16.01.2024

BITS F464: Machine Learning

MACHINE LEARNING OVERVIEW

Chittaranjan Hota, Sr. Professor
Dept. of Computer Sc. and Information Systems
hota@hyderabad.bits-pilani.ac.in

What is Machine Learning?

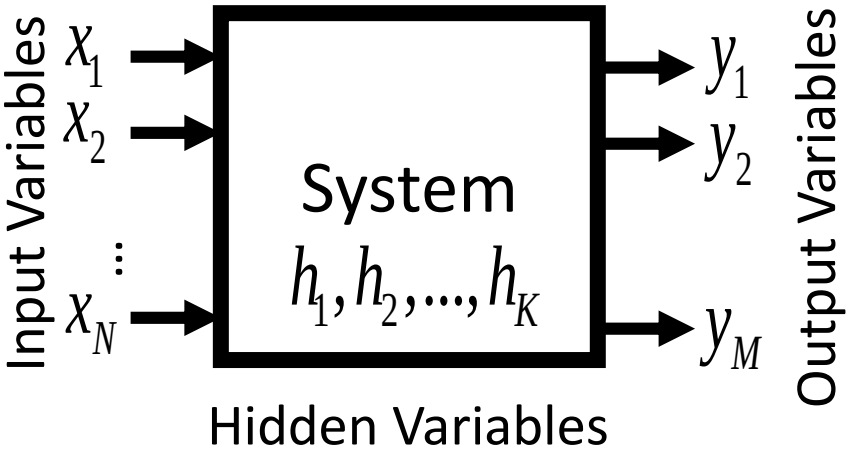

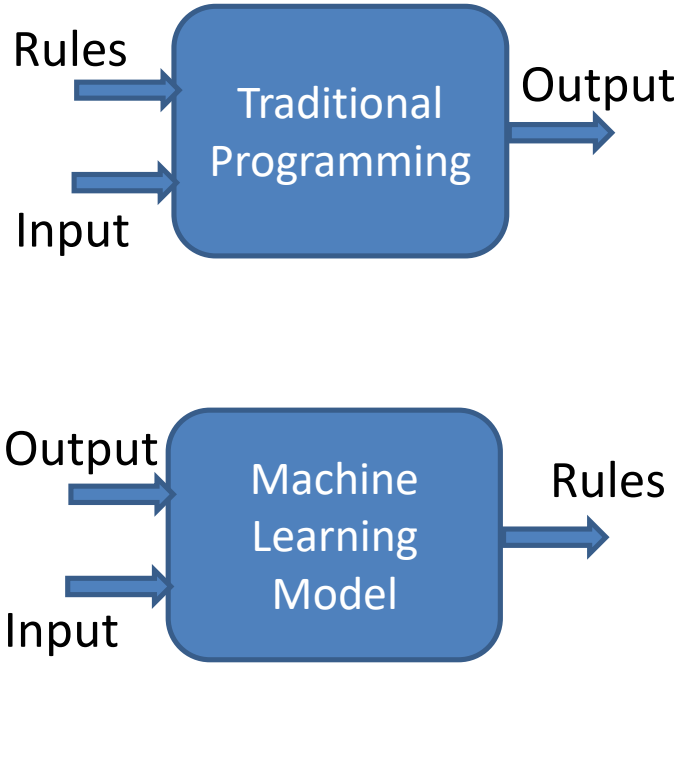

Simple Learning Process

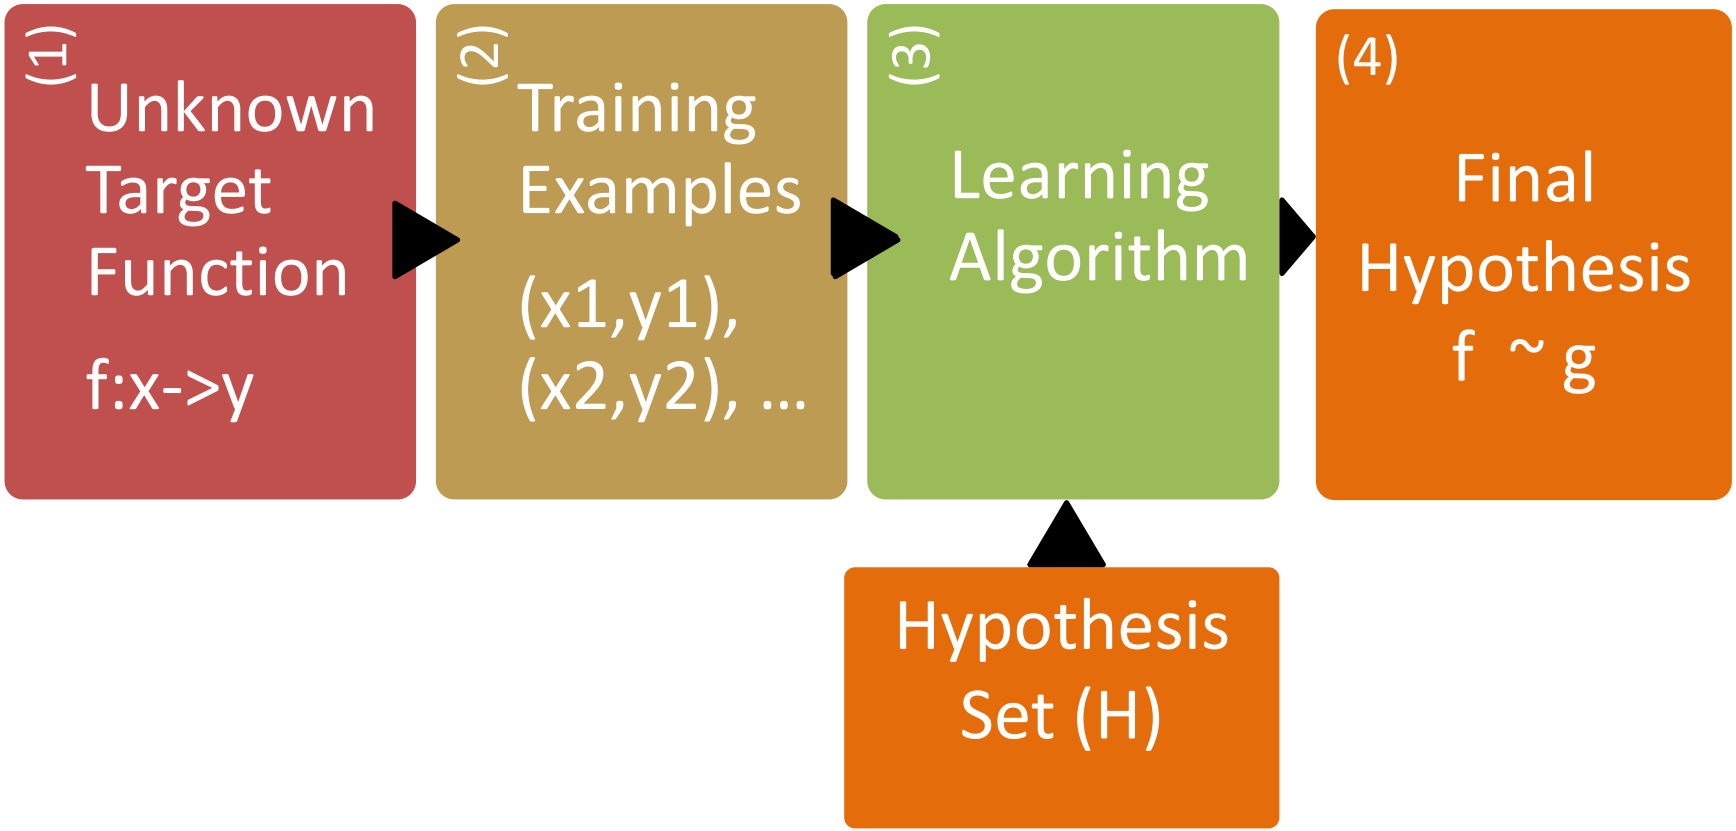

An Example: Step 1: Collecting the data

(Training Set)

Step 2: Designing the features

- Not a trivial task. Designing quality features could be very application dependent.
- For ex: Would you like to take "number of legs" as one feature to distinguish cats from dogs?

- A good one for our example:
  - size of nose, relative to the size of the head (ranging from small to big);
  - shape of ears (ranging from round to pointy).

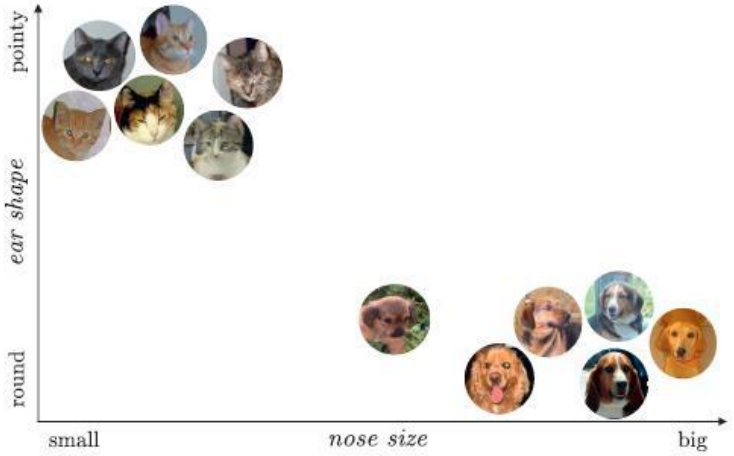

Called as a feature vector.

Step 2: Continued (Handwritten digits)

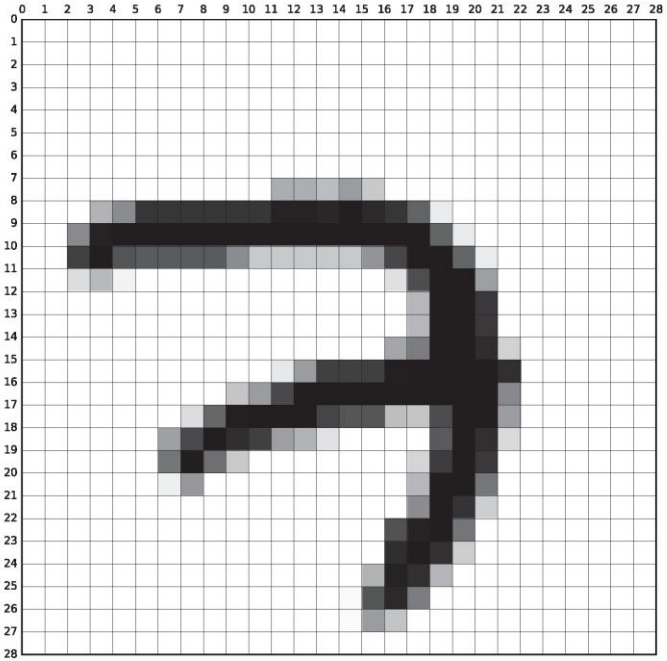

Img. Source: Alejandro B

(a) MNIST sample belonging to the digit '7'.

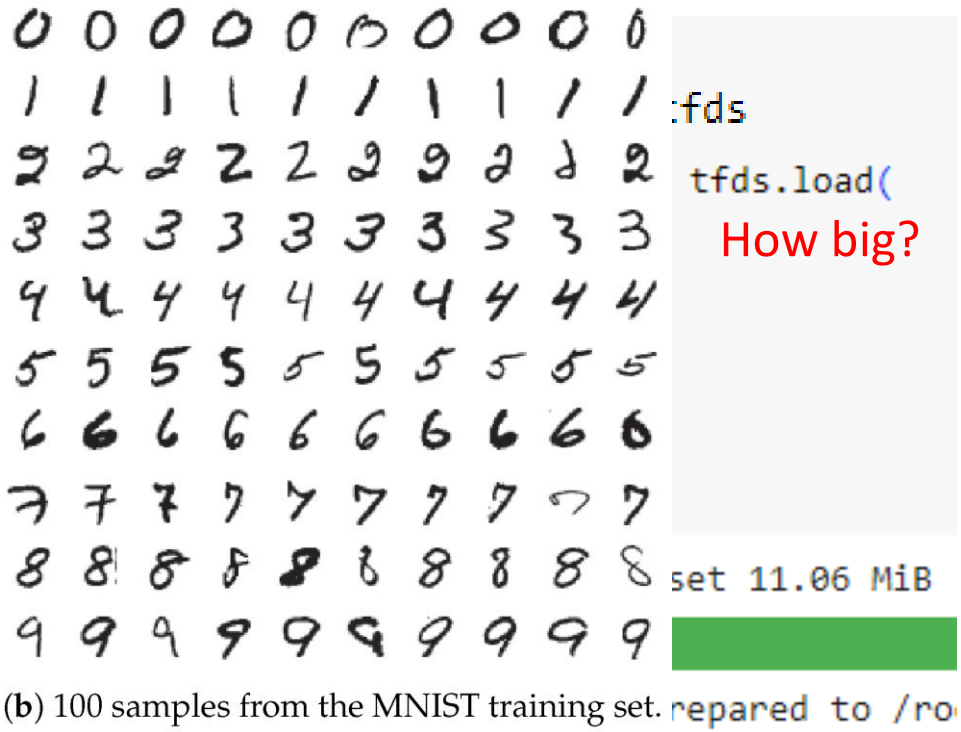

(b) 100 samples from the MNIST training set.

Features: Pixel values, Image size, Aspect Ratio, Normalized Pixel values, edges, … (manual) → Automatic: CNN: texture, shape, corners,…

Step 2: Continued (Type of Fish)

Salmon

Sea bass

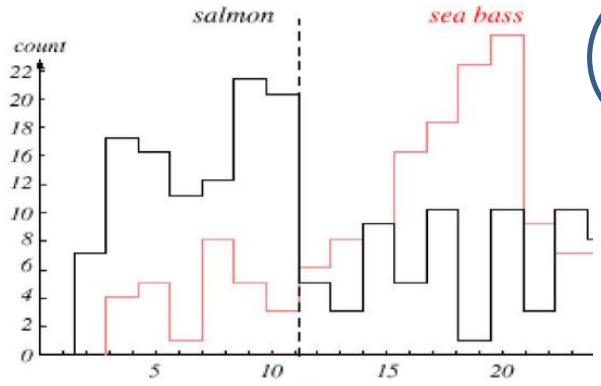

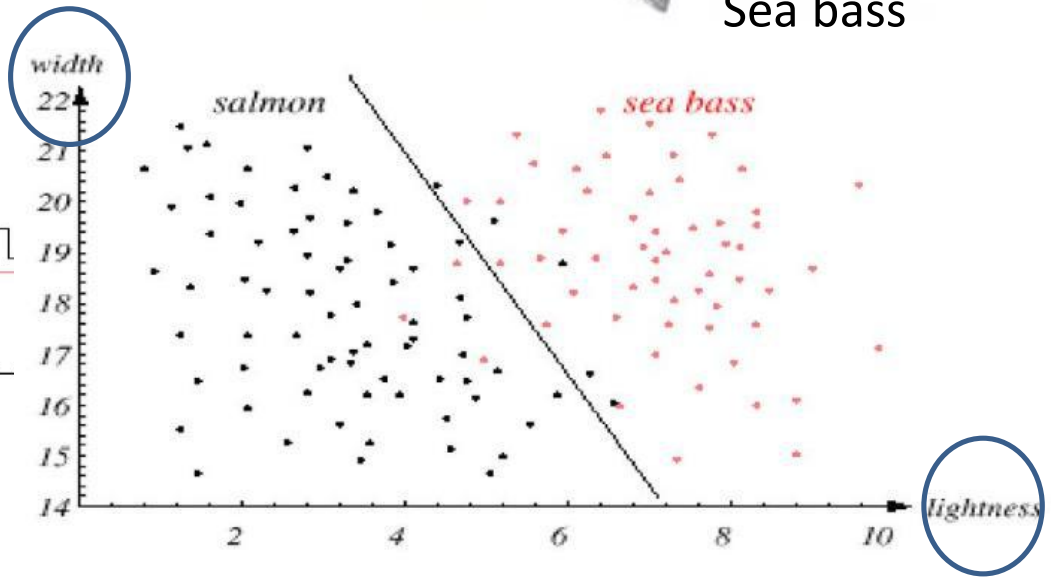

Fish → $x^T = [x_1 , x_2]$

lightness width

Step 3: Training the Model (Cats Vs Dogs)

- Now it is a simple geometric problem. Let the computer find out a Line (linear model) that separates cats from dogs.
- Equation of a line?
- How does the computer find out m and c?

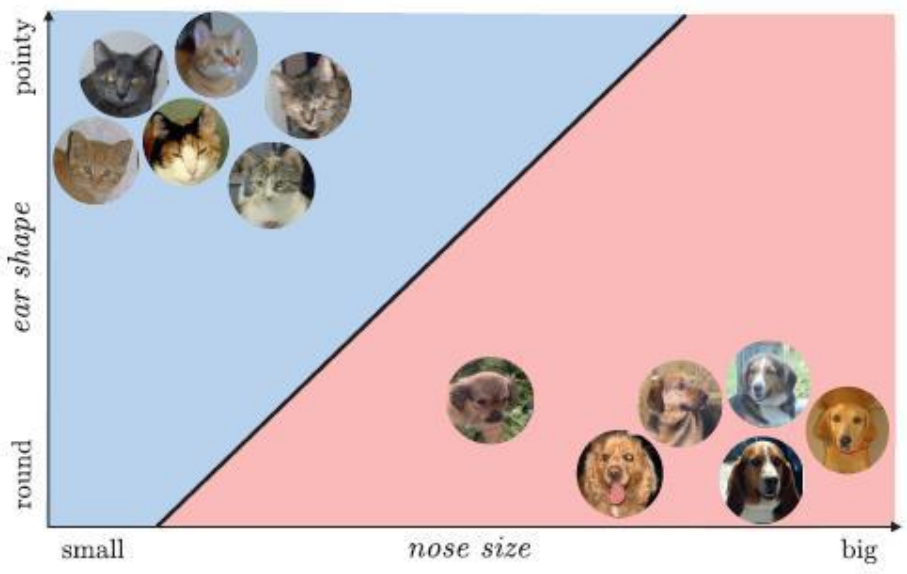

We could instead find a curve or nonlinear *model* that separates the data. In general, linear models are by far the most common choice in practice when features are designed properly.

Step 3 Continued (Fish Example)

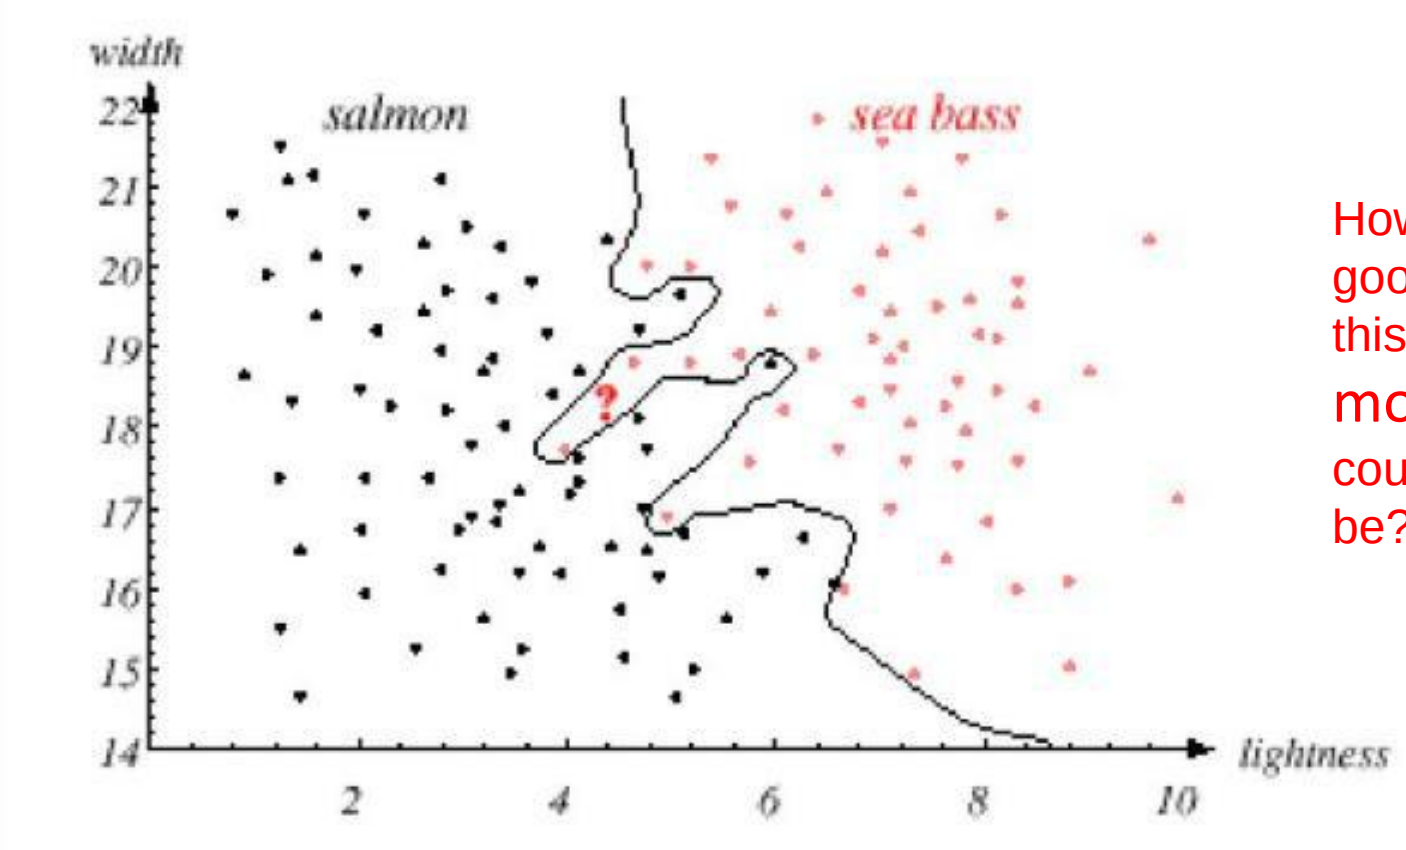

How good this model could be?

Img. Source: Sergios Theodoridis

Step 4: Testing the Model

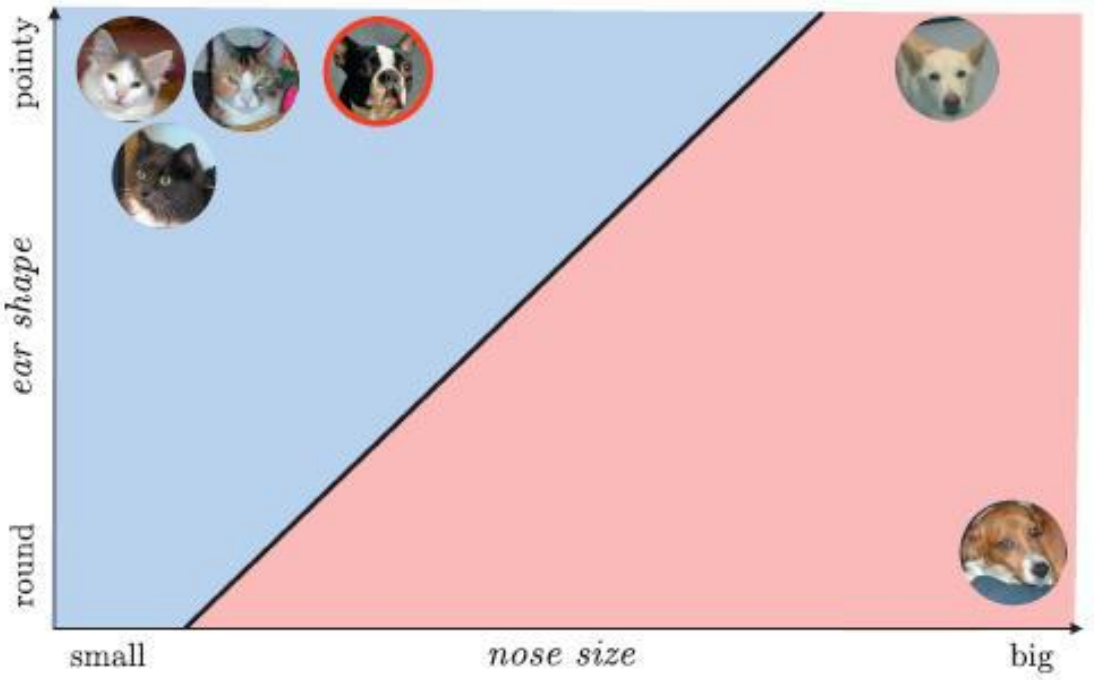

- What is the problem here?
- Can you list down few more discriminating features?

Types of Learning: Supervised Learning

- Correct Output known for each training example.
  - Classification: 1-of-N output (whether it is a Cat or a Dog?)
  - Regression: Real valued output (how many students will enroll into ML course next semester?)

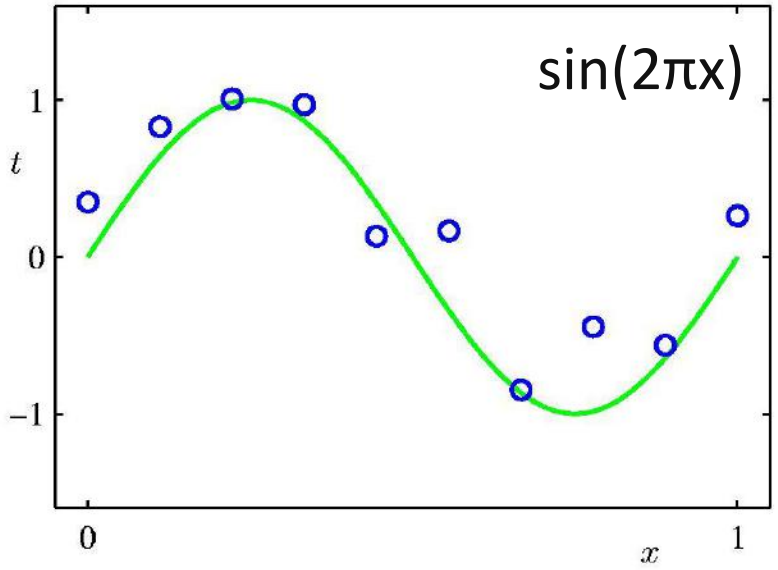

$$y(x, \mathbf{w}) = w_0 + w_1 x + w_2 x^2 + \ldots + w_M x^M = \sum_{j=0}^{M} w_j x^j$$

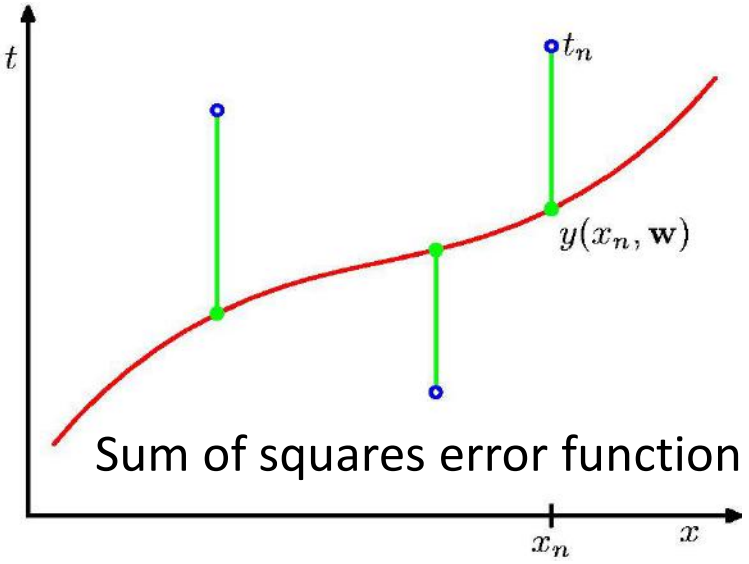

$$E(\mathbf{w}) = \frac{1}{2} \sum_{n=1}^{N} \{y(x_n, \mathbf{w}) - t_n\}^2$$

Model selection: What should be M?

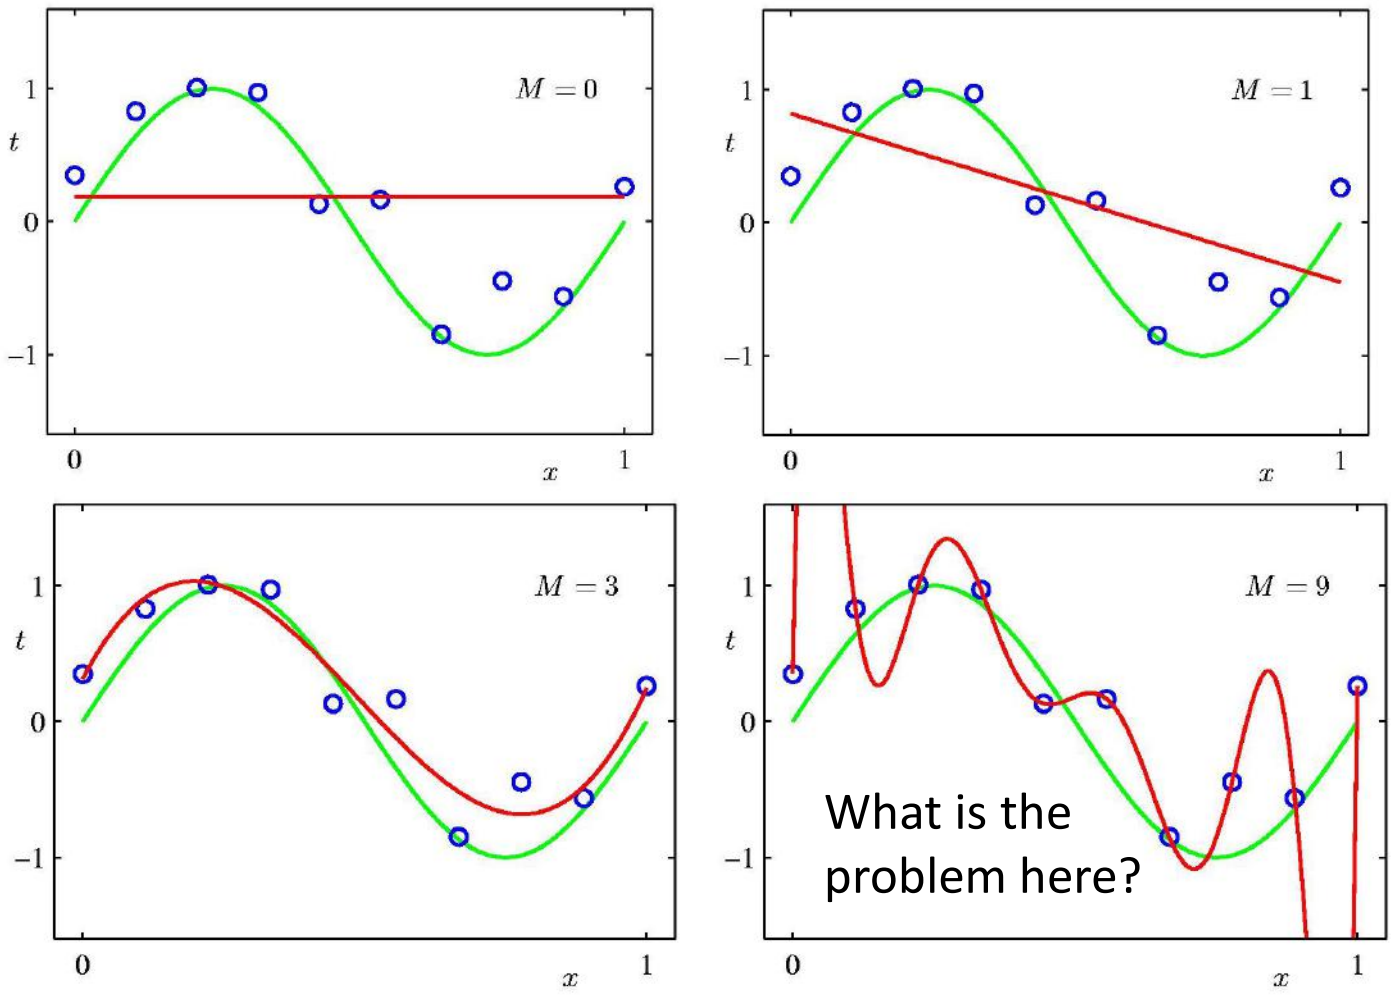

Over-fitting

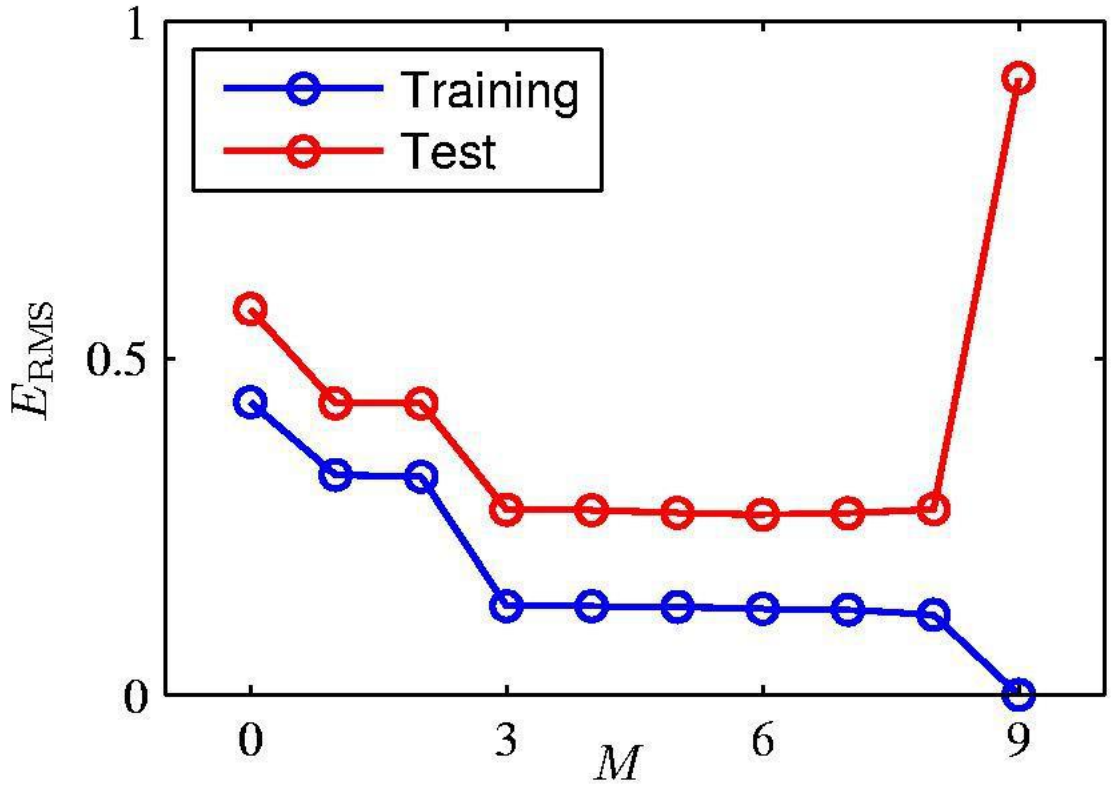

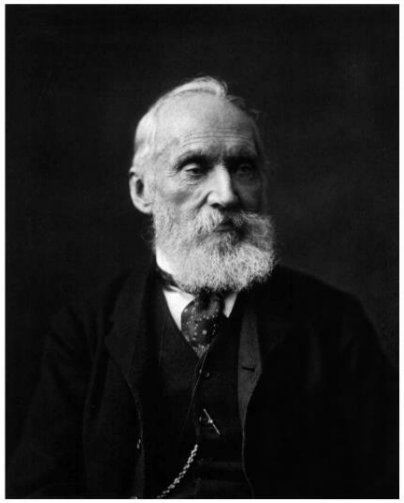

Solutions: later

Root-Mean-Square (RMS) Error: $E_{\mathrm{RMS}} = \sqrt{2E(\mathbf{w}^\star)/N}$

Unsupervised Learning

- Learns from data without human supervision.
- Using unlabeled data, these

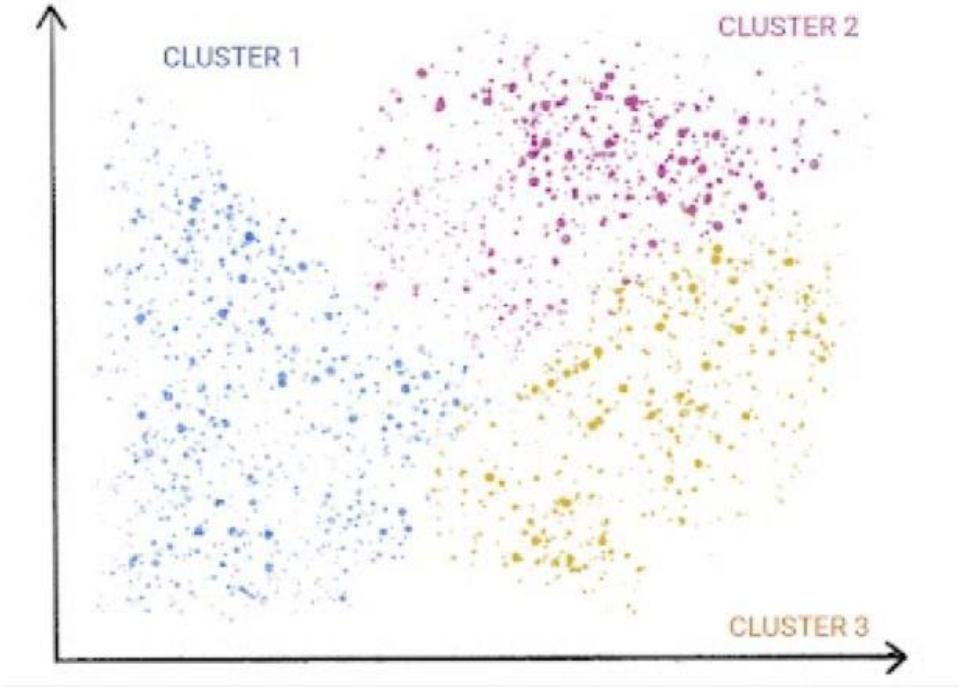

Img. Source: Quora

Reinforcement Learning

Learn action to maximize payoff

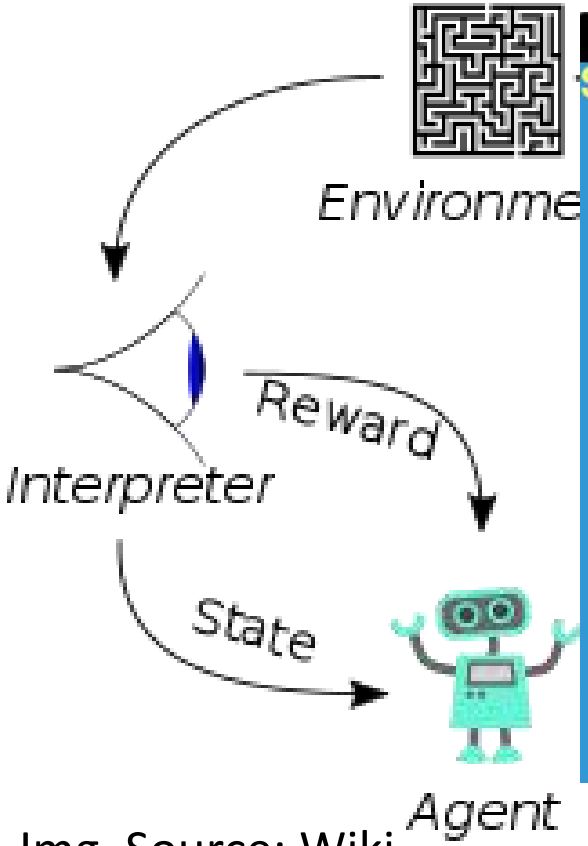

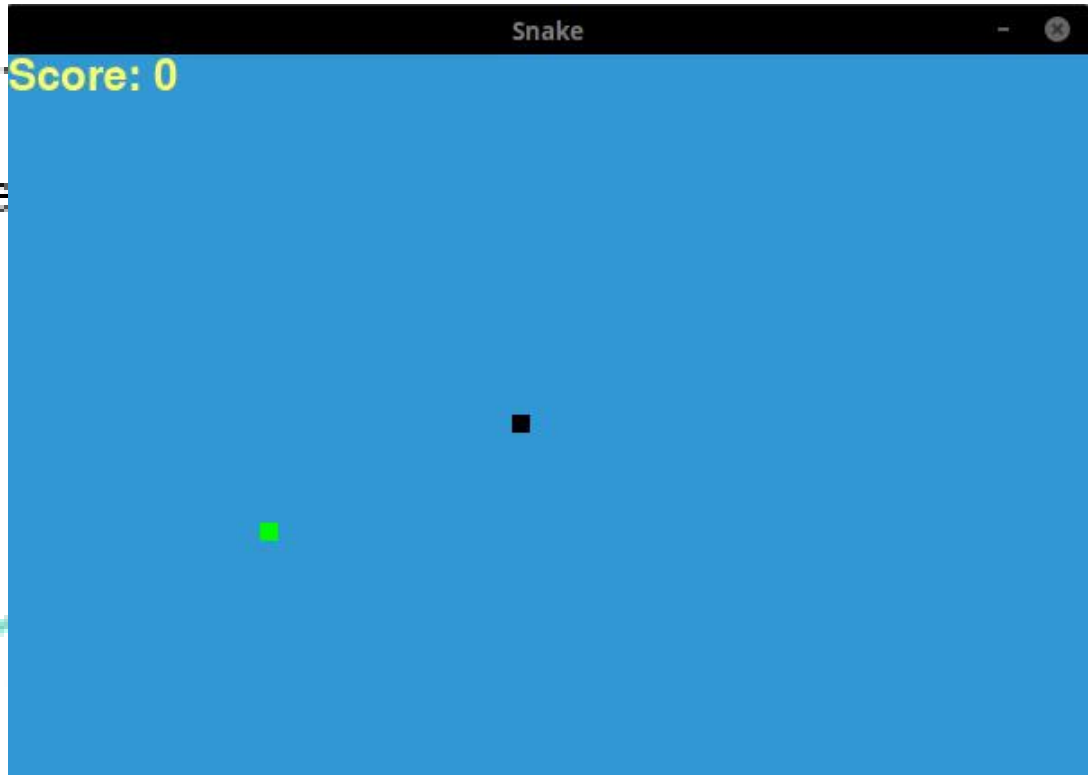

Img. Source: Wiki

Source: https://towardsdatascience.com/

Issues in Machine Learning

- What algorithms are available for learning a concept? How well do they perform?
- How much training data is sufficient to learn a concept with high confidence?
- How are the features generated?
- Are some training examples more useful than others?
- What are the best tasks for a system to learn?

Quiz for you

Q.1 If fig.a is the input to your ML model and fig.b is the output with names of people in the photograph, then what type of problem is this?

- Classification ✓
- Regression

(a)

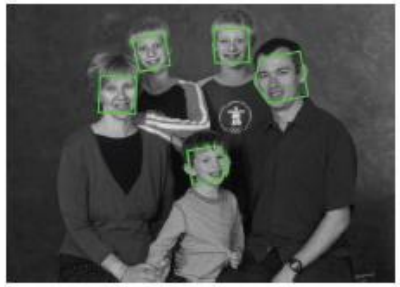

(b)

(Img. Source: Kevin Murphy)

Q.2 In the Cats and Dogs example that we discussed, what is being learnt by the model?

- Slope of the line and all the points (their coordinates) on the line.
- Slope of the line and the Intercept ✓

Q.3 Finding out who are the students in this class who play Cricket with an unlabelled dataset can be solved by using:

- Supervised Learning
- Un-supervised Learning ✓

Thank you!