

# Chapter 22

## Conclusion and Next Steps



### 22.1 Overview

The concepts of machine learning and artificial intelligence have become quite mainstream and the focus of a lot of cutting-edge research in science and technology. Although most of the concepts in machine learning are derived from the mathematical and statistical foundations that are quite hard to understand without having background in graduate-level courses in those topics, the artificially intelligent applications based on machine learning are becoming quite commonplace and accessible to anyone. This simplicity of applications has opened the doors to these topics for a huge number of engineers and analysts, which otherwise were only accessible to people doing their PhDs. This is good and bad in a way. This is *good* because these topics do have applications in technology that is used in every little aspect of our modern lives. The people who are building all this technology are not necessarily the people who are pursuing their PhDs in mathematics and statistics. Hence, when more people are armed with these powerful techniques, the faster is the pace of evolution of technology and revolution of lives of all the people. This is *bad* because oversimplification of the concepts can lead to misunderstanding and that can have catastrophically worse if not dangerous effects.

I have come across many colleagues and friends who had asked questions about different problems or applications of machine learning in day-to-day life. In spite of having a plethora of information on these topics on the Internet, in most cases, it creates more confusion than answering the questions. Many books, Wikipedia, and many other websites are dedicated on this topic. However, in my investigation, the material presented in most sources is focused on a small subset of topics and typically either goes too deep in theory or stays too superficial that leaves the readers searching for more clarity. Most of these sources are useful when you know exactly what you are looking for, but unfortunately, in most cases, that is not true. In such case, it is easy to get lost in the ocean of such information.

In recent years, the areas of machine learning and data science have emerged as stand-alone topics for graduate (Master's)-level courses. I have specifically organized this book to serve as a single reference textbook needed to create a course on the topic that can span one or two semesters.

Each individual aspect of this technology when looked in micro detail starts to seem obvious, trivial even. However, when all these individual components come together, the problem that is solved and the experience that is delivered are just magical. And that is what keeps me excited about this topic every single day. Hope to have passed this excitement through this book to every student!