

Introduction To Python Harvard University

This book is a mini-course for researchers in the atmospheric and oceanic sciences. "We assume readers will already know the basics of programming... in some other language." - Back cover.

Statistics, Data Mining, and Machine Learning in Astronomy is the essential introduction to the statistical methods needed to analyze complex data sets from astronomical surveys such as the Panoramic Survey Telescope and Rapid Response System, the Dark Energy Survey, and the Large Synoptic Survey Telescope. Now fully updated, it presents a wealth of practical analysis problems, evaluates the techniques for solving them, and explains how to use various approaches for different types and sizes of data sets. Python code and sample data sets are provided for all applications described in the book. The supporting data sets have been carefully selected from contemporary astronomical surveys and are easy to download and use. The accompanying Python code is publicly available, well documented, and follows uniform coding standards. Together, the data sets and code enable readers to reproduce all the figures and examples, engage with the different methods, and adapt them to their own fields of interest. An accessible textbook for students and an indispensable reference for researchers, this updated edition features new sections on deep learning methods, hierarchical Bayes modeling, and approximate Bayesian computation. The chapters have been revised throughout and the astroML code has been brought completely up to date. Fully revised and expanded Describes the most useful statistical and data-mining methods for extracting knowledge from huge and complex astronomical data sets Features real-world data sets from astronomical surveys Uses a freely available Python codebase throughout Ideal for graduate students, advanced undergraduates, and working astronomers

From the ads that track us to the maps that guide us, the twenty-first century runs on code. The business world is no different. Programming has become one of the fastest-growing topics at business schools around the world. An increasing number of MBAs are choosing to pursue careers in tech. For them and other professionals, having some basic coding knowledge is a must. This book is an introduction to programming with Python for MBA students and others in business positions who need a crash course. One of the most popular programming languages, Python is used for tasks such as building and running websites, data analysis, machine learning, and natural-language processing. Drawing on years of experience providing instruction in this material at Columbia Business School as well as extensive backgrounds in technology, entrepreneurship, and consulting, Mattan Griffel and Daniel Guetta teach the basics of programming from scratch. Beginning with fundamentals such as variables, strings, lists, and functions, they build up to data analytics and practical ways to derive value from large and complex datasets. They focus on business use cases throughout, using the real-world example of a major restaurant chain to offer a concrete look at what Python can do. Written for business students with no previous coding experience and those in business roles that include coding or working with coding teams, Python for MBAs is an indispensable introduction to a versatile and powerful programming language. Python for Everybody is designed to introduce students to programming and software development through the lens of exploring data. You can think of the Python programming language as your tool to solve data problems that are beyond the capability of a spreadsheet. Python is an easy to use and easy to learn programming language that is freely available on Macintosh, Windows, or Linux computers. So once you learn Python you can use it for the rest of your career without needing to purchase any software. This book uses the Python 3 language. The earlier Python 2 version of this book is titled "Python for Informatics: Exploring Information". There are free downloadable electronic copies of this book in various formats and supporting materials for the book at www.pythonlearn.com. The course materials are available to you under a Creative Commons License so you can adapt them to teach your own Python course.

Deep learning networks are getting smaller. Much smaller. The Google Assistant team can detect words with a model just 14 kilobytes in size—small enough to run on a microcontroller. With this practical book you'll enter the field of TinyML, where deep learning and embedded systems combine to make astounding things possible with tiny devices. Pete Warden and Daniel Situnayake explain how you can train models small enough to fit into any environment. Ideal for software and hardware developers who want to build embedded systems using machine learning, this guide walks you through creating a series of TinyML projects, step-by-step. No machine learning or microcontroller experience is necessary. Build a speech recognizer, a camera that detects people, and a magic wand that responds to gestures Work with Arduino and ultra-low-power microcontrollers Learn the essentials of ML and how to train your own models Train models to understand audio, image, and accelerometer data Explore TensorFlow Lite for Microcontrollers, Google's toolkit for TinyML Debug applications and provide safeguards for privacy and security Optimize latency, energy usage, and model and binary size

The new edition of an introductory text that teaches students the art of computational problem solving, covering topics ranging from simple algorithms to information visualization. This book introduces students with little or no prior programming experience to the art of computational problem solving using Python and various Python libraries, including PyLab. It provides students with skills that will enable them to make productive use of computational techniques, including some of the tools and techniques of data science for using computation to model and interpret data. The book is based on an MIT course (which became the most popular course offered through MIT's OpenCourseWare) and was developed for use not only in a conventional classroom but in in a massive open online course (MOOC). This new edition has been updated for Python 3, reorganized to make it easier to use for courses that cover only a subset of the material, and offers additional material including five new chapters. Students are introduced to Python and the basics of programming in the context of such computational concepts and techniques as exhaustive enumeration, bisection search, and efficient approximation algorithms. Although it covers such traditional topics as computational complexity and simple algorithms, the book focuses on a wide range of topics not found in most introductory texts, including information visualization, simulations to model randomness, computational techniques to understand data, and statistical techniques that inform (and misinform) as well as two related but relatively advanced topics: optimization problems and dynamic programming. This edition offers expanded material on statistics and machine learning and new chapters on Frequentist and Bayesian statistics.

"Optimizing and boosting your Python programming"--Cover.

Praise for the First Edition ". . . [this book] should be on the shelf of everyone interested in . . . longitudinal data analysis." —Journal of the American Statistical Association Features newly developed topics and applications of the analysis of longitudinal data Applied Longitudinal Analysis, Second Edition presents modern methods for analyzing data from longitudinal studies and

now features the latest state-of-the-art techniques. The book emphasizes practical, rather than theoretical, aspects of methods for the analysis of diverse types of longitudinal data that can be applied across various fields of study, from the health and medical sciences to the social and behavioral sciences. The authors incorporate their extensive academic and research experience along with various updates that have been made in response to reader feedback. The Second Edition features six newly added chapters that explore topics currently evolving in the field, including: Fixed effects and mixed effects models Marginal models and generalized estimating equations Approximate methods for generalized linear mixed effects models Multiple imputation and inverse probability weighted methods Smoothing methods for longitudinal data Sample size and power Each chapter presents methods in the setting of applications to data sets drawn from the health sciences. New problem sets have been added to many chapters, and a related website features sample programs and computer output using SAS, Stata, and R, as well as data sets and supplemental slides to facilitate a complete understanding of the material. With its strong emphasis on multidisciplinary applications and the interpretation of results, *Applied Longitudinal Analysis, Second Edition* is an excellent book for courses on statistics in the health and medical sciences at the upper-undergraduate and graduate levels. The book also serves as a valuable reference for researchers and professionals in the medical, public health, and pharmaceutical fields as well as those in social and behavioral sciences who would like to learn more about analyzing longitudinal data.

A fully updated tutorial on the basics of the Python programming language for science students Python is a computer programming language that has gained popularity throughout the sciences. This fully updated second edition of *A Student's Guide to Python for Physical Modeling* aims to help you, the student, teach yourself enough of the Python programming language to get started with physical modeling. You will learn how to install an open-source Python programming environment and use it to accomplish many common scientific computing tasks: importing, exporting, and visualizing data; numerical analysis; and simulation. No prior programming experience is assumed. This guide introduces a wide range of useful tools, including: Basic Python programming and scripting Numerical arrays Two- and three-dimensional graphics Animation Monte Carlo simulations Numerical methods, including solving ordinary differential equations Image processing Numerous code samples and exercises—with solutions—illustrate new ideas as they are introduced. This guide also includes supplemental online resources: code samples, data sets, tutorials, and more. This edition includes new material on symbolic calculations with SymPy, an introduction to Python libraries for data science and machine learning (pandas and sklearn), and a primer on Python classes and object-oriented programming. A new appendix also introduces command line tools and version control with Git.

Treat yourself to a lively, intuitive, and easy-to-follow introduction to computer programming in Python. The book was written specifically for biologists with little or no prior experience of writing code - with the goal of giving them not only a foundation in Python programming, but also the confidence and inspiration to start using Python in their own research. Virtually all of the examples in the book are drawn from across a wide spectrum of life science research, from simple biochemical calculations and sequence analysis, to modeling the dynamic interactions of genes and proteins in cells, or the drift of genes in an evolving population. Best of all, *Python for the Life Sciences* shows you how to implement all of these projects in Python, one of the most popular programming languages for scientific computing. If you are a life scientist interested in learning Python to jump-start your research, this is the book for you. What You'll Learn Write Python scripts to automate your lab calculations Search for important motifs in genome sequences Use object-oriented programming with Python Study mining interaction network data for patterns Review dynamic modeling of biochemical switches Who This Book Is For Life scientists with little or no programming experience, including undergraduate and graduate students, postdoctoral researchers in academia and industry, medical professionals, and teachers/lecturers. "A comprehensive introduction to using Python for computational biology... A lovely book with humor and perspective" -- John Novembre, Associate Professor of Human Genetics, University of Chicago and MacArthur Fellow "Fun, entertaining, witty and darn useful. A magical portal to the big data revolution" -- Sandro Santagata, Assistant Professor in Pathology, Harvard Medical School "Alex and Gordon's enthusiasm for Python is contagious" -- Glenys Thomson Professor of Integrative Biology, University of California, Berkeley

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

"This manual is part of the official reference documentation for Python, an object-oriented programming language created by Guido van Rossum. Python is free software. The term "free software" refers to your freedom to run, copy, distribute, study, change and improve the software. With Python you have all these freedoms. You can support free software by becoming an associate member of the Free Software Foundation. The Free Software Foundation is a tax-exempt charity dedicated to promoting

the right to use, study, copy, modify, and redistribute computer programs. It also helps to spread awareness of the ethical and political issues of freedom in the use of software. For more information visit the website www.fsf.org. The development of Python itself is supported by the Python Software Foundation. Companies using Python can invest in the language by becoming sponsoring members of this group. Donations can also be made online through the Python website. Further information is available at <http://www.python.org/psf/>."--Page 1.

This fast-paced introduction to Python moves from the basics to advanced concepts, enabling readers to gain proficiency quickly. Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With `fastai`, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of `fastai`, show you how to train a model on a wide range of tasks using `fastai` and `PyTorch`. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by `PyTorch` cofounder, Soumith Chintala

Artificial Intelligence with PythonPackt Publishing Ltd

The application of causal inference methods is growing exponentially in fields that deal with observational data. Written by pioneers in the field, this practical book presents an authoritative yet accessible overview of the methods and applications of causal inference. With a wide range of detailed, worked examples using real epidemiologic data as well as software for replicating the analyses, the text provides a thorough introduction to the basics of the theory for non-time-varying treatments and the generalization to complex longitudinal data.

A fully updated tutorial on the basics of the Python programming language for science students Python is a computer programming language that is rapidly gaining popularity throughout the sciences. This fully updated edition of *A Student's Guide to Python for Physical Modeling* aims to help you, the student, teach yourself enough of the Python programming language to get started with physical modeling. You will learn how to install an open-source Python programming environment and use it to accomplish many common scientific computing tasks: importing, exporting, and visualizing data; numerical analysis; and simulation. No prior programming experience is assumed. This tutorial focuses on fundamentals and introduces a wide range of useful techniques, including: Basic Python programming and scripting Numerical arrays Two- and three-dimensional graphics Monte Carlo simulations Numerical methods, including solving ordinary differential equations Image processing Animation Numerous code samples and exercises--with solutions--illustrate new ideas as they are introduced. Web-based resources also accompany this guide and include code samples, data sets, and more. This current edition brings the discussion of the Python language, Spyder development environment, and Anaconda distribution up to date. In addition, a new appendix introduces Jupyter notebooks.

"Princeton University Press published Imai's textbook, *Quantitative Social Science: An Introduction*, an introduction to quantitative methods and data science for upper level undergrads and graduates in professional programs, in February 2017. What is distinct about the book is how it leads students through a series of applied examples of statistical methods, drawing on real examples from social science research. The original book was prepared with the statistical software R, which is freely available online and has gained in popularity in recent years. But many existing courses in statistics and data sciences, particularly in some subject areas like sociology and law, use STATA, another general purpose package that has been the market leader since the 1980s. We've had several requests for STATA versions of the text as many programs use it by default. This is a "translation" of the original text, keeping all the current pedagogical text but inserting the necessary code and outputs from STATA in their place"--

Looks at ways to prevent and treat such disorders as dyspepsia, reflux disease, irritable bowel syndrome, constipation, and diarrhea.

I designed a learning system for myself that quadrupled my aptitude for learning computer languages. It worked so well for me that I've used it to teach coding to grandmothers, cab drivers, musicians, and 50,000 other newbies. Washington University research shows that a key teaching method I use--interactive recall practice--improves learning performance 400 percent. Computer languages are not inherently hard to understand, even for non-techies. Remembering is the problem. Research shows that you will remember everything if you're repeatedly asked to recall it. That's the beauty of flash cards. But technology offers an even better way to make information stick. With my book you get almost a thousand interactive exercises--they're free online--that embed the whole book in your memory. Algorithms check your work to make sure you know what you think you know. When you stumble, you do the exercise again. You keep trying until you know the chapter cold. The exercises keep you engaged, give you extra practice where you're shaky, and prepare you for each next step. Every lesson is built on top of a solid foundation that you and I have carefully constructed. Each individual step is small. But all the little steps add up to real knowledge--knowledge that you retain. You don't need to be a computer genius to learn Python. You just need to be smart about how you learn it.--Amazon.com description.

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

A gargantuan, mind-altering comedy about the Pursuit of Happiness in America Set in an addicts' halfway house and a tennis academy, and featuring the most endearingly screwed-up family to come along in recent fiction, *Infinite Jest* explores essential questions about what entertainment is and why it has come to so dominate our lives; about how our desire for entertainment affects our need to connect with other people; and about what the pleasures we choose say about who we are. Equal parts philosophical quest and screwball comedy, *Infinite Jest* bends every rule of fiction without sacrificing for a moment its own entertainment value. It is an exuberant, uniquely American exploration of the passions that make us human - and one of those rare books that renew the idea of what a novel can do. "The next step in fiction...Edgy, accurate, and darkly witty...Think Beckett, think Pynchon, think Gaddis. Think." --Sven Birkerts, *The Atlantic*

An introduction to coding for complete beginners, this friendly and accessible book will teach children the basics of Python (a widely used programming language), allowing them to get inside the code of their computer and create simple games and animations on screen.

One of the attractive aspects of C++ is that it offers good facilities for object-oriented programming (OOP), but, as a hybrid language, it also supports procedural programming. The significance of this for programmers is that it offers more flexibility allowing them to shift to object-oriented programming if and when they feel the need to do so. In this regard, C++ differs from some purely object-oriented languages, such as Smalltalk, Eiffel and Java. This book offers practical guidance on how to programme in both styles. The C++ language and its standard library have gone through a good many improvements and extensions during their evolution. This third edition has therefore been completely revised in accordance with the C++ language revision, which is embodied in the ANSI/ISO C++ Standard. For example, the new, important type string is used throughout the book and the Standard Template Library (STL) is introduced to readers at an early stage and discussed in more detail later on. All example programs and the solutions to the exercises can be downloaded from the website. <http://home.wxs.nl/~ammeraal/> Solutions for some of these exercises can be found in the appendix.

Python is an easy to learn, powerful programming language. It has efficient high-level data structures and a simple but effective approach to object-oriented programming. Python's elegant syntax and dynamic typing, together with its interpreted nature, make it an ideal language for scripting and rapid application development in many areas on most platforms. The Python interpreter and the extensive standard library are freely available in source or binary form for all major platforms from the Python Web site, <https://www.python.org/>, and may be freely distributed. The same site also contains distributions of and pointers to many free third party Python modules, programs and tools, and additional documentation. The Python interpreter is easily extended with new functions and data types implemented in C or C++ (or other languages callable from C). Python is also suitable as an extension language for customizable applications. This tutorial introduces the reader informally to the basic concepts and features of the python language and system. It helps to have a Python interpreter handy for hands-on experience, but all examples are self contained, so the tutorial can be read off-line as well. For a description of standard objects and modules, see [library-index](#). [reference-index](#) gives a more formal definition of the language. To write extensions in C or C++, read [extending-index](#) and [c-api-index](#). There are also several books covering Python in depth. This tutorial does not attempt to be comprehensive and cover every single feature, or even every commonly used feature. Instead, it introduces many of Python's most noteworthy features, and will give you a good idea of the language's flavor and style. After reading it, you will be able to read and write Python modules and programs, and you will be ready to learn more about the various Python library modules described in [library-index](#). The Glossary is also worth going through.

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 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

"Have you always wanted to learn computer programming but are afraid it'll be too difficult for you? Or perhaps you know other programming languages but are interested in learning the Python language fast? This book is for you"--Page 4 of cover.

Explore the mathematical computations and algorithms for image processing using popular Python tools and frameworks. Key Features Practical coverage of every image processing task with popular Python libraries Includes topics such as pseudo-coloring, noise smoothing, computing image descriptors Covers popular machine learning and deep learning techniques for complex image processing tasks Book Description Image processing plays an important role in our daily lives with various applications such as in social media (face detection), medical imaging (X-ray, CT-scan), security (fingerprint recognition) to robotics & space. This book will touch the core of image processing, from concepts to code using Python. The book will start from the classical image processing techniques and explore the evolution of image processing algorithms up to the recent advances in image processing or computer vision with deep learning. We will learn how to use image processing libraries such as PIL, scikit-image, and scipy ndimage in Python. This book will enable us to write code snippets in Python 3 and quickly implement complex image processing algorithms such as image enhancement, filtering, segmentation, object detection, and classification. We will be able to use machine learning models using the scikit-learn library and later explore deep CNN, such as VGG-19 with Keras, and we will also use an end-to-end deep learning model called YOLO for object detection. We will also cover a few advanced problems, such as

image inpainting, gradient blending, variational denoising, seam carving, quilting, and morphing. By the end of this book, we will have learned to implement various algorithms for efficient image processing. What you will learn Perform basic data pre-processing tasks such as image denoising and spatial filtering in Python Implement Fast Fourier Transform (FFT) and Frequency domain filters (e.g., Weiner) in Python Do morphological image processing and segment images with different algorithms Learn techniques to extract features from images and match images Write Python code to implement supervised / unsupervised machine learning algorithms for image processing Use deep learning models for image classification, segmentation, object detection and style transfer Who this book is for This book is for Computer Vision Engineers, and machine learning developers who are good with Python programming and want to explore details and complexities of image processing. No prior knowledge of the image processing techniques is expected.

Developed from celebrated Harvard statistics lectures, Introduction to Probability provides essential language and tools for understanding statistics, randomness, and uncertainty. The book explores a wide variety of applications and examples, ranging from coincidences and paradoxes to Google PageRank and Markov chain Monte Carlo (MCMC).

Additional

This book covers several of the statistical concepts and data analytic skills needed to succeed in data-driven life science research. The authors proceed from relatively basic concepts related to computed p-values to advanced topics related to analyzing highthroughput data. They include the R code that performs this analysis and connect the lines of code to the statistical and mathematical concepts explained.

Learn how to develop intelligent applications with sequential learning and apply modern methods for language modeling with neural network architectures for deep learning with Python's most popular TensorFlow framework. Key Features Train and deploy Recurrent Neural Networks using the popular TensorFlow library Apply long short-term memory units Expand your skills in complex neural network and deep learning topics Book Description Developers struggle to find an easy-to-follow learning resource for implementing Recurrent Neural Network (RNN) models. RNNs are the state-of-the-art model in deep learning for dealing with sequential data. From language translation to generating captions for an image, RNNs are used to continuously improve results. This book will teach you the fundamentals of RNNs, with example applications in Python and the TensorFlow library. The examples are accompanied by the right combination of theoretical knowledge and real-world implementations of concepts to build a solid foundation of neural network modeling. Your journey starts with the simplest RNN model, where you can grasp the fundamentals. The book then builds on this by proposing more advanced and complex algorithms. We use them to explain how a typical state-of-the-art RNN model works. From generating text to building a language translator, we show how some of today's most powerful AI applications work under the hood. After reading the book, you will be confident with the fundamentals of RNNs, and be ready to pursue further study, along with developing skills in this exciting field. What you will learn Use TensorFlow to build RNN models Use the correct RNN architecture for a particular machine learning task Collect and clear the training data for your models Use the correct Python libraries for any task during the building phase of your model Optimize your model for higher accuracy Identify the differences between multiple models and how you can substitute them Learn the core deep learning fundamentals applicable to any machine learning model Who this book is for This book is for Machine Learning engineers and data scientists who want to learn about Recurrent Neural Network models with practical use-cases. Exposure to Python programming is required. Previous experience with TensorFlow will be helpful, but not mandatory.

This educational book introduces emerging developers to computer programming through the Python software development language, and serves as a reference book for experienced developers looking to learn a new language or re-familiarize themselves with computational logic and syntax.

Python for Software Design is a concise introduction to software design using the Python programming language. The focus is on the programming process, with special emphasis on debugging. The book includes a wide range of exercises, from short examples to substantial projects, so that students have ample opportunity to practice each new concept.

A new edition of a graduate-level machine learning textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern topics in machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes several key aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first four chapters lay the theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct (PAC) learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel methods; boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata and languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise probability review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New material in the appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

A surprisingly simple way for students to master any subject--based on one of the world's most popular online courses and the bestselling book A Mind for Numbers A Mind for Numbers and its wildly popular online companion course "Learning How to Learn" have empowered more than two million learners of all ages from around the world to master subjects that they once struggled with. Fans often wish they'd discovered these learning strategies earlier and ask how they can help their kids master these skills as well. Now in this new book for kids and teens, the authors reveal how to make the most of time spent studying. We all have the tools to learn what might not seem to come naturally to us at first--the secret is to understand how the brain works so we can unlock its power. This book explains: • Why sometimes letting your mind wander is an important part of the learning process • How to avoid "rut think" in order to think outside the box • Why having a poor memory can be a good thing • The value of metaphors in developing understanding • A simple, yet powerful, way to stop procrastinating Filled with illustrations, application questions, and exercises, this book makes learning easy and fun.

The programming language Python was conceived in the late 1980s, [1] and its implementation was started in December 1989[2] by Guido

van Rossum at CWI in the Netherlands as a successor to the ABC (programming language) capable of exception handling and interfacing with the Amoeba operating system.[3] Van Rossum is Python's principal author, and his continuing central role in deciding the direction of Python is reflected in the title given to him by the Python community, Benevolent Dictator for Life (BDFL).[4][5] Python was named for the BBC TV show Monty Python's Flying Circus.[6] Python 2.0 was released on October 16, 2000, with many major new features, including a cycle-detecting garbage collector (in addition to reference counting) for memory management and support for Unicode. However, the most important change was to the development process itself, with a shift to a more transparent and community-backed process.[7] Python 3.0, a major, backwards-incompatible release, was released on December 3, 2008[8] after a long period of testing. Many of its major features have also been backported to the backwards-compatible Python 2.6 and 2.7.[9] In February 1991, van Rossum published the code (labeled version 0.9.0) to alt.sources.[10] Already present at this stage in development were classes with inheritance, exception handling, functions, and the core datatypes of list, dict, str and so on. Also in this initial release was a module system borrowed from Modula-3; Van Rossum describes the module as "one of Python's major programming units." [1] Python's exception model also resembles Modula-3's, with the addition of an else clause.[3] In 1994 comp.lang.python, the primary discussion forum for Python, was formed, marking a milestone in the growth of Python's userbase.[1] Python reached version 1.0 in January 1994. The major new features included in this release were the functional programming tools lambda, map, filter and reduce. Van Rossum stated that "Python acquired lambda, reduce(), filter() and map(), courtesy of a Lisp hacker who missed them and submitted working patches." [11] The last version released while Van Rossum was at CWI was Python 1.2. In 1995, Van Rossum continued his work on Python at the Corporation for National Research Initiatives (CNRI) in Reston, Virginia whence he released several versions. By version 1.4, Python had acquired several new features. Notable among these are the Modula-3 inspired keyword arguments (which are also similar to Common Lisp's keyword arguments) and built-in support for complex numbers. Also included is a basic form of data hiding by name mangling, though this is easily bypassed.[12] During Van Rossum's stay at CNRI, he launched the Computer Programming for Everybody (CP4E) initiative, intending to make programming more accessible to more people, with a basic "literacy" in programming languages, similar to the basic English literacy and mathematics skills required by most employers. Python served a central role in this: because of its focus on clean syntax, it was already suitable, and CP4E's goals bore similarities to its predecessor, ABC. The project was funded by DARPA.[13] As of 2007, the CP4E project is inactive, and while Python attempts to be easily learnable and not too arcane in its syntax and semantics, reaching out to non-programmers is not an active concern.[14] Here are what people are saying about the book: This is the best beginner's tutorial I've ever seen! Thank you for your effort. -- Walt Michalik The best thing i found was "A Byte of Python," which is simply a brilliant book for a beginner. It's well written, the concepts are well explained with self evident examples. -- Joshua Robin Excellent gentle introduction to programming #Python for beginners -- Shan Rajasekaran Best newbie guide to python -- Nickson Kaigi start to love python with every single page read -- Herbert Feutl perfect beginners guide for python, will give u key to unlock magical world of python

Combining GIS concepts and fundamental spatial thinking methodology with real programming examples, this book introduces popular Python-based tools and their application to solving real-world problems. It elucidates the programming constructs of Python with its high-level toolkits and demonstrates its integration with ArcGIS Theory. Filled with hands-on computer exercises in a logical learning workflow this book promotes increased interactivity between instructors and students while also benefiting professionals in the field with vital knowledge to sharpen their programming skills. Readers receive expert guidance on modules, package management, and handling shapefile formats needed to build their own mini-GIS. Comprehensive and engaging commentary, robust contents, accompanying datasets, and classroom-tested exercises are all housed here to permit users to become competitive in the GIS/IT job market and industry.

Learn web scraping and crawling techniques to access unlimited data from any web source in any format. With this practical guide, you'll learn how to use Python scripts and web APIs to gather and process data from thousands—or even millions—of web pages at once. Ideal for programmers, security professionals, and web administrators familiar with Python, this book not only teaches basic web scraping mechanics, but also delves into more advanced topics, such as analyzing raw data or using scrapers for frontend website testing. Code samples are available to help you understand the concepts in practice. Learn how to parse complicated HTML pages Traverse multiple pages and sites Get a general overview of APIs and how they work Learn several methods for storing the data you scrape Download, read, and extract data from documents Use tools and techniques to clean badly formatted data Read and write natural languages Crawl through forms and logins Understand how to scrape JavaScript Learn image processing and text recognition

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.

[Copyright: d25aca80b58f62faf8d0223a66f3c845](https://www.d25aca80b58f62faf8d0223a66f3c845)