Book Review: *Hands-on Machine Learning with Scikit-Learn, Keras, and Tensorflow*, 2nd edition by Aurélien Géron

O'Reilly Media, 2019, 600 pp., ISBN: 978-1-492-03264-9

Michael J. J. Douglass^{1,2}

Published online: 12 August 2020

© Australasian College of Physical Scientists and Engineers in Medicine 2020

Machine learning is rapidly becoming an essential skill for data scientists and it has been applied in most, if not all aspects of science including Medical Physics.

Before reading this book, I had previously worked through a similar textbook called “Deep Learning with Python” by François Chollet which gave me the skills to build deep learning models but with only a passing familiarity of the concepts beneath the models I was building. François Chollet in fact gives a glowing review of his own regarding this book on the back cover.

This textbook, while not walking through examples in the same level of detail as Chollet, definitely filled in the theoretical knowledge I felt I was missing. However, someone with a good python background can easily fill in the missing code to make their own full working examples. Having said that, the book also includes a link to a GitHub repository where the reader can download many practice data sets and Jupyter notebooks with complete machine learning examples to supplement their learning.

This book includes footnotes at the bottom of each page with useful references to the original papers which the concepts currently being discussed were first described, or sometimes, just humorous observations.

The author often explains challenging concepts with useful analogies which I found very helpful. One example of this which I think deserves repeating here was an analogy of dropout regularisation. Employees are asked to flip a coin to see if they should come into work. The company would

be forced to adapt the organisation to not rely on any single person to perform critical tasks forcing the work to be spread across several people. If one person quit or is on sick leave it wouldn't make much of a difference. Eliminating the bus factor in a department. Something worth exploring especially in this age of COVID-19 perhaps?

The book is divided into two parts, fundamentals of machine learning and neural networks/deep learning. At the end of each chapter is a list of exercises for the reader to evaluate what they have learnt in that chapter. The appendix of the book contains solutions to each of those exercises.

Chapter 1 has a very broad overview of what machine learning is, how it all started and where the author thinks it will continue. This chapter contains a very nice list of examples where machine learning could be applied and which chapter to read to guide the reader in designing their own applications. This chapter also includes sections on the major challenges of machine learning such as: limited training data, poor data quality, overfitting and under fitting.

Chapter 2 starts the reader on their own machine learning journey by giving a full worked example of a very typical regression style problem which includes steps such as: statistical analysis of your data, splitting the data into training and test sets, visualisation, data cleaning, feature engineering and finally, builds some simple regression models using python's Scikit-Learn module.

Chapter 3 continues this journey by introducing classification tasks and guides the reader through the “hello world” of classification tasks, predicting digits from the “MNIST dataset”.

Chapter 4 has an excellent discussion on the process of training various types of machine learning models and the various types of gradient descent optimisation algorithms.

Chapter 5 gives an overview of support vector machines, powerful and versatile machine learning models well suited for classification of complex small and medium size datasets.

✉ Michael J. J. Douglass
Michael.douglass@adelaide.edu.au

¹ Department of Medical Physics, Royal Adelaide Hospital, Adelaide, SA 5000, Australia

² School of Physical Sciences, University of Adelaide, Adelaide, SA 5005, Australia

Chapter 6 introduces Decision trees. This was an area of machine learning which I was not very familiar. This chapter helped me to understand the usefulness of decision trees compared with other machine learning techniques and their limitations (e.g. their sensitivity to changes in coordinate system). This chapter introduces the concept of white box vs. black box machine learning models. Decision trees are one example of a white box model because the decisions choices are visibly apparent. These concepts are built upon in chapter 7 with a discussion of ensemble learning and random forest models. The author makes a good analogy between posing a question to one expert or thousands of random people. The aggregated answer of the group of people may often be better than the expert. This is termed the “Wisdom of the crowd”. This is a concept in ensemble learning and highlights the power of random forest models.

Chapter 8 has a discussion of dimensionality reduction techniques and the “Curse of dimensionality”. This “curse” is very well explained, again, through simple analogies of the statistics of single points in multi-dimensional space. Principal component analysis is well explained in this chapter.

Chapter 9 covers unsupervised learning techniques including the various types of data clustering techniques such as K-Means.

Chapter 10 is the start of part II of the book and the beginning of the neural networks and deep learning discussion. This part of the book focusses on artificial neural network applications with Keras (a popular and high level machine learning API in Python) and Tensorflow. This chapter has a great discussion on the origins of artificial neural networks and their connection to biological neural networks. Although, the author is careful to emphasise that perhaps the analogy is still used too often and the relationship is not nearly as closely linked as most people think. There is also a great explanation of how backpropagation is used to train an artificial neural network.

This chapter then explains how to build some simple Keras machine learning models and gives recommendations for best practice. This chapter really highlights the subtleties of how Keras in particular works and what to look out for when training your model.

Chapter 11 begins the section on deep learning applications. It starts by describing some of the common issues with deep learning models and how to approach them, such as: limited training data and reducing training time. There is a brief but intriguing discussion on the developments of optimising techniques in deep learning such as optimisers utilising second order partial derivatives which is an emerging field but currently computationally infeasible due to memory requirements. There was a really interesting section on “Monte Carlo dropout” techniques which can boost existing

model performance without having to re-train them. This is a technique particularly relevant to applications which are “risk sensitive” such as radiation oncology where it is useful to know how confident the model is in its prediction.

Chapter 12 and 13 introduces TensorFlow for building custom models and some of the tools it has available for data pre-processing.

Chapter 14 covers computer vision applications with convolutional neural networks. It summarises some of the best performing convolutional neural networks to date such as AlexNet, VGGNet and ResNet and how to use transfer learning to build on these models for other applications. It also has some recommendations for open source image labelling tools for machine learning applications.

Chapter 15 and 16 discuss recurrent neural networks and their applications in natural language processing. There is a full worked example using recurrent neural networks to write new Shakespearean text from his existing works. This chapter was a lot of fun to read and work through the examples.

Chapter 17 focuses on auto encoders and generative adversarial networks (GANs). GANs have fascinated me ever since I first saw what they were capable of at the ACPSEM Machine Learning course in 2019. This chapter has some very nice worked examples on how to build your own GANs.

Chapter 18 focusses on reinforcement learning such as teaching a machine to play games.

Chapter 19 talks about deployment of your machine learning model, running machine learning models on embedded devices and how to speed up training and execution of your models with GPU acceleration.

With the addition of the resources provided in the GitHub repository, this book has everything you could want in a machine learning textbook. It provides an excellent starting point for someone who knows little or nothing about machine learning and wants to enter the field. It is also an excellent reference for someone who wants to build a specific application and needs a starting point to build on. This book’s strength is its vast exploration of all aspects of machine learning while explaining the nuisances of machine learning (particularly using python, Scikit-Learn and Keras) in practice. The advice given in each chapter will help you avoid some of the common pitfalls you might encounter on your own machine learning journey.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.