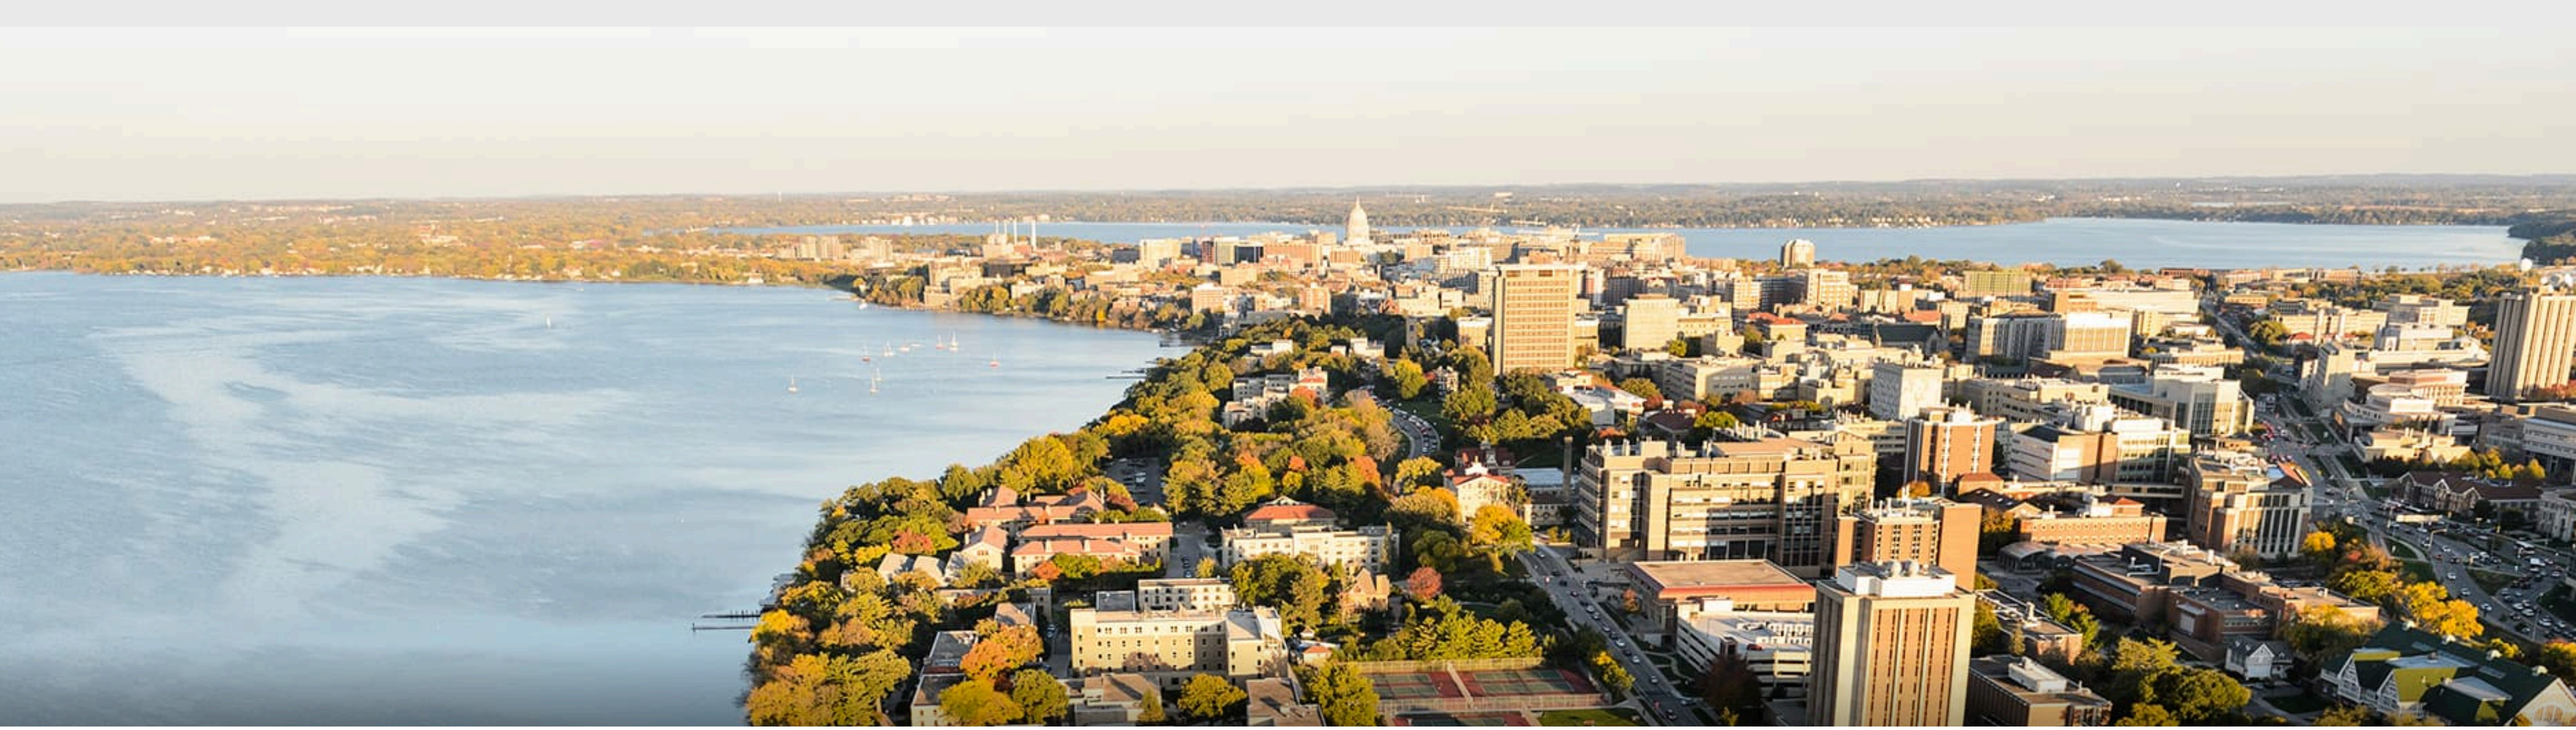

CS 540 Introduction to Artificial Intelligence
Machine Learning Overview

Sharon Yixuan Li
University of Wisconsin-Madison

Feb 16, 2021

Announcement

- HW **release dates** are now available on our website

HW4 (clustering)	Tuesday Feb 23
HW5 (regression)	Tuesday March 2
HW6 (neural networks)	Tuesday March 9
HW7 (deep learning)	Tuesday March 23
HW8 (game)	Tuesday April 6
HW9 (search)	Tuesday April 13
HW10 (RL)	Tuesday April 20

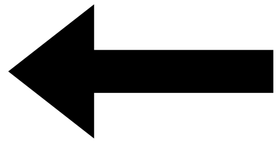

Next week :)

- HW3 review on Thursday

Today’s outline

- What is machine learning?
- Supervised Learning
  - Classification
  - Regression
- Unsupervised Learning
  - Clustering

Part I: What is machine learning?

https://www.youtube.com/watch?v=ukzFI9rgwfU

HUMANS LEARN FROM
PAST EXPERIENCES

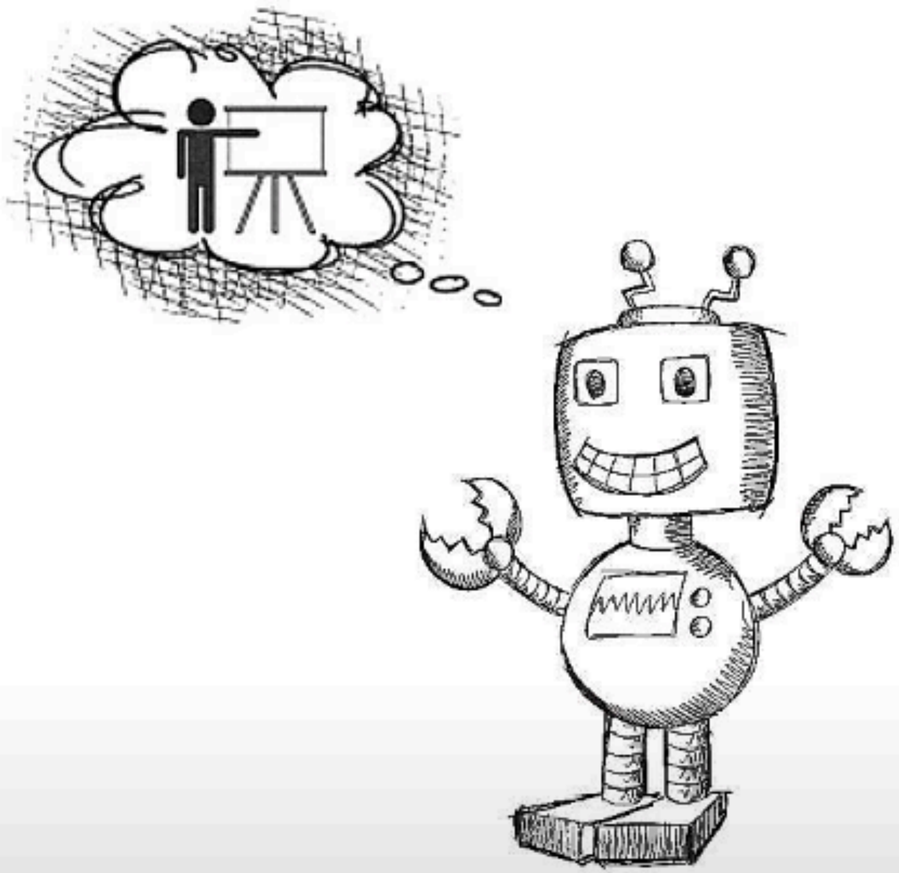

MACHINES FOLLOW INSTRUCTIONS
GIVEN BY HUMANS

What is machine learning?

- Arthur Samuel (1959): Machine learning is the field of study that gives the computer the ability to learn **without being explicitly programmed**.

Without Machine Learning

With Machine Learning

https://tung-dn.github.io/programming.html

What is machine learning?

- Arthur Samuel (1959): Machine learning is the field of study that gives the computer the ability to learn **without being explicitly programmed**.

- Tom Mitchell (1997): A computer program is said to learn from **experience E** with respect to some class of **tasks T** and **performance measure P**, if its performance at tasks in T as measured by P, improves with experience E.

Taxonomy of ML

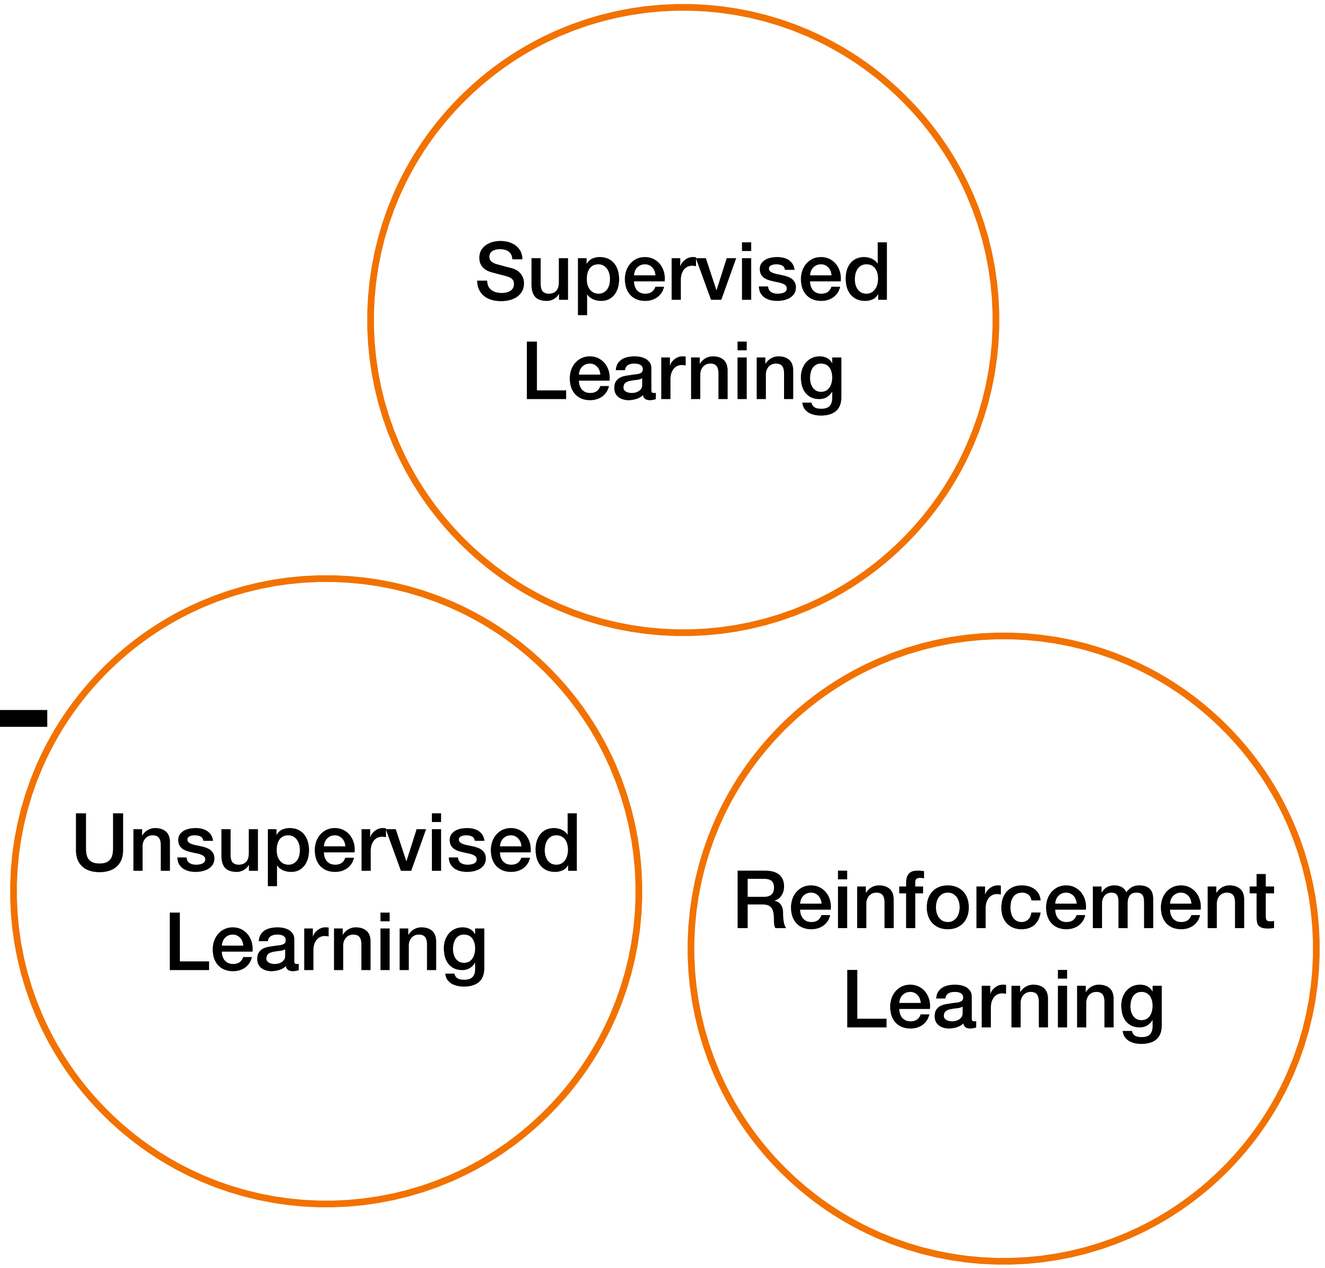

Part II: Supervised Learning

Example 1: Predict whether a user likes a song or not

model

Example 1: Predict whether a user likes a song or not

User Sharon

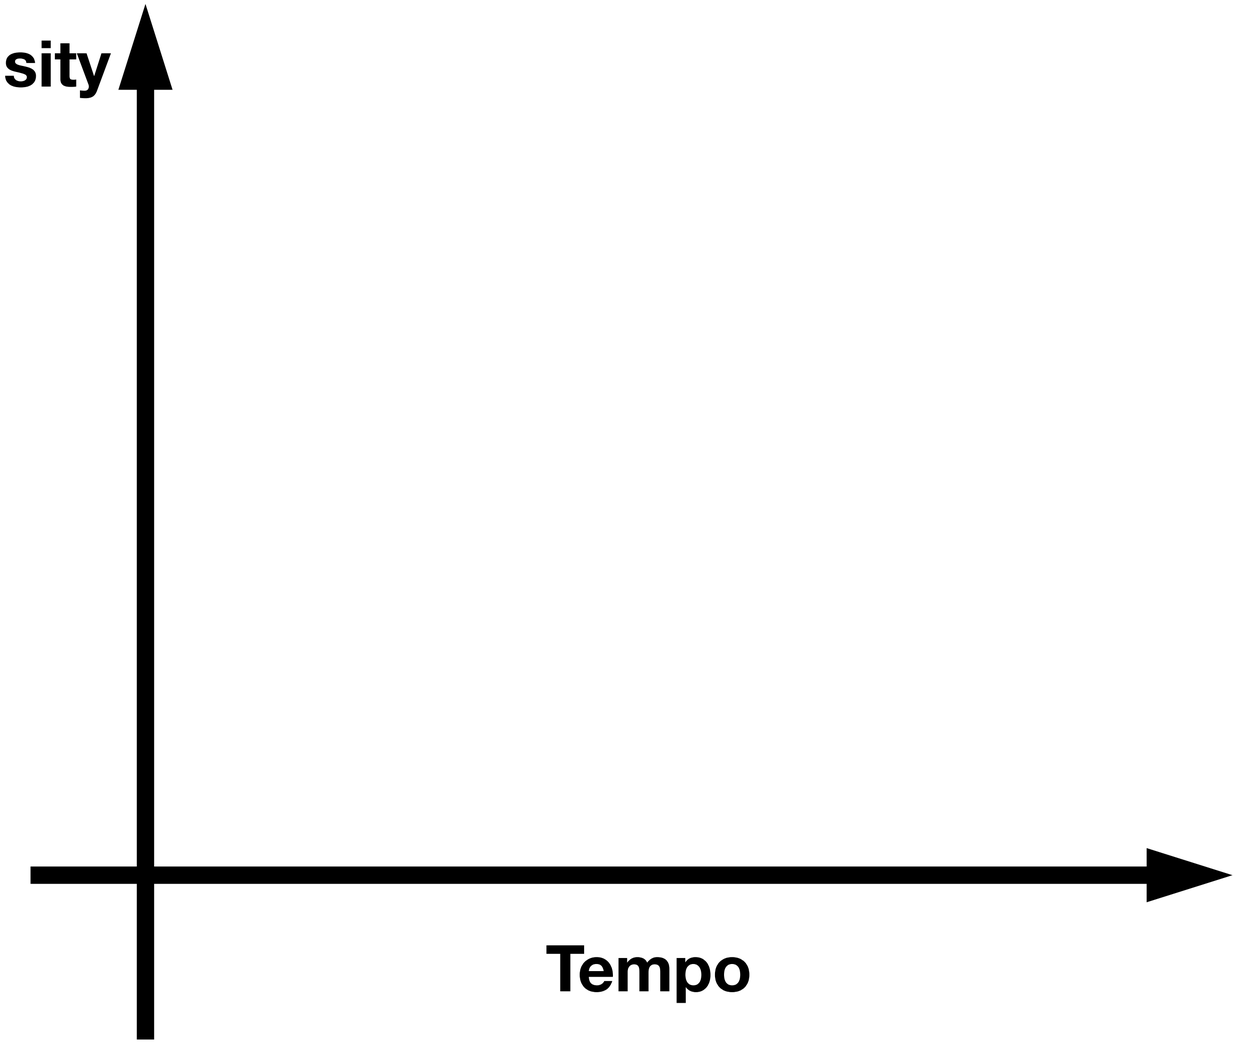

Example 1: Predict whether a user likes a song or not

User Sharon

DisLike

Like

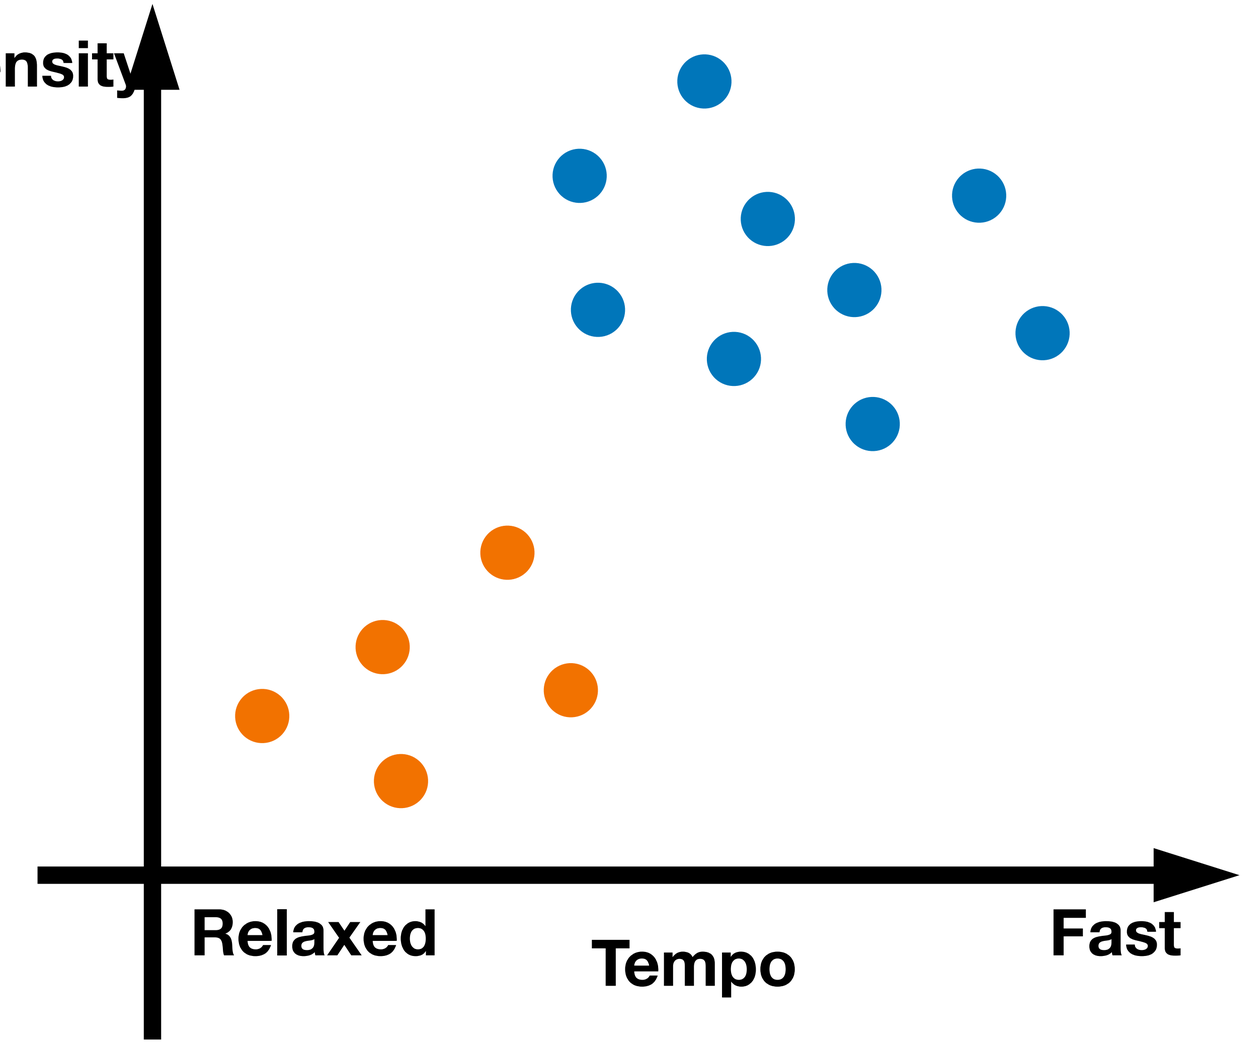

Example 1: Predict whether a user likes a song or not

User Sharon

DisLike

Like

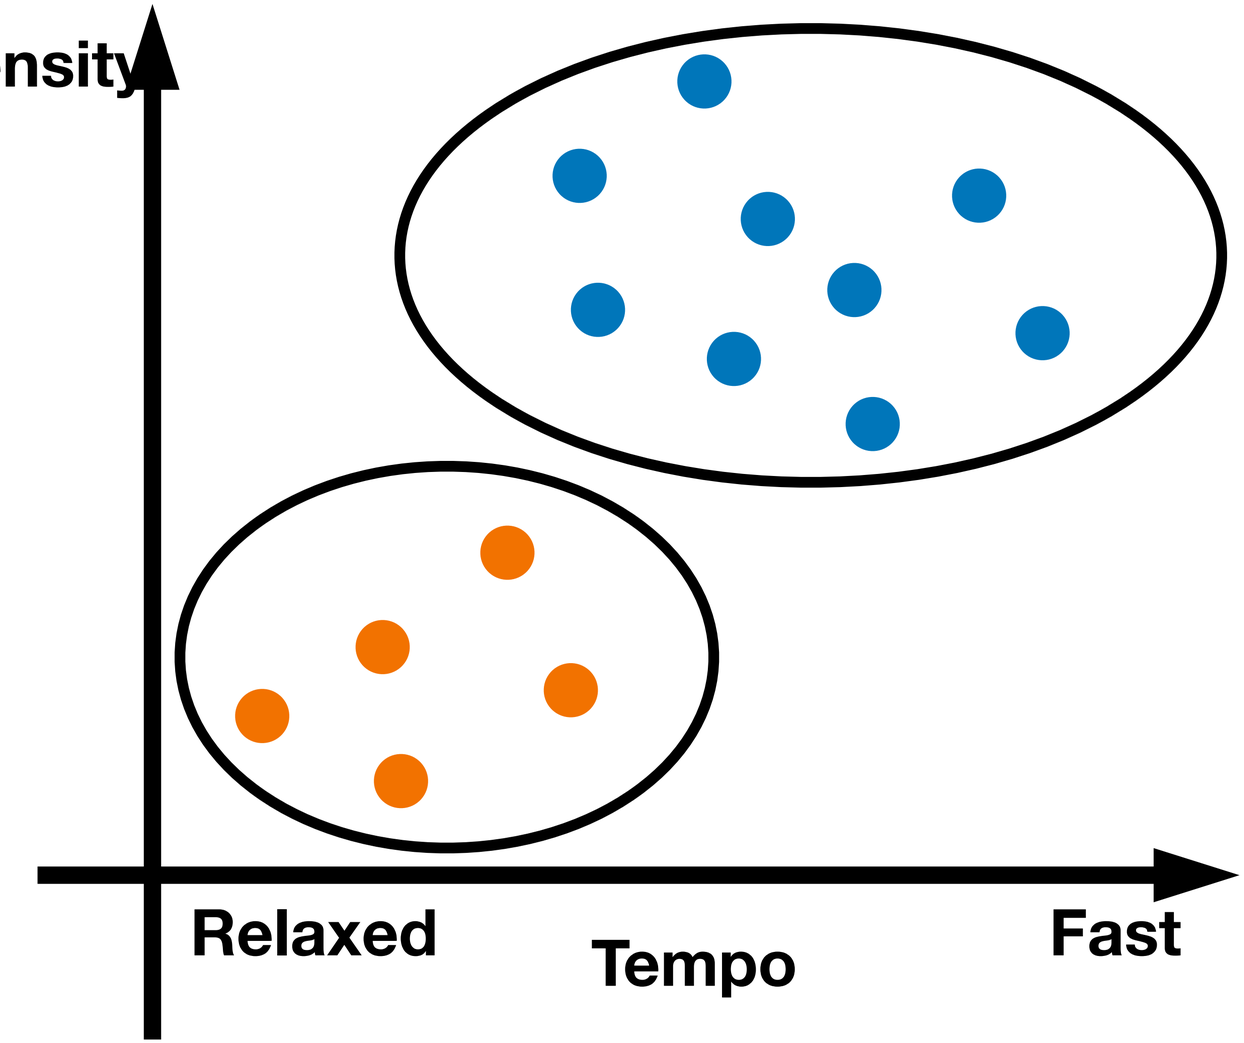

Example 1: Predict whether a user likes a song or not

User Sharon

DisLike

Like

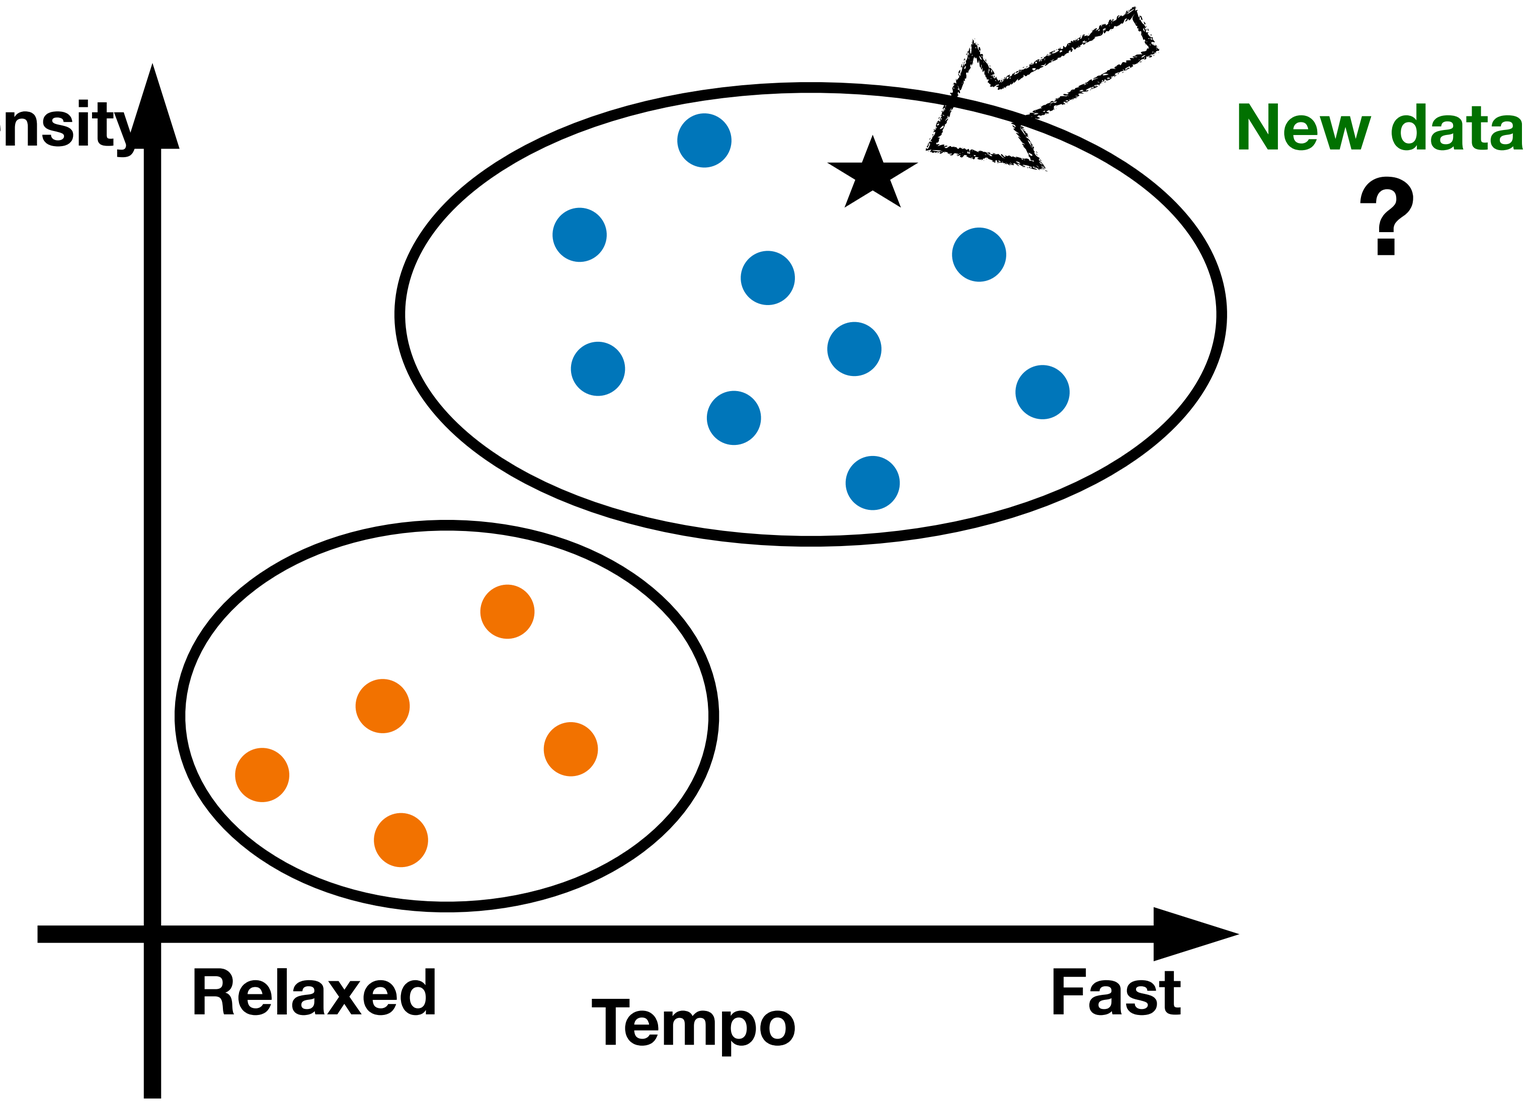

Example 1: Predict whether a user likes a song or not

User Sharon

DisLike

Like

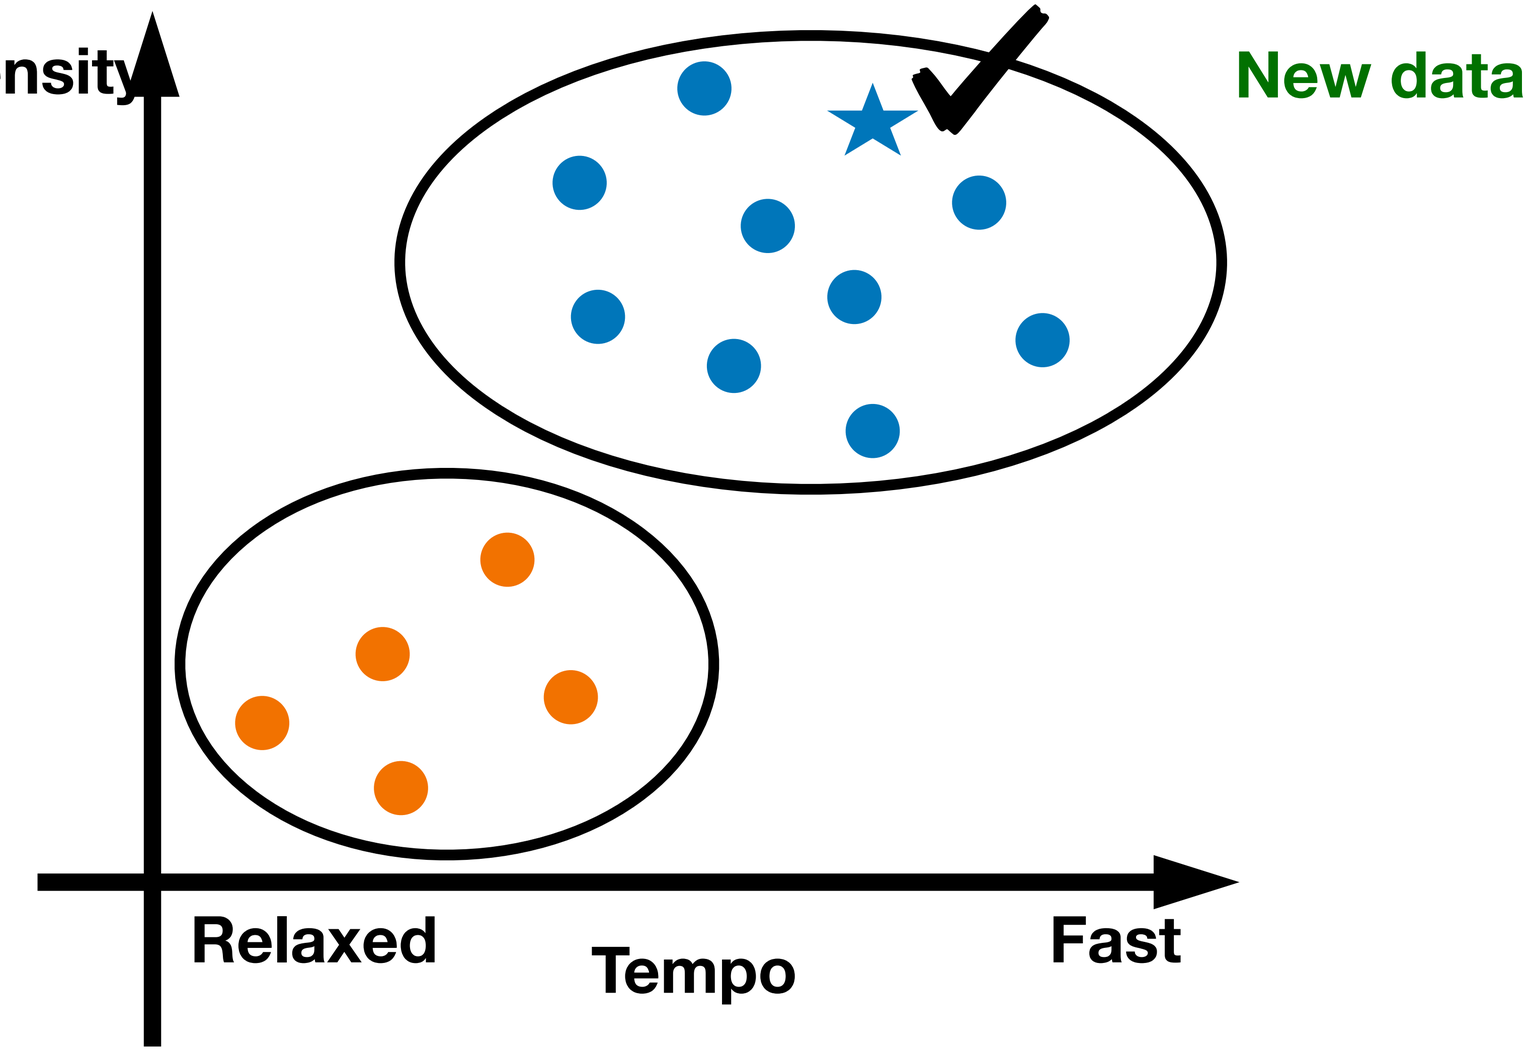

Example 2: Classify Images

http://www.image-net.org/

Example 2: Classify Images

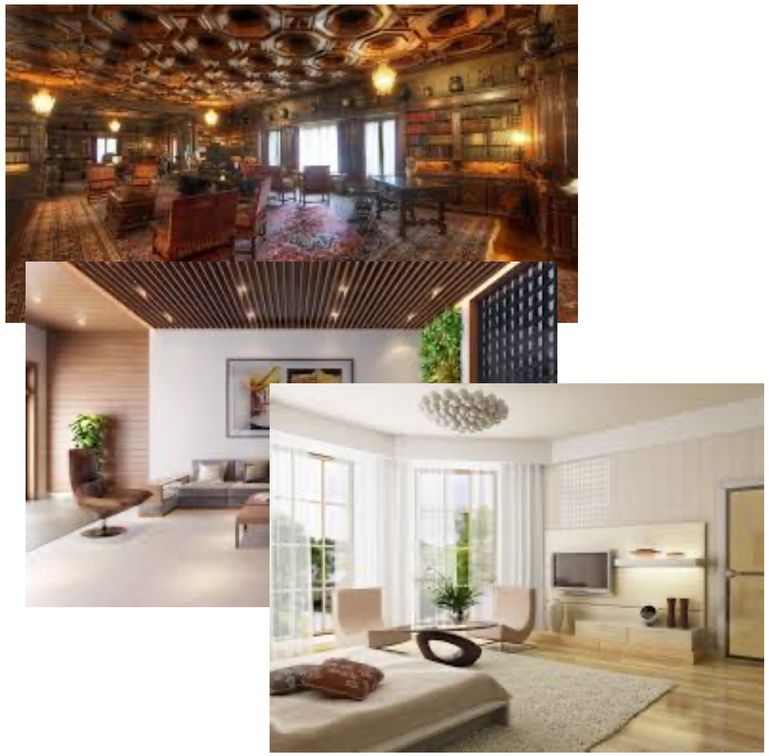

indoor

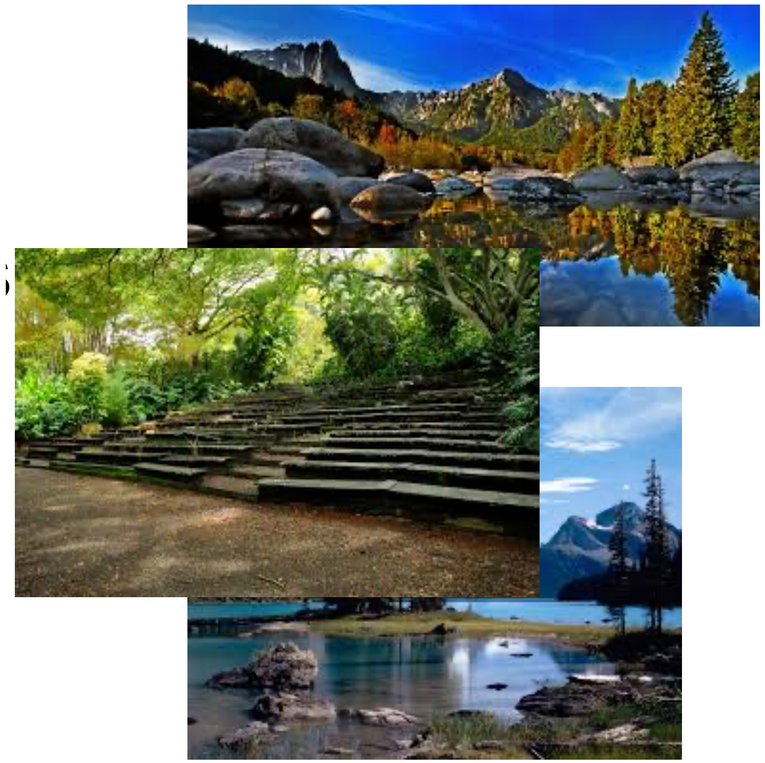

outdoor

Example 2: Classify Images

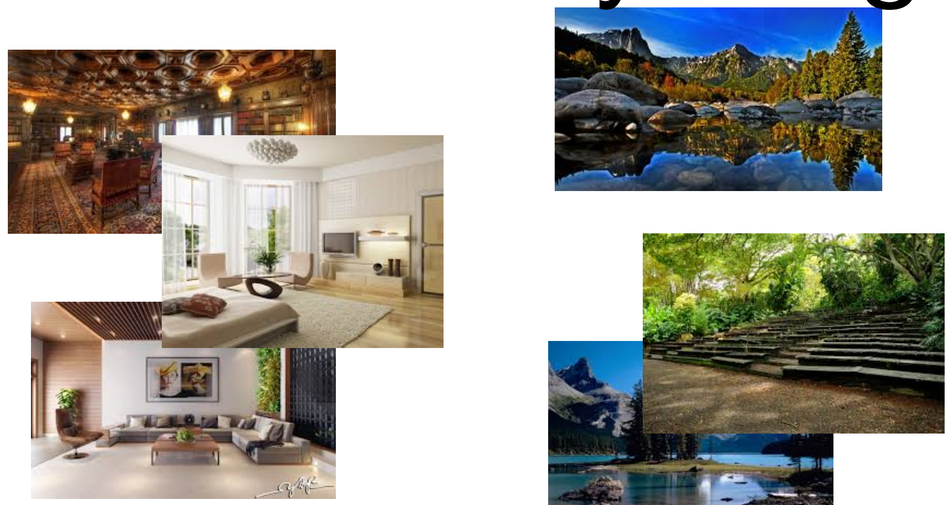

Training data

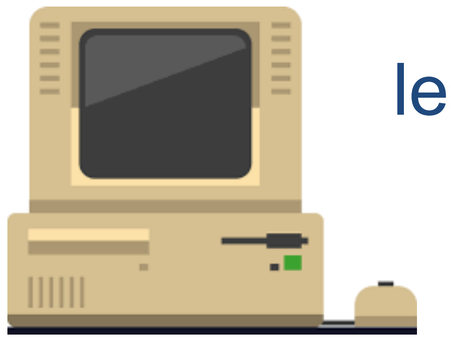

learning (i.e.,training)

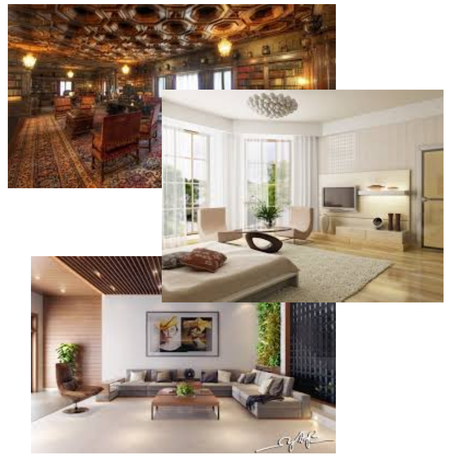

Training data

Label: outdoor

Label: indoor

Test data

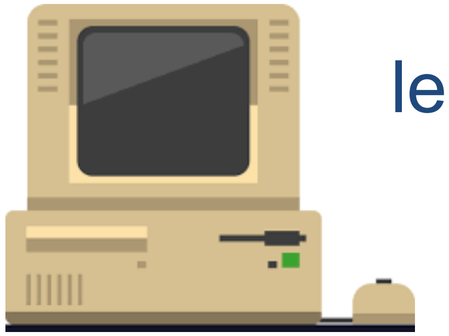

learning (i.e.,training)

testing

performance

How to represent data?

input data

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

d: feature dimension

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

x_1 Tempo

x_2 Intensity

There can be many features!

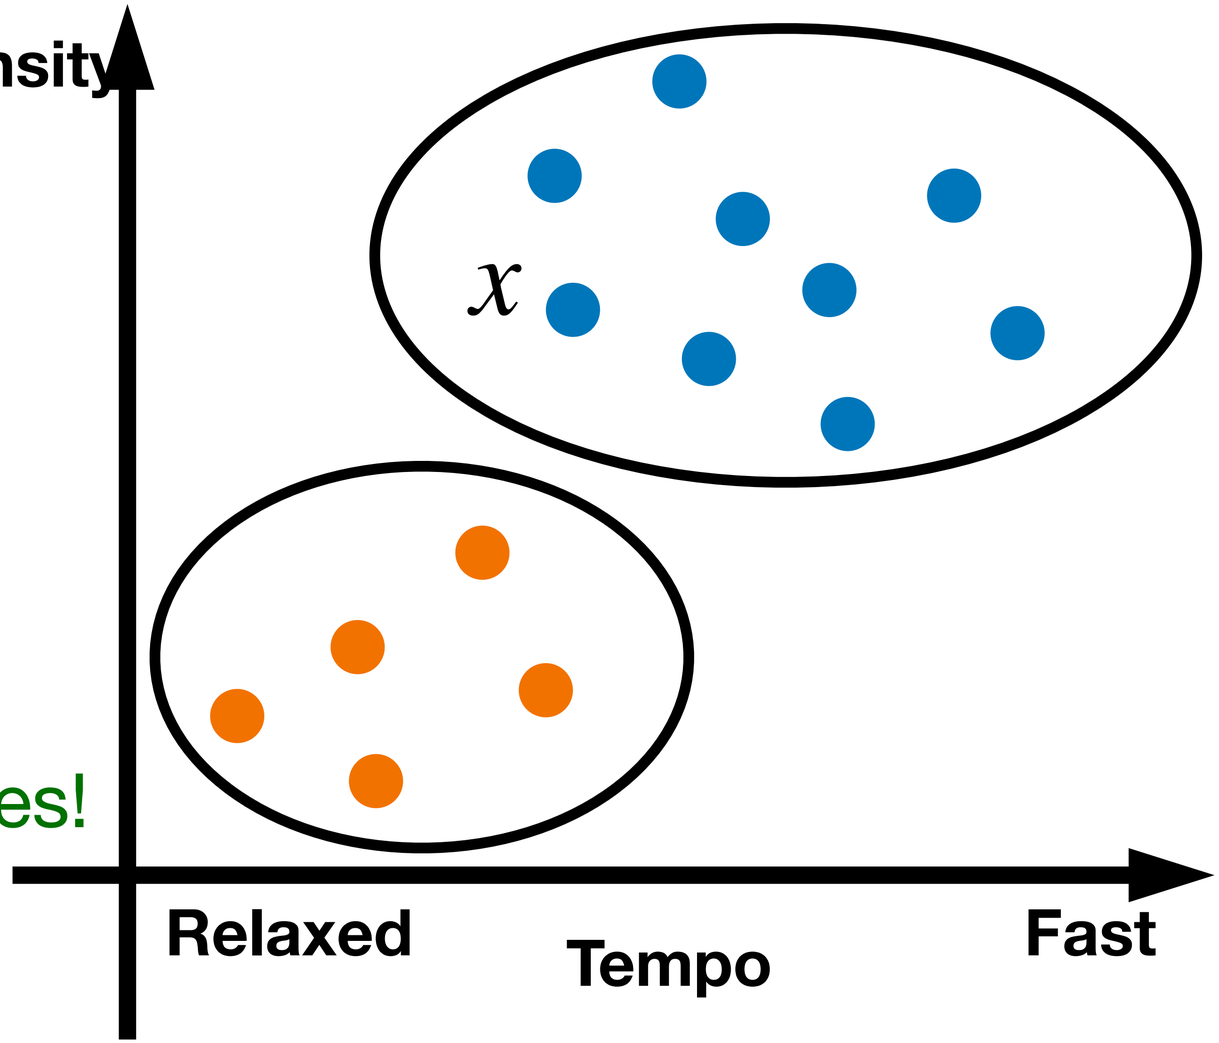

How to represent data?

Label

$y \in \{0,1\}$

Where “supervision” comes from

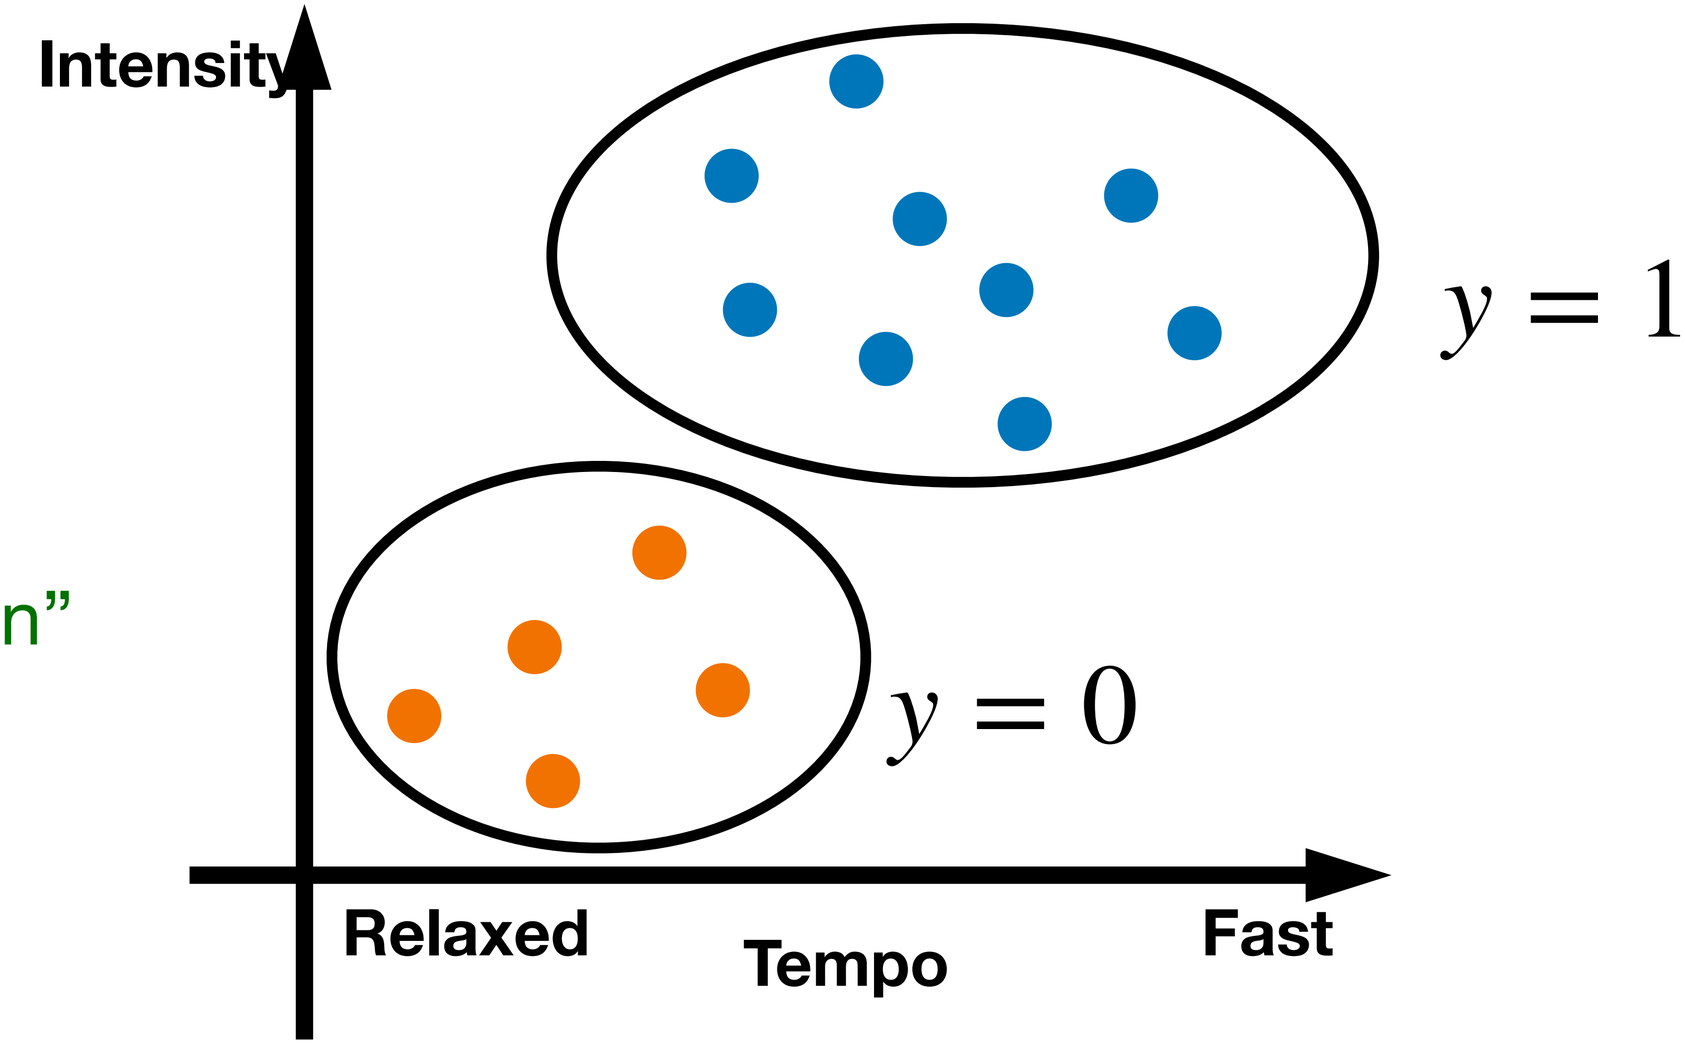

Represent various types of data

- Image
  - Pixel values

- Bank account
  - Credit rating, balance, # deposits in last day, week, month, year, #withdrawals

Two Types of Supervised Learning Algorithms

Classification

Regression

Example of regression: housing price prediction

Given: a dataset that contains n samples

$$(\mathbf{x}_1, y_2), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_n, y_n)$$

Task: if a residence has **x** squares feet, predict the price?

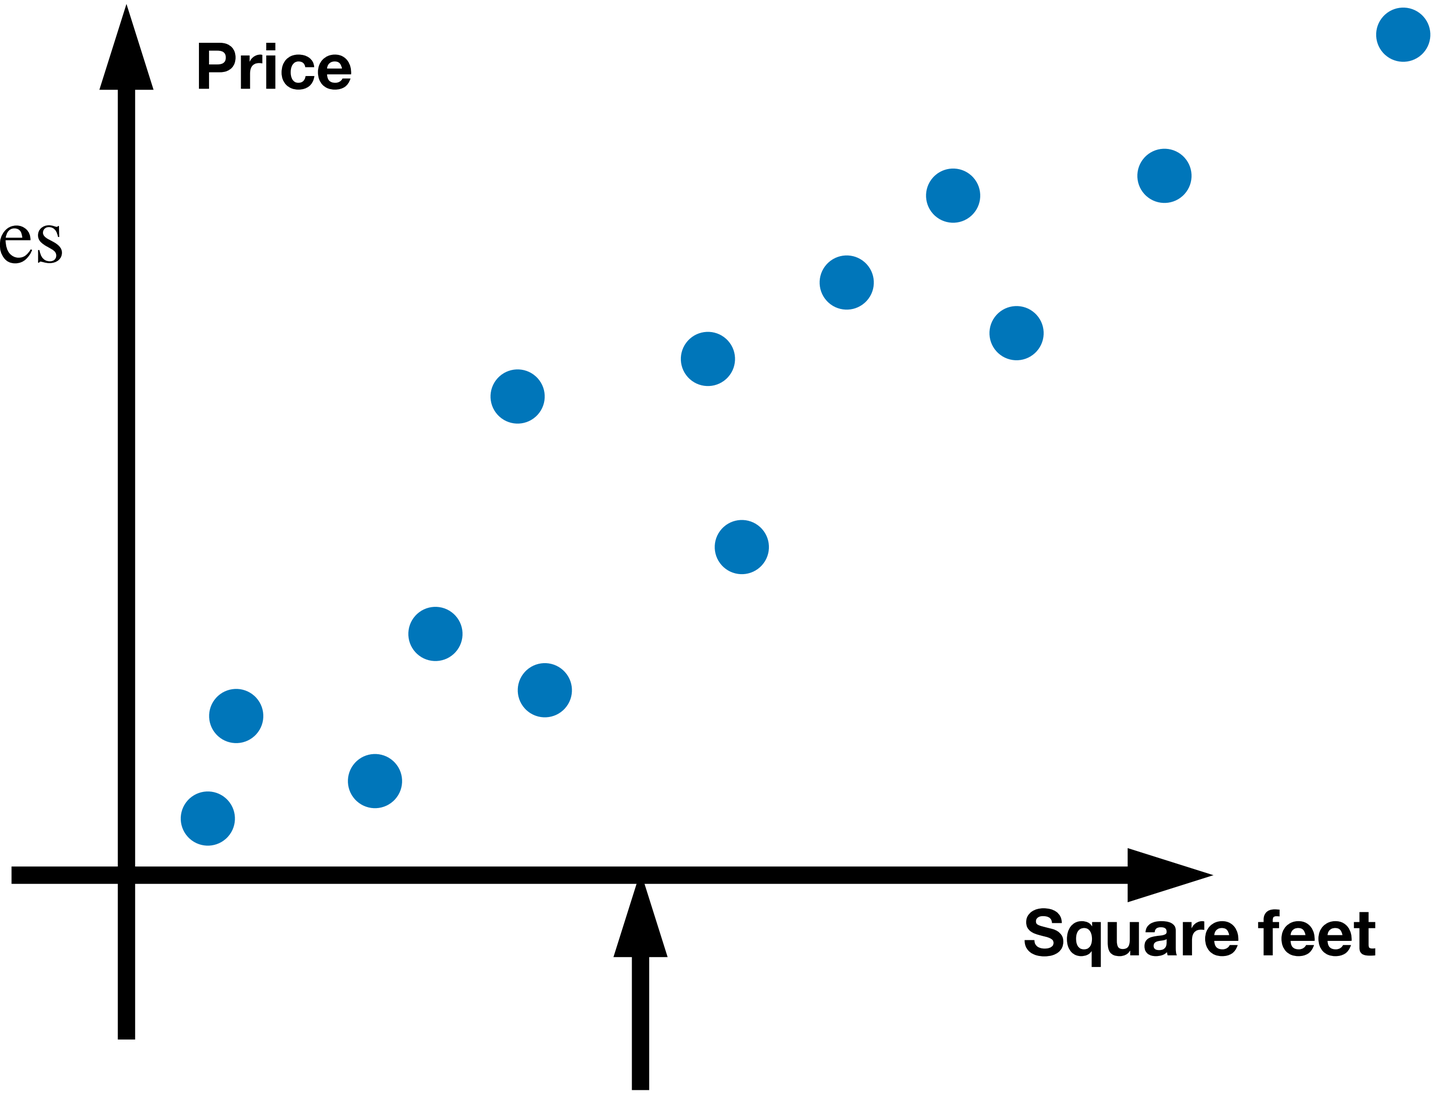

Example of regression: housing price prediction

Given: a dataset that contains n samples

$$(\mathbf{x}_1, y_2), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_n, y_n)$$

Task: if a residence has **x** squares feet, predict the price?

$$y \in \mathbb{R}$$

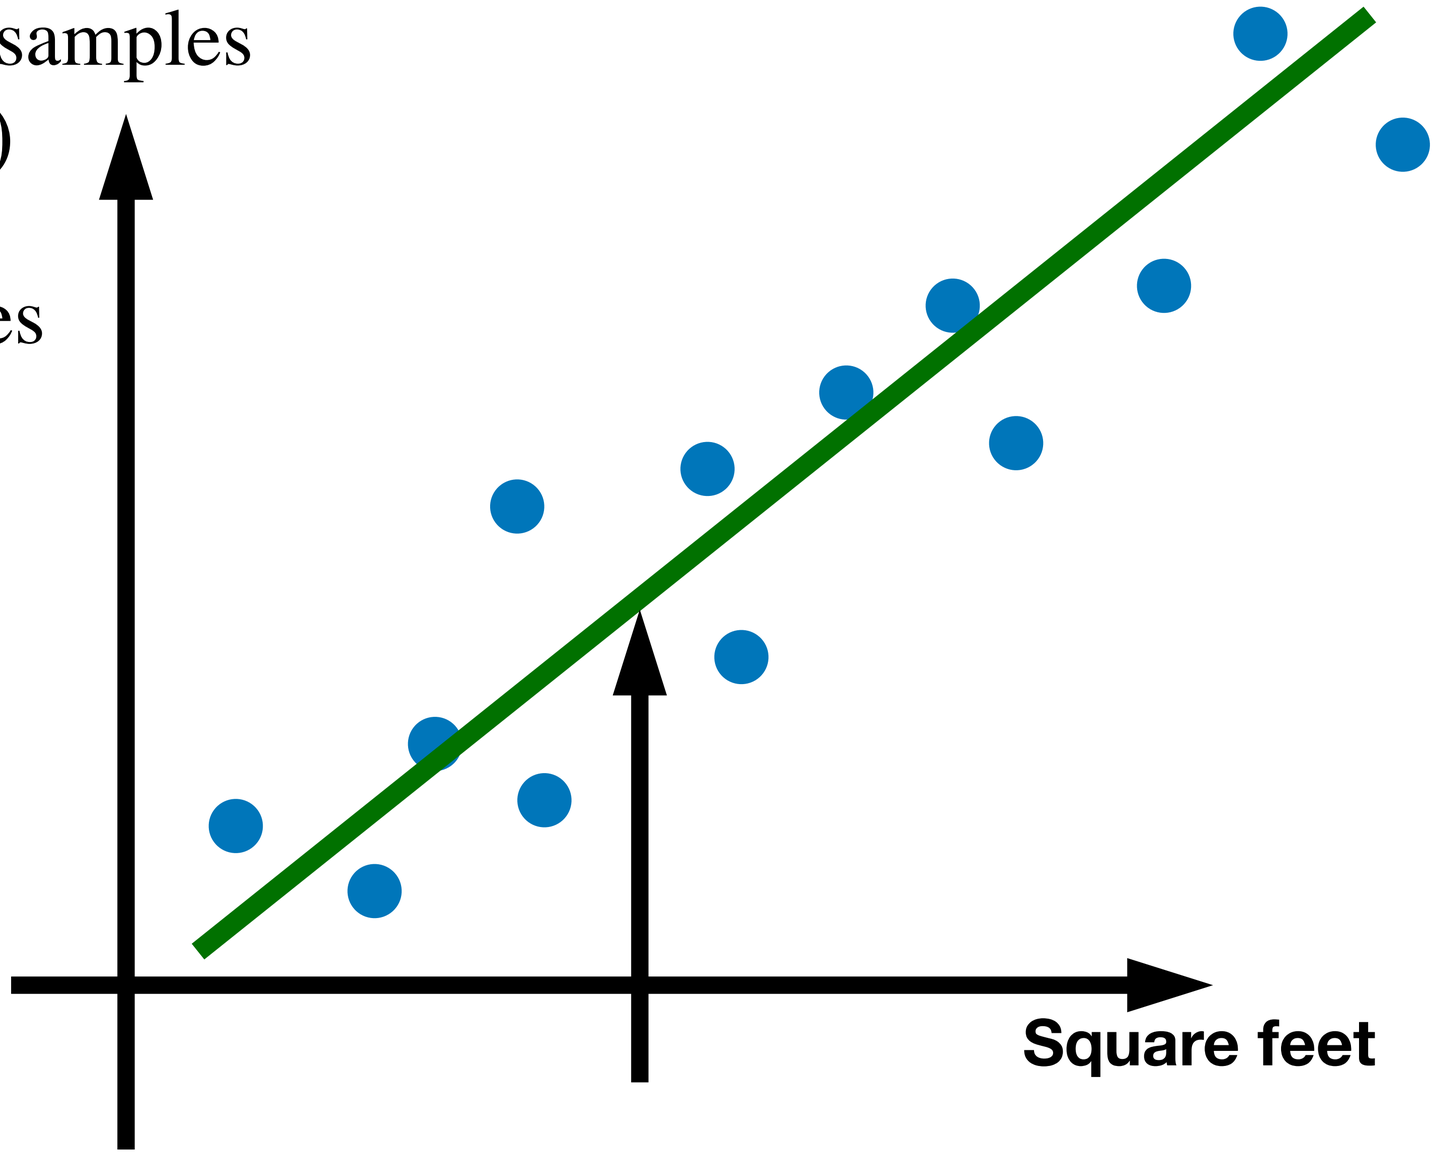

Example of regression: housing price prediction

Input with more features (e.g., lot size)

(size, lot size) → price

features/input
$x \in \mathbb{R}^2$

label/output
$y \in \mathbb{R}$

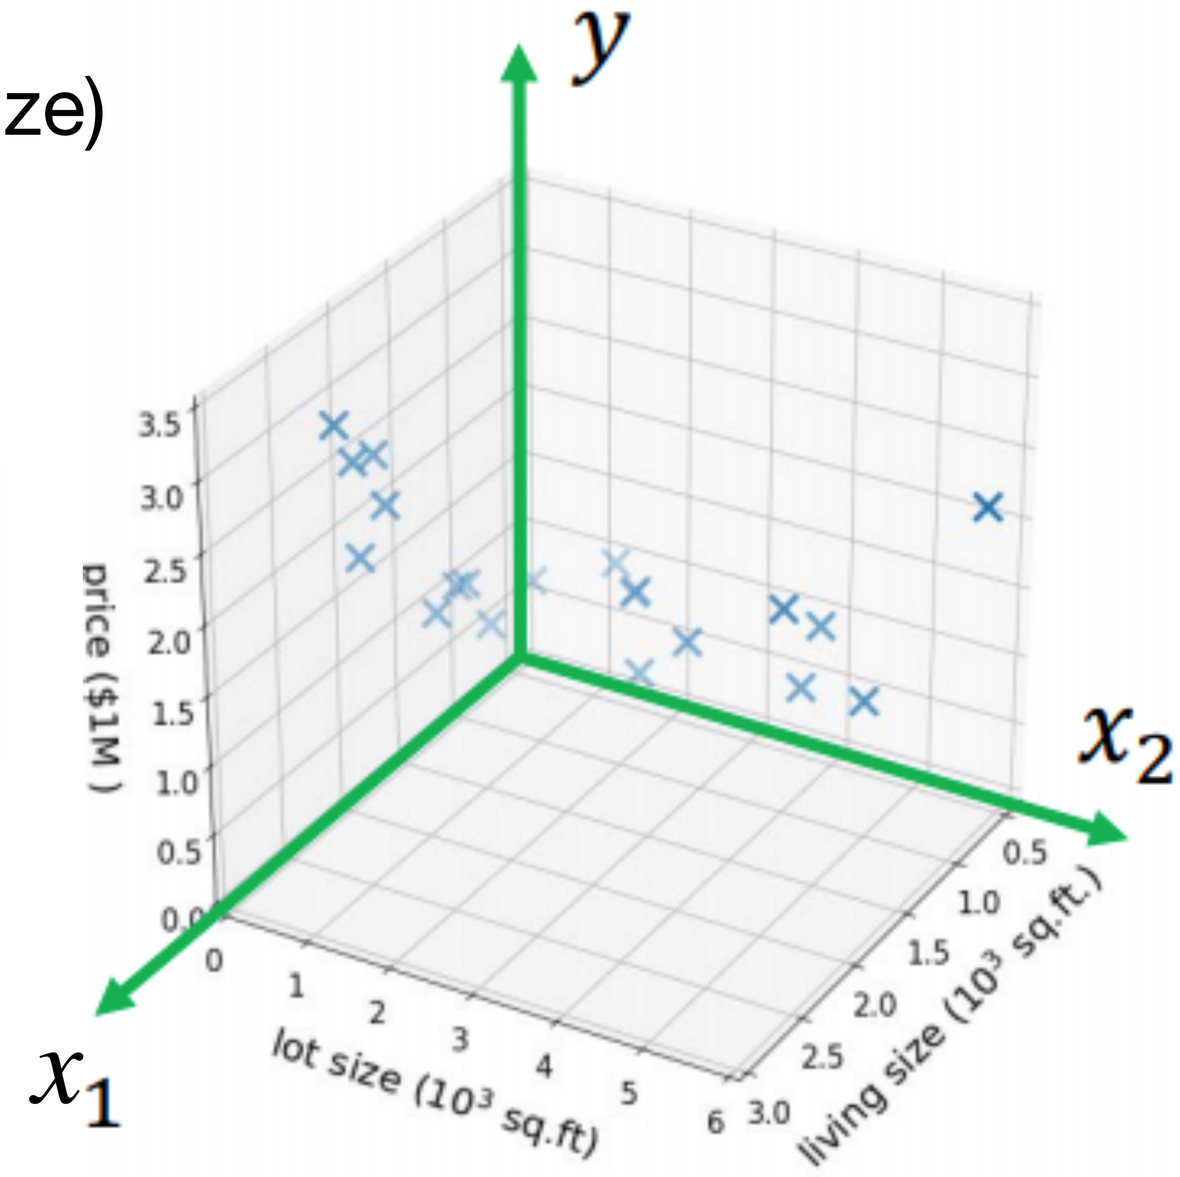

(credit: stanford CS229)

Supervised Learning: More examples

x = raw pixels of the image y = bounding boxes

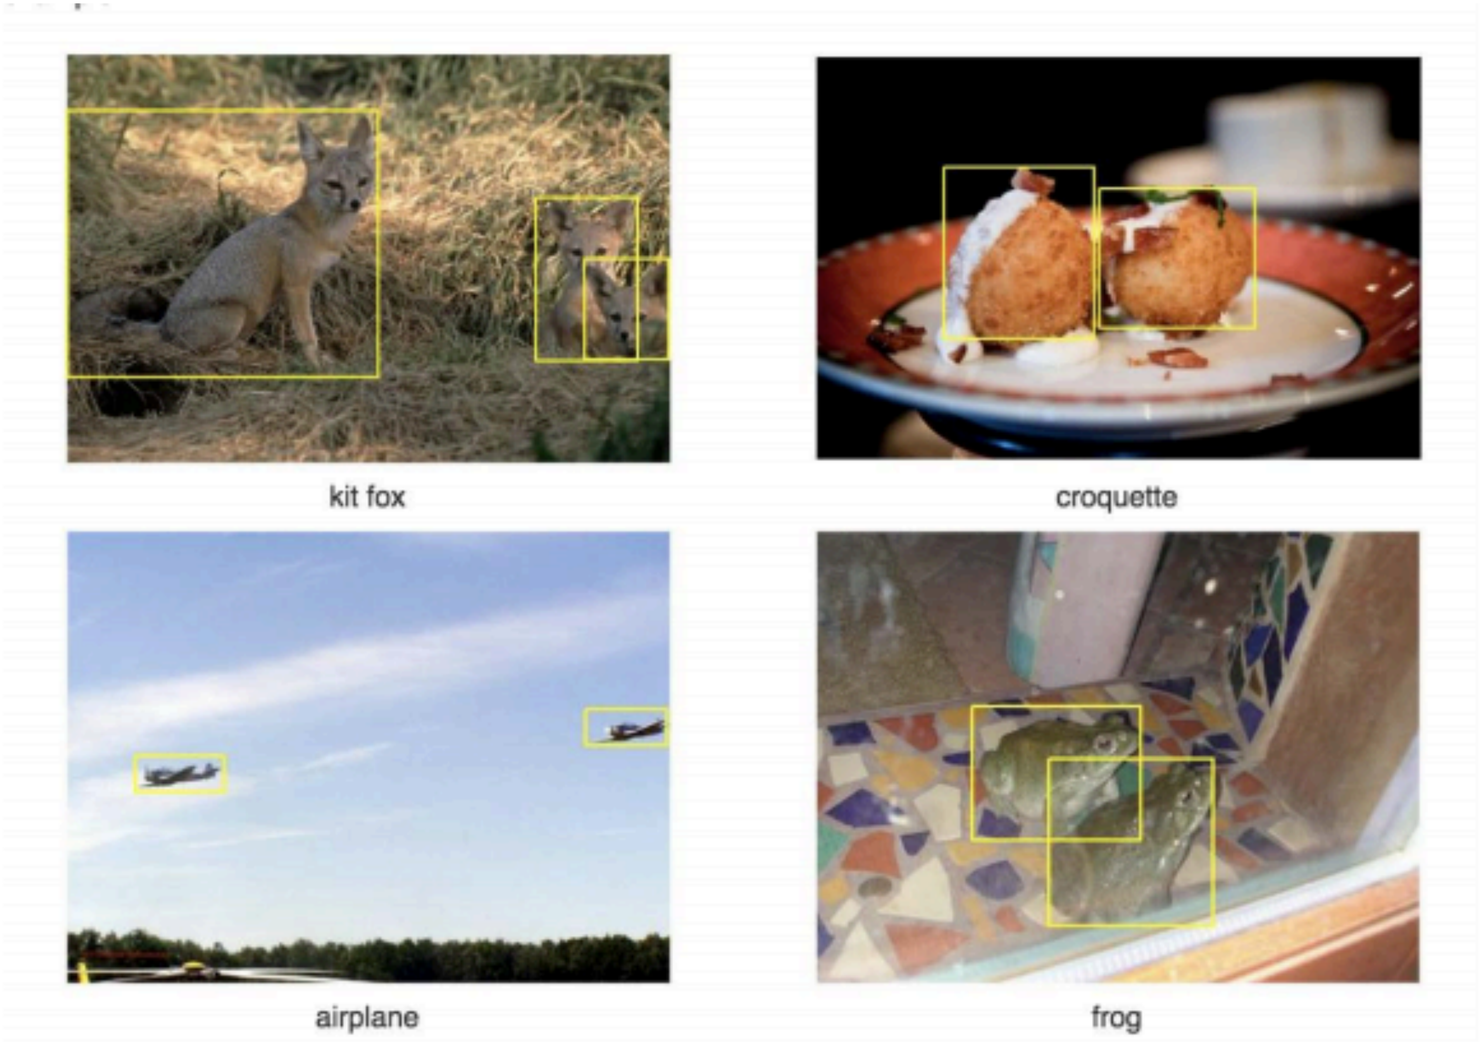

Russakovsky et al. 2015

Two Types of Supervised Learning Algorithms

Classification

- the label is a **discrete** variable

$$y \in \{1,2,3,...,K\}$$

Regression

- the label is a **continuous** variable

$$y \in \mathbb{R}$$

Training Data for Supervised Learning

Training data is a collection of input instances to the learning algorithm:

$$(\mathbf{x}_1, y_2), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_n, y_n)$$

input label

A training data is the "**experience**" given to a learning algorithm

Goal of Supervised Learning

Given training data

$$(\mathbf{x}_1, y_2), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_n, y_n)$$

Learn a function mapping $f : X \rightarrow Y,$ such that $f(\mathbf{x})$ predicts the label y on **future** data $\mathbf{x}$ (not in training data)

Goal of Supervised Learning

Training set error

- 0-1 loss for classification $\ell = \frac{1}{n}\sum_{i=1}^{n}(f(\mathbf{x}_i) \neq y_i)$
- Squared loss for regression: $\ell = \frac{1}{n}\sum_{i=1}^{n}(f(\mathbf{x}_i) - y_i)^2$

A learning algorithm optimizes the training objective

$$f^* = \arg\min \mathbb{E}_{(x,y)}\ell(f(x), y)$$

Details in upcoming lectures :)

Evaluate Machine Learning Model

Use test data (separate from training data)

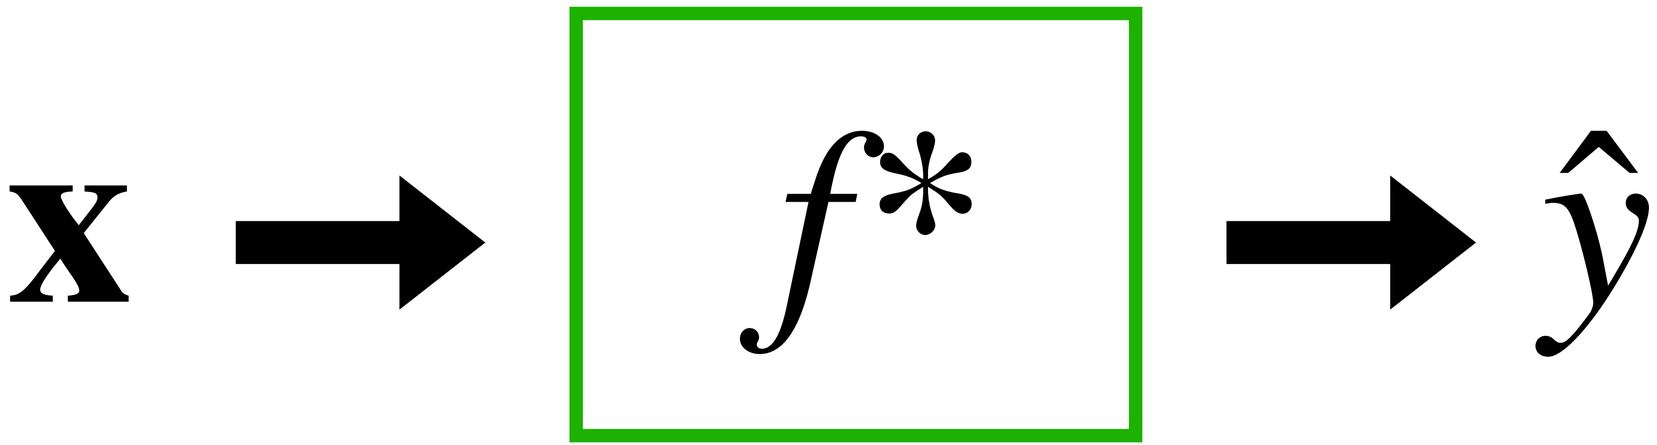

Part III: Unsupervised Learning

Unsupervised Learning

- Given: dataset contains **no label** $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$
- **Goal:** discover interesting patterns and structures in the data

Unsupervised Learning

- Given: dataset contains **no label** $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$
- **Goal:** discover interesting patterns and structures in the data

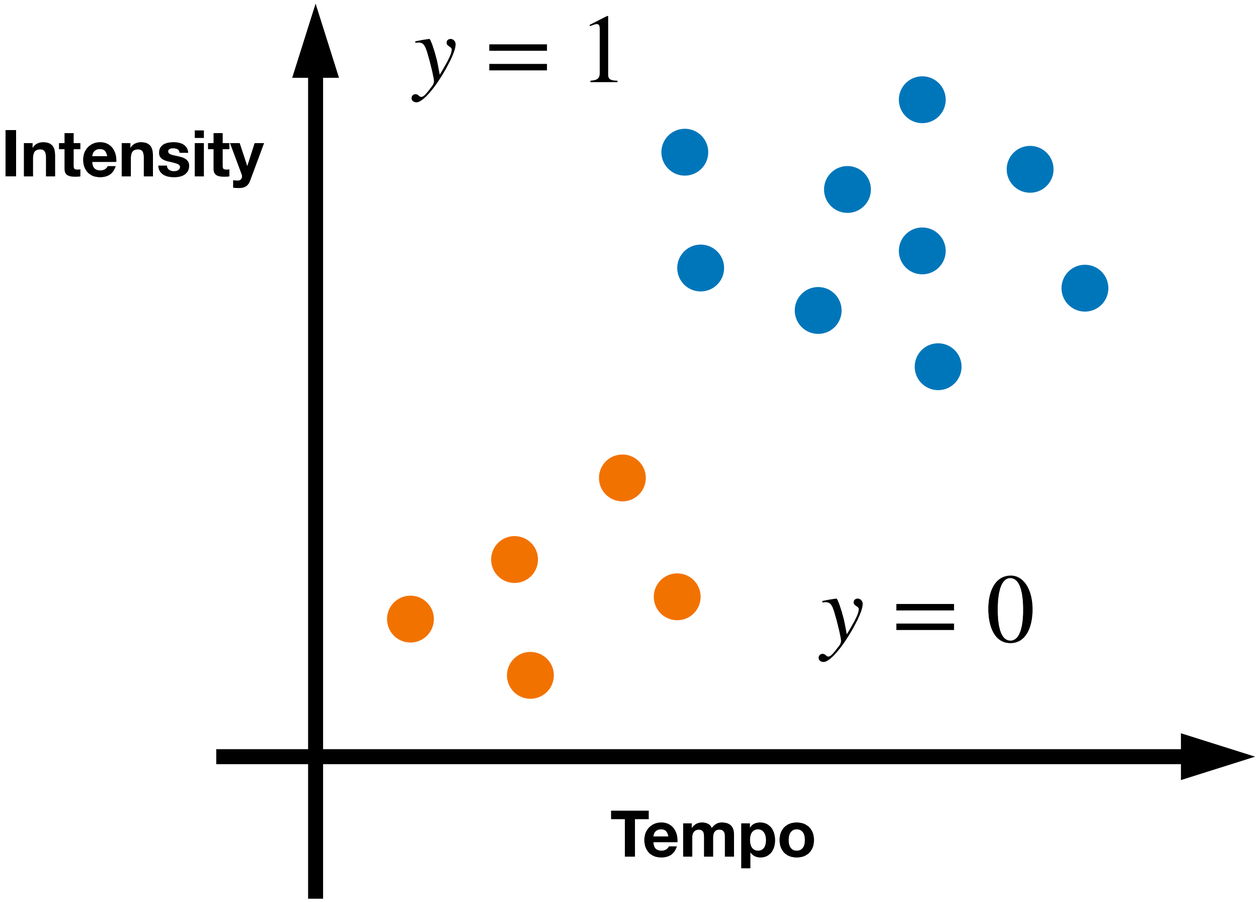

Unsupervised Learning

- Given: dataset contains **no label** $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$
- **Goal:** discover interesting patterns and structures in the data

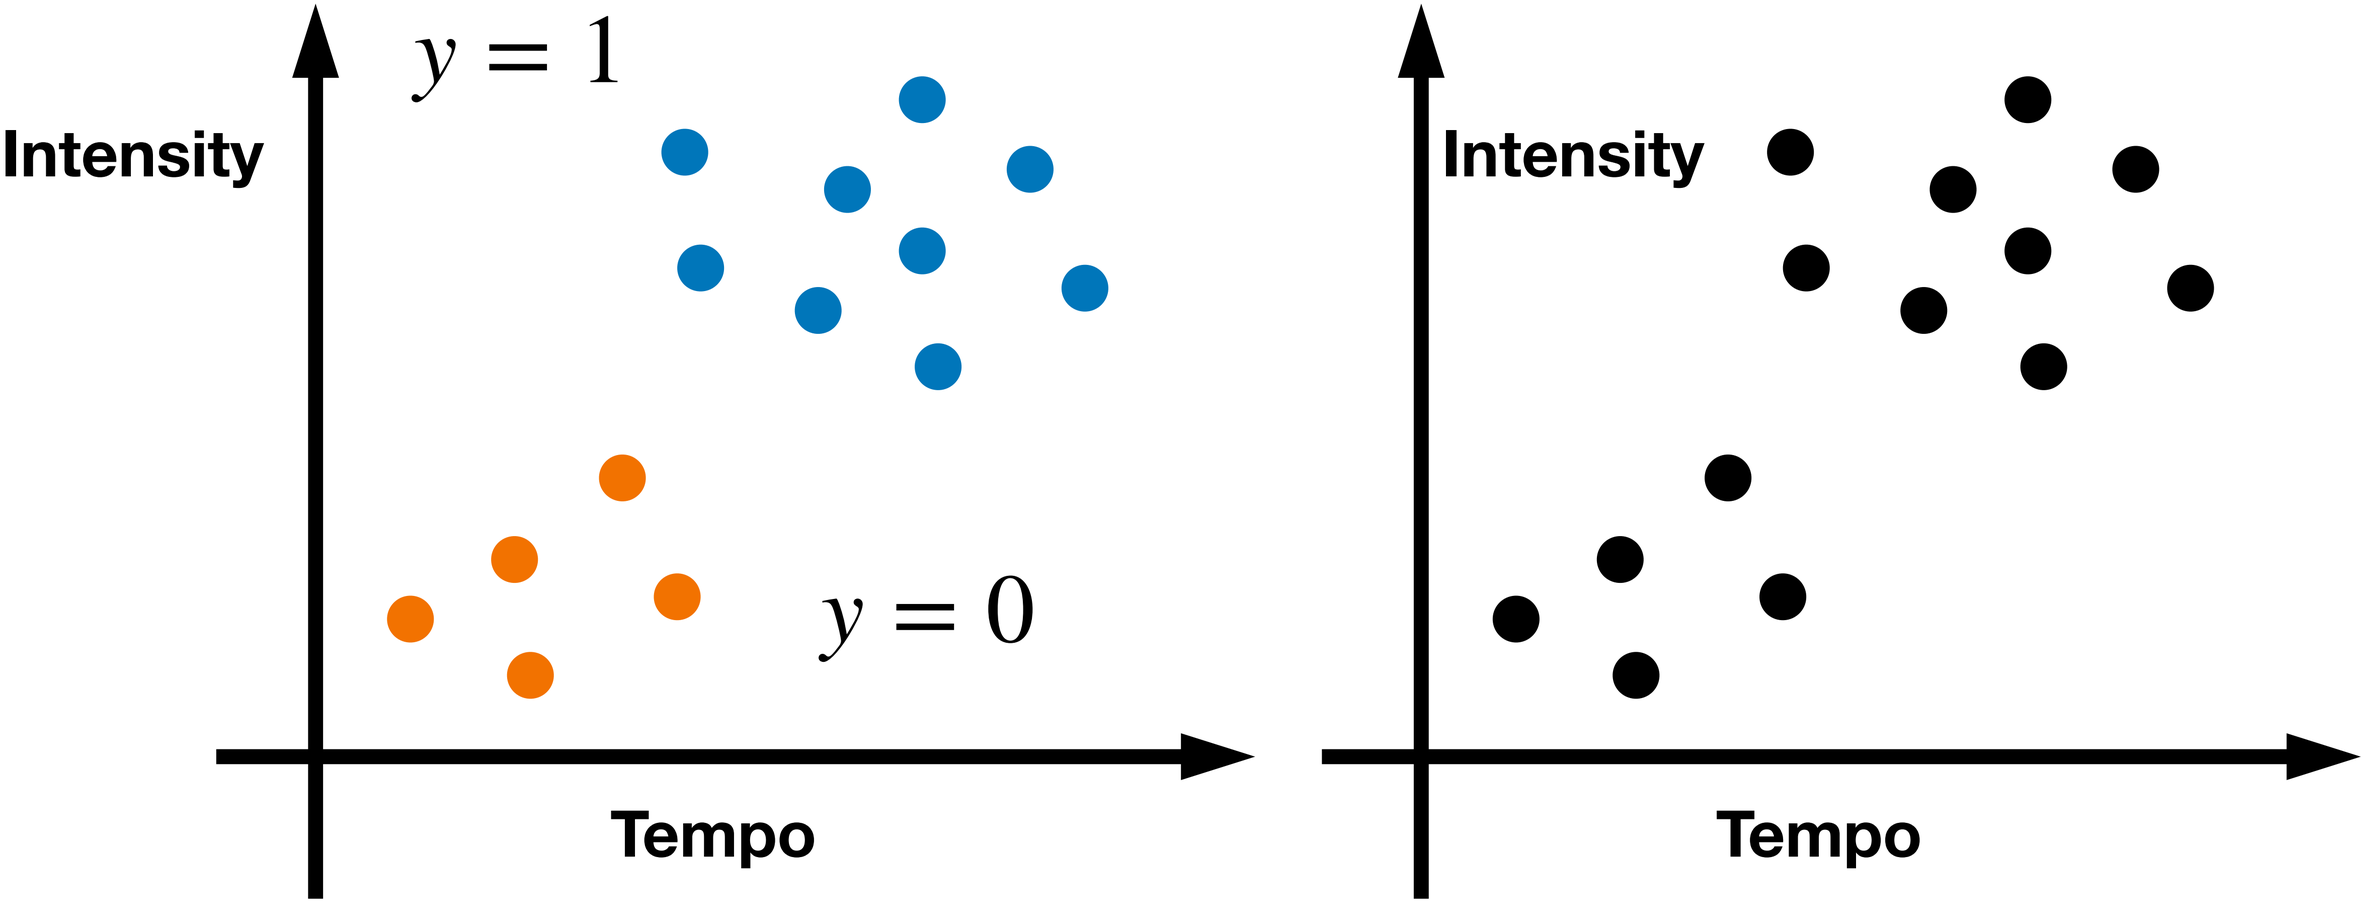

Clustering

- Given: dataset contains **no label** $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$
- **Output:** divides the data into clusters such that there are intra-cluster similarity and inter-cluster dissimilarity

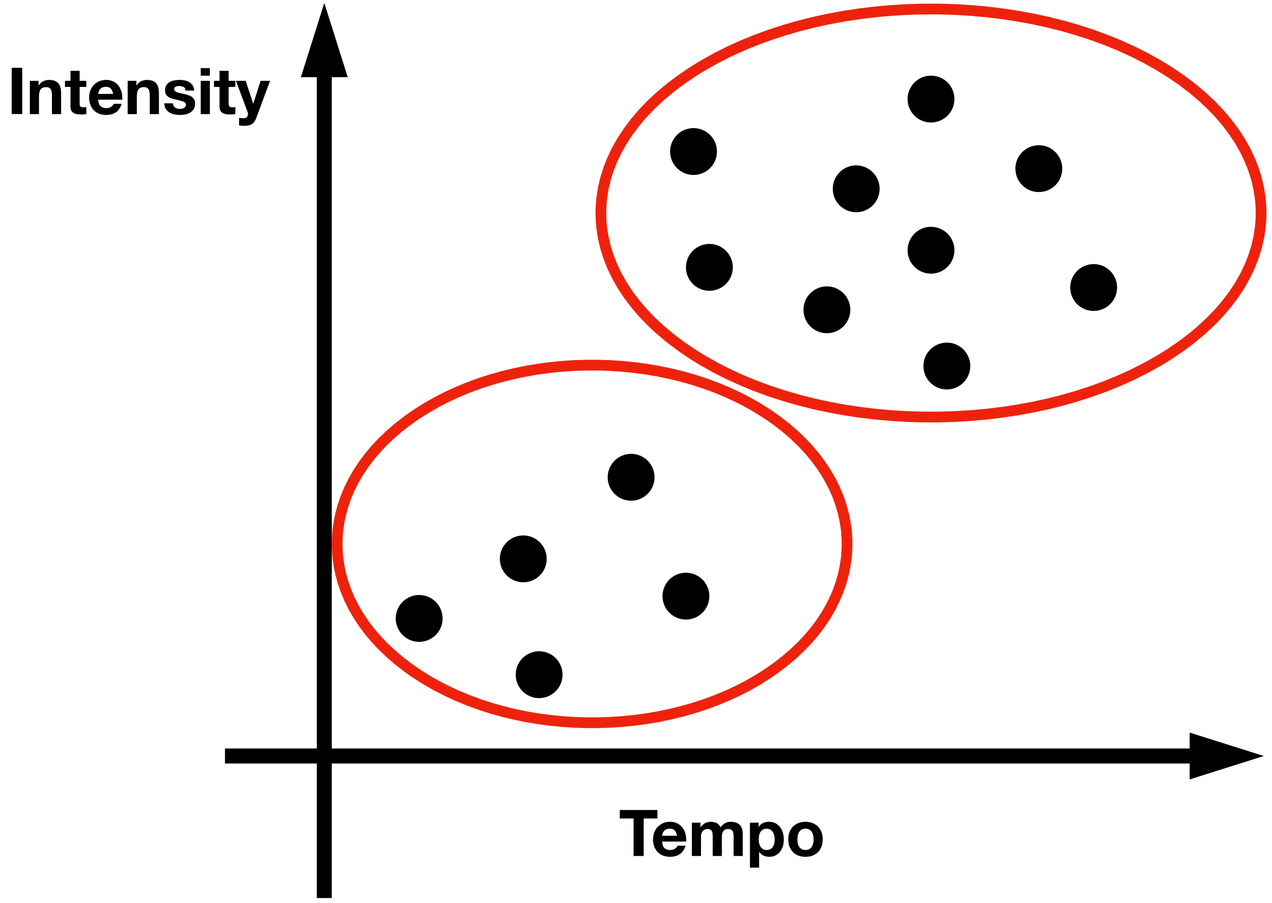

Clustering

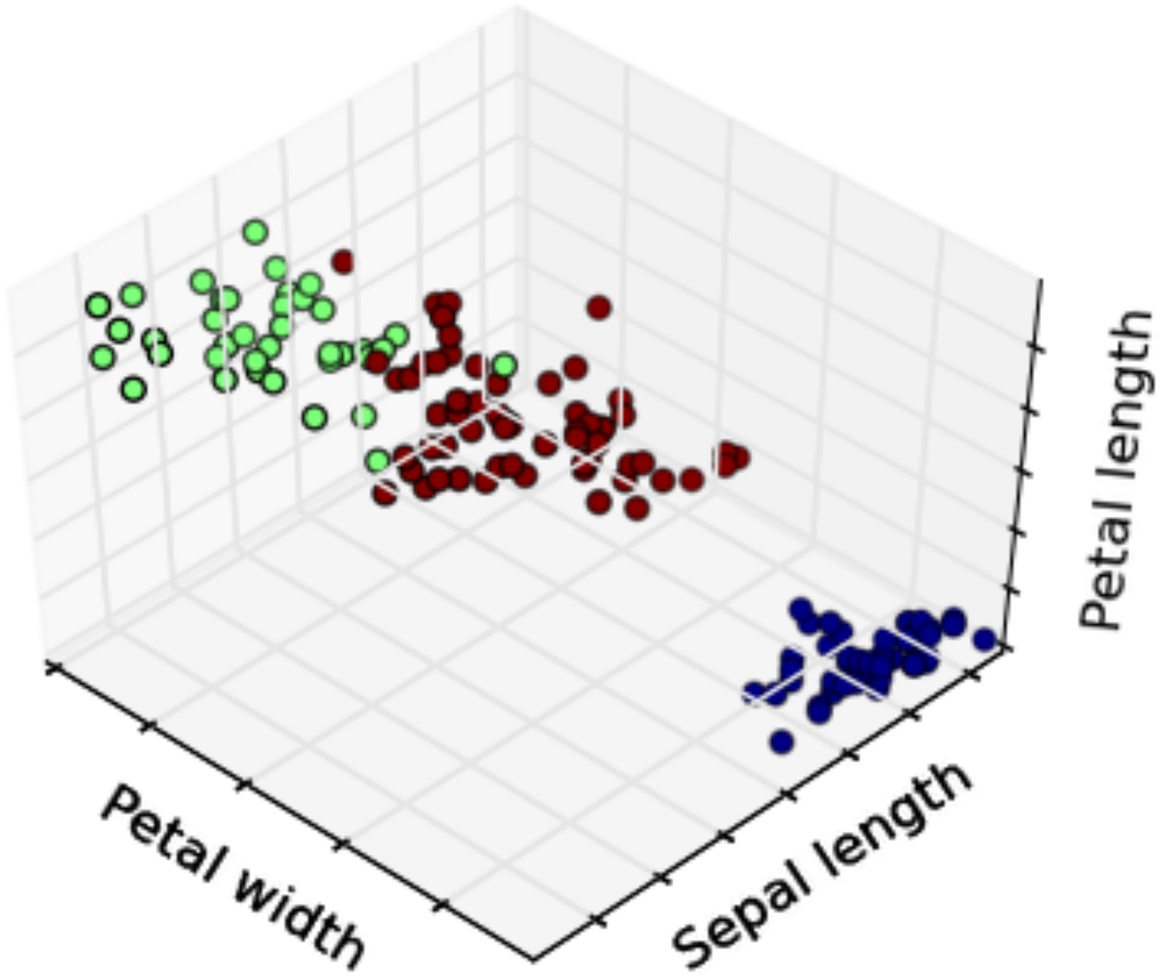

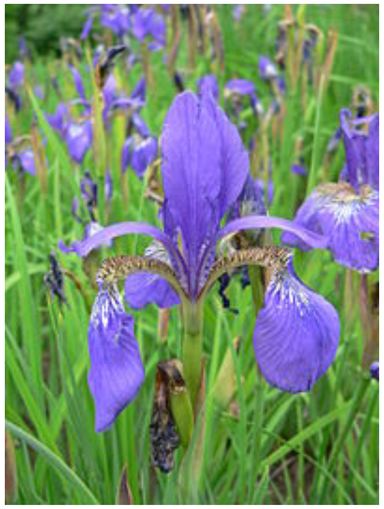

Clustering Irises using three different features

The colors represent clusters identified by the algorithm, not y's provided as input

Clustering

- You probably have >1000 digital photos stored on your phone
- After this class you will be able to organize them better (based on visual similarity)

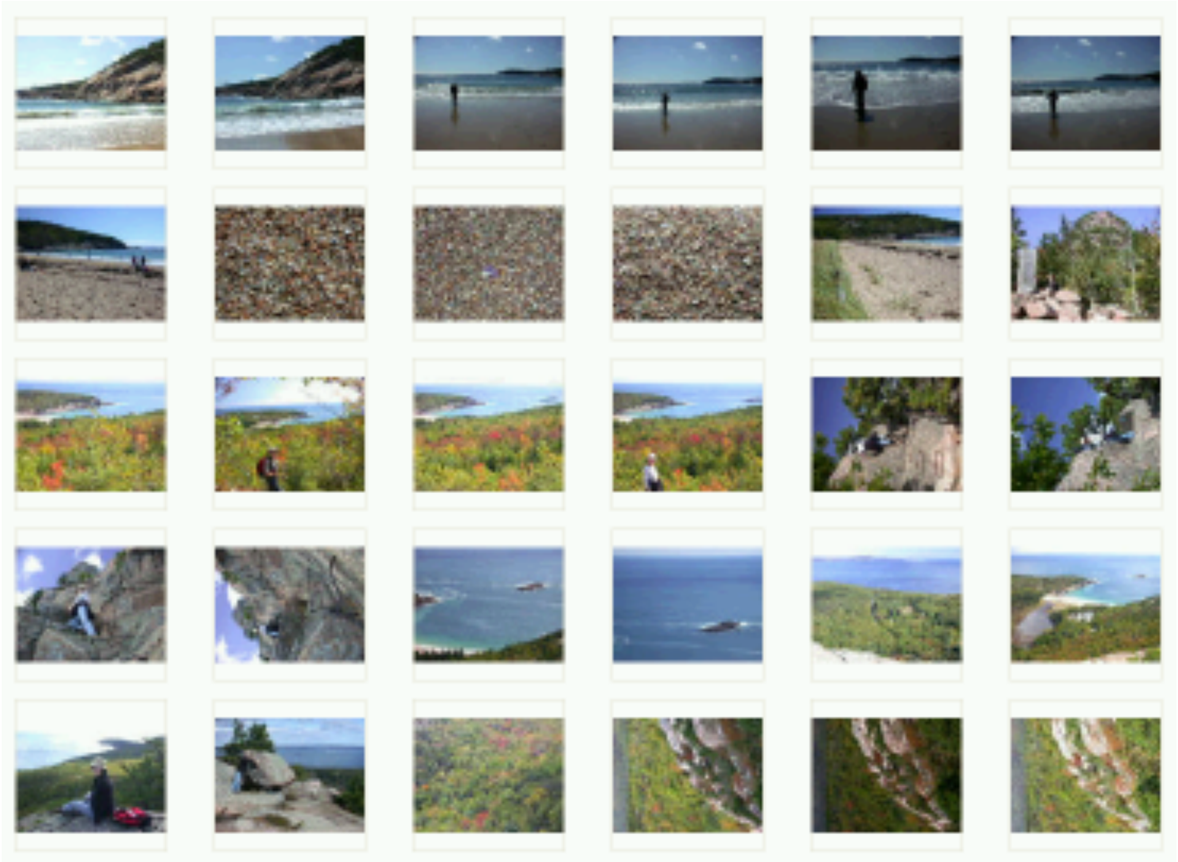

Clustering Genes

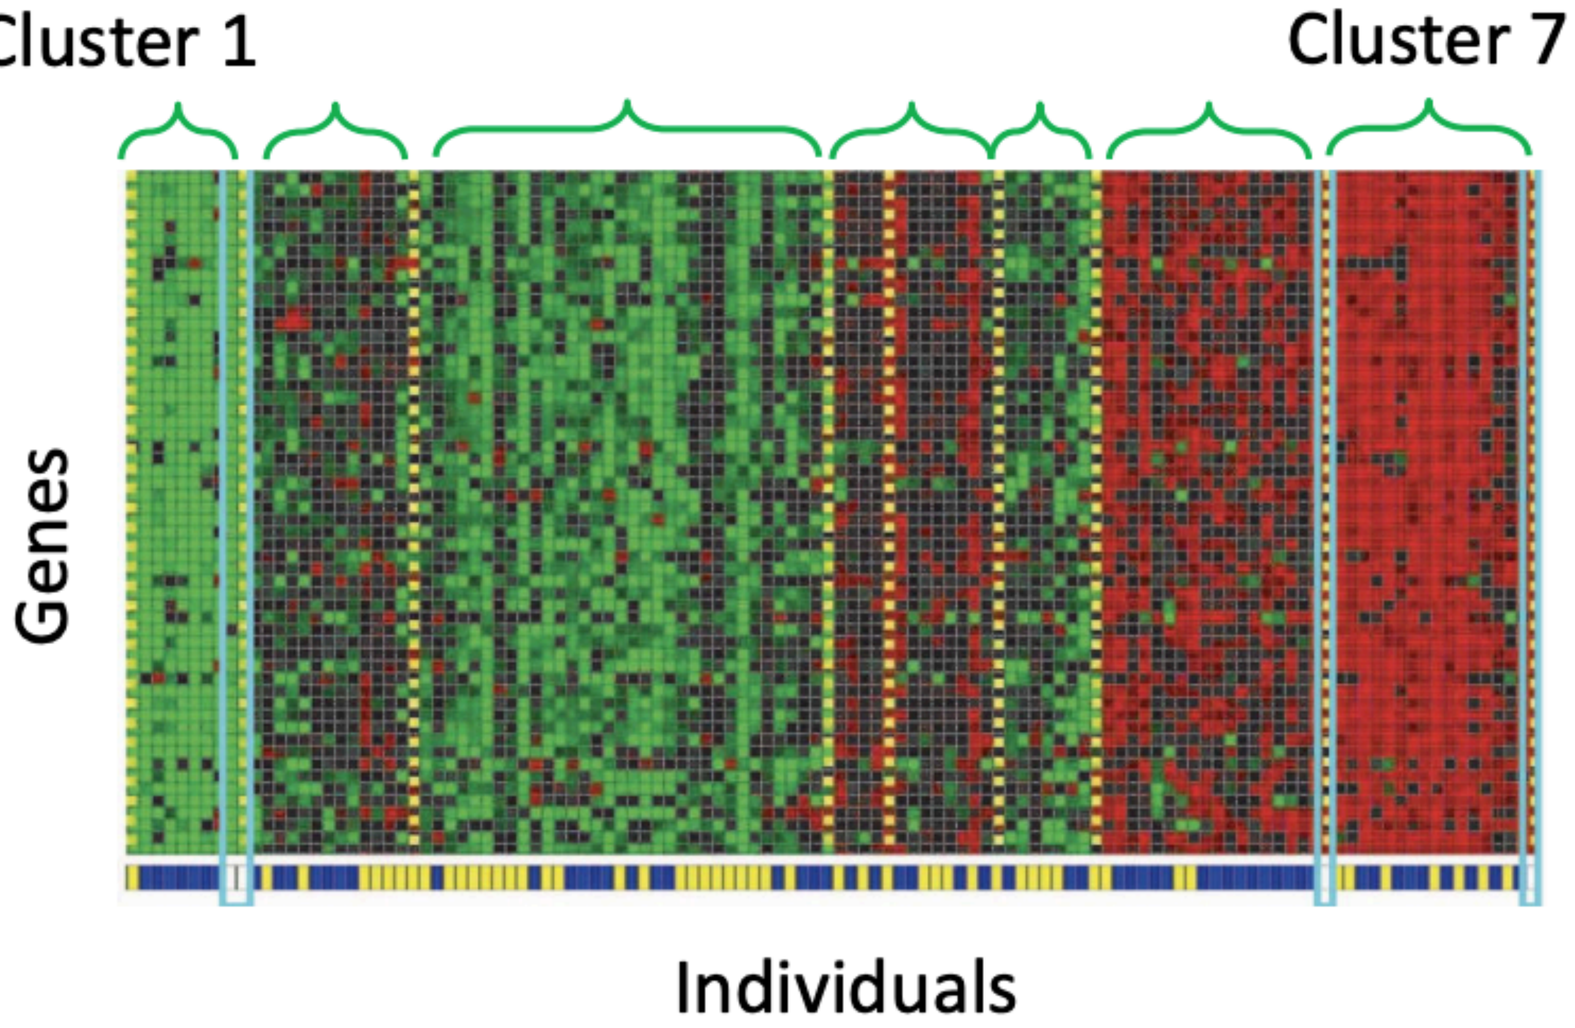

Identifying Regulatory Mechanisms using Individual Variation Reveals Key Role for Chromatin Modification. [Su-In Lee, Dana Pe'er, Aimee M. Dudley, George M. Church and Daphne Koller. '06]

Clustering Words with Similar Meanings

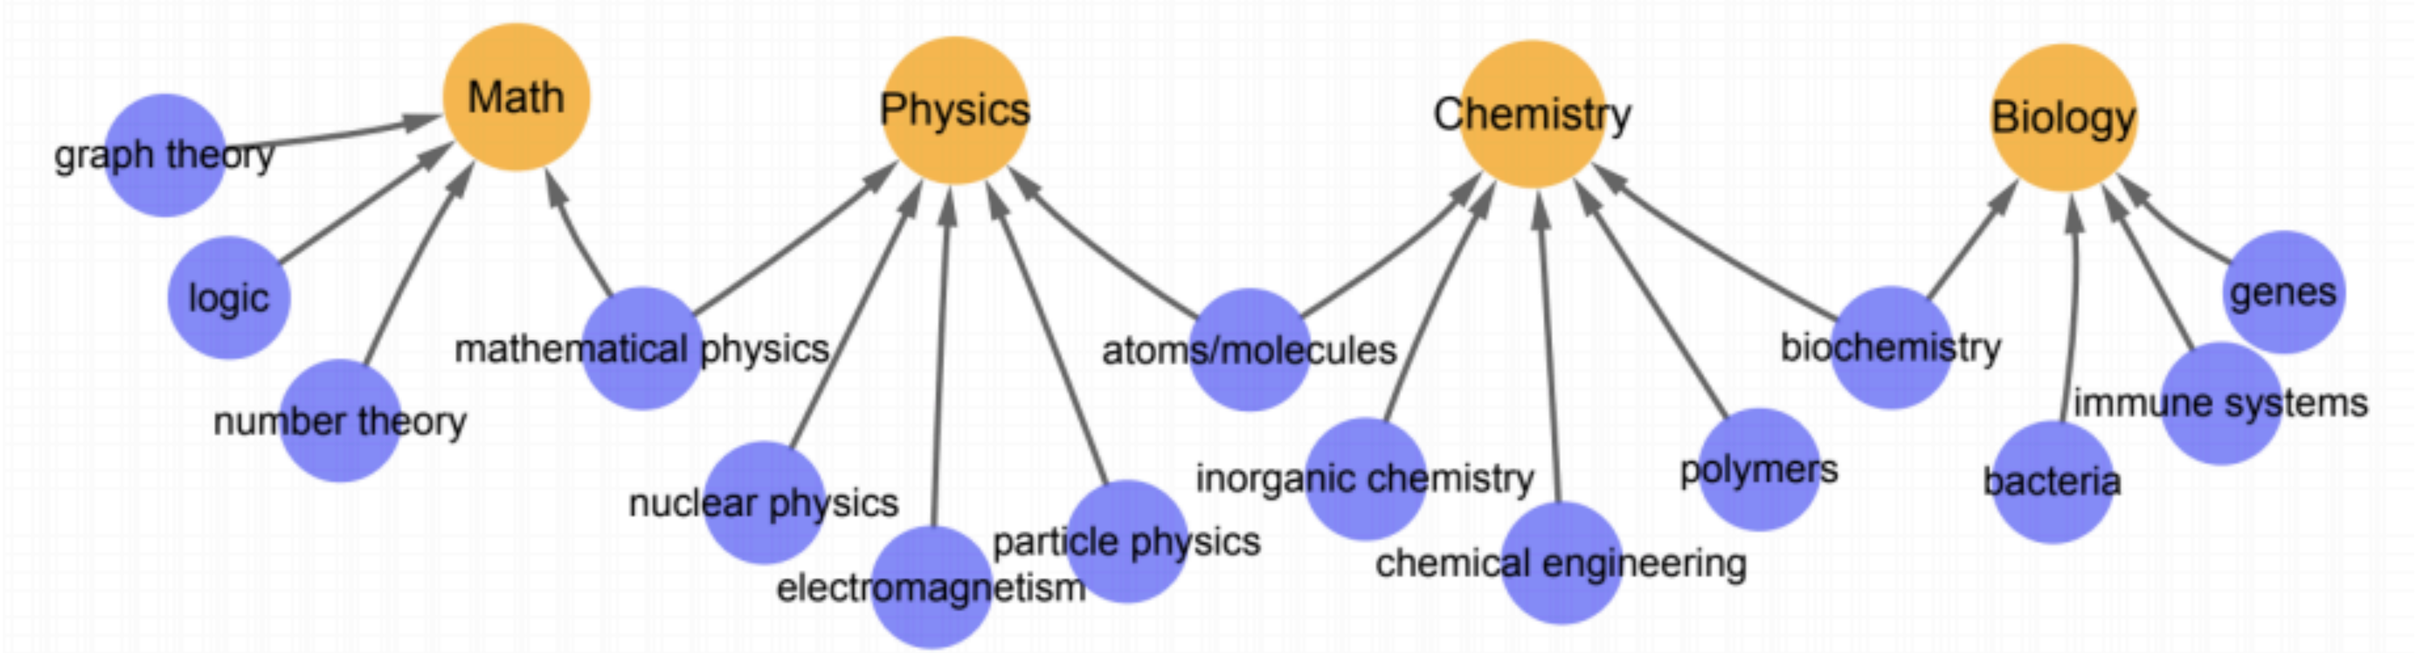

[Arora-Ge-Liang-M.-Risteski, TACL'17,18]

How do we perform clustering?

- Many clustering algorithms. We will look at the two most frequently used ones:
  - *K-means clustering*: we specify the desired number of clusters, and use an iterative algorithm to find them
  - *Hierarchical clustering*: we build a binary tree over the dataset

K-means clustering

- Very popular clustering method

- Don’t confuse it with k-NN classifier

- Input: a dataset $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$, and assume the number of clusters **k** is given

K-means clustering

Step 1: **Randomly** picking 2 positions as initial cluster centers (not necessarily a data point)

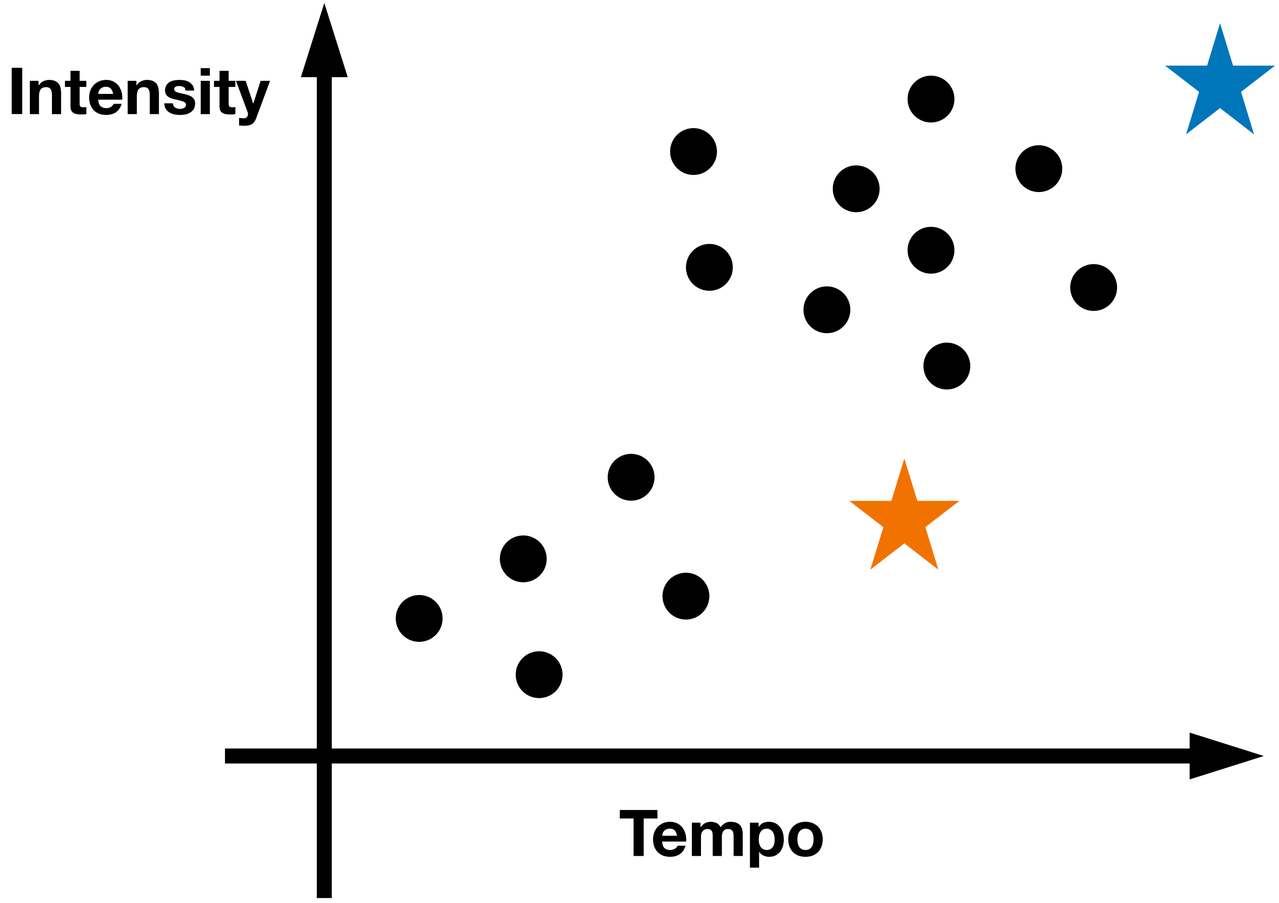

K-means clustering

Step 2: for each point x, determine its cluster: find the closest center in Euclidean space

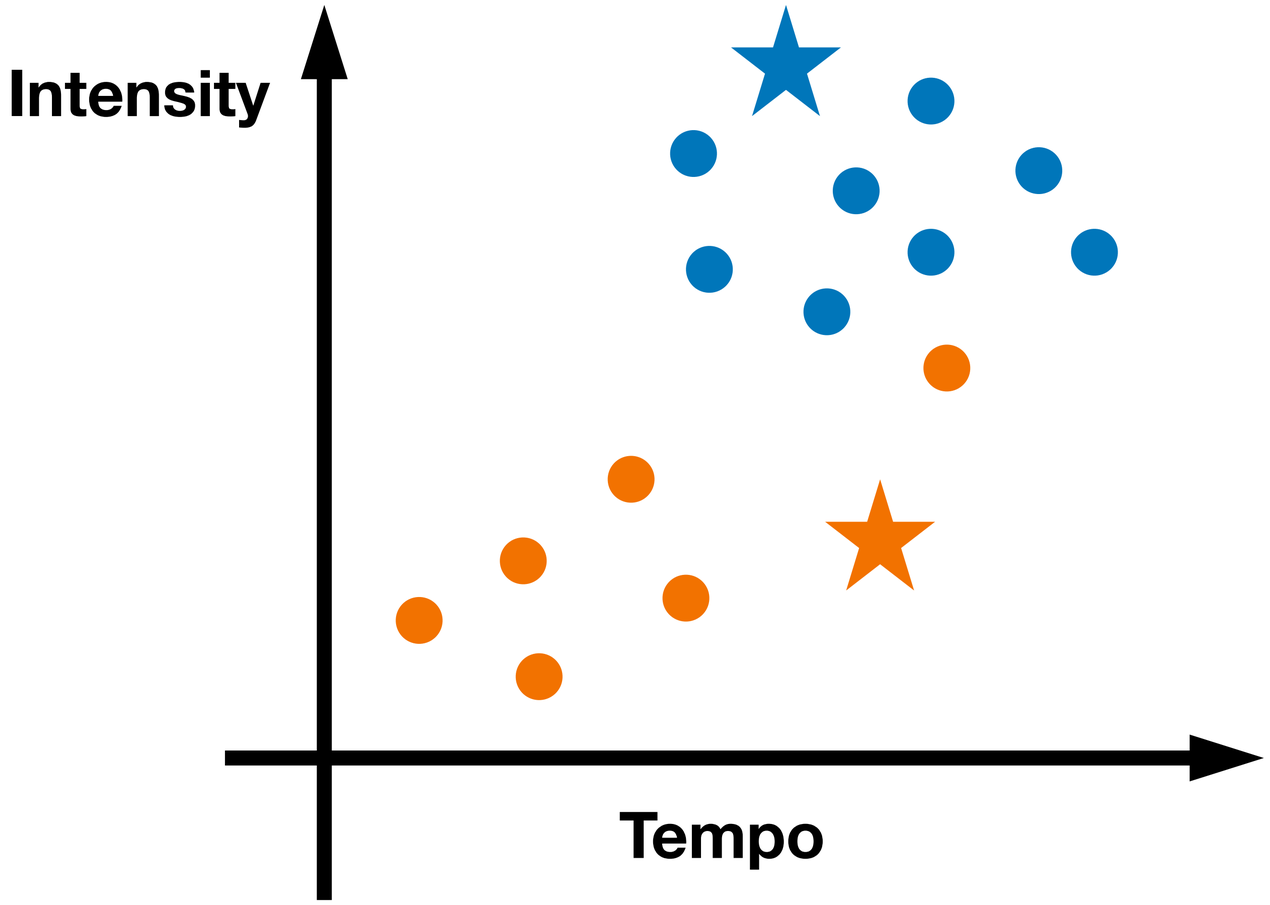

K-means clustering

Step 3: update all cluster centers as the centroids

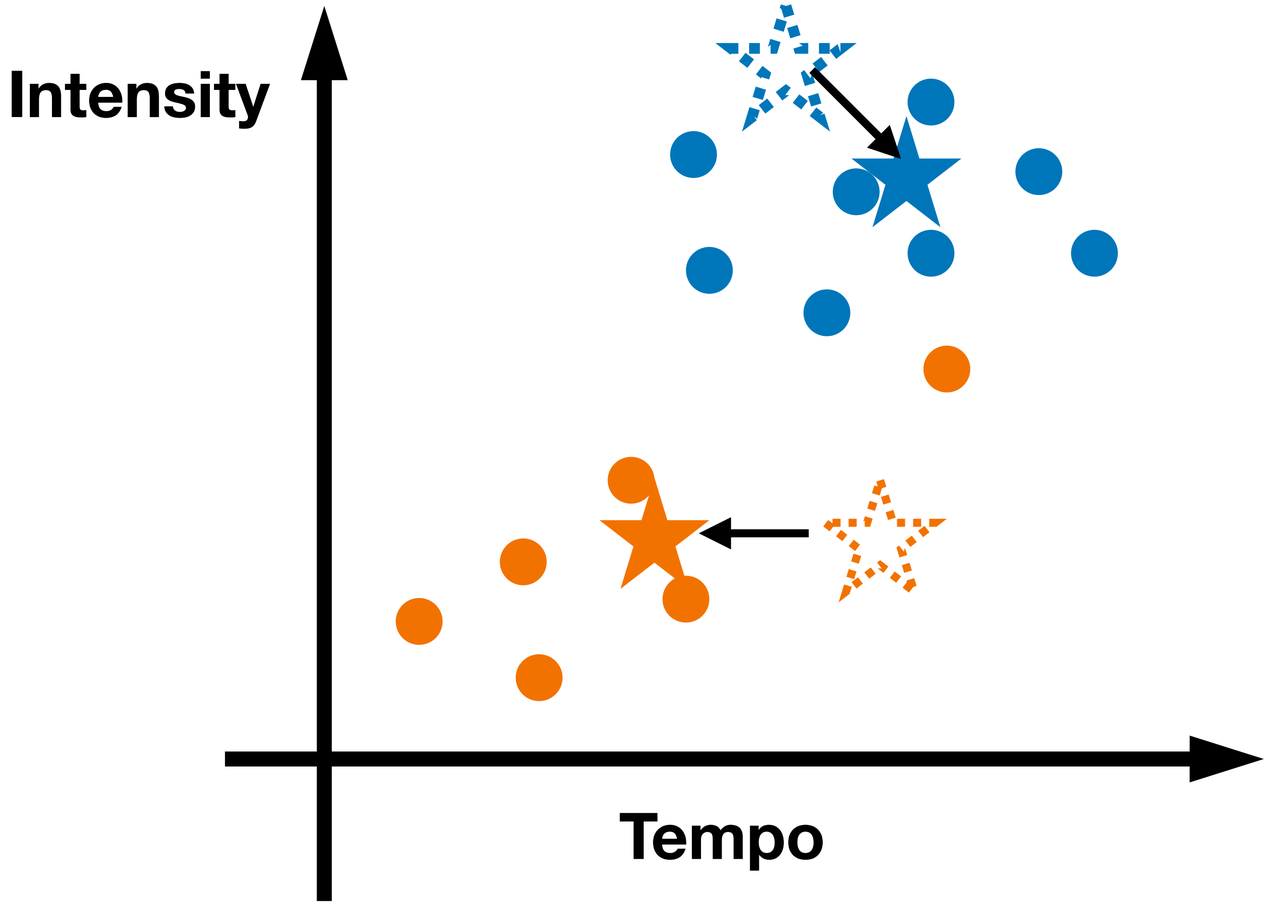

K-means clustering

Repeat step 2 & 3 until convergence

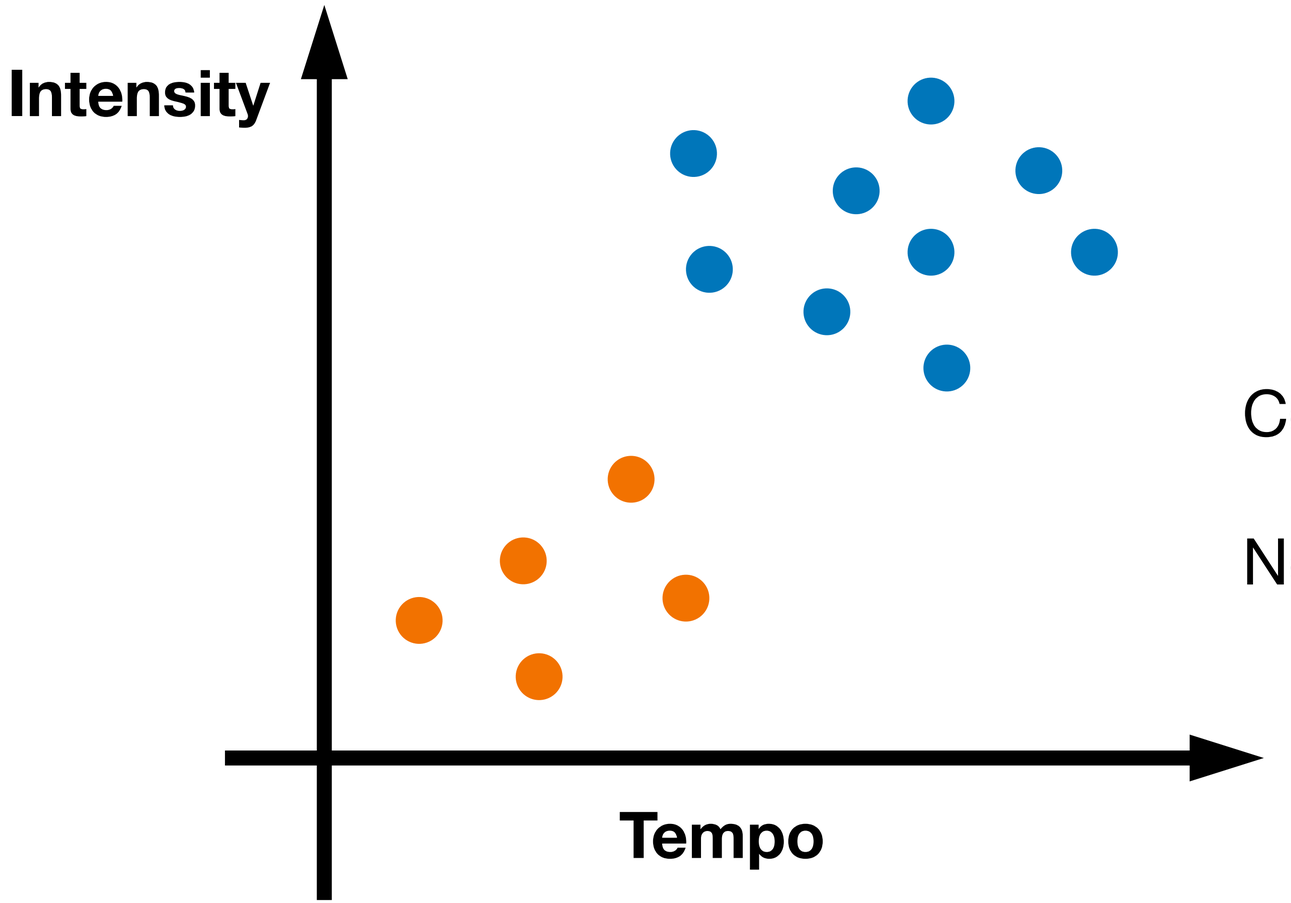

Converged solution!

No labels required!

K-means clustering: A demo

https://www.naftaliharris.com/blog/visualizing-k-means-clustering/

Hierarchical Clustering (more to follow next lecture)

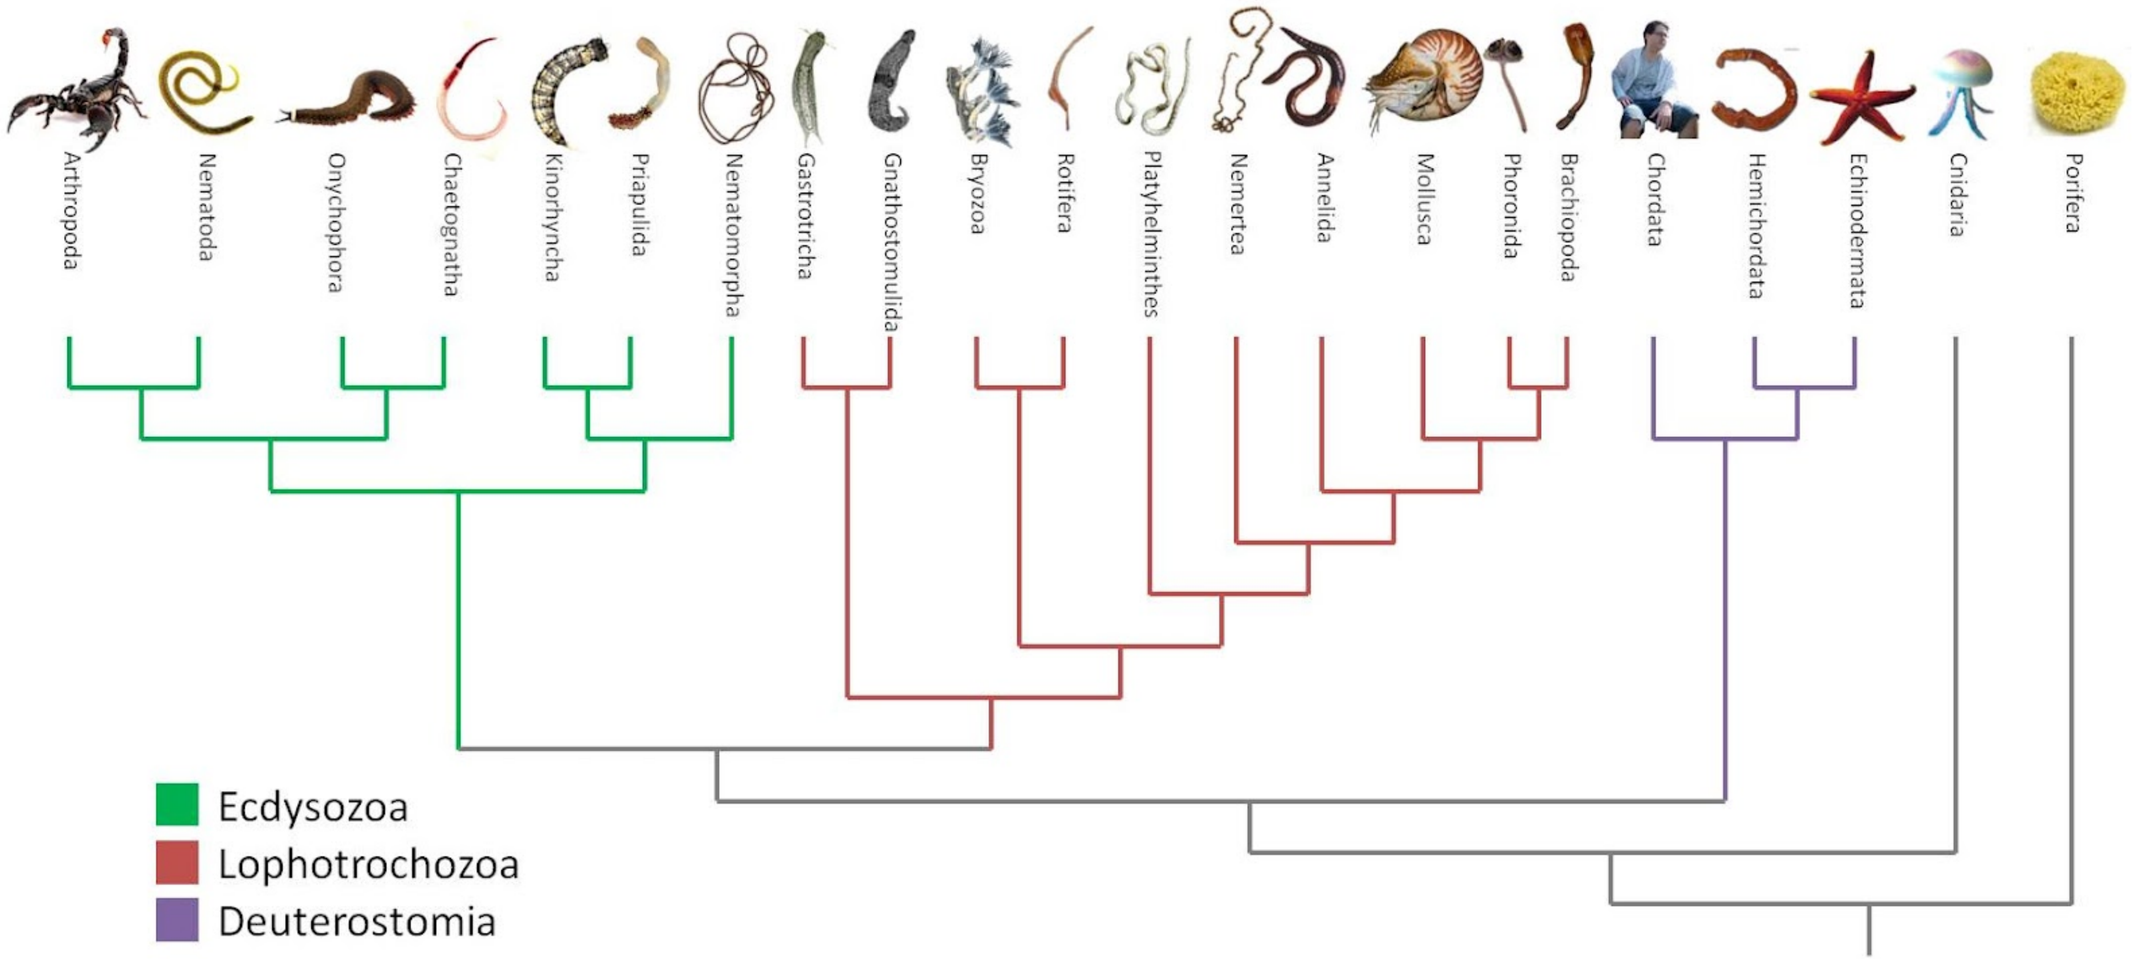

Today’s recap

- What is machine learning?
- Supervised Learning
  - Classification
  - Regression
- Unsupervised Learning
  - Clustering (Kmeans and hierarchical clustering)

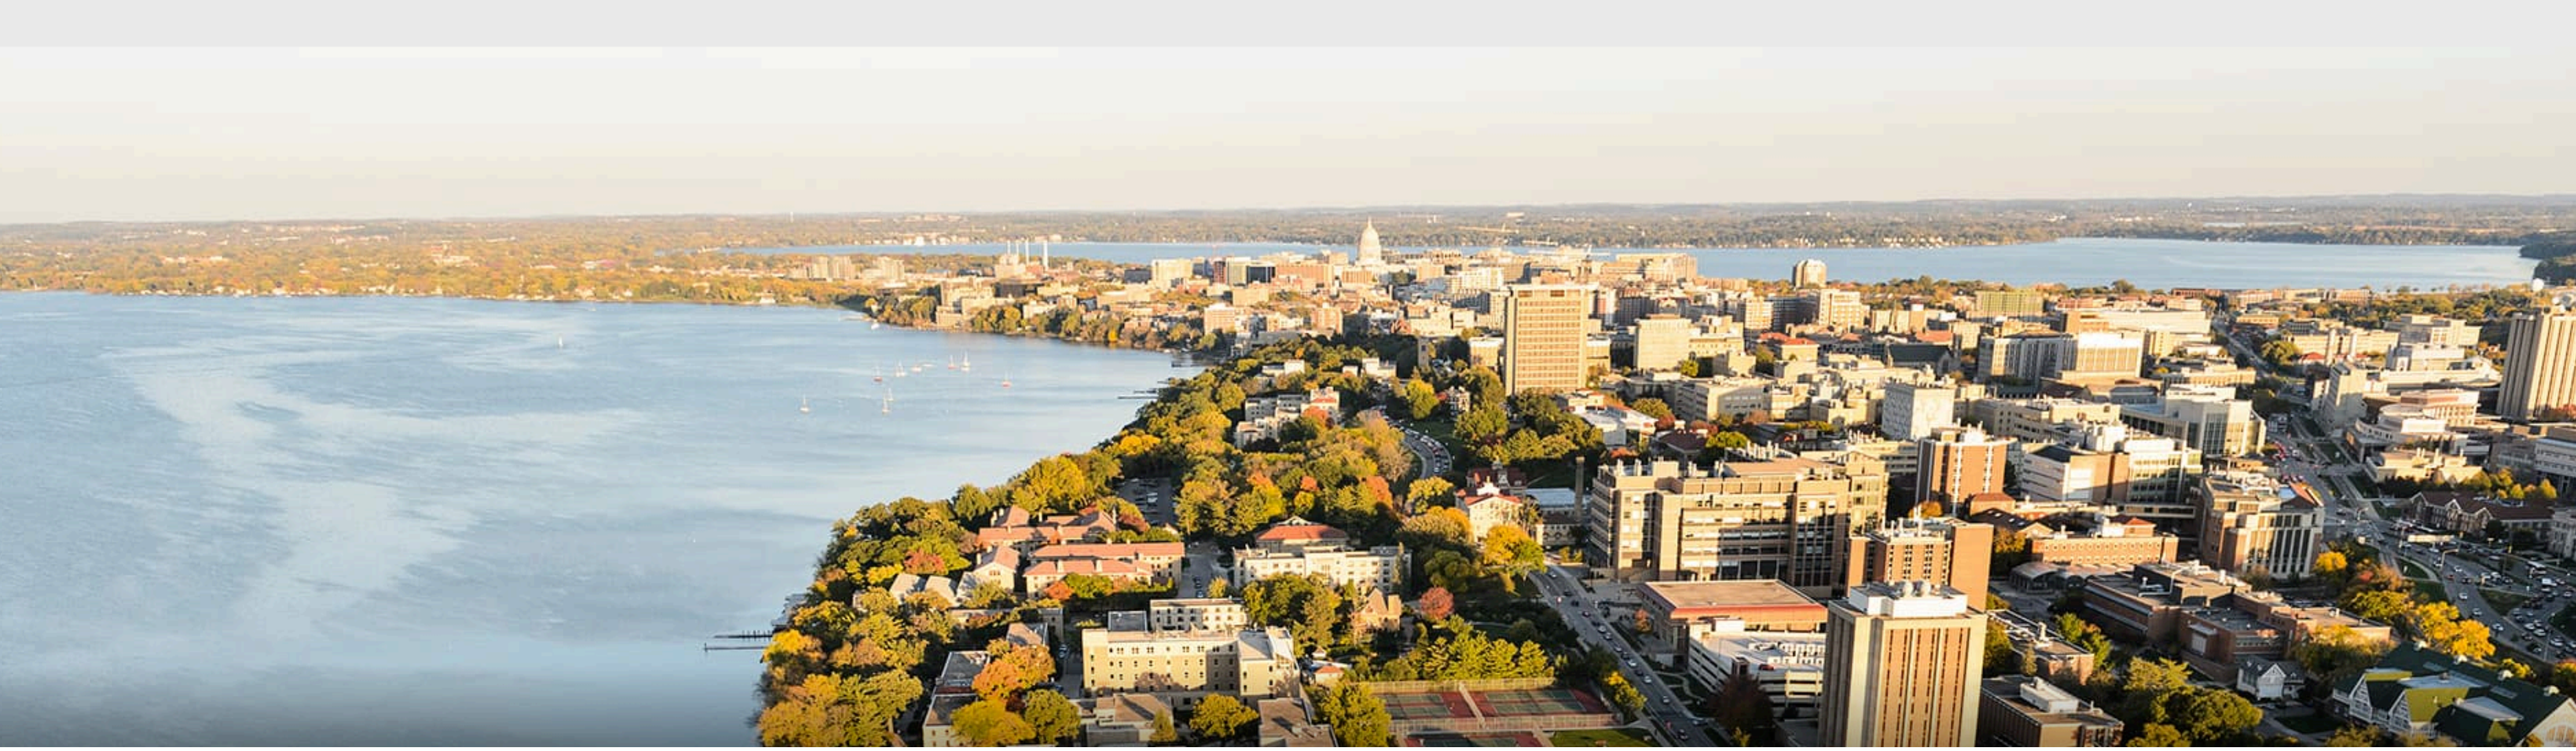

Thanks!