

MACHINE LEARNING AND AI

24-hour course

Patrice Latinne

Patrice Latinne is Partner at EY (aka Ernst & Young) where he has been leading the Data & Analytics department since 2014. In his 20+ years of experience, he has coached people, advised about possible strategies and delivered innovative Machine Learning & AI systems at various industries and business departments, including as Practice Partner at CSC. He holds a PhD in Machine Learning & AI from the Artificial Intelligence Lab of ULB (IRIDIA) and a MSc in electro-mechanics engineering.

Course Outline

Machine learning and more generally Artificial Intelligence have recently emerged as one of the cornerstone technologies of the so-called Fourth Industrial Revolution. Their popularity is at the scale of their potential impact on our Society: they are so promising but also so dangerous at the same time. Today, these domains of data science are not yet very well understood by most organizations and hence not very broadly and deeply applied.

In this course, we will examine how Machine Learning models & AI systems can be used or not in various real-world problems. In order to understand what this really means today and what efforts and risks it represents, the students will be guided step by step through the modern Machine Learning development lifecycle until the operationalization and monitoring of an AI system. We will “code” Machine Learning pipelines in Python or R as programming languages and we will also use “best-of-breed” open source and commercial ML workbench platforms for such developments.

Machine Learning & AI is a course intended for those who need to understand this advanced domain of data science and those who want to practice and be capable to translate a business problem into a ML/AI-driven solution.

This course will cover multiple aspects of Machine Learning & AI, from the theoretical perspective to the practical ones. We will gradually study a set of foundational concepts, methods and best practices as well as a set of popular algorithms that will be explained with examples or used properly by the participants themselves in exercises. This course will intensively cover supervised learning and unsupervised learning including a good introduction to deep learning. In order to broaden our view on the type of data that Machine Learning & AI can exploit, we will also introduce Natural Language Processing.

The class meetings will be a combination of lectures on the fundamental material, discussions of business applications from various industries, machine learning workshops and exercises.

At the end of the course, Participants will be able to:

- Participate integrally in a ML project as associate data scientist
- Discuss machine learning and AI systems intelligently with data scientists, machine learning engineers, architects, managers and business analysts
- Map business problems to right techniques with their advantages and drawbacks
- Acquire first programming skills in Python across the end-to-end ML development lifecycle
- Better understand proposals for machine learning projects and investments
- Consider the ratio feasibility / value of simple to complex AI architectures

Course Structure

I. Course Overview & Introduction to Machine Learning & AI

Introduction

Machine Learning & AI history in a nutshell

ML & AI landscape & challenges

Python programming for Machine Learning

Introduction to Project Groups

Business applications:

- *Predicting customer churn in Python*

II. Supervised Learning - Part I

Regression problems

Binary and multi-class, multi-label, multi-output classification problems

Linear regression and regularization

Decision tree, nearest neighbours, logistic regression

Model performance

Intro to SciKit-Learn

Business applications:

- *House pricing*
- *Car purchasing*

III. Supervised Learning - Part II

Support Vector Machines

Voting weak classifiers

Ensemble learning: random forest, boosting, stacking

Hyperparameters optimization

Model comparison

Business applications:

- *Handwritten text recognition*
- *Credit card fraud detection*
- *Propensity to buy*

IV. Unsupervised learning and semi-supervised learning

K-means, hierarchical clustering, density-based clustering

Anomaly detection techniques

Semi-supervised learning

Business applications:

- *Customer segmentation*
- *Sentiment analysis*

V. Deep Learning

Introduction to Deep Learning

Working with images

Introduction to Tensorflow and Keras

Example with a Convolutional Neural Network

Business applications (Computer Vision)

- *Object recognition*

- *Object detection*
- *Semantic segmentation*
- *Predict next character in a sentence*
- *Language translator*

VI. Artificial Intelligence systems I

What if you don't have enough labelled data?

Speech recognition

Chatbots

Business applications:

- *Sentiment analysis*
- *Virtual assistants*

VII. Artificial Intelligence systems II

Recommender Systems

Reinforcement learning

Ethics, fairness and interpretability

Cloud platforms for AI

AI system monitoring & operationalization

The measure of intelligence

Business applications:

- *Intelligent thermostat*
- *E-commerce warehouse flows optimization*
- *Self-driving car system*

VIII. Project Group Presentations and Q&A

Q&A

Conclusions