

Histograms H Just Maths

If you want a basic understanding of computer vision's underlying theory and algorithms, this hands-on introduction is the ideal place to start. You'll learn techniques for object recognition, 3D reconstruction, stereo imaging, augmented reality, and other computer vision applications as you follow clear examples written in Python. Programming Computer Vision with Python explains computer vision in broad terms that won't bog you down in theory. You get complete code samples with explanations on how to reproduce and build upon each example, along with exercises to help you apply what you've learned. This book is ideal for students, researchers, and enthusiasts with basic programming and standard mathematical skills. Learn techniques used in robot navigation, medical image analysis, and other computer vision applications Work with image mappings and transforms, such as texture warping and panorama creation Compute 3D reconstructions from several images of the same scene Organize images based on similarity or content, using clustering methods Build efficient image retrieval techniques to search for images based on visual content Use algorithms to classify image content and recognize objects Access the popular OpenCV library through a Python interface

The high-level language of R is recognized as one of the most powerful and flexible statistical software environments, and is rapidly becoming the standard setting for quantitative analysis, statistics and graphics. R provides free access to unrivalled coverage and cutting-edge applications, enabling the user to apply numerous statistical methods ranging from simple regression to time series or multivariate analysis. Building on the success of the author's bestselling Statistics: An Introduction using R, The R Book is packed with worked examples, providing an all inclusive guide to R, ideal for novice and more accomplished users alike. The book assumes no background in statistics or computing and introduces the advantages of the R environment, detailing its applications in a wide range of disciplines. Provides the first comprehensive reference manual for the R language, including practical guidance and full coverage of the graphics facilities. Introduces all the statistical models covered by R, beginning with simple classical tests such as chi-square and t-test. Proceeds to examine more advance methods, from regression and analysis of variance, through to generalized linear models, generalized mixed models, time series, spatial statistics, multivariate statistics and much more. The R Book is aimed at undergraduates, postgraduates and professionals in science, engineering and medicine. It is also ideal for students and professionals in statistics, economics, geography and the social sciences.

Familiar with traditional photography but confused by megapixels, white balance, electronic sensors, and all the other jargon associated with digital photography? Clear up the confusion, learn about the equipment and techniques, and get back to what's important - taking great pictures. This book provides an introduction to bird photography using a digital

SLR. Learn about digital equipment and the common characteristics of both digital and film. This book also covers the unique advantages of the digital medium, and the photographic techniques, fieldcraft, and post-processing techniques that are used to produce stunning images of the creatures that share our world. With a full-colour, liberally illustrated interior, this book is an essential addition to any photographer's library

Providing a much-needed bridge between elementary statistics courses and advanced research methods courses, *Understanding Advanced Statistical Methods* helps students grasp the fundamental assumptions and machinery behind sophisticated statistical topics, such as logistic regression, maximum likelihood, bootstrapping, nonparametrics, and Bayesian methods. The book teaches students how to properly model, think critically, and design their own studies to avoid common errors. It leads them to think differently not only about math and statistics but also about general research and the scientific method. With a focus on statistical models as producers of data, the book enables students to more easily understand the machinery of advanced statistics. It also downplays the "population" interpretation of statistical models and presents Bayesian methods before frequentist ones. Requiring no prior calculus experience, the text employs a "just-in-time" approach that introduces mathematical topics, including calculus, where needed. Formulas throughout the text are used to explain why calculus and probability are essential in statistical modeling. The authors also intuitively explain the theory and logic behind real data analysis, incorporating a range of application examples from the social, economic, biological, medical, physical, and engineering sciences. Enabling your students to answer the why behind statistical methods, this text teaches them how to successfully draw conclusions when the premises are flawed. It empowers them to use advanced statistical methods with confidence and develop their own statistical recipes. Ancillary materials are available on the book's website.

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

The Book of R is a comprehensive, beginner-friendly guide to R, the world's most popular programming language for statistical analysis. Even if you have no programming experience and little more than a grounding in the basics of mathematics, you'll find everything you need to begin using R effectively for statistical analysis. You'll start with the basics, like how to handle data and write simple programs, before moving on to more advanced topics, like producing statistical summaries of your data and performing statistical tests and modeling. You'll even learn how to create impressive data visualizations with R's basic graphics tools and contributed packages, like ggplot2 and ggvis, as well as interactive 3D visualizations using the rgl package. Dozens of hands-on exercises (with downloadable solutions) take you from theory to practice, as you learn:

- The fundamentals of programming in R, including how to write data frames, create functions, and use variables, statements, and loops
- Statistical concepts like exploratory data analysis, probabilities, hypothesis tests, and regression modeling, and how to execute them in R
- How to access R's thousands of functions, libraries, and data sets
- How to draw valid and useful conclusions from your data
- How to create publication-quality graphics of your results

Combining detailed explanations with real-world examples and exercises, this book will provide you with a solid understanding of both statistics and the depth of R's functionality. Make The Book of R your doorway into the growing world of data analysis.

A comprehensive and user-friendly introduction to statistics for behavioral science students—revised and updated Refined over seven editions by master teachers, this book gives instructors and students alike clear examples and carefully crafted exercises to support the teaching and learning of statistics for both manipulating and consuming data. One of the most popular and respected statistics texts in the behavioral sciences, the Seventh Edition of Introductory Statistics for the Behavioral Sciences has been fully revised. The new edition presents all the topics students in the behavioral sciences need in a uniquely accessible and easy-to-understand format, aiding in the comprehension and implementation of the statistical analyses most commonly used in the behavioral sciences. The Seventh Edition features:

- A continuous narrative that clearly explains statistics while tracking a common data set throughout, making the concepts unintimidating and memorable, and providing a framework that connects all of the topics and allows for easy comparison of different statistical analyses
- Coverage of important aspects of research design throughout the text, such as the "correlation is not causality" principle
- Updated and annotated SPSS output at the end of each chapter with step-by-step instructions
- Updated examples and exercises
- An expanded website, at www.wiley.com/go/welkowitz, with testbank, chapter quizzes, and PowerPoint slides for instructors, as well as a second website for students with additional basic math coverage, math review exercises, a study guide, a set of additional SPSS exercises, and more downloadable data sets

This is the first book to be devoted entirely to the application and development of flow techniques in cytogenetics. It provides comprehensive information on the use of flow cytometry and sorting for chromosome classification and purification. Cytogenetics and molecular biologists will find this book an invaluable reference source. Practical details for the preparation and analysis of chromosomes using flow cytometry Flow karyotyping for sensitive rapid analysis of chromosome normality and the detection of aberrant chromosomes Flow sorting as a source of chromosome-specific DNA for gene mapping and recombinant DNA libraries Construction and current status of chromosome-specific recombinant DNA libraries

Regression Analysis and Its Application: A Data-Oriented Approach answers the need for researchers and students who would like a better understanding of classical regression analysis. Useful either as a textbook or as a reference source, this book bridges the gap between the purely theoretical coverage of regression analysis and its practical application. The book presents regression analysis in the general context of data analysis. Using a teach-by-example format, it contains ten major data sets along with several smaller ones to illustrate the common characteristics of regression data and properties of statistics that are employed in regression analysis. The book covers model misspecification, residual analysis, multicollinearity, and biased regression estimators. It also focuses on data collection, model assumptions, and the interpretation of parameter estimates. Complete with an extensive bibliography, Regression Analysis and Its Application is suitable for statisticians, graduate and upper-level undergraduate students, and research scientists in biometry, business, ecology, economics, education, engineering, mathematics, physical sciences, psychology, and sociology. In addition, data collection agencies in the government and private sector will benefit from the book.

Nature-inspired algorithms such as cuckoo search and firefly algorithm have become popular and widely used in recent years in many applications. These algorithms are flexible, efficient and easy to implement. New progress has been made in the last few years, and it is timely to summarize the latest developments of cuckoo search and firefly algorithm and their diverse applications. This book will review both theoretical studies and applications with detailed algorithm analysis, implementation and case studies so that readers can benefit most from this book. Application topics are contributed by many leading experts in the field. Topics include cuckoo search, firefly algorithm, algorithm analysis, feature selection, image processing, travelling salesman problem, neural network, GPU optimization, scheduling, queuing, multi-objective manufacturing optimization, semantic web service, shape optimization, and others. This book can serve as an ideal reference for both graduates and researchers in computer science, evolutionary computing, machine learning, computational intelligence, and optimization, as well as engineers in business intelligence, knowledge management and information technology.

Collects short stories exploring themes of time and space travel, self-discovery, and science and technology.

Now, it is easier than ever before to understand complex mathematical concepts and formulas and how they relate to real-world business situations. All you have to do it apply the handy information you will find in Business Math For Dummies. Featuring practical practice problems to help you expand your skills, this book covers topics like using percents to calculate increases and

decreases, applying basic algebra to solve proportions, and working with basic statistics to analyze raw data. Find solutions for finance and payroll applications, including reading financial statements, calculating wages and commissions, and strategic salary planning. Navigate fractions, decimals, and percents in business and real estate transactions, and take fancy math skills to work. You'll be able to read graphs and tables and apply statistics and data analysis. You'll discover ways you can use math in finance and payroll investments, banking and payroll, goods and services, and business facilities and operations. You'll learn how to calculate discounts and markup, use loans and credit, and understand the ins and outs of math for business facilities and operations. You'll be the company math whiz in no time at all! Find out how to: Read graphs and tables Invest in the future Use loans and credit Navigate bank accounts, insurance, budgets, and payroll Calculate discounts and markup Measure properties and handle mortgages and loans Manage rental and commercial properties Complete with lists of ten math shortcuts to do in meetings and drive your coworkers nuts and ten tips for reading annual reports, *Business MathFor Dummies* is your one-stop guide to solving math problems in business situations.

Online Statistics: An Interactive Multimedia Course of Study is a resource for learning and teaching introductory statistics. It contains material presented in textbook format and as video presentations. This resource features interactive demonstrations and simulations, case studies, and an analysis lab. This print edition of the public domain textbook gives the student an opportunity to own a physical copy to help enhance their educational experience. This part I features the book *Front Matter*, Chapters 1-10, and the full Glossary. Chapters Include: I. Introduction, II. Graphing Distributions, III. Summarizing Distributions, IV. Describing Bivariate Data, V. Probability, VI. Research Design, VII. Normal Distributions, VIII. Advanced Graphs, IX. Sampling Distributions, and X. Estimation. *Online Statistics Education: A Multimedia Course of Study* (<http://onlinestatbook.com/>). Project Leader: David M. Lane, Rice University.

Waves in Oceanic and Coastal Waters describes the observation, analysis and prediction of wind-generated waves in the open ocean, in shelf seas, and in coastal regions with islands, channels, tidal flats and inlets, estuaries, fjords and lagoons. Most of this richly illustrated book is devoted to the physical aspects of waves. After introducing observation techniques for waves, both at sea and from space, the book defines the parameters that characterise waves. Using basic statistical and physical concepts, the author discusses the prediction of waves in oceanic and coastal waters, first in terms of generalised observations, and then in terms of the more theoretical framework of the spectral energy balance. He gives the results of established theories and also the direction in which research is developing. The book ends with a description of SWAN (Simulating Waves Nearshore), the preferred computer model of the engineering community for predicting waves in coastal waters.

A groundbreaking introduction to vectors, matrices, and least squares for engineering applications, offering a wealth of practical examples. *Statistical Inference via Data Science: A ModernDive into R and the Tidyverse* provides a pathway for learning about statistical inference using data science tools widely used in industry, academia, and government. It introduces the tidyverse suite of R packages, including the ggplot2 package for data visualization, and the dplyr package for data wrangling. After equipping readers with just enough of these data

science tools to perform effective exploratory data analyses, the book covers traditional introductory statistics topics like confidence intervals, hypothesis testing, and multiple regression modeling, while focusing on visualization throughout. Features: ? Assumes minimal prerequisites, notably, no prior calculus nor coding experience ? Motivates theory using real-world data, including all domestic flights leaving New York City in 2013, the Gapminder project, and the data journalism website, FiveThirtyEight.com ? Centers on simulation-based approaches to statistical inference rather than mathematical formulas ? Uses the infer package for "tidy" and transparent statistical inference to construct confidence intervals and conduct hypothesis tests via the bootstrap and permutation methods ? Provides all code and output embedded directly in the text; also available in the online version at moderndive.com This book is intended for individuals who would like to simultaneously start developing their data science toolbox and start learning about the inferential and modeling tools used in much of modern-day research. The book can be used in methods and data science courses and first courses in statistics, at both the undergraduate and graduate levels. Praise for the First Edition ". . . an excellent textbook . . . well organized and neatly written." —Mathematical Reviews ". . . amazingly interesting . . ." —Technometrics Thoroughly updated to showcase the interrelationships between probability, statistics, and stochastic processes, *Probability, Statistics, and Stochastic Processes, Second Edition* prepares readers to collect, analyze, and characterize data in their chosen fields. Beginning with three chapters that develop probability theory and introduce the axioms of probability, random variables, and joint distributions, the book goes on to present limit theorems and simulation. The authors combine a rigorous, calculus-based development of theory with an intuitive approach that appeals to readers' sense of reason and logic. Including more than 400 examples that help illustrate concepts and theory, the Second Edition features new material on statistical inference and a wealth of newly added topics, including: Consistency of point estimators Large sample theory Bootstrap simulation Multiple hypothesis testing Fisher's exact test and Kolmogorov-Smirnov test Martingales, renewal processes, and Brownian motion One-way analysis of variance and the general linear model Extensively class-tested to ensure an accessible presentation, *Probability, Statistics, and Stochastic Processes, Second Edition* is an excellent book for courses on probability and statistics at the upper-undergraduate level. The book is also an ideal resource for scientists and engineers in the fields of statistics, mathematics, industrial management, and engineering.

Introductory Statistics is designed for the one-semester, introduction to statistics course and is geared toward students majoring in fields other than math or engineering. This text assumes students have been exposed to intermediate algebra, and it focuses on the applications of statistical knowledge rather than the theory behind it. The foundation of this textbook is Collaborative Statistics, by Barbara Illowsky and Susan Dean. Additional topics, examples, and ample opportunities for practice have been added to each chapter. The development choices for this textbook were made with the guidance of many faculty members who are deeply involved in teaching this course. These choices led to innovations in art, terminology, and practical applications, all with a goal of increasing relevance and accessibility for students. We strove to make the discipline meaningful, so that students can draw from it a working knowledge that will enrich their future studies and help them make sense of the world around them. Coverage and Scope Chapter 1 Sampling and Data Chapter 2 Descriptive Statistics Chapter 3 Probability Topics Chapter 4 Discrete Random Variables Chapter 5 Continuous Random Variables Chapter 6 The Normal Distribution Chapter 7 The Central Limit Theorem Chapter 8 Confidence Intervals Chapter 9 Hypothesis Testing with One Sample Chapter 10 Hypothesis Testing with Two Samples Chapter 11 The Chi-Square Distribution Chapter 12 Linear Regression and Correlation Chapter 13 F Distribution and One-Way ANOVA

PCMag.com is a leading authority on technology, delivering Labs-based, independent reviews of the latest products and services. Our expert

industry analysis and practical solutions help you make better buying decisions and get more from technology.

The routine jobs of yesterday are being replaced by technology and/or shipped off-shore. In their place, job categories that require knowledge management, abstract reasoning, and personal services seem to be growing. The modern workplace requires workers to have broad cognitive and affective skills. Often referred to as "21st century skills," these skills include being able to solve complex problems, to think critically about tasks, to effectively communicate with people from a variety of different cultures and using a variety of different techniques, to work in collaboration with others, to adapt to rapidly changing environments and conditions for performing tasks, to effectively manage one's work, and to acquire new skills and information on one's own. The National Research Council (NRC) has convened two prior workshops on the topic of 21st century skills. The first, held in 2007, was designed to examine research on the skills required for the 21st century workplace and the extent to which they are meaningfully different from earlier eras and require corresponding changes in educational experiences. The second workshop, held in 2009, was designed to explore demand for these types of skills, consider intersections between science education reform goals and 21st century skills, examine models of high-quality science instruction that may develop the skills, and consider science teacher readiness for 21st century skills. The third workshop was intended to delve more deeply into the topic of assessment. The goal for this workshop was to capitalize on the prior efforts and explore strategies for assessing the five skills identified earlier. The Committee on the Assessment of 21st Century Skills was asked to organize a workshop that reviewed the assessments and related research for each of the five skills identified at the previous workshops, with special attention to recent developments in technology-enabled assessment of critical thinking and problem-solving skills. In designing the workshop, the committee collapsed the five skills into three broad clusters as shown below: Cognitive skills: nonroutine problem solving, critical thinking, systems thinking Interpersonal skills: complex communication, social skills, team-work, cultural sensitivity, dealing with diversity Intrapersonal skills: self-management, time management, self-development, self-regulation, adaptability, executive functioning Assessing 21st Century Skills provides an integrated summary of the presentations and discussions from both parts of the third workshop.

Upon publication, the first edition of the CRC Concise Encyclopedia of Mathematics received overwhelming accolades for its unparalleled scope, readability, and utility. It soon took its place among the top selling books in the history of Chapman & Hall/CRC, and its popularity continues unabated. Yet also unabated has been the d

This exciting book/disk package completely outfits the serious programmer for 3D graphics work in Windows. It explains 3D graphics programming fundamentals from a Windows perspective. Readers will understand the Windows graphic interface and how they can use a graphics library to create sophisticated graphics applications. Disk contains a professional 3D graphics library for Windows.

The papers in this volume provide a comprehensive account of the current methods and work in geostatistics, including recent theoretical developments and applications. Topics featured include: stochastic simulations, space-time modelling, and Bayesian framework.

This book, now in its third edition, offers a practical guide to the use of probability and statistics in experimental physics that is of value for both advanced undergraduates and graduate students. Focusing on applications and theorems and techniques actually used in experimental research, it includes worked problems with solutions, as well as homework exercises to aid understanding. Suitable for readers with no prior knowledge of statistical techniques, the book comprehensively discusses the topic and features a number of interesting and amusing applications that are often neglected. Providing an introduction to neural net techniques that encompasses deep learning, adversarial neural networks, and boosted decision trees, this new edition includes updated chapters with, for example, additions relating to generating and

characteristic functions, Bayes' theorem, the Feldman-Cousins method, Lagrange multipliers for constraints, estimation of likelihood ratios, and unfolding problems.

This manual was written to help the student make use of the statistical functions of the Texas Instruments TI-83 plus, TI-84 plus, and TI 89 graphing calculators.

Intelligent environments (IE) play an increasingly important role in many areas of our lives, including education, healthcare and the domestic environment. The term refers to physical spaces incorporating pervasive computing technology used to achieve specific goals for the user, the environment or both. This book presents the proceedings of the workshops of the 8th International Conference on Intelligent Environments (IE '12), held in Guanajuato, Mexico, in June 2012. The workshops which make up the conference range from regular lectures to practical sessions. They provide a forum for scientists, researchers and engineers from both industry and academia to engage in discussions on newly emerging or rapidly evolving topics in the field. Topics covered in the workshops include intelligent environments supporting healthcare and well-being; artificial intelligence techniques for ambient intelligence; large-scale intelligent environments; intelligent domestic robots; intelligent environment technology in education; multimodal interfaces applied in skills transfer, healthcare and rehabilitation; the reliability of intelligent environments and improving industrial automation using intelligent environments. IE can enrich user experience, better manage the environment's resources, and increase user awareness of that environment. This book will be of interest to all those whose work involves the application of intelligent environments.

Introduction to Data Science: Data Analysis and Prediction Algorithms with R introduces concepts and skills that can help you tackle real-world data analysis challenges. It covers concepts from probability, statistical inference, linear regression, and machine learning. It also helps you develop skills such as R programming, data wrangling, data visualization, predictive algorithm building, file organization with UNIX/Linux shell, version control with Git and GitHub, and reproducible document preparation. This book is a textbook for a first course in data science. No previous knowledge of R is necessary, although some experience with programming may be helpful. The book is divided into six parts: R, data visualization, statistics with R, data wrangling, machine learning, and productivity tools. Each part has several chapters meant to be presented as one lecture. The author uses motivating case studies that realistically mimic a data scientist's experience. He starts by asking specific questions and answers these through data analysis so concepts are learned as a means to answering the questions. Examples of the case studies included are: US murder rates by state, self-reported student heights, trends in world health and economics, the impact of vaccines on infectious disease rates, the financial crisis of 2007-2008, election forecasting, building a baseball team, image processing of hand-written digits, and movie recommendation systems. The statistical concepts used to answer the case study questions are only briefly introduced, so complementing with a probability and statistics textbook is highly recommended for in-depth understanding of these concepts. If you read and understand the chapters and complete the exercises, you will be prepared to learn the more advanced concepts and skills needed to become an expert.

This book constitutes the proceedings of the 19th International Conference on Cryptographic Hardware and Embedded Systems, CHES 2017, held in Taipei, Taiwan, in September 2017. The 33 full papers presented in this volume were carefully reviewed and selected from 130 submissions. The annual CHES conference highlights new results in the design and analysis of cryptographic hardware and software implementations. The workshop builds a valuable bridge between the research and cryptographic engineering communities and attracts participants from industry, academia, and government organizations.

'Easy as you Go' was originally intended to be solely a support for the learning of Mathematics. However it has evolved into something more than just that. The two volumes are packed with a total of 950 pages of mathematics, covering no less than 230 topics and containing a multitude of worked examples, equations and formulas, graphs and charts, tables, diagrams and illustrations. Together, the two volumes address all the significant issues encountered in First School, Secondary School and in Advanced studies, along with a plethora of anecdotal topics to capture the reader's imagination, and titivate their perhaps otherwise sanguine attitude towards Mathematics. 'Easy as you Go' is ideally suited to student, educator and parent alike because of its simplistic, down-to-earth and visual approach.

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.

While the prediction of observations is a forward problem, the use of actual observations to infer the properties of a model is an inverse problem. Inverse problems are difficult because they may not have a unique solution. The description of uncertainties plays a central role in the theory, which is based on probability theory. This book proposes a general approach that is valid for linear as well as for nonlinear problems. The philosophy is essentially probabilistic and allows the reader to understand the basic difficulties appearing in the resolution of inverse problems. The book attempts to explain how a method of acquisition of information can be applied to actual real-world problems, and many of the arguments are heuristic.

"We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document." —Soumith Chintala, co-creator of PyTorch

Key Features Written by PyTorch's creator and key contributors

- Develop deep learning models in a familiar Pythonic way
- Use PyTorch to build an image classifier for cancer detection
- Diagnose problems with your neural network and improve training with data augmentation

Purchase of the print book

includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

[Copyright: 862c11604e1ac5d9ea4944d2e9545dec](#)