

Machine Learning For Hackers Drew Conway

If you ally dependence such a referred **Machine Learning For Hackers Drew Conway** books that will meet the expense of you worth, get the enormously best seller from us currently from several preferred authors. If you want to witty books, lots of novels, tale, jokes, and more fictions collections are afterward launched, from best seller to one of the most current released.

You may not be perplexed to enjoy every book collections Machine Learning For Hackers Drew Conway that we will very offer. It is not roughly speaking the costs. Its nearly what you habit currently. This Machine Learning For Hackers Drew Conway , as one of the most involved sellers here will unconditionally be among the best options to review.

[Hack This](#) - John Baichtal 2011-10-06

Presents instructions for creating and enhancing a variety of projects, including a sandwich-making robot, a Twitter-monitoring Christmas tree, and a bronze-melting blast furnace.

R for Everyone - Jared P. Lander 2017-06-13

Statistical Computation for Programmers, Scientists, Quants, Excel Users, and Other Professionals Using the open source R language, you can build powerful statistical models to answer many of your most challenging questions. R has traditionally been difficult for non-statisticians to learn, and most R books assume far too much knowledge to be of help. R for Everyone, Second Edition, is the solution. Drawing on his unsurpassed experience teaching new users, professional data scientist Jared P. Lander has written the perfect tutorial for anyone new to statistical programming and modeling. Organized to make learning easy and intuitive, this guide focuses on the 20 percent of R functionality you'll need to accomplish 80 percent of modern data tasks. Lander's self-contained chapters start with the absolute basics, offering extensive hands-on practice and sample code. You'll download and install R; navigate and use the R environment; master basic program control, data import, manipulation, and visualization; and walk through several essential tests. Then, building on this foundation, you'll construct several complete models, both linear and nonlinear, and use some data mining techniques. After all this you'll make your code reproducible with LaTeX, RMarkdown, and Shiny. By the time you're done, you won't just know how to write R programs, you'll be ready to tackle the statistical problems you care about most. Coverage includes Explore R, RStudio, and R packages Use R for math: variable types, vectors, calling functions, and more Exploit data structures, including data.frames, matrices, and lists Read many different types of data Create attractive, intuitive statistical graphics Write user-defined functions Control program flow with if, ifelse, and complex checks Improve program efficiency with group manipulations Combine and reshape multiple datasets Manipulate strings using R's facilities and regular expressions Create normal, binomial, and Poisson probability distributions Build linear, generalized linear, and nonlinear models Program basic statistics: mean, standard deviation, and t-tests Train machine learning models Assess the quality of models and variable selection Prevent overfitting and perform variable selection, using the Elastic Net and Bayesian methods Analyze univariate and multivariate time series data Group data via K-means and hierarchical clustering Prepare reports, slideshows, and web pages with knitr Display interactive data with RMarkdown and htmlwidgets Implement dashboards with Shiny Build reusable R packages with devtools and Rcpp Register your product at informit.com/register for convenient access to downloads, updates, and corrections as they become available.

Business unIntelligence - Dr. Barry Devlin 2013-10-01

Business intelligence (BI) used to be so simple—in theory anyway. Integrate and copy data from your transactional systems into a specialized relational database, apply BI reporting and query tools and add business users. Job done. No longer. Analytics, big data and an array of diverse technologies have changed everything. More importantly, business is insisting on ever more value, ever faster from information and from IT in general. An emerging biz-tech ecosystem demands that business and IT work together. Business unIntelligence reflects the new reality that in today's socially complex and rapidly changing world, business decisions must be based on a combination of rational and intuitive thinking. Integrating cues from diverse information sources and tacit knowledge, decision makers create unique meaning to innovate heuristically

at the speed of thought. This book provides a wealth of new models that business and IT can use together to design support systems for tomorrow's successful organizations. Dr. Barry Devlin, one of the earliest proponents of data warehousing, goes back to basics to explore how the modern trinity of information, process and people must be reinvented and restructured to deliver the value, insight and innovation required by modern businesses. From here, he develops a series of novel architectural models that provide a new foundation for holistic information use across the entire business. From discovery to analysis and from decision making to action taking, he defines a fully integrated, closed-loop business environment. Covering every aspect of business analytics, big data, collaborative working and more, this book takes over where BI ends to deliver the definitive framework for information use in the coming years. As the person who defined the conceptual framework and physical architecture for data warehousing in the 1980s, Barry Devlin has been an astute observer of the movement he initiated ever since. Now, in Business unIntelligence, Devlin provides a sweeping view of the past, present, and future of business intelligence, while delivering new conceptual and physical models for how to turn information into insights and action. Reading Devlin's prose and vision of BI are comparable to reading Carl Sagan's view of the cosmos. The book is truly illuminating and inspiring. --Wayne Eckerson, President, BI Leader Consulting Author, "Secrets of Analytical Leaders: Insights from Information Insiders"

[Democratizing Innovation](#) - Eric Von Hippel 2006-02-17

The process of user-centered innovation: how it can benefit both users and manufacturers and how its emergence will bring changes in business models and in public policy. Innovation is rapidly becoming democratized. Users, aided by improvements in computer and communications technology, increasingly can develop their own new products and services. These innovating users—both individuals and firms—often freely share their innovations with others, creating user-innovation communities and a rich intellectual commons. In Democratizing Innovation, Eric von Hippel looks closely at this emerging system of user-centered innovation. He explains why and when users find it profitable to develop new products and services for themselves, and why it often pays users to reveal their innovations freely for the use of all. The trend toward democratized innovation can be seen in software and information products—most notably in the free and open-source software movement—but also in physical products. Von Hippel's many examples of user innovation in action range from surgical equipment to surfboards to software security features. He shows that product and service development is concentrated among "lead users," who are ahead on marketplace trends and whose innovations are often commercially attractive. Von Hippel argues that manufacturers should redesign their innovation processes and that they should systematically seek out innovations developed by users. He points to businesses—the custom semiconductor industry is one example—that have learned to assist user-innovators by providing them with toolkits for developing new products. User innovation has a positive impact on social welfare, and von Hippel proposes that government policies, including R&D subsidies and tax credits, should be realigned to eliminate biases against it. The goal of a democratized user-centered innovation system, says von Hippel, is well worth striving for. An electronic version of this book is available under a Creative Commons license.

Mastering Machine Learning with Python in Six Steps - Manohar Swamynathan 2019-10-01

Explore fundamental to advanced Python 3 topics in six steps, all designed to make you a worthy practitioner. This updated version's approach is based on the "six degrees of separation" theory, which

states that everyone and everything is a maximum of six steps away and presents each topic in two parts: theoretical concepts and practical implementation using suitable Python 3 packages. You'll start with the fundamentals of Python 3 programming language, machine learning history, evolution, and the system development frameworks. Key data mining/analysis concepts, such as exploratory analysis, feature dimension reduction, regressions, time series forecasting and their efficient implementation in Scikit-learn are covered as well. You'll also learn commonly used model diagnostic and tuning techniques. These include optimal probability cutoff point for class creation, variance, bias, bagging, boosting, ensemble voting, grid search, random search, Bayesian optimization, and the noise reduction technique for IoT data. Finally, you'll review advanced text mining techniques, recommender systems, neural networks, deep learning, reinforcement learning techniques and their implementation. All the code presented in the book will be available in the form of iPython notebooks to enable you to try out these examples and extend them to your advantage. What You'll Learn Understand machine learning development and frameworks Assess model diagnosis and tuning in machine learning Examine text mining, natural language processing (NLP), and recommender systems Review reinforcement learning and CNN Who This Book Is For Python developers, data engineers, and machine learning engineers looking to expand their knowledge or career into machine learning area.

Fundamentals of Deep Learning - Nikhil Buduma 2017-05-25

With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks Learn how to train feed-forward neural networks Use TensorFlow to implement your first neural network Manage problems that arise as you begin to make networks deeper Build neural networks that analyze complex images Perform effective dimensionality reduction using autoencoders Dive deep into sequence analysis to examine language Learn the fundamentals of reinforcement learning

Natural Language Processing for the Semantic Web - Diana Maynard 2016-12-13

This book introduces core natural language processing (NLP) technologies to non-experts in an easily accessible way, as a series of building blocks that lead the user to understand key technologies, why they are required, and how to integrate them into Semantic Web applications. Natural language processing and Semantic Web technologies have different, but complementary roles in data management. Combining these two technologies enables structured and unstructured data to merge seamlessly. Semantic Web technologies aim to convert unstructured data to meaningful representations, which benefit enormously from the use of NLP technologies, thereby enabling applications such as connecting text to Linked Open Data, connecting texts to each other, semantic searching, information visualization, and modeling of user behavior in online networks. The first half of this book describes the basic NLP processing tools: tokenization, part-of-speech tagging, and morphological analysis, in addition to the main tools required for an information extraction system (named entity recognition and relation extraction) which build on these components. The second half of the book explains how Semantic Web and NLP technologies can enhance each other, for example via semantic annotation, ontology linking, and population. These chapters also discuss sentiment analysis, a key component in making sense of textual data, and the difficulties of performing NLP on social media, as well as some proposed solutions. The book finishes by investigating some applications of these tools, focusing on semantic search and visualization, modeling user behavior, and an outlook on the future.

Machine Learning for Hackers - Drew Conway 2012-02-15

Presents algorithms that enable computers to train themselves to automate tasks, focusing on specific problems such as prediction, optimization, and classification.

Machine Learning for Kids - Dale Lane 2021-01-19

A hands-on, application-based introduction to machine learning and artificial intelligence (AI) that guides

young readers through creating compelling AI-powered games and applications using the Scratch programming language. Machine learning (also known as ML) is one of the building blocks of AI, or artificial intelligence. AI is based on the idea that computers can learn on their own, with your help. Machine Learning for Kids will introduce you to machine learning, painlessly. With this book and its free, Scratch-based, award-winning companion website, you'll see how easy it is to add machine learning to your own projects. You don't even need to know how to code! As you work through the book you'll discover how machine learning systems can be taught to recognize text, images, numbers, and sounds, and how to train your models to improve their accuracy. You'll turn your models into fun computer games and apps, and see what happens when they get confused by bad data. You'll build 13 projects step-by-step from the ground up, including:

- Rock, Paper, Scissors game that recognizes your hand shapes
- An app that recommends movies based on other movies that you like
- A computer character that reacts to insults and compliments
- An interactive virtual assistant (like Siri or Alexa) that obeys commands
- An AI version of Pac-Man, with a smart character that knows how to avoid ghosts

NOTE: This book includes a Scratch tutorial for beginners, and step-by-step instructions for every project. Ages 12+

Natural Language Annotation for Machine Learning - James Pustejovsky 2012-10-11

Create your own natural language training corpus for machine learning. Whether you're working with English, Chinese, or any other natural language, this hands-on book guides you through a proven annotation development cycle—the process of adding metadata to your training corpus to help ML algorithms work more efficiently. You don't need any programming or linguistics experience to get started. Using detailed examples at every step, you'll learn how the MATTER Annotation Development Process helps you Model, Annotate, Train, Test, Evaluate, and Revise your training corpus. You also get a complete walkthrough of a real-world annotation project. Define a clear annotation goal before collecting your dataset (corpus) Learn tools for analyzing the linguistic content of your corpus Build a model and specification for your annotation project Examine the different annotation formats, from basic XML to the Linguistic Annotation Framework Create a gold standard corpus that can be used to train and test ML algorithms Select the ML algorithms that will process your annotated data Evaluate the test results and revise your annotation task Learn how to use lightweight software for annotating texts and adjudicating the annotations This book is a perfect companion to O'Reilly's Natural Language Processing with Python.

The Cathedral & the Bazaar - Eric S. Raymond 2001-02-01

Open source provides the competitive advantage in the Internet Age. According to the August Forrester Report, 56 percent of IT managers interviewed at Global 2,500 companies are already using some type of open source software in their infrastructure and another 6 percent will install it in the next two years. This revolutionary model for collaborative software development is being embraced and studied by many of the biggest players in the high-tech industry, from Sun Microsystems to IBM to Intel. The Cathedral & the Bazaar is a must for anyone who cares about the future of the computer industry or the dynamics of the information economy. Already, billions of dollars have been made and lost based on the ideas in this book. Its conclusions will be studied, debated, and implemented for years to come. According to Bob Young, "This is Eric Raymond's great contribution to the success of the open source revolution, to the adoption of Linux-based operating systems, and to the success of open source users and the companies that supply them." The interest in open source software development has grown enormously in the past year. This revised and expanded paperback edition includes new material on open source developments in 1999 and 2000. Raymond's clear and effective writing style accurately describing the benefits of open source software has been key to its success. With major vendors creating acceptance for open source within companies, independent vendors will become the open source story in 2001.

Numerical Python - Robert Johansson 2018-12-24

Leverage the numerical and mathematical modules in Python and its standard library as well as popular open source numerical Python packages like NumPy, SciPy, FiPy, matplotlib and more. This fully revised edition, updated with the latest details of each package and changes to Jupyter projects, demonstrates how to numerically compute solutions and mathematically model applications in big data, cloud computing, financial engineering, business management and more. Numerical Python, Second Edition, presents many brand-new case study examples of applications in data science and statistics using Python, along with

extensions to many previous examples. Each of these demonstrates the power of Python for rapid development and exploratory computing due to its simple and high-level syntax and multiple options for data analysis. After reading this book, readers will be familiar with many computing techniques including array-based and symbolic computing, visualization and numerical file I/O, equation solving, optimization, interpolation and integration, and domain-specific computational problems, such as differential equation solving, data analysis, statistical modeling and machine learning. What You'll Learn Work with vectors and matrices using NumPy Plot and visualize data with Matplotlib Perform data analysis tasks with Pandas and SciPy Review statistical modeling and machine learning with statsmodels and scikit-learn Optimize Python code using Numba and Cython Who This Book Is For Developers who want to understand how to use Python and its related ecosystem for numerical computing.

Python for Data Analysis - Wes McKinney 2017-09-25

Get complete instructions for manipulating, processing, cleaning, and crunching datasets in Python. Updated for Python 3.6, the second edition of this hands-on guide is packed with practical case studies that show you how to solve a broad set of data analysis problems effectively. You'll learn the latest versions of pandas, NumPy, IPython, and Jupyter in the process. Written by Wes McKinney, the creator of the Python pandas project, this book is a practical, modern introduction to data science tools in Python. It's ideal for analysts new to Python and for Python programmers new to data science and scientific computing. Data files and related material are available on GitHub. Use the IPython shell and Jupyter notebook for exploratory computing Learn basic and advanced features in NumPy (Numerical Python) Get started with data analysis tools in the pandas library Use flexible tools to load, clean, transform, merge, and reshape data Create informative visualizations with matplotlib Apply the pandas groupby facility to slice, dice, and summarize datasets Analyze and manipulate regular and irregular time series data Learn how to solve real-world data analysis problems with thorough, detailed examples

Learning OpenCV 3 - Adrian Kaehler 2016-12-14

"This book provides a working guide to the C++ Open Source Computer Vision Library (OpenCV) version 3.x and gives a general background on the field of computer vision sufficient to help readers use OpenCV effectively."--Preface.

Envisioning Information - Edward R. Tufte 1990

Escaoping flatland. Micro/Macro readings. Layering and separation. Small multiples. Color and information. Narratives of Space and time. Epilogue.

Doing Data Science - Cathy O'Neil 2013-10-09

Now that people are aware that data can make the difference in an election or a business model, data science as an occupation is gaining ground. But how can you get started working in a wide-ranging, interdisciplinary field that's so clouded in hype? This insightful book, based on Columbia University's Introduction to Data Science class, tells you what you need to know. In many of these chapter-long lectures, data scientists from companies such as Google, Microsoft, and eBay share new algorithms, methods, and models by presenting case studies and the code they use. If you're familiar with linear algebra, probability, and statistics, and have programming experience, this book is an ideal introduction to data science. Topics include: Statistical inference, exploratory data analysis, and the data science process Algorithms Spam filters, Naive Bayes, and data wrangling Logistic regression Financial modeling Recommendation engines and causality Data visualization Social networks and data journalism Data engineering, MapReduce, Pregel, and Hadoop Doing Data Science is collaboration between course instructor Rachel Schutt, Senior VP of Data Science at News Corp, and data science consultant Cathy O'Neil, a senior data scientist at Johnson Research Labs, who attended and blogged about the course.

Machine Learning for Absolute Beginners - Oliver Theobald 2018

"The manner in which computers are now able to mimic human thinking to process information is rapidly exceeding human capabilities in everything from chess to picking the winner of a song contest. In the modern age of machine learning, computers do not strictly need to receive an 'input command' to perform a task, but rather 'input data'. From the input of data they are able to form their own decisions and take actions virtually as a human world. But given it is a machine, it can consider many more scenarios and execute far more complicated calculations to solve complex problems. This is the element that excites data

scientists and machine learning engineers the most. The ability to solve complex problems never before attempted. This book will dive in to introduce machine learning, and is ideal for beginners starting out in machine learning."--page 4 of cover.

How Smart Machines Think - Sean Gerrish 2019-10-22

Everything you've always wanted to know about self-driving cars, Netflix recommendations, IBM's Watson, and video game-playing computer programs. The future is here: Self-driving cars are on the streets, an algorithm gives you movie and TV recommendations, IBM's Watson triumphed on Jeopardy over puny human brains, computer programs can be trained to play Atari games. But how do all these things work? In this book, Sean Gerrish offers an engaging and accessible overview of the breakthroughs in artificial intelligence and machine learning that have made today's machines so smart. Gerrish outlines some of the key ideas that enable intelligent machines to perceive and interact with the world. He describes the software architecture that allows self-driving cars to stay on the road and to navigate crowded urban environments; the million-dollar Netflix competition for a better recommendation engine (which had an unexpected ending); and how programmers trained computers to perform certain behaviors by offering them treats, as if they were training a dog. He explains how artificial neural networks enable computers to perceive the world—and to play Atari video games better than humans. He explains Watson's famous victory on Jeopardy, and he looks at how computers play games, describing AlphaGo and Deep Blue, which beat reigning world champions at the strategy games of Go and chess. Computers have not yet mastered everything, however; Gerrish outlines the difficulties in creating intelligent agents that can successfully play video games like StarCraft that have evaded solution—at least for now. Gerrish weaves the stories behind these breakthroughs into the narrative, introducing readers to many of the researchers involved, and keeping technical details to a minimum. Science and technology buffs will find this book an essential guide to a future in which machines can outsmart people.

Feature Engineering for Machine Learning - Alice Zheng 2018-03-23

Feature engineering is a crucial step in the machine-learning pipeline, yet this topic is rarely examined on its own. With this practical book, you'll learn techniques for extracting and transforming features—the numeric representations of raw data—into formats for machine-learning models. Each chapter guides you through a single data problem, such as how to represent text or image data. Together, these examples illustrate the main principles of feature engineering. Rather than simply teach these principles, authors Alice Zheng and Amanda Casari focus on practical application with exercises throughout the book. The closing chapter brings everything together by tackling a real-world, structured dataset with several feature-engineering techniques. Python packages including numpy, Pandas, Scikit-learn, and Matplotlib are used in code examples. You'll examine: Feature engineering for numeric data: filtering, binning, scaling, log transforms, and power transforms Natural text techniques: bag-of-words, n-grams, and phrase detection Frequency-based filtering and feature scaling for eliminating uninformative features Encoding techniques of categorical variables, including feature hashing and bin-counting Model-based feature engineering with principal component analysis The concept of model stacking, using k-means as a featurization technique Image feature extraction with manual and deep-learning techniques

Practical Machine Learning with Python - Dipanjan Sarkar 2017-12-20

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering

and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

Machine Learning with R - Brett Lantz 2013-10-25

Written as a tutorial to explore and understand the power of R for machine learning. This practical guide that covers all of the need to know topics in a very systematic way. For each machine learning approach, each step in the process is detailed, from preparing the data for analysis to evaluating the results. These steps will build the knowledge you need to apply them to your own data science tasks. Intended for those who want to learn how to use R's machine learning capabilities and gain insight from your data. Perhaps you already know a bit about machine learning, but have never used R; or perhaps you know a little R but are new to machine learning. In either case, this book will get you up and running quickly. It would be helpful to have a bit of familiarity with basic programming concepts, but no prior experience is required.

Programming Collective Intelligence - Toby Segaran 2007-08-16

Want to tap the power behind search rankings, product recommendations, social bookmarking, and online matchmaking? This fascinating book demonstrates how you can build Web 2.0 applications to mine the enormous amount of data created by people on the Internet. With the sophisticated algorithms in this book, you can write smart programs to access interesting datasets from other web sites, collect data from users of your own applications, and analyze and understand the data once you've found it. Programming Collective Intelligence takes you into the world of machine learning and statistics, and explains how to draw conclusions about user experience, marketing, personal tastes, and human behavior in general -- all from information that you and others collect every day. Each algorithm is described clearly and concisely with code that can immediately be used on your web site, blog, Wiki, or specialized application. This book explains: Collaborative filtering techniques that enable online retailers to recommend products or media Methods of clustering to detect groups of similar items in a large dataset Search engine features -- crawlers, indexers, query engines, and the PageRank algorithm Optimization algorithms that search millions of possible solutions to a problem and choose the best one Bayesian filtering, used in spam filters for classifying documents based on word types and other features Using decision trees not only to make predictions, but to model the way decisions are made Predicting numerical values rather than classifications to build price models Support vector machines to match people in online dating sites Non-negative matrix factorization to find the independent features in a dataset Evolving intelligence for problem solving -- how a computer develops its skill by improving its own code the more it plays a game Each chapter includes exercises for extending the algorithms to make them more powerful. Go beyond simple database-backed applications and put the wealth of Internet data to work for you. "Bravo! I cannot think of a better way for a developer to first learn these algorithms and methods, nor can I think of a better way for me (an old AI dog) to reinvigorate my knowledge of the details." -- Dan Russell, Google "Toby's book does a great job of breaking down the complex subject matter of machine-learning algorithms into practical, easy-to-understand examples that can be directly applied to analysis of social interaction across the Web today. If I had this book two years ago, it would have saved precious time going down some fruitless paths." -- Tim Wolters, CTO, Collective Intellect

Deep Learning - Josh Patterson 2017-07-28

Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their open-source DeepLearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J. Dive into machine learning concepts in general, as well as deep learning in particular Understand how deep networks evolved from neural network fundamentals Explore the major deep network architectures, including Convolutional and Recurrent Learn how to map specific deep networks to the right problem Walk through the fundamentals of tuning general neural networks and specific deep network architectures Use vectorization techniques for different data types with DataVec, DL4J's workflow tool Learn how to use DL4J natively on Spark and Hadoop

The Artist in the Machine - Arthur I. Miller 2020-11-10

An authority on creativity introduces us to AI-powered computers that are creating art, literature, and music that may well surpass the creations of humans. Today's computers are composing music that sounds "more Bach than Bach," turning photographs into paintings in the style of Van Gogh's Starry Night, and even writing screenplays. But are computers truly creative—or are they merely tools to be used by musicians, artists, and writers? In this book, Arthur I. Miller takes us on a tour of creativity in the age of machines. Miller, an authority on creativity, identifies the key factors essential to the creative process, from "the need for introspection" to "the ability to discover the key problem." He talks to people on the cutting edge of artificial intelligence, encountering computers that mimic the brain and machines that have defeated champions in chess, Jeopardy!, and Go. In the central part of the book, Miller explores the riches of computer-created art, introducing us to artists and computer scientists who have, among much else, unleashed an artificial neural network to create a nightmarish, multi-eyed dog-cat; taught AI to imagine; developed a robot that paints; created algorithms for poetry; and produced the world's first computer-composed musical, Beyond the Fence, staged by Android Lloyd Webber and friends. But, Miller writes, in order to be truly creative, machines will need to step into the world. He probes the nature of consciousness and speaks to researchers trying to develop emotions and consciousness in computers. Miller argues that computers can already be as creative as humans—and someday will surpass us. But this is not a dystopian account; Miller celebrates the creative possibilities of artificial intelligence in art, music, and literature.

Applied Natural Language Processing with Python - Taweh Beysolow II 2018-09-11

Learn to harness the power of AI for natural language processing, performing tasks such as spell check, text summarization, document classification, and natural language generation. Along the way, you will learn the skills to implement these methods in larger infrastructures to replace existing code or create new algorithms. Applied Natural Language Processing with Python starts with reviewing the necessary machine learning concepts before moving onto discussing various NLP problems. After reading this book, you will have the skills to apply these concepts in your own professional environment. What You Will Learn Utilize various machine learning and natural language processing libraries such as TensorFlow, Keras, NLTK, and Gensim Manipulate and preprocess raw text data in formats such as .txt and .pdf Strengthen your skills in data science by learning both the theory and the application of various algorithms Who This Book Is For You should be at least a beginner in ML to get the most out of this text, but you needn't feel that you need be an expert to understand the content.

Bandit Algorithms for Website Optimization - John White 2013

When looking for ways to improve your website, how do you decide which changes to make? And which changes to keep? This concise book shows you how to use Multiarmed Bandit algorithms to measure the real-world value of any modifications you make to your site. Author John Myles White shows you how this powerful class of algorithms can help you boost website traffic, convert visitors to customers, and increase many other measures of success. This is the first developer-focused book on bandit algorithms, which were previously described only in research papers. You'll quickly learn the benefits of several simple algorithms—including the epsilon-Greedy, Softmax, and Upper Confidence Bound (UCB) algorithms—by

working through code examples written in Python, which you can easily adapt for deployment on your own website. Learn the basics of A/B testing—and recognize when it's better to use bandit algorithms Develop a unit testing framework for debugging bandit algorithms Get additional code examples written in Julia, Ruby, and JavaScript with supplemental online materials

Hacking the Hacker - Roger A. Grimes 2017-04-18

Meet the world's top ethical hackers and explore the tools of the trade Hacking the Hacker takes you inside the world of cybersecurity to show you what goes on behind the scenes, and introduces you to the men and women on the front lines of this technological arms race. Twenty-six of the world's top white hat hackers, security researchers, writers, and leaders, describe what they do and why, with each profile preceded by a no-experience-necessary explanation of the relevant technology. Dorothy Denning discusses advanced persistent threats, Martin Hellman describes how he helped invent public key encryption, Bill Cheswick talks about firewalls, Dr. Charlie Miller talks about hacking cars, and other cybersecurity experts from around the world detail the threats, their defenses, and the tools and techniques they use to thwart the most advanced criminals history has ever seen. Light on jargon and heavy on intrigue, this book is designed to be an introduction to the field; final chapters include a guide for parents of young hackers, as well as the Code of Ethical Hacking to help you start your own journey to the top. Cybersecurity is becoming increasingly critical at all levels, from retail businesses all the way up to national security. This book drives to the heart of the field, introducing the people and practices that help keep our world secure. Go deep into the world of white hat hacking to grasp just how critical cybersecurity is Read the stories of some of the world's most renowned computer security experts Learn how hackers do what they do—no technical expertise necessary Delve into social engineering, cryptography, penetration testing, network attacks, and more As a field, cybersecurity is large and multi-faceted—yet not historically diverse. With a massive demand for qualified professional that is only going to grow, opportunities are endless. Hacking the Hacker shows you why you should give the field a closer look.

Principles of Strategic Data Science - Dr Peter Prevos 2019-06-03

Take the strategic and systematic approach to analyze data to solve business problems Key FeaturesGain detailed information about the theory of data scienceAugment your coding knowledge with practical data science techniques for efficient data analysis Learn practical ways to strategically and systematically use dataBook Description Principles of Strategic Data Science is created to help you join the dots between mathematics, programming, and business analysis. With a unique approach that bridges the gap between mathematics and computer science, this book takes you through the entire data science pipeline. The book begins by explaining what data science is and how organizations can use it to revolutionize the way they use their data. It then discusses the criteria for the soundness of data products and how to best visualize information. As you progress, you'll discover the strategic aspects of data science by learning the five-phase framework that enables you to enhance the value you extract from data. The final chapter of the book discusses the role of a data science manager in helping an organization take the data-driven approach. By the end of this book, you'll have a good understanding of data science and how it can enable you to extract value from your data. What you will learnGet familiar with the five most important steps of data scienceUse the Conway diagram to visualize the technical skills of the data science teamUnderstand the limitations of data science from a mathematical and ethical perspectiveGet a quick overview of machine learningGain insight into the purpose of using data science in your workUnderstand the role of data science managers and their expectationsWho this book is for This book is ideal for data scientists and data analysts who are looking for a practical guide to strategically and systematically use data. This book is also useful for those who want to understand in detail what is data science and how can an organization take the data-driven approach. Prior programming knowledge of Python and R is assumed.

What Technology Wants - Kevin Kelly 2011-09-27

From the author of the New York Times bestseller *The Inevitable*— a sweeping vision of technology as a living force that can expand our individual potential In this provocative book, one of today's most respected thinkers turns the conversation about technology on its head by viewing technology as a natural system, an extension of biological evolution. By mapping the behavior of life, we paradoxically get a glimpse at where technology is headed-or "what it wants." Kevin Kelly offers a dozen trajectories in the coming decades for

this near-living system. And as we align ourselves with technology's agenda, we can capture its colossal potential. This visionary and optimistic book explores how technology gives our lives greater meaning and is a must-read for anyone curious about the future.

Machine Learning in Action - Peter Harrington 2012-04-03

Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce

Hackers - Steven Levy 2010-05-19

This 25th anniversary edition of Steven Levy's classic book traces the exploits of the computer revolution's original hackers -- those brilliant and eccentric nerds from the late 1950s through the early '80s who took risks, bent the rules, and pushed the world in a radical new direction. With updated material from noteworthy hackers such as Bill Gates, Mark Zuckerberg, Richard Stallman, and Steve Wozniak, *Hackers* is a fascinating story that begins in early computer research labs and leads to the first home computers. Levy profiles the imaginative brainiacs who found clever and unorthodox solutions to computer engineering problems. They had a shared sense of values, known as "the hacker ethic," that still thrives today. *Hackers* captures a seminal period in recent history when underground activities blazed a trail for today's digital world, from MIT students finagling access to clunky computer-card machines to the DIY culture that spawned the Altair and the Apple II.

Thoughtful Machine Learning - Matthew Kirk 2014-09-26

Learn how to apply test-driven development (TDD) to machine-learning algorithms—and catch mistakes that could sink your analysis. In this practical guide, author Matthew Kirk takes you through the principles of TDD and machine learning, and shows you how to apply TDD to several machine-learning algorithms, including Naive Bayesian classifiers and Neural Networks. Machine-learning algorithms often have tests baked in, but they can't account for human errors in coding. Rather than blindly rely on machine-learning results as many researchers have, you can mitigate the risk of errors with TDD and write clean, stable machine-learning code. If you're familiar with Ruby 2.1, you're ready to start. Apply TDD to write and run tests before you start coding Learn the best uses and tradeoffs of eight machine learning algorithms Use real-world examples to test each algorithm through engaging, hands-on exercises Understand the similarities between TDD and the scientific method for validating solutions Be aware of the risks of machine learning, such as underfitting and overfitting data Explore techniques for improving your machine-learning

models or data extraction

The Mueller Report - Robert S. Mueller 2019-04-26

This is the full Mueller Report, as released on April 18, 2019, by the U.S. Department of Justice. A reprint of the report exactly as it was issued by the government, it is without analysis or commentary from any other source and with nothing subtracted except for the material redacted by the Department of Justice. The mission of the Mueller investigation was to examine Russian interference in the 2016 Presidential election, consisting of possible links, or "collusion," between the Donald Trump campaign and the Russian government of Vladimir Putin as well as any allegations of obstruction of justice in this regard. It was also intended to detect and prosecute, where warranted, any other crimes that surfaced during the course of the investigation. The report consists of a detailed summary of the various investigations and inquiries that the Special Counsel and colleagues carried out in these areas. The investigation was initiated in the aftermath of the firing of FBI Director James Comey by Donald Trump on May 9, 2017. The FBI, under Director Comey, had already been investigating links between Russia and the Trump campaign. Mueller submitted his report to Attorney General William Barr on March 22, 2019, and the Department of Justice released the redacted report one month later.

Learning to Love Data Science - Mike Barlow 2015-10-27

Until recently, many people thought big data was a passing fad. "Data science" was an enigmatic term. Today, big data is taken seriously, and data science is considered downright sexy. With this anthology of reports from award-winning journalist Mike Barlow, you'll appreciate how data science is fundamentally altering our world, for better and for worse. Barlow paints a picture of the emerging data space in broad strokes. From new techniques and tools to the use of data for social good, you'll find out how far data science reaches. With this anthology, you'll learn how: Analysts can now get results from their data queries in near real time Indie manufacturers are blurring the lines between hardware and software Companies try to balance their desire for rapid innovation with the need to tighten data security Advanced analytics and low-cost sensors are transforming equipment maintenance from a cost center to a profit center CIOs have gradually evolved from order takers to business innovators New analytics tools let businesses go beyond data analysis and straight to decision-making Mike Barlow is an award-winning journalist, author, and communications strategy consultant. Since launching his own firm, Cumulus Partners, he has represented major organizations in a number of industries.

AI and Machine Learning for Coders - Laurence Moroney 2020-10-01

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow Serving

Machine Learning for Hackers - Drew Conway 2012-02-13

If you're an experienced programmer interested in crunching data, this book will get you started with machine learning—a toolkit of algorithms that enables computers to train themselves to automate useful tasks. Authors Drew Conway and John Myles White help you understand machine learning and statistics tools through a series of hands-on case studies, instead of a traditional math-heavy presentation. Each chapter focuses on a specific problem in machine learning, such as classification, prediction, optimization, and recommendation. Using the R programming language, you'll learn how to analyze sample datasets and write simple machine learning algorithms. Machine Learning for Hackers is ideal for programmers from any background, including business, government, and academic research. Develop a naïve Bayesian classifier to determine if an email is spam, based only on its text Use linear regression to predict the

number of page views for the top 1,000 websites Learn optimization techniques by attempting to break a simple letter cipher Compare and contrast U.S. Senators statistically, based on their voting records Build a "whom to follow" recommendation system from Twitter data

Machine Learning For Dummies - John Paul Mueller 2021-02-09

One of Mark Cuban's top reads for better understanding A.I. (inc.com, 2021) Your comprehensive entry-level guide to machine learning While machine learning expertise doesn't quite mean you can create your own Turing Test-proof android—as in the movie *Ex Machina*—it is a form of artificial intelligence and one of the most exciting technological means of identifying opportunities and solving problems fast and on a large scale. Anyone who masters the principles of machine learning is mastering a big part of our tech future and opening up incredible new directions in careers that include fraud detection, optimizing search results, serving real-time ads, credit-scoring, building accurate and sophisticated pricing models—and way, way more. Unlike most machine learning books, the fully updated 2nd Edition of *Machine Learning For Dummies* doesn't assume you have years of experience using programming languages such as Python (R source is also included in a downloadable form with comments and explanations), but lets you in on the ground floor, covering the entry-level materials that will get you up and running building models you need to perform practical tasks. It takes a look at the underlying—and fascinating—math principles that power machine learning but also shows that you don't need to be a math whiz to build fun new tools and apply them to your work and study. Understand the history of AI and machine learning Work with Python 3.8 and TensorFlow 2.x (and R as a download) Build and test your own models Use the latest datasets, rather than the worn out data found in other books Apply machine learning to real problems Whether you want to learn for college or to enhance your business or career performance, this friendly beginner's guide is your best introduction to machine learning, allowing you to become quickly confident using this amazing and fast-developing technology that's impacting lives for the better all over the world.

Machine Learning for Email - Drew Conway 2011-10-27

This compact book explores standard tools for text classification, and teaches the reader how to use machine learning to decide whether a e-mail is spam or ham (binary classification), based on raw data from The SpamAssassin Public Corpus. Of course, sometimes the items in one class are not created equally, or we want to distinguish among them in some meaningful way. The second part of the book will look at how to not only filter spam from our email, but also placing "more important" messages at the top of the queue. This is a curated excerpt from the upcoming book "Machine Learning for Hackers."

What Is Data Science? - Mike Loukides 2011-04-10

We've all heard it: according to Hal Varian, statistics is the next sexy job. Five years ago, in *What is Web 2.0*, Tim O'Reilly said that "data is the next Intel Inside." But what does that statement mean? Why do we suddenly care about statistics and about data? This report examines the many sides of data science -- the technologies, the companies and the unique skill sets. The web is full of "data-driven apps." Almost any e-commerce application is a data-driven application. There's a database behind a web front end, and middleware that talks to a number of other databases and data services (credit card processing companies, banks, and so on). But merely using data isn't really what we mean by "data science." A data application acquires its value from the data itself, and creates more data as a result. It's not just an application with data; it's a data product. Data science enables the creation of data products.

Machine Learning and Security - Clarence Chio 2018-01-26

Can machine learning techniques solve our computer security problems and finally put an end to the cat-and-mouse game between attackers and defenders? Or is this hope merely hype? Now you can dive into the science and answer this question for yourself! With this practical guide, you'll explore ways to apply machine learning to security issues such as intrusion detection, malware classification, and network analysis. Machine learning and security specialists Clarence Chio and David Freeman provide a framework for discussing the marriage of these two fields, as well as a toolkit of machine-learning algorithms that you can apply to an array of security problems. This book is ideal for security engineers and data scientists alike. Learn how machine learning has contributed to the success of modern spam filters Quickly detect anomalies, including breaches, fraud, and impending system failure Conduct malware analysis by extracting useful information from computer binaries Uncover attackers within the network by finding

patterns inside datasets Examine how attackers exploit consumer-facing websites and app functionality

Translate your machine learning algorithms from the lab to production Understand the threat attackers pose to machine learning solutions