



University of Cyprus – MSc Artificial Intelligence

MAI644 – COMPUTER VISION

Lecture 11: Visual Recognition – Image Classification

Melinos Averkiou


CYENS Centre of Excellence

University of Cyprus - Department of Computer Science

m.averkiou@cyens.org.cy



CYENS
CENTRE OF EXCELLENCE



Last time

- Visual Recognition Tasks
- Introduction to segmentation and clustering
- Agglomerative clustering
- K-means clustering
- Mean-shift clustering
- Efficient Graph-based image segmentation

Today's Agenda

- A simple Image Classification pipeline
 - Classification overview
- K-nearest neighbor algorithm
 - kNN: algorithm
 - kNN: analysis

[material based on Niebles-Krishna]



Today's Agenda

- A simple Image Classification pipeline
 - Classification overview
- K-nearest neighbor algorithm
 - kNN: algorithm
 - kNN: analysis



Visual recognition: a classification framework

- Apply a prediction function to a feature representation of the image to get the desired output:

$$f(\text{apple image}) = \text{"apple"}$$

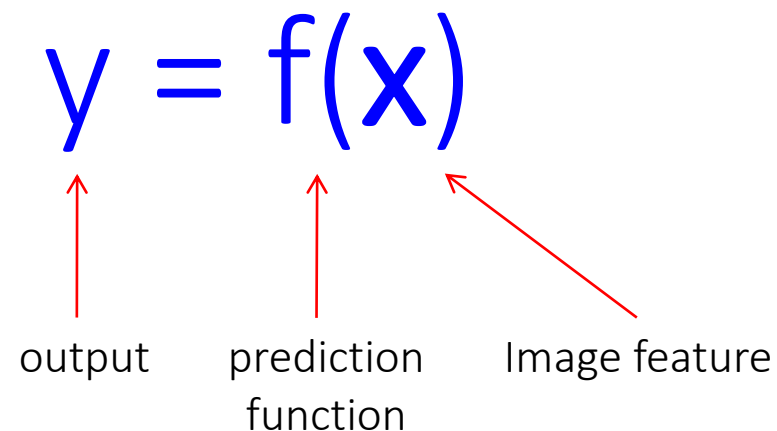
$$f(\text{tomato image}) = \text{"tomato"}$$

$$f(\text{cow image}) = \text{"cow"}$$

Dataset: ETH-80, by B. Leibe

Slide credit: L. Lazebnik

The machine learning framework




- **Training:** given a *training set* of labeled examples $\{(x_1, y_1), \dots, (x_N, y_N)\}$, estimate the prediction function f by minimizing the prediction error on the training set
- **Testing:** apply f to a never-before-seen *test example* x and output the predicted value $y = f(x)$

Slide credit: L. Lazebnik

Classification

- Assign input feature vector to one of two or more classes
- Any decision rule divides input space into *decision regions* separated by *decision boundaries*



Slide credit: L. Lazebnik

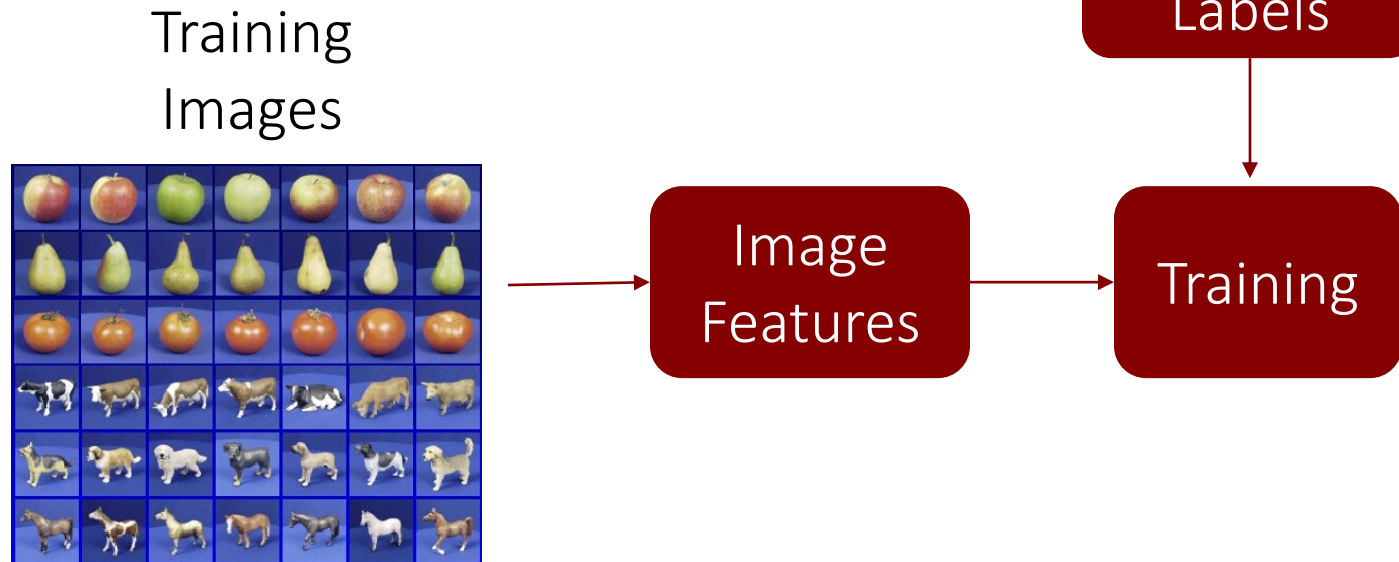
A simple pipeline - Training

Training Images




Image Features

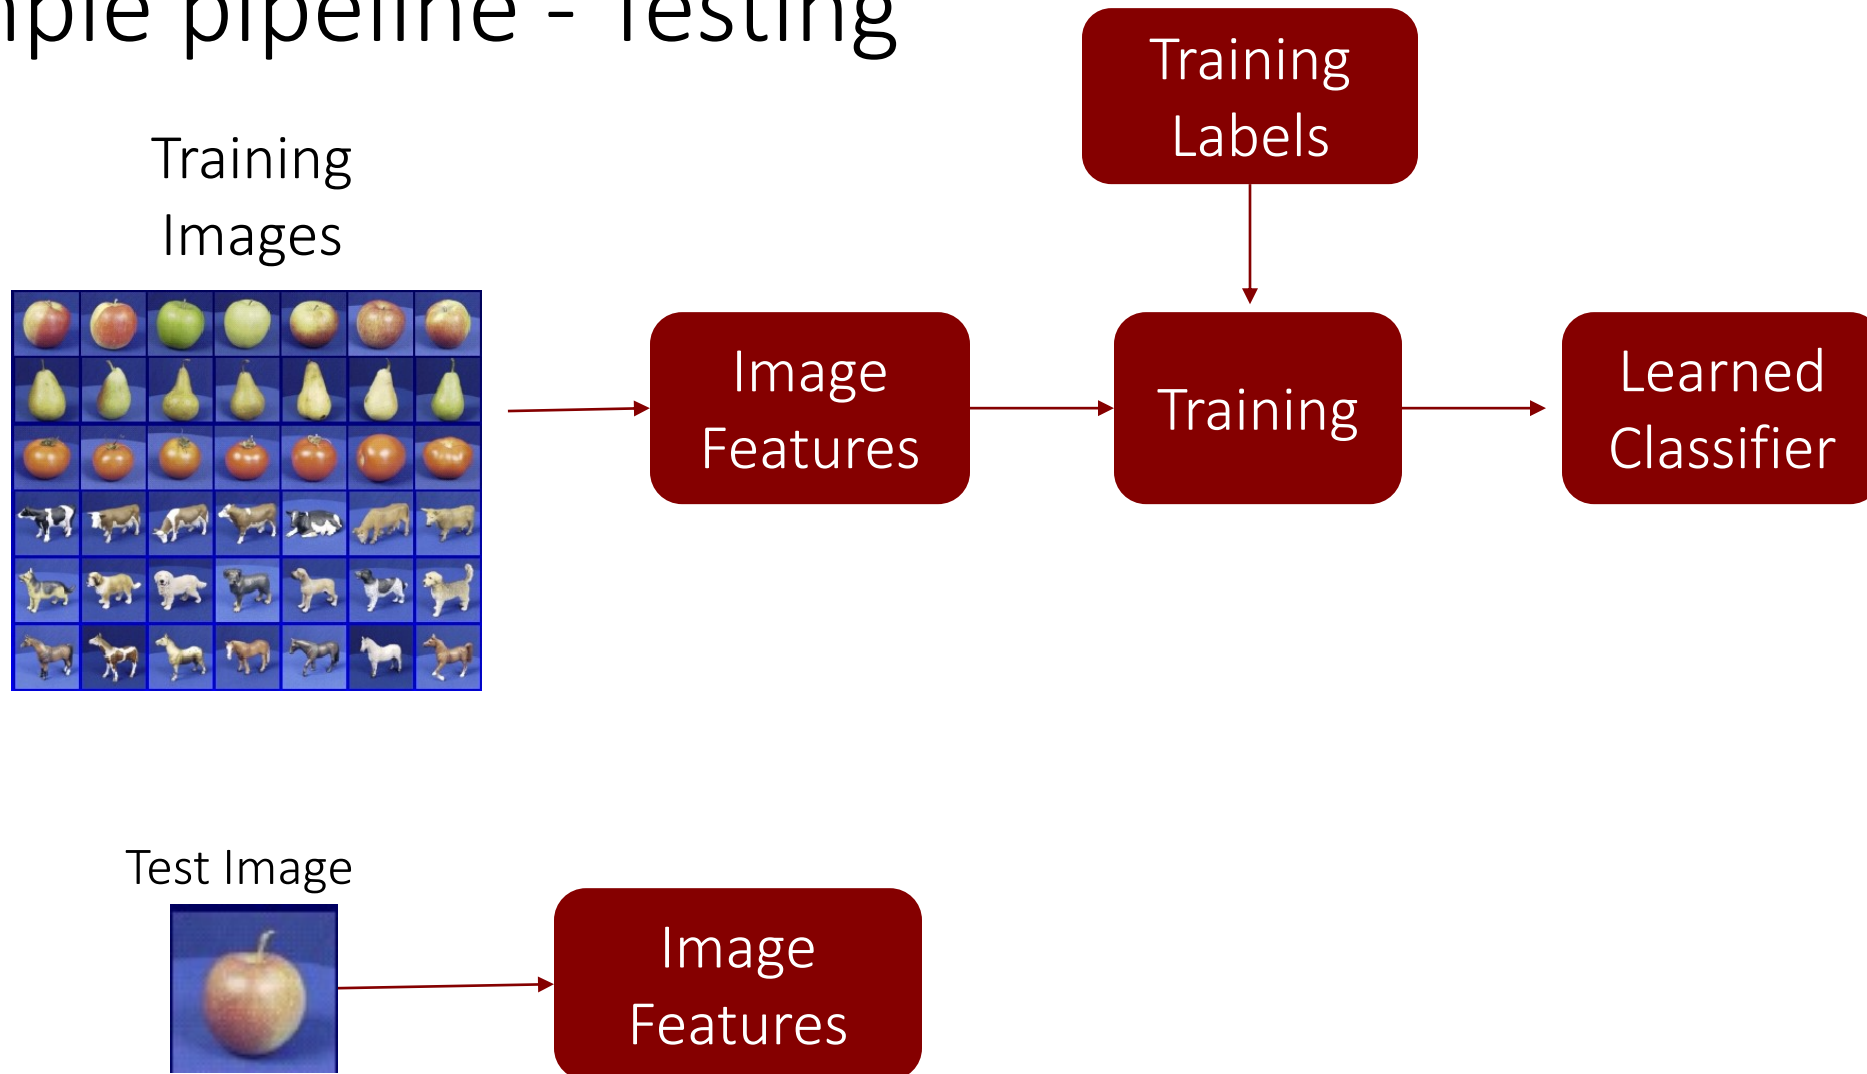
A simple pipeline - Training



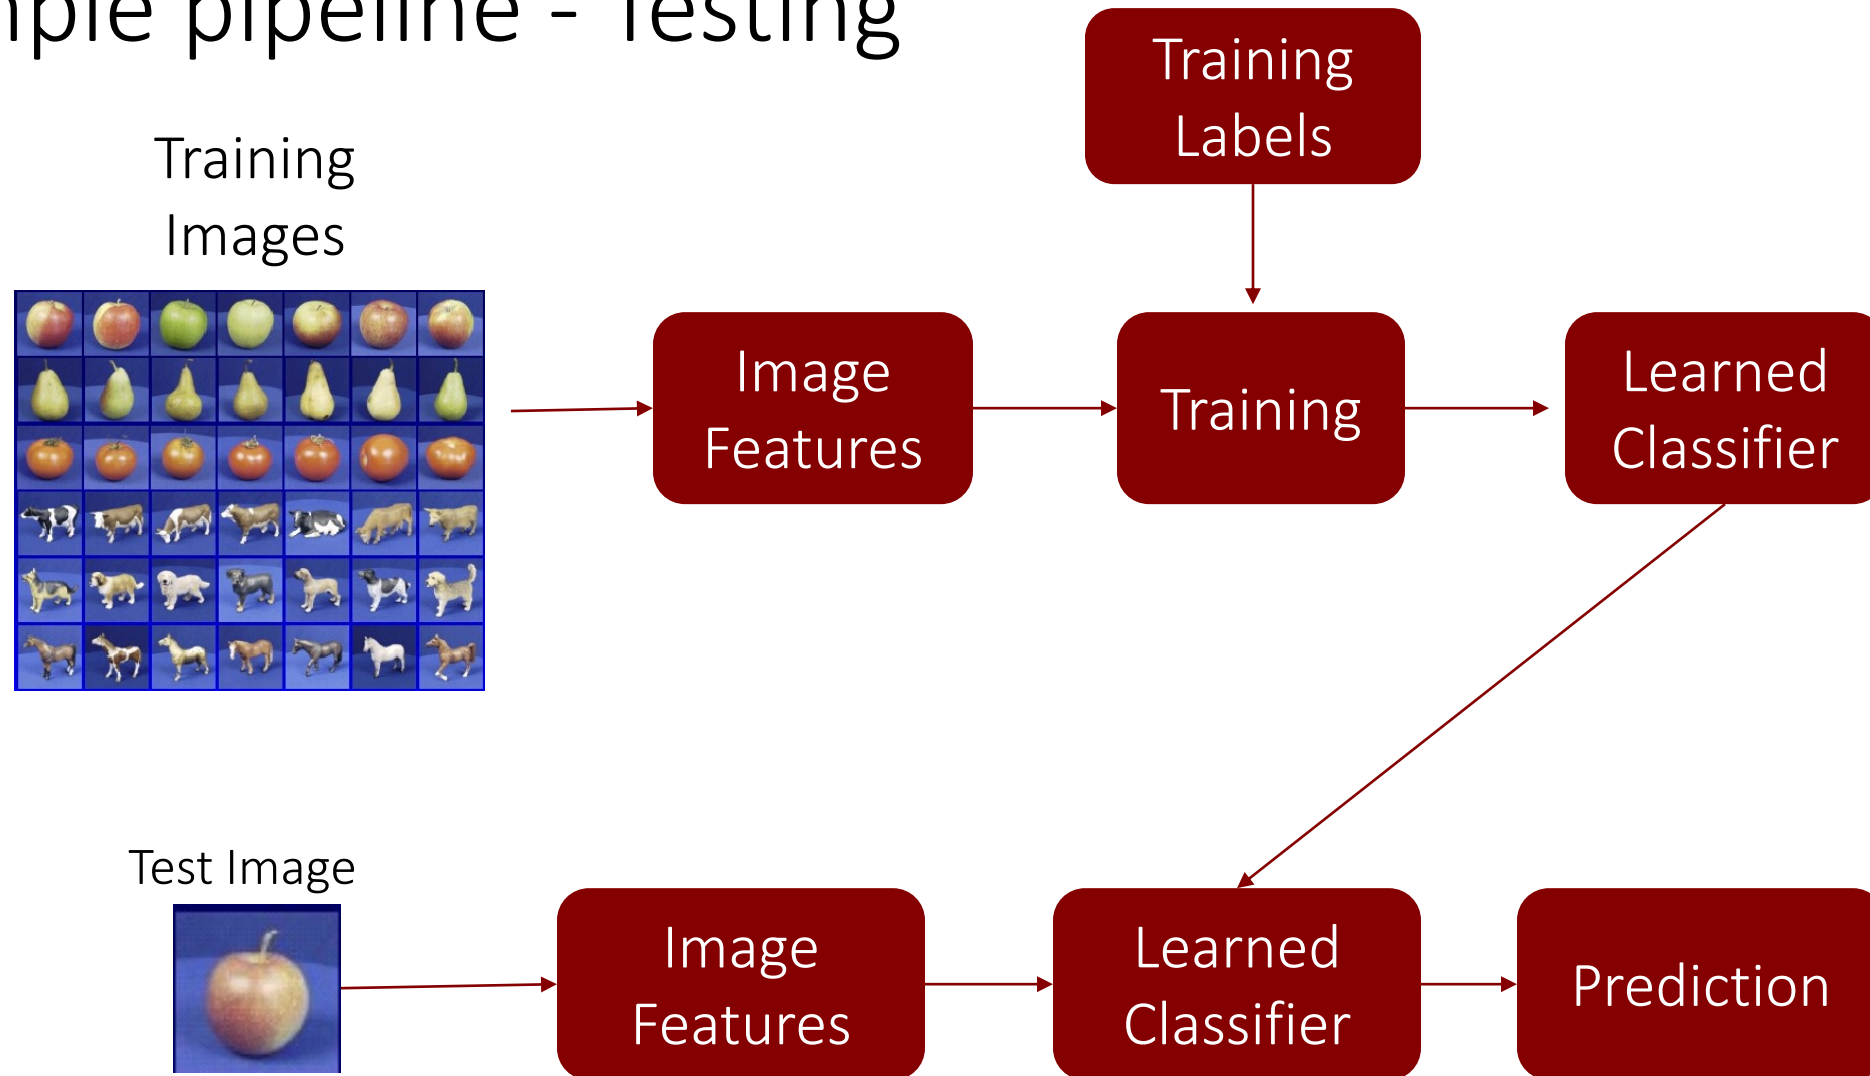
A simple pipeline - Training



A simple pipeline - Testing



A simple pipeline - Testing



A simple pipeline - Training

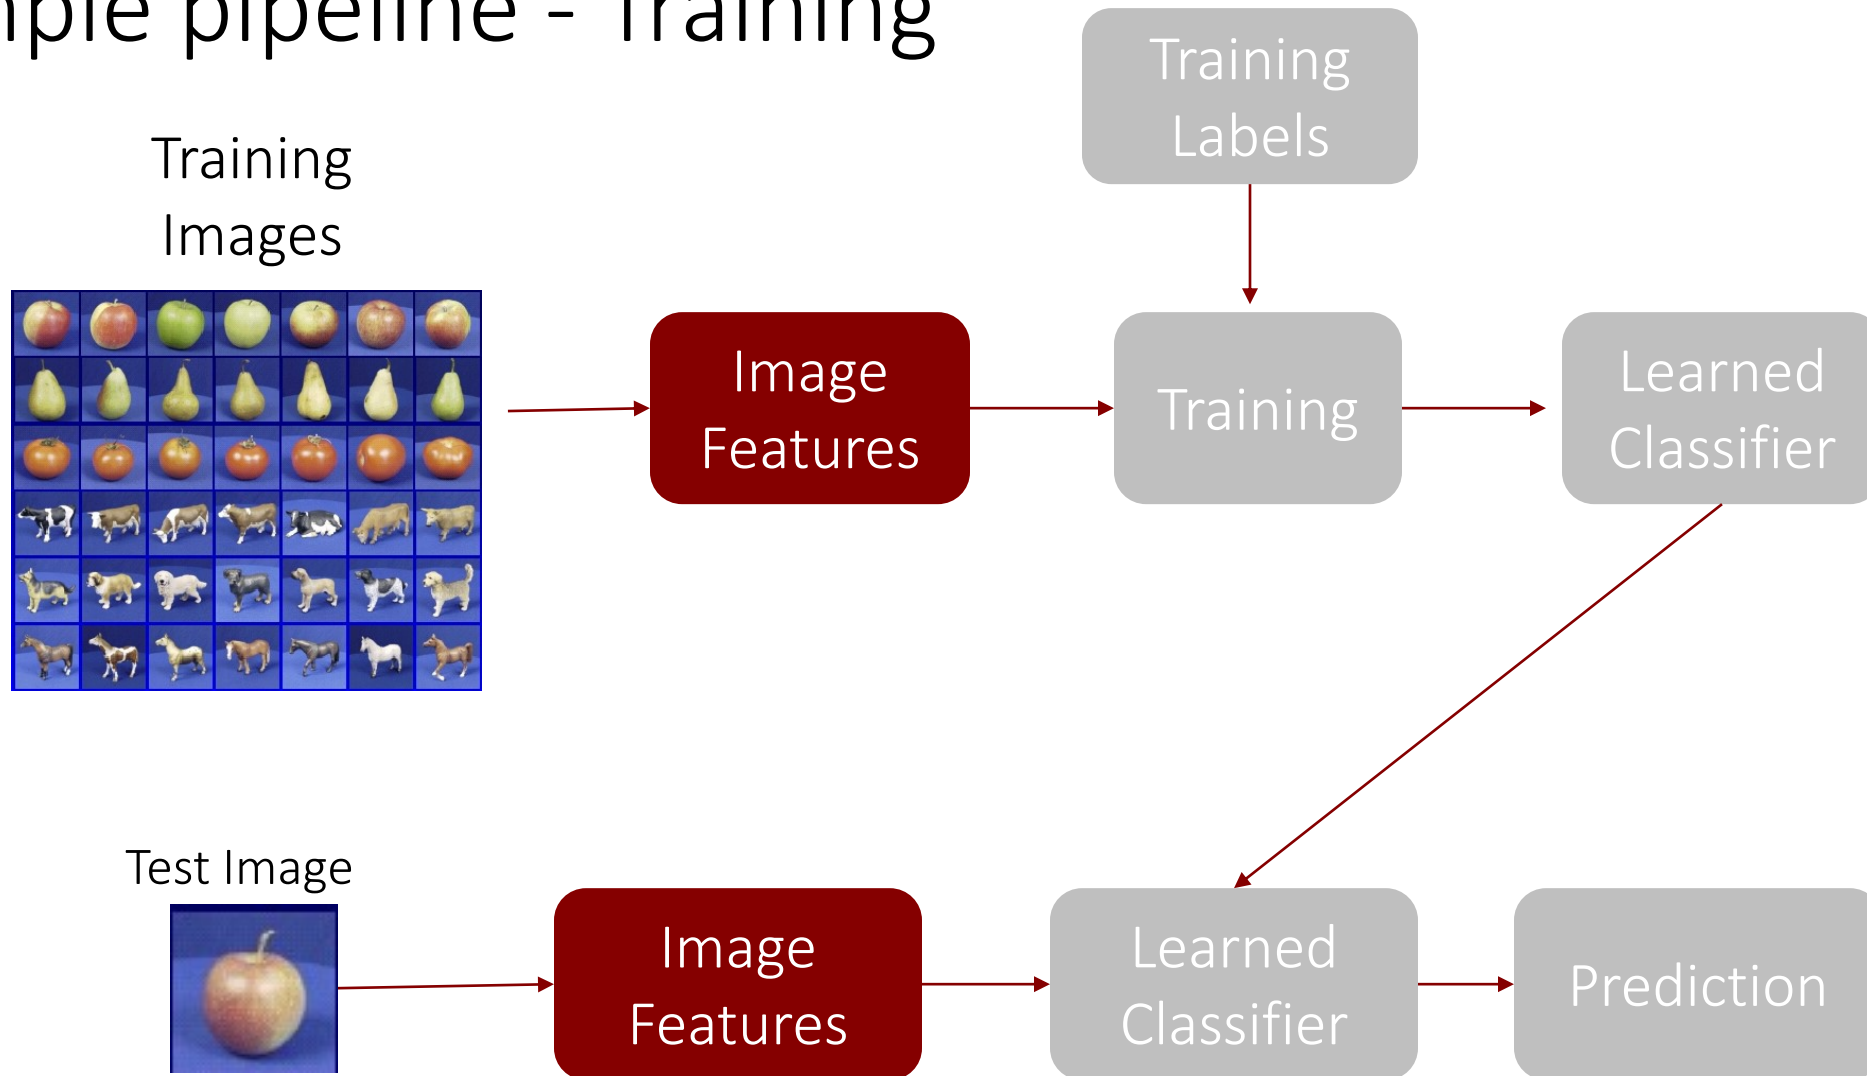



Image Features

Input image



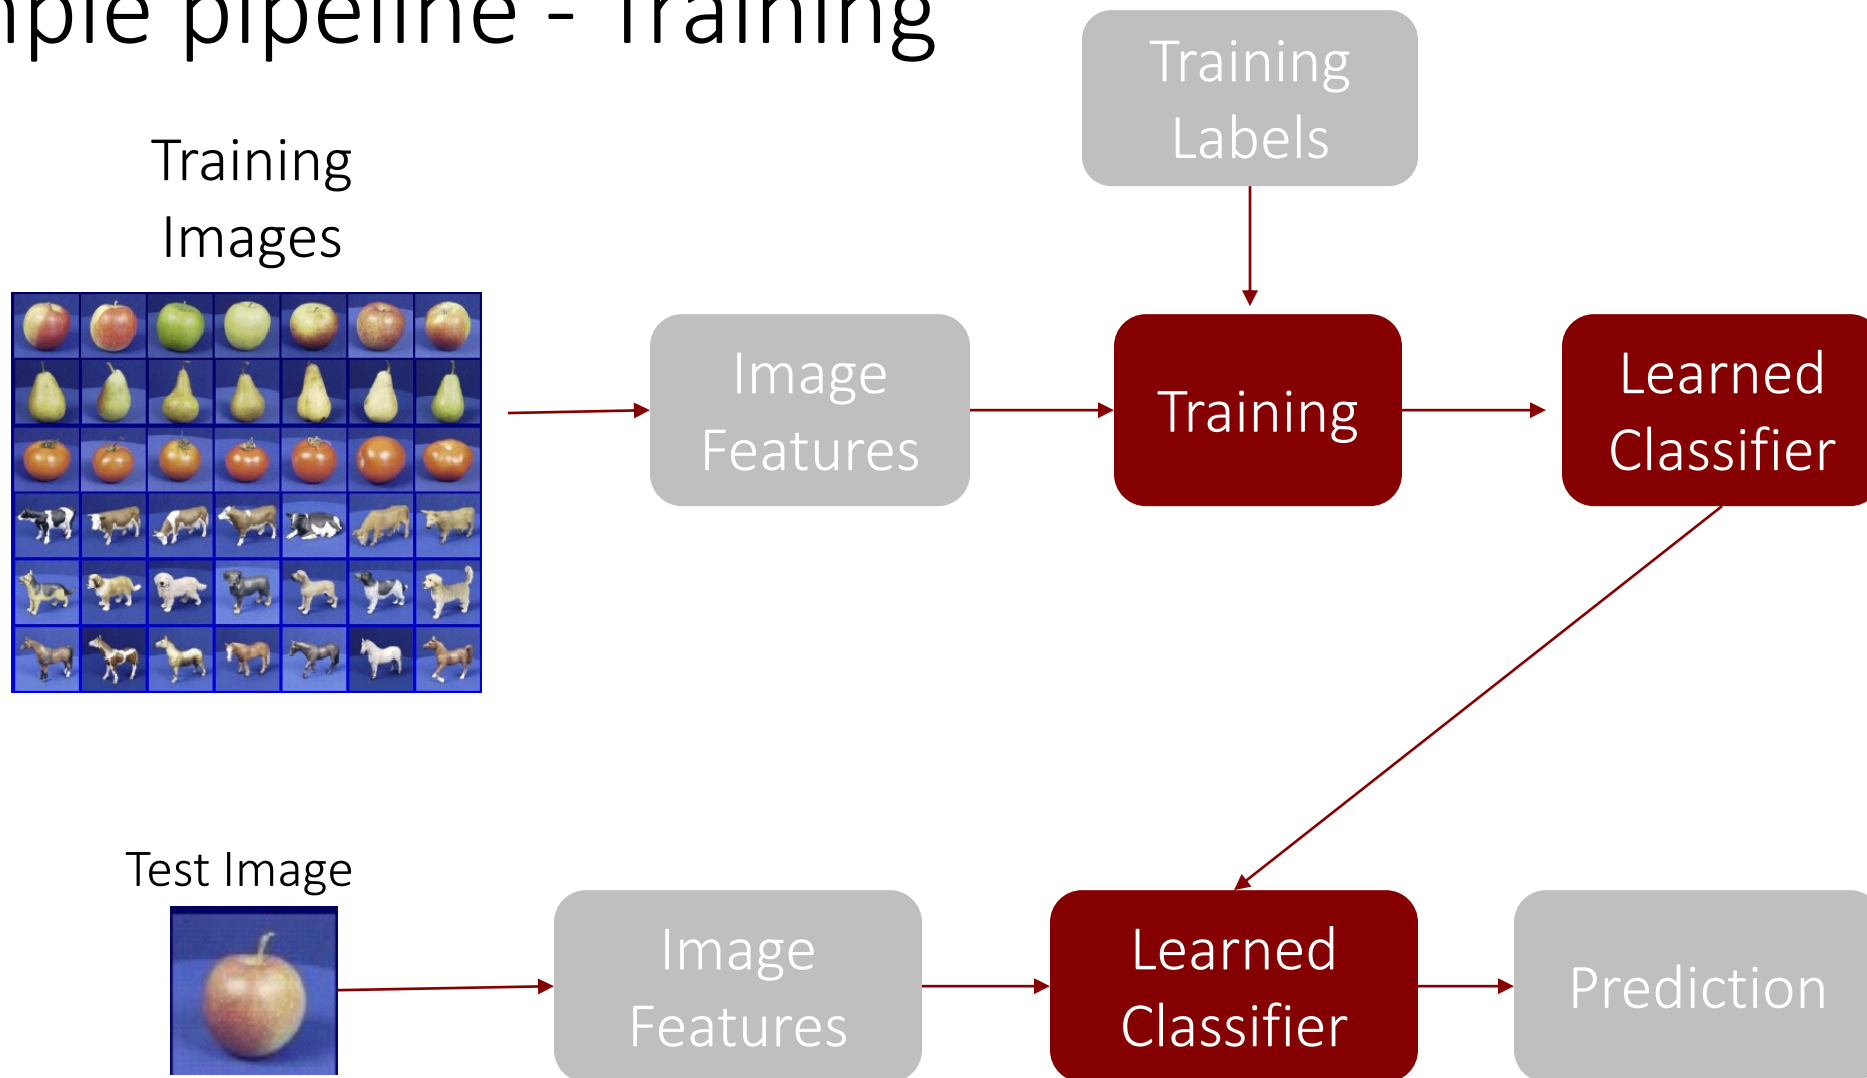
Many features to choose from

- Histogram of Color
- Histogram of Gradients
- SIFT
- Bag of words
- etc...



[Image From ETH-80 dataset: Analyzing Appearance and Contour Based Methods for Object Categorization - by Leibe et al 2003](#)

A simple pipeline - Training



Many classifiers to choose from

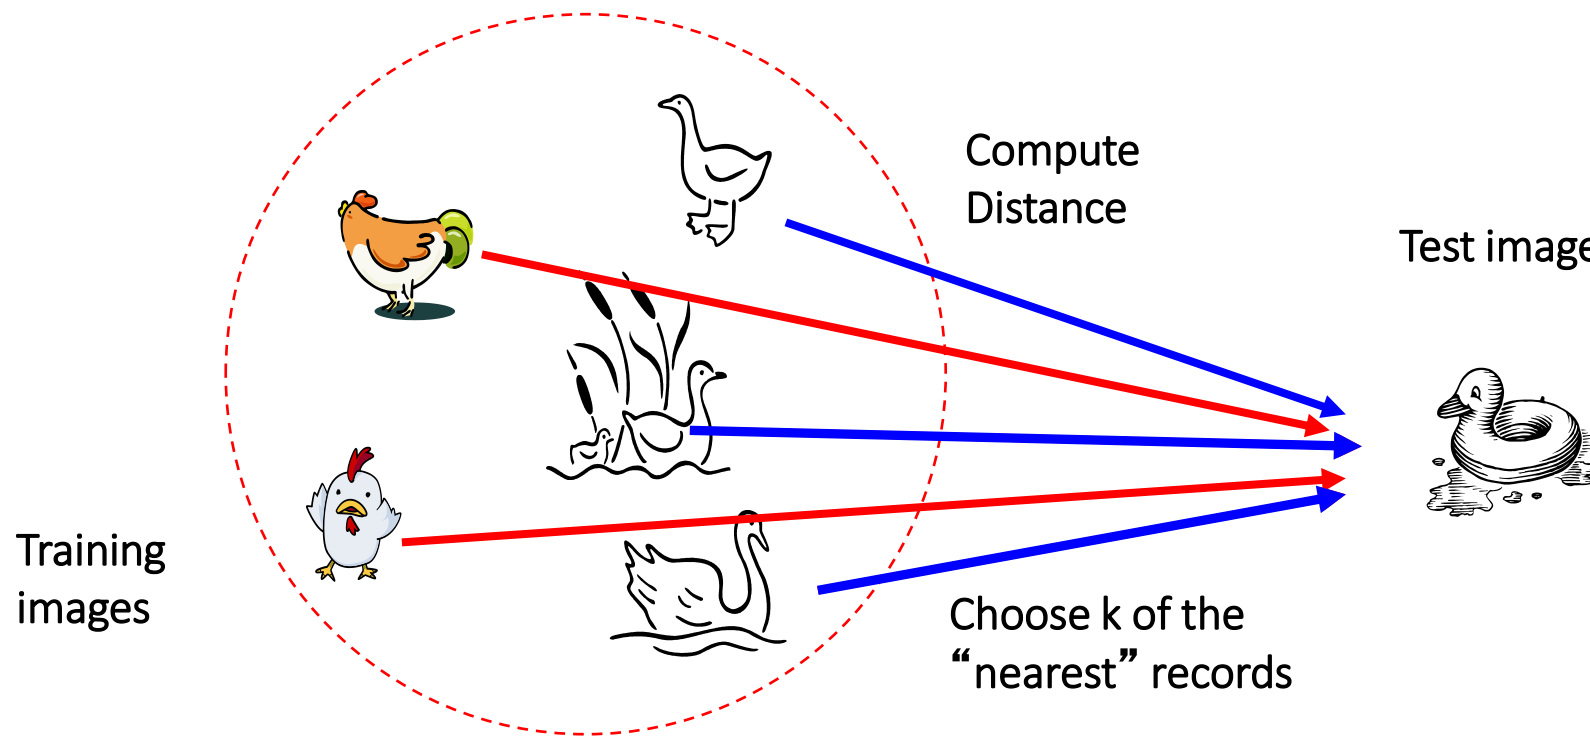
- **K-nearest neighbor**
- SVM
- Neural networks
- Naïve Bayes
- Bayesian network
- Logistic regression
- Randomized Forests
- Boosted Decision Trees
- Restricted Boltzmann Machines
- Etc.

Which is the best one?

Slide credit: D. Hoiem

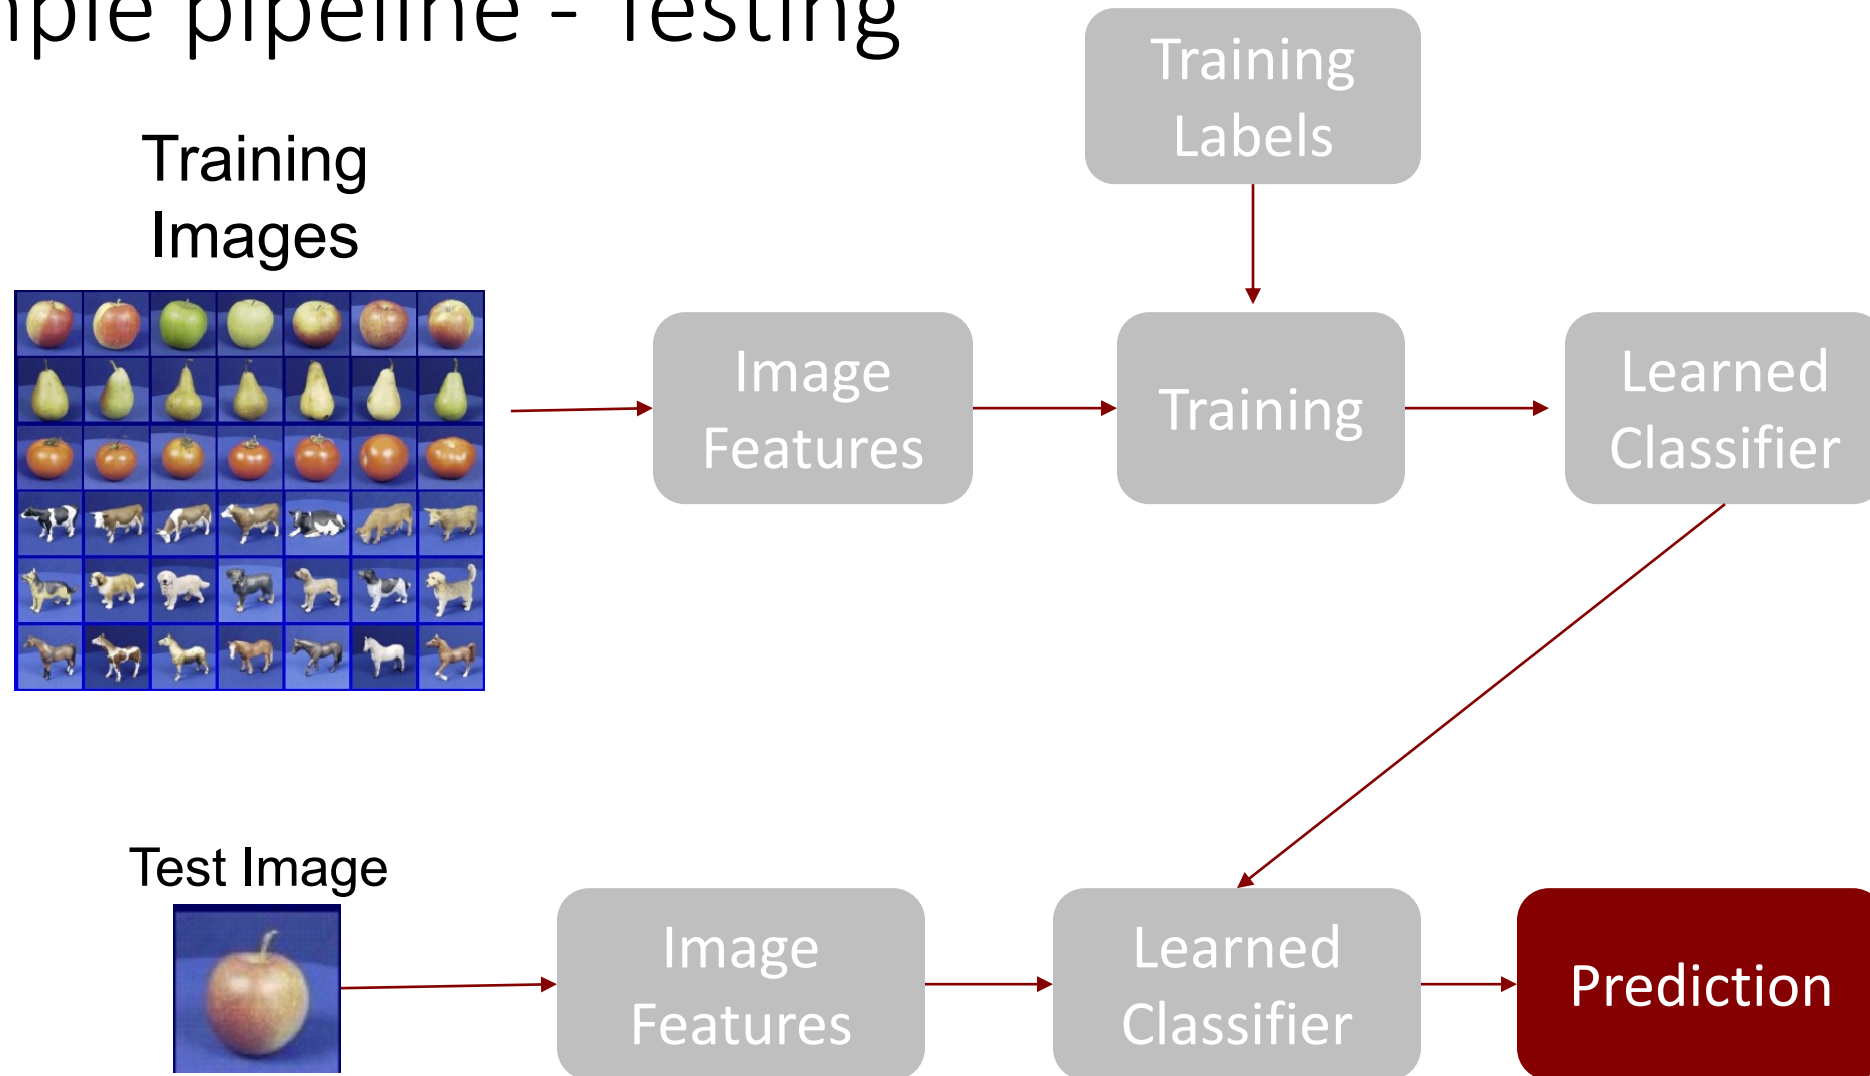
Classifiers: Nearest neighbor

Assign label of nearest training data point to each test data point




Source: N. Goyal

A simple pipeline - Testing




Classifiers: Nearest neighbor



Slide credit: L. Lazebnik

Classifiers: Nearest neighbor




Slide credit: L. Lazebnik

Today's Agenda

- A simple Image Classification pipeline
 - Classification overview
- K-nearest neighbor algorithm
 - kNN: algorithm
 - kNN: analysis

K-nearest neighbor

- Algorithm (training):
 1. Store all training data points x_i with their corresponding category labels y_i
- Algorithm (testing):
 1. We are given a new test point x
 2. Compute distance to all training data points
 3. Select k training points closest to x
 4. Assign x to label y that is most common among the k nearest neighbors.




- Distance measurement: – Euclidean


$$Dist(X^n, X^m) = \sqrt{\sum_{i=1}^D (X_i^n - X_i^m)^2}$$

Where X^n and X^m are the n-th and m-th data points


1-nearest neighbor



3-nearest neighbor



5-nearest neighbor



Today's Agenda


- A simple Image Classification pipeline
 - Classification overview
- K-nearest neighbor algorithm
 - kNN: algorithm
 - kNN: analysis

K-NN: a very useful algorithm

- Simple, a good one to try first
- Very flexible decision boundaries

K-NN: issues to keep in mind


- Choosing the value of k:
 - If too small, sensitive to noise points
 - If too large, neighborhood may include points from other classes



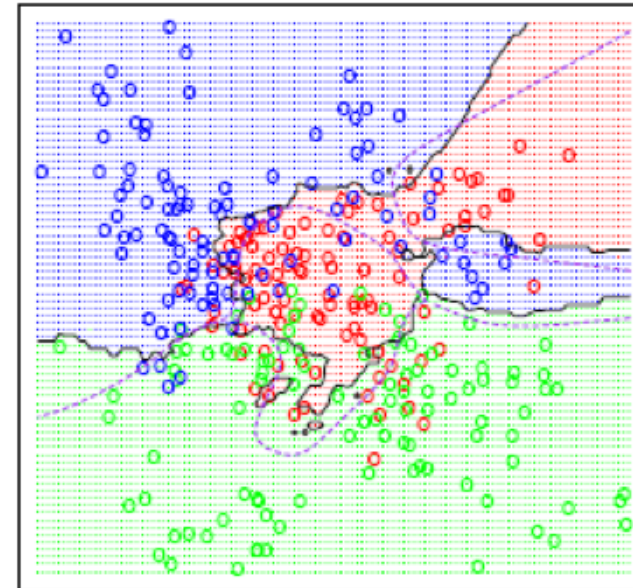
K-NN: issues to keep in mind

- Choosing the value of k :
 - If too small, sensitive to noise points
 - If too large, neighborhood may include points from other classes

K=1



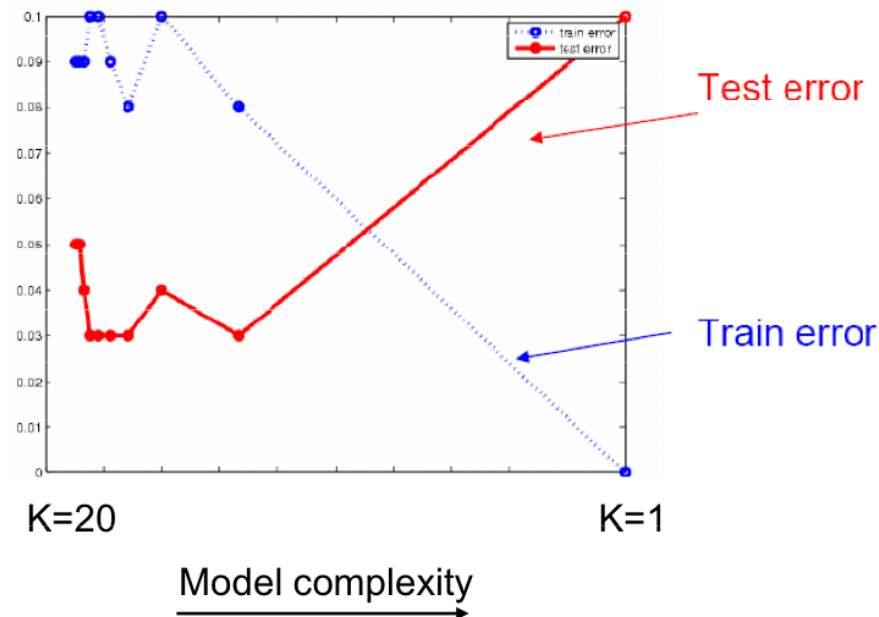
K=15





K-NN: issues to keep in mind

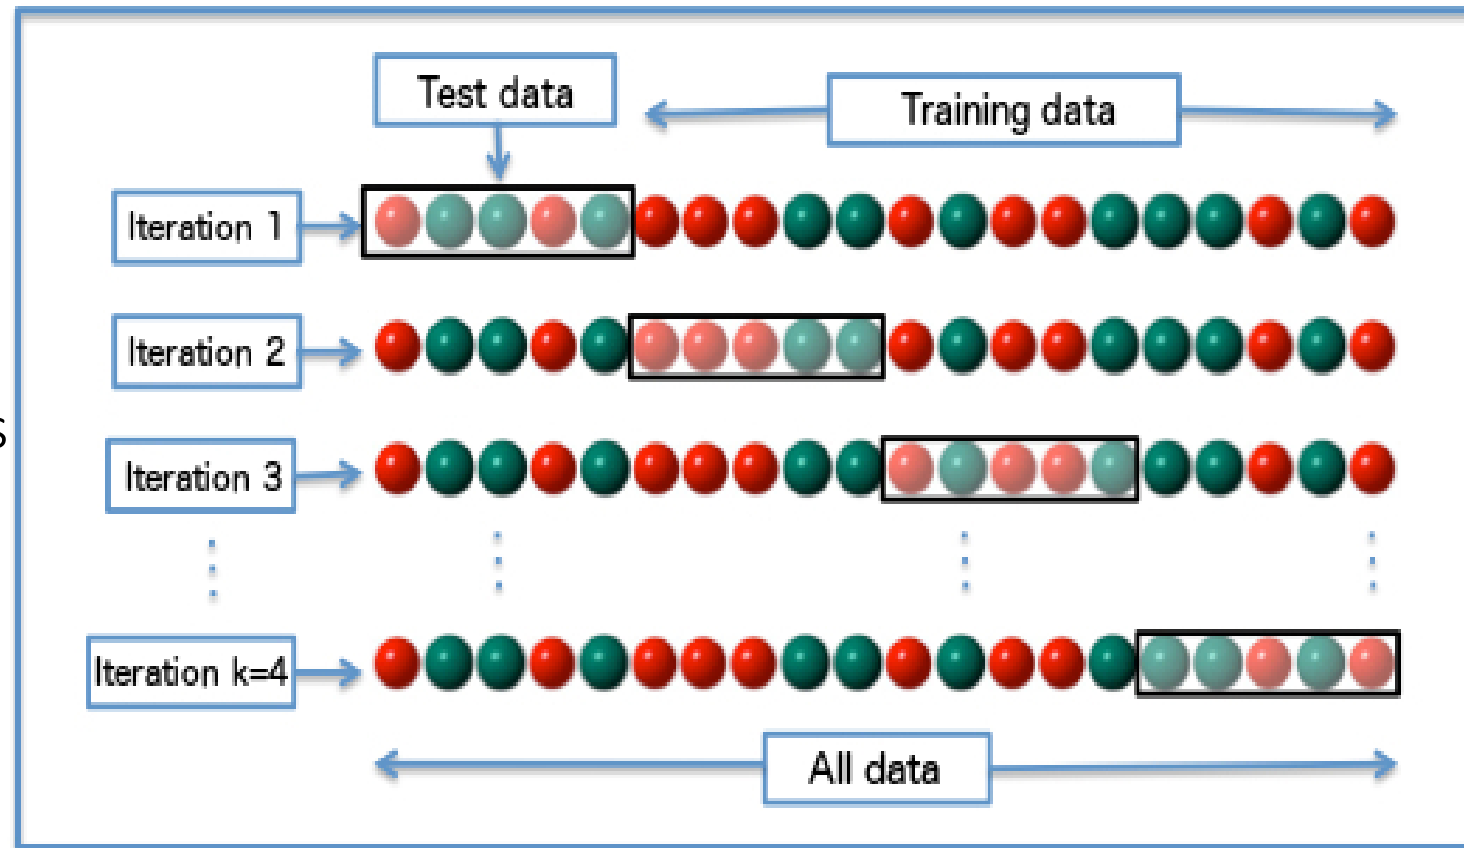
- Choosing the value of k :
 - If too small, sensitive to noise points
 - If too large, neighborhood may include points from other classes
 - **Solution**: cross validate!



Cross validation

- For each value of k in the nearest neighbors algorithm:
- Create multiple train/test splits
 - For each split:
 - Measure performance
- Average performance over all splits

- Select k with best average performance



K-NN: issues to keep in mind

- Choosing the value of k :
 - If too small, sensitive to noise points
 - If too large, neighborhood may include points from other classes
 - **Solution**: cross validate!
- Can produce counter-intuitive results (using Euclidean measure)

Euclidean measure

1 1 1 1 1 1 1 1 1 1 0

0 1 1 1 1 1 1 1 1 1 1

$d = 1.4142$

VS

1 0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0 1

$d = 1.4142$

K-NN: issues to keep in mind


- Choosing the value of k :
 - If too small, sensitive to noise points
 - If too large, neighborhood may include points from other classes
 - **Solution**: cross validate!
- Can produce counter-intuitive results (using Euclidean measure)
 - **Solution**: normalize the vectors to unit length

K-NN: issues to keep in mind

- Choosing the value of k :
 - If too small, sensitive to noise points
 - If too large, neighborhood may include points from other classes
 - Solution: cross validate!
- Can produce counter-intuitive results (using Euclidean measure)
 - Solution: normalize the vectors to unit length
- Curse of Dimensionality


Curse of dimensionality

- Assume 5000 points uniformly distributed in the unit (hyper)cube and we want to apply 5-NN. Suppose our query point is at the origin.
 - In 1-dimension, we must go a distance of $5/5000=0.001$ on average to capture 5 nearest neighbors.
 - In 2 dimensions, we must go $\sqrt{0.001}$ to get a square that contains 0.001 of the volume.
 - In d dimensions, we must go $(0.001)^{1/d}$



Curse of dimensionality

- Assume 5000 points uniformly distributed in the unit hypercube and we want to apply 5-NN. Suppose our query point is at the origin.
 - In 1-dimension, we must go a distance of $5/5000=0.001$ on average to capture 5 nearest neighbors.
 - In 2 dimensions, we must go $\sqrt{0.001}$ to get a square that contains 0.001 of the volume.
 - In d dimensions, we must go $(0.001)^{1/d}$



K-NN: issues to keep in mind

- Choosing the value of k :
 - If too small, sensitive to noise points
 - If too large, neighborhood may include points from other classes
 - **Solution**: cross validate!
- Can produce counter-intuitive results (using Euclidean measure)
 - **Solution**: normalize the vectors to unit length
- Curse of Dimensionality
 - **Solution**: no good one – need to get more data

MAI4CAREU

Master programmes in Artificial
Intelligence 4 Careers in Europe



CYENS
CENTRE OF EXCELLENCE



Thank you.

