

A technical introduction to fundamental machine learning concepts for engineers and developers

This course introduces you to fundamental machine learning (ML) concepts.

If you want to get a head start into ML with both a theoretical foundation, but also hands-on coding, this course is for you.

There are a lot of introductions to ML online and many of them are really good, but they tend to focus way too much on the algorithms or mathematical theory.

This course is given from the perspective of a seasoned machine learning practitioner which means that there will be a good balance of theory and practice. Rather than getting lost in mathematical definitions, we will look at issues that arise in real settings and how you solve them.

You could spend hours just trying to understand the intricacies of a specific model, but then you would miss the big picture. There is a whole chain of steps that you need to take and knowledge that you need to have in order to implement a successful ML project, and this course makes sure you get the whole picture.

During the day you will also get the opportunity to solve a problem on your own, with guidance from the instructor.

After the course you will have access to both the slides and all code, so that you can continue to practice on your own.

Course organization

This introduction will be given as a full day course at our training facility with the following blocks:

Introduction
Lecture
Coffee break
Interactive coding example
Lunch break
Lecture
Break
Lab
Wrap-up

Course contents

We will cover the two most common forms of ML – supervised and unsupervised learning. You will learn the difference between the two, and their respective advantages and limitations.

You will also learn how to handle the most common forms that your data may come in – structured and unstructured. This covers tabular data (databases, spreadsheets), but also text, images, and time series.

We will go through how you can train a model from scratch, but also how you can take advantage of the latest breakthroughs from giants like Facebook and Google and use transfer learning to reach state-of-the-art performance for your own models.

Crucially, we will not just focus on training an ML model, but we also incorporate the just as important steps of data preparation and model evaluation, which are frequently overlooked. But how would you expect your model to perform if it cannot fully make use of the data? How would you explain to your boss that you put a model into production without understanding its potential flaws?

In this course we will use the Python programming language, a de facto standard in the world of machine learning. It offers a high degree of freedom for prototyping, while offering great library support for all the leading frameworks, and still performing excellently (essential code paths are run as optimised machine code or on the GPU). If you do not know Python, you will still be able to follow along as long as you have some programming experience.

About the instructor

Kristoffer Röshammar is co-founder and senior developer at Tenfifty. Kristoffer has thirteen years of professional machine learning experience, specialising in natural language processing.



About Tenfifty

Tenfifty is a company that develops custom services using AI, data science and crawling. Whether you have text in any language, voice, image or just numeric data, we have deployed solutions into production for two decades.



TENFIFTY
TAILORED AI SOLUTIONS