Artificial Intelligence and Machine Learning for Business.

A No-Nonsense Guide to Data Driven Technologies

Steven Finlay

Relativistic

Relativistic Books

First published in Great Britain by Relativistic

ISBN-13: 978-1-9997303-0-7

ISBN-10: 1999730305

No plants or animals were mistreated in the writing of this book.

Cover image: Thanks ESA/Herschel/PACS/MESS Key Programme Supernova Remnant Team; NASA, ESA and Allison Loll/Jeff Hester (Arizona State University).

To Sam and Ruby.

CONTENTS

Foreword

56 minutes one way. 112 minutes for the round trip. That's the time it takes my commuter train to travel from my home town of Preston to the great British city of Manchester and back. I find this is an ideal time to catch up on a bit of reading. Therefore, I thought that a concise no-nonsense book about machine learning, which could be read in about this time, would be useful to people with not much free time on their hands.

An understanding of machine learning is important because it is having a huge impact across many aspects of our lives. In particular, it is driving the explosion in "Artificial intelligence" applications in areas such as language translation, autonomous robots and medical diagnosis, to name but a few.

Machine learning is also having a direct impact on many everyday business functions. Automated systems based on machine learning are replacing numerous tasks that where once undertaken by people. This is giving organizations that embrace machine learning a competitive advantage over their rivals because of the efficiency savings and improved customer service that such systems can deliver.

This is the second edition of the book (or the "extended cut" as I like to call it) which was previously titled 'Predictive Analytics in 56 Minutes.' The first edition was intended to be read in just under an hour based on the average person reading at about 250 words a minute. This revised version is substantially longer, covers a wider range of topics and is an overall much more meaty and satisfying read. With 112 minutes to play with, that equates to almost 30,000 words. Including preliminary matter and appendices the total text is

considerably longer, but I hope that the average reader is able to digest the main part of the book in the allotted time.

1 Introduction

Do you have a smart phone or a credit card? Do you buy stuff from supermarkets or play computer games? Are you employed or use health care? If your answer to any of these questions is yes, then machine learning will be having an impact on your life in one way or another. This is because it is the primary tool organizations use to leverage the data they hold about you in order to decide how they are going to deal with you. They use tools derived using machine learning to inform them about how you are likely to behave under different circumstances, and hence the way that they should treat you in order to maximize their (and sometimes your) objectives.

Machine learning is also being used to enhance our everyday experiences via "artificially intelligent" machines and computer interfaces. Amazon's echo, Apple's Siri and Google Translate are just three well known products that demonstrate the fruits of machine learning. Similar artificial intelligence tools, based on machine learning, are being incorporated into everyday devices at an ever increasing rate and are supporting or supplanting human experts in many everyday tasks. For example, insurance underwriting and medical diagnosis.

This concise text provides a managerial (i.e. non-technical and no complex formulas) overview of machine learning, what it is and how it's used. No prior knowledge is assumed. To put it another way, if you can read and write and do basic arithmetic (there is a bit of arithmetic, but that not much!) then you should be OK with the material in this book.

A good question to ask at this point is: why do I need to know

about these things? One reason is personal. Decision making systems, based on machine learning, are widely used by organizations to decide how to treat you, your friends and your family. They use these tools to decide if you will receive a great offer or a poor one, if you should be placed at the front or the back of the queue, if you will be subject to a tax audit or treated as a suspect in a criminal case. Therefore, it's not a bad idea to know something about these things so that you can understand why an organization may have treated you in one way and not in another.

The other reason to learn about machine learning, and the one that is the main focus of this book, is that it is now a mainstream business tool. Not that long ago, machine learning was the domain of few nerdy specialists working mainly in academia, financial services or large marketing departments. These days, anyone whose day-to-day role requires them to deal with people and make decisions about them on a regular basis may find themselves using the fruits of machine learning. Those organizations that can use machine learning intelligently to solve business problems, improve efficiency and cut costs will benefit at the expense of their rivals.

This doesn't mean that you need to learn all the technical details that a machine learning specialist (a data scientist) needs to know. However, having a working knowledge of what machine learning is and how it can be used to help organizations deliver better products and services will be beneficial. Not least, because in order to make effective use of machine learning it needs to be focused on business objectives to address specific problems that organizations face.

If on the other hand you happen to be an equation quoting, formula juggling, bad ass mathematical genius who thinks they know all there is to know about artificial intelligence and machine learning, then this book may also have some value for you too. Possibly even more than those who know nothing at all. Why? Because if all you care about are the theoretical aspects of machine learning then you face a real risk of hitting a brick wall when it comes to delivering useful solutions in the minefield that is the real world; a world populated with social, ethical and political issues.

This, together with a growing raft of privacy and data protection legislation, could derail your solutions no matter how good they are mathematically. Without consideration of these "soft issues", then the best case is that the solutions you develop don't get to be deployed. The worst case scenario is that you design a machine learning solution that lands you in court because it unfairly discriminates against minorities, women or some other group of people. Maybe you can skip a few of the earlier chapters, but you should certainly read the later ones.

To get the most out of machine learning, data scientists need to engage with business users to understand their problems. Data scientists also need to understand an organization's culture, and its approach to the adoption of new ideas, technologies and working practices. Legal and regulatory issues in the region(s) in which the their clients operate also require due consideration. It doesn't matter how good a solution is in terms of cutting edge hardware and software, if it's not aligned with an organization's business objectives and operational processes, then it's all just a waste of time and money.

Lots of solution suppliers can bamboozle you with their fancy tech and the latest terminology, which is often just a rebranding of last year's tech with a new twist. However, the ones that add value will be those who spend time understanding how you and your organization work. They will then determine if and how their machine learning solutions can be used to improve what you do, and explain this to you in simple language which you can understand without needing to reach for Wikipedia.

Successful machine learning is a two way thing. Data scientists need to know something about your organization and what it does, and you need to understand a little bit about machine learning. Without this joint understanding it's unlikely that you or your organization will be able to realize the full benefits that machine learning has to offer.

OK. So what will you learn from reading this book? The key topics that we are going to cover in the chapters following this one are:

- What machine learning and artificial intelligence are.
- The sort of things organizations use machine learning and artificial intelligence for.
- What a predictive model looks like.
- The relationship between machine learning and "Big Data".
- The people and tools needed to apply machine learning.
- How to apply machine learning to improve business processes and the bottom line.
- Legal and ethical issues which need to be considered when using predictive models, developed using machine learning.
- The current limitations of machine learning and artificial intelligence.
- How advanced forms of machine learning are applied to drive artificial intelligence applications, such as object recognition and language translation.

Recommended further reading and a glossary of common machine learning/artificial intelligence terms are provided in Appendices B and C respectively.

2 What are Machine Learning and Artificial Intelligence (AI)?

Machine learning is the use of mathematical procedures (**algorithms**) to analyze data. The aim is to discover useful patterns (relationships or correlations) between different items of data. Once the relationships have been identified, these can be used to make inferences about the behavior of new cases when they present themselves. In essence, this is analogous to the way people learn. We observe what goes on around us and draw conclusions from our experiences about how the world works. We then apply what we have learnt to help us deal with new situations that we find ourselves in. The more we experience and learn, the better our ability to make decisions becomes.

One application of machine learning is object recognition. The goal is to develop systems which can identify everyday objects from images the system is presented with. The data used to develop an object recognition system consists of pictures of different objects such as chairs, umbrellas, washing machines, and so on. Each picture presented to the machine learning algorithm is labeled to identify what type of object it contains[1]. For each type of object there may be hundreds or thousands of different images, representing alternative forms of that object from different perspectives (you'd be surprised at just how many variants of an umbrella there are!)

By analyzing the different images, machine learning algorithms identify that certain objects are associated with certain features (patterns). Chairs tend to have protuberance (legs) coming from a flat, often squarish base (the seat). They are also differentiated from

stools by having a back rest. Washing machines tend to be cube shaped with knobs on, and are never green or orange (please, please let me know if you ever come across somewhere where I can buy an orange washing machine!) Likewise, umbrellas are long and thin (when closed) are often, but not always, black, and so on.

One of the most common, and arguably the first application of machine learning, is prediction. It's about using machine learning techniques to determine something that you don't currently know, based on the information that you currently have available. The patterns that one finds relate to the relationships between behaviors and outcomes. Very often this relates to people's past behavior and what they then subsequently went on to do. Having identified the relationships that exist, it is then possible to make predictions about someone's future behavior based on their current state of being. If you give me a sample of peoples' previous purchasing history, I can utilize machine learning to identify patterns in their purchase behavior. I can then use these patterns to predict what goods someone is likely to buy next; i.e. future purchases are the outcome that I want to predict. This allows me to target them with tailored marketing offers for those specific products.

Using machine learning for prediction is sometimes referred to as predictive modelling or **Predictive Analytics** (PA). In fact, predictive analytics is such a common application of machine learning that many people (rightly or wrongly) often use the two terms interchangeably.

Predicting the future behavior of individuals is what people usually associate with machine learning, but there are other situations and problems to which machine learning can be applied. All you need is some unknown event or thing that you want to determine (predict), and this could be in the past, present or future. Doctors examine their patients, carry out tests and question them about their symptoms in order to gather evidence (data). They then use the data that they have gathered to come to a view as to what they think is wrong with the patient. They are not making a prediction about the patient's future health, but trying to work out

what's wrong with them today. Doctors can do this with a high degree of accuracy because they cross reference the patient information that they have obtained against what they have learnt from years of training and practice. In other words, they are looking for where the patient's symptoms correlate with their knowledge of known illnesses. Machine learning can be applied in the same way. Given a host of detailed information about the symptoms of different illnesses, machine learning can be used to estimate the probability that someone has a certain condition, based on the symptoms that they present.

Another way to think about machine learning/predictive analytics is as a method of reducing uncertainty. There are a whole host of possible outcomes that could occur in any given situation. Machine learning won't tell you with absolute certainty which outcome will occur, but it can provide some insight into the likelihood, or odds, of each outcome.

You may know that when someone goes grocery shopping they often buy bread, wine and chicken, but with machine learning you can determine that there is say, an 80% chance that the next product they buy is bread, a 15% chance that they buy wine and a 5% chance that they buy chicken. Therefore, if you want to encourage them to make their next purchase in your store, you are far more likely to win their custom with a bread offer rather than a wine or chicken offer.

---

1 This type of machine learning, where the algorithm is told what each item or event is, is called supervised learning. Unsupervised learning is where the individual results are not disclosed. The only information available is the overall success rate. Not surprisingly, supervised learning tends to be more widely used in practice.