




Machine Learning

Table of contents

Welcome!	1
Real results	2
Change when appearance starts	2
Your instructor	2
Is this course for me?	3
The software you will be using	3
Why Machine Learning?	4
What you will learn	4
Bibliography	5

💡 Welcome!


Course homepage




This course is designed with some coding experience who want to learn how to apply machine learning to practical problems.

Welcome!

Machine Learning



- Build and train deep learning models for computer vision, natural language processing, tabular analysis, and collaborative filtering problems
- Create random forests and regression models
- Deploy models
- Use PyTorch, the world's fastest growing deep learning software, plus popular libraries like fastai and Hugging Face



You don't need any special hardware or software — we'll show you how to use free resources for both building and deploying models. You don't need any university math either — we'll teach you the calculus and linear algebra you need during the course.

💡 Get started

Start watching lesson 1 now!

Real results

Our videos have been viewed over 6,000,000 times already! Take a look at the dozens of testimonials about our book and course by alumni, top academics, and industry experts.

'Deep Learning is for everyone' we see in Chapter 1, Section 1 of this book, and while other books may make similar claims, this book delivers on the claim. The authors have extensive knowledge of the field but are able to describe it in a way that is perfectly suited for a reader with experience in programming but not in machine learning. The book shows examples first, and only covers theory in the context of concrete examples. For most people, this is the best way to learn. The book does an impressive job of covering the key applications of deep learning in computer vision, natural language processing, and tabular data processing, but also covers key topics like data ethics that some other books miss. Altogether, this is one of the best sources for a programmer to become proficient in deep learning.

Change when appearance starts

- This is list item 1
- This is list item 2
- This is list item 3

Your instructor

I am Ghassem



Figure 1: Ghassem Tofghi

💡 Get started

Start watching lesson 1 now!

Is this course for me?

It doesn't matter if you don't come from a technical or a mathematical background (though it's okay if you do too!); we wrote this course to make deep learning accessible to as many people as possible. The only prerequisite is that you know how to code (a year of experience is enough), preferably in Python, and that you have at least followed a high school math course.

Deep learning is a computer technique to extract and transform data—with use cases ranging from human speech recognition to animal imagery classification—by using multiple layers of neural networks. A lot of people assume that you need all kinds of hard-to-find stuff to get great results with deep learning, but as you'll see in this course, those people are wrong. Here's a few things you *absolutely don't need* to do world-class deep learning:

Myth (don't need)	Truth
Lots of math	Just high school math is sufficient
Lots of data	We've seen record-breaking results with <50 items of data
Lots of expensive computers	You can get what you need for state of the art work for free

💡 Get started

Start watching lesson 1 now!

The software you will be using

In this course, you'll be using PyTorch, fastai, Hugging Face Transformers, and Gradio.

We've completed hundreds of machine learning projects using dozens of different packages, and many different programming languages. At fast.ai, we have written courses using most of the main deep learning and machine learning packages used today. We spent over a thousand hours testing PyTorch before deciding that we would use it for future courses, software development, and research. PyTorch is now the world's fastest-growing deep learning library and is already used for most research papers at top conferences.

PyTorch works best as a low-level foundation library, providing the basic operations for higher-level functionality. The fastai library is one of the most popular libraries for adding this higher-level functionality on top of PyTorch. In this course, as we go deeper and deeper into the foundations of deep learning, we will also go deeper and deeper into the layers of fastai.

Transformers is a popular library focused on natural language processing (NLP) using *transformers models*. In the course you'll see how to create a cutting-edge transformers model using this library to detect similar concepts in patent applications.



Why Machine Learning?

Deep learning has power, flexibility, and simplicity. That's why we believe it should be applied across many disciplines. These include the social and physical sciences, the arts, medicine, finance, scientific research, and many more. Here's a list of some of the thousands of tasks in different areas at which deep learning, or methods heavily using deep learning, is now the best in the world:

- **Natural language processing (NLP)** Answering questions; speech recognition; summarizing documents; classifying documents; finding names, dates, etc. in documents; searching for articles mentioning a concept
- **Computer vision** Satellite and drone imagery interpretation (e.g., for disaster resilience); face recognition; image captioning; reading traffic signs; locating pedestrians and vehicles in autonomous vehicles
- **Medicine** Finding anomalies in radiology images, including CT, MRI, and X-ray images; counting features in pathology slides; measuring features in ultrasounds; diagnosing diabetic retinopathy
- **Biology** Folding proteins; classifying proteins; many genomics tasks, such as tumor-normal sequencing and classifying clinically actionable genetic mutations; cell classification; analyzing protein/protein interactions
- **Image generation** Colorizing images; increasing image resolution; removing noise from images; converting images to art in the style of famous artists
- **Recommendation systems** Web search; product recommendations; home page layout
- **Playing games** Chess, Go, most Atari video games, and many real-time strategy games
- **Robotics** Handling objects that are challenging to locate (e.g., transparent, shiny, lacking texture) or hard to pick up
- **Other applications** Financial and logistical forecasting, text to speech, and much more...

What you will learn


After finishing this course you will know:

- How to train models that achieve state-of-the-art results in:
 - Computer vision, including image classification (e.g., classifying pet photos by breed)
 - Natural language processing (NLP), including document classification (e.g., movie review sentiment analysis) and phrase similarity
 - Tabular data with categorical data, continuous data, and mixed data
 - Collaborative filtering (e.g., movie recommendation)
- How to turn your models into web applications, and deploy them
- Why and how deep learning models work, and how to use that knowledge to improve the accuracy, speed, and reliability of your models
- The latest deep learning techniques that really matter in practice
- How to implement stochastic gradient descent and a complete training loop from scratch

Here are some of the techniques covered (don't worry if none of these words mean anything to you yet—you'll learn them all soon):



- Random forests and gradient boosting
- Affine functions and nonlinearities
- Parameters and activations
- Transfer learning
- Stochastic gradient descent (SGD)
- Data augmentation
- Weight decay
- Image classification
- Entity and word embeddings
- And much more

 Get started

Start watching lesson 1 now!

Bibliography