

[Read and download] Data Smart: Using Data Science to Transform Information into Insight

Data Smart: Using Data Science to Transform Information into Insight

John W. Foreman

ePub | *DOC | audiobook | ebooks | Download PDF



 Download

 Read Online

#50292 in eBooks 2013-10-31 2013-10-31 File Name: B00F0WRXI0 | File size: 60.Mb

John W. Foreman : Data Smart: Using Data Science to Transform Information into Insight before purchasing it in order to gauge whether or not it would be worth my time, and all praised Data Smart: Using Data Science to Transform Information into Insight:

69 of 69 people found the following review helpful. Great Book, Lousy References
By Lou Gutnick
I would go as far as to say that the book is brilliant. First, a drop about me from the standpoint of this book. I have been an IT professional for many years specializing in programming, database, and MS Office add-ons. Part of my job entails self enrichment, that is, expand my working knowledge in areas potentially important for my job. I chose Foreman's book to help with this task for a number of reasons: a) Data Science is a hot area and my company does have a Data Science group, b) I have lots of data experience under my belt - I felt that it would be nice for once to get some useful information from the data, and c) I have a really good Excel background - so I figured that Foreman's approach would be perfect for me - little did I know that I would seriously add to my Excel bag of tricks. The author makes the

assumptions that: a) the reader is somewhat technical, b) he knows nothing about Data Science, and c) he is relatively comfortable working in Excel. Reading the book is a joy because Foreman has a cozy, chummy style. He definitely doesn't throw all the technical stuff at the reader rat-tat-tat machine gun style like many other authors. Instead, Foreman gently introduces his topics and then ramps up technical details carefully. This most definitely helps the learning process. Speaking of learning, by the end of the you will have learned important concepts in "machine learning" and I believe that you will be ready for the next step. I sure was. I found the topics interesting and I wanted to learn more. This is where the book's only problem area comes into play - the next step. Foreman has 3 references - one good, but minor, one terrible, and the other is inappropriate. Let me explain. Foreman recommends a free resource as a follow-on to his Forecasting Chapter. This is a good reference, but I believe that Forecasting is a minor topic in Data Science, unless, of course, Forecasting becomes your thing. Foreman's main reference is: "Data Mining with R" by Luis Torgo. Foreman recommends this as the next step after his book. I tried to read this several times, but couldn't. It certainly wasn't my next step. The other reference, "The Elements of Statistical Learning" by Trevor Hastie, et. al, is totally inappropriate for Data Science newbies. You can check out the reviews for this book and you'll see that you need a pretty serious background in statistics to get anything out of that reference. In fact, the author Hastie says as much in his next book "An Introduction to Statistical Learning- with Applications in R". This is the appropriate next step, but I'll get to that in a moment. Here are my recommendations: A. Read Foreman's book and follow along with him in working through the Excel spreadsheets. This is a first step in getting comfortable with Machine Learning. B. Take the Coursera courses: 1) Machine Learning Foundations: A Case Study Approach, and 2) Machine Learning: Regression. The courses are free unless you want completion certificates, in which case there is nominal cost. C. Now you are ready for: An Introduction to Statistical Learning: with Applications in R (Springer Texts in Statistics) This book is also available for free by the authors - check online. 232 of 241 people found the following review helpful. Insightful, practical, and colorful. Perspective from a biased reviewer. By Evan Miller Disclaimer: I served as a paid technical editor for Data Smart. I am not affiliated with the publisher, but I did receive a small fee for double-checking the book's mathematical content before it went to press. I also went to elementary school with the author. So as you read the rest of the review, keep in mind that this reviewer's judgment could be clouded by my lifelong allegiance to Lookout Mountain Elementary School, as well as the Scarface-esque pile of one dollar bills currently sitting on my kitchen table. Anyway, books about "Data" seem to fit into one of the following categories: * Extremely technical gradate-level mathematics books with lots of Greek letters and summation signs * Pie-in-the-sky business bestsellers about how "Data" is going to revolutionize the world as we know it. (I call these "Moneyball" books) * Technical books about the hottest new "Big Data" technology such as R and Hadoop Data Smart is none of these. Unlike "Moneyball" books, Data Smart contains enough practical information to actually start performing analyses. Unlike most textbooks, it doesn't get bogged down in mathematical notation. And unlike books about R or the distributed data blah-blah du jour, all the examples use good old Microsoft Excel. It's geared toward competent analysts who are comfortable with Excel and aren't afraid of thinking about problems in a mathematical way. It's goal isn't to "revolutionize" your business with million-dollar software, but rather to make incremental improvements to processes with accessible analytic techniques. I don't work at a big company, so I can't attest to the number of dollars your company will save by applying the book's methods. But I can attest that the author makes difficult mathematical concepts accessible with his quirky sense of humor and gift for metaphor. For example, I previously had not been exposed to the nitty-gritty of clustering techniques. After a couple of hours with the clustering chapters, which include illuminating diagrams and spreadsheet formulas, I felt like I had a good handle on the concepts, and would feel comfortable implementing the ideas in Excel -- or any other language, for that matter. What I like most about the book is that it doesn't try to wave a magic data wand to cure all of your company's ills. Instead it focuses on a few areas where data and analytic techniques can deliver a concrete benefit, and gives you just enough to get started. In particular: * Optimization techniques (Ch. 4) can systematically reduce the cost of manufacturing inputs * Clustering techniques (Ch. 2 and 5) can deliver insights into customer behavior * Predictive techniques (Ch. 3, 6, and 7) can increase margins with better predictions of uncertain outcomes * Forecasting techniques (Ch. 8) can reduce waste with better demand planning It may take some creativity to figure out how to apply the methods to your own business processes, but all of the techniques are "tried and true" in the sense of being widely deployed at large companies with big analytics budgets and teams of Ph.D.'s on staff. This book's contribution is to make these techniques available to anyone with a little background in applied mathematics and a copy of Excel. For that reason, despite the absence of glitter and/or Jack Welch on the book's cover, I think Data Smart is an important business book. I had a few criticisms of the book as I was reading drafts, but almost all of them were addressed before the final revision. For the sake of completeness, I'll tell you what they were. Some of the chapters ran on a bit long, but these have been split up into manageable pieces. The Optimization chapter is a bit of a doozie, and used to be at the very beginning, but the reader can now "warm up" with some easier chapters on clustering and simple Bayesian techniques. The Regression chapter originally didn't discuss Receiver Operating Characteristic curves, which are important for evaluating predictive models visually, but now ROC curves are abundant. Only one real criticism from me remains: I would have liked to see more on quantile regression, which is only mentioned in passing. It's a great technique for dealing with outlier-

heavy data. The book by Koenker has good but highly mathematical coverage, and I would have loved to see this subject given the Foreman treatment. But, you can't have everything, and I suppose John needs to leave some material for Data Smart 2: The Spreadsheet of Doom. In sum, Data Smart is a well-written and engaging guide to getting new insights from data using familiar tools. The techniques aren't really cutting-edge -- in fact, most have been around for decades -- but to my knowledge this is the first time they've been presented in a way that Excel-slinging business analysts can apply the methods without needing her own team of operations researchers and data scientists. If you're not sure whether the book's sophistication is on par with your own skills, you can download a complete sample chapter (as well as example spreadsheets) from the author's website. One last thing: unlike many books with a technical bent, the prose is engaging and extremely clear. I think this can be traced to John's childhood. When John misbehaved, his father (who is a professor of English) would punish John by forcing him to read a novel by Charles Dickens. Minor infractions resulted in A Christmas Carol being meted out, and when he was really bad he had to read Great Expectations. This is a true story which you should ask John about if you see him at a book-signing event.

93 of 94 people found the following review helpful. Reminds you that technical books can be insightful and fun to read.

By Jim Vallandingham

When I began to read the introduction for this book, after receiving it as a gift - I was a bit disheartened. I am not one of personas listed in the "Who Are You" section - a CEO or VP of an online startup, a beginner BI analyst. Instead, I am a software developer specializing in data visualization and data analysis. Furthermore, Excel is far from my preferred research tool of choice. I like code instead of screenshots. Python, Ruby, and R are where I turn when I want to look at data. *Even* with this mismatch of intended audience, I found myself engrossed in this book, reading it cover to cover in a few days. Data Smart is a wonderful resource. The use of Excel as a primary means for exploring data science concepts is surprisingly effective. It strips away all the code magic. You can't rely on SciKit-learn, or Weka, or even proper functions when all you have are cells and sheets. Instead, it provides a way for John Foreman to break down these complex concepts into the fundamental components that make them tick. You start to see the patterns between seemingly disparate technologies that are actually built off the same few bits of logic. Things start to click. The writing and real-world situations are really what make it fun and worth reading through and enjoying the ride. John's style hits the sweet spot between clarity and comical. Each chapter is well scoped. You understand the rationale behind why someone might want to use the particular tool being described to solve the problem at hand. The whimsy and flare added by the author moves the plot along at a good pace. The problems are simple enough to wrap your head around - but not toys. The datasets generated for this book must have taken a while to curate. The book is really fun to read. I think for me this book provides a great reminder of the landscape of data science tools, as well as a story-telling process to describe and relate these tools to non-programmy non-programmers. Even if you aren't a startup CEO... yet - this book is worth having on your shelf. Check it out today!

Data Science gets thrown around in the press like it's magic. Major retailers are predicting everything from when their customers are pregnant to when they want a new pair of Chuck Taylors. It's a brave new world where seemingly meaningless data can be transformed into valuable insight to drive smart business decisions. But how does one exactly do data science? Do you have to hire one of these priests of the dark arts, the "data scientist," to extract this gold from your data? Nope. Data science is little more than using straight-forward steps to process raw data into actionable insight. And in Data Smart, author and data scientist John Foreman will show you how that's done within the familiar environment of a spreadsheet. Why a spreadsheet? It's comfortable! You get to look at the data every step of the way, building confidence as you learn the tricks of the trade. Plus, spreadsheets are a vendor-neutral place to learn data science without the hype. But don't let the Excel sheets fool you. This is a book for those serious about learning the analytic techniques, the math and the magic, behind big data. Each chapter will cover a different technique in a spreadsheet so you can follow along: Mathematical optimization, including non-linear programming and genetic algorithms Clustering via k-means, spherical k-means, and graph modularity Data mining in graphs, such as outlier detection Supervised AI through logistic regression, ensemble models, and bag-of-words models Forecasting, seasonal adjustments, and prediction intervals through monte carlo simulation Moving from spreadsheets into the R programming language You get your hands dirty as you work alongside John through each technique. But never fear, the topics are readily applicable and the author laces humor throughout. You'll even learn what a dead squirrel has to do with optimization modeling, which you no doubt are dying to know.

From the Back Cover "Data Smart makes modern statistic methods and algorithms understandable and easy to implement. Slogging through textbooks and academic papers is no longer required!" —Patrick Crosby, Founder of StatHat first CTO at OkCupid "When Mr. Foreman interviewed for a job at my company, he arrived dressed in a 'Kentucky Colonel' kind of suit and spoke about nonsensical things like barbecue, lasers, and orange juice pulp. Then, he explained how to de-mystify and solve just about any complex 'big data' problem in our company with simple spreadsheets. No server clusters, mainframes, or Hadoop-a-ma-jigs. Just Excel. I hired him on the spot. After reading this book, you too will learn how to use math and basic spreadsheet formulas to improve your business or, at the very

least, how to trick senior executives into hiring you as their data scientist." mdash;Ben Chestnut, Founder CEO of MailChimp "You need a John Foreman on your analytics team. But if you can't have John, then reading this book is the next best thing." mdash;Patrick Lennon, Director of Analytics, The Coca-Cola Company Most people are approaching data science all wrong. Here's how to do it right. Not to disillusion you, but data scientists are not mystical practitioners of magical arts. Data science is something you can do. Really. This book shows you the significant data science techniques, how they work, how to use them, and how they benefit your business, large or small. It's not about coding or database technologies. It's about turning raw data into insight you can act upon, and doing it as quickly and painlessly as possible. Roll up your sleeves and let's get going. Relax mdash; it's just a spreadsheet Visit the companion website at www.wiley.com/go/datasmart to download spreadsheets for each chapter, and follow them as you learn about: Artificial intelligence using the general linear model, ensemble methods, and naive Bayes Clustering via k-means, spherical k-means, and graph modularity Mathematical optimization, including non-linear programming and genetic algorithms Working with time series data and forecasting with exponential smoothing Using Monte Carlo simulation to quantify and address risk Detecting outliers in single or multiple dimensions Exploring the data-science-focused R language About the AuthorJohn W. Foreman is Chief Data Scientist for MailChimp.com, where he leads a data science product development effort called the Email Genome Project. As an analytics consultant, John has created data science solutions for The Coca-Cola Company, Royal Caribbean International, Intercontinental Hotels Group, Dell, the Department of Defense, the IRS, and the FBI.