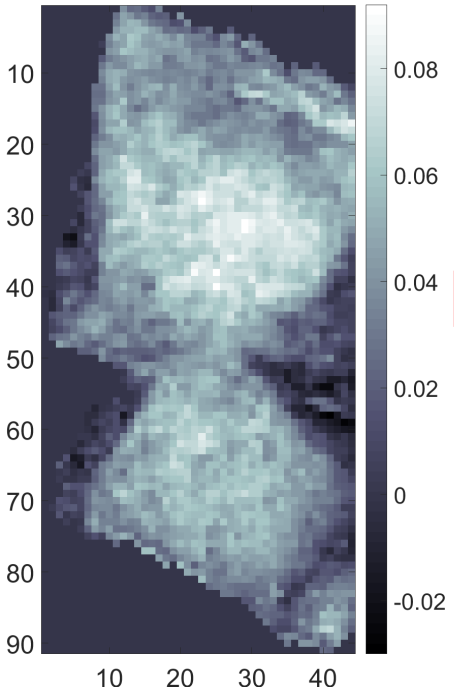

$I_1 \in \mathbb{R}^{r_1 \times c_1}$

Feature Extraction Algorithm

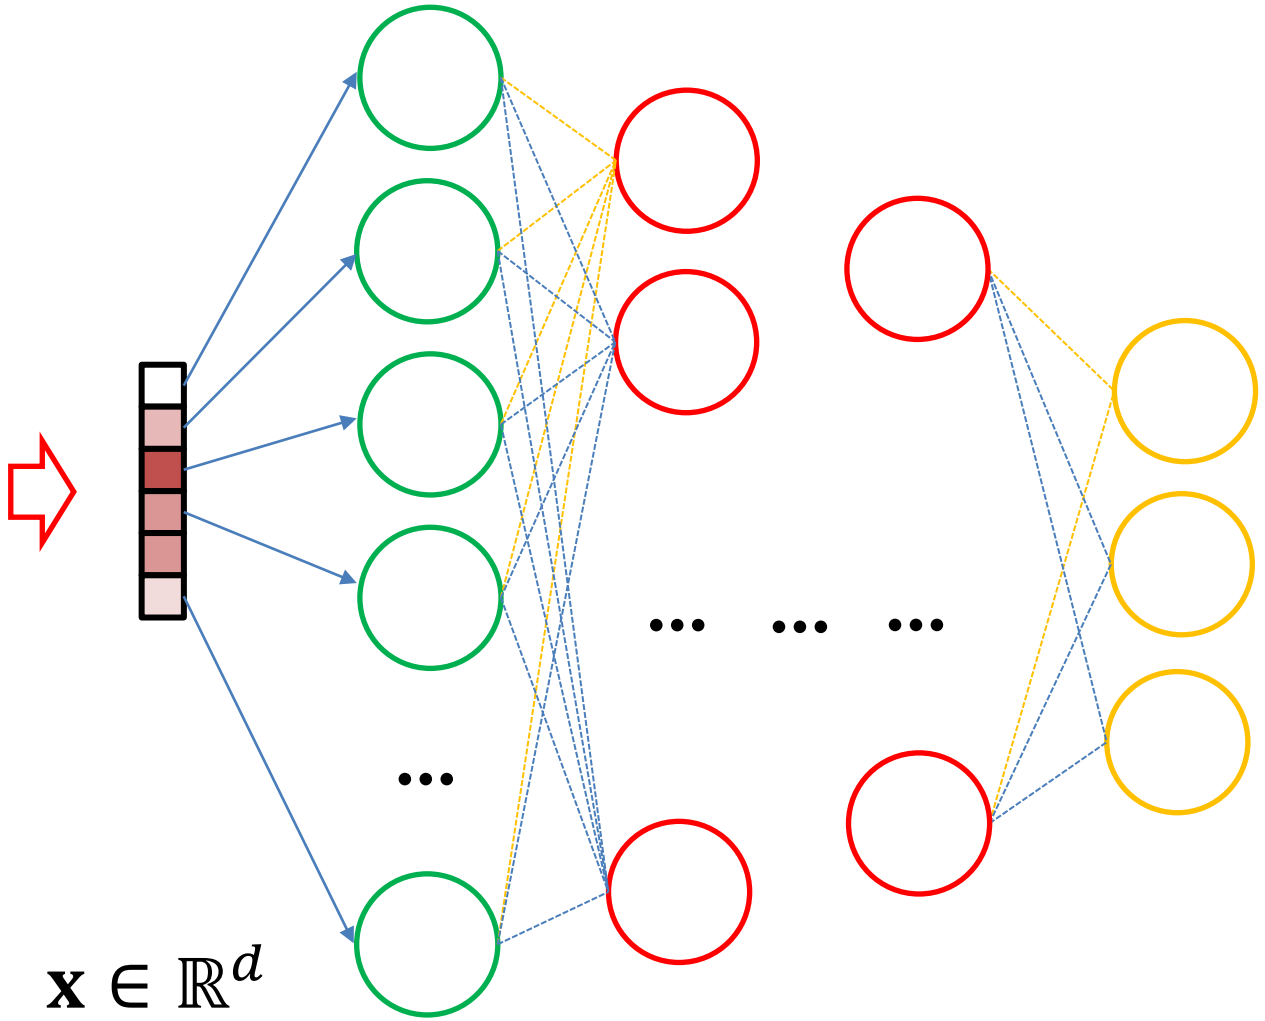

$\mathbf{x} \in \mathbb{R}^d$

input layer

Hidden layer(s)

Output Layer

Neural Networks

Input layer: Same size of the feature vector

Input image

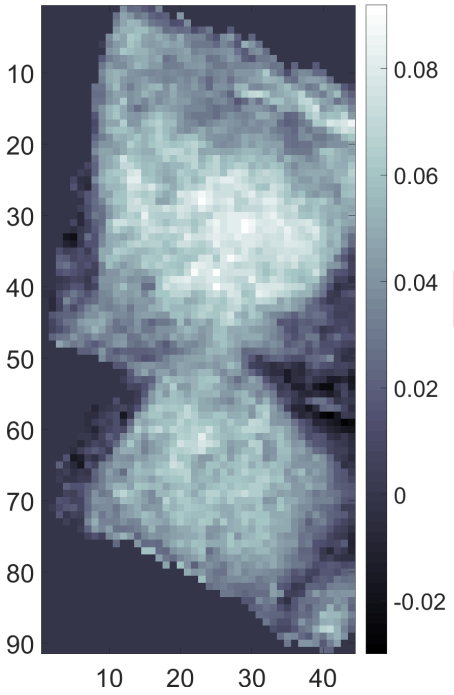

$I_1 \in \mathbb{R}^{r_1 \times c_1}$

Feature Extraction Algorithm

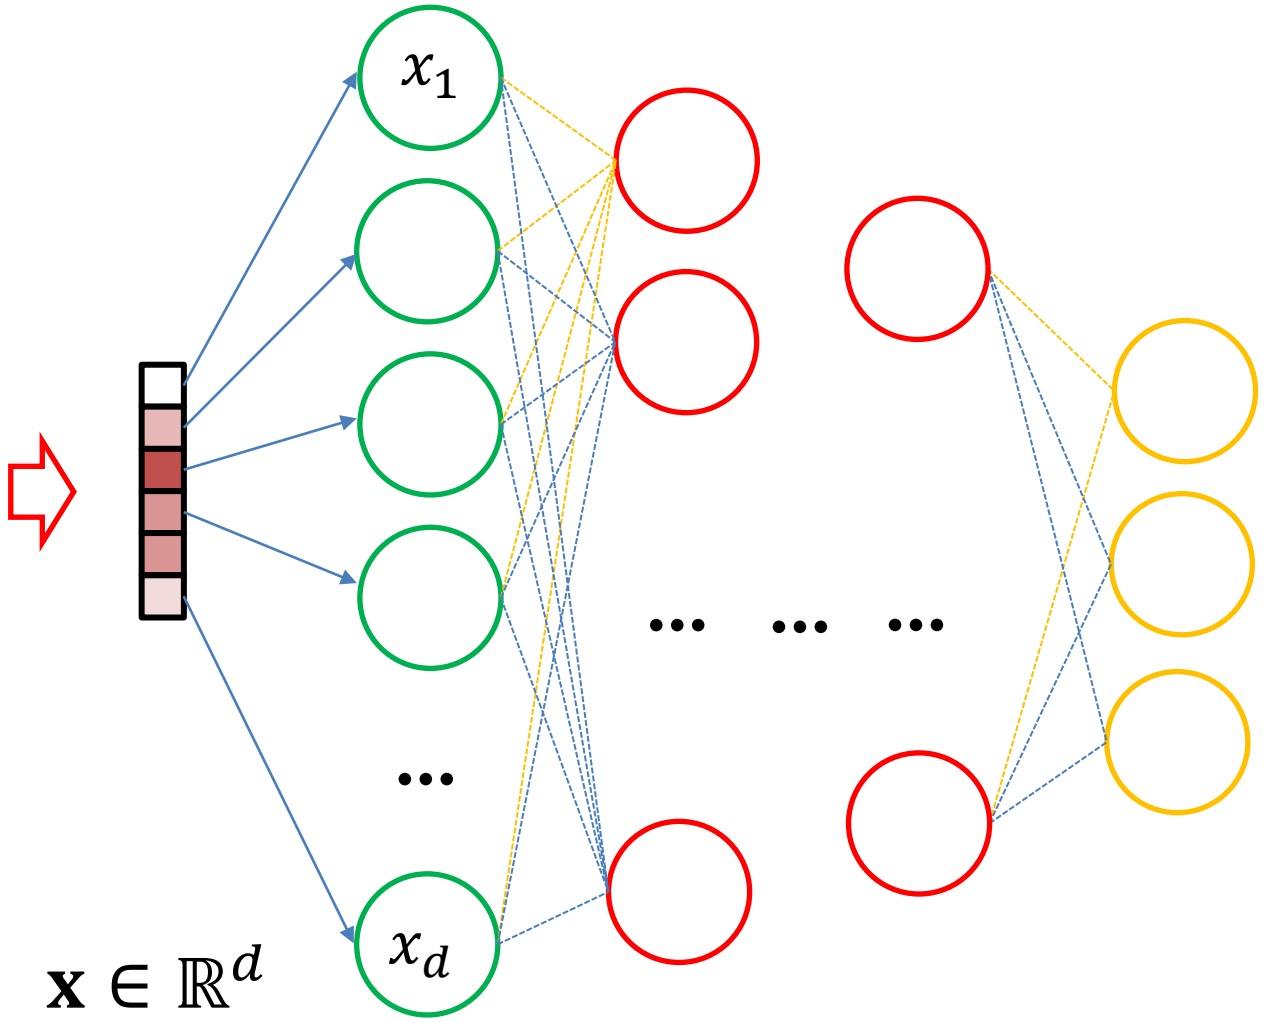

$\mathbf{x} \in \mathbb{R}^d$

input layer

Hidden layer(s)

Output Layer

Neural Networks

Output layer: Same size as the number of classes

Input image

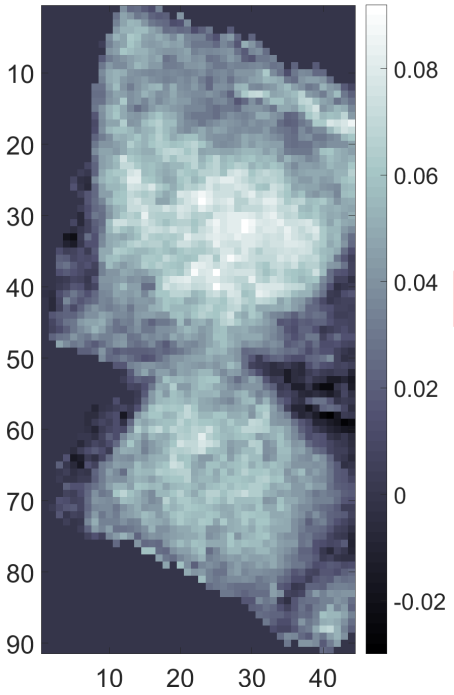

$I_1 \in \mathbb{R}^{r_1 \times c_1}$

Feature Extraction Algorithm

$\mathbf{x} \in \mathbb{R}^d$

$P(t = \text{"doub."} | \boldsymbol{x})$

$P(t = \text{"env."} | \boldsymbol{x})$

$P(t = \text{"parc."} | \boldsymbol{x})$

input layer

Hidden layer(s)

Output Layer

Neural Networks

Hidden layers: arbitrary size

Input image

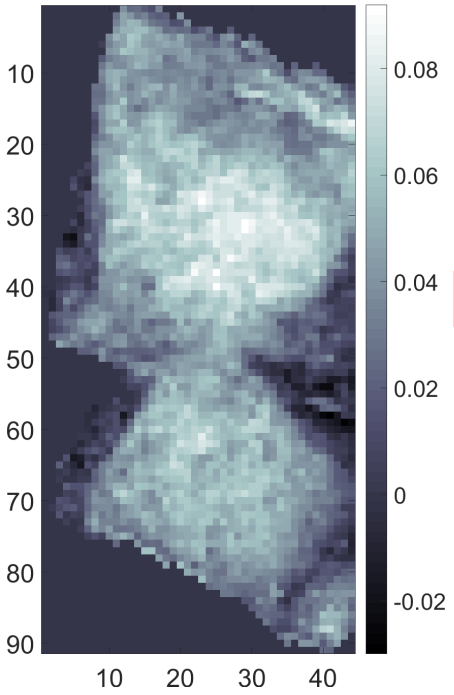

$I_1 \in \mathbb{R}^{r_1 \times c_1}$

Feature Extraction Algorithm

$\mathbf{x} \in \mathbb{R}^d$

input layer

Hidden layer(s)

Output Layer

Image Classification by Hand Crafted Features

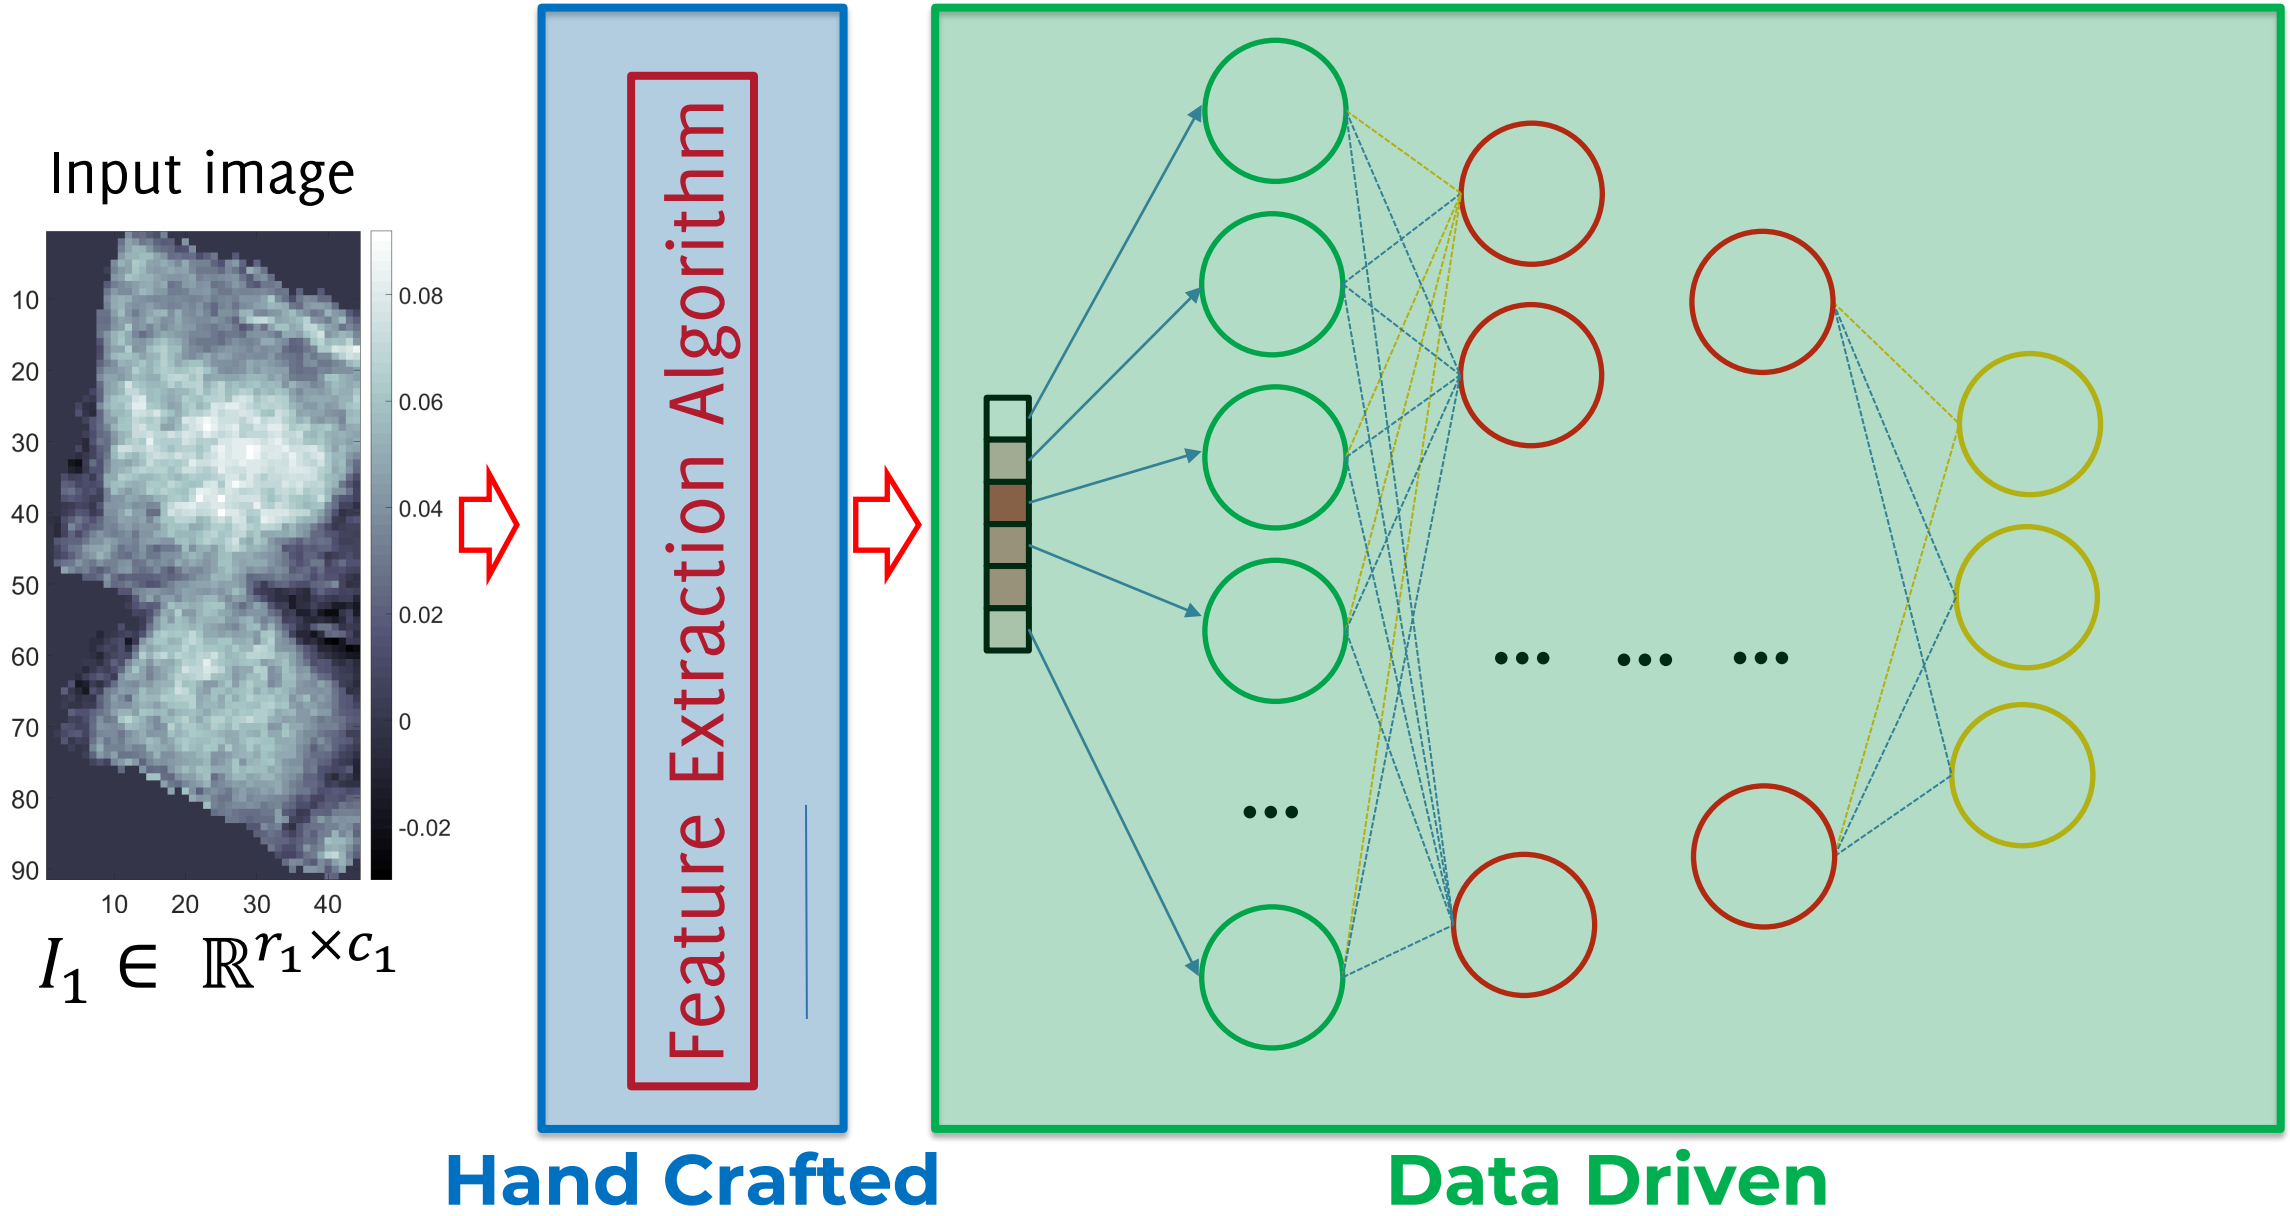

Hand Crafted Featues, pros:

- **Exploit a priori** / expert **information**
- Features are **interpretable** (you might understand why they are not working)
- You can **adjust features** to improve your performance
- **Limited amount of training data** needed
- You can give more relevance to some features

Hand Crafted Featues, cons:

- Requires a lot of **design/programming efforts**
- **Not viable** in many **visual recognition** tasks that are easily performed by humans (e.g. when dealing with natural images)
- **Risk of overfitting** the training set used in the feature design
- **Not very general and "portable"**

What is Deep Learning after all?

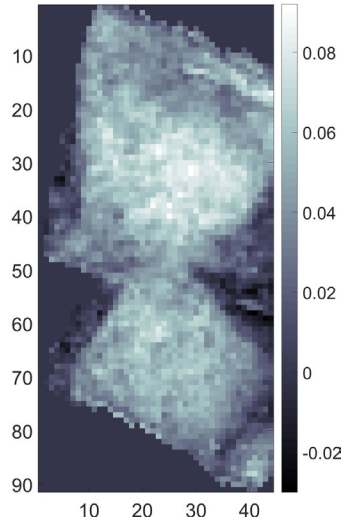

Machine learns how to take the Iris apart

Hand-crafted Features

Learned Classifier

double

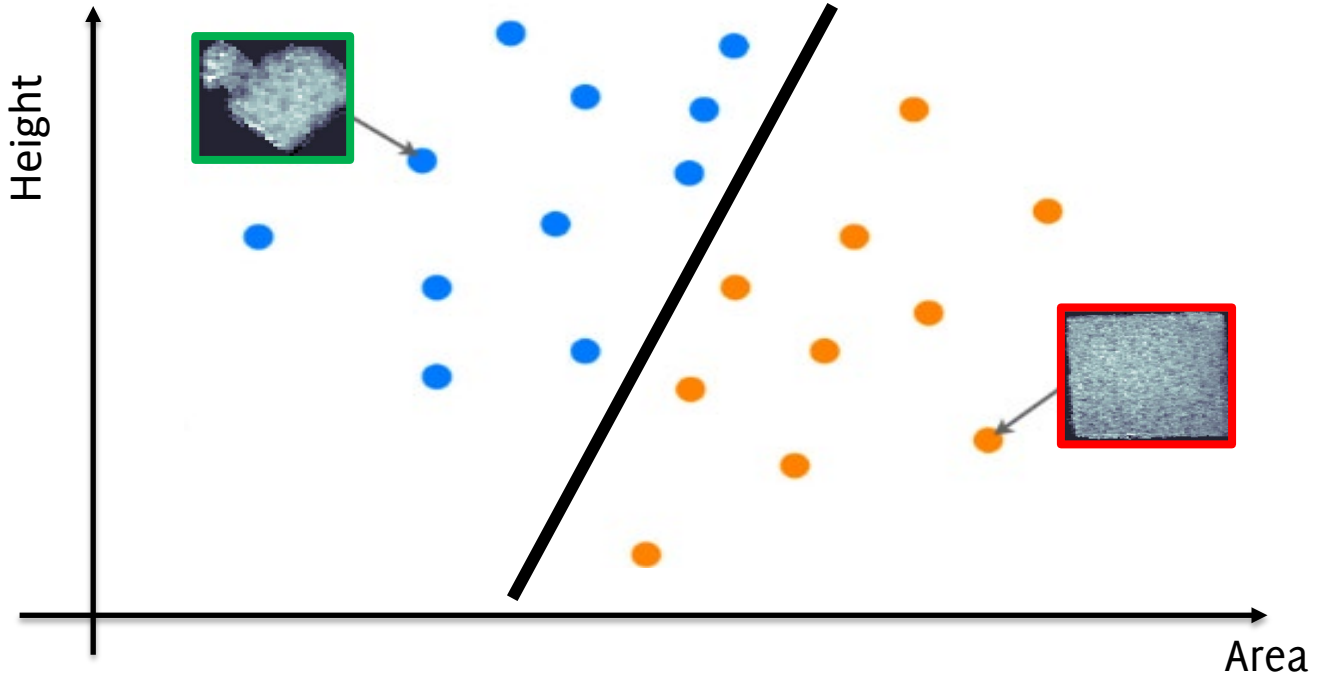

What is Deep Learning after all?

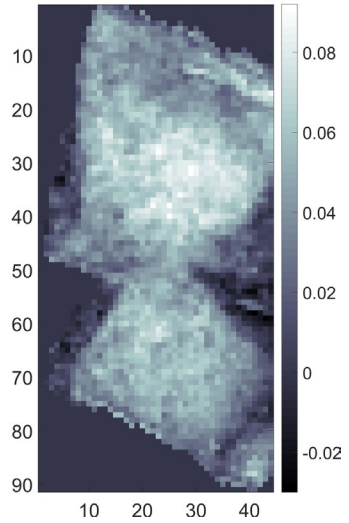

Machine learns how to take the Iris apart

Hand-crafted Features

Learned Classifier

double

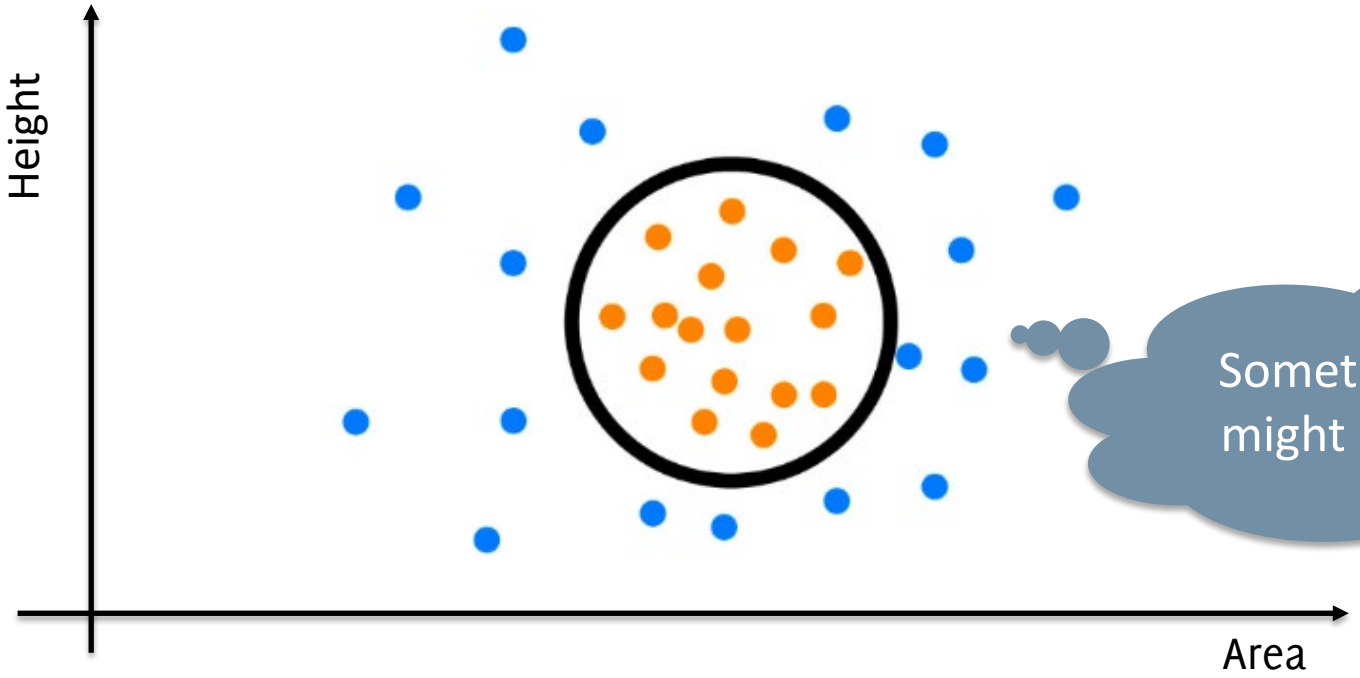

Sometimes the decision might be more complex

What is Deep Learning after all?

Hand-crafted Features

Learned Classifier

double

Machine learns how to take the Iris apart

Sometimes the decision might be Impossible!

What is Deep Learning after all?

This happens if you do not know which features to extract!!!

Hand-crafted Features

Learned Classifier

double

Height

Area

Sometimes the decision might be Impossible!

Data Driven Features

… the advent of Deep Learning

Data-Driven Features

Input image

$I_1 \in \mathbb{R}^{r_1 \times c_1}$

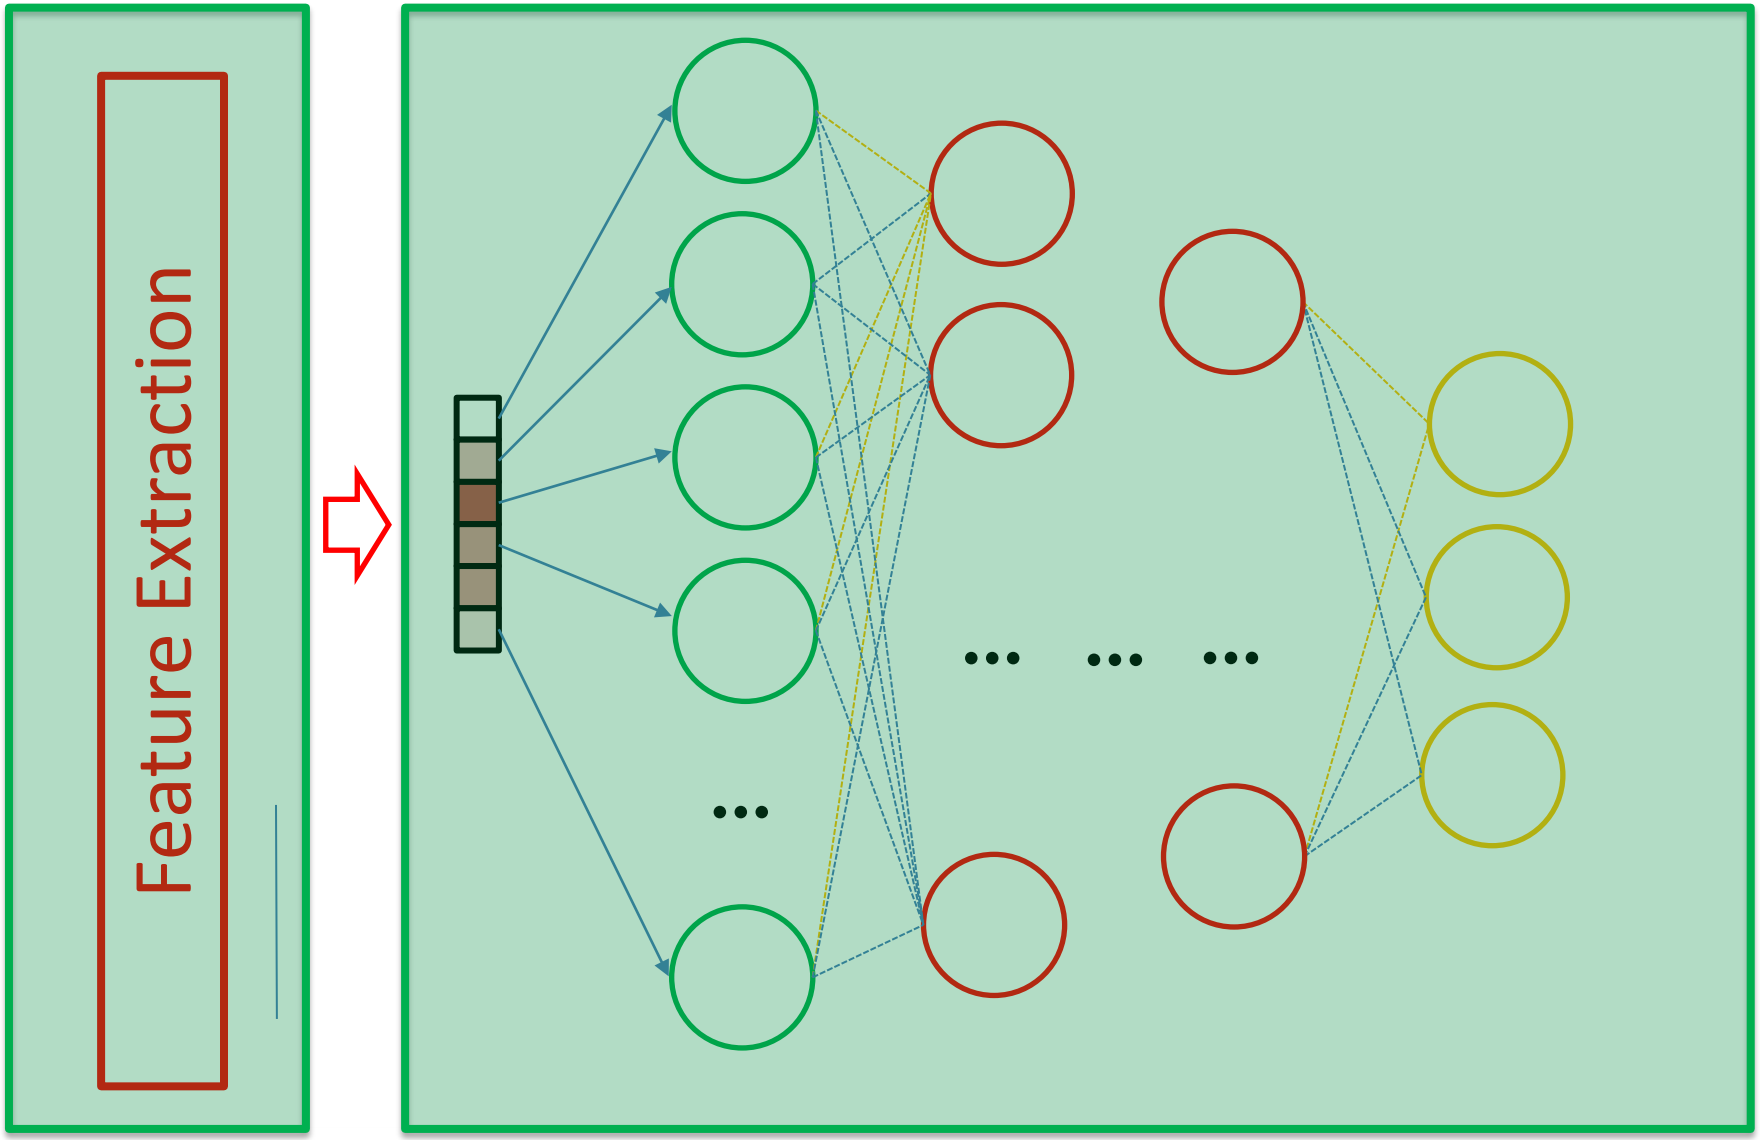

What is Deep Learning after all?

Optimized for the task!

Machine Learned Features

Learned Classifier

double

Easier to learn!

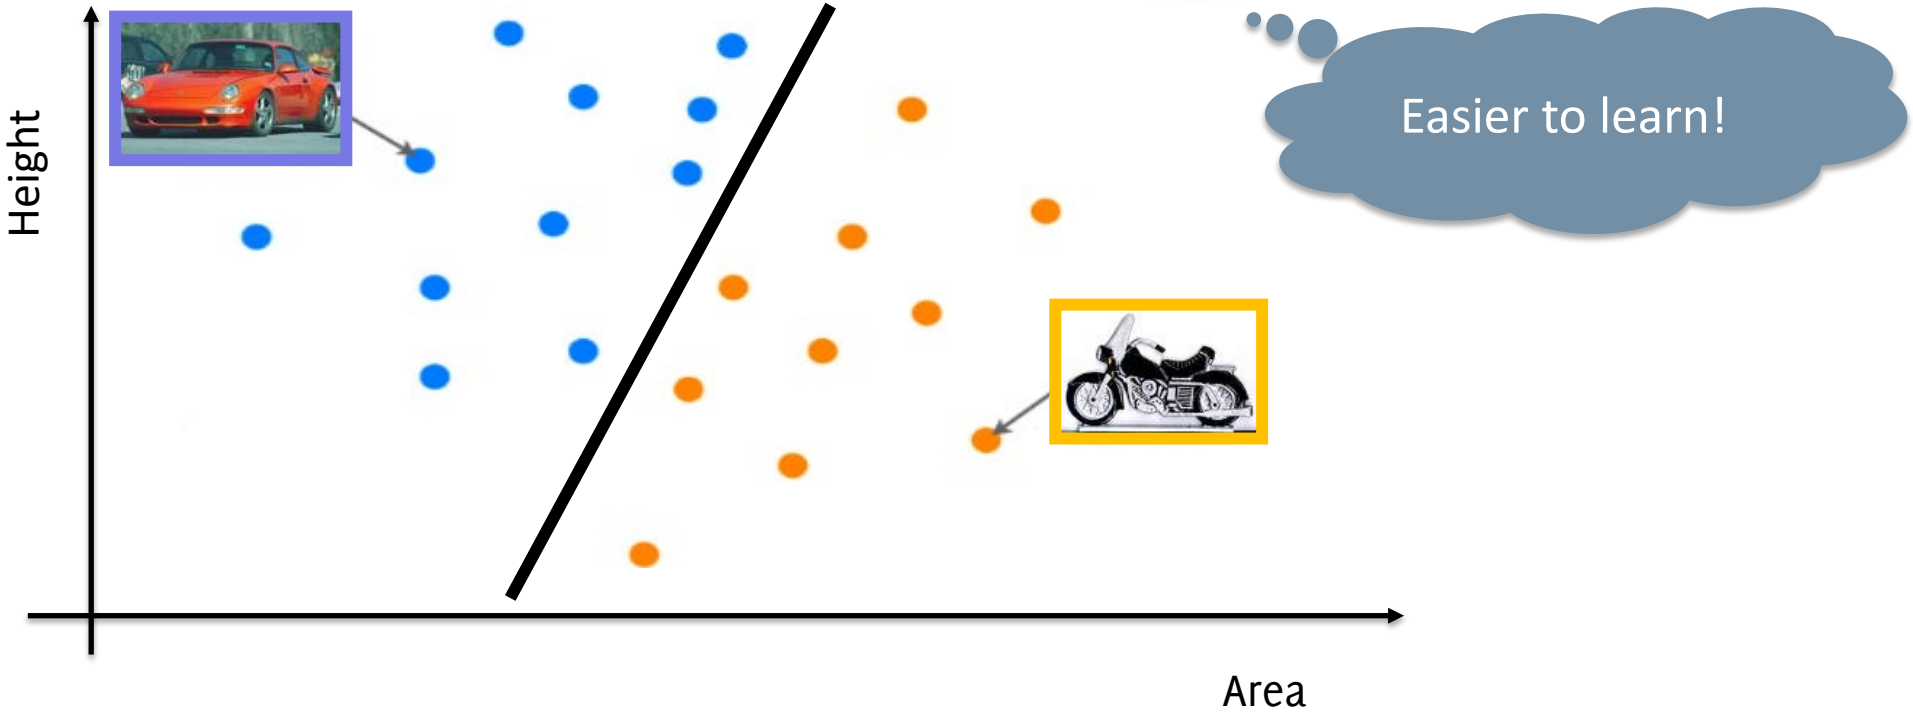

What is Deep Learning after all?

Learn from data!

Hierarchical representation optimized for the task!

Learned features

Learned features

Learned features

Learned Classifier

double

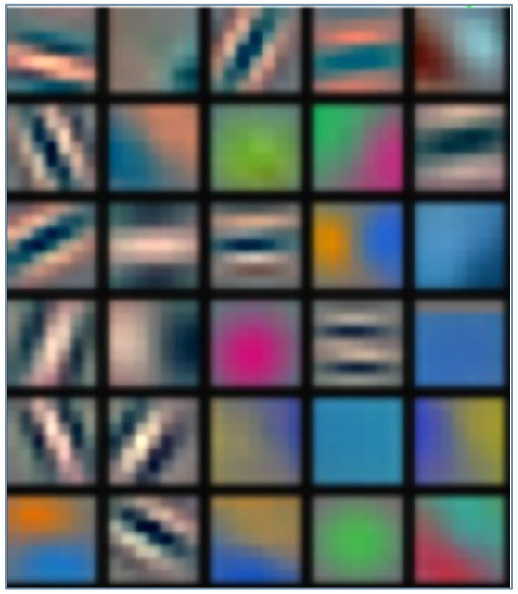

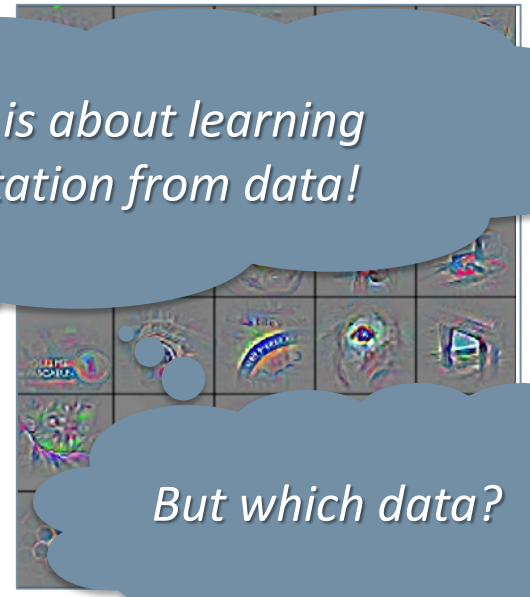

Deep Learning is about learning data representation from data!

But which data?

Deep Learning

a Breakthrough in Visual Recognition

Image Classfiication on Imagenet

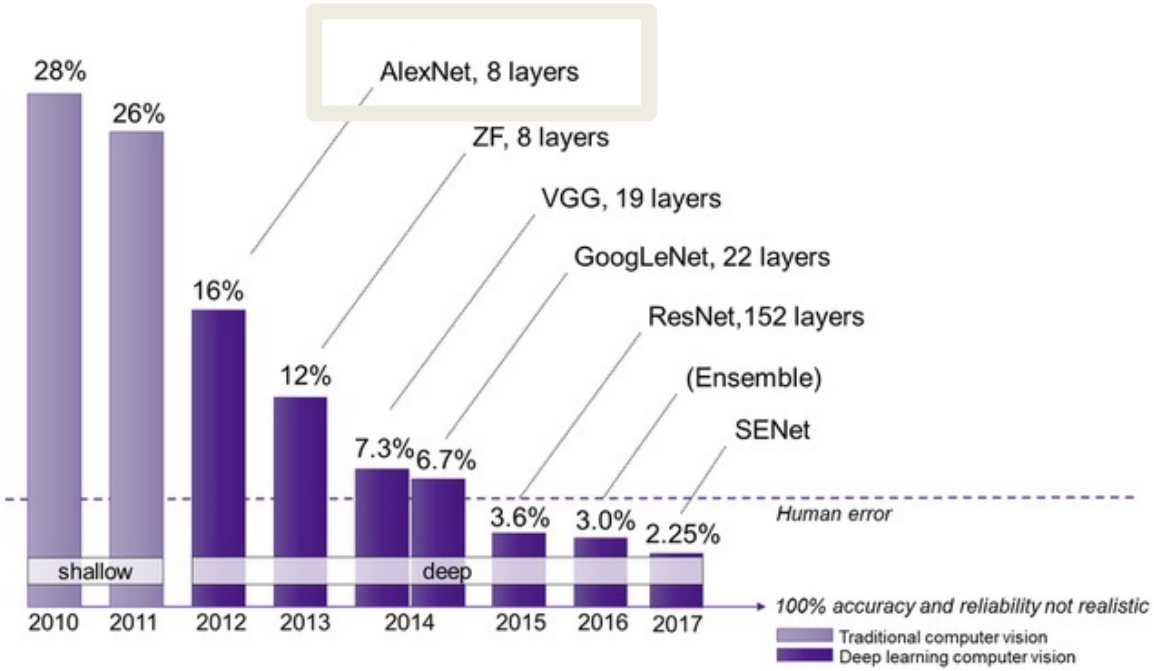

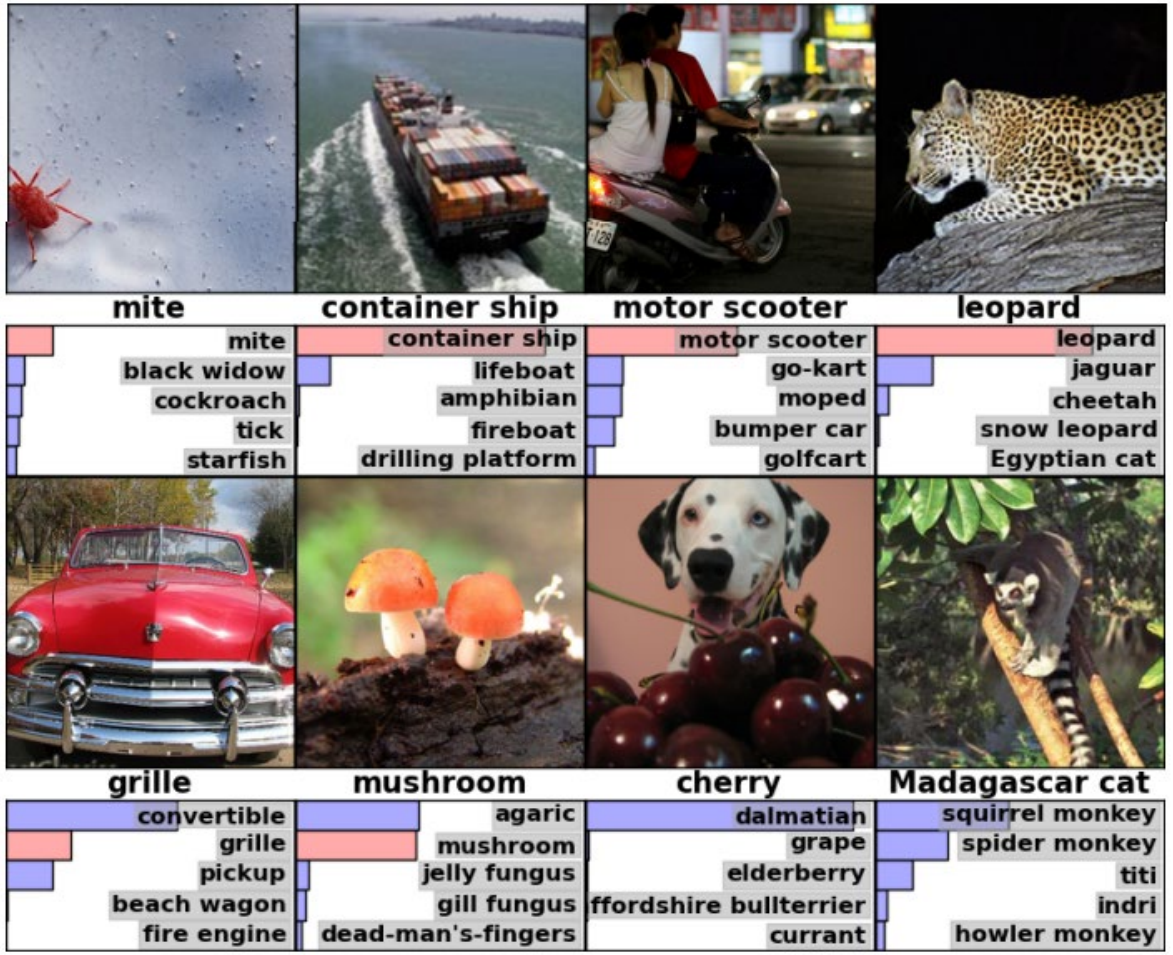

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems* 25 (2012).

The impact of Deep Learning in Visual Recognition

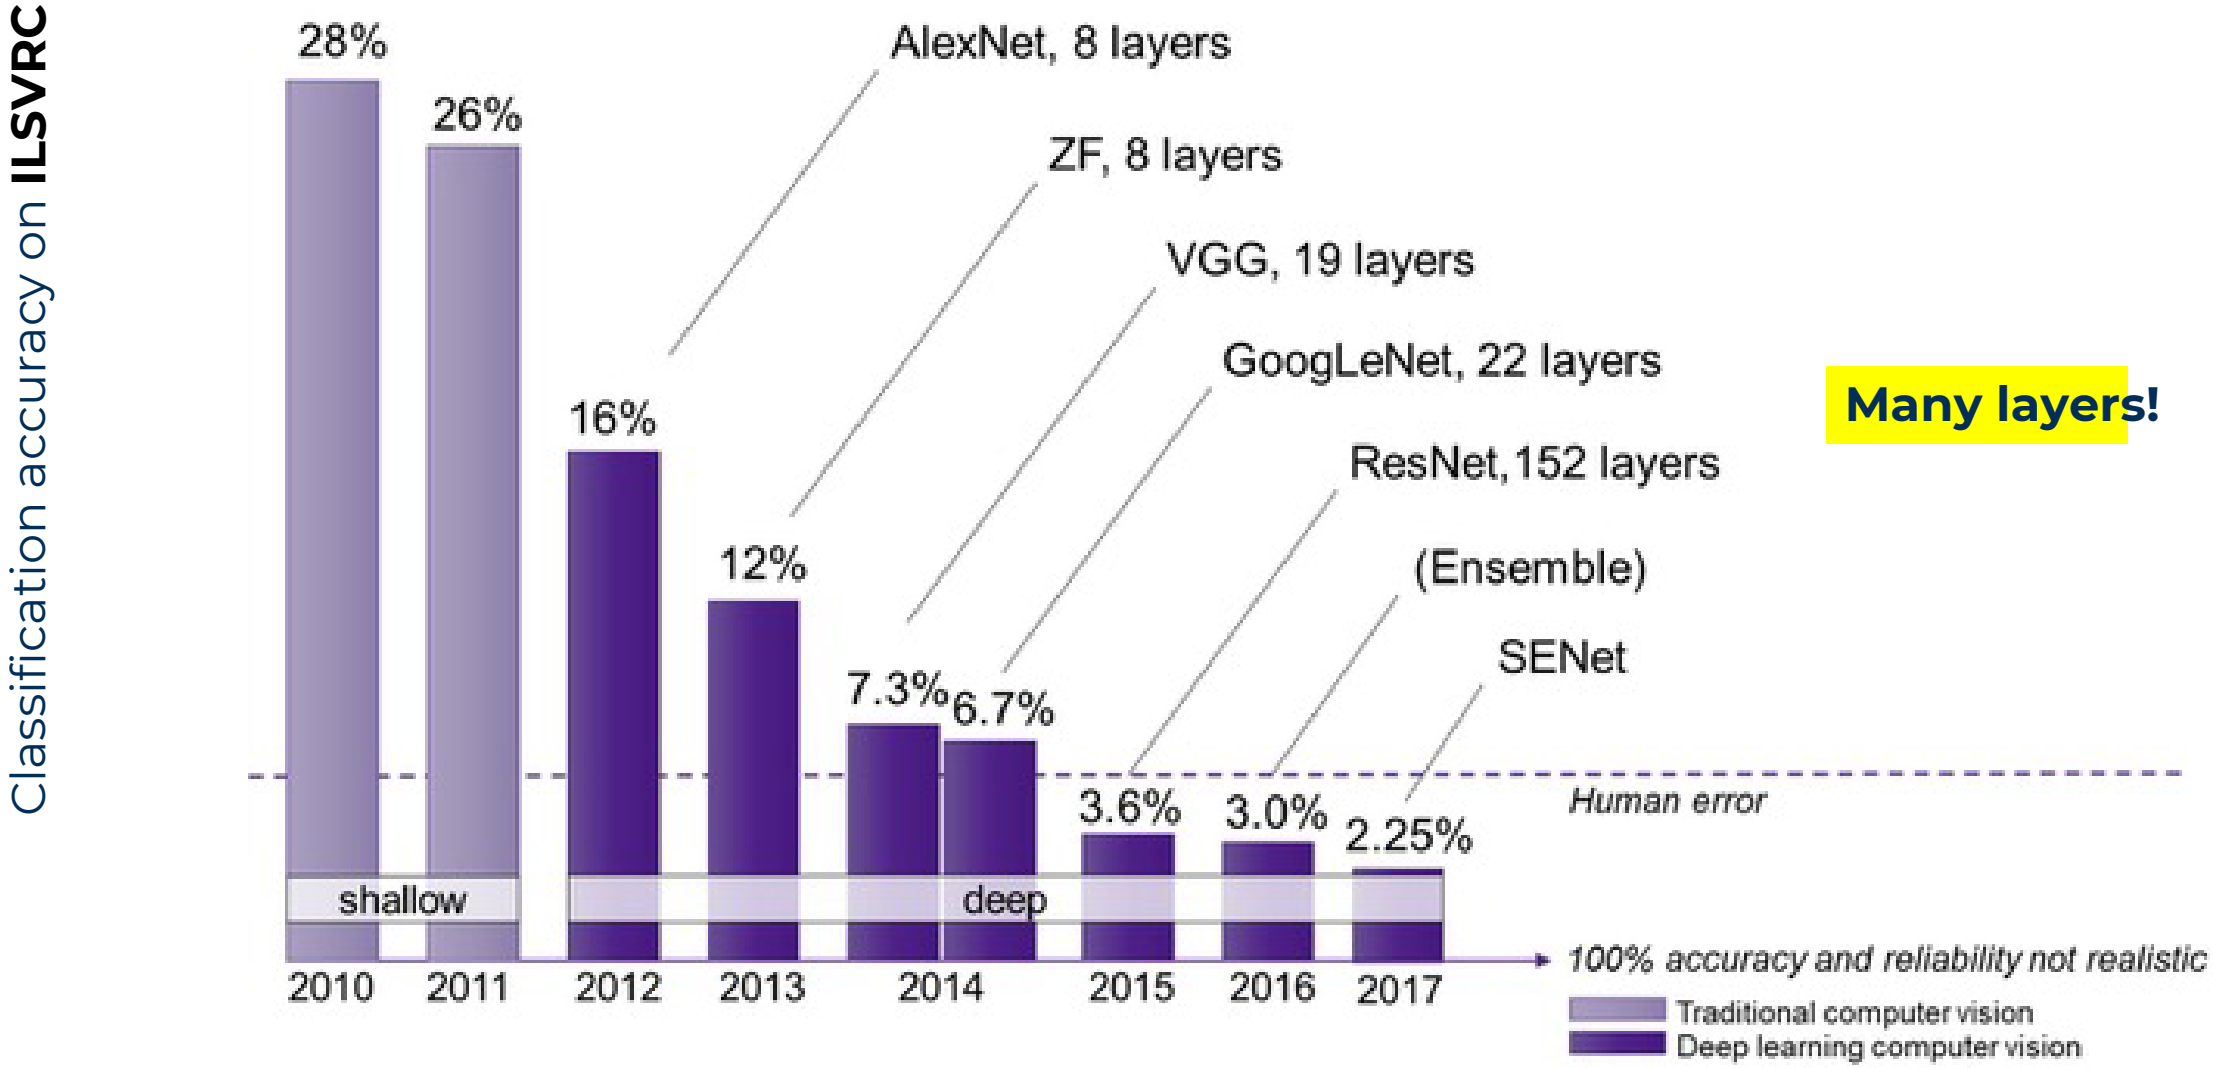

Many layers!

ILSVCR: ImageNet Large Scale Visual Recognition Challenge

The impact of Deep Learning in Visual Recognition

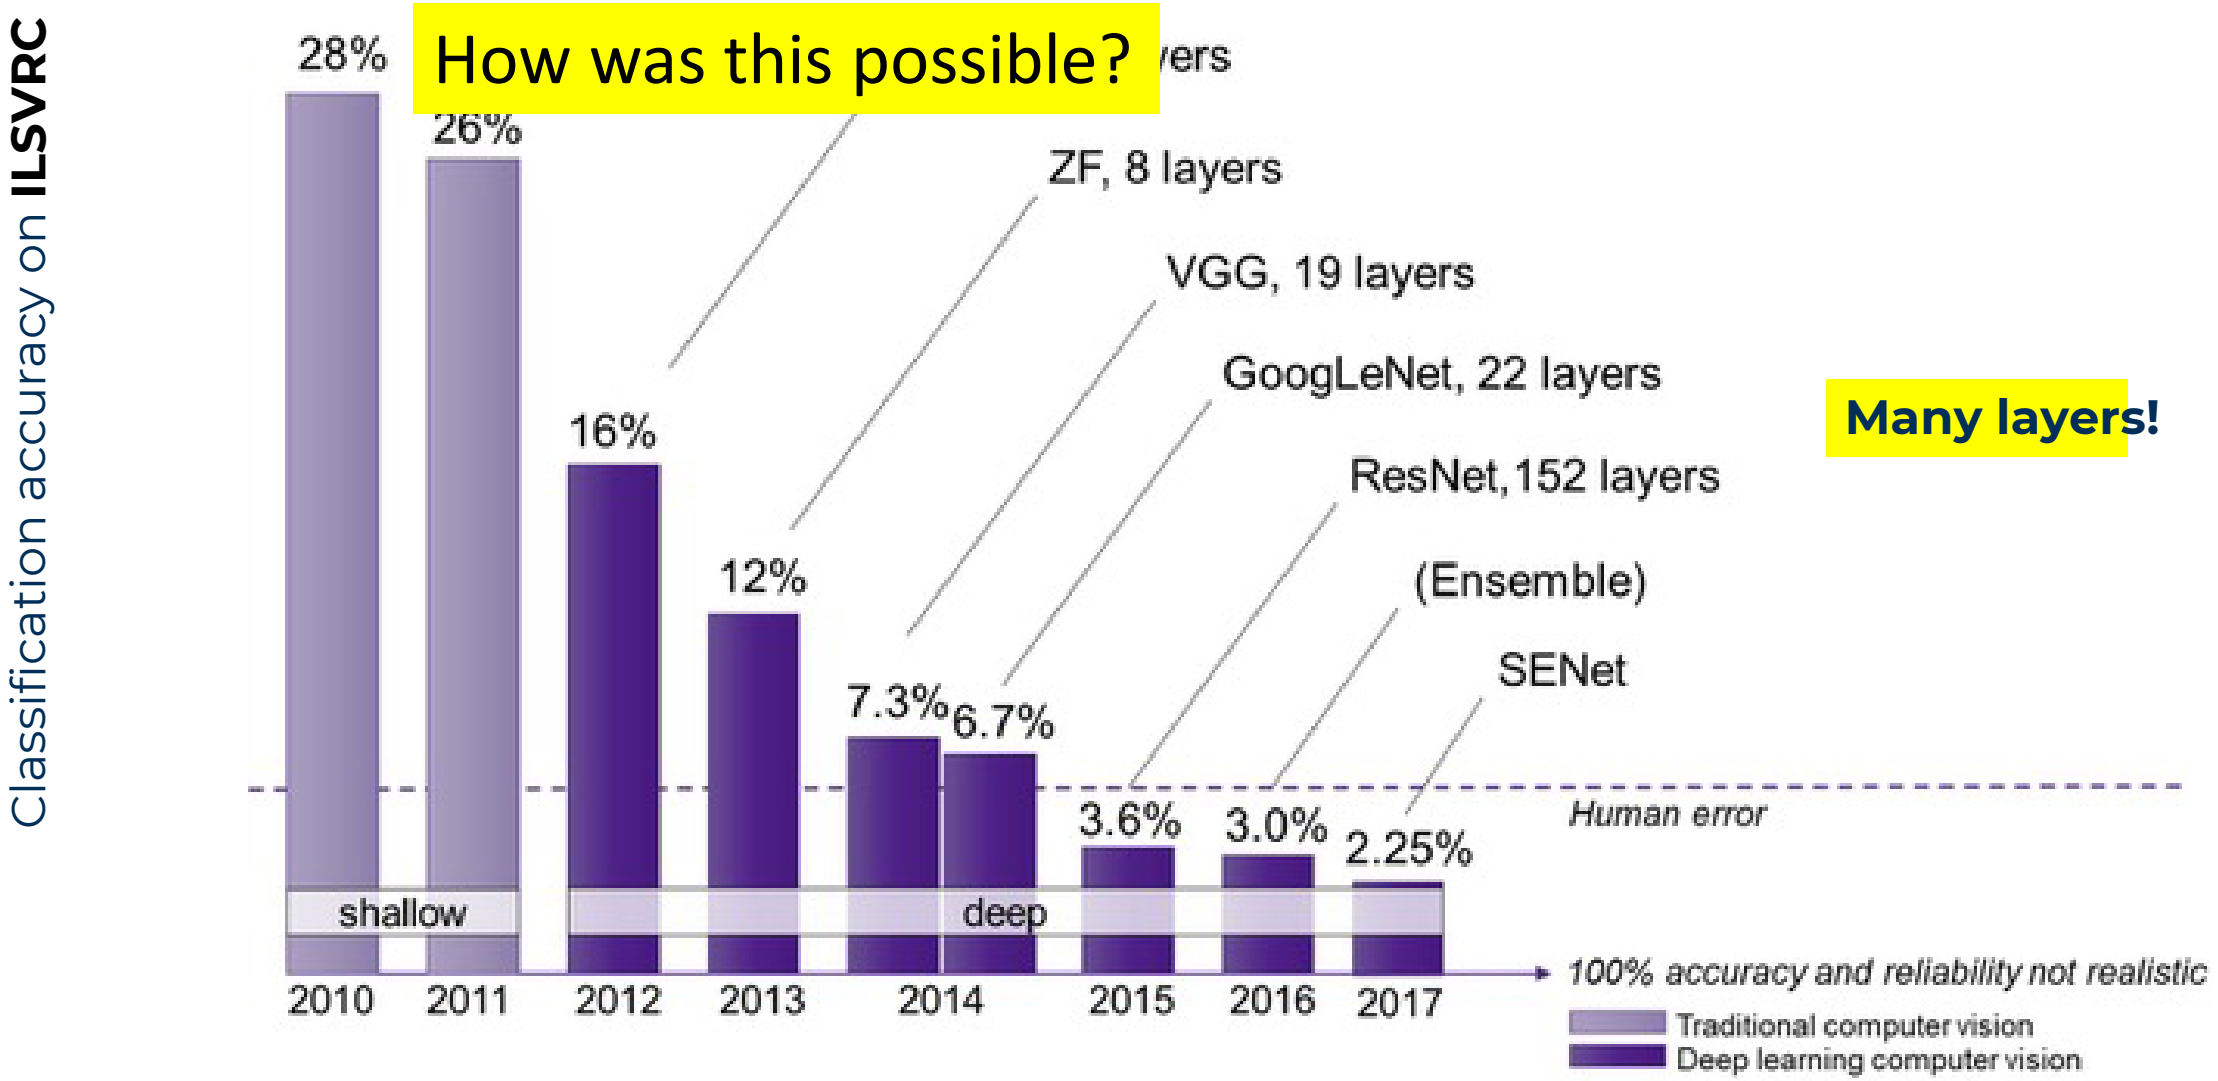

ILSVCR: ImageNet Large Scale Visual Recognition Challenge

Large Collections of Annotated Data

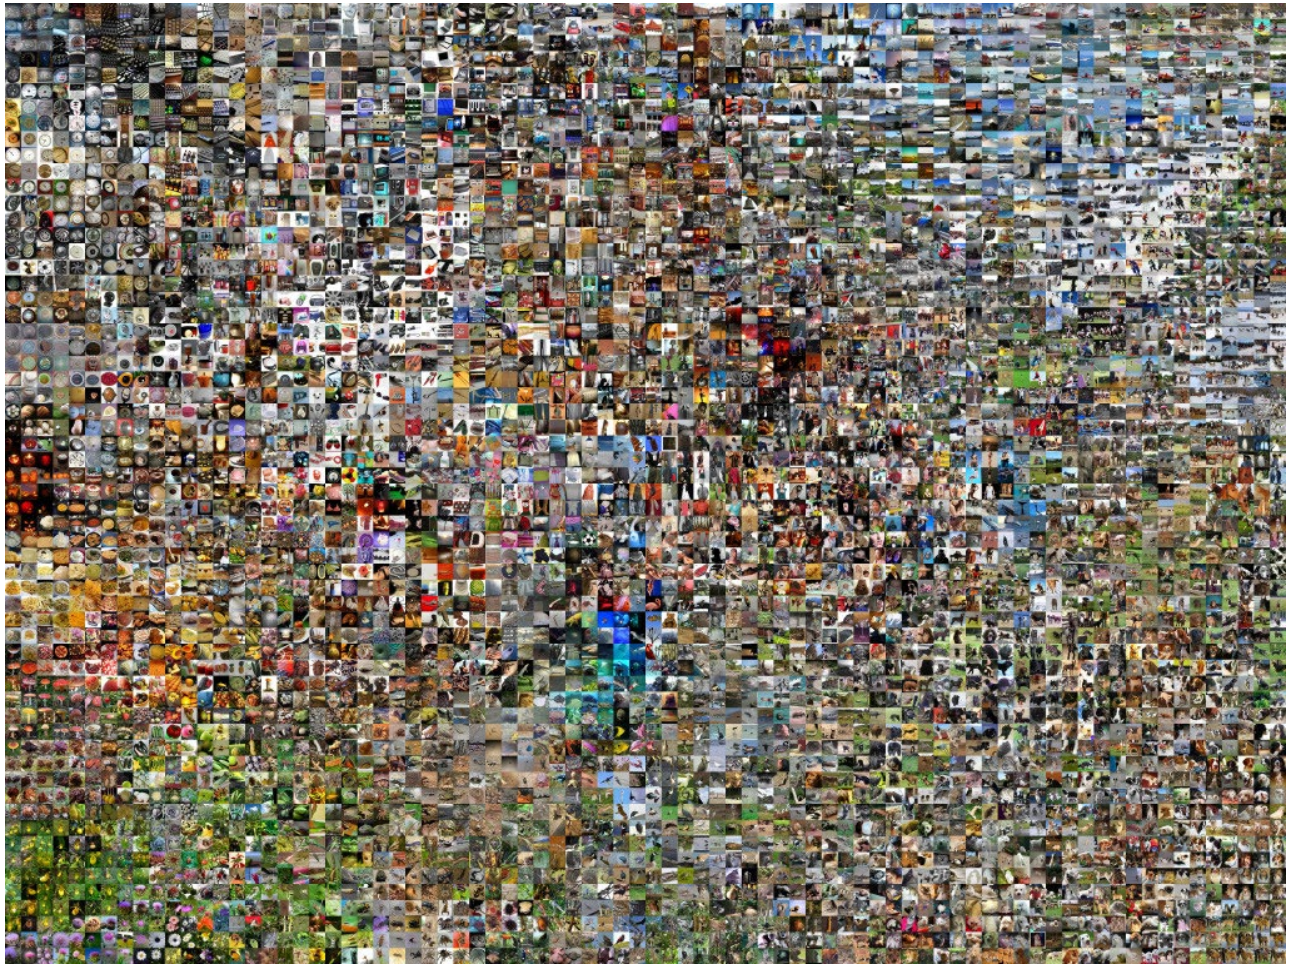

The ImageNet project is a large visual database designed for use in visual object recognition software research. ***More than 14 million images*** *have been hand-annotated by the project to indicate what objects are pictured and in at least one million of the images, bounding boxes are also provided.[3]* ***ImageNet contains more than 20,000 categories***

From Wikipedia October 2021

J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li and L. Fei-Fei, ImageNet: A Large-Scale Hierarchical Image Database. *CVPR, 2009.*

Parallel Computing Architectures

And more recently…. Software libraries

TensorFlow

PyTorch

Google LLC, Public domain, via Wikimedia Commons

PyTorch, BSD <http://opensource.org/licenses/bsd-license.php>, via Wikimedia Commons

And of course "New" Network Architectures…

…but these were around since ‘97

You will learn to read this!
(required to pass the exam)

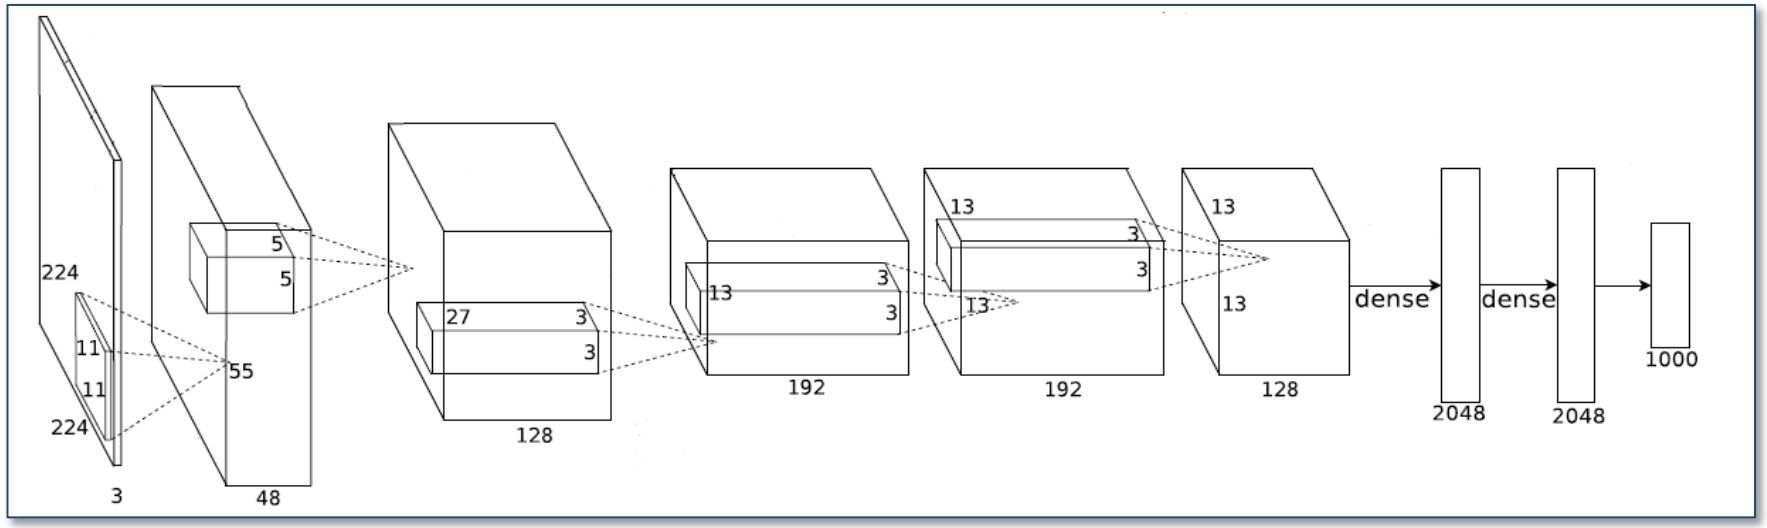

LeCun, Y., Bottou, L., Bengio, Y., Haffner, P. *“Gradient-based learning applied to document recognition”* Proceedings of the IEEE, 1998 86(11), 2278-2324.

Fathers of the Deep Learning Revolution Receive ACM A.M. Turing Award

Bengio, Hinton and LeCun Ushered in Major Breakthroughs in Artificial Intelligence

https://awards.acm.org/about/2018-turing

Advanced Visual Recognition Problems with DL

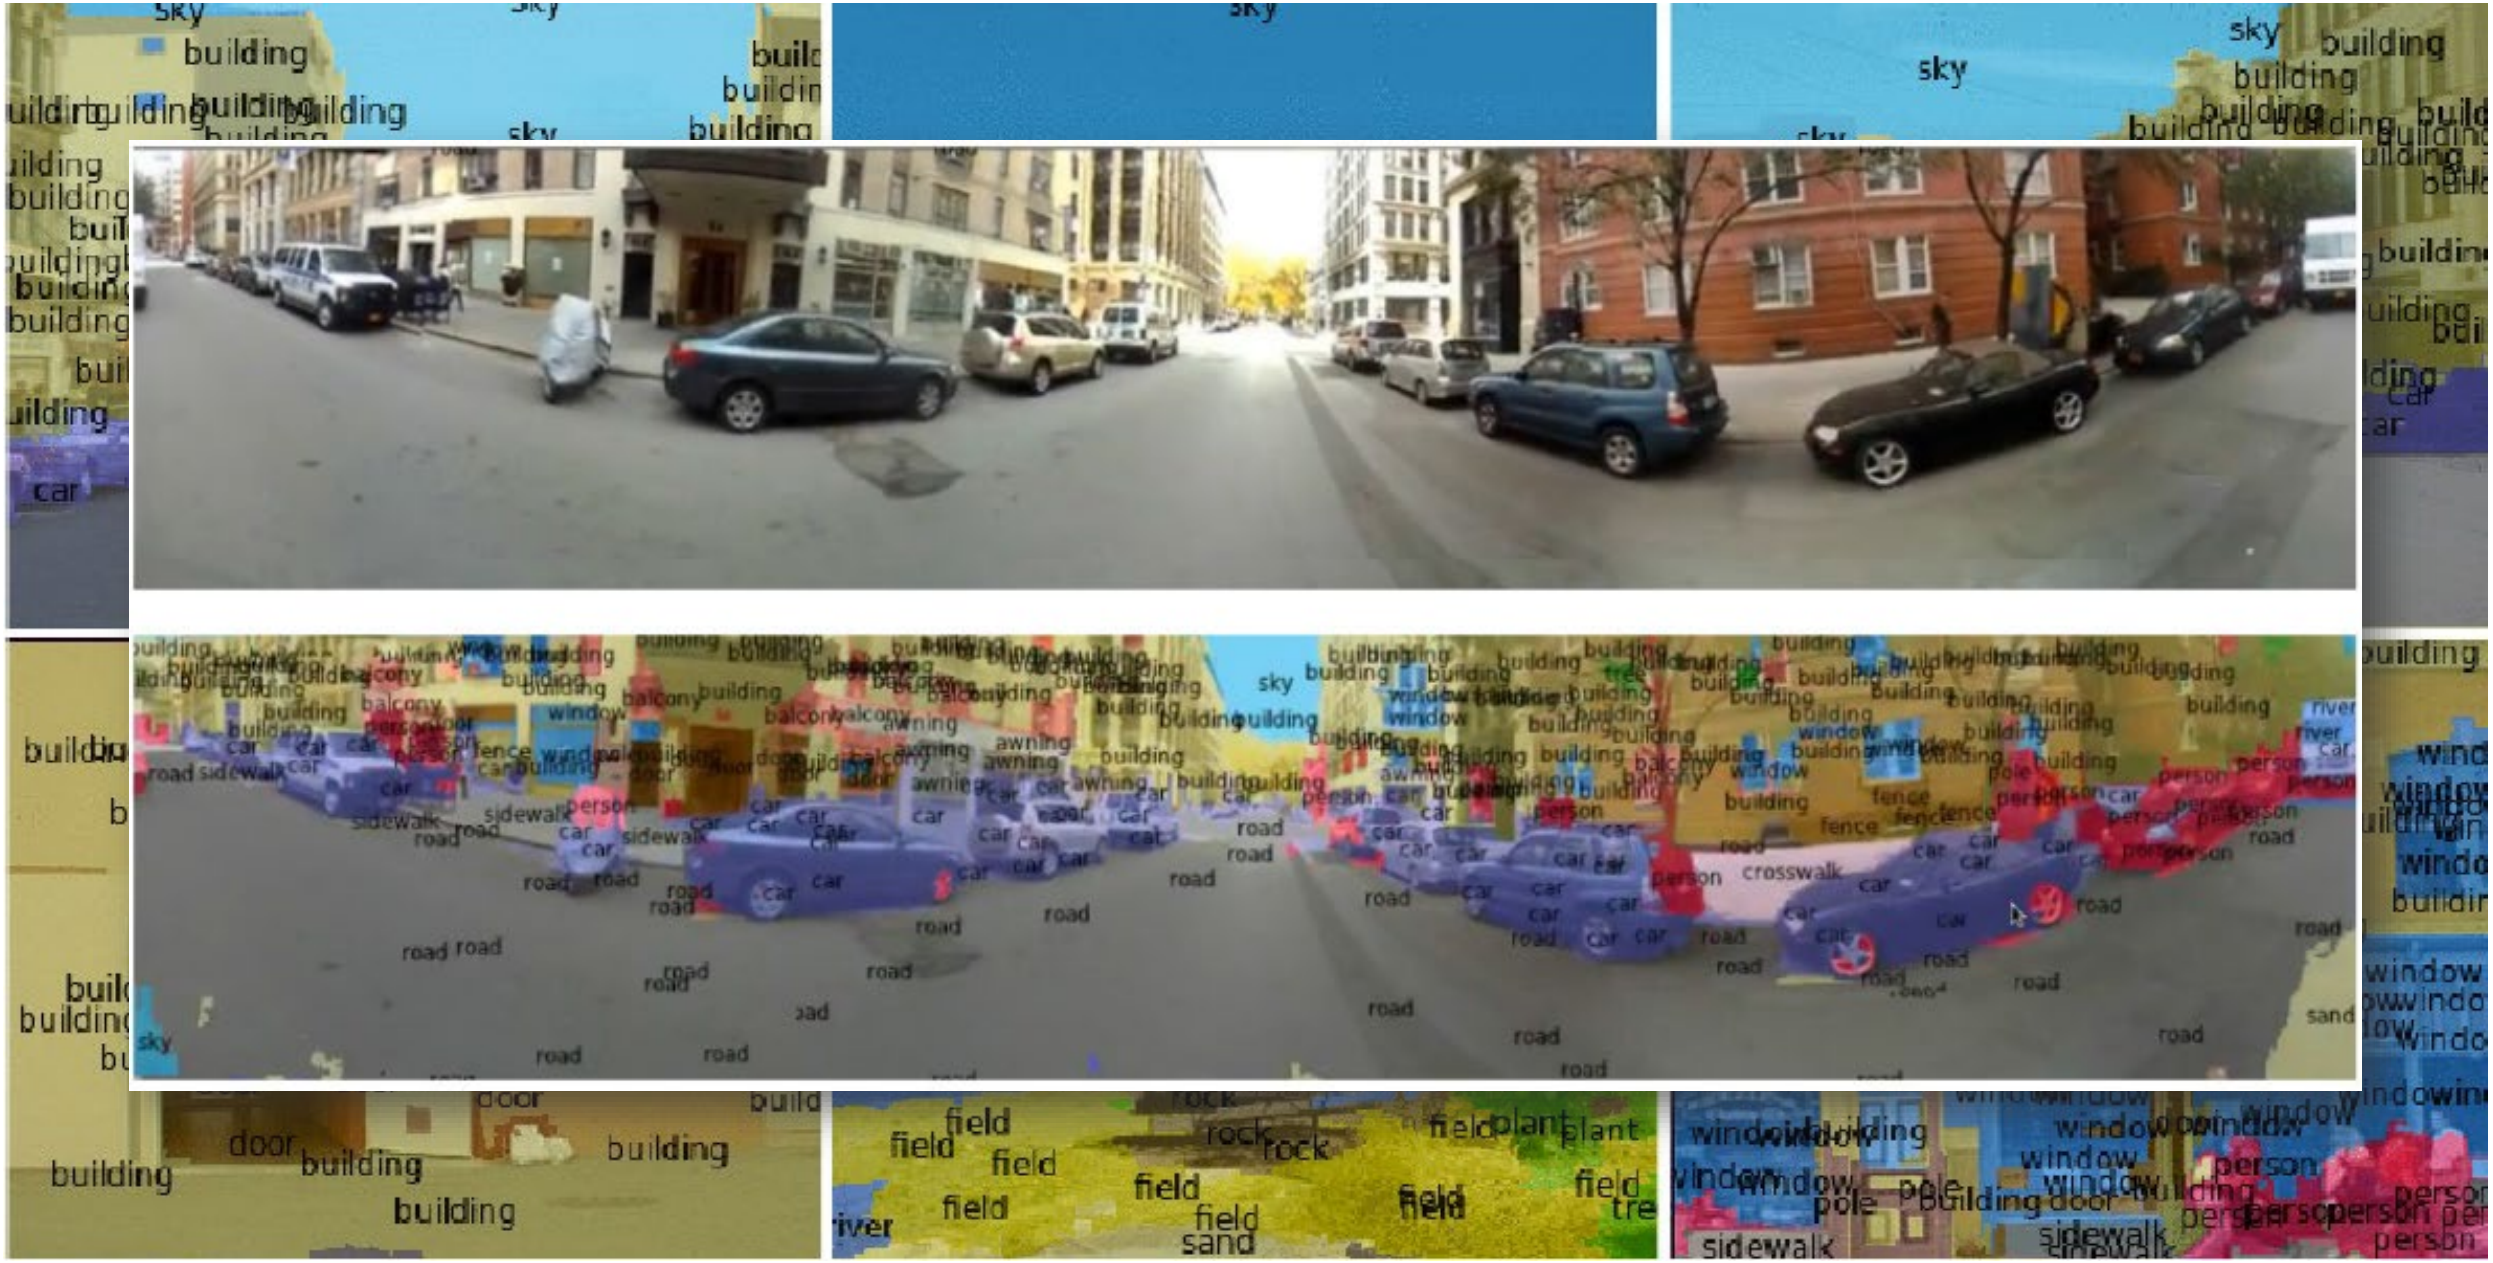

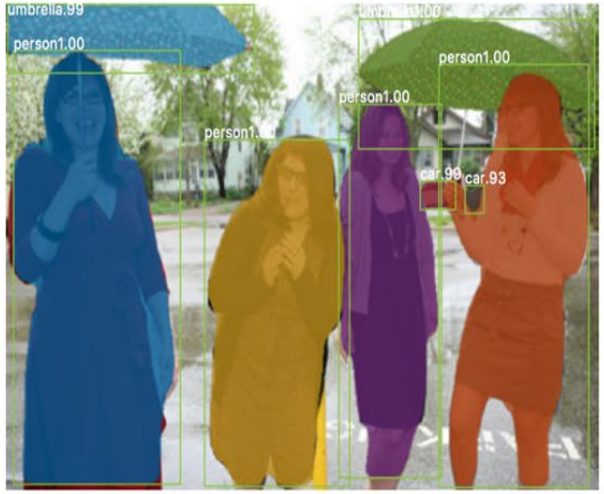
umbrella.99
person1.00
person1.00
person1.00
person1.00
car.99
car.93

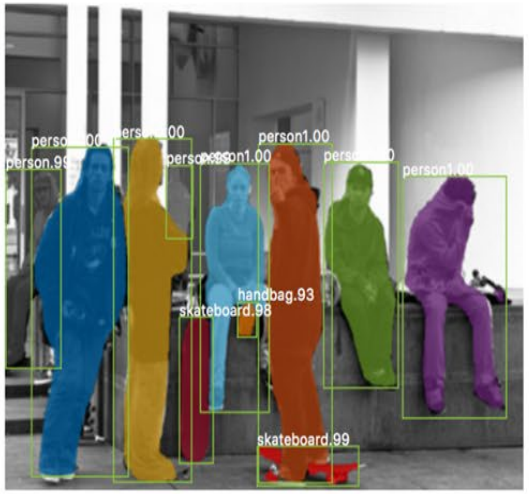
person1.00
person1.00
handbag.93
skateboard.98
skateboard.99

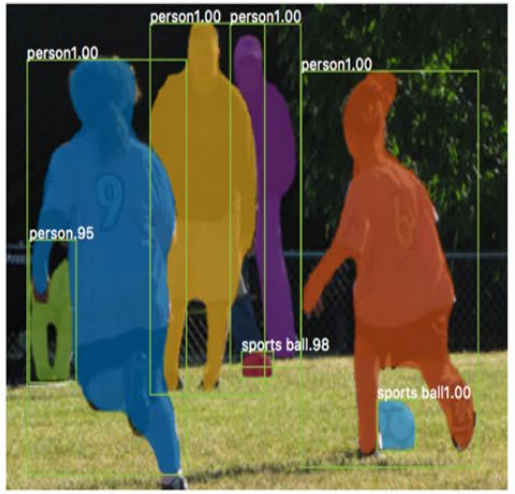
person1.00
person1.00
person1.00
person1.00
person.95
sports ball.98
sports ball1.00

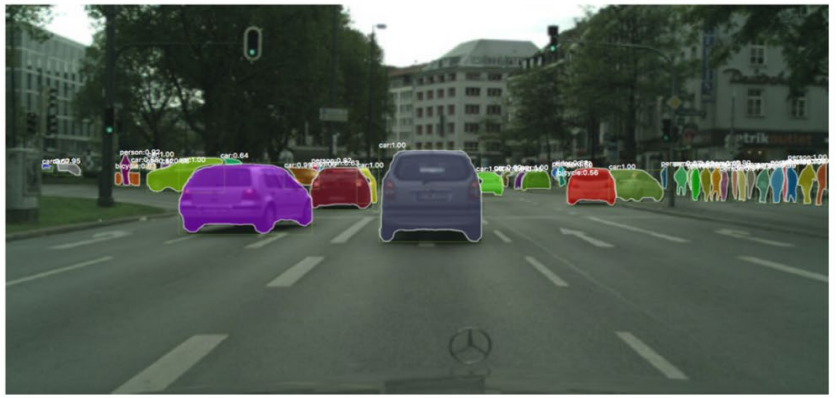


imax PHOTGRAPHY

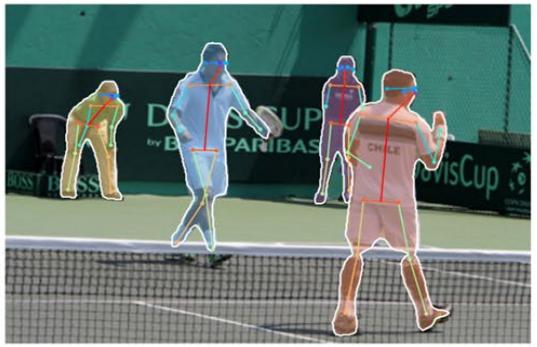

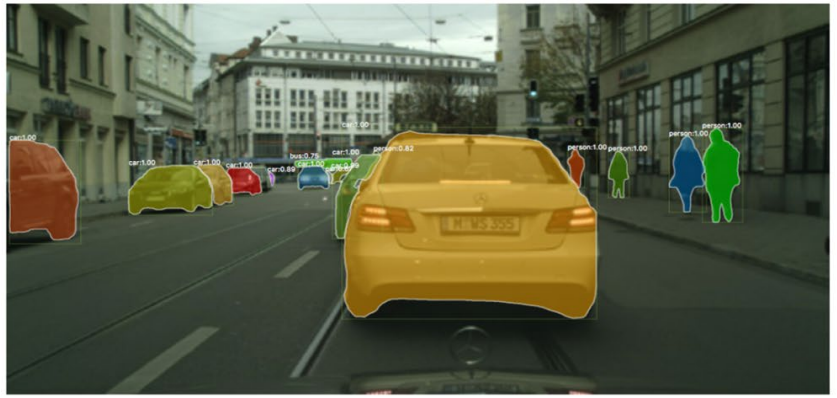

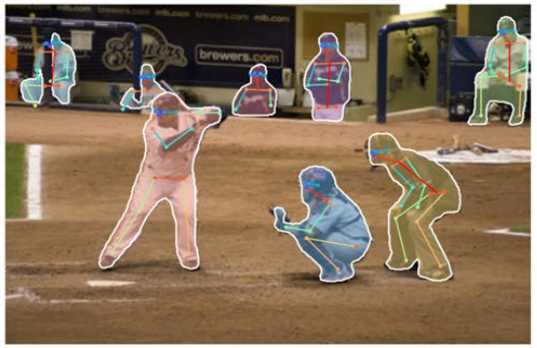

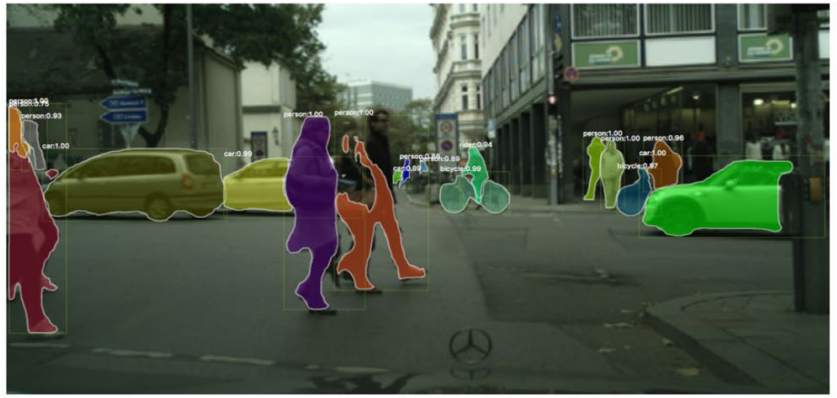






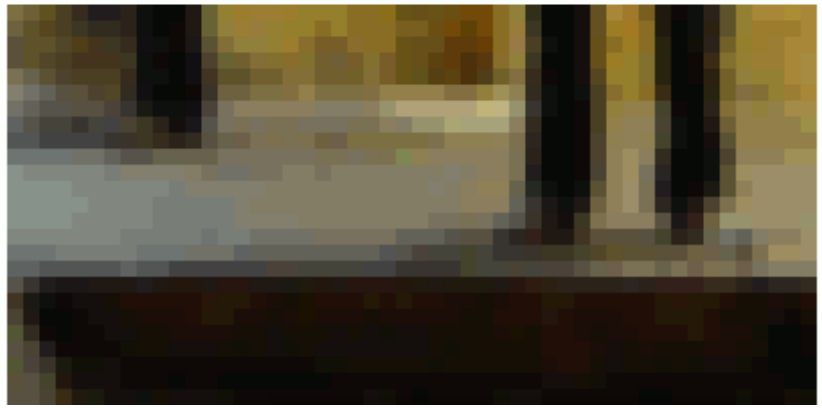


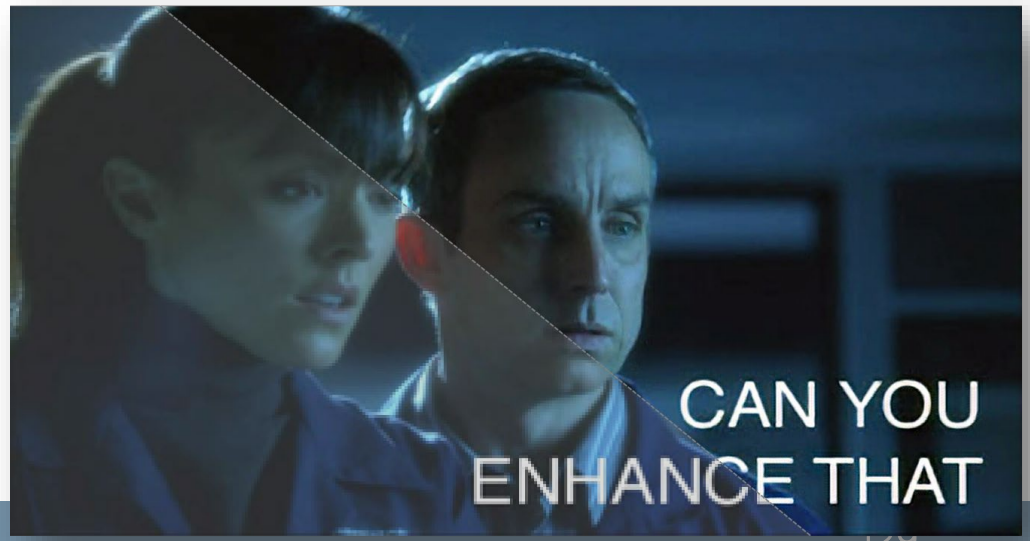

https://github.com/alexjc/neural-enhance

Stile Transfer

https://github.com/luanfujun/deep-photo-styletransfer

https://github.com/jcjohnson/neural-style

https://github.com/jcjohnson/fast-neural-style

https://ml4a.github.io/ml4a/style_transfer/

Image Captioning

"little girl is eating piece of cake."

"black cat is sitting on top of suitcase."

Andrej Karpathy, Li Fei-Fei "Deep Visual-Semantic Alignments for Generating Image Descriptions" CVPR 2015

Generative Adversarial Networks (these people do not exist)

Tero Karras, Samuli Laine, Timo Aila «A Style-Based Generator Architecture for Generative Adversarial Networks” **CVPR 2019**

Image Generation by Generative Adversarial Networks

Text description	This flower has petals that are white and has pink shading	This flower has a lot of small purple petals in a dome-like configuration	This flower has long thin yellow petals and a lot of yellow anthers in the center	This flower is pink, white, and yellow in color, and has petals that are striped	This flower is white and yellow in color, with petals that are wavy and smooth	This flower has upturned petals which are thin and orange with rounded edges	This flower has petals that are dark pink with white edges and pink stamen
256x256 StackGAN							

Text description	This bird is red and brown in color, with a stubby beak	The bird is short and stubby with yellow on its body	A bird with a medium orange bill white body gray wings and webbed feet	This small black bird has a short, slightly curved bill and long legs	A small bird with varying shades of brown with white under the eyes	A small yellow bird with a black crown and a short black pointed beak	This small bird has a white breast, light grey head, and black wings and tail
256x256 StackGAN							


Midjourney Bot BOT Today at 2:32 PM
Pope Francis wearing a long white puffer coat --v 5 - @a2jess


The Infinite Dude
@TheInfiniteDude
Replying to @okeefe_reborn
EXCLUSIVE:
🚨Trump Arrested in FBI Mar A Lago raid this evening.
2:57 PM · Mar 18, 2023 · 656.3K Views

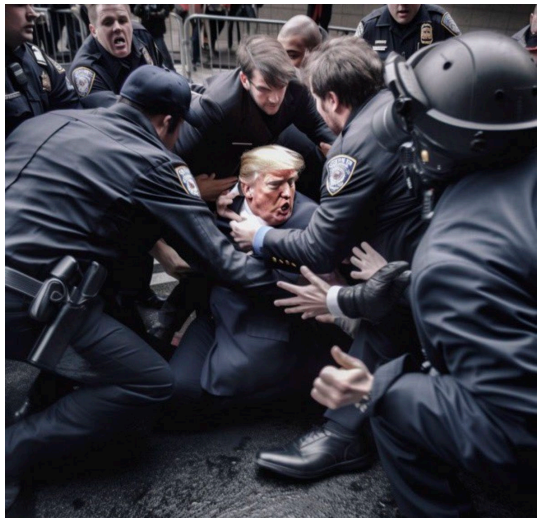

Even though sometimes it fails…

google.com/search?q=smoothing&oq=smoothing&aqs=chrome..69i57j0i512l9.3285j0j7&sourceid=chrome&ie=UTF-8

Google

smoothing

All Images Maps Videos News More Tools

About 289,000,000 results (0.66 seconds)

https://en.wikipedia.org › wiki › Smoothing

Smoothing - Wikipedia

In **smoothing**, the data points of a signal are modified so individual points higher than the adjacent points (presumably because of noise) are reduced, ...

Exponential smoothing · Additive smoothing · Smoothing spline · Edge-preserving

https://it.wikipedia.org › wiki › Lisci... · Translate this page

Lisciamento - Wikipedia

In statistica ed elaborazione digitale delle immagini, il lisciamento (traduzione letterale dell'inglese **smoothing**) o, meglio, perequazione di un insieme ...

People also ask

What do you mean by smoothing?

What is the purpose of smoothing?

How can smoothing be done?

What is the difference between smoothing and smoothening?

Feedback

Smoothing

In statistics and image processing, to smooth a data set is to create an approximating function that attempts to capture important patterns in the data, while leaving out noise or other fine-scale structures/rapid phenomena. Wikipedia

Feedback

42%

Even though sometimes it fails…