

Workshop Machine Learning in TensorFlow/Keras Fundamentals

- **Formato do curso:** Live training
- **Preço:** 935€
- **Duração:** 14 horas

This Tensorflow/Keras fundamentals course is designed for users who want to dive into the realms of Artificial Intelligence.

Esta formação é ministrada em Inglês.

Em parceria com a entidade acreditada:



Pré-requisitos

- Knowledge of fundamental concepts of Python for Data Science

Nice to have:

- Basic Linear Algebra
 - Vectors
 - Vectorial spaces
 - Matrix operations
 - Factorizations
- Basic Probabilities and Statistics
 - Histograms, Heatmaps
 - Variance, Covariance, Correlation
 - Distributions: Gaussian, Poisson, etc.
 - Bayes Theory

Programa

Introduction

- What is ML?
- Where can I find it in real life?
- Why now?
- What are the three main categories of ML?
 - Supervised learning
 - Unsupervised learning
 - Reinforcement learning (demo)
- ML pipeline

Machine Learning with sci-kit

- ML pipeline review
- Scikit Python Library
- Data representation
 - Feature matrix
 - Target array
 - Iris dataset example
- Estimator API
- Linear Regression
 - Simple Linear Regression
 - Model Evaluation
 - Polynomial Regression
- Selecting the best model
- The bias-variance trade-off
- Logistic Regression
 - Who survives the Titanic?
- Naive Bayes
 - Gaussian Naive Bayes
 - Multinomial Naive Bayes
 - Categorical Naive Bayes
- k Nearest Neighbours
- k-Means Clustering
- Dimensionality reduction
 - Principal Components Analysis (PCA)
 - Singular Value Decomposition (SVD)
- Decision Trees
- Random Forests

Hands-on Lab:

- Playing around with different values affecting the bias and the variance, calculating precision, recall, F1 and F2-scores, comparing different models on the training and testing accuracies

- Doing a little bit of data preprocessing, analyzing the difference between categorical and numerical data, plotting some relevant statistical values and visually inspecting the correlation between features

Neural Networks in Tensorflow/Keras

- Artificial Neural Networks (ANNs)
 - Neurons
 - Layers
 - Activation Functions
 - More vocabulary
- Popular Frameworks
- Keras
- Linear Regression
 - Defining Models in Keras
 - Training and predicting
- Fashion MNIST example

Hands-on Lab:

- Creating our first custom neural network model
- Choosing the number of layers and the number of neurons per layer
- Tweaking the learning rate
- Training the neural network on real world data

Convolutional Neural Networks (peek)

- Motivation behind CNNs
- CNN Building blocks
 - Convolution Layers
 - Pooling Layers
- CNNs in Keras
- Data Augmentation
- Architectures

NLP using Deep Learning

- Spam detector
- Sentiment analyzer
- Autocomplete

Reinforcement Learning

- Frozen Lake demo
- Flappy Bird demo

Recommender Systems

- Data preparation

- Cosine distance
- SVD for recommender systems
- Autoencoder demo