

CIVIL-226

Introduction to machine learning for engineers

Alahi Alexandre

| Cursus | Sem. | Type |
|-------------------|------|------|
| Civil Engineering | BA4 | Obl. |
| HES - GC | E | Obl. |

| | |
|----------------------------|---------------------|
| Language of teaching | English |
| Credits | 5 |
| Session | Summer |
| Semester | Spring |
| Exam | During the semester |
| Workload | 150h |
| Weeks | 14 |
| Hours | 5 weekly |
| Courses | 2 weekly |
| Exercises | 3 weekly |
| Number of positions | |

Summary

Machine learning is one of the fundamental building blocks of the Computational Thinking education at EPFL.

Content

This class will focus on the basics of Machine Learning. Students will code in Python.

CONTENT

The class will cover the following concepts:

- Machine learning basics

- Supervised vs Unsupervised
- Regression vs. Classification
- Underfitting vs Overfitting
- Bias vs. Variance
- Parametric vs Non-parametric approaches
- Discriminative vs Generative models
- Shallow vs Deep learning

- Shallow supervised learning

- Linear Regression
- Logistic Regression
- Naive Bayes
- Gaussian naive
- Trees
- Random forest
- Ensemble/bagging/boosting
- Support Vector Machine

- Deep learning

- Neural networks overview
- Type Layers
- Activation functions
- Backpropagation

- Unsupervised learning / dimensionality reduction

- SVD
- PCA
- Embeddings
- Autoencoders

- Unsupervised learning / clustering

- K-means
- Gaussian mixture
- DBSCAN

- Role of input

- Feature engineering
- Role of representation
- Handling different types of features
- Missing values
- Feature expansion

- Role of supervision

- Loss functions
- Multitask learning

- Role of optimization

- Stochastic/Gradient descent

- Recipe/tips for training

- Weight initialization
- Data augmentation
- Regularization techniques
- Transfer learning

- ML ethics

Keywords

Machine learning, Computational Thinking, Artificial intelligence

Learning Prerequisites

Required courses

CS-119(h)

Linear algebra

Basic programming skills (labs will use Python).

Learning Outcomes

By the end of the course, the student must be able to:

- Define the following basic machine learning problems: regression, classification, clustering, dimensionality reduction
- Explain the main differences between them
- Implement algorithms for these machine learning models

- Optimize the main trade-offs such as overfitting, and computational cost vs accuracy
- Implement machine learning methods for real-world problems, and rigorously evaluate their performance using cross-validation. Experience common pitfalls and how to overcome them.
- Finally, civil students will know the basics of Machine learning, and how they can use it in their fields of interest.

Teaching methods

Lectures and lab exercises.

Assessment methods

Lab homeworks: 20%

Midterm: 20%

Final project: 30%

Final exam: 30%

Supervision

| | |
|--------------|-----|
| Office hours | Yes |
| Assistants | Yes |
| Forum | Yes |