

# Lesson 01

CSC357 Machine Learning

13 January 2020

# Chapter 1

## Outline

### I. The Fundamentals of Machine Learning

1. The Machine Learning Landscape
2. End-to-End Machine Learning Project
3. Classification
4. Training Models
5. Support Vector Machines
6. Decision Trees
7. Ensemble Learning and Random Forests
8. Dimensionality Reduction
9. Unsupervised Learning Techniques

### II. Neural Networks and Deep Learning

10. Introduction to Artificial Neural Networks with Keras
11. Training Deep Neural Networks
12. Custom Models and Training with TensorFlow
13. Loading and Preprocessing Data with TensorFlow
14. Deep Computer Vision Using Convolutional Neural Networks
  
15. Processing Sequences Using RNNs and CNNs
16. Natural Language Processing with RNNs and Attention
17. Representation Learning and Generative Learning  
Using Autoencoders and GANs
18. Reinforcement Learning
19. Training and Deploying TensorFlow Models at Scale

## Chapter 2

# Questions

1. What is the difference between an initialism and an acronym?
2. The titles of our book’s chapters include some initialisms. What are they? What do they mean?
3. Imagine that you are writing a book about machine learning. You might want to teach your readers...
  - important applications
  - general concepts
  - how specific algorithms work
  - how to use specific software tools
  - steps and techniques that are important when working with several different algorithms and software tools

Can you guess at Aurélien Géron’s purpose in some of his chapters?

4. Do the titles of Chapters 6 and 7 suggest a relationship between the topics of the two chapters? Explain.
5. The title of Chapter 9 identifies one kind of machine learning. You might infer from this title that there is another kind of machine learning that is not found in the title of any chapter. What is that other kind of machine learning?
6. The titles of our book’s chapters include the names of some software libraries. Identify these libraries.
7. The words “natural language” appear in the table of contents. What is a natural language? What kinds of languages are not natural languages?

8. “Classification” is the subject of Chapter 3. Search on the Internet with the words “machine learning classification irises.” What is the purpose of the machine learning examples that you find? Can you generalize to identify a common goal for machine Learning?
9. Classification is one common goal for machine learning. What is another?
10. The word “model” appears in the titles of Chapters 4, 12, and 19. What is model? This is a hard question—you will might have to search on the Internet with several different combinations of words/questions before you find a satisfactory answer.
11. The word “training” appears in the titles of Chapters 4, 11, 12, and 19. What does it mean to train a machine learning model? Again, this is a big subject and we are at the very beginning of our study of machine learning. Tell us just enough to help us get started—we want a sense of how “training” has a special meaning in the field that we are studying.
12. The word “convolutional” is part of the title of Chapter 14. Find out what a mathematical convolution is. We could spend the rest of the year studying this idea. Right now we just want a simple statement of the idea—something that you could share with a classmate at lunch.