

Machine Learning For Hackers Drew Conway

Get up and running, and productive in no time with MLflow using the most effective machine learning engineering approach

Key Features Explore machine learning workflows for stating ML problems in a concise and clear manner using MLflow Use MLflow to iteratively develop a ML model and manage it Discover and work with the features available in MLflow to seamlessly take a model from the development phase to a production environment

Book Description MLflow is a platform for the machine learning life cycle that enables structured development and iteration of machine learning models and a seamless transition into scalable production environments. This book will take you through the different features of MLflow and how you can implement them in your ML project. You will begin by framing an ML problem and then transform your solution with MLflow, adding a workbench environment, training infrastructure, data management, model management, experimentation, and state-of-the-art ML deployment techniques on the cloud and premises. The book also explores techniques to scale up your workflow as well as performance monitoring techniques. As you progress, you'll discover how to create an operational dashboard to manage machine learning systems. Later, you will learn how you can use MLflow in the AutoML, anomaly detection, and deep learning context with the help of use cases. In addition to this, you will understand how to use machine learning platforms for local development as well as for cloud and managed environments. This book will also show you how to use MLflow in non-Python-based languages such as R and Java, along with covering approaches to extend MLflow with Plugins. By the end of this machine learning book, you will be able to produce and deploy reliable machine learning algorithms using MLflow in multiple environments. What you will learn

Develop your machine learning project locally with MLflow's different features Set up a centralized MLflow tracking server to manage multiple MLflow experiments Create a model life cycle with MLflow by creating custom models Use feature streams to log model results with MLflow Develop the complete training pipeline infrastructure using MLflow features Set up an inference-based API pipeline and batch pipeline in MLflow Scale large volumes of data by integrating MLflow with high-performance big data libraries

Who this book is for This book is for data scientists, machine learning engineers, and data engineers who want to gain hands-on machine learning engineering experience and learn how they can manage an end-to-end machine learning life cycle with the help of MLflow. Intermediate-level knowledge of the Python programming language is expected.

Algorithms are a fundamental building block of artificial intelligence - and, increasingly, society - but our legal institutions have largely failed to recognize or respond to this reality. The Cambridge Handbook of the Law of Algorithms, which features contributions from US, EU, and Asian legal scholars, discusses the specific challenges algorithms pose not only to current law, but also - as algorithms replace people as decision makers - to the foundations of society itself. The work includes wide coverage of the law as it relates to algorithms, with chapters analyzing how human biases have crept into algorithmic decision-making about who receives housing or credit, the length of sentences for defendants convicted of crimes, and many other decisions that impact constitutionally protected groups. Other issues covered in the work include the impact of algorithms on the law of free speech, intellectual property, and commercial and human rights law.

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

After a long time of neglect, Artificial Intelligence is once again at the center of most of our political, economic, and socio-cultural debates. Recent advances in the field of Artificial Neural Networks have led to a renaissance of dystopian and utopian speculations on an AI-rendered future. Algorithmic technologies are deployed for identifying potential terrorists through vast surveillance networks, for producing sentencing guidelines and recidivism risk profiles in criminal justice systems, for demographic and psychographic targeting of bodies for advertising or propaganda, and more generally for automating the analysis of language, text, and images. Against this background, the aim of this book is to discuss the heterogenous conditions, implications, and effects of modern AI and Internet technologies in terms of their political dimension: What does it mean to critically investigate efforts of net politics in the age of machine learning algorithms? Equipped with the latest updates, this third edition of Python Machine Learning By Example provides a comprehensive course for ML enthusiasts to strengthen their command of ML concepts, techniques, and algorithms.

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

"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.

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.

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 Featured by Tableau as the first of "7 Books About Machine Learning for Beginners." Ready to spin up a virtual GPU instance and smash through petabytes of data? Want to add 'Machine Learning' to your LinkedIn profile? Well, hold on there... Before you embark on your journey, there are some high-level theory and statistical principles to weave through first. But rather than spend \$30-\$50 USD on a thick textbook, you may want to read this book first. As a clear and concise alternative, this book provides a high-level introduction to machine learning, free downloadable code exercises, and video demonstrations. Machine Learning for Absolute Beginners Third Edition has been written and designed for absolute beginners. This means plain-English explanations and no coding experience required. Where core algorithms are introduced, clear explanations and visual examples are added to make it easy to follow along at home. This new edition also features extended chapters with quizzes, free supplementary online video tutorials for coding models in Python, and downloadable resources not included in the Second Edition. Readers of the Second Edition should not feel compelled to purchase this Third Edition. Disclaimer: If you have passed the 'beginner' stage in your study of machine learning and are ready to tackle coding and deep learning, you would be well served with a long-format textbook. If, however, you are yet to reach that Lion King moment - as a fully grown Simba looking over the Pride Lands of Africa - then this is the book to gently hoist you up and give a clear lay of the land. In this step-by-step guide you will learn: - How to download free datasets- What tools and machine learning libraries you need- Data scrubbing techniques, including one-hot encoding, binning and dealing with missing data- Preparing data for analysis, including k-fold Validation- Regression analysis to create trend lines- k-Means Clustering to find new relationships- The basics of Neural Networks- Bias/Variance to improve your machine learning model- Decision Trees to decode classification, and- How to build your first Machine Learning Model to predict house values using Python Frequently Asked Questions Q: Do I need programming experience to complete this e-book? A: This e-book is designed for absolute beginners, so no programming experience is required. However, two of the later chapters introduce Python to demonstrate an actual machine learning model, so you will see some programming used in this book. Q: I have already purchased the Second Edition of Machine Learning for Absolute Beginners, should I purchase this Third Edition? A: As the same topics from the Second Edition are covered in the Third Edition, you may be better served reading a more advanced title on machine learning. If you have purchased a previous edition of this book and wish to get access to the free video tutorials, please email the author. Q: Does this book include everything I need to become a machine learning expert? A: Unfortunately, no. This book is designed for readers taking their first steps in machine learning and further learning will be required beyond this book to master machine learning.

Longlisted for the National Book Award New York Times Bestseller A former Wall Street quant sounds an alarm on the mathematical models that pervade modern life -- and threaten to rip apart our social fabric We live in the age of the algorithm. Increasingly, the decisions that affect our lives--where we go to school, whether we get a car loan, how much we pay for health insurance--are being made not by humans, but by mathematical models. In theory, this should lead to greater fairness: Everyone is judged according to the same rules, and bias is eliminated. But as Cathy O'Neil reveals in this urgent and necessary book, the opposite is true. The models being used today are opaque, unregulated, and uncontestable, even when they're wrong. Most troubling, they reinforce discrimination: If a poor student can't get a loan because a lending model deems him too risky (by virtue of his zip code), he's then cut off from the kind of education that could pull him out of poverty, and a vicious spiral ensues. Models are propping up the lucky and punishing the downtrodden, creating a "toxic cocktail for democracy." Welcome to the dark side of Big Data. Tracing the arc of a person's life, O'Neil exposes the black box models that shape our future, both as individuals and as a society. These "weapons of math destruction" score teachers and students, sort r sum s, grant (or deny) loans, evaluate workers, target voters, set parole, and monitor our health. O'Neil calls on modelers to take more responsibility for their algorithms and on policy makers to regulate their use. But in the end, it's up to us to become more savvy about the models that govern our lives. This important book empowers us to ask the tough questions, uncover the truth, and demand change. -- Longlist for National Book Award (Non-Fiction) -- Goodreads, semi-finalist for the 2016 Goodreads Choice Awards (Science and Technology) -- Kirkus, Best Books of 2016 -- New York Times, 100 Notable Books of 2016 (Non-Fiction) -- The Guardian, Best Books of 2016 -- WBUR's "On Point," Best Books of 2016: Staff Picks -- Boston Globe, Best Books of 2016, Non-Fiction

A practical introduction perfect for final-year undergraduate and graduate students without a solid background in linear algebra and calculus.

For many researchers, Python is a first-class tool mainly because of its libraries for storing, manipulating, and gaining insight from data. Several resources exist for individual pieces of this data science stack, but only with the Python Data Science Handbook do you get them all—IPython, NumPy, Pandas, Matplotlib, Scikit-Learn, and other related tools. Working scientists and data crunchers familiar with reading and writing Python code will find this comprehensive desk reference ideal for tackling day-to-day issues: manipulating, transforming, and cleaning data; visualizing different types of data; and using data to build statistical or machine learning models. Quite simply, this is the must-have reference for scientific computing in Python. With this handbook, you'll learn how to use: IPython and Jupyter: provide computational environments for data scientists using Python NumPy: includes the ndarray for efficient storage and manipulation of dense data arrays in Python Pandas: features the DataFrame for efficient storage and manipulation of labeled/columnar data in Python Matplotlib: includes capabilities for a flexible range of data visualizations in Python Scikit-Learn: for efficient and clean Python implementations of the most important and established machine learning algorithms

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+ The Data Science Handbook is a curated collection of 25 candid, honest and insightful interviews conducted with some of the world's top data scientists. In this book, you'll hear how the co-creator of the term 'data scientist' thinks about career and personal success. You'll hear from a young woman who created her own data scientist curriculum, subsequently landing her a role in the field. Readers of this book will be left with war stories, wisdom and

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.

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

For over a decade, Andrew "bunnie" Huang, one of the world's most esteemed hackers, has shaped the fields of hacking and hardware, from his cult-classic book Hacking the Xbox to the open-source laptop Novena and his mentorship of various hardware startups and developers. In The Hardware Hacker, Huang shares his experiences in manufacturing and open hardware, creating an illuminating and compelling career retrospective. Huang's journey starts with his first visit to the staggering electronics markets in Shenzhen, with booths overflowing with capacitors, memory chips, voltmeters, and possibility. He shares how he navigated the overwhelming world of Chinese factories to bring chumby, Novena, and Chibitronics to life, covering everything from creating a Bill of Materials to choosing the factory to best fit his needs. Through this collection of personal essays and interviews on topics ranging from the legality of reverse engineering to a comparison of intellectual property practices between China and the United States, bunnie weaves engineering, law, and society into the tapestry of open hardware. With highly detailed passages on the ins and outs of manufacturing and a comprehensive take on the issues associated with open source hardware, The Hardware Hacker is an invaluable resource for aspiring hackers and makers.

A double is haunting the world--the double of abstraction, the virtual reality of information, programming or poetry, math or music, curves or colorings upon which the fortunes of states and armies, companies and communities now depend. The bold aim of this book is to make manifest the origins, purpose, and interests of the emerging class responsible for making this new world--for producing the new concepts, new perceptions, and new sensations out of the stuff of raw data. "A Hacker Manifesto" deftly defines the fraught territory between the ever more strident demands by drug and media companies for protection of their patents and copyrights and the pervasive popular culture of file sharing and pirating. This vexed ground, the realm of so-called "intellectual

property," gives rise to a whole new kind of class conflict, one that pits the creators of information--the hacker class of researchers and authors, artists and biologists, chemists and musicians, philosophers and programmers--against a possessing class who would monopolize what the hacker produces. Drawing in equal measure on Guy Debord and Gilles Deleuze, "A Hacker Manifesto" offers a systematic restatement of Marxist thought for the age of cyberspace and globalization. In the widespread revolt against commodified information, McKenzie Wark sees a utopian promise, beyond the property form, and a new progressive class, the hacker class, who voice a shared interest in a new information commons.

The second edition of a comprehensive introduction to machine learning approaches used in predictive data analytics, covering both theory and practice. Machine learning is often used to build predictive models by extracting patterns from large datasets. These models are used in predictive data analytics applications including price prediction, risk assessment, predicting customer behavior, and document classification. This introductory textbook offers a detailed and focused treatment of the most important machine learning approaches used in predictive data analytics, covering both theoretical concepts and practical applications. Technical and mathematical material is augmented with explanatory worked examples, and case studies illustrate the application of these models in the broader business context. This second edition covers recent developments in machine learning, especially in a new chapter on deep learning, and two new chapters that go beyond predictive analytics to cover unsupervised learning and reinforcement learning.

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

This book takes a pragmatic and hype-free approach to explaining artificial intelligence and how it can be utilised by businesses today. At the core of the book is a framework, developed by the author, which describes in non-technical language the eight core capabilities of Artificial Intelligence (AI). Each of these capabilities, ranging from image recognition, through natural language processing, to prediction, is explained using real-life examples and how they can be applied in a business environment. It will include interviews with executives who have successfully implemented AI as well as CEOs from AI vendors and consultancies. AI is one of the most talked about technologies in business today. It has the ability to deliver step-change benefits to organisations and enables forward-thinking CEOs to rethink their business models or create completely new businesses. But most of the real value of AI is hidden behind marketing hyperbole, confusing terminology, inflated expectations and dire warnings of 'robot overlords'. Any business executive that wants to know how to exploit AI in their business today is left confused and frustrated. As an advisor in Artificial Intelligence, Andrew Burgess regularly comes face-to-face with business executives who are struggling to cut through the hype that surrounds AI. The knowledge and experience he has gained in advising them, as well as working as a strategic advisor to AI vendors and consultancies, has provided him with the skills to help business executives understand what AI is and how they can exploit its many benefits. Through the distilled knowledge included in this book business leaders will be able to take full advantage of this most disruptive of technologies and create substantial competitive advantage for their companies.

"A globe-spanning investigation into the Transhumanist movement, considering the tech billionaires, scientific luminaries, and DIY body-hackers attempting to prolong, improve, and ultimately transcend the limits of human life"--

A substantially revised fourth edition of a comprehensive textbook, including new coverage of recent advances in deep learning and neural networks. The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Machine learning underlies such exciting new technologies as self-driving cars, speech recognition, and translation applications. This substantially revised fourth edition of a comprehensive, widely used machine learning textbook offers new coverage of recent advances in the field in both theory and practice, including developments in deep learning and neural networks. The book covers a broad array of topics not usually included in introductory machine learning texts, including supervised learning, Bayesian decision theory, parametric methods, semiparametric methods, nonparametric methods, multivariate analysis, hidden Markov models, reinforcement learning, kernel machines, graphical models, Bayesian estimation, and statistical testing. The fourth edition offers a new chapter on deep learning that discusses training, regularizing, and structuring deep neural networks such as convolutional and generative adversarial networks; new material in the chapter on reinforcement learning that covers the use of deep networks, the policy gradient methods, and deep reinforcement learning; new material in the chapter on multilayer perceptrons on autoencoders and the word2vec network; and discussion of a popular method of dimensionality reduction, t-SNE. New appendixes offer background material on linear algebra and optimization. End-of-chapter exercises help readers to apply concepts learned. Introduction to Machine Learning can be used in courses for advanced undergraduate and graduate students and as a reference for professionals.

Machine Learning for Hackers Case Studies and Algorithms to Get You Started"O'Reilly Media, Inc."

The second edition of a bestselling textbook, *Using R for Introductory Statistics* guides students through the basics of R, helping them overcome the sometimes steep learning curve. The author does this by breaking the material down into small, task-oriented steps. The second edition maintains the features that made the first edition so popular, while updating data, examples, and changes to R in line with the current version. See *What's New in the Second Edition*: Increased emphasis on more idiomatic R provides a grounding in the functionality of base R. Discussions of the use of RStudio helps new R users avoid as many pitfalls as possible. Use of knitr package makes code easier to read and therefore easier to reason about. Additional information on computer-intensive approaches motivates the traditional approach. Updated examples and data make the information current and topical. The book has an accompanying package, *UsingR*, available from CRAN, R's repository of user-contributed packages. The package contains the data sets mentioned in the text (`data(package="UsingR")`), answers to selected problems (`answers()`), a few demonstrations (`demo()`), the errata (`errata()`), and sample code from the text. The topics of this text line up closely with traditional teaching progression; however, the book also highlights computer-intensive approaches to motivate the more traditional approach. The authors emphasize realistic data and examples and rely on visualization techniques to gather insight. They introduce statistics and R seamlessly, giving students the tools they need to use R and the information they need to navigate the sometimes complex world of statistical computing.

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

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

Every day, businesses, investors, and consumers are grappling with the seismic changes technology has brought to the banking and finance industry. *The Money Hackers* is the dramatic story of fintech's major players and explores how these disruptions are transforming even money itself. Whether you've heard of fintech or not, it's already changing your life. Have you ever “Venmoed” someone? Do you think of investing in Bitcoin—even though you can't quite explain what it is? If you've deposited a check using your iPhone, that's fintech. And if you've gone to a bank branch and discovered it has been closed and shuttered for good, odds are that's because of fintech too. *The Money Hackers* focuses on some of fintech's most powerful disruptors—a ragtag collection of financial outsiders and savants—and uses their incredible stories to explain not just how the technology works, but how the Silicon Valley thinking behind the technology, ideas like friction, hedonic adaptation, democratization, and disintermediation, is having a drastic effect on the entire banking and finance industry. Turn to this remarkable new guide to: Feel empowered with the knowledge needed to spot the opportunities the next wave of fintech disruptions will bring. Understand the critical pain points that fintech is resolving, through a profile of the major finsurgents behind the disruption. Topic areas include Friction (featuring founders of Venmo), Aggregate and Automate (featuring Adam Dell, founder of Open Table and brother of Michael Dell), and Rise of the Machines (featuring Jon Stein, founder of robo-advisor Betterment). Learn about some of the larger-than-life characters behind the fintech movement. *The Money Hackers* tells the fascinating story of fintech—how it began, and where it is likely taking us.

You've decided to tackle machine learning - because you're job hunting, embarking on a new project, or just think self-driving cars are cool. But where to start? It's easy to be intimidated, even as a software developer. The good news is that it doesn't have to be that hard. Master machine learning by writing code one line at a time, from simple learning programs all the way to a true deep learning system. Tackle the hard topics by breaking them down so they're easier to understand, and build your confidence by getting your hands dirty. Peel away the obscurities of machine learning, starting from scratch and going all the way to deep learning. Machine learning can be intimidating, with its reliance on math and algorithms that most programmers don't encounter in their regular work. Take a hands-on approach, writing the Python code yourself, without any libraries to obscure what's really going on. Iterate on your design, and add layers of complexity as you go. Build an image recognition application from scratch with supervised learning. Predict the future with linear regression. Dive into gradient descent, a fundamental algorithm that drives most of machine learning. Create perceptrons to classify data. Build neural networks to tackle more complex and sophisticated data sets. Train and refine those networks with backpropagation and batching. Layer the neural networks, eliminate overfitting, and add convolution to transform your neural network into a true deep learning system. Start from the beginning and code your way to machine learning mastery. *What You Need*: The examples in this book are written in Python, but don't worry if you don't know this language: you'll pick up all the Python you need very quickly. Apart from that, you'll only need your computer, and your code-adept brain.

An introduction to computational complexity theory, its connections and interactions with mathematics, and its central role in the natural and social sciences, technology, and philosophy *Mathematics and Computation* provides a broad, conceptual overview of computational complexity theory—the mathematical study of efficient computation. With important

practical applications to computer science and industry, computational complexity theory has evolved into a highly interdisciplinary field, with strong links to most mathematical areas and to a growing number of scientific endeavors. Avi Wigderson takes a sweeping survey of complexity theory, emphasizing the field's insights and challenges. He explains the ideas and motivations leading to key models, notions, and results. In particular, he looks at algorithms and complexity, computations and proofs, randomness and interaction, quantum and arithmetic computation, and cryptography and learning, all as parts of a cohesive whole with numerous cross-influences. Wigderson illustrates the immense breadth of the field, its beauty and richness, and its diverse and growing interactions with other areas of mathematics. He ends with a comprehensive look at the theory of computation, its methodology and aspirations, and the unique and fundamental ways in which it has shaped and will further shape science, technology, and society. For further reading, an extensive bibliography is provided for all topics covered. Mathematics and Computation is useful for undergraduate and graduate students in mathematics, computer science, and related fields, as well as researchers and teachers in these fields. Many parts require little background, and serve as an invitation to newcomers seeking an introduction to the theory of computation. Comprehensive coverage of computational complexity theory, and beyond High-level, intuitive exposition, which brings conceptual clarity to this central and dynamic scientific discipline Historical accounts of the evolution and motivations of central concepts and models A broad view of the theory of computation's influence on science, technology, and society Extensive bibliography

What others in the trenches say about *The Pragmatic Programmer*... "The cool thing about this book is that it's great for keeping the programming process fresh. The book helps you to continue to grow and clearly comes from people who have been there." —Kent Beck, author of *Extreme Programming Explained: Embrace Change* "I found this book to be a great mix of solid advice and wonderful analogies!" —Martin Fowler, author of *Refactoring and UML Distilled* "I would buy a copy, read it twice, then tell all my colleagues to run out and grab a copy. This is a book I would never loan because I would worry about it being lost." —Kevin Ruland, Management Science, MSG-Logistics "The wisdom and practical experience of the authors is obvious. The topics presented are relevant and useful.... By far its greatest strength for me has been the outstanding analogies—tracer bullets, broken windows, and the fabulous helicopter-based explanation of the need for orthogonality, especially in a crisis situation. I have little doubt that this book will eventually become an excellent source of useful information for journeymen programmers and expert mentors alike." —John Lakos, author of *Large-Scale C++ Software Design* "This is the sort of book I will buy a dozen copies of when it comes out so I can give it to my clients." —Eric Vought, Software Engineer "Most modern books on software development fail to cover the basics of what makes a great software developer, instead spending their time on syntax or technology where in reality the greatest leverage possible for any software team is in having talented developers who really know their craft well. An excellent book." —Pete McBreen, Independent Consultant "Since reading this book, I have implemented many of the practical suggestions and tips it contains. Across the board, they have saved my company time and money while helping me get my job done quicker! This should be a desktop reference for everyone who works with code for a living." —Jared Richardson, Senior Software Developer, iRenaissance, Inc. "I would like to see this issued to every new employee at my company...." —Chris Cleeland, Senior Software Engineer, Object Computing, Inc. "If I'm putting together a project, it's the authors of this book that I want. . . . And failing that I'd settle for people who've read their book." —Ward Cunningham

Straight from the programming trenches, *The Pragmatic Programmer* cuts through the increasing specialization and technicalities of modern software development to examine the core process--taking a requirement and producing working, maintainable code that delights its users. It covers topics ranging from personal responsibility and career development to architectural techniques for keeping your code flexible and easy to adapt and reuse. Read this book, and you'll learn how to Fight software rot; Avoid the trap of duplicating knowledge; Write flexible, dynamic, and adaptable code; Avoid programming by coincidence; Bullet-proof your code with contracts, assertions, and exceptions; Capture real requirements; Test ruthlessly and effectively; Delight your users; Build teams of pragmatic programmers; and Make your developments more precise with automation. Written as a series of self-contained sections and filled with entertaining anecdotes, thoughtful examples, and interesting analogies, *The Pragmatic Programmer* illustrates the best practices and major pitfalls of many different aspects of software development. Whether you're a new coder, an experienced programmer, or a manager responsible for software projects, use these lessons daily, and you'll quickly see improvements in personal productivity, accuracy, and job satisfaction. You'll learn skills and develop habits and attitudes that form the foundation for long-term success in your career. You'll become a Pragmatic Programmer.

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

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."

A concise and self-contained introduction to causal inference, increasingly important in data science and machine learning. The mathematization of causality is a relatively recent development, and has become increasingly important in data science and machine learning. This book offers a self-contained and concise introduction to causal models and how to learn them from data. After explaining the need for causal models and discussing some of the principles underlying causal inference, the book teaches readers how to use causal models: how to compute intervention distributions, how to infer causal models from observational and interventional data, and how causal ideas could be exploited for classical machine learning problems. All of these topics are discussed first in terms of two variables and then in the more general multivariate case. The bivariate case turns out to be a particularly hard problem for causal learning because there are no conditional independences as used by classical methods for solving multivariate cases. The authors consider analyzing statistical asymmetries between cause and effect to be highly instructive, and they report on their decade of intensive research into this problem. The book is accessible to readers with a background in machine learning or statistics, and can be used in graduate courses or as a reference for researchers. The text includes code snippets that can be copied and pasted, exercises, and an appendix with a summary of the most important technical concepts.

Hack your antivirus software to stamp out future vulnerabilities The Antivirus Hacker's Handbook guides you through the process of reverse engineering antivirus software. You explore how to detect and exploit vulnerabilities that can be leveraged to improve future software design, protect your network, and anticipate attacks that may sneak through your antivirus' line of defense. You'll begin building your knowledge by diving into the reverse engineering process, which details how to start from a finished antivirus software program and work your way back through its development using the functions and other key elements of the software. Next, you leverage your new knowledge about software development to evade, attack, and exploit antivirus software—all of which can help you strengthen your network and protect your data. While not all viruses are damaging, understanding how to better protect your computer against them can help you maintain the integrity of your network. Discover how to reverse engineer your antivirus software Explore methods of antivirus software evasion Consider different ways to attack and exploit antivirus software Understand the current state of the antivirus software market, and get recommendations for users and vendors who are leveraging this software The Antivirus Hacker's Handbook is the essential reference for software reverse engineers, penetration testers, security researchers, exploit writers, antivirus vendors, and software engineers who want to understand how to leverage current antivirus software to improve future applications.

This hands-on book will help you make your machine learning models fairer, safer, and more reliable and in turn improve business outcomes. Every chapter introduces a new mission where you learn how to apply interpretation methods to realistic use cases with methods that work for any model type as well as methods specific for deep neural networks. Unlock today's statistical controversies and irreproducible results by viewing statistics as probing and controlling errors. A comprehensive guide to getting well-versed with the mathematical techniques for building modern deep learning architectures Key Features Understand linear algebra, calculus, gradient algorithms, and other concepts essential for training deep neural networks Learn the mathematical concepts needed to understand how deep learning models function Use deep learning for solving problems related to vision, image, text, and sequence applications Book Description Most programmers and data scientists struggle with mathematics, having either overlooked or forgotten core mathematical concepts. This book uses Python libraries to help you understand the math required to build deep learning (DL) models. You'll begin by learning about core mathematical and modern computational techniques used to design and implement DL algorithms. This book will cover essential topics, such as linear algebra, eigenvalues and eigenvectors, the singular value decomposition concept, and gradient algorithms, to help you understand how to train deep neural networks. Later chapters focus on important neural networks, such as the linear neural network and multilayer perceptrons, with a primary focus on helping you learn how each model works. As you advance, you will delve into the math used for regularization, multi-layered DL, forward propagation, optimization, and backpropagation techniques to understand what it takes to build full-fledged DL models. Finally, you'll explore CNN, recurrent neural network (RNN), and GAN models and their application. By the end of this book, you'll have built a strong foundation in neural networks and DL mathematical concepts, which will help you to confidently research and build custom models in DL. What you will learn Understand the key mathematical concepts for building neural network models Discover core multivariable calculus concepts Improve the performance of deep learning models using optimization techniques Cover optimization algorithms,

from basic stochastic gradient descent (SGD) to the advanced Adam optimizer Understand computational graphs and their importance in DL Explore the backpropagation algorithm to reduce output error Cover DL algorithms such as convolutional neural networks (CNNs), sequence models, and generative adversarial networks (GANs) Who this book is for This book is for data scientists, machine learning developers, aspiring deep learning developers, or anyone who wants to understand the foundation of deep learning by learning the math behind it. Working knowledge of the Python programming language and machine learning basics is required.

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

[Copyright: 60afcaebceb66d49d973ab0d33579fbf](#)