

High Performance Python Practical Performant Programming For Humans

Your Python code may run correctly, but you need it to run faster. Updated for Python 3, this expanded edition shows you how to locate performance bottlenecks and significantly speed up your code in high-data-volume programs. By exploring the fundamental theory behind design choices, High Performance Python helps you gain a deeper understanding of Python's implementation. How do you take advantage of multicore architectures or clusters? Or build a system that scales up and down without losing reliability? Experienced Python programmers will learn concrete solutions to many issues, along with war stories from companies that use high-performance Python for social media analytics, productionized machine learning, and more. Get a better grasp of NumPy, Cython, and profilers Learn how Python abstracts the underlying computer architecture Use profiling to find bottlenecks in CPU time and memory usage Write efficient programs by choosing appropriate data structures Speed up matrix and vector computations Use tools to compile Python down to machine code Manage multiple I/O and computational operations concurrently Convert multiprocessing code to run on local or remote clusters Deploy code faster using tools like Docker

With the resurgence of neural networks in the 2010s, deep learning has become essential for machine learning practitioners and even many software engineers. This book provides a comprehensive introduction for data scientists and software engineers with machine learning experience. You'll start with deep learning basics and move quickly to the details of important advanced architectures, implementing everything from scratch along the way. Author Seth Weidman shows you how neural networks work using a first principles approach. You'll learn how to apply multilayer neural networks, convolutional neural networks, and recurrent neural networks from the ground up. With a thorough understanding of how neural networks work mathematically, computationally, and conceptually, you'll be set up for success on all future deep learning projects. This book provides: Extremely clear and thorough mental models—accompanied by working code examples and mathematical explanations—for understanding neural networks Methods for implementing multilayer neural networks from scratch, using an easy-to-understand object-oriented framework Working implementations and clear-cut explanations of convolutional and recurrent neural networks Implementation of these neural network concepts using the popular PyTorch framework

Easy to understand and fun to read, this updated edition of Introducing Python is ideal for beginning programmers as well as those new to the language. Author Bill Lubanovic takes you from the basics to more involved and varied topics, mixing tutorials with cookbook-style code recipes to explain concepts in Python 3. End-of-chapter exercises help you practice what you've learned. You'll gain a strong foundation in the language,

including best practices for testing, debugging, code reuse, and other development tips. This book also shows you how to use Python for applications in business, science, and the arts, using various Python tools and open source packages.

The Hitchhiker's Guide to Python takes the journeyman Pythonista to true expertise. More than any other language, Python was created with the philosophy of simplicity and parsimony. Now 25 years old, Python has become the primary or secondary language (after SQL) for many business users. With popularity comes diversity—and possibly dilution. This guide, collaboratively written by over a hundred members of the Python community, describes best practices currently used by package and application developers. Unlike other books for this audience, The Hitchhiker's Guide is light on reusable code and heavier on design philosophy, directing the reader to excellent sources that already exist.

Advanced R

High Performance Spark

Best Practices for Scaling and Optimizing Apache Spark

Robust Python

The Big Ideas Behind Reliable, Scalable, and Maintainable Systems

Explore high-performance parallel computing with CUDA

Discover the right way to code in Python. This book provides the tips and techniques you need to produce cleaner, error-free, and eloquent Python projects. Your journey to better code starts with understanding the importance of formatting and documenting your code for maximum readability, utilizing built-in data structures and Python dictionary for improved maintainability, and working with modules and meta-classes to effectively organize your code. You will then dive deep into the new features of the Python language and learn how to effectively utilize them. Next, you will decode key concepts such as asynchronous programming, Python data types, type hinting, and path handling. Learn tips to debug and conduct unit and integration tests in your Python code to ensure your code is ready for production. The final leg of your learning journey equips you with essential tools for version management, managing live code, and intelligent code completion. After reading and using this book, you will be proficient in writing clean Python code and successfully apply these principles to your own Python projects. What You'll Learn Use

the right expressions and statements in your Python code Create and assess Python Dictionary Work with advanced data structures in Python Write better modules, classes, functions, and metaclasses Start writing asynchronous Python immediately Discover new features in Python Who This Book Is For Readers with a basic Python programming knowledge who want to improve their Python programming skills by learning right way to code in Python.

Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

This book constitutes revised selected papers from 7 workshops that were held in conjunction with the ISC High Performance 2016 conference in Frankfurt, Germany, in June 2016. The 45 papers presented in this volume were carefully reviewed and selected for inclusion in this book. They stem from the following workshops: Workshop on Exascale Multi/Many Core Computing Systems, E-MuCoCoS; Second International Workshop on Communication Architectures at Extreme Scale, ExaComm; HPC I/O in the Data Center Workshop, HPC-IODC; International Workshop on OpenPOWER for HPC, IWOPH; Workshop on the Application Performance on Intel Xeon Phi - Being Prepared for KNL and Beyond, IXPUG; Workshop on Performance and Scalability of Storage Systems, WOPSSS; and International

Workshop on Performance Portable Programming Models for Accelerators, P3MA. Build real-world applications with Python 2.7, CUDA 9, and CUDA 10. We suggest the use of Python 2.7 over Python 3.x, since Python 2.7 has stable support across all the libraries we use in this book. Key Features Expand your background in GPU programming—PyCUDA, scikit-cuda, and Nsight Effectively use CUDA libraries such as cuBLAS, cuFFT, and cuSolver Apply GPU programming to modern data science applications Book Description Hands-On GPU Programming with Python and CUDA hits the ground running: you'll start by learning how to apply Amdahl's Law, use a code profiler to identify bottlenecks in your Python code, and set up an appropriate GPU programming environment. You'll then see how to "query" the GPU's features and copy arrays of data to and from the GPU's own memory. As you make your way through the book, you'll launch code directly onto the GPU and write full blown GPU kernels and device functions in CUDA C. You'll get to grips with profiling GPU code effectively and fully test and debug your code using Nsight IDE. Next, you'll explore some of the more well-known NVIDIA libraries, such as cuFFT and cuBLAS. With a solid background in place, you will now apply your new-found knowledge to develop your very own GPU-based deep neural network from scratch. You'll then explore advanced topics, such as warp shuffling, dynamic parallelism, and PTX assembly. In the final chapter, you'll see some topics and applications related to GPU programming that you may wish to pursue, including AI, graphics, and blockchain. By the end of this book, you will be able to apply GPU programming to problems related to data science and high-performance computing. What you will learn Launch GPU code directly from Python Write effective and efficient GPU kernels and device functions Use libraries such as cuFFT, cuBLAS, and cuSolver Debug and profile your code with Nsight and Visual Profiler Apply GPU programming to datascience problems Build a GPU-based deep neuralnetwork from scratch Explore advanced GPU hardware features, such as warp shuffling Who this book is for Hands-On GPU Programming with Python and CUDA is for developers and data scientists who want to learn the basics of effective GPU programming to improve performance using Python code. You should have an understanding of first-year college or university-level engineering mathematics and physics, and have some experience with Python as well as in any C-based

programming language such as C, C++, Go, or Java.

The Hitchhiker's Guide to Python

Introducing Python

Using Asyncio in Python

ISC High Performance 2016 International Workshops, ExaComm, E-MuCoCoS, HPC-IODC, IXPUG, IWOPH, P³MA, VHPC, WOPSSS, Frankfurt, Germany, June 19-23, 2016, Revised Selected Papers

Building Data Science Applications with FastAPI

Beginning Python Games Development, Second Edition

High Performance MySQL

Break into the powerful world of parallel GPU programming with this down-to-earth, practical guide Designed for professionals across multiple industrial sectors, Professional CUDA C Programming presents CUDA -- a parallel computing platform and programming model designed to ease the development of GPU programming -- fundamentals in an easy-to-follow format, and teaches readers how to think in parallel and implement parallel algorithms on GPUs. Each chapter covers a specific topic, and includes workable examples that demonstrate the development process, allowing readers to explore both the "hard" and "soft" aspects of GPU programming. Computing architectures are experiencing a fundamental shift toward scalable parallel computing motivated by application requirements in industry and science. This book demonstrates the challenges of efficiently utilizing compute resources at peak performance, presents modern techniques for tackling these challenges, while increasing accessibility for professionals who are not necessarily parallel programming experts. The CUDA programming model and tools empower developers to write high-performance applications on a scalable, parallel computing platform: the GPU. However, CUDA itself can be difficult to learn without extensive programming experience. Recognized CUDA authorities John Cheng, Max Grossman, and Ty McKercher guide readers through essential GPU programming skills and best practices in Professional CUDA C Programming, including: CUDA Programming Model GPU Execution Model GPU Memory model Streams, Event and Concurrency Multi-GPU Programming CUDA Domain-Specific Libraries Profiling and Performance Tuning The book makes complex CUDA concepts easy to understand for anyone with knowledge of basic software development with exercises designed to be both readable and high-performance. For the professional seeking entrance to parallel computing and the high-performance computing community, Professional CUDA C Programming is an invaluable

resource, with the most current information available on the market.

Boost the performance of your Haskell applications using optimization, concurrency, and parallel programming About This Book Explore the benefits of lazy evaluation, compiler features, and tools and libraries designed for high performance Write fast programs at extremely high levels of abstraction Work through practical examples that will help you address the challenges of writing efficient code Who This Book Is For To get the most out of this book, you need to have a working knowledge of reading and writing basic Haskell. No knowledge of performance, optimization, or concurrency is required. What You Will Learn Program idiomatic Haskell that's also surprisingly efficient Improve performance of your code with data parallelism, inlining, and strictness annotations Profile your programs to identify space leaks and missed opportunities for optimization Find out how to choose the most efficient data and control structures Optimize the Glasgow Haskell Compiler and runtime system for specific programs See how to smoothly drop to lower abstractions wherever necessary Execute programming for the GPU with Accelerate Implement programming to easily scale to the cloud with Cloud Haskell In Detail Haskell, with its power to optimize the code and its high performance, is a natural candidate for high performance programming. It is especially well suited to stacking abstractions high with a relatively low performance cost. This book addresses the challenges of writing efficient code with lazy evaluation and techniques often used to optimize the performance of Haskell programs. We open with an in-depth look at the evaluation of Haskell expressions and discuss optimization and benchmarking. You will learn to use parallelism and we'll explore the concept of streaming. We'll demonstrate the benefits of running multithreaded and concurrent applications. Next we'll guide you through various profiling tools that will help you identify performance issues in your program. We'll end our journey by looking at GPGPU, Cloud and Functional Reactive Programming in Haskell. At the very end there is a catalogue of robust library recommendations with code samples. By the end of the book, you will be able to boost the performance of any app and prepare it to stand up to real-world punishment. Style and approach This easy-to-follow guide teaches new practices and techniques to optimize your code, and then moves towards more advanced ways to effectively write efficient Haskell code. Small and simple practical examples will help you test the concepts yourself, and you will be able to easily adapt them for any application.

As Python continues to grow in popularity, projects are becoming larger and more complex. Many Python developers are now taking an interest in high-level software design patterns such as

hexagonal/clean architecture, event-driven architecture, and the strategic patterns prescribed by domain-driven design (DDD). But translating those patterns into Python isn't always straightforward. With this hands-on guide, Harry Percival and Bob Gregory from MADE.com introduce proven architectural design patterns to help Python developers manage application complexity—and get the most value out of their test suites. Each pattern is illustrated with concrete examples in beautiful, idiomatic Python, avoiding some of the verbosity of Java and C# syntax. Patterns include: Dependency inversion and its links to ports and adapters (hexagonal/clean architecture) Domain-driven design's distinction between entities, value objects, and aggregates Repository and Unit of Work patterns for persistent storage Events, commands, and the message bus Command-query responsibility segregation (CQRS) Event-driven architecture and reactive microservices

Does it seem like your Python projects are getting bigger and bigger? Are you feeling the pain as your codebase expands and gets tougher to debug and maintain? Python is an easy language to learn and use, but that also means systems can quickly grow beyond comprehension. Thankfully, Python has features to help developers overcome maintainability woes. In this practical book, author Patrick Viafore shows you how to use Python's type system to the max. You'll look at user-defined types, such as classes and enums, and Python's type hinting system. You'll also learn how to make Python extensible and how to use a comprehensive testing strategy as a safety net. With these tips and techniques, you'll write clearer and more maintainable code. Learn why types are essential in modern development ecosystems Understand how type choices such as classes, dictionaries, and enums reflect specific intents Make Python extensible for the future without adding bloat Use popular Python tools to increase the safety and robustness of your codebase Evaluate current code to detect common maintainability gotchas Build a safety net around your codebase with linters and tests

The Python Book

Python High Performance

Understanding Python's Asynchronous Programming Features

The Art of Scientific Python

Learn Ruthlessly Effective Automation

Architecture Patterns with Python

Practices of the Python Pro

Welcome to Scientific Python and its community. If you're a scientist who programs with Python, this practical guide not only teaches

Where To Download High Performance Python Practical Performant Programming For Humans

parts of SciPy and libraries related to it, but also gives you a taste for beautiful, easy-to-read code that you can use in practice. You'll see elegant code that's clear, concise, and efficient at executing the task at hand. Throughout the book, you'll work with examples from the Python ecosystem, using code that illustrates principles outlined in the book. Using actual scientific data, you'll work on real-world problems with NumPy, Pandas, scikit-image, and other Python libraries. Explore the NumPy array, the data structure that underlies numerical scientific computing. Perform quantile normalization to ensure that measurements fit a specific distribution. Represent separate regions in an image with a RegionAdapter. Convert temporal or spatial data into frequency domain data with the Fast Fourier Transform. Solve sparse matrix problems, including image segmentations, with SciPy's sparse module. Perform linear algebra by using SciPy packages. Explore image alignment (registration) with the skimage module. Process large datasets with Python data streaming primitives and the Toolz library.

The second edition of this best-selling Python book (over 500,000 copies sold!) uses Python 3 to teach even the technically uninclined how to write programs that do in minutes what would take hours to do by hand. There is no prior programming experience required and the book is accessible to novices and majors and geeks alike. If you've ever spent hours renaming files or updating hundreds of spreadsheet cells, you know how tedious these tasks can be. But what if you could have your computer do them for you? In this fully revised second edition of the best-selling classic *Automate the Boring Stuff with Python*, you'll learn how to use Python to write programs that do in minutes what would take you hours to do by hand--no prior programming experience required. You'll learn the basics of Python and explore Python's rich library of modules for performing specific tasks, like scraping data off websites, reading PDF and Word documents, and automating clicking and typing tasks. The second edition of this international fan favorite includes a new chapter on input validation, as well as tutorials on automating Gmail and Google Sheets, plus tips on automatically updating CSV files. You'll learn how to create programs that effortlessly perform useful feats of automation to:

- Search for text in a file or across multiple files
- Create, update, and rename files and folders
- Search the Web and download online content
- Update and format data in Excel spreadsheets of any size
- Send a watermark, and encrypt PDFs
- Send email responses and text notifications
- Fill out online forms

Step-by-step instructions walk you through each program, and updated practice projects at the end of each chapter challenge you to improve those programs and use your newfound skills to solve similar tasks. Don't spend your time doing work a well-trained monkey could do. Even if you've never written a line of code, you can make your computer do the grunt work. Learn how in *Automate the Boring Stuff with Python, 2nd Edition*.

An Essential Reference for Intermediate and Advanced R Programmers *Advanced R* presents useful tools and techniques for attacking non-trivial programming problems, helping you avoid mistakes and dead ends. With more than ten years of experience programming in R, the author shares the elegance, beauty, and flexibility at the heart of R. The book develops the necessary skills to produce quality code that can be used in a wide range of circumstances. You will learn:

- The fundamentals of R, including standard data types and functions
- Functional programming as a useful framework for solving wide classes of problems
- The positives and negatives of metaprogramming
- How to write fast, memory-efficient code

This book is for current R users who want to become R programmers but also shows existing programmers what's special about R. Intermediate R programmers can learn new strategies for solving diverse problems while programmers from other languages can learn the details of R and understand the way it does.

High Performance Python Practical Performant Programming for Humans O'Reilly Media

Learning Java

Automate the Boring Stuff with Python, 2nd Edition

Python for DevOps

Clean Python

HIGH PERFORMANCE PYTHON.

Cython

Deep Learning from Scratch

The agile development movement represents the latest advances in tools and techniques intended to boost developer productivity. This is the first book to apply these sought after principles to Python developers, introducing both the tools and techniques built and supported by the Python community. Authored by Jeff Younker, who is perhaps best known for his creation of a popular Python testing framework, this book is sure to be a hit among readers who may have reached their limits of knowledge regarding the Python language, yet are seeking to improve their understanding of how sound processes can boost productivity to unparalleled heights.

Summary Professional developers know the many benefits of writing application code that's clean, well-organized, and easy to maintain. By learning and following established patterns and best practices, you can take your code and your career to a new level. With Practices of the Python Pro, you'll learn to design professional-level, clean, easily maintainable software at scale using the incredibly popular programming language, Python. You'll find easy-to-grok examples that use pseudocode and Python to introduce software development best practices, along with dozens of instantly useful techniques that will help you code like a pro. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Professional-quality code does more than just run without bugs. It's clean, readable, and easy to maintain. To step up from a capable Python coder to a professional developer, you need to learn industry standards for coding style, application design, and development process. That's where this book is indispensable. About the book Practices of the Python Pro teaches you to design and write professional-quality software that's understandable, maintainable, and extensible. Dane Hillard is a Python pro who has helped many dozens of developers make this step, and he knows what it takes. With helpful examples and exercises, he teaches you when, why, and how to modularize your code, how to improve quality by reducing complexity, and much more. Embrace these core principles, and your code will become easier for you and others to read, maintain, and reuse. What's inside Organizing large Python projects Achieving the right levels of abstraction Writing clean, reusable code Inheritance and composition Considerations for testing and performance About the reader For readers familiar with the basics of Python, or another OO language. About the author Dane Hillard has spent the majority of his development career using Python to build web applications. Table of Contents: PART 1 WHY IT ALL MATTERS 1 ; The bigger picture PART 2 FOUNDATIONS OF DESIGN 2 ; Separation of concerns 3 ; Abstraction and encapsulation 4 ; Designing for high performance 5 ; Testing your software PART 3 NAILING DOWN LARGE SYSTEMS 6 ; Separation of concerns in practice 7 ; Extensibility and flexibility 8 ; The rules (and exceptions) of inheritance 9 ; Keeping things

lightweight 10 ; Achieving loose coupling PART 4 WHAT'S NEXT? 11 ; Onward and upward

Your Python code may run correctly, but you need it to run faster. By exploring the fundamental theory behind design choices, the updated edition of this practical guide, expanded and enhanced for Python 3, helps you gain a deeper understanding of Python's implementation. You'll learn how to locate performance bottlenecks and significantly speed up your code in high-data-volume programs. How can you take advantage of multicore architectures or clusters? Or build a system that can scale up and down without losing reliability? Experienced Python programmers and students alike will learn concrete solutions to these and other issues, along with war stories from companies that use high-performance Python for social media analytics, productionized machine learning, and other situations. Get a better grasp of NumPy, Cython, and profilers Learn how Python abstracts the underlying computer architecture Understand performant pandas Use profiling to find bottlenecks in CPU time and memory usage Write efficient programs by choosing appropriate data structures Speed up matrix, vector, and even tensor computations Use tools to compile Python down to machine code, on CPUs and GPUs Manage multiple I/O and computational operations concurrently Convert multiprocessing code to run on a local or remote cluster Deploy code faster using tools like Docker Solve large problems while using less RAM Get real-life stories and lessons from Python programmers How can you realize MySQL's full power? With High Performance MySQL, you'll learn advanced techniques for everything from setting service-level objectives to designing schemas, indexes, and queries to tuning your server, operating system, and hardware to achieve your platform's full potential. This guide also teaches database administrators safe and practical ways to scale applications through replication, load balancing, high availability, and failover. Updated to reflect recent advances in cloud- and self-hosted MySQL, InnoDB performance, and new features and tools, this revised edition helps you design a relational data platform that will scale with your business. You'll learn best practices for database security along with hard-earned lessons in both performance and database stability. Dive into MySQL's architecture, including key facts about its storage engines Learn how server configuration works with your hardware and deployment choices Make query performance part of your software delivery process Examine enhancements to MySQL's replication and high availability Compare different MySQL offerings in managed cloud environments Explore MySQL's full stack optimization from application-side configuration to server tuning Turn traditional database management tasks into automated processes

Effective Computation in Physics

Collecting Data from the Modern Web

Haskell High Performance Programming

Enabling Test-Driven Development, Domain-Driven Design, and Event-Driven Microservices

An Introduction to Real-World Programming with Java

Practical Performant Programming for Humans

Python Parallel Programming Cookbook

Learn how to use Python to create efficient applications About This Book Identify the bottlenecks in your applications and solve them using the best profiling techniques Write efficient numerical code in NumPy, Cython, and Pandas Adapt your programs to run on multiple processors and machines with parallel programming Who This Book Is For The book is aimed at Python developers who want to improve the performance of their application. Basic knowledge of Python is expected What You Will Learn Write efficient numerical code with the NumPy and Pandas libraries Use Cython and Numba to achieve native performance Find bottlenecks in your Python code using profilers Write asynchronous code using Asyncio and RxPy Use Tensorflow and Theano for automatic parallelism in Python Set up and run distributed algorithms on a cluster using Dask and PySpark In Detail Python is a versatile language that has found applications in many industries. The clean syntax, rich standard library, and vast selection of third-party libraries make Python a wildly popular language. Python High Performance is a practical guide that shows how to leverage the power of both native and third-party Python libraries to build robust applications. The book explains how to use various profilers to find performance bottlenecks and apply the correct algorithm to fix them. The reader will learn how to effectively use NumPy and Cython to speed up numerical code. The book explains concepts of concurrent programming and how to implement robust and responsive applications using Reactive programming. Readers will learn how to write code for parallel architectures using Tensorflow and Theano, and use a cluster of computers for large-scale computations using technologies such as Dask and PySpark. By the end of the book, readers will have learned to achieve performance and scale from their Python applications.

Style and approach A step-by-step practical guide filled with real-world use cases and examples

Create distributed applications with clever design patterns to solve complex problems Key Features Set up and run distributed algorithms on a cluster using Dask and PySpark Master skills to accurately implement concurrency in your code Gain practical experience of Python design patterns with real-world examples Book Description This Learning Path shows you how to leverage the power of both native and third-party Python libraries for building robust and responsive applications. You will learn about profilers and reactive programming, concurrency and parallelism, as well as tools for making your apps quick and efficient. You will discover how to write code for parallel architectures using TensorFlow and Theano, and use a cluster of computers for large-scale computations using technologies such as Dask and PySpark. With the knowledge of how Python design patterns work, you will be able to clone objects, secure interfaces, dynamically choose algorithms, and accomplish much more in high performance computing. By the end of this Learning Path, you will have the skills and confidence to build engaging models that quickly offer efficient solutions to your problems. This Learning Path includes content from the following Packt products: Python High Performance - Second Edition by Gabriele Lanaro Mastering Concurrency in Python by Quan Nguyen Mastering Python Design Patterns by Sakis Kasampalis What you will learn Use NumPy and pandas to import and manipulate datasets Achieve native performance with Cython and Numba Write asynchronous code using asyncio and RxPy Design highly scalable programs with application scaffolding Explore abstract methods to maintain data consistency Clone objects using the prototype pattern Use the adapter pattern to make incompatible interfaces compatible Employ the strategy pattern to dynamically choose an algorithm Who this book is for This Learning Path is specially designed for Python developers who want to build high-performance applications and learn about single core and multi-core programming, distributed concurrency, and Python design patterns. Some experience with Python programming language will help you get the most out of this Learning Path.

Master efficient parallel programming to build powerful applications using Python About This Book Design and implement efficient parallel

Where To Download High Performance Python Practical Performant Programming For Humans

software Master new programming techniques to address and solve complex programming problems Explore the world of parallel programming with this book, which is a go-to resource for different kinds of parallel computing tasks in Python, using examples and topics covered in great depth Who This Book Is For Python Parallel Programming Cookbook is intended for software developers who are well versed with Python and want to use parallel programming techniques to write powerful and efficient code. This book will help you master the basics and the advanced of parallel computing. What You Will Learn Synchronize multiple threads and processes to manage parallel tasks Implement message passing communication between processes to build parallel applications Program your own GPU cards to address complex problems Manage computing entities to execute distributed computational tasks Write efficient programs by adopting the event-driven programming model Explore the cloud technology with Django and Google App Engine Apply parallel programming techniques that can lead to performance improvements In Detail Parallel programming techniques are required for a developer to get the best use of all the computational resources available today and to build efficient software systems. From multi-core to GPU systems up to the distributed architectures, the high computation of programs throughout requires the use of programming tools and software libraries. Because of this, it is becoming increasingly important to know what the parallel programming techniques are. Python is commonly used as even non-experts can easily deal with its concepts. This book will teach you parallel programming techniques using examples in Python and will help you explore the many ways in which you can write code that allows more than one process to happen at once. Starting with introducing you to the world of parallel computing, it moves on to cover the fundamentals in Python. This is followed by exploring the thread-based parallelism model using the Python threading module by synchronizing threads and using locks, mutex, semaphores queues, GIL, and the thread pool. Next you will be taught about process-based parallelism where you will synchronize processes using message passing along with learning about the performance of MPI Python Modules. You will then go on to learn the asynchronous parallel programming model using the Python asyncio module along with handling exceptions. Moving on, you will discover distributed computing with Python, and learn how to install a broker, use Celery Python Module, and create a worker. You will also understand the StarCluster framework, Pycsp, Scoop, and Disco modules in Python. Further on, you will learn GPU programming with Python using the PyCUDA module along with evaluating performance limitations. Next you will get acquainted with the cloud computing concepts in Python, using Google App Engine (GAE), and building your first application with GAE. Lastly, you will learn about grid computing concepts in Python and using PyGlobus toolkit, GFTP and GASS COPY to transfer files, and service monitoring in PyGlobus. Style and approach A step-by-step guide to parallel programming using Python, with recipes accompanied by one or more programming examples. It is a practically oriented book and has all the necessary underlying parallel computing concepts.

Build software that combines Python ' s expressivity with the performance and control of C (and C++). It ' s possible with Cython, the compiler and hybrid programming language used by foundational packages such as NumPy, and prominent in projects including Pandas, h5py, and scikits-learn. In this practical guide, you ' ll learn how to use Cython to improve Python ' s performance—up to 3000x— and to wrap C and C++ libraries in Python with ease. Author Kurt Smith takes you through Cython ' s capabilities, with sample code and in-depth practice exercises. If you ' re just starting with Cython, or want to go deeper, you ' ll learn how this language is an essential part of any performance-oriented Python programmer ' s arsenal. Use Cython ' s static typing to speed up Python code Gain hands-on experience using Cython features to boost your numeric-heavy Python Create new types with Cython—and see how fast object-oriented programming in Python can be Effectively organize Cython code into separate modules and packages without sacrificing performance Use Cython to give Pythonic interfaces to C and C++

Where To Download High Performance Python Practical Performant Programming For Humans

libraries Optimize code with Cython 's runtime and compile-time profiling tools Use Cython 's prange function to parallelize loops transparently with OpenMP

Hands-On GPU Programming with Python and CUDA

Professional CUDA C Programming

Practical Programming for Total Beginners

Mastering Large Datasets with Python

Build high performance, concurrent, and multi-threaded apps with Python using proven design patterns

High Performance Computing

Elegant SciPy

If you're new to Java—or new to programming—this best-selling book will guide you through the language features and APIs of Java 11. With fun, compelling, and realistic examples, authors Marc Loy, Patrick Niemeyer, and Daniel Leuck introduce you to Java fundamentals—including its class libraries, programming techniques, and idioms—with an eye toward building real applications. You'll learn powerful new ways to manage resources and exceptions in your applications—along with core language features included in recent Java versions. Develop with Java, using the compiler, interpreter, and other tools Explore Java's built-in thread facilities and concurrency package Learn text processing and the powerful regular expressions API Write advanced networked or web-based applications and services

"If you're an experienced Python programmer, High Performance Python will guide you through the various routes of code optimization. You'll learn how to use smarter algorithms and leverage peripheral technologies, such as numpy, cython, cpython, and various multi-threaded and multi-node strategies. There's a lack of good learning and reference material available if you want to learn Python for highly computational tasks. Because of it, fields from physics to biology and systems infrastructure to data science are hitting barriers. They need the fast prototyping nature of Python, but too few people know how to wield it"—Publisher's description.

Much has changed in technology over the past decade. Data is hot, the cloud is ubiquitous, and many organizations need some form of automation. Throughout these transformations, Python has become one of the most popular languages in the world. This practical resource shows you how to use Python for everyday Linux systems administration tasks with today's most useful DevOps tools, including Docker, Kubernetes, and Terraform. Learning how to interact and automate with Linux is essential for millions of professionals. Python makes it much easier. With this book, you'll learn how to develop software and solve problems using containers, as well as how to

monitor, instrument, load-test, and operationalize your software. Looking for effective ways to "get stuff done" in Python? This is your guide. Python foundations, including a brief introduction to the language How to automate text, write command-line tools, and automate the filesystem Linux utilities, package management, build systems, monitoring and instrumentation, and automated testing Cloud computing, infrastructure as code, Kubernetes, and serverless Machine learning operations and data engineering from a DevOps perspective Building, deploying, and operationalizing a machine learning project

Summary Modern data science solutions need to be clean, easy to read, and scalable. In *Mastering Large Datasets with Python*, author J.T. Wolohan teaches you how to take a small project and scale it up using a functionally influenced approach to Python coding. You'll explore methods and built-in Python tools that lend themselves to clarity and scalability, like the high-performing parallelism method, as well as distributed technologies that allow for high data throughput. The abundant hands-on exercises in this practical tutorial will lock in these essential skills for any large-scale data science project. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Programming techniques that work well on laptop-sized data can slow to a crawl—or fail altogether—when applied to massive files or distributed datasets. By mastering the powerful map and reduce paradigm, along with the Python-based tools that support it, you can write data-centric applications that scale efficiently without requiring codebase rewrites as your requirements change. About the book *Mastering Large Datasets with Python* teaches you to write code that can handle datasets of any size. You'll start with laptop-sized datasets that teach you to parallelize data analysis by breaking large tasks into smaller ones that can run simultaneously. You'll then scale those same programs to industrial-sized datasets on a cluster of cloud servers. With the map and reduce paradigm firmly in place, you'll explore tools like Hadoop and PySpark to efficiently process massive distributed datasets, speed up decision-making with machine learning, and simplify your data storage with AWS S3. What's inside An introduction to the map and reduce paradigm Parallelization with the multiprocessing module and pathos framework Hadoop and Spark for distributed computing Running AWS jobs to process large datasets About the reader For Python programmers who need to work faster with more data. About the author J. T. Wolohan is a lead data scientist at Booz Allen Hamilton, and a PhD researcher at Indiana University, Bloomington. Table of Contents: PART 1 1 ; Introduction 2 ; Accelerating large

dataset work: Map and parallel computing 3 ; Function pipelines for mapping complex transformations 4 ; Processing large datasets with lazy workflows 5 ; Accumulation operations with reduce 6 ; Speeding up map and reduce with advanced parallelization PART 2 7 ; Processing truly big datasets with Hadoop and Spark 8 ; Best practices for large data with Apache Streaming and mrjob 9 ; PageRank with map and reduce in PySpark 10 ; Faster decision-making with machine learning and PySpark PART 3 11 ; Large datasets in the cloud with Amazon Web Services and S3 12 ; MapReduce in the cloud with Amazon's Elastic MapReduce

Foundations of Agile Python Development

Develop, manage, and deploy efficient machine learning applications with Python

Field Guide to Research with Python

Designing Data-Intensive Applications

Software Architecture with Python

Advanced Python Programming

Best Practices for Development

Apache Spark is amazing when everything clicks. But if you haven't seen the performance improvements you expected, or still don't feel confident enough to use Spark in production, this practical book is for you. Authors Holden Karau and Rachel Warren demonstrate performance optimizations to help your Spark queries run faster and handle larger data sizes, while using fewer resources. Ideal for software engineers, data engineers, developers, and system administrators working with large-scale data applications, this book describes techniques that can reduce data infrastructure costs and developer hours. Not only will you gain a more comprehensive understanding of Spark, you'll also learn how to make it sing. With this book, you'll explore: How Spark SQL's new interfaces improve performance over SQL's RDD data structure The choice between data joins in Core Spark and Spark SQL Techniques for getting the most out of standard RDD transformations How to work around performance issues in Spark's key/value pair paradigm Writing high-performance Spark code without Scala or the JVM How to test for functionality and performance when applying suggested improvements Using Spark MLlib and Spark ML machine learning libraries Spark's Streaming components and external community packages

Architect and design highly scalable, robust, clean, and highly performant applications in Python About This Book Identify design issues and make the necessary adjustments to achieve improved performance Understand practical architectural quality attributes from the perspective of a practicing engineer and architect using Python Gain knowledge of architectural principles and how they can be used to provide accountability and

rationale for architectural decisions Who This Book Is For This book is for experienced Python developers who are aspiring to become the architects of enterprise-grade applications or software architects who would like to leverage Python to create effective blueprints of applications. What You Will Learn Build programs with the right architectural attributes Use Enterprise Architectural Patterns to solve scalable problems on the Web Understand design patterns from a Python perspective Optimize the performance testing tools in Python Deploy code in remote environments or on the Cloud using Python Secure architecture applications in Python In Detail This book starts off by explaining how Python fits into an application architecture. As you move along, you will understand the architecturally significant demands and how to determine them. Later, you'll get a complete understanding of the different architectural quality requirements that help an architect to build a product that satisfies business needs, such as maintainability/reusability, testability, scalability, performance, usability, and security. You will use various techniques such as incorporating DevOps, Continuous Integration, and more to make your application robust. You will understand when and when not to use object orientation in your applications. You will be able to think of the future and design applications that can scale proportionally to the growing business. The focus is on building the business logic based on the business process documentation and which frameworks are to be used when. We also cover some important patterns that are to be taken into account while solving design problems as well as those in relatively new domains such as the Cloud. This book will help you understand the ins and outs of Python so that you can make those critical design decisions that not just live up to but also surpass the expectations of your clients. Style and approach Filled with examples and use cases, this guide takes a no-nonsense approach to help you with everything it takes to become a successful software architect.

If you're among the Python developers put off by asyncio's complexity, it's time to take another look. Asyncio is complicated because it aims to solve problems in concurrent network programming for both framework and end-user developers. The features you need to consider are a small subset of the whole asyncio API, but picking out the right features is the tricky part. That's where this practical book comes in. Veteran Python developer Caleb Hattingh helps you gain a basic understanding of asyncio's building blocks—enough to get started writing simple event-based programs. You'll learn why asyncio offers a safer alternative to preemptive multitasking (threading) and how this API provides a simple way to support thousands of simultaneous socket connections. Get a critical comparison of asyncio and threading for concurrent network programming Take an asyncio walk-through, including a quickstart guide for hitting the ground looping with event-based programming Learn the difference between asyncio features for end-user developers and those for framework developers Understand asyncio's new async/await language syntax, including coroutines and task and future APIs Get detailed case studies (with

code) of some popular asyncio-compatible third-party libraries

Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

A Guide for Python Programmers

Clean Code in Python

Web Scraping with Python

Practical Deep Learning for Cloud, Mobile, and Edge

Building with Python from First Principles

With PyGame

Elegant Coding in Python

Tackle inefficiencies and errors the Pythonic way Key Features Enhance your coding skills using the new features introduced in Python 3.9 Implement the refactoring techniques and SOLID principles in Python Apply microservices to your legacy systems by implementing practical techniques Book Description Experienced professionals in every field face several instances of disorganization, poor readability, and testability due to unstructured code. With updated code and revised content aligned to the new features of Python 3.9, this second edition of Clean Code in Python will provide you with all the tools you need to overcome these obstacles and manage your projects successfully. The book begins by describing the basic elements of writing clean code and how it plays a key role in Python programming. You will learn about writing efficient and readable code using the Python standard library and best practices for software design. The book discusses object-oriented programming in Python and shows you how to use objects with descriptors and generators. It will also show you the design principles of software testing and how to resolve problems by implementing software design patterns in your code. In the concluding chapter, we break down a monolithic application into a microservices-based one starting from the code as the basis for a solid platform. By the end of this clean code book, you will be proficient in applying industry-approved coding practices to design clean, sustainable, and readable real-world

Python code. What you will learn Set up a productive development environment by leveraging automatic tools Leverage the magic methods in Python to write better code, abstracting complexity away and encapsulating details Create advanced object-oriented designs using unique features of Python, such as descriptors Eliminate duplicated code by creating powerful abstractions using software engineering principles of object-oriented design Create Python-specific solutions using decorators and descriptors Refactor code effectively with the help of unit tests Build the foundations for solid architecture with a clean code base as its cornerstone Who this book is for This book is designed to benefit new as well as experienced programmers. It will appeal to team leads, software architects and senior software engineers who would like to write Pythonic code to save on costs and improve efficiency. The book assumes that you have a strong understanding of programming

The Python Book Discover the power of one of the fastest growing programming languages in the world with this insightful new resource The Python Book delivers an essential introductory guide to learning Python for anyone who works with data but does not have experience in programming. The author, an experienced data scientist and Python programmer, shows readers how to use Python for data analysis, exploration, cleaning, and wrangling. Readers will learn what in the Python language is important for data analysis, and why. The Python Book offers readers a thorough and comprehensive introduction to Python that is both simple enough to be ideal for a novice programmer, yet robust to be useful for those more experienced in the language. The book assists budding programmers to gradually increase their skills as they move through the book, always with an understanding of what they are covering and why it is useful. Used by major companies like Google, Facebook, Instagram, Spotify, and more, Python promises to remain central to the programming landscape for years to come. Containing a thorough discussion of Python programming topics like variables, equalities and comparisons, tuple and dictionary data types, while and for loops, and if statements, readers will also learn: How to use highly useful Python programming libraries, including Pandas and Matplotlib How to write Python functions and classes How to write and use Python scripts To deal with different data types within Python Perfect for statisticians, computer scientists, software programmers, and practitioners working in private industry and medicine, The Python Book will also be of interest to students in any of the aforementioned fields. As it assumes no programming experience or knowledge, the book is ideal for those who work with data and want to learn to use Python to enhance their work.

More physicists today are taking on the role of software developer as part of their research, but software development isn't always easy or obvious, even for physicists. This practical book teaches essential software development skills to help you automate and accomplish nearly any aspect of research in a physics-based field. Written by two PhDs in nuclear engineering, this book includes practical examples drawn from a working knowledge of physics concepts. You'll learn how to use the Python programming language to perform everything from collecting and analyzing data to building software and publishing your results. In four parts, this book includes: Getting Started: Jump into Python, the command line, data containers, functions, flow control and logic, and classes and objects Getting It Done: Learn about regular expressions, analysis and visualization, NumPy, storing data in files and HDF5, important data structures in physics, computing in parallel, and deploying software Getting It Right: Build pipelines and software, learn to use local and remote version control, and debug and test your code Getting It Out There: Document your code, process and publish your findings, and collaborate efficiently; dive into software licenses, ownership, and copyright procedures

* Covers low-level networking in Python —essential for writing a new networked application protocol. * Many working examples

demonstrate concepts in action -- and can be used as starting points for new projects. * Networked application security is demystified. * Exhibits and explains multitasking network servers using several models, including forking, threading, and non-blocking sockets. * Features extensive coverage of Web and E-mail. Describes Python's database APIs.

Foundations of Python Network Programming

Develop maintainable and efficient code

Real-World AI & Computer-Vision Projects Using Python, Keras & TensorFlow

High Performance Python

Modern Computing in Simple Packages

Parallelize and Distribute Your Python Code

Get well-versed with FastAPI features and best practices for testing, monitoring, and deployment to run high-quality and robust data science applications

Key FeaturesCover the concepts of the FastAPI framework, including aspects relating to asynchronous programming, type hinting, and dependency injectionDevelop efficient RESTful APIs for data science with modern PythonBuild, test, and deploy high performing data science and machine learning systems with FastAPIBook Description FastAPI is a web framework for building APIs with Python 3.6 and its later versions based on standard Python-type hints. With this book, you'll be able to create fast and reliable data science API backends using practical examples. This book starts with the basics of the FastAPI framework and associated modern Python programming language concepts. You'll be taken through all the aspects of the framework, including its powerful dependency injection system and how you can use it to communicate with databases, implement authentication and integrate machine learning models. Later, you'll cover best practices relating to testing and deployment to run a high-quality and robust application. You'll also be introduced to the extensive ecosystem of Python data science packages. As you progress, you'll learn how to build data science applications in Python using FastAPI. The book also demonstrates how to develop fast and efficient machine learning prediction backends and test them to achieve the best performance. Finally, you'll see how to implement a real-time face detection system using WebSockets and a web browser as a client. By the end of this FastAPI book, you'll have not only learned how to implement Python in data science projects but also how to maintain and design them to meet high programming standards with the help of FastAPI. What you will learnExplore the basics of modern Python and async I/O programmingGet to grips with basic and advanced concepts of the FastAPI frameworkImplement a FastAPI dependency to efficiently run a machine learning modelIntegrate a simple face detection algorithm in a FastAPI backendIntegrate common Python data science libraries in a web backendDeploy a performant and reliable web backend for a data science applicationWho this book is for This Python data science book is for data scientists and software developers interested in gaining knowledge of FastAPI and its ecosystem to build data science applications. Basic knowledge of data science and machine learning concepts and how to apply them in Python is recommended.

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

Where To Download High Performance Python Practical Performant Programming For Humans

Beginning Python Games Development, Second Edition teaches you how to create compelling games using Python and the PyGame games development library. It will teach you how to create visuals, do event handling, create 3D games, add media elements, and integrate OpenGL into your Python game. In this update to the first ever book to cover the popular open source PyGame games development library, you'll stand to gain valuable technical insights and follow along with the creation of a real-world, freely downloadable video game. Written by industry veterans and Python experts Will McGugan and Harrison Kinsley, this is a comprehensive, practical introduction to games development in Python. You can also capitalize upon numerous tips and tricks the authors have accumulated over their careers creating games for some of the world's largest game developers.