

Machine Learning: Andrew NG's course from coursera

Have you ever wondered how handwriting recognition, music recommendation or spam-classification work? The answer is Machine Learning.

I've taken this year a course about Machine Learning from [coursera](#). Andrew NG's course is derived from his CS229 Stanford course. It's my first mooc so I can't compare with another one but one thing is sure: this course is very interesting for someone who likes algorithms.

Presentation

During the first week the professor explains some mathematical concepts needed for the rest of the course and the weeks after he presents different algorithms really used by data scientists. Each week, the course is divided into 3 parts:

- the lecture video (1h-1h30/week) that presents a new algorithm. The recorded video is divided into 10-minutes-long parts so that one can easily watch the course on multiple days. Each part contains some (ungraded) quizzes.
- the review questions (quizzes) that needs to be answered by the end of the week (10 minutes)
- some programming exercises related to the video lecture (in Octave, a free implementation of Matlab). The programming exercises need to be sent by the end of the week. The correction is automatic and one can re-submit his homework if he was wrong. The exercises are not very difficult if you listen seriously to the lecture. The pattern is always the same: understanding the pre-written code and finding the missing 3-4 lines of code in order to make the algorithm work. (2-5h/week). The exercises could take more time you don't come from a scientific background (involving computer science and mathematics) especially for some tricky subjects like the neural network backpropagation algorithm during week 5.

This course is progressive so even for someone working (which is my case) the

workload is doable

Review

This course is just an introduction of machine learning so don't expect to be an expert at the end of it. Andrew NG presents (extremely well) different topics of the field so that you'll have a good overview of machine learning and will be able to look deeper into it.

In my opinion here are the drawbacks of this course:

- the biggest drawback is the will of Andrew Ng to make the course very accessible and by doing so, there are too many equations that come from thin air. This makes the course difficult to really understand for someone who wants to have a deep understanding of the mathematics/algorithms presented by Andrew NG. Thankfully, the real CS229 Stanford lectures are available on Youtube.
- The same problem appears during the exercises (and it's even worse). Since Ng tries to simplify the course, the exercises are also too simplified so that it's possible to finish them without understanding the related algorithm presented in the lecture. But I understand that for a mooc Ng couldn't give exercises that would have required more than 10 hours to be solved.
- Another drawback is the feeling that the course is very "mechanical". By that I mean that Andrew NG recorded the videos on 2011, developed the exercises the same year and then "left". But since it's free I can't really complain about it.
- At last, as a professional developer I think this course should focus more on the problems of time and memory complexity.

But still, if you like algorithms and mathematics, you should definitely take this course. This course gave me a good understanding of how some problems are solved by data scientists and what are the pros and cons of the presented algorithms. For a mooc, I'm quite impressed by the quality of the course.

Here is a link to the course: [coursera](#).

Course overview

If you're interested, here is an overview of the main topics showed during the course:

Week 1

- Introduction
- Linear Regression with One Variable
- (Optional) Linear Algebra Review

Week 2

- Linear Regression with Multiple Variables
- Octave Tutorial

Week 3

- Logistic Regression
- Regularization

Week 4

- Neural Networks: Representation

Week 5

- Neural Networks: Learning

Week 6

- Advice for Applying Machine Learning
- Machine Learning System Design

Week 7

- Support Vector Machines (SVMs)

Week 8

- Clustering
- Dimensionality Reduction

Week 9

- Anomaly Detection
- Recommender Systems

Week 10

- Large-Scale Machine Learning
- Example of an application of machine learning