Lecture 13

15.12.2025

Today's plan and announcements

- Hour 1
  - Roberto Castello from Swiss Data Science Center on applications of AI/ML in Swiss industries
- Hour 2
  - Course review
- Exercise hour
  - Problem set 6
- Reminder: python homework due December 17

Introduction

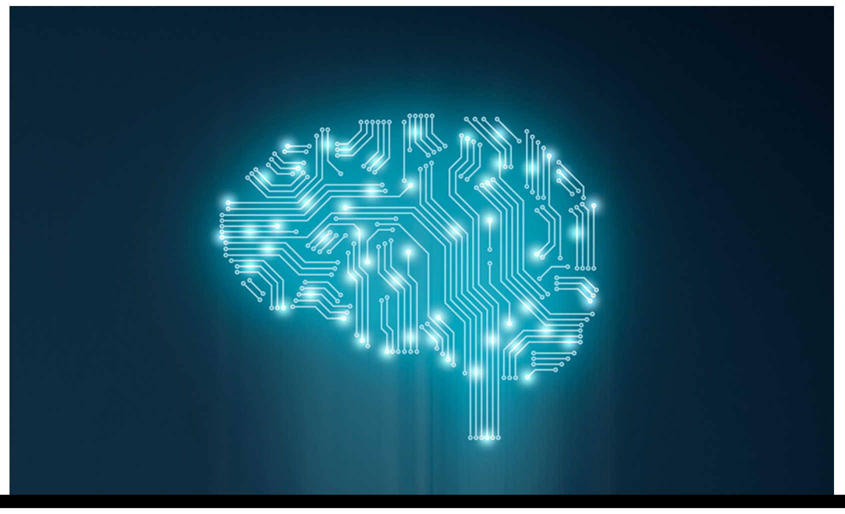

Linear regression

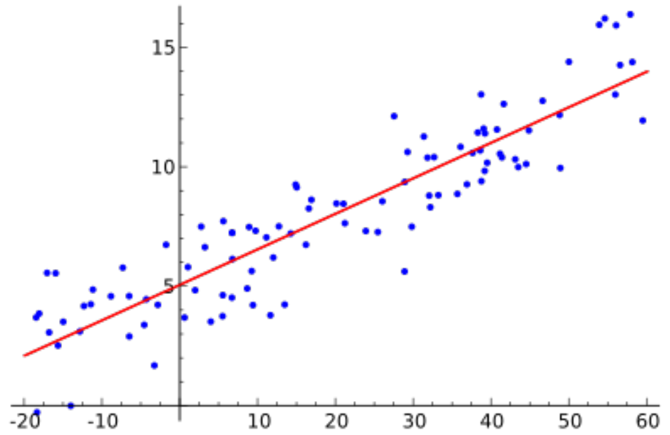

Logistic regression

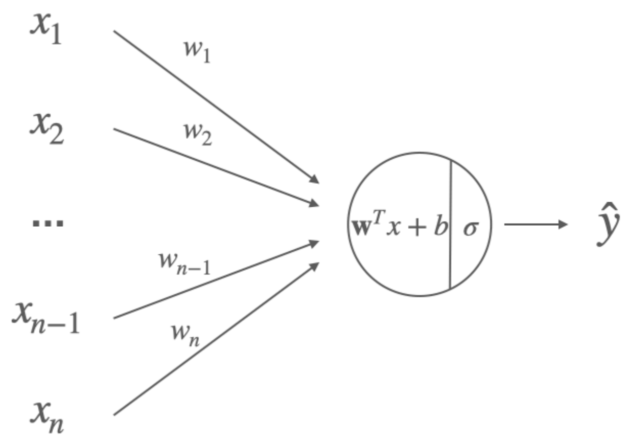

Neural networks

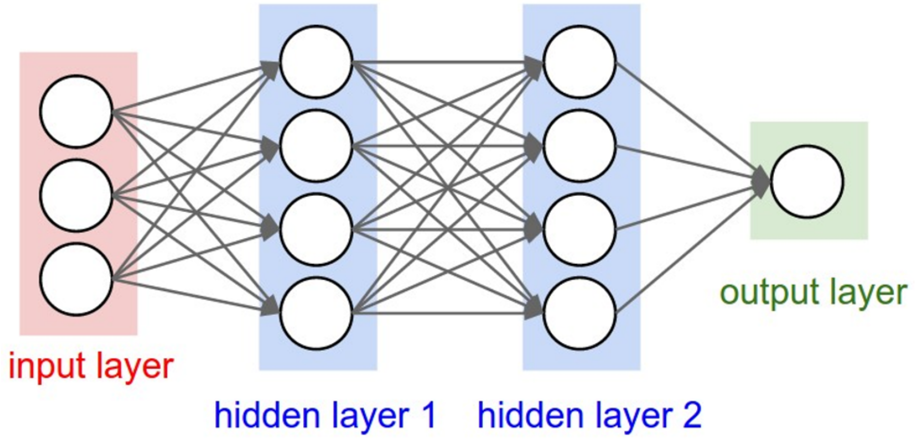

Convolutional neural networks

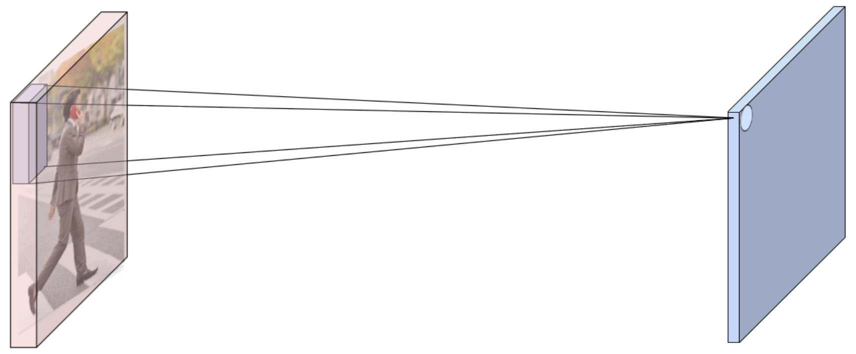

AI & sustainability

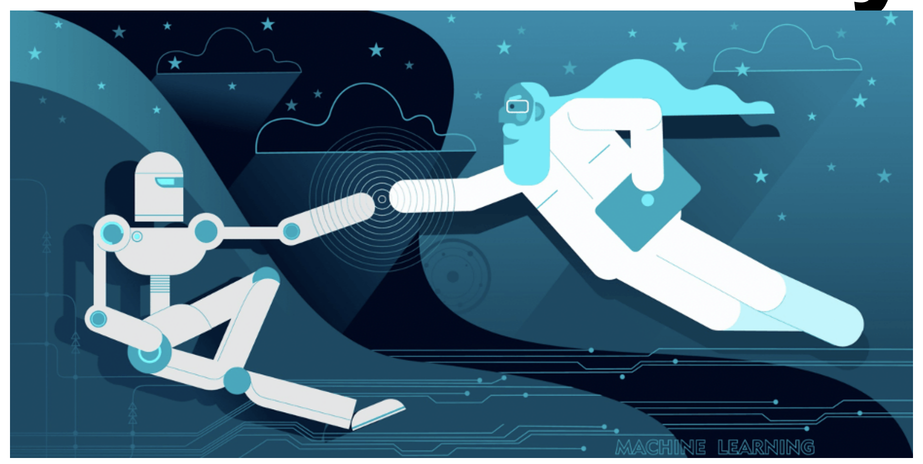

Reinforcement learning

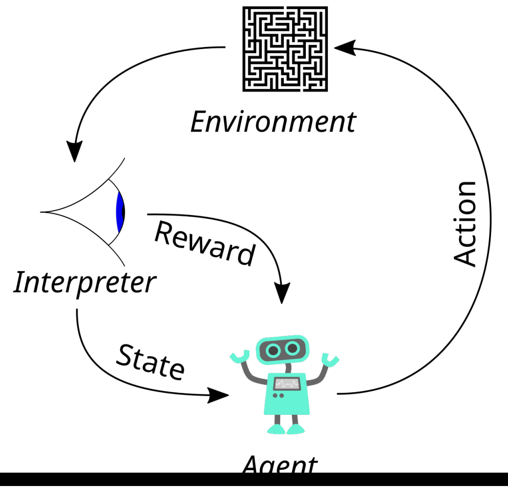

Recurrent neural networks

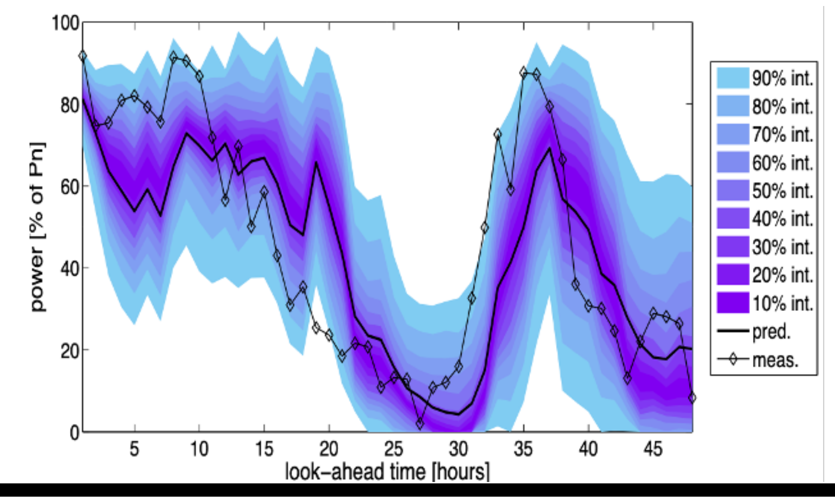

Naive Bayes

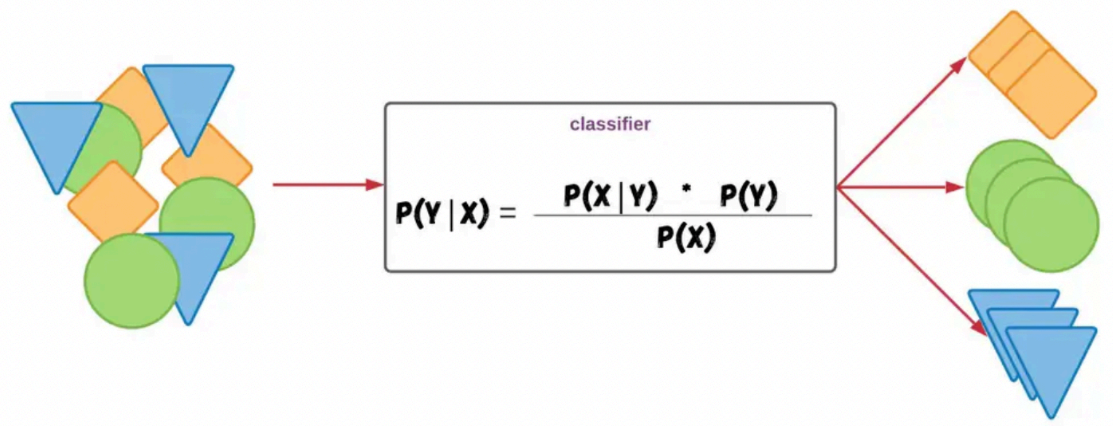

KNN

Dimensionality reduction

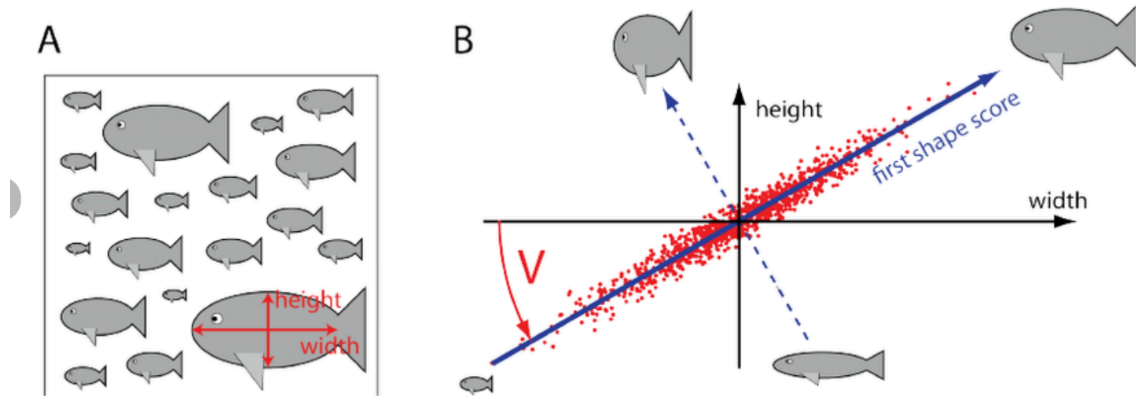

Clustering

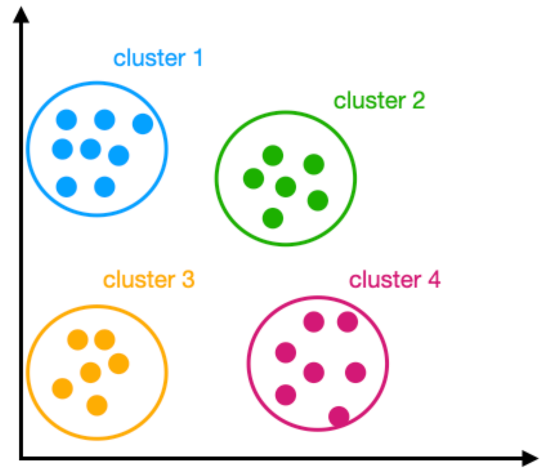

Review: K-means

- Cluster data into K groups, by trying to minimize the objective , data $\{x^i\}_{i=1}^N$

$$L(\mu^1, \ldots, \mu^K, c^1, \ldots, c^N) = \frac{1}{N} \sum_{i=1}^{N} \|x^i - \mu^{c^i}\|_2^2$$

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

- Decision variables: Cluster means $\{\mu^1, \ldots, \mu^K\} \subset \mathbb{R}^d$, cluster each data point $x^i \in \mathbb{R}^d$ belongs to $c^i \in \{1,2,\ldots,K\}$, $i = 1, 2, \ldots, N$
- Non-convex optimization: generally, we don't find the global optimum
  - K-means algorithm is an iterative heuristic that converges to a local minimum

Choosing number of clusters

- Run the K-means algorithm for $K = 2,3,\ldots,\bar{K}$
- Evaluate the objective function upon algorithm convergence, WSS: within cluster sum-of-squares)

$$L(\mu^1, \ldots, \mu^K, c^1, \ldots, c^N) = \frac{1}{N}\sum_{i=1}^{N} \|x^i - \mu^{c^i}\|_2^2$$

- Plot the WSS
- Choose K when marginal improvement

  $\Delta(K) = L(K) - L(K+1)$ drops significantly

  → written as a function of # of cluster

- Note: this is a heuristic. Some curves may not even have such an elbow point.

Elbow method

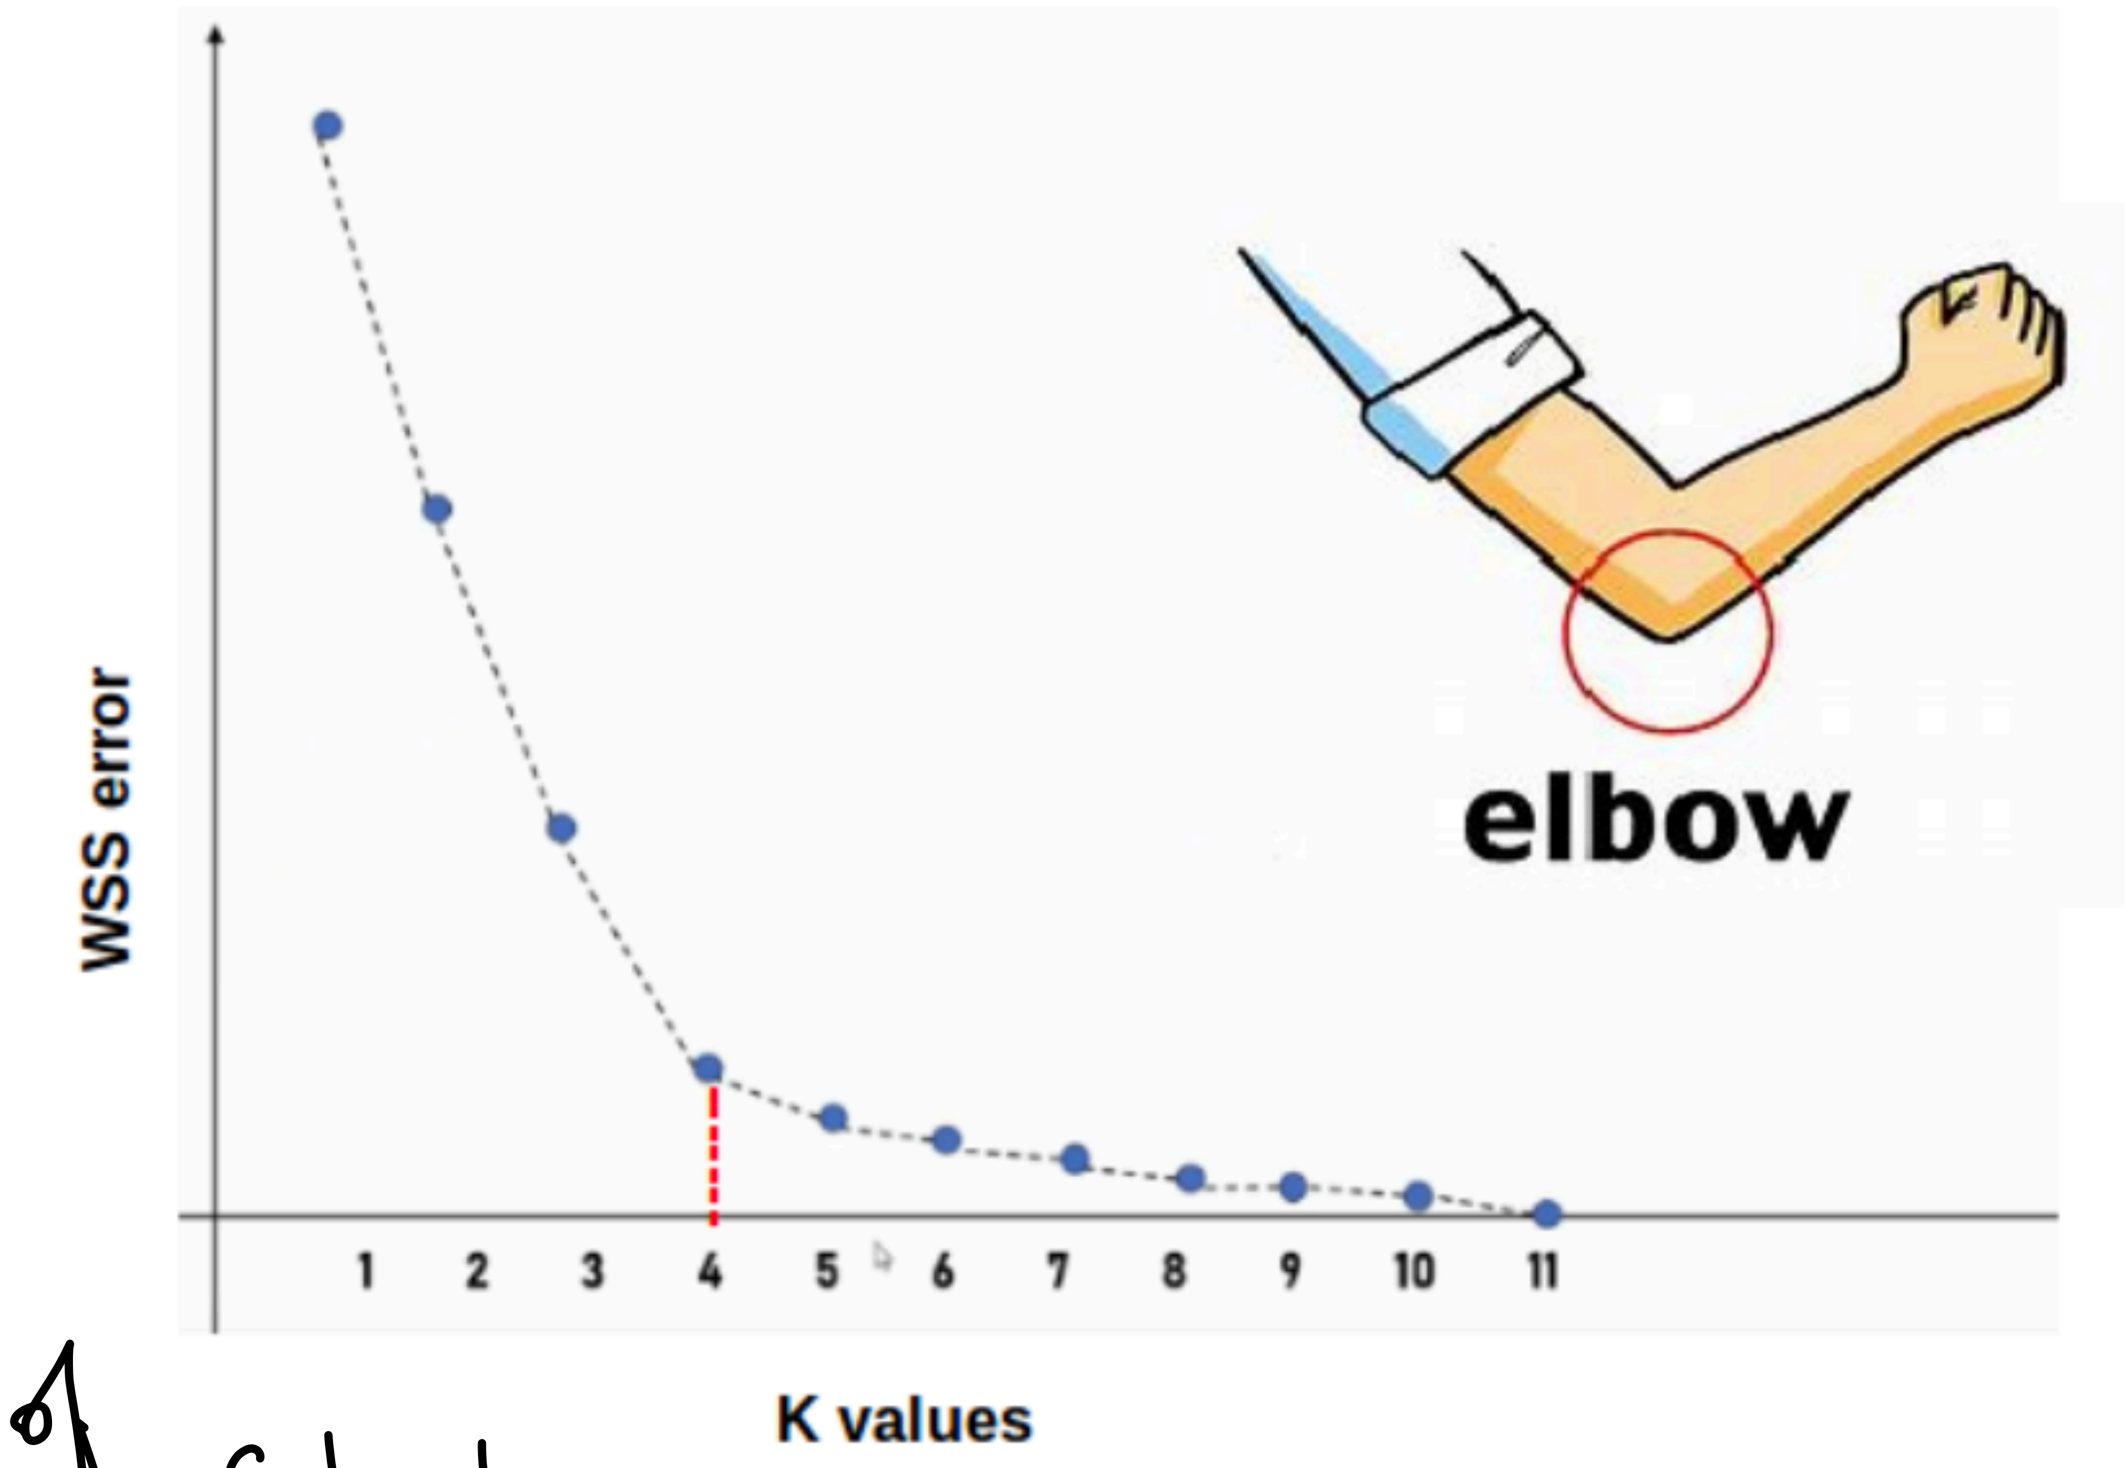

Review of the course

Tasks we learned

- Regression , $\{x^i, y^i\}_{i=1}^N$, $y^i \in \mathbb{R}^m$
- Classification $\{x^i, y^i\}_{i=1}^N$, $y^i \in \{1, \dots, K\}$

supervised learned

- Reinforcement learning : find a control policy $\pi : S \to A$ that optimizes a reward for a dynamical system through interacting with a simulator, or the real system
- Time-series prediction :

$\{x_t\}_{t=0}^T$ want to predict $x_{T+1}, \dots x_{T+\tau}$

- Dimensionality reduction
- Clustering

unsupervised learning

$\{x^i\}_{i=1}^N$, $x^i \in \mathbb{R}^d$

Approaches we learned for each task

- Regression : linear regression, neural network, kNN
- Classification : logistic regression, neural network, kNN, Naive Bayes

- Reinforcement learning : policy gradient

- Time-series prediction : recurrent neural networks

- Dimensionality reduction : PCA, briefly saw autoencoder
- Clustering : k-means & kernelized k-means

Pillar 1 - Linear algebra

- Fundamental: data is represented as vectors, matrices
- Linearity also makes ML scalable (NN, CNN, RNN)
- Linear maps (operators)
  - Linear/logistic regression (even with nonlinear features) $w_1 \phi_1(x) + w_2 \phi_2(x) + \dots + w_p \phi_p(x)$
  - Neural network operations in each layer before nonlinear activation
  - Convolution
- Distance metric : ℓ_1 distance, ℓ_2 distance
- Eigenvalues, eigenvectors : PCA

Pillar 2 - Optimization

- Almost all methods we looked at involved optimizing a loss (cost/objective) function
  - Optimizing an objective $\min_{w} L(w)$, optimizer: $\arg\min_{w} L(w)$
  - Example losses: $\underbrace{\text{mean-square loss}}_{\text{regression}}$, $\underbrace{\text{cross-entropy loss}}_{\text{classification}}$
  - Convexity: key property enabling efficiently finding a global optimum
  - Gradient descent: common approach in ML

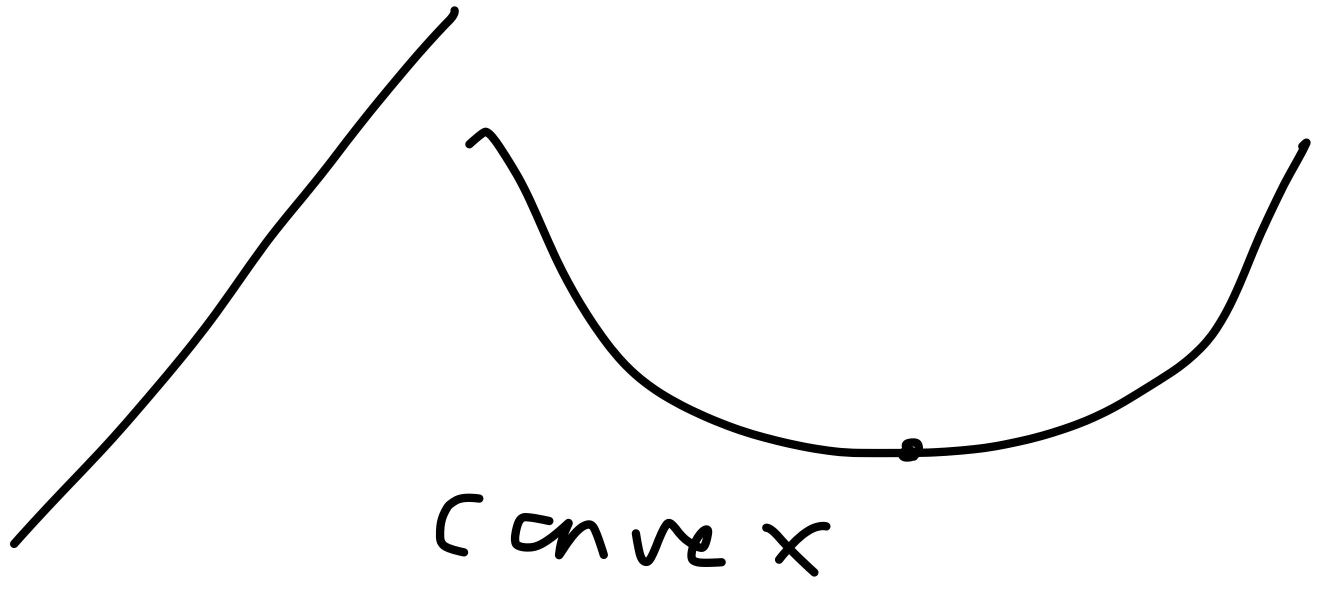

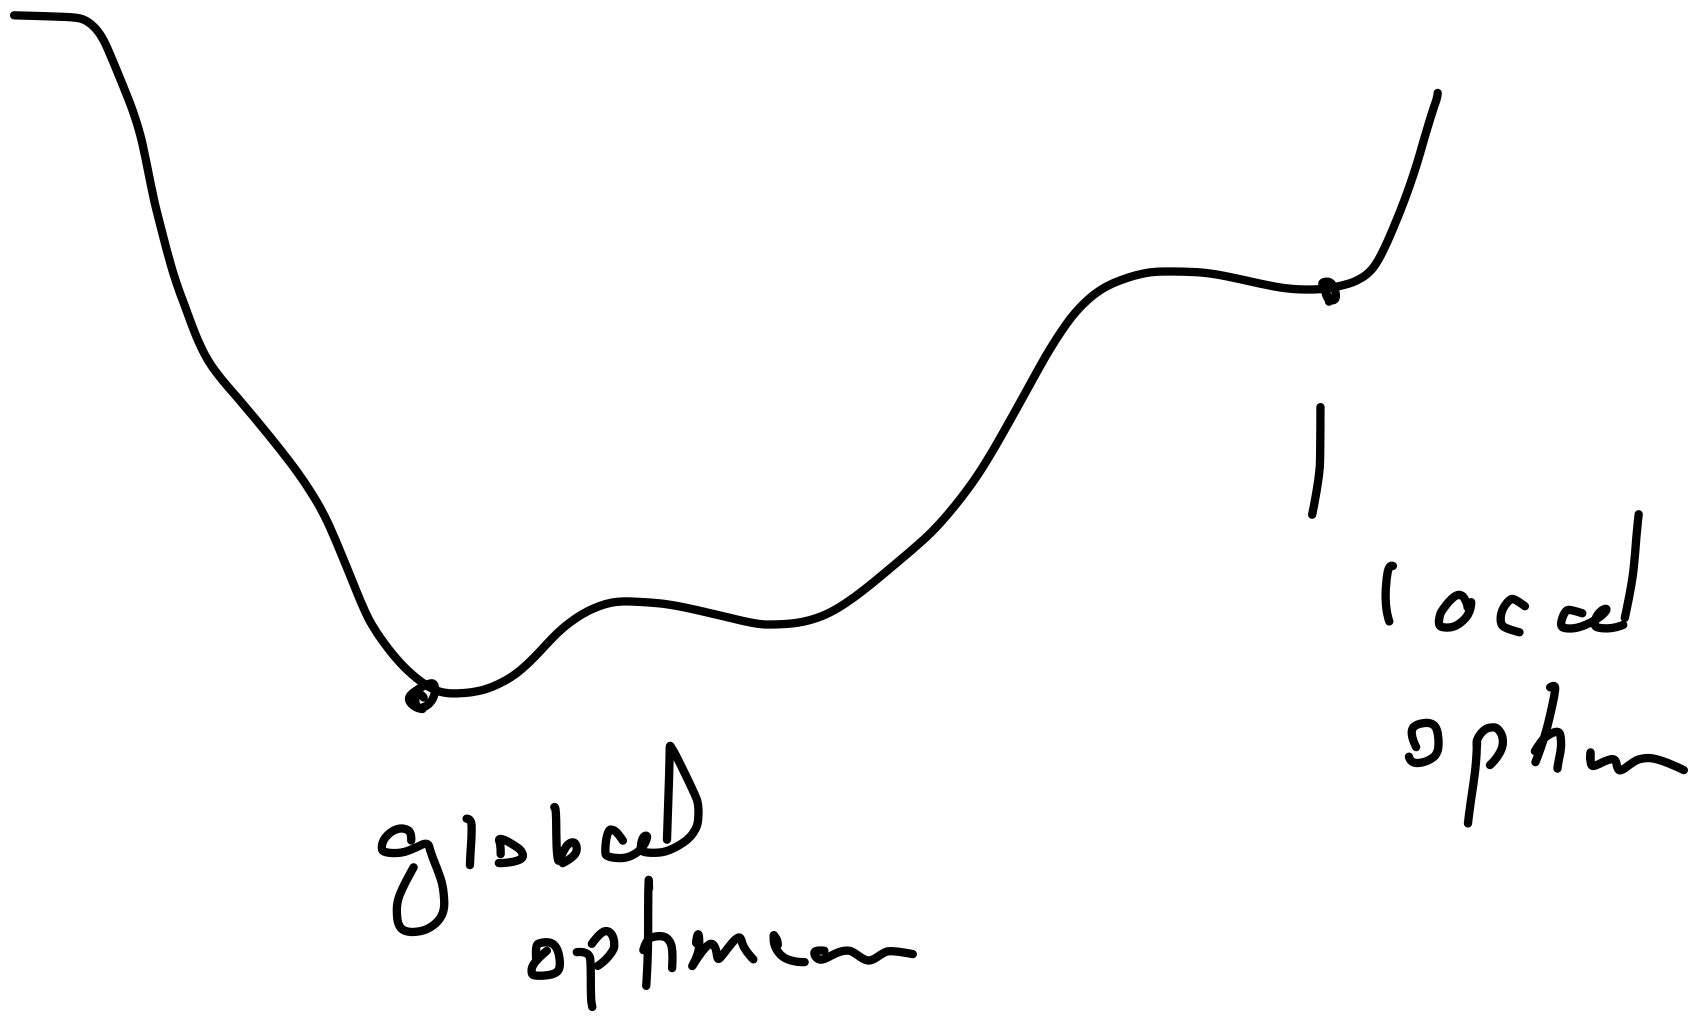

Pillar 3 - Probability and statistics

- Given data, how likely is an outcome? How uncertain is a prediction?
  - Almost all ML approaches we look at also have a probabilistic interpretation

- Probability distribution, conditional distribution, Baye’s rule
- Empirical estimation of expectation, empirical distribution
- Sample mean, variance, quantiles

Pillar 4 - Generalization and learning theory

- We didn't cover this theory but it answers questions such as: how much data we need to learn a given model class, and why we cannot keep reducing both the test & training error. It also motivates the following concepts

- Train, validate, test split of the dataset
- Cross-validation

- Overfitting/underfitting
- Regularization

Pillar 5 - machine learning practice

- Writing pseudo-code, example: gradient descent, k-means, kNN, PCA
- Python and libraries (you had a homework, will not be on the final exam)
- Data inspection and cleaning, eg: normalization/standardization, outliers, missing entries
- Feature engineering, eg: nonlinear transformations
- Evaluation: accuracy, error rate, confusion matrix, mean-square error

- regression

classification

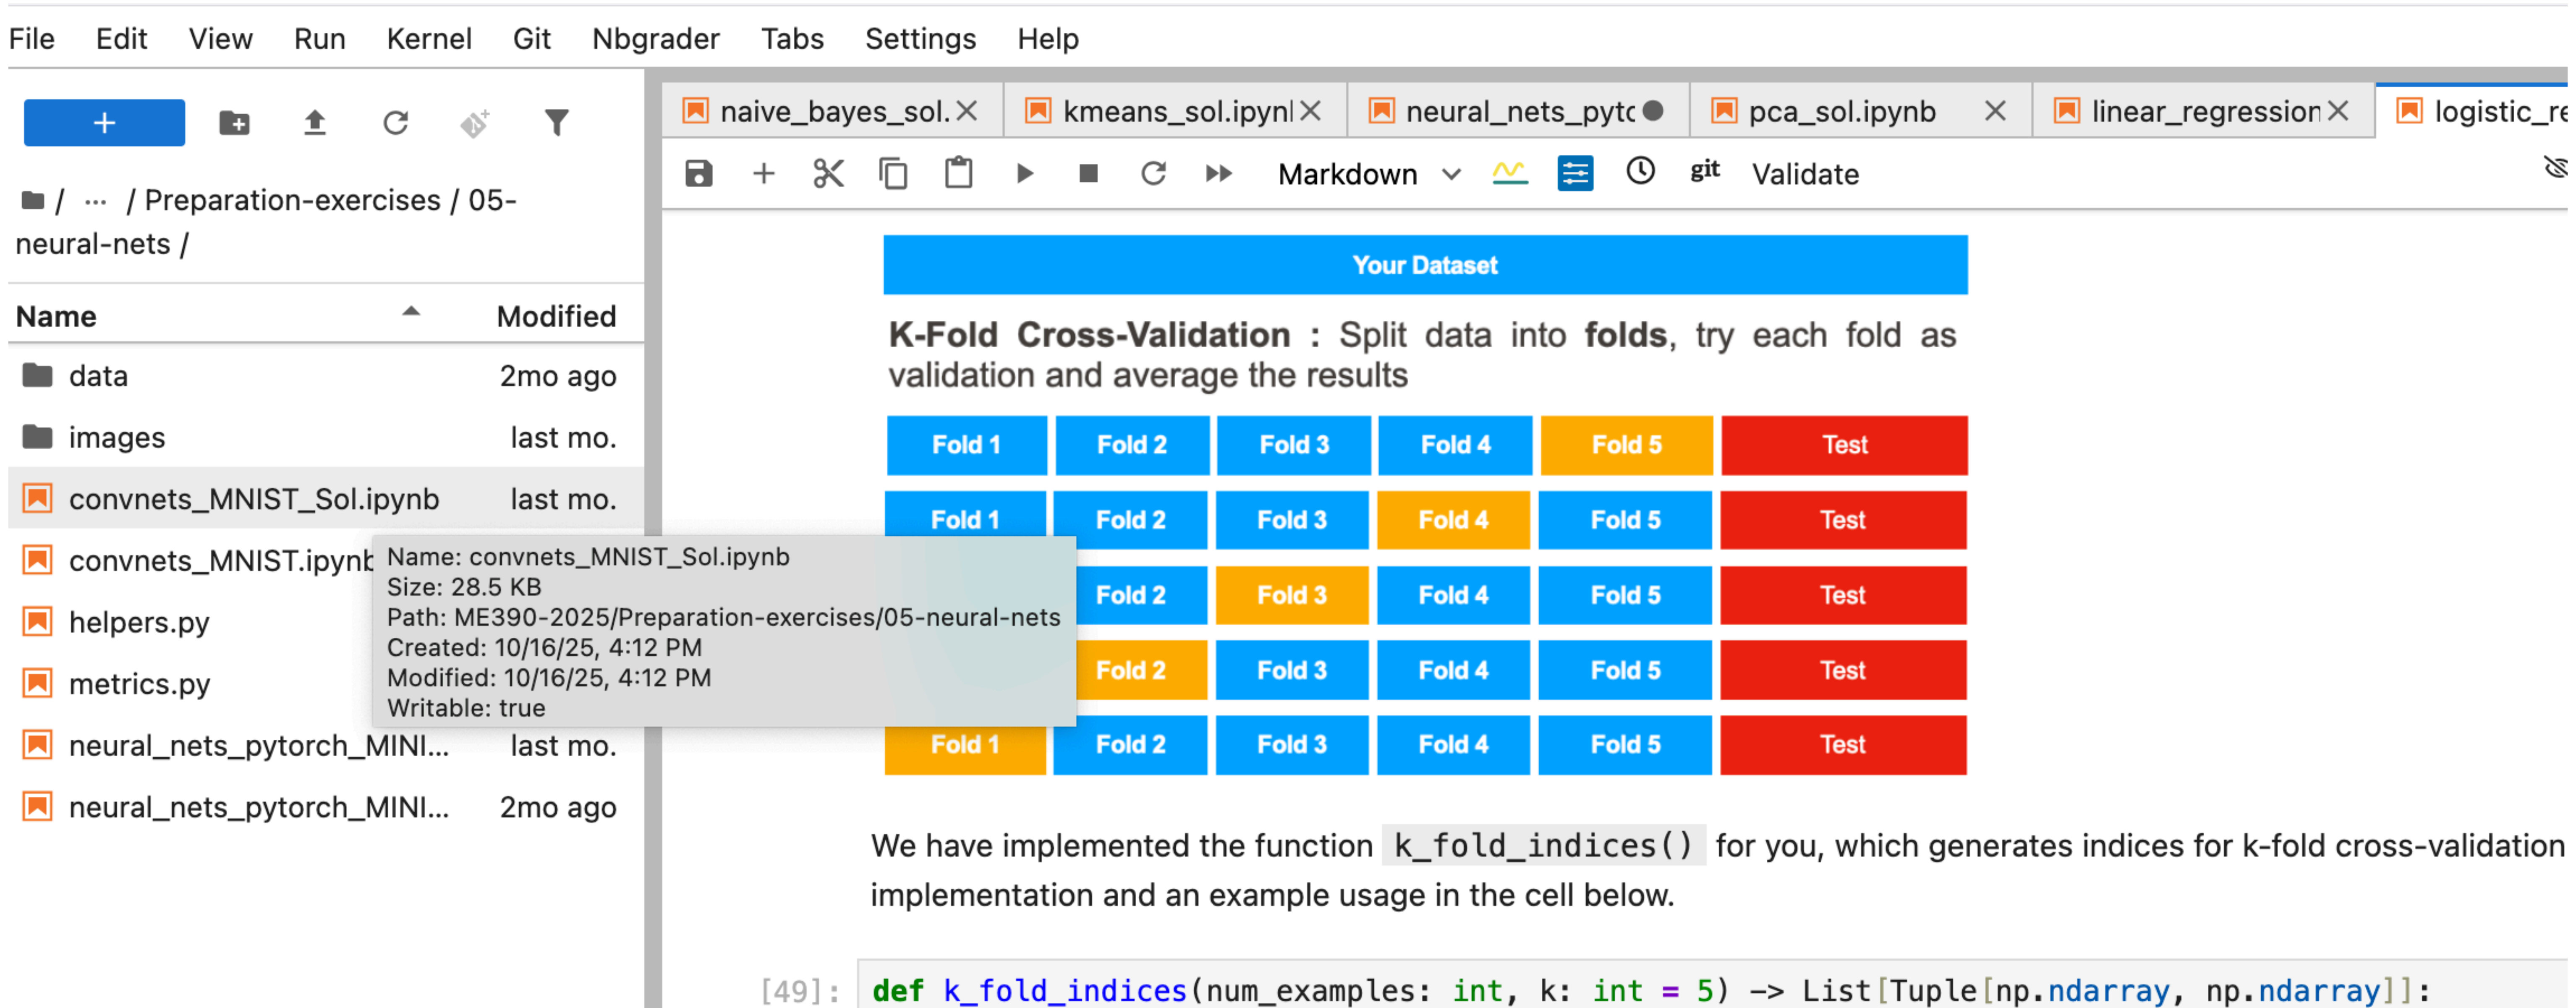

AI sustainability/ethics

- We reflected in class on the following questions (see class answers on Moodle link)

1. What are the main challenges our societies face and where do you see a positive or negative role for AI? As for society, consider not just locally but also globally.
2. Who is using AI and to what purpose? Who is developing AI and with what motivation?
3. How should we govern AI to ensure it benefits humanity and to minimize its risk?

- You can see additional videos (for bonus on exam) on Moodle this week

EPFL

Final exam

- Format: see Moodle course front page
- Date, time and location: see your ISA
- Preparation:
  - Course slides, problem sets, quizzes, python exercises, no programming on final
  - Previous year exams
- Questions? post on EdDiscussions
  - Course staff break: December 20th - January 5th

Course format

In-person lectures, in-person exercise hours.

Assessment

2 in-person quizzes, during exercise hour 10% each, one programming assignment 10%, one end of semester written final exam (70%). The quiz and assignment grades are counted if they help your final grade. So, your final grade is calculated as follows:

final grade = max(70% final + 10% q1 + 10% q2 + 10% a1, 80% final + 10% q1 + 10% q2, ..., 100% final)

Above, q1,q2, refer to quizzes 1, 2, and and a1 refers to assignment 1. It follows that your final grade is the maximum among 8 possible combination of quiz grades, assignment grade and the final exam grade.

Quizzes are on 15.10, 26.11 during the exercise hour. They are 20 minutes and no aid is allowed (no books/notes, no electronics). The assignment should be handed in by 17.12.

The final exam is closed-book. You are allowed one printed cheatsheet, where you can use a double-sided page to write any material from the course.

FORUM
Announcements

EXTERNAL TOOL
Discussions.

FOLDER
Past exams and quizzes

Summary and course takeaways

- Introduction to artificial intelligence (AI) and the machine learning (ML) approach to AI
  - You learned the math behind key topics commonly used in ML
  - You applied the methods learned to real-world datasets
- There are many more things to learn
- Best way to master ML is strengthening one's mathematical background (4 out of 5 pillars)
- When choosing what you will work on, reflect:
  - What we are doing? Why? For whom? Short-term/long-term, local/global impact