AI DEEP DIVE

BY DATA SCIENCE RETREAT

CURRICULUM

CURRICULUM

We, curriculum designers for AI deep dive, wanted two things:

(1) to provide you with expertise and confidence in building machine learning solutions to real problems.

(2) To help you start your machine learning career with a significant portfolio of finished products. We believe that a personal brand that shows your growth and possibilities as a professional is the single most decisive ingredient for your future success.

We want to make sure the course fits your current level, while at the same time pushing you to the next level. It's a delicate equilibrium that we achieve through (1) practical courses, (2) guided projects and (3) lectures that introduce you to cutting edge Machine Learning/AI.

CURRICULUM

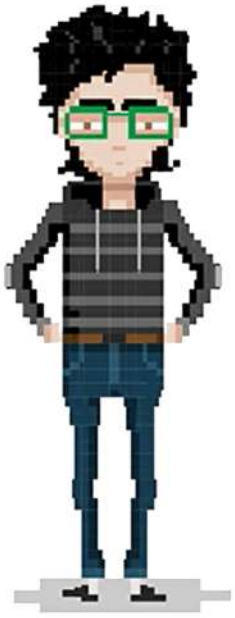

The syllabus works in three blocks: :

The first block consists of classes by experienced professionals that will get you up to speed as AI practitioner. You'll learn and gain a deep understanding of the essential concepts while working on exciting, real-life tasks.

The second block consists of on two projects that will provide you with the experience and confidence necessary to develop whole AI products from scratch. We have chosen the two most foundational and sought-after domains in modern-day machine learning: computer vision and natural language processing. Those two projects will make up the first part of your final portfolio.

CURRICULUM

In the third block, the training wheels fall off and you'll work on your own project. We'll help you decide on an exciting application that showcases your expertise, and we will guide you through the process so that you end up with an impressive final product that you can show to the world, and that will finish up your personal portfolio.

The last two parts will focus on building complete products and developing your portfolio and personal brand but, at the same time, we will offer a series of lectures by leading researchers and practitioners in AI.

PRACTICAL COURSES

FIRST BLOCK

Advanced Python and Pydata Stack

Going deeper into Python and the Pydata stack

A taste of linear algebra in numpy/pytorch

Setting up a working environment

Pandas

32 HOURS

Ensuring everything is right

8 HOURS

Data exploration

Tests

Enforcing constraints with engarde

Automatic feature extraction with featuretools

Fundamentals of Machine Learning

Understanding a selection of classical machine learning algorithms

Cross-validation, regularization

16 HOURS

PRACTICAL COURSES

FIRST BLOCK

How to win in Business and Kaggle

How to get better performance in Kaggle competitions

How to communicate results in business

16 HOURS

Neural networks from Scratch

Implementing a simple NN and backprop using Numpy

8 HOURS

Modern SQL

Retrieving data from relational databases in the 21st century

8 HOURS

Applied ML using Spark

Data wrangling and ML models using Spark

24 HOURS

PRACTICAL COURSES

FIRST BLOCK

Advanced Pipelines

How to build data pipelines in sklearn

Introduction to testing in data pipelines and ML

12 HOURS

Models in Production

16 HOURS

Palladium framework

Heuristics to decide when to retrain

Training on GPUs and going to production without them

Microservices

Building microservices

Building containers

flask

API design

8 HOURS

SPECIALIZATION THROUGH PROJECTS

SECOND BLOCK

First project: Going deep with Computer Vision

50 HOURS

In this part, we begin our journey through deep learning products. There are a lot of new, exciting applications and state-of-the-art models that we will be discussing in our lectures while you build a state of the art deep learning network that can locate different objects in a set of images.

Along the way, you'll be practicing the following skills:

- **Introduction to Pytorch**
- **Engineering skills: how to write clean code, how to set up a working repository**
- **Preprocessing and transforming images**
- **Deciding what architecture to use: convolutions, pooling, inception modules,...**
- **Finding the correct hyper-parameters and loss functions**
- **Transfer learning: leveraging pre-trained networks to boost your product**
- **Visualizing activations and weights**
- **Measuring and reporting performance**

SPECIALIZATION THROUGH PROJECTS

SECOND BLOCK

Second Project: Natural Language Processing

50 HOURS

Deep Learning NLP is currently a less mature field than other domains: whereas there are pretty effective techniques to deal with many problems in areas such as image processing, most of the challenges in NLP are still open to active research. Precisely for that reason, it is an exciting field that also happens to be of immediate practical importance. We will explore this promising area and discuss a whole array of impressive applications as we build a smart chatbot.

Skills that you'll practice as you build your model include:

- **Representing language: embeddings, their geometry, how to avoid bias**
- **Recurrent neural network architectures: plain, GRUs, LSTMs**
- **Attention in RNNs**
- **Solving NLP problems: sentiment analysis, captioning, summarizing, translation, ...**
- **How to build a chatbot that actually works; what is actually possible: the necessity of domain-specific solutions**
- **New frontiers in NLP**

CREATIVE PORTFOLIO PROJECT

LAST BLOCK

250 HOURS

We will assist you in finding a project that helps you brand yourself as an AI specialist and data scientist. You will apply many of the skills that you have developed during the previous parts to make sure your project is top notch.

As part of the project, we will also be discussing the following topics:

- **Deciding on a personal project**
- **Debugging and troubleshooting your data product**
- **Deploying your model**
- **Career advice: how to present your portfolio, how to brand yourself, what to do next**

AI DEEP DIVE

BY DATA SCIENCE RETREAT

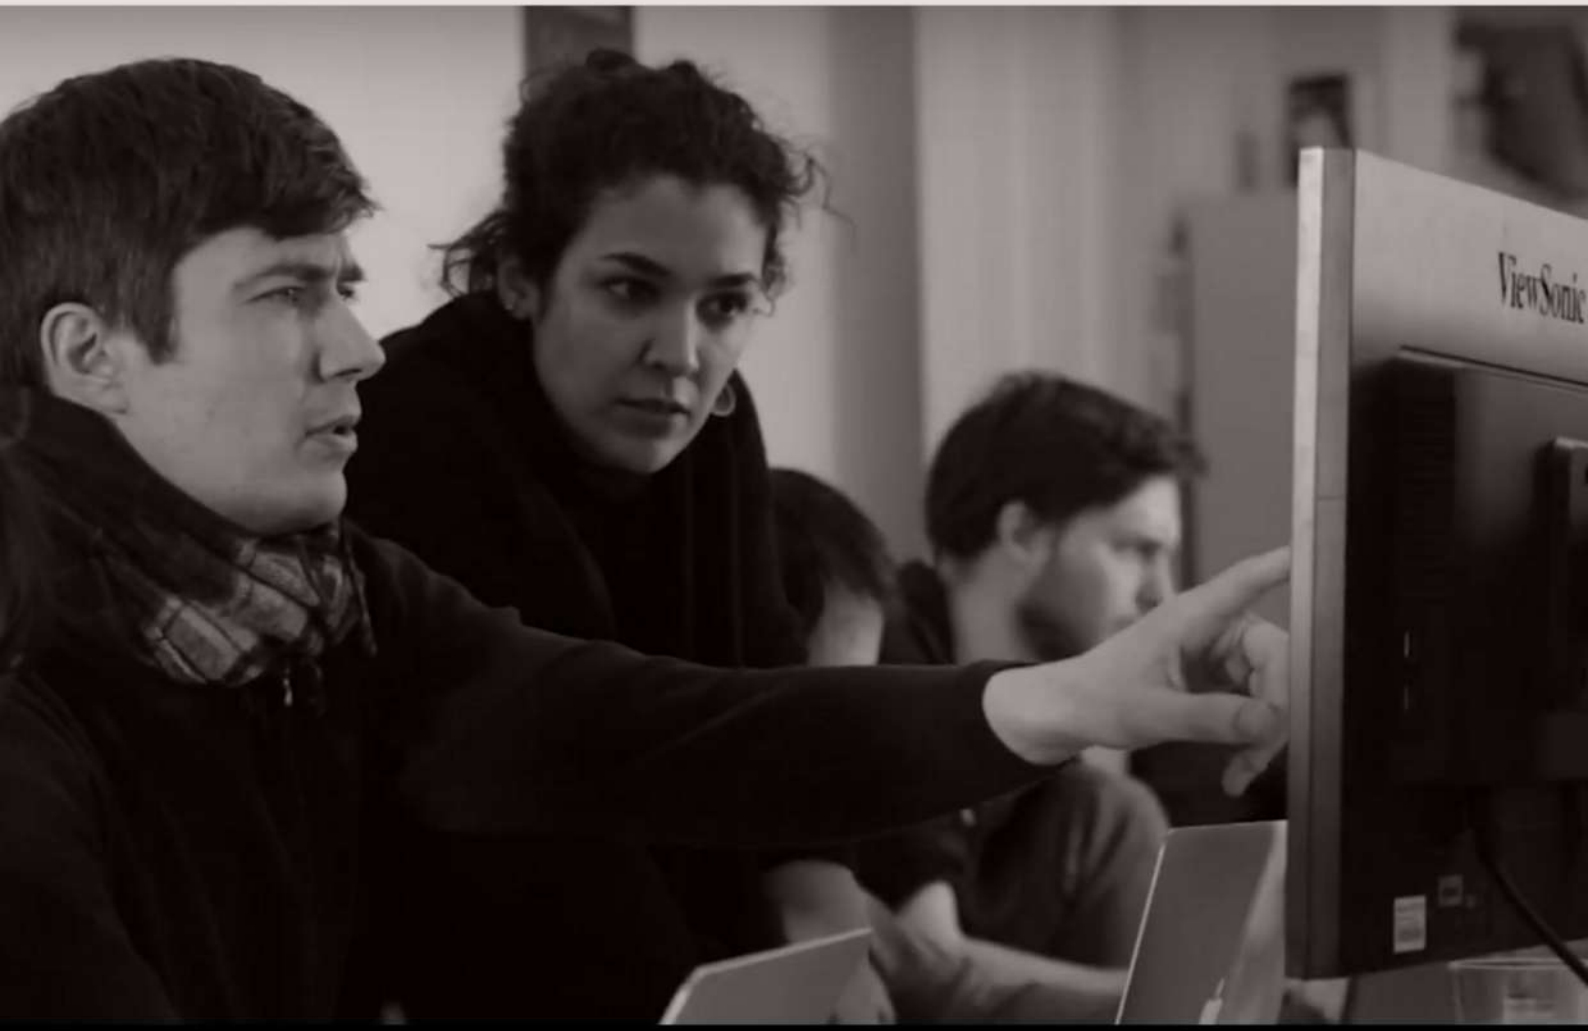

APPLY NOW

aideepdive.com