Course duration

- 2 days

Course Benefits

- Obtain a blueprint of the complete process for deploying a deep learning application: from environment setup to model deployment.
- Obtain a hands-on introduction to TensorFlow and Keras, popular technologies for building production-grade deep learning models.
- Get started with an example web application that uses an HTTP API interface to retrieve model predictions.

Course Outline

1. Introduction to Neural Networks and Deep Learning
    1. What Are Neural Networks?
    2. Configuring a Deep Learning Environment
2. Model Architecture
    1. Choosing the Right Model Architecture
    2. Using Keras as a TensorFlow Interface
3. Model Evaluation and Evaluation
    1. Model Evaluation
    2. Hyperparameter Optimization
4. Productization
    1. Handling New Data
    2. Deploying a Model as a Web Application

Class Materials

Each student will receive a comprehensive set of materials, including course notes and all the class examples.